HP
Switch Software
Advanced Traffic Management Guide

Abstract

Applicable Products

HP Switch 3500-24 (J9470A)
HP Switch 3500-48 (J9472A)
HP Switch 3500-24-PoE (J9471A)
HP Switch 3500-48-PoE (J9473A)
HP Switch 3500yl-24G-PWR (J8692A)
HP Switch 3500yl-48G-PWR (J8693A)
HP Switch 3800–24G-PoE+–2SFP+ (J9573A)
HP Switch 3800–48G-PoE+-4SFP+ (J9574A)
HP Switch 3800–24G-2SFP+ (J9575A)
HP Switch 3800–48G-4SFP+ (J9576A)
HP Switch 3800–24GS-2XG tl (J9584A)
HP Switch 3800–24G-2XGT tl (J9585A)
HP Switch 3800–48G-4XGT tl (J9586A)
HP Switch 3800–24G-2XGT-PoE+ tl (J9587A)
HP Switch 3800–48G-4XGT-PoE+ tl (J9588A)
HP Switch 3800 4–port Stacking Module (J9577A)
HP Switch 5406z (J8697A)
HP Switch 5406zl-48G-PoE+ (J9447A)
HP Switch 5412zl (J8698A)
HP Switch 5412zl-96G-PoE+ (J9448A)
HP Switch 6200yl-24G (J8992A)
HP Switch 8206zl (J9475A)
HP Switch 8212zl (J8715A/B)
HP Switch 6600-24G (J9263A)
HP Switch 6600-24G-4XG (J9264A)
HP Switch 6600-24G-24XG (J9265A)
HP Switch 6600-48G (J9451A)
HP Switch 6600-48G-4XG (J9452A)

* * 5 9 9 8 - 1 4 7 7 * *

HP Part Number: 5998-1477 Rev. B
Published: July 2012
Edition: Software version K.15.09

Acknowledgments

Warranty

For the software end user license agreement and the hardware limited warranty information for HP Networking products, visit www.hp.com/networking/support.

Contents

1 Static Virtual LANs (VLANs)

Command Syntax	Description	Default	CLI reference page	Menu reference page
show vlans	Displays VLAN configuration		16	
show vlans *vid*				
show vlans ports *port-list*				
show vlan ports *port-list* [detail]	Displays VLAN memberships		18	
show vlans custom [port *port-list*] column-list	Customizes show vlan output		21	
max-vlans *1-2048*	Changes the number of VLANs allowed on a switch		23	
primary-vlan [*vid* | *asciiname-string*]	Assigns the primary VLAN		24	
vlan [*vid* | *asciiname-string*]	Creates a new static VLAN		27	
no vlan *vid*	Deletes a static VLAN		29	
static-vlan *vlan-id*	Converts a dynamic to a static VLAN		30	
[no] vlan *vid*	Configures static VLAN per-port settings		30	26
[no] management-vlan [*vlan-id* | *vlan-name*]	Makes an existing VLAN the management VLAN	Disabled	(page 37)	
vlan *vid* qos priority *0- 7*	Prioritizes Voice VLAN QoS	1 (normal)	41	
[no] ip-recv-mac-address *mac-address* [interval *seconds*]	Configures a VLAN MAC address with heartbeat interval	60 seconds	42	
show ip-recv-mac-address	Displays the VLAN MAC address configuration		42	

General steps for using VLANs

VLANs enable grouping users by logical function instead of physical location. They make managing bandwidth usage within networks possible by:

- Allowing grouping of 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.

This chapter describes static VLANs configured for port-based or protocol-based operation.

Static VLANs are configured with a name, VLAN ID number (VID), and port members. For *dynamic* VLANs, see "GVRP" (page 67). 802.1Q compatibility enables you to assign each switch port to multiple VLANs.

Some recommended steps to take for using VLANs:

1. Plan your VLAN strategy and create a map of the logical topology. Include consideration for the interaction between VLANs and other features such as Spanning Tree Protocol, port trunking, and IGMP. See "Effects of VLANs on other switch features" (page 62). If you plan on using dynamic VLANs, include the port configuration planning necessary to support this feature, see "GVRP" (page 67).

   By default, the switches covered in this guide are 802.1Q VLAN-enabled, allow for up to 256 static VLANs, and 2048 total static and dynamic VLANs.
2. Configure at least one VLAN in addition to the default VLAN.
3. Assign the desired switch ports to the new VLANs.
4. If you are managing VLANs with SNMP in an IP network, the VLAN through which you are managing the switch must have an IP address. For information on the procedure and restrictions when you configure an IP address on a VLAN interface, see Table 2 (page 43).

Configuring VLANs

The Menu interface enables configuration and display of port-based VLANs only. 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 up to 2048 VLANs, with up to 4094 VIDs, by adding new VLAN names and VIDs, and then assigning one or more ports to each VLAN.

Per-port static VLAN configuration options

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

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

Figure 1 Comparing per-port VLAN options with and without GVRP

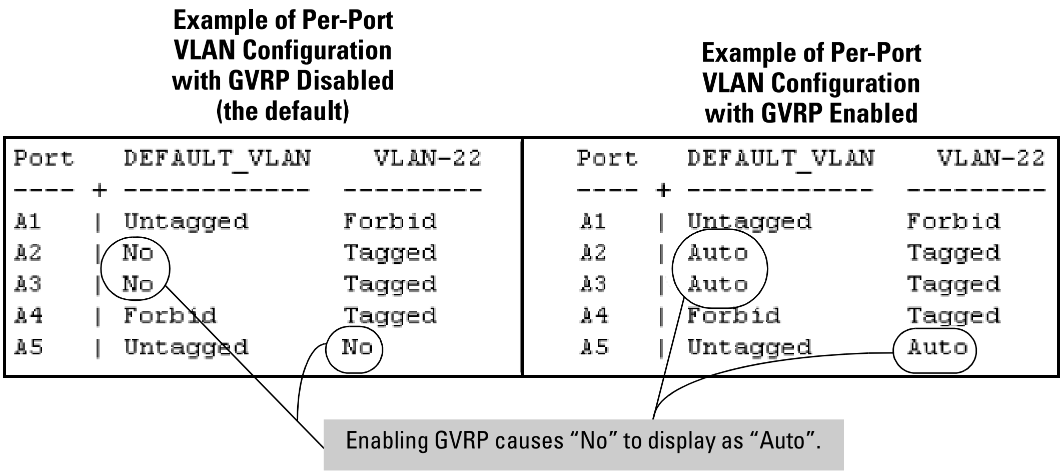
Example of Per-Port VLAN Configuration with GVRP Disabled (the default)
Example of Per-Port VLAN Configuration with GVRP Enabled

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

```
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 1 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.

`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; allows the port to dynamically join any advertised VLAN that has the same VID
	`Auto:`
	When GVRP is enabled on the switch; 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 and protocol-based VLAN parameters (CLI)

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 up to 255 additional static VLANs by adding new VLAN names, and then assigning one or more ports to each VLAN.

The switch accepts a maximum of 2048 VLANs with VIDs numbered up to 4094 This must include the default VLAN and any dynamic VLANs the switch creates if you enable GVRP.

NOTE: Each port can be assigned to multiple VLANs by using VLAN tagging. See “802.1Q VLAN tagging” (page 54).

Displaying a switch's VLAN configuration (CLI)

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
```

Maximum VLANs to support

Shows the number of VLANs the switch can currently support. Default is 256, Maximum 2048)

Primary VLAN

See "The pimary VLAN" (page 58).

Management VLAN

See "The secure Management VLAN" (page 59).

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 "Voice VLANs" (page 61).

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.

Example

Example 1 Displaying VLAN listing with GVRP enabled

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" (page 67).

```
HP 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.

Descriptions of items displayed by the command are:

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.

Examples

Example 2 Displaying VLAN ports (cumulative listing)

```
HP 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
```

Example 3 Displaying VLAN ports (detailed listing)

```
HP 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.

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” (page 67).

Voice

Indicates whether a port-based VLAN is configured as a voice VLAN. See “Voice VLANs” (page 61).

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.

Examples

Example 4 Displaying information for a specific static VLAN

```
HP 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
```

Example 5 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.

```
HP 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

Examples

Example 6 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.

```
HP 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
```

Example 7 Wrapping column headers

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

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

 VLANID
 ------
 1
 33
 44

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

 VL
 --
 1
 33
 44
```

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.

Example 8 Using a VLAN alias

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

HP 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
```

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:

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

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

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

Syntax:

`max-vlans` *`1-2048`*

In the default VLAN configuration, the switch allows a maximum of 256 VLANs. Use this command to specify the maximum VLANs to allow, and specify any value from 1 to 2048.

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.

Example

Example 9 Changing the number of allowed VLANs

The following example shows the command sequence for changing the number of VLANs allowed to 10. Note that you can execute the commands to `write memory` and `boot` at another time.

```
HP Switch(config)# max-vlans 10
This command will take effect after saving the configuration
and rebooting the system.
HP Switch(config)# write memory
HP Switch(config)# boot
Device will he rebooted, do you want to continue [y/n]? y
```

Assigning the Primary VLAN (CLI)

Syntax:

`primary-vlan  vid |` *`ascii-name-string`*

In the default VLAN configuration, the port-based default VLAN (`DEFAULT_VLAN`) is the Primary VLAN. This command allows reassignment of the Primary VLAN function to an existing, port-based, static VLAN.

The switch will not reassign the Primary VLAN function to a protocol VLAN.

NOTE: If you reassign the Primary VLAN to a non-default VLAN, you cannot later delete that VLAN from the switch until you 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`.

Example

Example 10 Re-assigning, renaming and displaying the VLAN command sequence

The following example shows how to re-assign 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):

```
HP Switch(config)# primary-vlan 22
HP Switch(config)# vlan 22 name 22-Primary
HP 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
```

Adding or editing VLAN names (Menu)

Use this procedure to add a new VLAN or to edit the name of an existing VLAN.

1. From the Main Menu select **2. Switch Configuration** —> **8. VLAN Menu ...** —> **2. VLAN Names**

   If multiple VLANs are not yet configured, you will see a screen similar to Figure 2 (page 25).

   Figure 2 The default VLAN names screen

   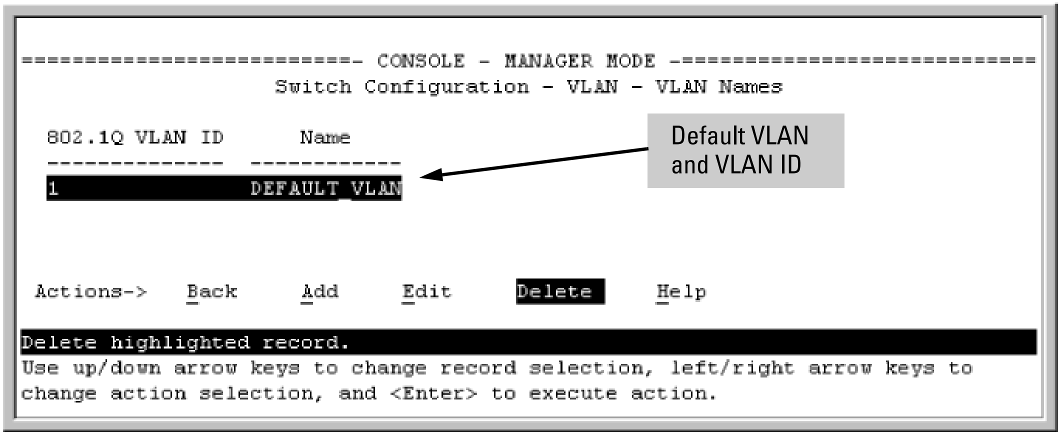

```
==========================- CONSOLE - MANAGER MODE -============================
                  Switch Configuration - VLAN - VLAN Names

  802.1Q VLAN ID      Name
  --------------    ------------
  1                 DEFAULT_VLAN



 Actions->   Back     Add     Edit     Delete     Help

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

   Default VLAN and VLAN ID

2. Press **A** (for Add).

   You will be prompted for a new VLAN name and VLAN ID:

   802.1Q VLAN ID :
   1 Name : _

3. Type a VID (VLAN ID number). This can be any number from 2 to 4094 that is not already being used by another VLAN (the switch reserves 1 for the default VLAN).

   NOTE: A VLAN must have the same VID in every switch in which you configure that same VLAN. GVRP dynamically extends VLANs with correct VID numbering to other switches; see “GVRP” (page 67) .

4. Press ↓ key to move the cursor to the **Name** line and enter the VLAN name, using up to 12 characters with no spaces. Press **Enter**.

   NOTE: Do not use the following characters in VLAN names: **@, #, $, ^, &, ***, **(**, and **)**.

5. Press **S** (for Save).

   The VLAN Names screen appears with the new VLAN listed.

   Figure 3 VLAN Names screen with a new VLAN added

   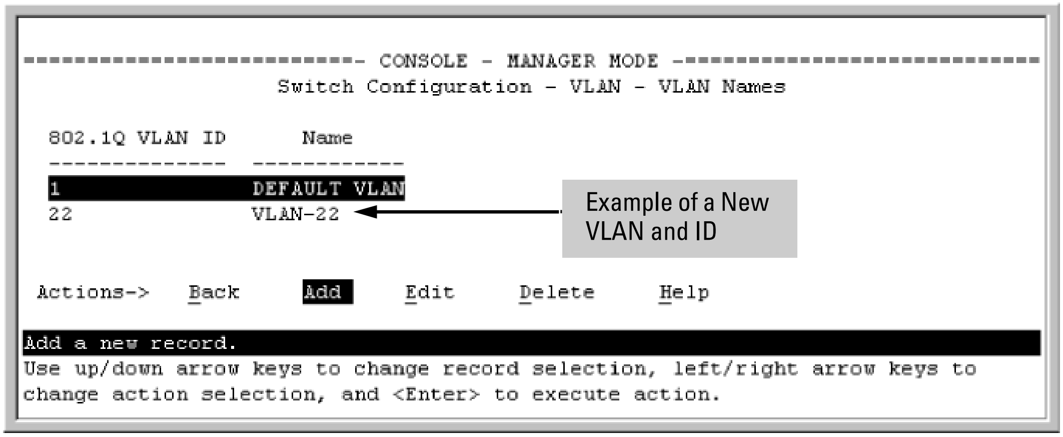

```
==========================- CONSOLE - MANAGER MODE -============================
                  Switch Configuration - VLAN - VLAN Names

  802.1Q VLAN ID      Name
  --------------    ------------
  1                 DEFAULT_VLAN
  22                VLAN-22

 Actions->   Back     Add     Edit     Delete     Help

Add a new record.
Use up/down arrow keys to change record selection, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

   Example of a New VLAN and ID

6. Repeat steps 2 through 5 to add more VLANs.

   You can add VLANs until you reach the number specified in the **Maximum VLANs to support** field on the VLAN Support screen. This includes any VLANs added dynamically due toGVRP operation.

Return to the VLAN Menu to assign ports to the new VLAN, as described in "Adding or changing a VLAN port assignment (Menu)" (page 33).

Changing VLAN support settings (Menu)

The following procedure provides instructions for changing the maximum number of VLANs to support, changing the primary VLAN selection, and enabling or disabling dynamic VLANs.

1. From the Main Menu select: **2. Switch Configuration** —> **8. VLAN Menu …** —> **1. VLAN Support**

   You see the following screen:

Figure 4 The default VLAN Support screen

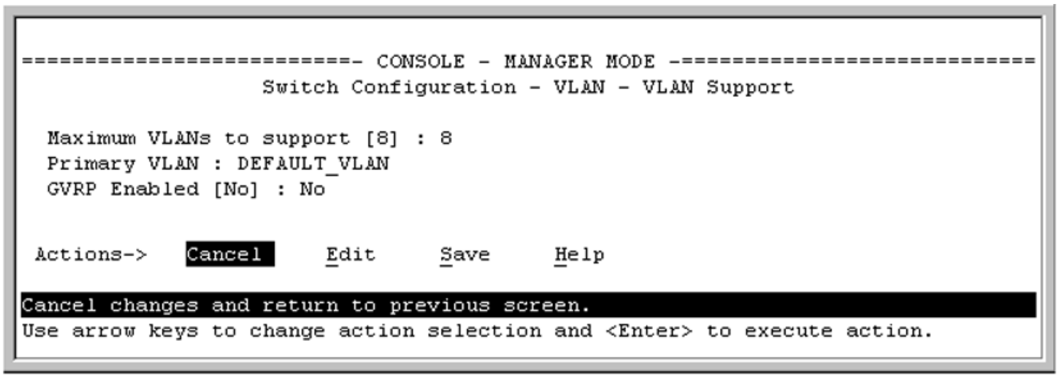

```
=========================- CONSOLE - MANAGER MODE -============================
                   Switch Configuration - VLAN - VLAN Support

  Maximum VLANs to support [8] : 8
  Primary VLAN : DEFAULT_VLAN
  GVRP Enabled [No] : No


 Actions->   Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

2. Press **E** (for Edit), and then do one or more of the following:
   - To change the maximum number of VLANs, enter the new number (1 - 2048 allowed; default 256).
   - To designate a different VLAN as the Primary VLAN, select the **Primary VLAN** field and use the space bar to select from the existing options. The Primary VLAN must be a static, port-based VLAN.
   - To enable or disable dynamic VLANs, select the **GVRP Enabled** field and use the Space bar to toggle between options. For GVRP information, see "GVRP" (page 67).

   NOTE: For optimal switch memory utilization, set the number of VLANs at the number you will likely be using or a few more. If you need more VLANs later, you can increase this number, but a switch reboot will be required at that time.

3. Press **Enter** and then **S** to save the VLAN support configuration and return to the VLAN Menu screen.

   If you changed the value for Maximum **VLANs to support**, an asterisk appears next to the **VLAN Support** option; see Figure 5 (page 27).

Figure 5 VLAN menu screen indicating the need to reboot the switch

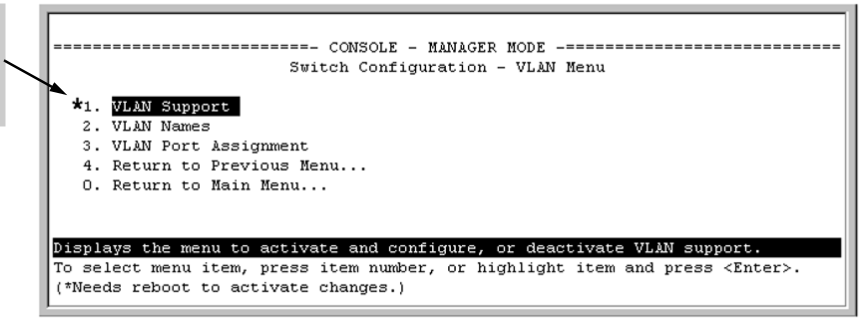

```
==========================- CONSOLE - MANAGER MODE -============================
                          Switch Configuration - VLAN Menu

  *1. VLAN Support
   2. VLAN Names
   3. VLAN Port Assignment
   4. Return to Previous Menu...
   0. Return to Main Menu...


Displays the menu to activate and configure, or deactivate VLAN support.
To select menu item, press item number, or highlight item and press <Enter>.
(*Needs reboot to activate changes.)
```

An asterisk indicates you must reboot the switch to implement the new Maximum VLANs setting.

- If you changed the VLAN Support option, you must reboot the switch before the maximum VLANs change takes effect. You can go on to configure other VLAN parameters first, but you must reboot the switch when you finish.
- If you did not change the VLAN Support option, a reboot is not necessary.

4. Press **0** to return to the Main Menu.

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 already 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 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 "Voice VLANs" (page 61).

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.

Examples

Example 11 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" (page 59).

```
HP Switch(config)# vlan 100
HP Switch(config)# show vlans

 Status and Counters - VLAN Information
 Maximum VLANs to support : 8
 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
```

Example 12 Changing the VLAN context level

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

```
HP Switch (vlan-100)# vlan default_vlan
HP Switch(vlan-1) _
```

Deleting a static VLAN (CLI)

Syntax:

```
no vlan vid
```

△ **CAUTION:** Prior to deleting a static VLAN, re-assign all ports in the VLAN to another VLAN.

Example

Example 13 Deleting a static VLAN

Using Figure 3 (page 25), 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.)

```
HP 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
HP Switch(config)#
```

Converting a dynamic VLAN to a static VLAN (CLI)

Syntax:

`static-vlan` *`vlan-id`*

Converts a dynamic, port-based VLAN membership to a 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" (page 67).

Example

Example 14 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:

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

Configuring static VLAN per-port settings (CLI)

Syntax:

`[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" (page 67).

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" (page 67).

Examples

Example 15 Changing the VLAN name and set ports to tagged

Suppose 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:

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

Example 16 Moving the context level

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

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

Example 17 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:

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

- or -

At the VLAN 100 context level, use:

```
HP Switch(vlan-100)# no tagged a1-a5
```

NOTE: You cannot use these commands with dynamic VLANs. Attempting to do so results in the message VLAN already exists and no change occurs.

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 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:

[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.

If QinQ is enabled, `svlan` can be configured as well.

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.

Examples

Example 18 Displaying a VLAN disabled for Layer 3

```
HP 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      :

  VLAN                 | IP Config  IP Address      Subnet Mask     Proxy ARP
  -------------------- + ---------- --------------- --------------- ---------
  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.

Example 19 Displaying IPv6 Layer 3 status for a VLAN

```
HP 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.

Adding or changing a VLAN port assignment (Menu)

Ports not specifically assigned to a VLAN are automatically in the default VLAN.

1. From the Main Menu select: **2. Switch Configuration** —> **8. VLAN Menu …** —> **3. VLAN Port Assignment**

   You will see a screen similar to the following:

   Figure 6 Port-based VLAN port assignment screen in the menu interface

   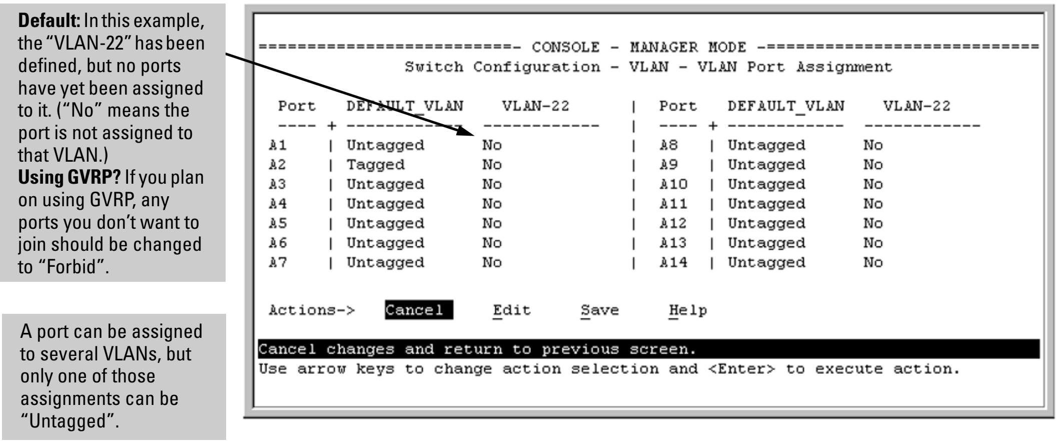

```
=========================- CONSOLE - MANAGER MODE -============================
           Switch Configuration - VLAN - VLAN Port Assignment

  Port   DEFAULT_VLAN    VLAN-22        |  Port   DEFAULT_VLAN    VLAN-22
  ---- + ------------  ------------     |  ---- + ------------  ------------
 A1    | Untagged      No               |  A8   | Untagged      No
 A2    | Tagged        No               |  A9   | Untagged      No
 A3    | Untagged      No               |  A10  | Untagged      No
 A4    | Untagged      No               |  A11  | Untagged      No
 A5    | Untagged      No               |  A12  | Untagged      No
 A6    | Untagged      No               |  A13  | Untagged      No
 A7    | Untagged      No               |  A14  | Untagged      No

 Actions->   Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

   Default: In this example, the "VLAN-22" has been defined, but no ports have yet been assigned to it. ("No" means the port is not assigned to that VLAN.)
   Using GVRP? If you plan on using GVRP, any ports you don't want to join should be changed to "Forbid".

   A port can be assigned to several VLANs, but only one of those assignments can be "Untagged".

   NOTE: The "VLAN Port Assignment" screen displays up to 32 static, port-based VLANs in ascending order, by VID. If the switch configuration includes more than 32 such VLANs, use the following CLI command to list data on VLANs having VIDs numbered sequentially higher than the first 32.

   ```
   show vlans [ VID | ports [ port-list ]]
   ```

2. To change a port's VLAN assignment:
   a. Press **E** (for Edit).
   b. Use the arrow keys to select a VLAN assignment you want to change.
   c. Press the Space bar to make your assignment selection (**No**, **Tagged**, **Untagged** , or **Forbid**. For information on VLAN tags, see "802.1Q VLAN tagging" (page 54).
   d. If you are finished assigning ports to VLANs, press **Enter** and then **S** (for Save) to activate the changes and return to the Configuration menu. (The console then returns to the VLAN menu.)
3. Return to the Main menu.

   NOTE: For GVRP Operation: If you enable GVRP on the switch, **No** converts to **Auto**, which allows the VLAN to dynamically join an advertised VLAN that has the same VID.

   Untagged VLANs

   Only one untagged VLAN is allowed per port. Also, there must be at least one VLAN assigned to each port. In the factory default configuration, all ports are assigned to the default VLAN (DEFAULT_VLAN).

For ports A4 and A5 to belong to both DEFAULT_VLAN and VLAN-22, and ports A6 and A7 to belong only to VLAN-22, use the settings in "The default VLAN names screen" (page 25). This example assumes that the default GVRP setting is `disabled` and that you do not plan to enable GVRP later.

Example 20 Displaying port-based VLAN assignments for specific ports

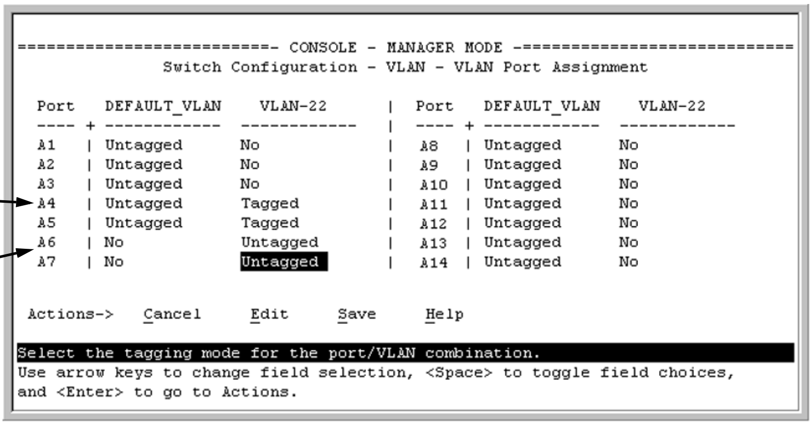

```
==========================- CONSOLE - MANAGER MODE -============================
                Switch Configuration - VLAN - VLAN Port Assignment

  Port   DEFAULT_VLAN    VLAN-22       |  Port   DEFAULT_VLAN    VLAN-22
  ---- + ------------  ------------    |  ---- + ------------  ------------
  A1   | Untagged     No               |  A8   | Untagged     No
  A2   | Untagged     No               |  A9   | Untagged     No
  A3   | Untagged     No               |  A10  | Untagged     No
  A4   | Untagged     Tagged           |  A11  | Untagged     No
  A5   | Untagged     Tagged           |  A12  | Untagged     No
  A6   | No           Untagged         |  A13  | Untagged     No
  A7   | No           Untagged         |  A14  | Untagged     No


 Actions->   Cancel     Edit     Save     Help

Select the tagging mode for the port/VLAN combination.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Ports A4 and A5 are assigned to both VLANs.

Ports A6 and A7 are assigned only to VLAN-22.

All other ports are assigned only to the Default VLAN.

Example of correcting an unsupported configuration

The problem

In Example 21 (page 35), the MAC address table for Switch 8000M will sometimes record the switch as accessed on port A1 (VLAN 1), and other times as accessed on port B1 (VLAN 2):

Example 21 An invalid configuration for single-forwarding to multiple-forwarding database devices in a multiple VLAN environment

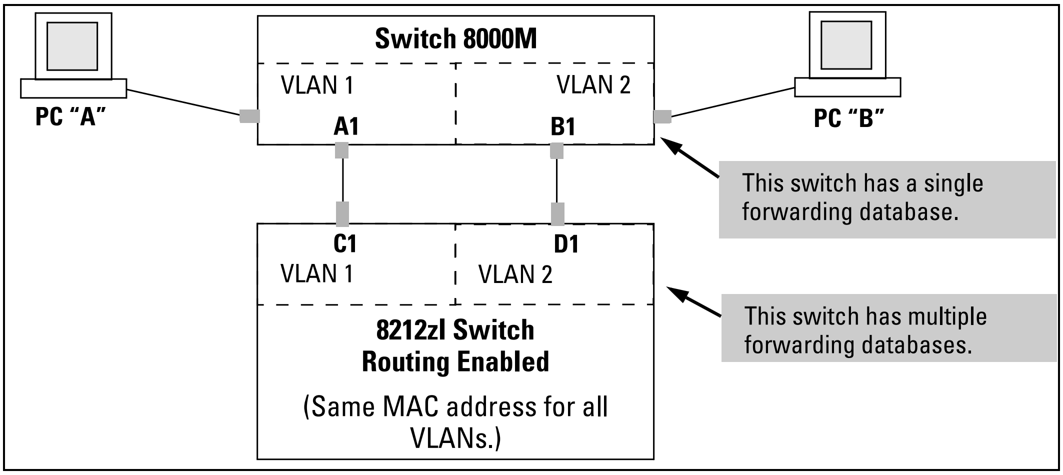

This switch has a single forwarding database.

This switch has multiple forwarding databases.

Here, PC A sends an IP packet to PC B.

1. The packet enters VLAN 1 in the Switch 8000 with the 8212zl switch's MAC address in the destination field. Because the 8000M 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 8212zl switch. The 8212zl switch then routes the packet through the VLAN 2 link to the 8000M, which forwards the packet on to PC "B". Because the 8000M received the packet from the 8212zl switch on VLAN 2 (port "B1"), the 8000M's single forwarding database records the 8212zl switch as being on port "B1" (VLAN 2).
2. PC "A" now sends a second packet to PC "B". The packet again enters VLAN 1 in the Switch 8000 with the 8212zl switch's MAC address in the destination field. However, this time the Switch 8000M's single forwarding database indicates that the 8212zl is on port B1 (VLAN 2), and the 8000M drops the packet instead of forwarding it.

3. Later, the 8212zl switch transmits a packet to the 8000M through the VLAN 1 link, and the 8000M updates its address table to indicate that the 8212zl switch is on port A1 (VLAN 1) instead of port B1 (VLAN 2). Thus, the 8000M's information on the location of the 8212zl switch changes over time. For this reason, the 8000M discards some packets directed through it for the 8212zl switch, resulting in poor performance and the appearance of an intermittent or broken link.

The solution

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

Now, the 8000M forwarding database always lists the 8212zl MAC address on port A1, and the 8000M will send traffic to either VLAN on the 8212zl.

Example 22 A solution for single-forwarding to multiple-forwarding database devices in a multiple VLAN environment

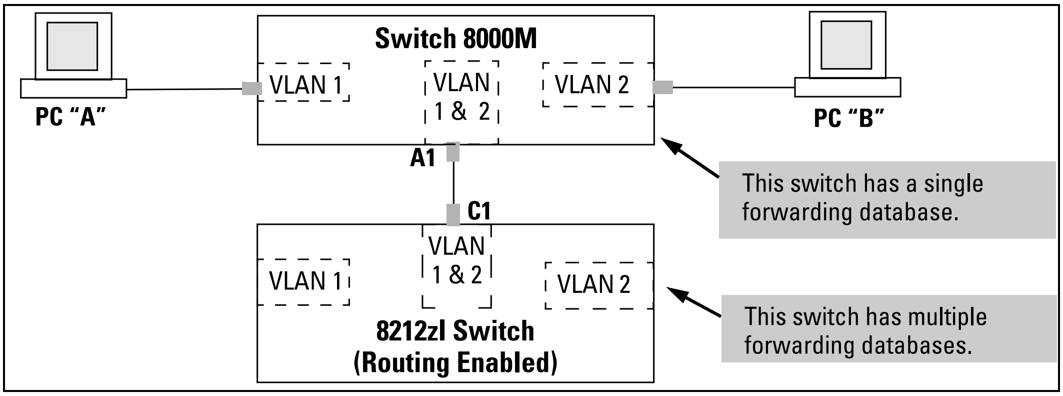
PC "A"
Switch 8000M
VLAN 1
VLAN 1 & 2
VLAN 2
PC "B"
A1
C1
This switch has a single forwarding database.
VLAN 1
VLAN 1 & 2
VLAN 2
8212zl Switch (Routing Enabled)
This switch has multiple forwarding databases.

Connecting an HP Switch to another 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 Table 4 (page 53). 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:

Example 23 A valid topology for devices having multiple forwarding databases in a multiple VLAN environment

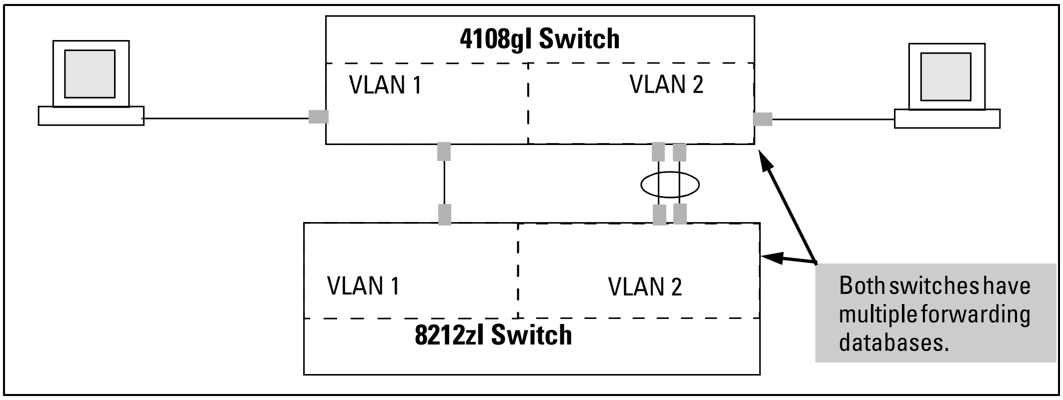

Configuring a secure Management VLAN (CLI)

Preparation

1. Determine a VID and VLAN name suitable for your Management VLAN.
2. Plan your topology to use HP switches that support Management VLANs. See "The secure Management VLAN" (page 59).
3. Include only the following ports:
   - Ports to which you will connect authorized management stations, such as Port A7 in Example 44 (page 60).
   - Ports on one switch that you will use to extend the Management VLAN to ports on other HP switches, such as ports A1 and Example 44 (page 60).
4. Half-duplex repeaters dedicated to connecting management stations to the Management VLAN can also be included in this topology. Note that 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:

`[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.

Examples

Example 24 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 HP switch that supports the Management VLAN feature.

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

Example 25 Configuration Example

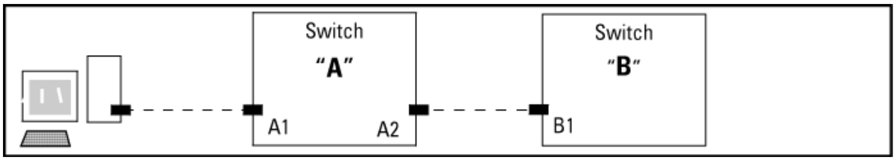

Obtaining an IP address using DHCP (CLI)

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.

Examples

Example 26 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 of 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.

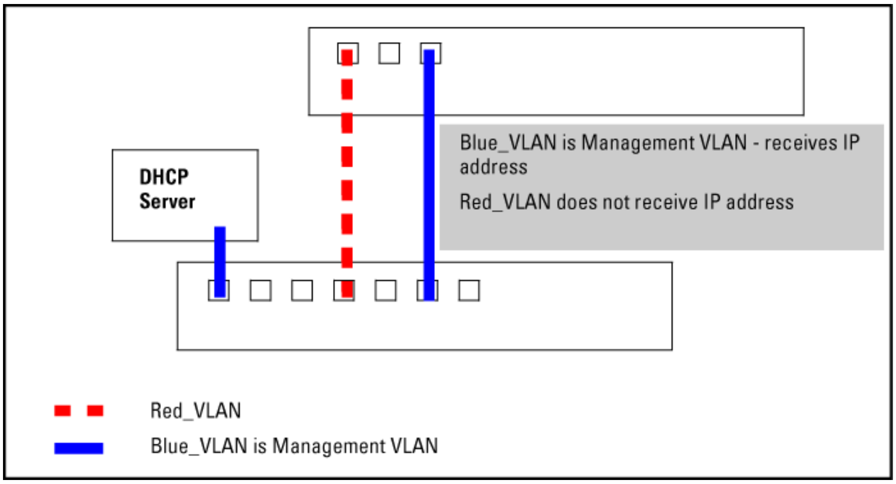
DHCP Server
Blue_VLAN is Management VLAN - receives IP address
Red_VLAN does not receive IP address
Red_VLAN
Blue_VLAN is Management VLAN

Example 27 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.

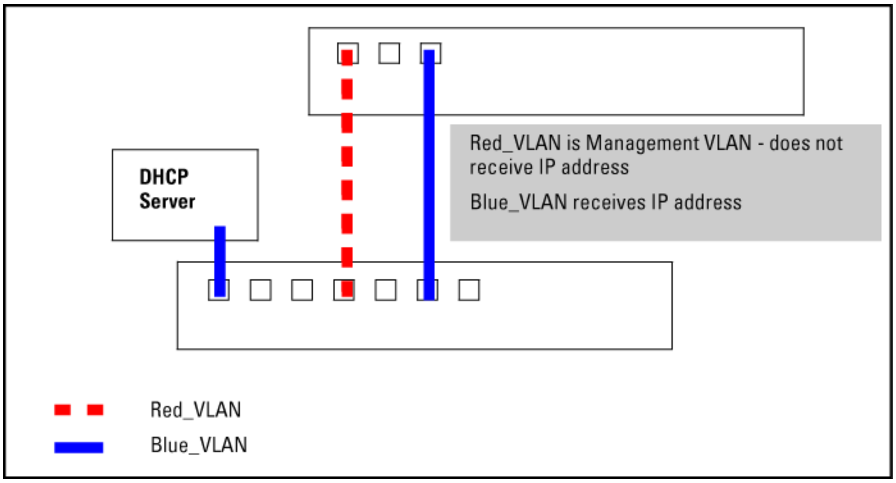
DHCP Server
Red_VLAN is Management VLAN - does not receive IP address
Blue_VLAN receives IP address
Red_VLAN
Blue_VLAN

Example 28 No Management VLANs configured

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

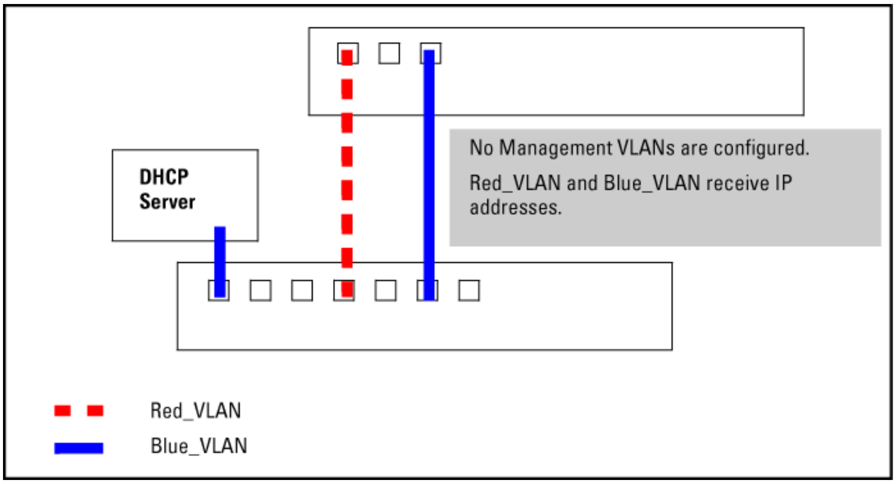
DHCP Server
No Management VLANs are configured.
Red_VLAN and Blue_VLAN receive IP addresses.
Red_VLAN
Blue_VLAN

Example 29 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.

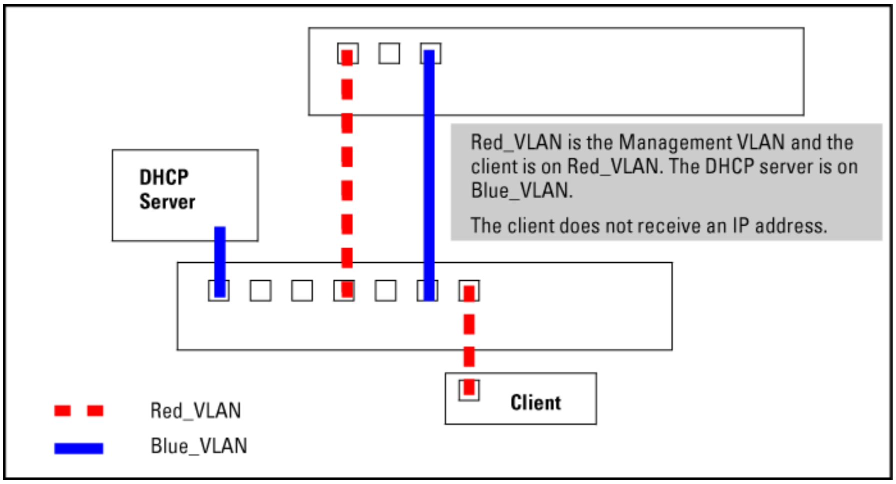
DHCP Server
Red_VLAN is the Management VLAN and the client is on Red_VLAN. The DHCP server is on Blue_VLAN.
The client does not receive an IP address.
Client
Red_VLAN
Blue_VLAN

Example 30 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.

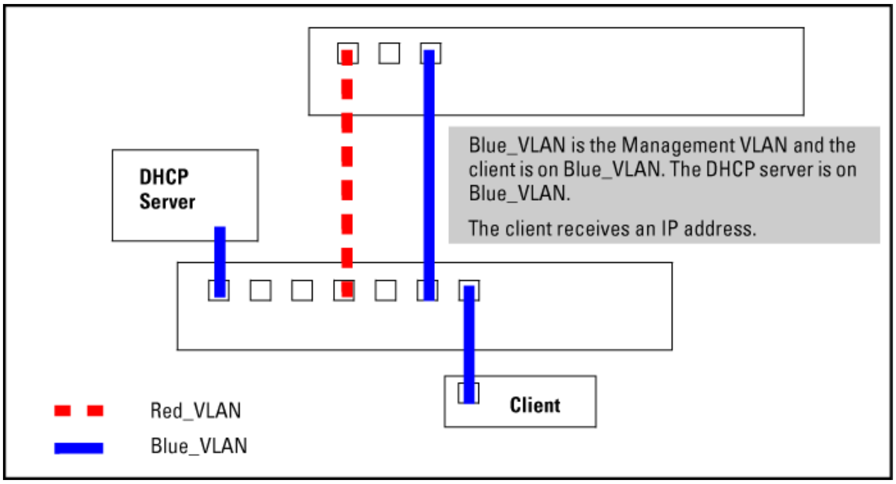
DHCP Server
Blue_VLAN is the Management VLAN and the client is on Blue_VLAN. The DHCP server is on Blue_VLAN.
The client receives an IP address.
Client
Red_VLAN
Blue_VLAN

Disabling the Management feature (CLI)

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

Example 31 Disabling the secure management feature

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

```
HP Switch (config)# no management-vlan 100
HP Switch (config)# no management-vlan my_vlan
```

For more information, see "The secure Management VLAN" (page 59).

Prioritizing voice VLAN QoS (CLI) (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:

```
HP Switch(config) # vlan 10 qos priority 7
HP 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: Managing bandwidth effectively" (page 204).

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 "Voice VLANs" (page 61).

Configuring a VLAN MAC address with heartbeat interval (CLI)

When installing HP 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 an HP routing switch.
- Optionally configure the time interval to use for sending heartbeat packets with the configured MAC address.

Syntax:

[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 HP switch.

`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
```

Example

Example 32 Displaying a VLAN MAC address

```
HP 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 static VLAN operation

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. Thus, all ports passing traffic for a particular subnet address should be configured to the same VLAN. Cross-domain broadcast traffic in the switch is eliminated and bandwidth is saved by not allowing packets to flood out all ports.

Table 2 Comparative operation of port based and protocol based VLANs

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 the 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 *Basic Operation Guide* for the switch.	You can configure IP addresses on all protocol VLANs. However, 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, which 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 tconfigured 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 *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 2 Comparative operation of port based and protocol based VLANs *(continued)*

Function	Port-Based VLANs	Protocol-Based VLANs
Tagged VLAN Membership	A port can be a tagged member of any port-based VLAN (see above).	A port can be a taggedmember of any protocol-based VLAN (see above).
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 configuration enables IP routing. 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 [ e |` *`port-list`* `]]`	`vlan` *`vid`* `protocol [ ipx | ipv4 | ipv6 | arp | appletalk | sna | netbeui ]`
		`vlan` *`vid`*`[ tagged | untagged [ e |` *`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" (page 58).

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, for IPv4 traffic.
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 switches covered in this guide allow up to 2048 (vids up to 4094) VLANs of all types. UsingVLAN tagging, ports can belong to multiple VLANs of all types. Enabling routing on the switch enables it route IPv4 traffic between port-based VLANs and between port-based VLANs and IPv4protocol 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 IPv4 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; see Table 2 (page 43).

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" (page 58).

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 pimary VLAN" (page 58).

Secure Management VLAN

This optional, port-based VLAN establishes an isolated network for managing HP 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 pimary VLAN" (page 58).

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 "Voice VLANs" (page 61).

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" (page 52).

The default VLAN

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

Example 33 A switch in the default VLAN configuration

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

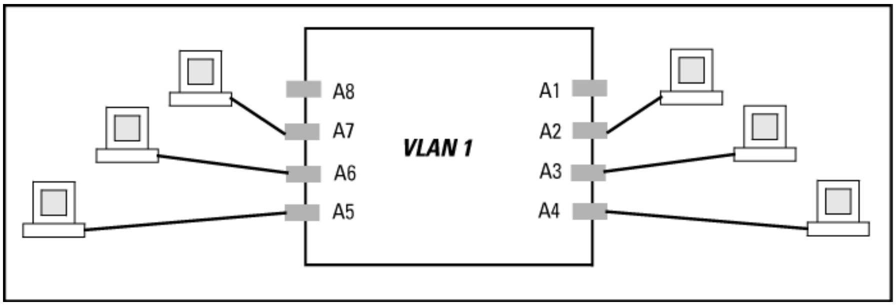

Multiple port-based VLANs

In Example 34 (page 46), routing within the switch is disabled (the default). This means that 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.

Note that 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.

Example 34 A switch with multiple VLANs configured and internal routing disabled

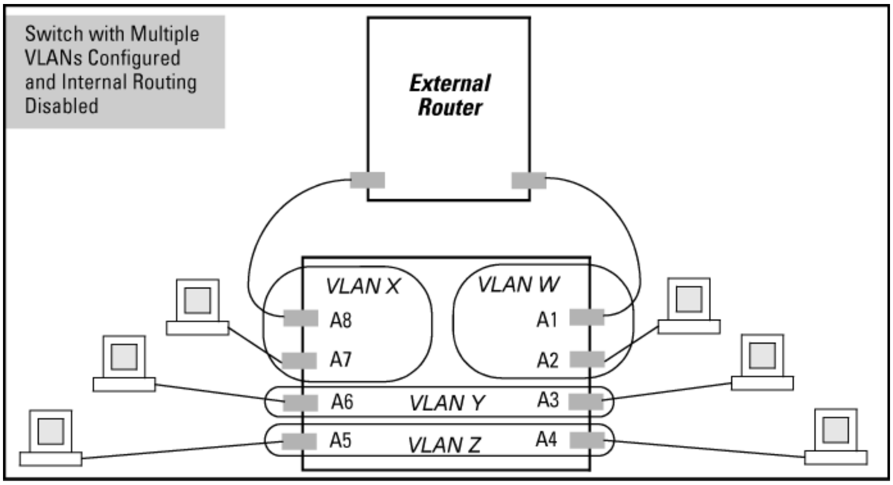
Switch with Multiple VLANs Configured and Internal Routing Disabled
External Router
VLAN X
A8
A7
VLAN W
A1
A2
A6
VLAN Y
A3
A5
VLAN Z
A4

Protocol VLAN environment

Example 34 (page 46) 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 3 Options for routing between VLAN types in the switch

		Port-Based	IPX	IPv4	IPv6	ARP	AppleTalk	SNA[1]	NETbeui[1]
Port-Based		Yes	—	Yes	—	—	—	—	—
Protocol	IPX	—	Yes[2]	—	—	—	—	—	—
	IPX4	Yes	—	Yes	—	—	—	—	—
	IPV6	—	—	—	Yes[2]	—	—	—	—
	ARP	—	—	—	—	Yes[2]	—	—	—
	AppleTalk	—	—	—	—	—	Yes[2]	—	—
	SNA	—	—	—	—	—	—	—	—
	NETbeui	—	—	—	—	—	—	—	—

[1] Not a routable protocol type. End stations intended to receive traffic in these protocols must be attached to the same physical network.

[2] Requires an external router to route between VLANs.

Overlapping (Tagged) VLANs

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 I.D. (VID) of the VLAN on which the packet was generated.

Example 35 Overlapping VLANs using the same server

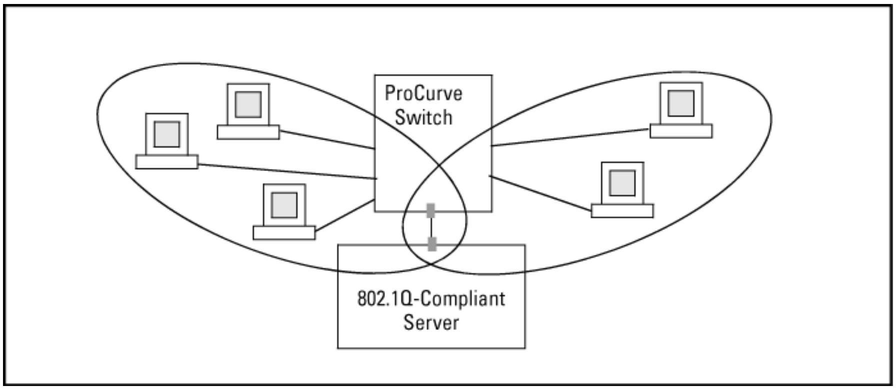

Similarly, using 802.1Q-compliant switches, you can connect multiple VLANs through a single switch-to-switch link.

Example 36 Connecting multiple VLANs through the same link

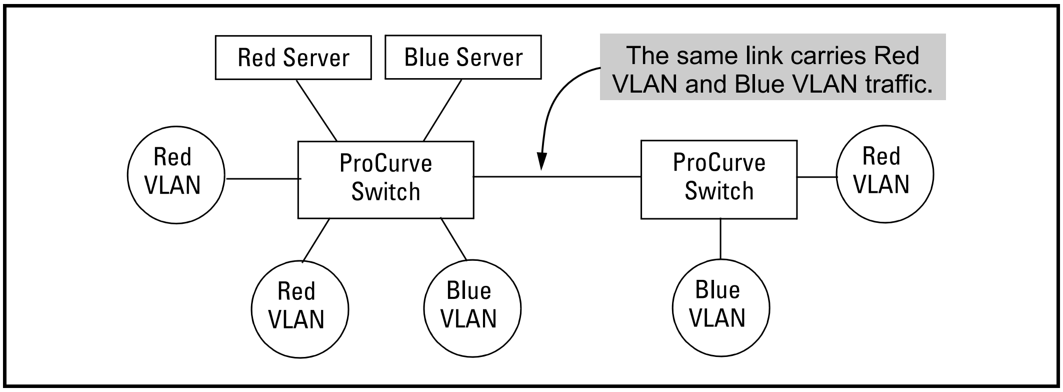

Introducing tagged VLAN technology into networks running untagged VLANs

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. This means that on the 802.1Q-compliant device, separate ports (configured as untagged) must be used to connect separate VLANs to non-802.1Q devices.

Example 37 Tagged and untagged VLAN technology in the same network

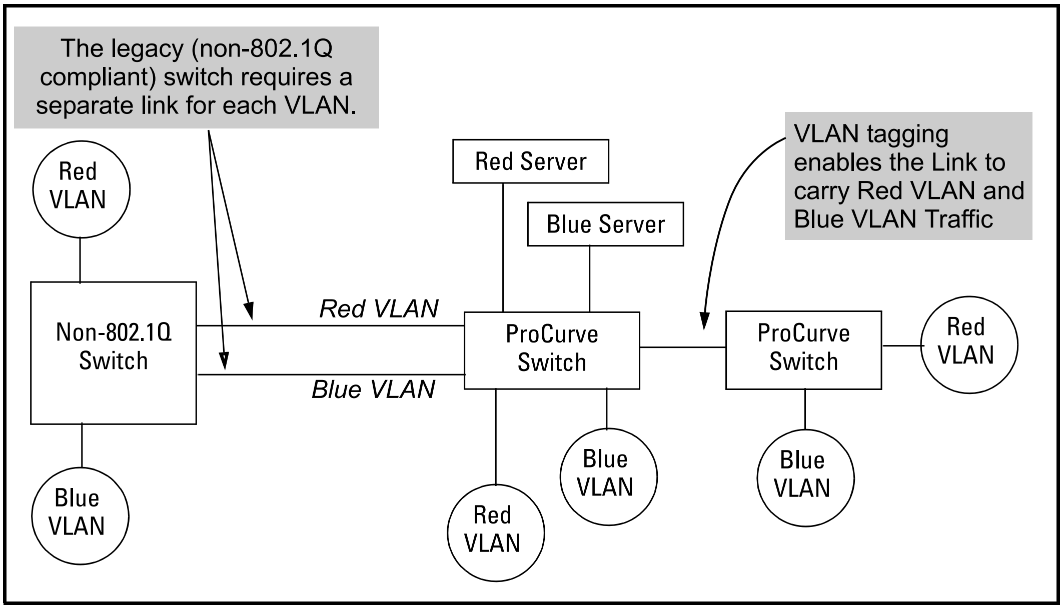
The legacy (non-802.1Q compliant) switch requires a separate link for each VLAN.
Red VLAN
Red Server
Blue Server
VLAN tagging enables the Link to carry Red VLAN and Blue VLAN Traffic
Red VLAN
Non-802.1Q Switch
ProCurve Switch
ProCurve Switch
Red VLAN
Blue VLAN
Blue VLAN
Blue VLAN
Blue VLAN
Red VLAN

VLAN Operating Rules

Disabled overlapping subnet configuration

Previous software versions allowed configuration of VLAN IP addresses in overlapping subnets, which can cause incorrect routing of packets and result in IP communication failure. As of software version K.15.09, overlapping subnet configurations are no longer allowed. An overlapping subnet is determined by the configuration order. The subnet that is configured first is valid, but any subsequent IP addresses that overlap are not allowed.

When the switch is booted into software version K.15.09 or later, and the configuration file includes overlapping subnets, the following occurs:

- The event log provides an error message in the format:

  ```
  ip: VLANx : IP initialization failed for vlan x.
  ```

  For a multinetted VLAN (multiple IP addresses assigned to the VLAN), only the IP addresses that are overlapping subnets are removed. The other IP addresses on the VLAN are retained and function correctly. The error message can be somewhat misleading; the IP addresses on the VLAN that are not overlapping are initialized correctly.

- The output of the `show ip` command correctly indicates that the overlapping IP address does not exist on the VLANs that have error messages in the event log.
- The output of the `show running-config` command incorrectly indicates that the overlapping IP address is configured. For example, in the following output, the IP address shown in VLAN6 is not actually configured on the VLAN; it has been removed.

  Example

  ### Example 38 An IP address that is not actually configured on the VLAN

  ```
  HP Switch(config)# show running-config

  .
  .
  .
    vlan 5
       name “VLAN5”
       ip address 11.22.33.1 255.0.0.0
       exit
    vlan 6
       name “VLAN6”
       ip address 11.23.34.1 255.255.255.0
       exit
  ```

The information is retained in the config file to allow you to boot up the switch and have it function as it did when it was configured with earlier software that allows overlapping subnets. This occurs because the overlapping IP address has been removed and is not visible to the switch. To resolve this:

If you attempt to remove the overlapping subnet from the VLAN, the switch displays an error message similar to:

```
The IP address ip address is not configured on this VLAN
```

This occurs because the overlapping IP address has been removed and is not visible to the switch. To resolve this:

- Enter the `show ip` command to determine which addresses are visible to the switch.
- Remove the erroneous IP addresses from the config file by entering the `no ip address` command to remove all the IP addresses from the specific VLAN. Be sure to document the other valid IP addresses on that VLAN so they can be restored after removing the erroneous IP addresses from the config file.

If you go back to a software version prior to K.15.09 before removing the overlapping IP address, the prior software version enables the overlapping IP subnet.

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:

```
HP 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

To Change 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.

Inbound Tagged Packets

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

Untagged Packet Forwarding

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 7 Untagged VLAN operation

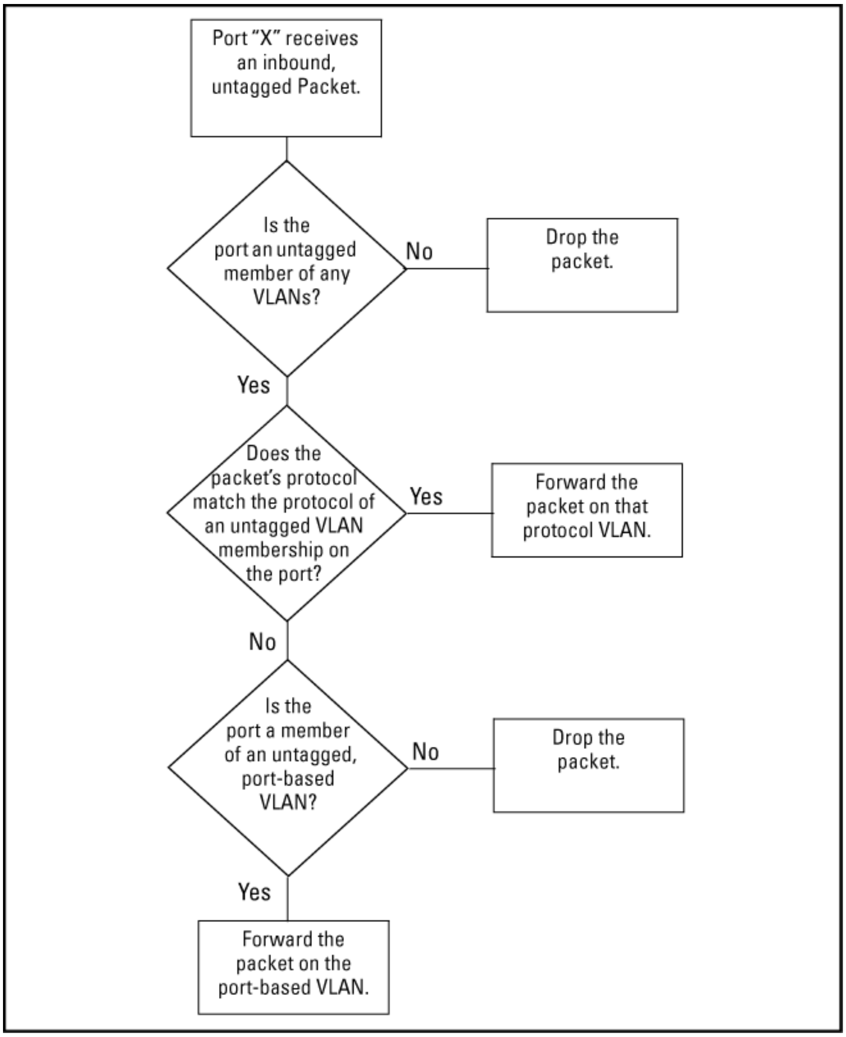

Port "X" receives an inbound, untagged Packet.
Is the port an untagged member of any VLANs?
No
Drop the packet.
Yes
Does the packet's protocol match the protocol of an untagged VLAN membership on the port?
Yes
Forward the packet on that protocol VLAN.
No
Is the port a member of an untagged, port-based 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 8 Tagged VLAN operation

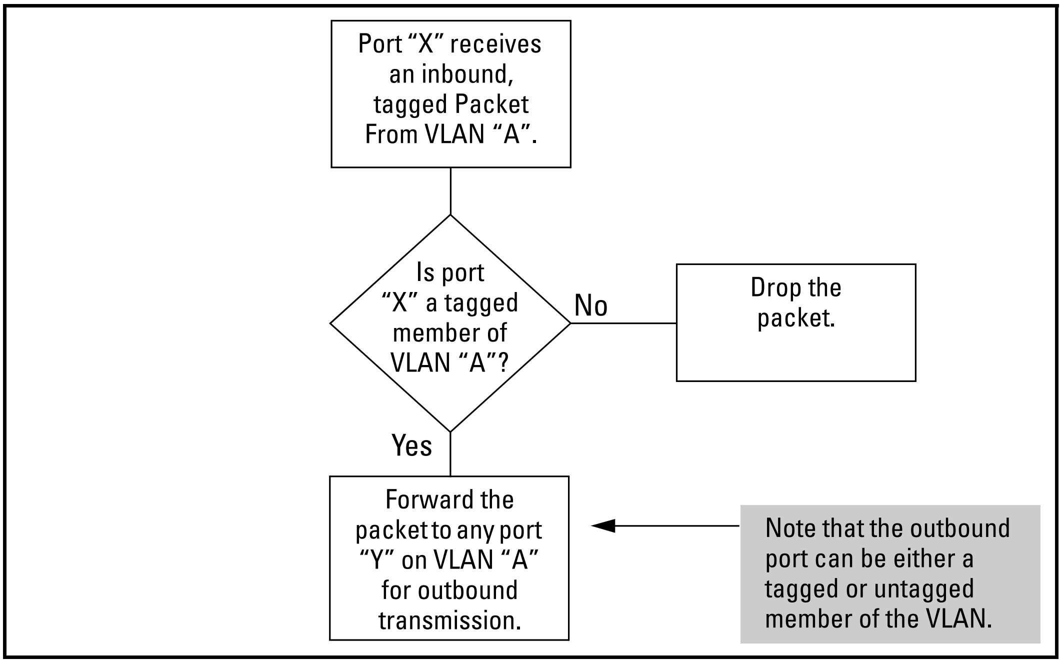
Port "X" receives an inbound, tagged Packet From VLAN "A".
Is port "X" a tagged member of VLAN "A"?
No
Drop the packet.
Yes
Forward the packet to any port "Y" on VLAN "A" for outbound transmission.
Note that the outbound port can be either a tagged or untagged member of the VLAN.

See also "Multiple VLAN considerations" (page 52).

△ **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. HP 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.

Example 39 VLANs using ports from the same port-bank for each VLAN

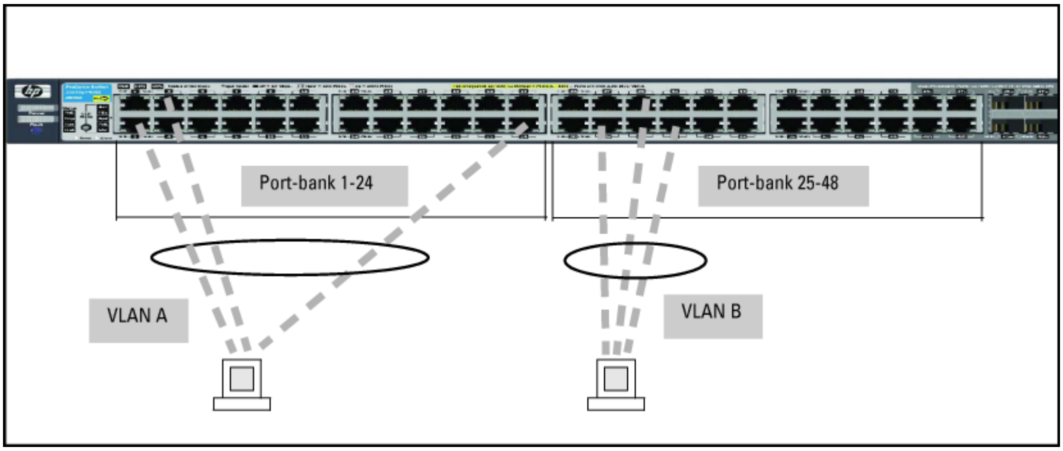

Multiple VLAN considerations

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 folloiwng table illustrates the functional difference between the two database types.

Table 4 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.		

Table 5 Forwarding database structure for managed HP switches

Multiple forwarding databases	Single forwarding database
Series 8200zl switches	Switch 1600M/2400M/2424M
Switch 6600	Switch 4000M/8000M
Series 6400cl switches	Series 2500 switches
Switch 6200yl	Switch 2000
Switch 6108	Switch 800T
Series 5400zl switches	
Series 5300xl switches	
Series 4200vl switches	
Series 4100gl switches	
Series 3800 switches	
Series 3500 switches	
Series 3500yl switches	
Series 3400cl switches	
Switch 2810	
Series 2800 switches	
Series 2600/2600-PWR switches	
Series 2510 switches	
*To determine whether other vendors' devices use single-forwarding or multiple-forwarding database architectures, see the documentation provided for those devices.	

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.

If (1) two types of switches connect through multiple ports or trunks belonging to different VLANs, and (2) routing is enabled on the switch having the multiple forwarding database then, on the switch having the single forwarding database, the port and VLAN record it maintained for the connected multiple-forwarding-database switch on the switch having the single forwarding database, maintains for the connected multiple-forwarding-database switch can frequently change. This causes poor performance and the appearance of an intermittent or broken connection. See "Example of correcting an unsupported configuration" (page 35).

802.1Q VLAN tagging

- The switch requires VLAN tagging on a given port if more than one VLAN of the same type uses the port. 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 under any circumstances.

- The switch requires VLAN tagging on a given port if the port will be receiving inbound, tagged VLAN traffic that should be forwarded. Even if the port belongs to only one VLAN, it forwards inbound tagged traffic only if it is a tagged member of that VLAN.
- 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.

Examples

Example 40 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.

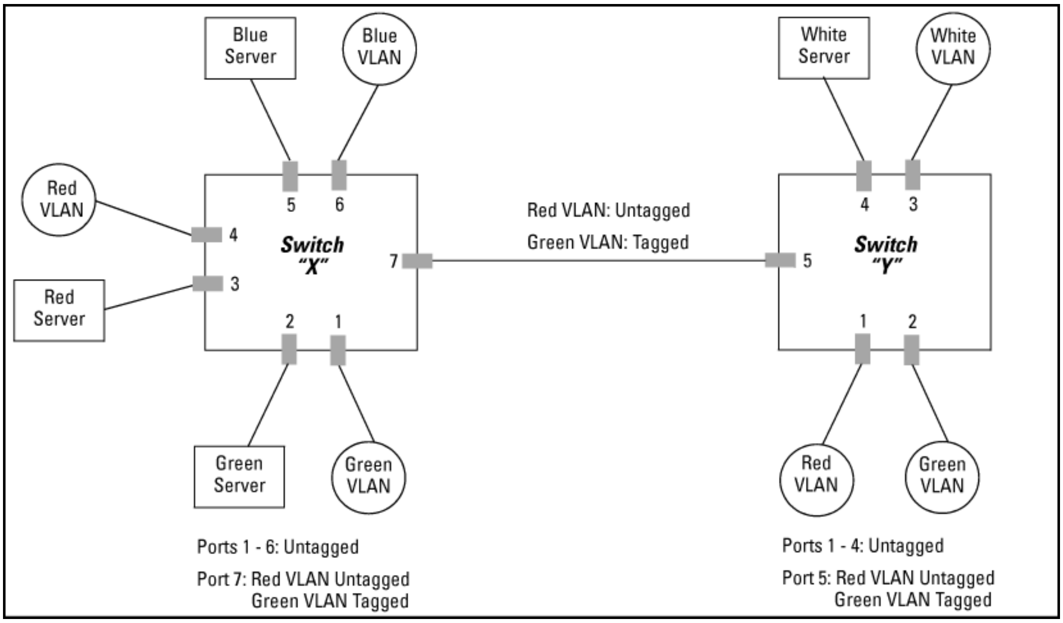

In switch X:

- 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 Example 41 "VLAN ID numbers assigned in the VLAN names screen", the Red VLAN must be untagged on port X7 and Y5 and the Green VLAN must be tagged on port X7 and Y5, or vice-versa.

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 in which it is 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.

Example 41 VLAN ID numbers assigned in the VLAN names screen

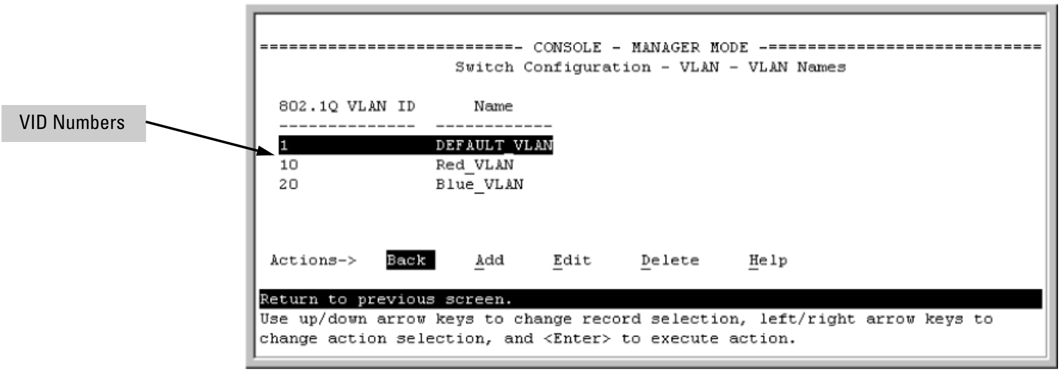

```
=========================- CONSOLE - MANAGER MODE -============================
                    Switch Configuration - VLAN - VLAN Names

  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.
```

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. See above.	A port can be a tagged member of any protocol-based VLAN. See above.
Note: 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 Operating Rules" (page 48):

- "Inbound Tagged Packets"
- "Untagged Packet Forwarding" and Figure 7 (page 51)
- "Tagged Packet Forwarding" and Figure 8 (page 52)

Example 42 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.

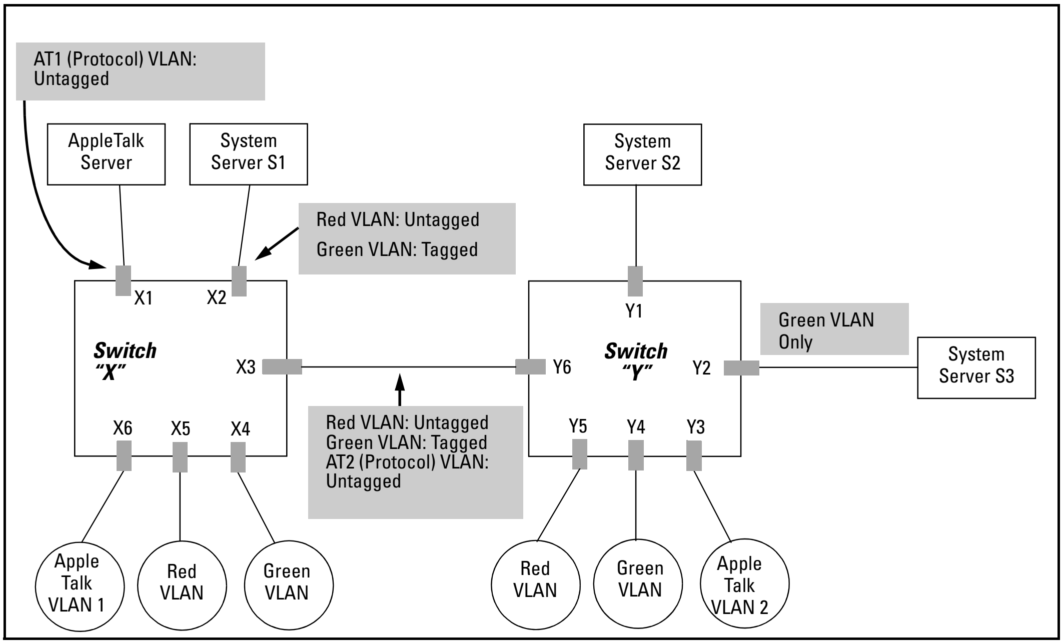

- The VLANs assigned to ports X4 - X6 and Y2 - Y5 can all be untagged because there is only one VLAN assigned per port.
- Port X1 has two AppleTalk VLANs assigned, which means that one VLAN assigned to this port can be untagged and the other must be tagged.
- Ports X2 and Y1 have two port-based VLANs 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.

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[1]	No[1]	Y1	No[1]	No[1]	Untagged	Tagged
X2	No[1]	No[1]	Untagged	Tagged	Y2	No[1]	No[1]	No[1]	Untagged
X3	No[1]	Untagged	Untagged	Tagged	Y3	No[1]	Untagged	No[1]	No[1]
X4	No[1]	No[1]	No[1]	Untagged	Y4	No[1]	No[1]	No[1]	Untagged
X5	No[1]	No[1]	Untagged	No[1]	Y5	No[1]	No[1]	Untagged	No[1]
X6	Untagged	No[1]	No[1]	No[1]	Y6	No	Untagged	Untagged	Tagged

[1] No means 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.

NOTE: VLAN configurations onports 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."

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 static VLAN per-port settings (CLI)" (page 30).

The pimary 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.

To change the Primary VLAN configuration, see "Changing VLAN support settings (Menu)" (page 26).

The secure Management VLAN

Configuring a secure Management VLAN creates an isolated network for managing the HP switches that support this feature. Access to a secure Management VLAN and the switch's management functions (Menu and CLI), 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.

Examples

Example 43 Potential security breaches in a network

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

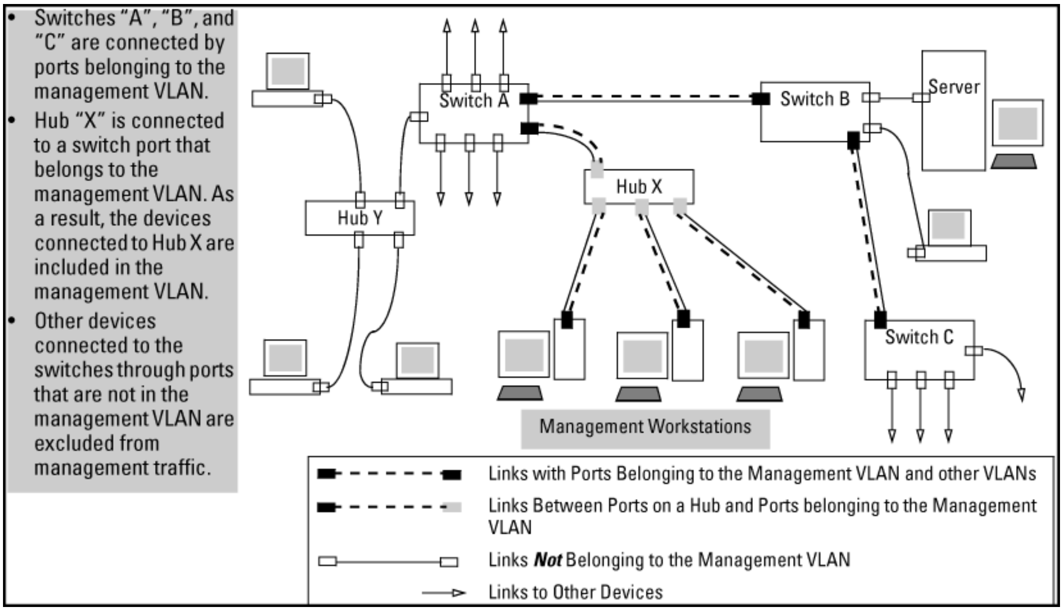

- Switches "A", "B", and "C" are connected by ports belonging to the management VLAN.
- Hub "X" is connected to a switch port that belongs to the management VLAN. As a result, the devices connected to Hub X are included in the management VLAN.
- Other devices connected to the switches through ports that 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

Example 44 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.

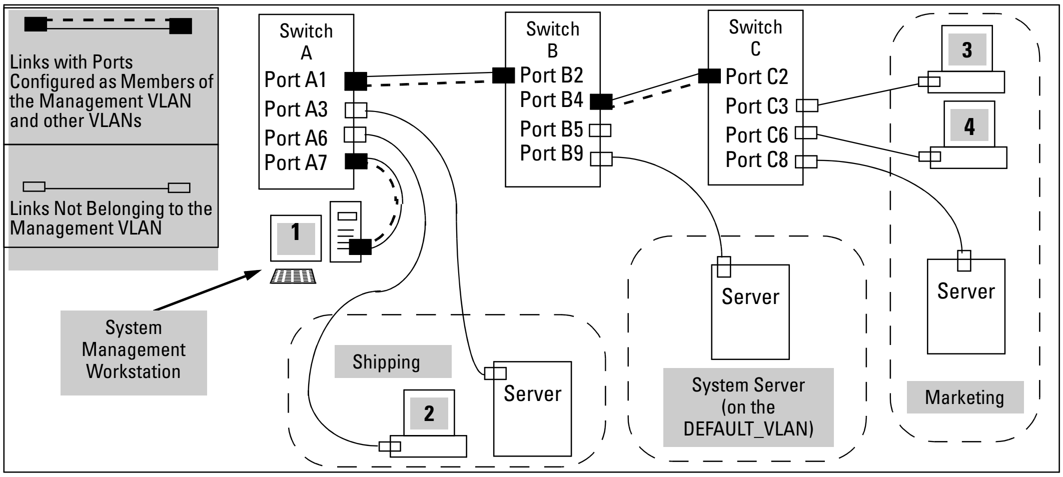
Links with Ports Configured as Members of the Management VLAN and other VLANs
Links Not Belonging to the Management VLAN
System Management Workstation
1
Switch A
Port A1
Port A3
Port A6
Port A7
Switch B
Port B2
Port B4
Port B5
Port B9
Switch C
Port C2
Port C3
Port C6
Port C8
3
4
Server
Server
Server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

Table 6 VLAN membership in Example 44 "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)" (page 37) 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.
- If you implement a Management VLAN in a switch mesh environment, all meshed ports on the switch will be members of 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. This can also occur where meshing is configured and the Management VLAN is configured on a separate link.
- 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 *Management and Configuration Guide* for your switch.

Example 45 Inadvertently blocking a Management VLAN link by implementing spanning tree

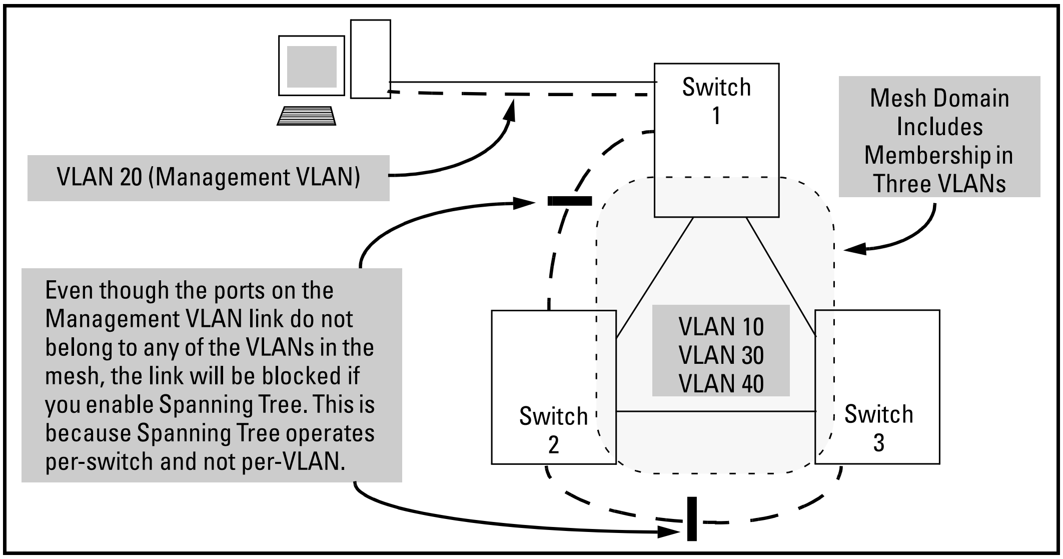
Switch 1
Mesh Domain Includes Membership in Three VLANs
VLAN 20 (Management VLAN)
Even though the ports on the Management VLAN link do not belong to any of the VLANs in the mesh, the link will be blocked if you enable Spanning Tree. This is because Spanning Tree operates per-switch and not per-VLAN.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

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.

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 on a per-VLAN basis.

For single-instance operation, this means that 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” (page 78).

Note that Spanning Tree operates differently in different devices. For example, in the (obsolete, non-802.1Q) HP Switch 2000 and the HP Switch 800T, Spanning Tree operates on a per-VLAN basis, 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. Likewise, you can assign an IP address to the VLAN interface, and 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, such as the Switch 4000M, some cabling restrictions apply. For more on this topic, see "Multiple VLAN considerations" (page 52).

Port trunks

When assigning a port trunk to a VLAN, all ports in the trunk are automatically assigned to the same VLAN. Do not split trunk members across multiple VLANs. A port trunk is tagged, untagged, or excluded from a VLAN in the same way as individual, untrunked ports.

Port monitoring

If you designate a port on the switch for network monitoring, this port will appear in the PortVLAN 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 *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 *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, note that 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.

- Prior to deleting a static VLAN, you must first re-assign 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) " (page 27).
- Protocol-based VLANs, port-based VLANs and LLDP radio port VLANs cannot run concurrently with RPVST+.

Migrating Layer 3 VLANs using VLAN MAC configuration

HP 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 HP routing switch.

VLAN MAC address reconfiguration

HP 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 allows you to 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 HP routing switch.

Packets transmitted from the routing switch (packets originating from the router and forwarded packets) use the original HP 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 an HP Switch, you can swap the physical port of a router to the HP 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 that is configured for each VLAN interface.

Outgoing VLAN traffic

This uses the MAC address of the HP Sswitch 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 "gracious" 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

NOTE: The Virtual Router Redundancy Protocol (VRRP) is not supported on VLAN interfaces on which the MAC address for incoming traffic has been reconfigured.

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 HP 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 stream 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 HP 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 HP Switch unicast MAC address in the destination field. This MAC address is assigned to the HP Switch and is not used by other non-HP 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 HP 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.

Operating notes

- The `ip-recv-mac-address` command allows you to configure only one MAC address for a specified VLAN. If you re-enter the command to configure another MAC address, the previously configured MAC address is overwritten.
- Enter the `no` form of the command to remove a configured MAC address and restore the default MAC address of the HP switch.
- When you configure a VLAN MAC address, you may also specify a heartbeat interval. The `interval` *`seconds`* parameter is optional.

- After you configure a VLAN MAC address:
  - IP router and MAC ARP replies to other VLAN devices contain the user-defined MAC address as the Ethernet sender hardware address.
  - Outbound VLAN traffic contains the HP Switch MAC address, not the configured MAC address, as the source MAC address in packet headers.
- Immediately after you configure a VLAN MAC address or remove a configured MAC address, a gratuitous ARP message is broadcast on the connected segment to announce the change of the IP-to-MAC address binding to all connected IP-based equipment.
- A configured VLAN MAC address supports proxy ARP and gracious ARP.
- A new MIB variable, `ifRcvAddressTable`, is introduced to support VLAN MAC configuration.
- You cannot configure a VLAN MAC address using the WebAgent or menu interface. You must use the CLI.
- VRRP is not supported on a VLAN interface with a user-configured MAC address.

Example

Example 46 Configuring a MAC address

The following example shows how to configure a MAC address on VLAN 101.

```
HP Switch# configure terminal
HP Switch(config)# vlan 101
HP Switch(vlan-101)# ip-recv-mac-address 0060b0-e9a200 interval 100
```

2 GVRP

Command syntax	Description	Default	CLI reference page	Menu reference page
`show gvrp`	Shows whether GVRP is disabled, and the current settings for the maximum number of VLANs and the current Primary VLAN.		68	69
`show vlans`	Lists static and dynamic VLANs on a GVRP-enabled switch.		71	

`gvrp`
`no gvrp`
	Enables or disables GVRP on the switch	Disabled	70	69

`interface` *`port-list`* `unknown-vlans [` *`learn`* `|` *`block`* `|` *`disable`* `]`	Controls how individual ports handle advertisements for new VLANs	Learn	70	69
`static` *`dynamic-vlan-id`*	Converts a dynamic VLAN to a static VLAN		72	

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.

Example 47 Forwarding advertisements and dynamic joining

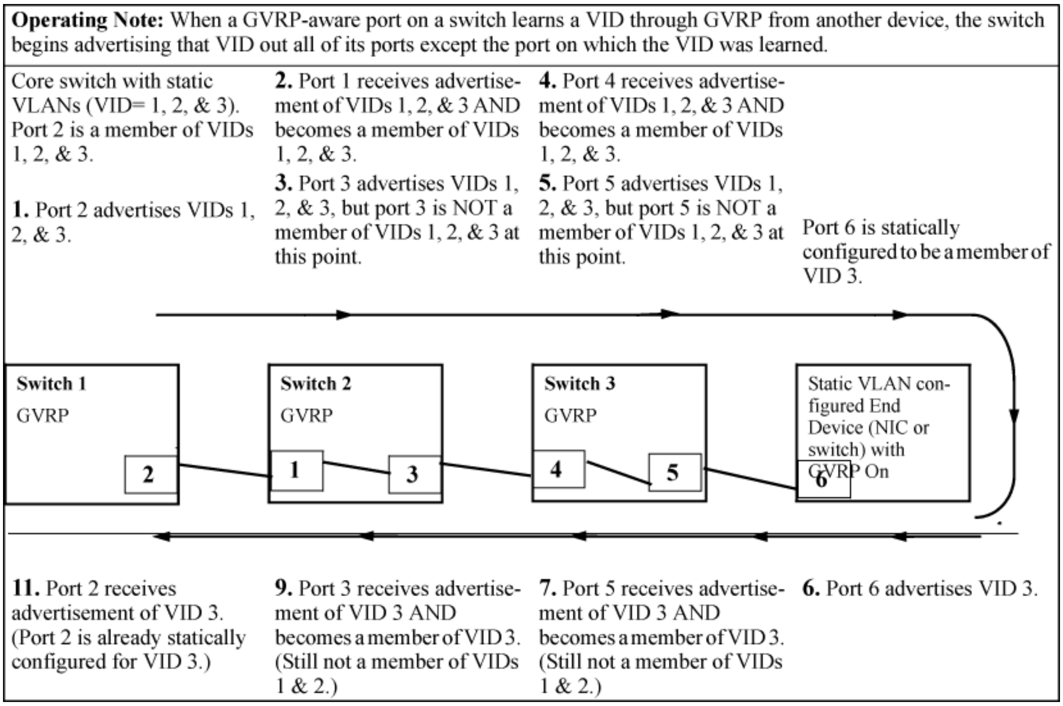

Operating Note: When a GVRP-aware port on a switch learns a VID through GVRP from another device, the switch begins advertising that VID out all of its ports except the port on which the VID was learned.

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.

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

9. 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.

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

If a static VLAN is configured on at least one port of a switch, 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:

1. Determine the VLAN topology required for each segment (broadcast domain) on the network.
2. Determine the VLANs that must be static and the VLANs that can be dynamically propagated.
3. Determine the devices on which static VLANs must be manually created in order to propagate VLANs throughout the segment.
4. Determine security boundaries and how the individual ports in the segment will handle dynamic VLAN advertisements. (See Table 7 (page 71) and Table 8 (page 76))
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 the static VLANs on the switches needed, along with the per-VLAN parameters (Tagged, Untagged, Auto, and Forbid—see Table 8 (page 76))on each port.
7. Dynamic VLANs will now appear automatically, according to the configuration options you have chosen.
8. Convert dynamic VLANs to static VLANs where dynamic VLANs need to become permanent.

Displaying the switch's current GVRP configuration (CLI)

Syntax:

```
show gvrp
```

Shows whether GVRP is disabled, along with the current settings for the maximum number of VLANs and the current Primary VLAN.

Examples

Example 48 Displaying GVRP status with GVRP disabled

```
HP Switch(config)# show gvrp

GVRP support

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

Example 49 Displaying GVRP status with GVRP enabled

This example shows the listing 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).

```
HP 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
```

Viewing and configuring GVRP (Menu)

1. From the Main Menu, select: **2. Switch Configuration...** —> **8. VLAN Menu...** —> **1. VLAN Support**

   Figure 9 The VLAN Support screen (default configuration)

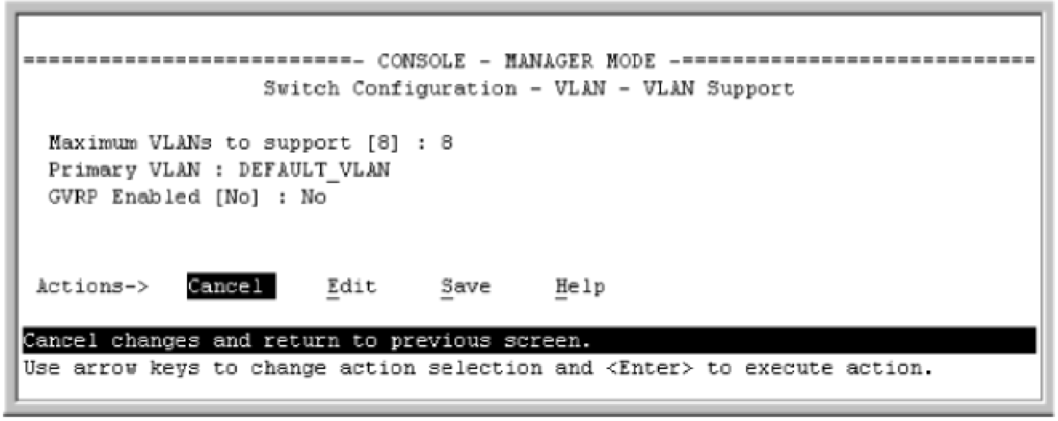

```
=========================- CONSOLE - MANAGER MODE -============================
                     Switch Configuration - VLAN - VLAN Support

  Maximum VLANs to support [8] : 8
  Primary VLAN : DEFAULT_VLAN
  GVRP Enabled [No] : No


 Actions->  Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

2. Do the following to enable GVRP and display the Unknown VLAN fields:
   a. Press **E** (for Edit).
   b. Use ↓ to move the cursor to the **GVRP Enabled** field.

c. Press the Space bar to select **Yes**.

d. Press ↓ again to display the **Unknown VLAN** fields.

Example 50 Default settings for handling advertisements

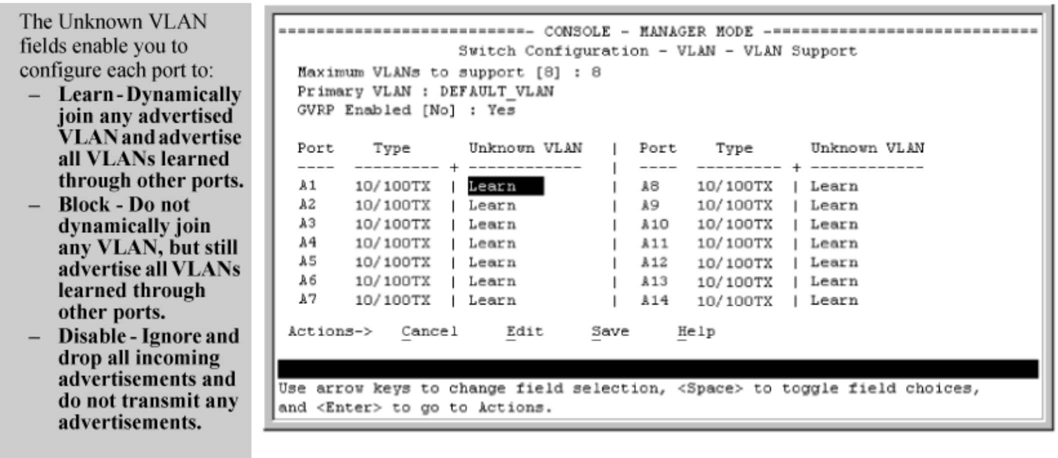

The Unknown VLAN fields enable you to configure each port to:

- **Learn - Dynamically join any advertised VLAN and advertise all VLANs learned through other ports.**
- **Block - Do not dynamically join any VLAN, but still advertise all VLANs learned through other ports.**
- **Disable - Ignore and drop all incoming advertisements and do not transmit any advertisements.**

```
=========================- CONSOLE - MANAGER MODE -===========================
                  Switch Configuration - VLAN - VLAN Support
  Maximum VLANs to support [8] : 8
  Primary VLAN : DEFAULT_VLAN
  GVRP Enabled [No] : Yes

  Port   Type      Unknown VLAN  |  Port   Type      Unknown VLAN
  ----   --------- + ----------- |  ----   --------- + ------------
  A1     10/100TX  | Learn       |  A8     10/100TX  | Learn
  A2     10/100TX  | Learn       |  A9     10/100TX  | Learn
  A3     10/100TX  | Learn       |  A10    10/100TX  | Learn
  A4     10/100TX  | Learn       |  A11    10/100TX  | Learn
  A5     10/100TX  | Learn       |  A12    10/100TX  | Learn
  A6     10/100TX  | Learn       |  A13    10/100TX  | Learn
  A7     10/100TX  | Learn       |  A14    10/100TX  | Learn

 Actions->   Cancel     Edit     Save     Help

Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

3. Use the arrow keys to select the port you want, and the Space bar to select the Unknown VLAN option for any ports you want to change.
4. When you finish making configuration changes, press **Enter**, then **S** (for Save) to save your changes to the Startup-Config file.

To view or configure static VLANs for GVRP operation, see "VLAN Operating Rules" (page 48)

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" (page 54).

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 this 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 in order to control how one or more ports handle advertisements. Use at either the Manager or interface context level for a port.

Example

Example 51 Changing the Unknown VLAN field

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

```
HP 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.

Example

Example 52 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:

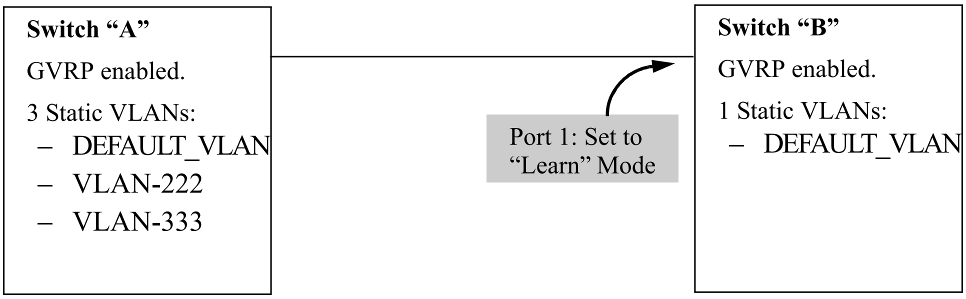

The `show vlans` command lists the dynamic (and static) VLANs in switch B after it has learned and joined VLAN-222 and VLAN-333.

```
Switch-B> show vlans

 Status and Counters - VLAN Information

   VLAN support : Yes
   Maximum VLANs to support : 8
   Primary VLAN : DEFAULT_VLAN

   802.1Q 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 dynamic-vlan-id
```

Example 53 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.

```
HP Switch(config)# static 333
```

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 chances for errors in VLAN configuration 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 if desired. Also GVRP can 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, and in this guide a GVRP BPDU is termed an *advertisement*. On a switch, advertisements are sent outbound from ports to the devices directly connected to those ports.

GVRP operating notes

- A dynamic VLAN must be converted to a static VLAN before it can have an IP address.
- On the switches covered in this guide, GVRP can be enabled only if `max vlans` is set to no more than 256 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 Menu interface, click on 2. Switch Configuration … | 8. VLAN Menu | 1. VLAN Support. 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 not GVRP-aware. This is because a half-duplex repeater or a switch that is not GVRP-aware will flood the GVRP (multicast) advertisement packets out all ports.
- Rebooting a switch on which adynamic 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.
- While GVRP is enabled on the switch, you cannot apply any ACLs to VLANs configured on the same switch.
- GVRP (GARP VLAN Registration Protocol) is an application of GARP (Generic Attribute Registration Protocol.) It is defined in the IEEE 802.1Q standard, and GARP is defined in the IEEE 802.1D-1998 standard.
- GVRP cannot run concurrently with RPVST+.

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.

Example 54 GVRP operation

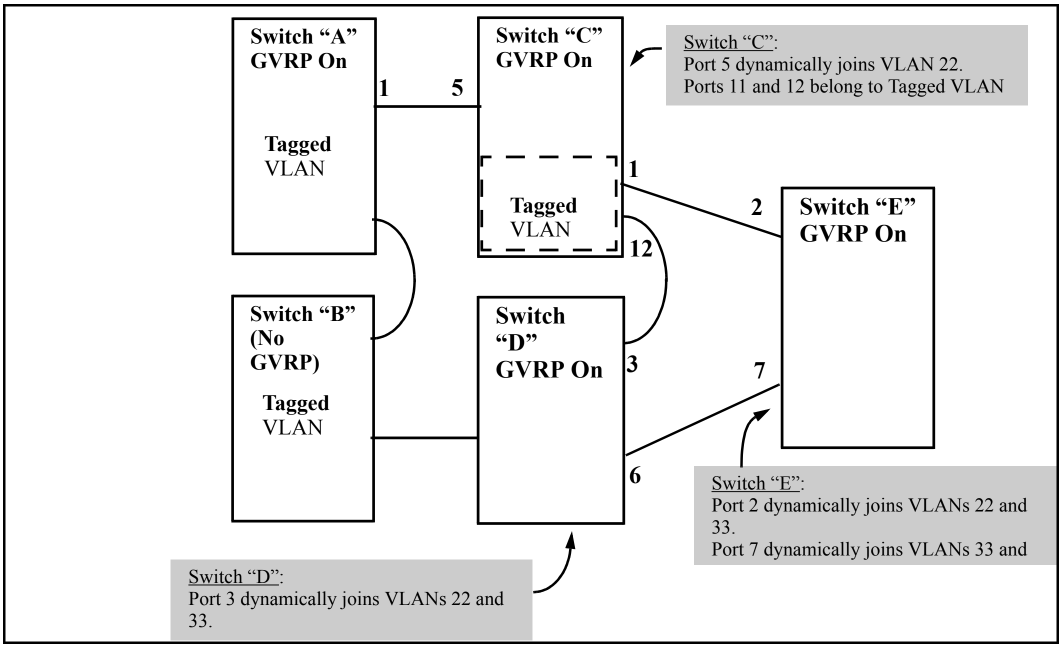
Switch “A”
GVRP On
Tagged
VLAN
1
5
Switch “C”
GVRP On
Tagged
VLAN
1
12
Switch “C”:
Port 5 dynamically joins VLAN 22.
Ports 11 and 12 belong to Tagged VLAN
2
Switch “E”
GVRP On
Switch “B”
(No
GVRP)
Tagged
VLAN
Switch
“D”
GVRP On
3
7
6
Switch “E”:
Port 2 dynamically joins VLANs 22 and
33.
Port 7 dynamically 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 “Per-port options for dynamic VLAN advertising and joining” (page 75).
- 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. Note that 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.

Example

Example 55 GVRP unknown VLAN settings

Suppose that in Example 54 (page 74), 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 and the menu interface VLAN Support screen show a switch's current GVRP configuration, including the Unknown VLAN settings.

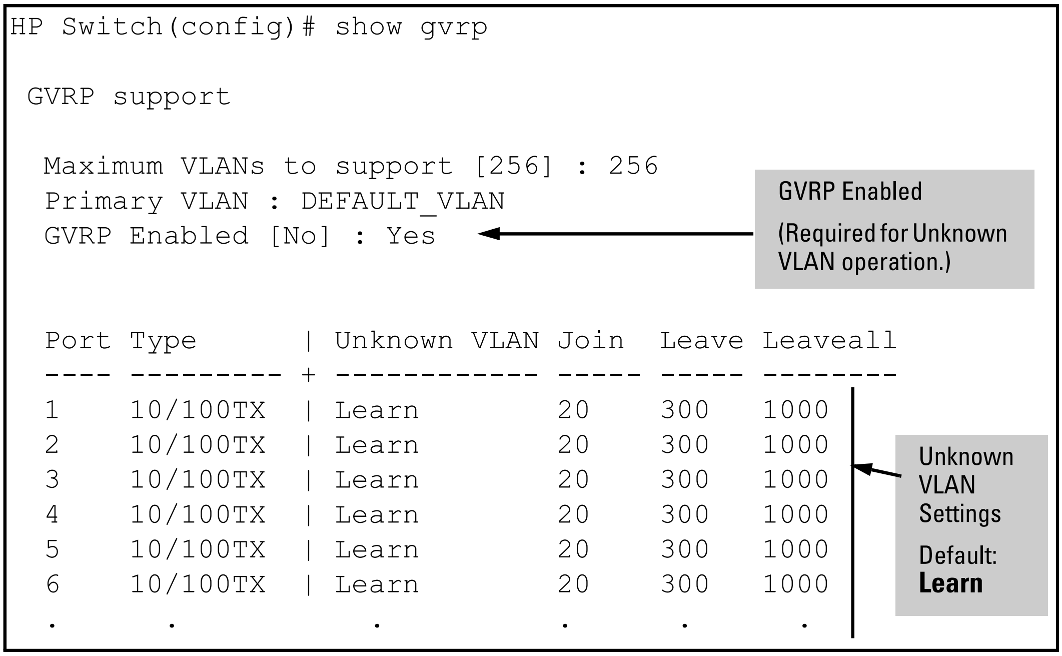

```
HP Switch(config)# show gvrp

 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

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 Table 8 (page 76).

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 Table 8 (page 76).

Table 8 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. • Will not join any advertised VLANs. • Will not advertise VLANs.	The port: • Will not become a member of the specified VLAN. • Will ignore GVRP PDUs. • Will not join any dynamic VLANs. • Will not advertise VLANs.	The port: • Will not become a member of this VLAN. • Will ignore GVRP PDUs. • 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 static VLAN per-port settings (CLI)" (page 30) (for the CLI) or "Adding or changing a VLAN port assignment (Menu)" (page 33) (for the menu).

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 Table 8 (page 76), 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, using either the menu interface or the CLI.

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

GVRP and VLAN access control

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.

3 Multiple instance spanning tree operation

Command syntax	Description	Default value	CLI reference page

`spanning-tree mode mstp`
`spanning-tree clear-debug-counters`
	Specifies that spanning tree will run in MSTP mode		86

`[no] spanning-tree config-name` *`ascii-string`*	Resets the configuration name of the MST region in which the switch resides	A text string using the switch's MAC address	86
`spanning-tree config-revision` *`revision-number`*	Sets the revision number designated for the MST region in which you want the switch to reside	0	87
`spanning-tree force-version [` *`stp-compatible`* `|` *`rstp-operation`* `|` *`mstp-operation`* `]`	Sets the spanning tree compatibility mode		87
`spanning-tree forward-delay`	Sets the time in seconds the switch waits between transitioning from listening to learning and from learning to forwarding states	15	88
`[no]spanning-tree legacy-mode`	Forces spanning tree to operate in legacy (802.!D) mode	Native mode: MSTP	88
`spanning-tree legacy-path-cost`	Forces spanning tree to operate with legacy (802.!D) path cost values	802.1t	88
`spanning-tree hello-time` *`1..10`*	Sets the time in seconds between transmissions of BPDUs for all ports on the switch configured with the Global option. (the default)	2	88

Command syntax	Description	Default value	CLI reference page
`spanning-tree max-hops` *`hop-count`*	Resets the number of hops allowed for BPDUs in an MST region	20	89
`spanning-tree maximum age`	Sets the maximum age (in seconds) for received STP information before it is discarded	20	89
`spanning-tree pending [ apply |` *`config-name`* `|` *`config-revision`* `|` *`instance`* `| reset ]`	Manipulates the pending MSTP configuration		89
`spanning-tree priority` *`priority-multiplier`*	Sets the switch (bridge) priority for a region, which determines its priority as the spanning tree root switch		90
[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	Disabled	90
[no] `spanning-tree` *`port-list`* `admin-edge-port`	Allows specified port(s) to transition immediately to a forwarding state	Disabled	91
[no] `spanning-tree` *`port-list`* `auto-edge-port`	Supports the automatic identification of edge ports	Enabled	91
`spanning-tree` *`port-list hello-time`* `[ global |` *`1 - 10`* `]`	Specifies the interval (in seconds) between periodic BPDU transmissions by the designated ports	2	91
`spanning-tree` *`port-list`* `mcheck`	Forces a port to send RST/MST BPDUs for 3 seconds		92

Command syntax	Description	Default value	CLI reference page
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	Auto	92
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	92
spanning-tree *port-list*priority *priority-multiplier*	Determines the priority of specified port(s) for use in forwarding		93
spanning-tree *port-list* root-guard	Enables root guard on specified port(s)	Disabled	93
spanning-tree *port-list* tcn-guard	Causes specified port(s) to stop propagating received topology change notifications and topology changes to other ports	Disabled	93
[no] spanning-tree [*port-list* | all] bpdu-filter	Enables or disables BPDU filtering	Disabled	94
spanning-tree show *port* configuration	Displays BPDU filtering information		94
[no]spanning-tree *port-list* bpdu-protection	Enables or disables BPDU protection	Disabled	95
[no] spanning-tree *port-list* bpdu-protection-timeout *timeout*	Sets the duration of time (in seconds) when protected ports receiving unauthorized BPDUs will remain disabled	0	95
[no] spanning-tree trap errant-bpdu	Enables and disables the	Disabled	95

Command syntax	Description	Default value	CLI reference page
	sending of errant BPDU traps		
show spanning-tree bpdu-protection	Displays BPDU protection status		96
[no] spanning-tree *port-list* pvst-protection	Enables and disables PVST protection	Disabled	96
[no] spanning-tree *port-list pvst-filter*	Enables and disables PVST filtering	Disabled	97
[no] spanning-tree bpdu-protection-timeout *timeout*	Re-enables ports manually	0	97
show spanning-tree pvst-filter	Displays which ports have PVST filtering enabled		98
show spanning-tree pvst-protection	Displays which ports have PVST protection enabled		98
[no] spanning-tree instance *1..16* vlan *vid* [*vid...vid*]	Configures MSTP instance parameters		98
spanning-tree instance *1..16* priority *priority-multiplier*	Sets the bridge priority for an instance		99
[no] spanning-tree instance *1..16* vlan *vid* [*vid...vid*]	Creates a new MST instance (MSTI) and moves the specified VLANs from the IST to the MSTI	instance (MSTPI): none	100
spanning-tree instance *1..16 port-list* path-cost [auto | *1..200000000*]	Assigns an individual port cost for the specified MST instance	auto	100
spanning-tree instance *1..16 port-list* priority *priority-multiplier*	Sets the switch (bridge) priority for the specified ports in the specified MST instance		100
spanning-tree *port-list* priority *priority-multiplier*	Sets the switch (bridge) priority for the specified ports for the IST (Instance 0) of the region in	priority: 32768 (multiplier: 8)	101

Command syntax	Description	Default value	CLI reference page
	which the switch resides		
[no] spanning-tree	Enables or disables MSTP spanning tree operation	Disabled	102
[no] spanning-tree pending [apply | *config-name* | *config-revision* | *instance* | reset]	Exchanges the currently active MSTP configuration with the current pending MSTP configuration		102
[no] spanning-tree instance *1..16* vlan *vid* [*vid...vid*]	Pre-configures VLANs in an MST instance		104
show spanning-tree	Displays MSTP statistics		107
show spanning-tree *port-list*			
show spanning-tree detail			
show spanning-tree *port-list*detail			
show spanning-tree instance [ist | *1..16*]			
show spanning-tree instance [ist | *1..16*] detail			
show spanning-tree *port-list* instance [ist | *1..16*] detail			
show spanning-tree config	Displays the MSTP configuration		110
show spanning-tree *port-list*config			
show spanning-tree config instance [ist | *1..16*]			
show spanning-tree *port-list* config instance [ist | *1..16*]			
show spanning-tree mst-config			
show spanning-tree pending [instance | mst-config] instance [ist | *1..16*]			
[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*]]	Configures loop protection	send-disable	114
show loop-protect *port-list*	Displays loop protection status		115
show spanning-tree root-history	Troubleshoots an MSTP configuration		120
show spanning-tree debug counters			
show spanning-tree debug-counters instance *instance-id*			
show spanning-tree debug-counters instance *instance-id* ports *port-list*			
[no] spanning-tree trap { errant-bpdu | loop-guard | new-root | root-guard }			

Overview

NOTE: For information on configuring RPVST+, see the *Rapid per-VLAN spanning tree operation* chapter in this guide.

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 (see "Configuring loop protection" (page 114)).

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.

Example

Suppose 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:

Example 56 A multiple spanning tree application

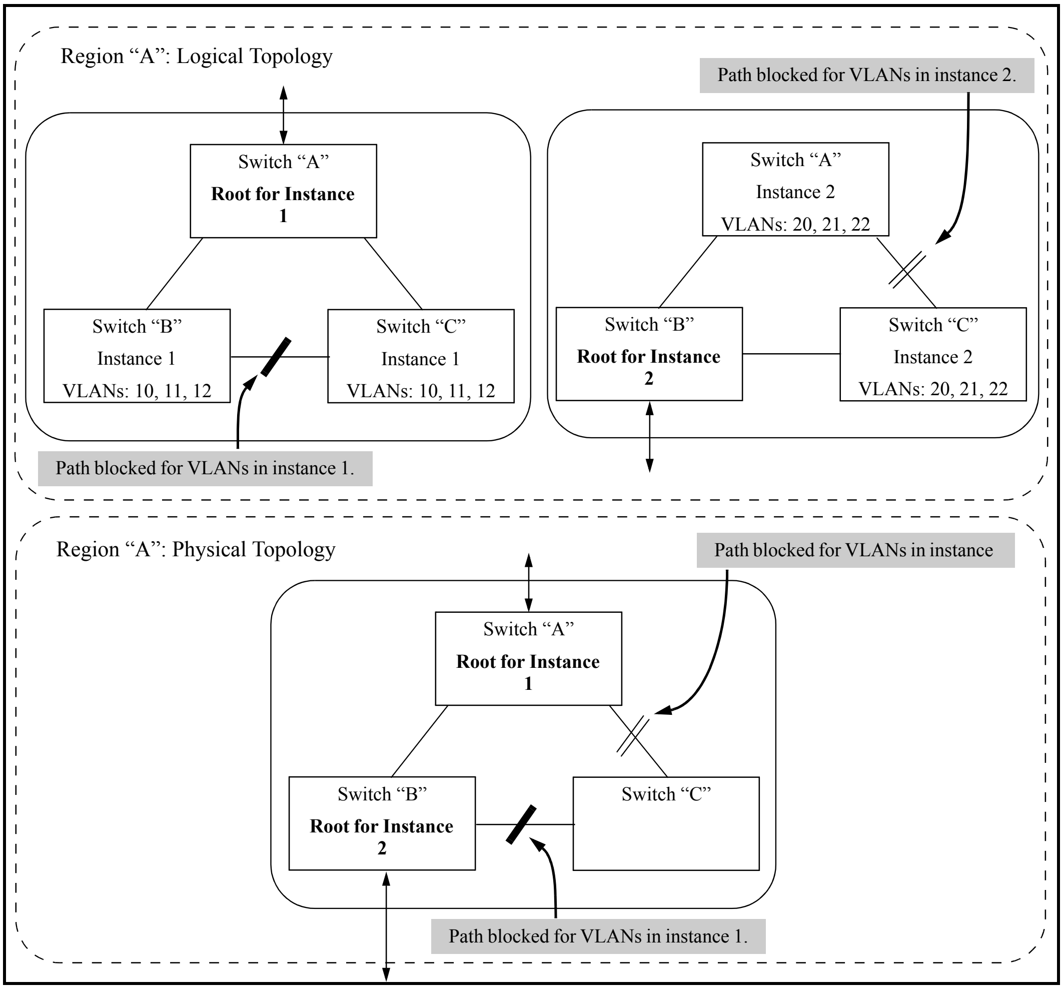

Planning an MSTP application

Before configuring MSTP, keep in mind the following tips and considerations:

- Ensure that the VLAN configuration in your network supports all of 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 inter-regional links
  - Any IST (Internal Spanning Tree) or MST 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" (page 133).)
- 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.

Configuring MSTP at a glance

The general steps for configuring MSTP via the CLI are:

1. Configure MSTP global parameters. This involves:
   - Selecting MSTP as the spanning tree mode:
     `spanning-tree mode mstp`
   - Clearing spanning tree debug counters:

     `spanning-tree clear-debug-counters`
   - Specifying required parameters for MST region identity:

     `Region Name:spanning-tree config-name`

     `Region Revision Number:spanning-tree config-revision`
   - Optionally, specifying MSTP parameter changes for region settings:

     HP 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`
   - 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. Note that inappropriate changes to these settings can result in severely degraded network performance. For this reason, HP 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. HP recommends that you use the default settings for these parameters and apply changes on a per-port basis only where a non-default 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 unmap the VLAN from the current instance, then add the VLAN to the other instance. (While the VLAN is unmapped 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. HP recommends that you apply changes on a per-port basis only where a non-default setting is clearly indicated by the circumstances of individual links. For example, you might want to set the path cost value for the ports used by a specific MST instance. Use the following command:

   ```
   spanning-tree instance 1..16 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. For configuring spanning tree settings on individual ports, see "Configuring MSTP per-port parameters" (page 91).

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:

```
[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" (page 85), 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 misordering in the switched LAN. To allow MSTP switches to support applications and protocols that may be sensitive to frame duplication and misordering, 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 transitioning from listening to learning and from learning to forwarding states.

Range: 4 - 30

Default: 15 seconds

Setting spanning tree to operate in 802. ID legacy mode

Syntax:

`[no] spanning-tree legacy-mode`

Forces spanning tree to operate in legacy (802.!D) 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. ID legacy path cost values

Syntax:

`spanning-tree legacy-path-cost`

Forces spanning tree to operate with legacy (802.!D) 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.

Note that 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). The ability to 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:

```
[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. Note that this command is designed to be used in conjunction with the `spanning-tree bpdu-filter` command (see "Configuring BPDU filtering" (page 94)) and the `bpdu-protection` command (see "Enabling and disabling BPDU protection" (page 95)).

`errant-bpdu`

Enables SNMP notification when an errant BPDU is received. Designed for use with BPDU filtering (See "Configuring BPDU filtering" (page 94)).

`loop-guard`

Enables SNMP notification when a loop guard inconsistency is detected. Designed for use with the Loop Guard option (See "STP loop guard" (page 116)).

`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. See "Denying a port the role of root port" (page 93).

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, HP recommends that you use the default settings for these parameters and apply changes on a per-port basis only where a non-default 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:

[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:

[no] `spanning-tree` *`port-list`* `auto-edge-port`

Supports 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: 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`

See the note (page 136) for information on compatibility with devices running 802.1D STP for the path cost values

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.

`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-tree port-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-tree port-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

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:

```
[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 (see "About BPDU protection" (page 136) for details).

△ **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).

Example

Example 57 Configuring BPDU filtering

To configure BPDU filtering on port a9, enter:

```
HP Switch(config)# spanning-tree a9 bpdu-filter
```

Viewing BPDU filtering

Syntax:

```
spanning-tree show port configuration
```

Displays the BPDU filter state.

Examples

Example 58 Displaying BPDU filter status using the `show spanning tree` command

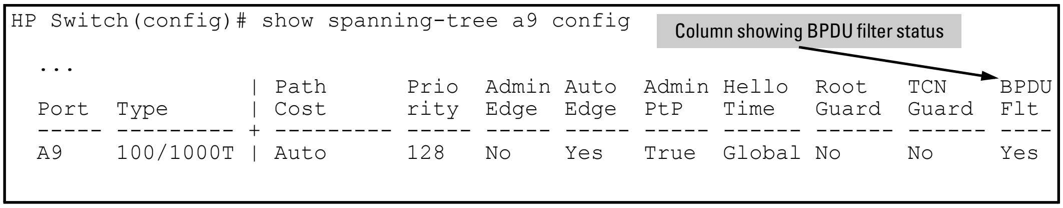

```
HP Switch(config)# show spanning-tree a9 config

  ...
                   | Path      Prio  Admin Auto  Admin Hello  Root   TCN    BPDU
  Port  Type       | Cost      rity  Edge  Edge  PtP   Time   Guard  Guard  Flt
  ----- ---------- + --------- ----- ----- ----- ----- ------ ------ ------ ----
  A9    100/1000T  | Auto      128   No    Yes   True  Global No     No     Yes
```

Column showing BPDU filter status

Example 59 Displaying BPDU filters using the `show configuration` command

This example shows how BPDU filters per port are displayed as separate entries of the spanning tree category within the configuration file.

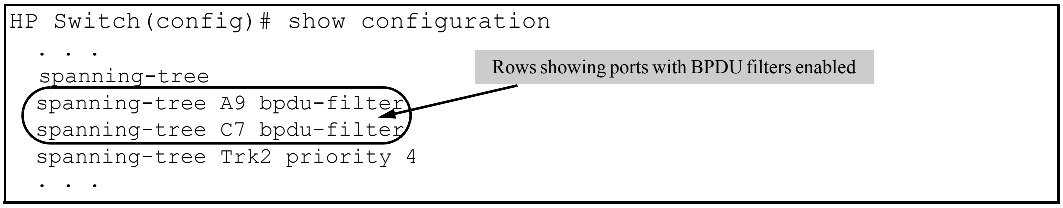

```
HP Switch(config)# show configuration
  . . .
  spanning-tree
  spanning-tree A9 bpdu-filter
  spanning-tree C7 bpdu-filter
  spanning-tree Trk2 priority 4
  . . .
```

Rows showing ports with BPDU filters enabled

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.

Example

Example 60 Configuring BPDU protection

To configure BPDU protection on ports 1 to 10 with SNMP traps enabled, enter:

```
HP Switch(config)# spanning-tree 1-10 bpdu protection
HP 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.

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 10 Displaying BPDU protection status

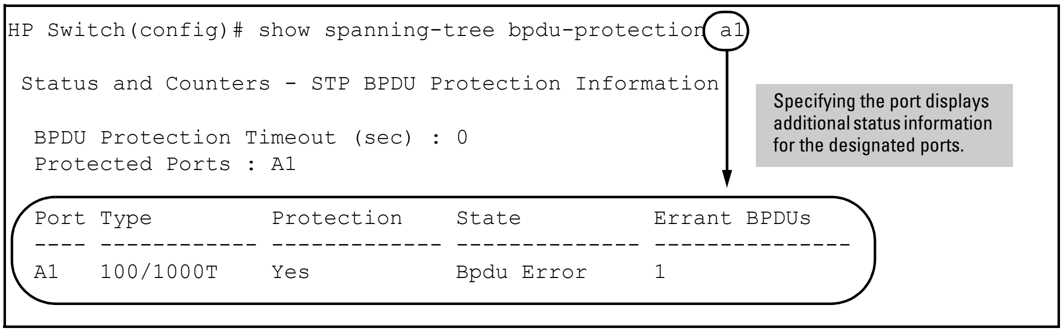

```
HP Switch(config)# show spanning-tree bpdu-protection a1

 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 11 Displaying BPDU filters using the `show configuration` command

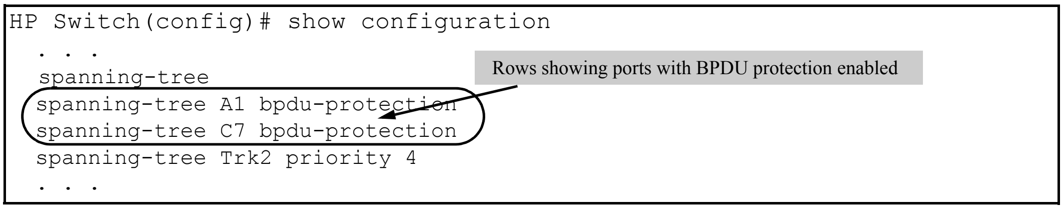

```
HP Switch(config)# show configuration
  . . .
  spanning-tree
  spanning-tree A1 bpdu-protection
  spanning-tree C7 bpdu-protection
  spanning-tree Trk2 priority 4
  . . .
```

Rows showing ports with BPDU protection enabled

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.

Example

Example 61 Enabling PVST protection

To enable the PVST protection feature on ports 4 through 8, enter:

```
HP Switch(config)# spanning-tree 4-8 pvst-protection
```

To disable the PVST protection feature on a port, for example, port 4, enter:

```
HP 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

Example

Example 62 Enabling PVST filtering on a port

```
HP 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:

```
HP Switch(config)# spanning-tree bpdu-protection-timeout 120
```

Displaying ports configured with PVST protection and filtering

Example 63 Displaying all ports with PVST protection enabled

```
HP 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
```

Example 64 Displaying all ports with PVST filtering enabled

```
HP 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

Example 65 Displaying if PVST protection is enabled (Yes)

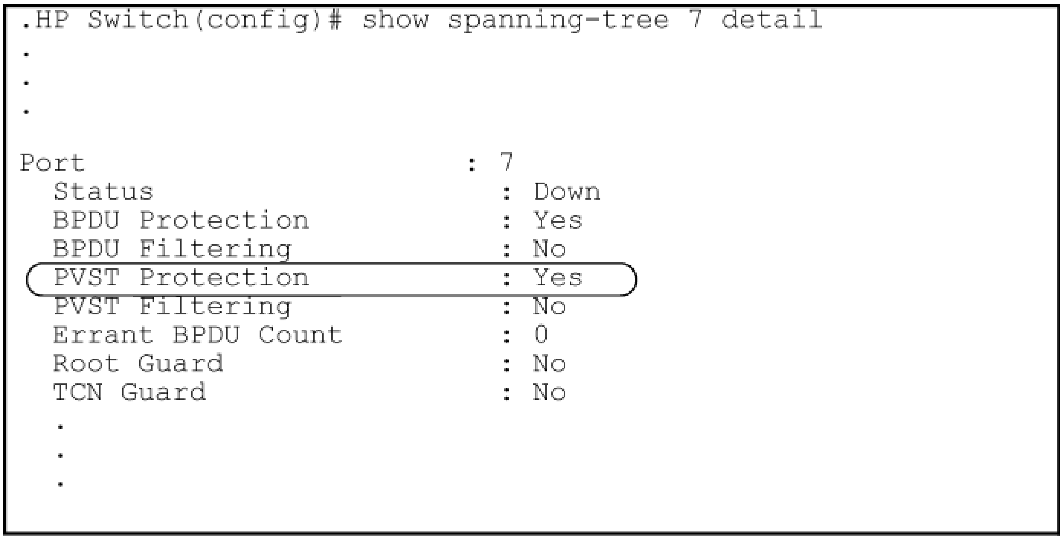

```
.HP Switch(config)# show spanning-tree 7 detail
.
.
.

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
  .
  .
  .
```

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:

[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, 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 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.

Configuring MST instance per-port parameters

Assigning a port cost for an MST instance

Syntax:

```
spanning-tree instance  1..16 port-list path-cost [ auto
| 1..200000000 ]
```

Assigns an individual port cost 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

Example

Example 66 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

Example

Example 67 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, ensure 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.

Disabliing 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 (see "Enabling an entire MST region at once or exchanging one region configuration for another" (page 102)).

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 1..16 vlan [ vid | vid-range ]`

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:

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.

Preconfiguring an MSTP regional topology

Starting in software release 13.*x*. *x* , 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:

1. Remove all MSTP mappings from the configuration file, then reconfigure the instance mapping after running the desired software version.
2. 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. See "Saving the current configuration before a software upgrade" (page 106).

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:

[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.

Example

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.

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

On HP 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. (See Example 68 “Mapping VLANs with the range option where all VLANs are included” and Example 69 “Mapping VLANs on other HP switches”)

Example 68 “Mapping VLANs with the range option where all VLANs are included” shows an example of an MSTP instance configured with the VLAN range option. All the VLANs are included in the instance whether they exist or not. Example 69 “Mapping VLANs on other HP switches” shows an example of an MSTP instance configured on another HP switch. Only VLANs 1, 5, and 7 are included in the instance.

Example 68 Mapping VLANs with the range option where all VLANs are included

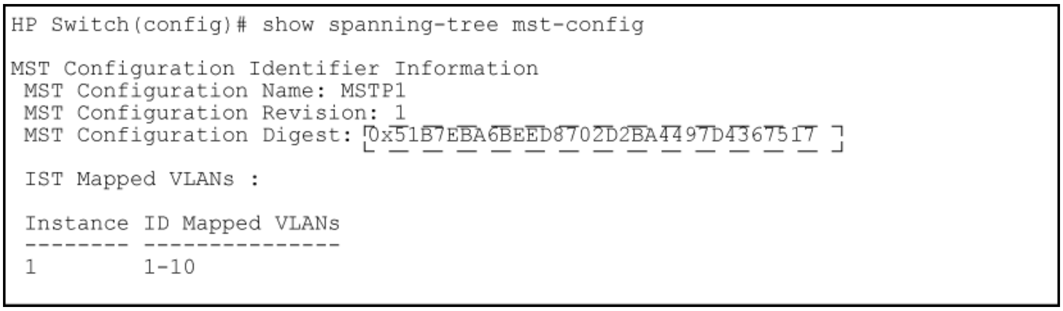

```
HP Switch(config)# show spanning-tree mst-config

MST Configuration Identifier Information
 MST Configuration Name: MSTP1
 MST Configuration Revision: 1
 MST Configuration Digest: 0x51B7EBA6BEED8702D2BA4497D43675l7

 IST Mapped VLANs :

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

The Configuration Digest value shown in Example 69 “Mapping VLANs on other HP switches” is not the same as in Example 68 “Mapping VLANs with the range option where all VLANs are included”, indicating that these switches do not operate in the same instance.

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

Example 69 Mapping VLANs on other HP switches

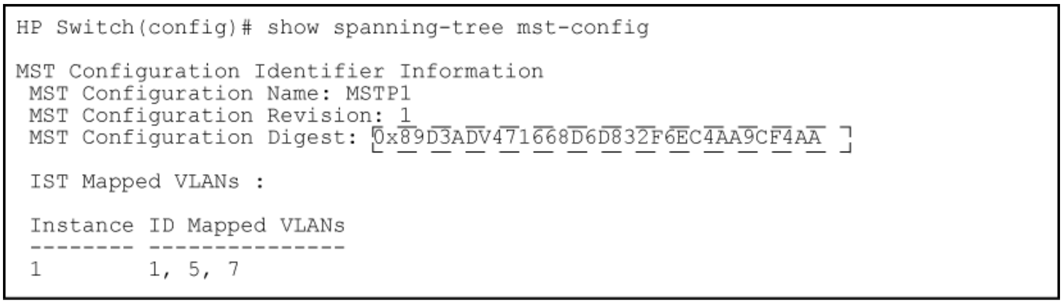

```
HP Switch(config)# show spanning-tree mst-config

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
```

See “Operating notes for the VLAN configuration enhancement” (page 135).

Saving the current configuration before a software upgrade

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

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

```
HP 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, enter this command:

```
HP 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.

```
HP 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.

   ```
   HP 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), enter this command:

   ```
   HP Switch(config)# boot system flash secondary config configK1243.cfg
   ```

   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.

   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.

Displaying 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.

Displaying 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`*

Dsplays 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`

Example 70 Displaying a common spanning tree status

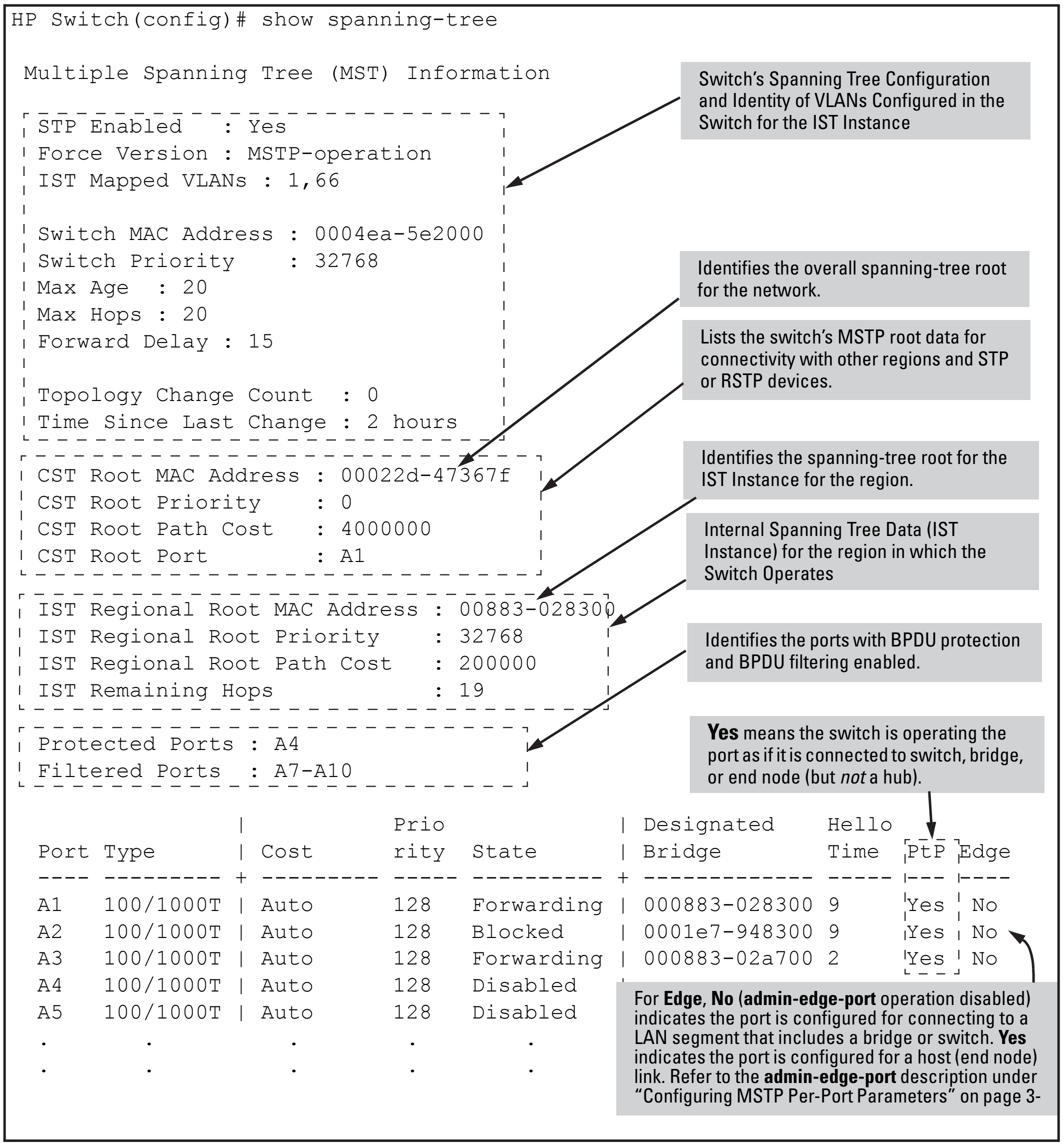

```
HP Switch(config)# show spanning-tree

 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-

Displaying 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.

Example

Example 71 Displaying port information

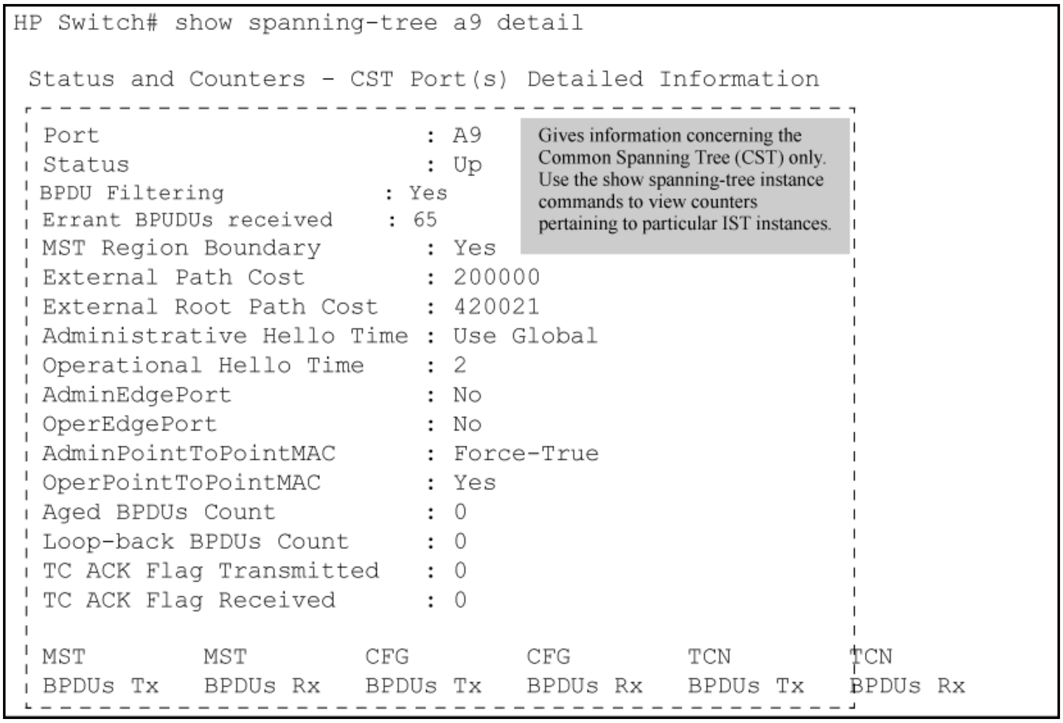

```
HP Switch# show spanning-tree a9 detail

 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.

Displaying 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.

Example

Example 72 Displaying 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.

```
HP 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
 .        .       .         .          .         .
```

Displaying the MSTP configuration

Displaying 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. For information on these parameters, see "Configuring MSTP per-port parameters" (page 91).

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`

Example

Example 73 Displaying the switch's global spanning tree configuration

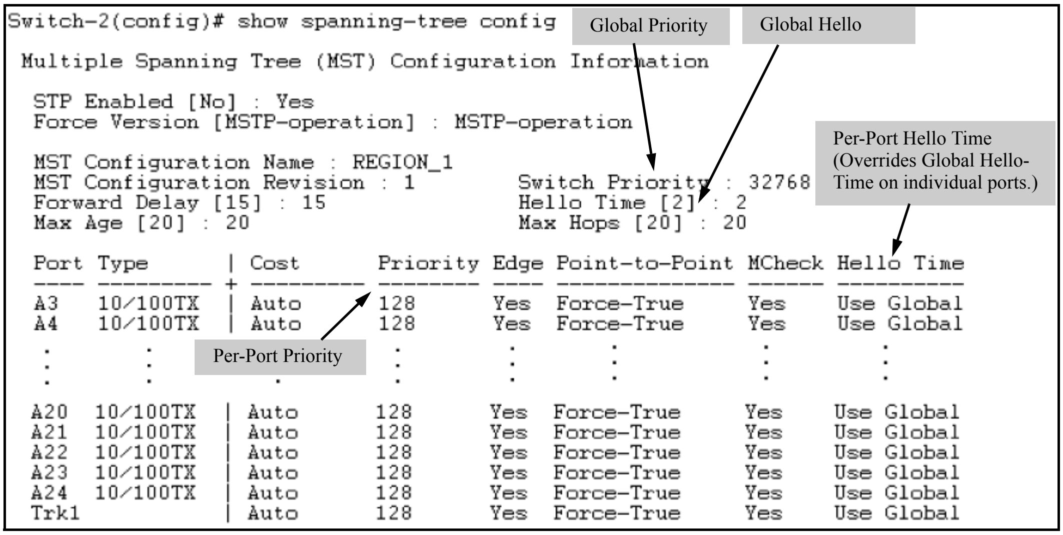

```
Switch-2(config)# show spanning-tree config

 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

Displaying 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 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.

Example

Example 74 Displaying port data

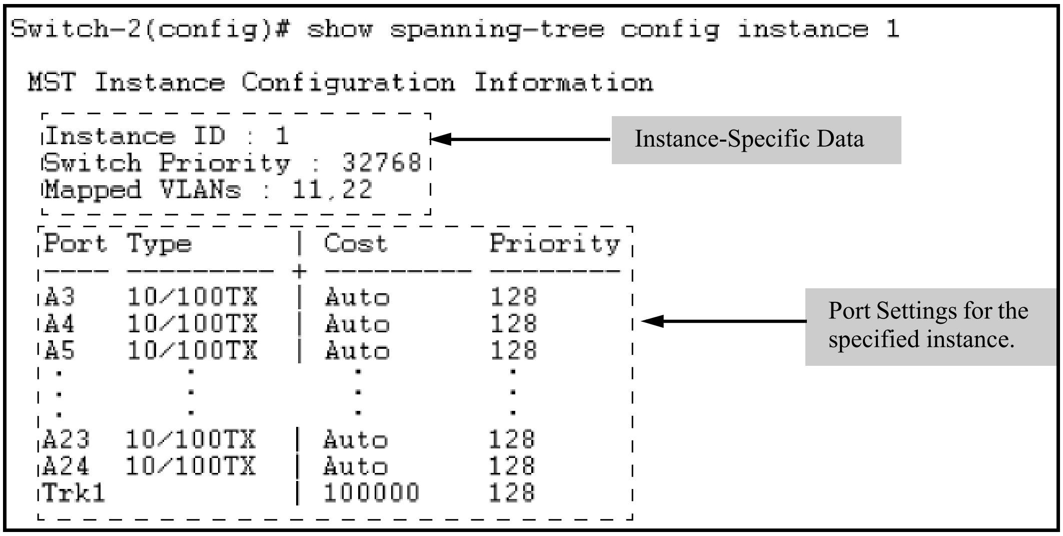

```
Switch-2(config)# show spanning-tree config instance 1

 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:

```
HP Switch(config)# show spanning-tree a20-a24,trk1 config instance 1
```

Displaying 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. (See Example 75 (page 113).)

Example 75 Displaying a region-level configuration

```
HP 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
```

Displaying the pending MSTP configuration

This command displays the MSTP configuration the switch will implement if you execute the `spanning tree pending apply` command. See "Enabling an entire MST region at once or exchanging one region configuration for another" (page 102).

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.

Example

Example 76 Displaying a pending configuration

```
HP 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
```

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*]

Configures per-port loop protection on the switch.

`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` *`vlan-id-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.

1. Configure port mode with this command:

   ```
   HP 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:

   ```
   HP 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:

   ```
   HP 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:

1. Configure VLAN mode with the command:

   ```
   HP 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:

   ```
   HP 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.

Example 77 Changing modes for loop protection

```
HP Switch(config)# loop-protect mode port
Any Loop Protect enabled VLAN will be deleted.
Do you want to continue [Y/N]? n
```

Displaying 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.

Example

Example 78 Displaying loop protection information for port mode

```
HP 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
```

Displaying 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.

Example

Example 79 Displaying loop protection information for VLAN mode

```
HP 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
```

For more information, see "Loop protection" (page 138).

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 12 Loop creation with transmission failure

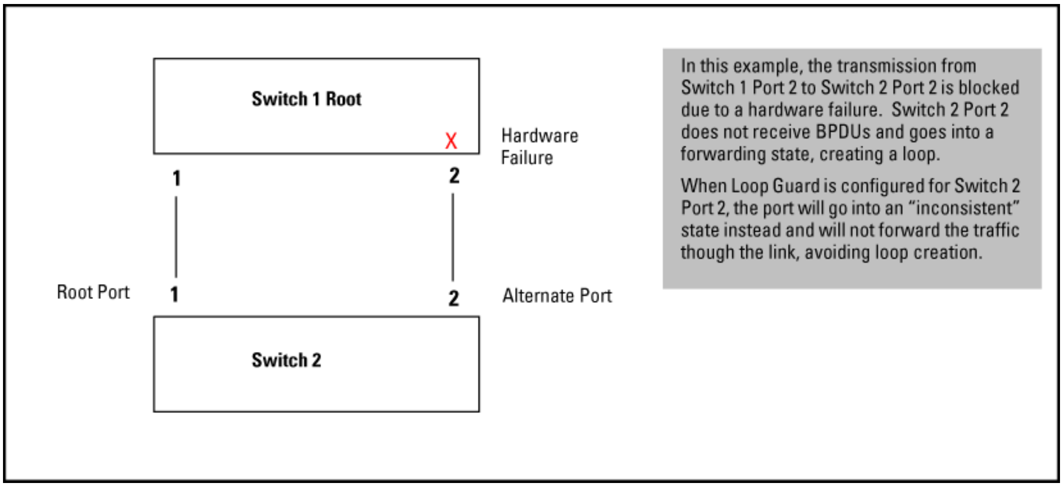
Switch 1 Root
X
Hardware Failure
1
2
Root Port
1
2
Alternate Port
Switch 2
In this example, the transmission from Switch 1 Port 2 to Switch 2 Port 2 is blocked due to a hardware failure. Switch 2 Port 2 does not receive BPDUs and goes into a forwarding state, creating a loop.
When Loop Guard is configured 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:

`[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.

Examples

Example 80 Enabling spanning tree loop guard on Port 2 and displaying the port's status

```
HP Switch(config)# spanning-tree 2 loop-guard
HP 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     |
```

Example 81 Displaying summary spanning tree configuration information

```
HP 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
.
.
.
```

Example 82 Displaying detailed spanning tree configuration information

```
HP 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
```

Example 83 Displaying spanning tree configuration information for a single port

```
HP 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 |
```

Troubleshooting an MSTP configuration

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

Displaying 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.

Examples

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.

Example 84 Displaying `show spanning-tree root-history` CST output

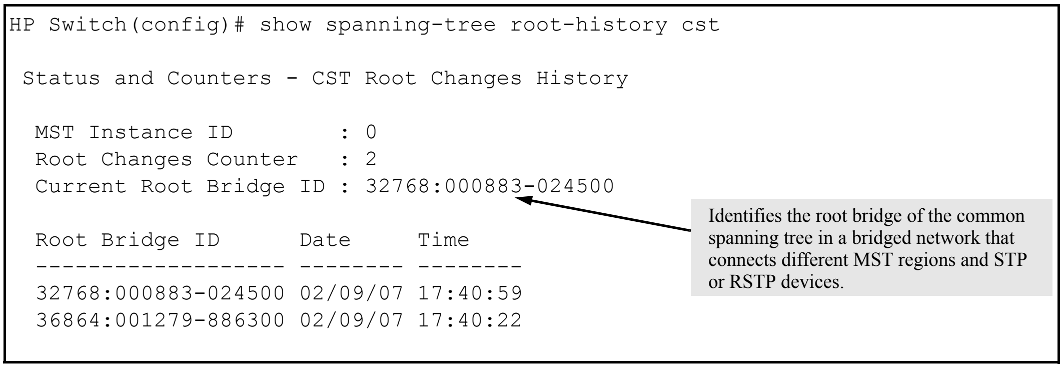

```
HP Switch(config)# show spanning-tree root-history cst

 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.

Example 85 Displaying `show spanning-tree root-history` IST output

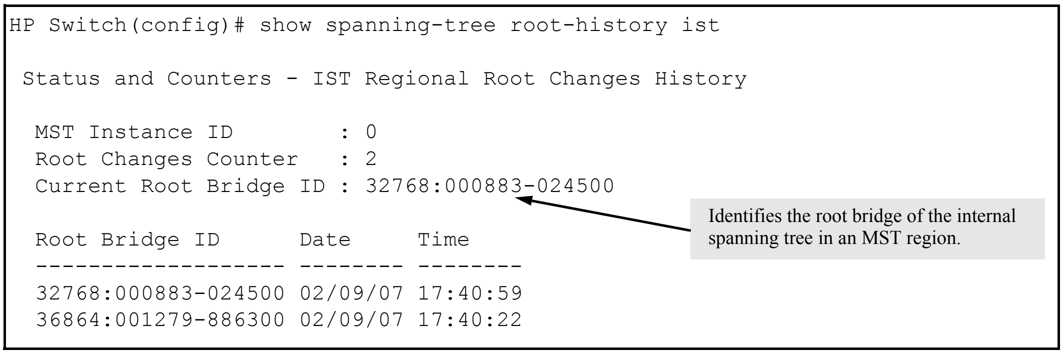

```
HP Switch(config)# show spanning-tree root-history ist

 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.

Example 86 Displaying `show spanning-tree root-history` MSTI output

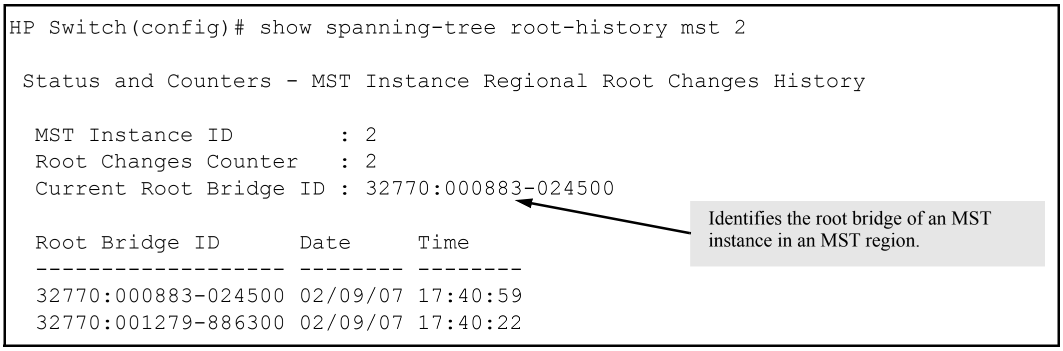

```
HP Switch(config)# show spanning-tree root-history mst 2

 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 displaying trap configuration

Syntax

`[no] spanning-tree trap { errant-bpdu | loop-guard | new-root | root-guard }`

Enables or disables SNMP traps. See “Enabling SNMP traps” (page 90)

Syntax

`show spanning-tree traps`

Displays the current spanning tree trap configuration on the switch.

Exanple

Example 87 Displaying spanning tree traps in their default configuration

```
HP 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
```

Displaying 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.

Example

Example 88 Displaying output for debug counters

The following example shows sample output of the `show spanning-tree debug-counters` command for all ports.

```
HP 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
```

Displaying 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.

Example

Example 89 Displaying 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.

```
HP 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
```

Displaying 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`.

Example

Example 90 Displaying 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.

```
HP 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
```

Example 91 Displaying 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.

```
HP 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 9 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 9 MSTP debug command output: field descriptions *(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.
`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 a 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.

Table 9 MSTP debug command output: field descriptions *(continued)*

Field	Displays the number of...
`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.
`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 10 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.

About MSTP

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 13 An MSTP network with legacy STP and RSTP devices connected

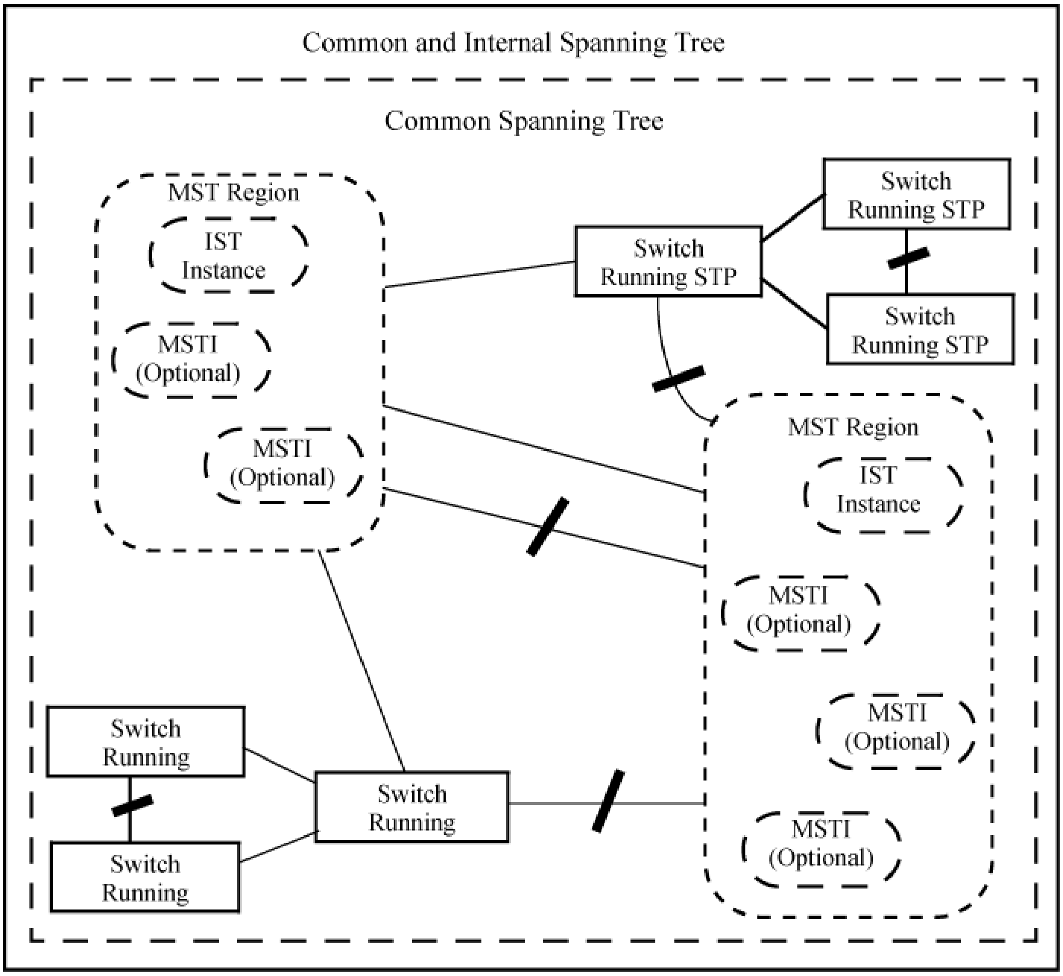

How MSTP operates

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. (See “Enabling an entire MST region at once or exchanging one region configuration for another” (page 102).)

NOTE: The switch automatically senses port identity and type, and automatically defines spanning tree parameters for each type, as well as parameters that apply across the switch. Although these parameters can be adjusted, HP 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 just 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.

For MSTP information beyond what is provided in this manual, see the IEEE 802.1s standard.

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 Figure 14 (page 132) each instance has a different forwarding path.

Figure 14 Active topologies built by three independent MST instances

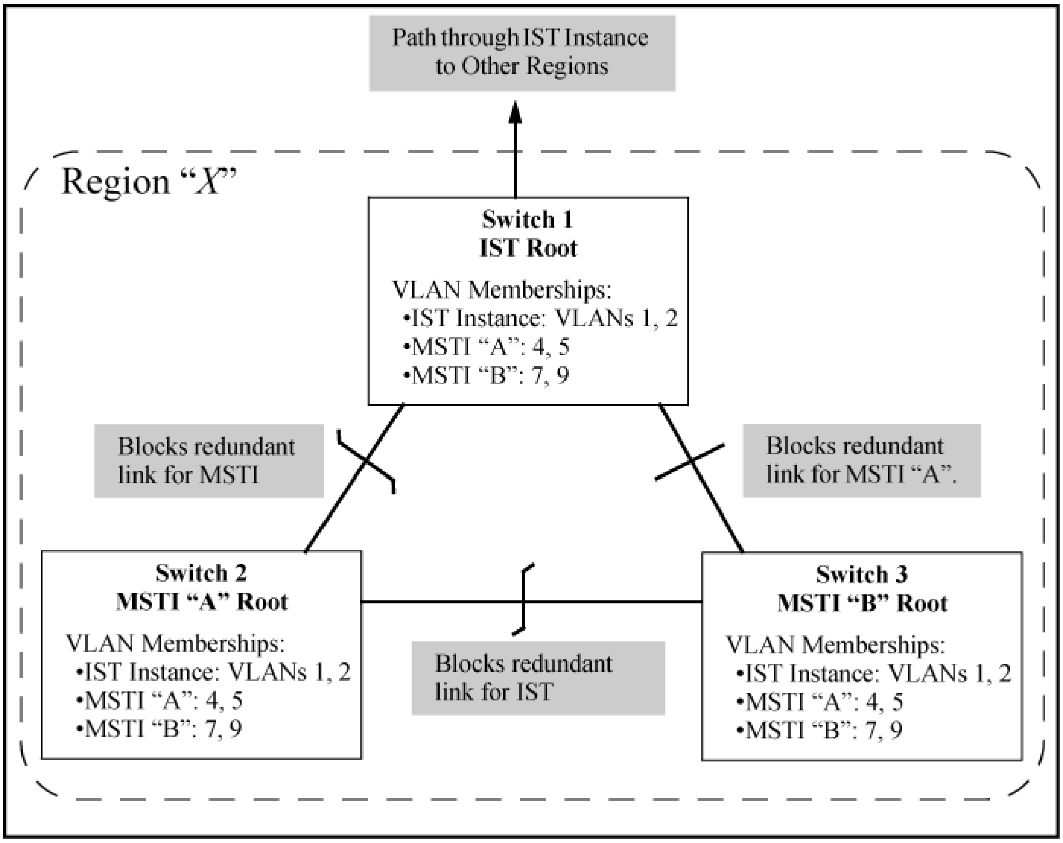
Path through IST Instance to Other Regions
Region "X"
Switch 1
IST Root
VLAN Memberships:
•IST Instance: VLANs 1, 2
•MSTI "A": 4, 5
•MSTI "B": 7, 9
Blocks redundant link for MSTI
Blocks redundant link for MSTI "A".
Switch 2
MSTI "A" Root
VLAN Memberships:
•IST Instance: VLANs 1, 2
•MSTI "A": 4, 5
•MSTI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

All MSTP switches (as well as 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 Figure 13 (page 130).)

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.

Example 92 Using a trunked link to support multiple VLAN connectivity within the same MST instance

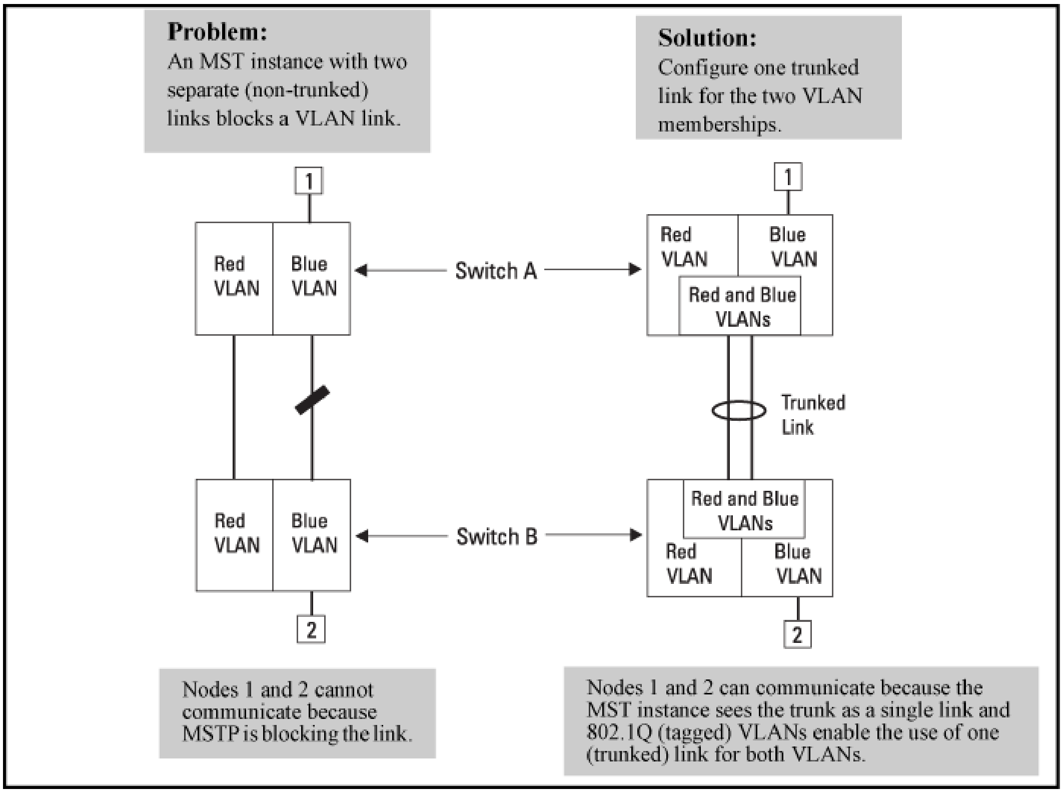
Problem:
An MST instance with two separate (non-trunked) links blocks a VLAN link.
Solution:
Configure one trunked link for the two VLAN memberships.
1
Red VLAN
Blue VLAN
Switch A
Red and Blue VLANs
Trunked Link
Switch B
2
Nodes 1 and 2 cannot communicate because MSTP is blocking the link.
Nodes 1 and 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.

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, HP 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.

Operating rules

- All switches in a region must be configured with the same set of VLANs, as well as 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. For more information, see "Configuring MST instance parameters" (page 98).
- 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.

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.

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. See information on force version on "Setting the spanning tree compatibility mode" (page 87).
- 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 (see "Setting spanning tree to operate with 802. ID legacy path cost values" (page 88)). See also the Note on Path Cost below (page 136).

NOTE: 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 below.

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.

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 in Example 93 “BPDU protection enabled at the network edge”.

Example 93 BPDU protection enabled at the network edge

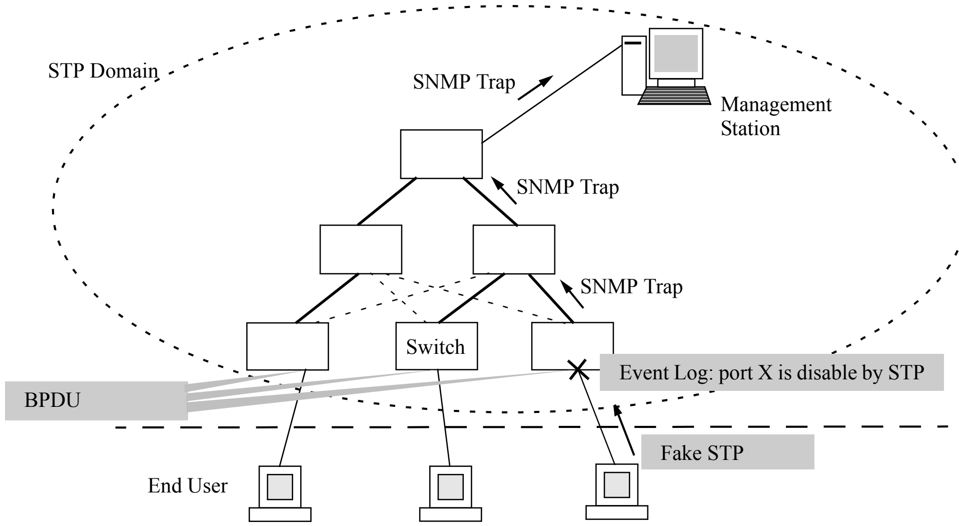

PVST protection and filtering

NOTE: These options are available for switches that support the MSTP protocol only. They are not supported for switches running RSTP.

PVST protection

If an HP 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 `hpSwitchStpPortErrantBpduCounter` is incremented. The PVST protection feature is enabled per-port.

Figure 15 PVST switch being isolated after sending a PVST BPDU

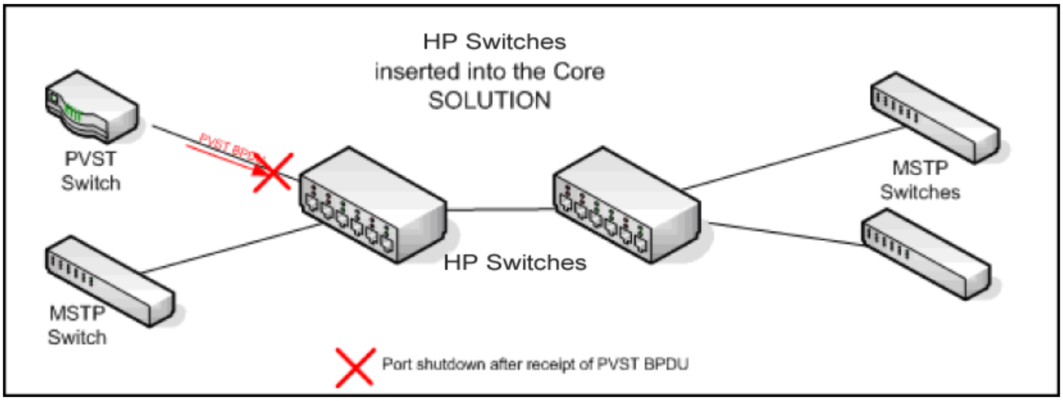

NOTE: This is similar to the BPDU Guard feature where BPDU protection is applied to edge ports connected to end user devices that do not run STP. If STP BPDU packets are received on a protected port, the feature will disable that port and alert the network manager via an SNMP trap.

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 `hpSwitchStpPortErrantBpduCounter` 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
```

Loop protection

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.

Example 94 Loop protection enabled in preference to STP

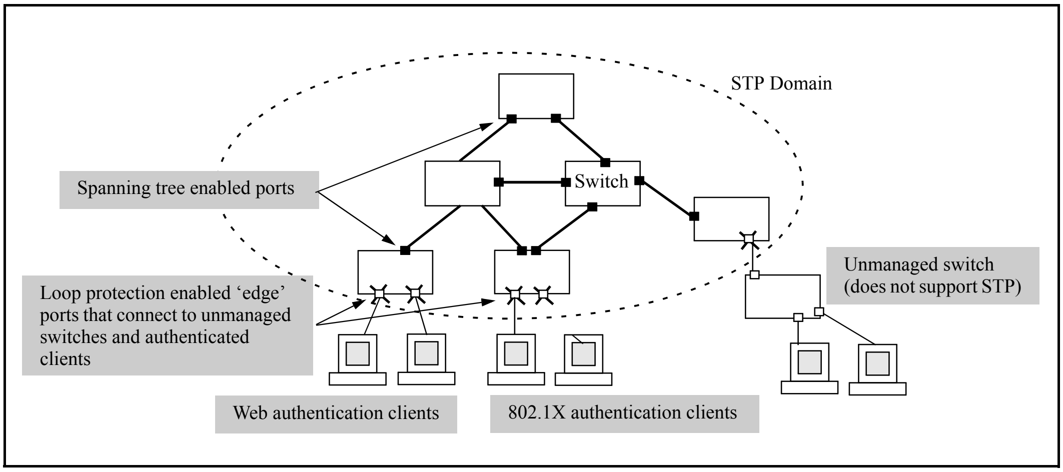
STP Domain
Spanning tree enabled ports
Switch
Unmanaged switch
(does not support STP)
Loop protection enabled ‘edge’ ports that connect to unmanaged switches and authenticated clients
Web authentication clients
802.1X authentication clients

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 a 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.

4 Rapid per-VLAN spanning tree (RPVST+) operation

Command syntax	Description	Default value	CLI page reference
`[no] spanning-tree mode[ mstp | [rapid-pvst]]`	Specifies that spanning tree will run in MSTP (default) or RPVST+ mode.	MST	145
`spanning-tree extend system-id`	Creates a unique bridge identifier for each VLAN.		145
`[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.	Disabled	146
`[no] spanning-tree bpdu-protection-timeout` *timeout*	Configures the duration of time when protected ports receiving unauthorized BPDUs will remain disabled.	0	146
`spanning-tree vlan` *vid* `hello-time` *1...10*	Sets the time in seconds between transmissions of BPDUs on the specified VLAN(s) when the switch is root for those VLAN(s).	2	146
`spanning-tree vlan forward-delay` *4...30*	Sets the time in seconds the switch waits before transitioning from listening to learning and from learning	15	146

Command syntax	Description	Default value	CLI page reference
	to forwarding states.		
`spanning-tree vlan vid-list maximum age 6...40`	Sets the maximum age in seconds of received STP information before it is discarded for specified VLAN(s).	20	146
`spanning-tree vlan vid-list priority 0...15`	Sets the switch (bridge) priority for the designated VLAN.		146
`[no] spanning-tree vlan vid-list root { primary | secondary }`	Specifies the switch as the primary or secondary root bridge for the specified VLAN(s).	Determined by lowest MAC address	147
`[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).	Auto	147
`[no] spanning-tree port port-# vlan vid-list priority priority`	Sets the port priority for the specified VLANs.	8	147
`[no]spanning-tree port-list admin-edge-port`	Enables or disables `admin-edge-port` on ports connected to end nodes.	Disabled	148
`[no] spanning-tree port-list auto-edge-port`	Enables or disables automatic identification of edge ports.	Enabled	148
`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	148
`spanning-tree port-list root-guard`	When enabled, causes a port to not be selected as the root port even if it receives	Disabled	149

Command syntax	Description	Default value	CLI page reference
	superior STP BPDUs.		
spanning-tree *port-list* tcn-guard	When enabled, causes a port to stop propagating received topology change notifications and topology changes to other ports.	Disabled	149
[no] spanning-tree [enable | disable]	Globally enables or disables RPVST+ on all VLANs on the switch.	Disabled	149
spanning-tree vlan *vid list* [enable | disable]	Enables or disables RPVST+ on the specified VLAN(s).	Disabled	149
[no] spanning-tree *port-list* bpdu-filter	Enables or disables BPDU filtering on the specified port(s).	Disabled	150
[no] spanning-tree *port-list* bpdu-protection	Enables or disables BPDU protection on the specified port(s).	Disabled	152
[no] spanning-tree *port-list* loop-guard	Enables or disables STP Loop Guard.	Disabled	154
show spanning-tree *port-list* configuration	Displays spanning tree configuration for specified ports.		150
show spanning-tree bpdu-protection [*port-list*]	Displays a summary or per-port BPDU protection status information.		153
show spanning-tree	Displays spanning tree and VLAN global statistics.		158
show spanning-tree vlan *vlan-id*	Displays detailed		158

Command syntax	Description	Default value	CLI page reference
	spanning tree information for a VLAN and the ports belonging to the specified VLAN.		
`show spanning-tree port-list`	Displays spanning tree status for designated port(s).		159
`show spanning-tree port-list vlan vlan-id`	Displays detailed information for specified port(s) in the specified VLAN .		160
`show spanning-tree system-limits rapid-pvst`	Displays RPVST+ VLAN and virtual port (vPort) status on the switch.		162
`show spanning-tree config vlan vlan-id`	Displays the spanning tree port parameter settings for only the specified VLAN.		167
`show spanning-tree port-list config`	Displays the spanning tree port parameter settings (global and per VLAN) for only the specified port(s) and/or trunk(s).		168
`show spanning-tree port-list config vlan vlan-id`	Displays the spanning tree port parameter settings per port per VLAN.		169
`show spanning-tree root-history vlan vlan-id`	Displays the change history (up to 10 history entries) for the root bridge in the specified RPVST+ topology.		170
`[no] spanning-tree trap { errant-bpdu | loop-guard | new-root | root-guard }`	Enables or disables SNMP traps for errant-BPDU,	Disabled	90

Command syntax	Description	Default value	CLI page reference
	loop guard, new root, and root guard event notifications.		
`show spanning-tree traps`	Displays the spanning tree trap configuration.		171
`show spanning-tree debug-counters`	Displays aggregate values of all RPVST+ debug counters.		171
`show spanning-tree debug vlan vlan-id`	Displays aggregate values of all RPVST+ debug counters for a specified VLAN.		172
`show spanning-tree debug ports` *`port-list`*`vlan` *`vlan-id`*	Displays aggregate values of all RPVST+ debug counters on one or more ports used by a specified VLAN.		173
`spanning-tree clear-debug-counters [ports` *`port-list`*`] [vlan` *`vid-list`*`]`	Clears all spanning tree debug counters unless specific ports and/or VLANs are specified.		176

`[no] debug rpvst [event [filter 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.		176

For conceptual information on RPVST+, see "About RPVST+" (page 177).

Overview

NOTE: For information on configuring basic and multiple instance spanning tree, see the *Multiple instance spanning tree operation* chapter in this guide.

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.

Configuring RPVST+ at a glance

The general steps for configuring RPVST+ via the CLI are:

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:
   - One or more selected VLANs: `spanning-tree vlan` *`vid-list`*
   - One or more selected VLANs: `spanning-tree vlan` *`vid-list`*
   - 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:

- “Configuring BPDU filtering” (page 150)
- “Allowing traffic on VLAN ID (PVID) mismatched links” (page 153)
- “Configuring STP loop guard” (page 154)

Configuring RPVST+

Selecting RPVST+ as the spanning tree mode

Syntax:

`[no] spanning-tree mode [ mstp | rapid-pvst ]`

Specifies that spanning tree will run in MSTP (default) or RPVST+ mode.

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 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. See "Allowing traffic on VLAN ID (PVID) mismatched links" (page 153).

Default: Disabled

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

See also "Configuring and managing BPDU protection" (page 151)

Configuring per-VLAN spanning tree

Syntax:

`spanning-tree vlan` *`vid`* `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 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-list`* `maximum 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). The ability to 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:

[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:

[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:

[no] `spanning-tree port` *`port-number`* `vlan` *`vid-list`* `priority` *`priority`*

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:

[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:

[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

See "Configuring BPDU filtering" (page 150).

Syntax:

[no] `spanning tree` *`port-list`* `bpdu-protection`

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/or 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 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 list disable
[no] spanning-tree vlan vid list enable
```

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 on all VLANs where the port is a member, and stay locked in the spanning tree forwarding state. All other ports will maintain their role.

Syntax:

`[no] spanning-tree [ port-list | all ] bpdu-filter`

Enables/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.

Sample scenarios in which this feature may be used:

- 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, redundant links) using these ports. If you suddenly have a high load, disconnect the link and disable the BPDU filter (using the `no` command).

Example

To configure BPDU filtering on ports 23 and 24, enter:

```
HP Switch(config)# spanning-tree 23,24 bpdu-filter
```

Viewing BPDU filtering

Syntax:

`show spanning-tree port-list configuration`

Displays the BPDU's filter state.

Examples

Example 95 Displaying BPDU filtering for specific ports within the config file

This example shows how BPDU filter state is displayed for ports 23 and 24 within the configuration file.

```
HP Switch# show spanning-tree 23,24 config
Spanning Tree Information
STP Enabled [No] : Yes
Mode : RPVST
Switch MAC Address : 0024a8-d60b80
RPVST Enabled VLANs : 10,20

      Admin Auto Admin Root Loop TCN BPDU BPDU
Port  Edge  Edge PtP   Grd  Grd  Grd Flt  Guard
----- ----- ---- ----- ---- ---- --- ---- -----
23    No    Yes  True  No   No   No  Yes  No
24    No    Yes  True  No   No   No  Yes  No
```

Example 96 Displaying BPDU filtering as separate entries of the spanning tree category within the running config file

This example shows how BPDU filters per port are displayed as separate entries of the spanning tree category within the configuration file.

```
HP Switch(config)# show running-config
Running configuration:
. . .
spanning-tree
spanning-tree 23 bpdu-filter
spanning-tree 24 bpdu-filter
spanning-tree mode rapid-pvst
. . .
```

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 Figure 16 (page 152).

Figure 16 BPDU protection enabled at the network edge

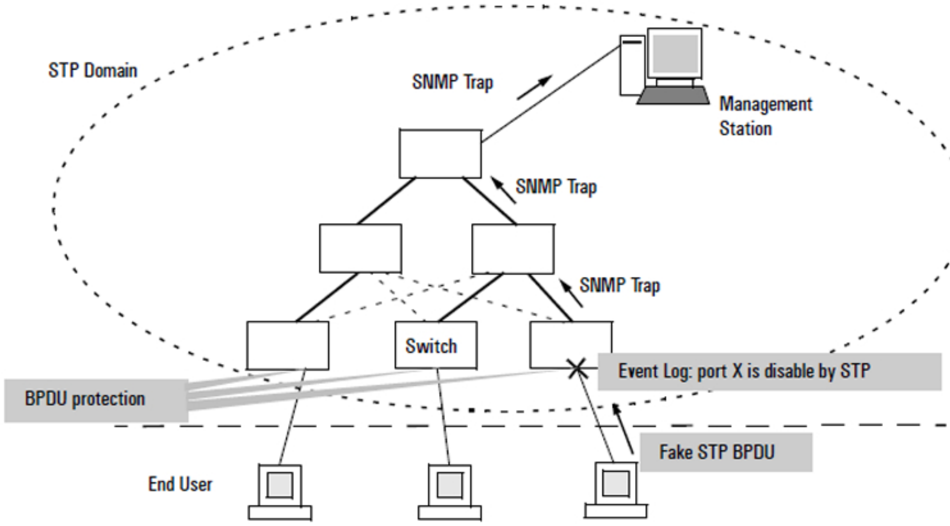

The following commands allow you to configure BPDU protection on VLANs for which the port is a member.

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

For an example of using this command, see "Re-enabling a port blocked by BPDU protection" (page 153).

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.

Example

Example 97 Displaying BPDU protection status for specific ports

```
HP 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:

```
HP Switch(config)# spanning-tree bpdu-protection-timeout 60
```

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

Example

Table 11 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[1]
Untagged on VLAN 10	Untagged on VLAN 20	Forward[1, 2]
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.

[2] If both sides (ports) of the link are untagged to different VLANs, but the VLAN on the switch on one end of the link is not RPVST+-enabled, untagged RPVST+ frames received on that switch port (where RPVST+ is disabled) would be forwarded to any other ports belonging to the inbound VLAN.

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:

[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

Examples

Figure 17 Loop creation with transmission failure

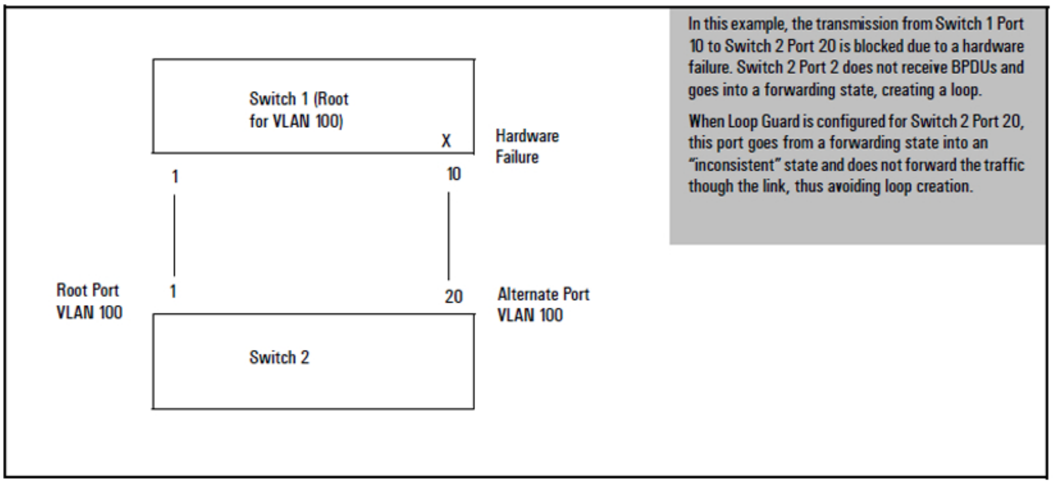

In this example, the transmission from Switch 1 Port 10 to Switch 2 Port 20 is blocked due to a hardware failure. Switch 2 Port 2 does not receive BPDUs and goes into a forwarding state, creating a loop.

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.

Example 98 Before configuring loop guard

Before configuring Loop Guard on port 20, the status of VLAN 20 appears as follows:

```
HP 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
```

Example 99 After configuring loop guard

This example shows that, by executing `spanning-tree 20 loop-guard`, loop guard has been configured on port 20 of Switch 2:

```
HP 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
```

Example 100 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.

```
HP 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
```

Example 101 Displaying config file with loop guard enabled

The following example displays `show spanning-tree config` output with loop guard enabled on Port 20:

```
HP 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
```

Displaying 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).

Displaying RPVST+ global statistics

Displaying global and VLAN spanning tree status

Syntax:

`show spanning-tree`

Displays the switch's global and VLAN spanning tree status.

Example

Example 102 Displaying the switch's global and VLAN spanning tree status

```
HP 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
```

Displaying 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.

Example

Example 103 Displaying status for a specific VLAN

```
HP Switch# 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
```

Displaying 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`

Example

Example 104 Displaying status for a specific port list

```
HP 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
```

Displaying 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`*.

Example

Example 105 Displaying status per-port per-VLAN

```
HP 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
```

Displaying 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.

Examples

Example 106 Displaying BPDU status in show spanning tree output

```
HP 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
```

Example 107 Displaying BPDU protection status on specific ports

```
HP 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
```

Displaying 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.

Examples

Example 108 Displaying RPVST+ VLAN and vPort system limits

```
HP Switch# 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
```

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 in addition to 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. (See Table 12 (page 164).) **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.

Example 109 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]	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.

Exceeding a vPort recommended maximum

In a modular switch, if the vPort count for a given module exceeds the recommended limit for that module, a warning message is displayed in the CLI and an Event Log message is generated. Also, the total vPort count on a switch cannot exceed the maximum vPort count for the switch.

Table 12 Maximum VLANs and vPorts per-switch or per-module

Platform	Maximum PVST-Enabled VLANs	Maximum vPorts per device (User-Assignable)	Maximum vPorts per Chassis Module or Fixed-Port Switch
5400/8200	400	2000	V1 Modules: 200 V2 Modules: 400
3500-24 6200-24 6600-24	400	400	200.
3500-48	400	400	Ports 1 - 24: 200 Ports 25 - 48: 200
3800-24 3800-48 (Standalone or Stack)	400 (Same maximum for standalone or stack.)	2000 (Same maximum for standalone or stack.)	For each stack member: Ports 1 - 24: 400 Ports 25 - 48: 400

NOTE: The output of `show spanning-tree system-limits rapid-pvst` shows a Maximum Allowed Virtual Ports value as a larger number than the values quoted in this table. This is because each port on the switch belongs to at least one VLAN (VLAN-1 by default) and this is a system created vPort instance. For example, on a 3500-24 in the factory default configuration, the switch is already configured with 24 vPorts because all 24 ports are members of VLAN-1. The user assigned vPort instances are in addition to these system assigned vPort instances, meaning that 400 vPort instances can be created over and above the system-created vPort instances. The `show spanning-tree system-limits rapid-pvst` command includes the system-created vPort instances and the user-assignable vPort instances when calculating the maximum vPorts on the switch. In the 3500-24 case, the maximum vPorts would be displayed as 424 (24 system + 400 user-assignable).

Examples

Example 110 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
```

Example 111 Calculating per-module vPorts on chassis switches

In addition to 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

Displaying the RPVST+ configuration

Displaying 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). See “Configuring per-VLAN spanning tree” (page 146).

Example

Example 112 Displaying the global RPVST+ configuration

```
HP 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
```

Displaying 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.

Example

Example 113 Displaying the global RPVST+ configuration per VLAN

```
HP 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
```

Displaying the global RPVST+ configuration per port

Syntax:

`show spanning-tree` *`port-list`* `config`

Lists the spanning tree port parameter settings (global and per VLAN) for only the specified port(s) and/or 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`

Example

Example 114 Displaying the global RPVST+ configuration per port

```
HP 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
```

Displaying the global RPVST+ configuration per port per VLAN

Syntax:

`show spanning-tree` *`port-list`* `config vlan` *`vlan-id`*

Lists the spanning tree port parameter settings per port per VLAN.

Example

Example 115 Displaying the global RPVST+ configuration per port per VLAN

```
HP 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
```

Troubleshooting an RPVST+ configuration

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

Displaying the change history of root bridges

Syntax:

`show spanning-tree root-history vlan` *`vlan-id`*

Displays he 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.

Example

Example 116 Displaying the change history of root bridges

```
HP 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 displaying trap configuration

Syntax:

`[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.

`errant-bpdu`

Enables SNMP notification when an errant BPDU is received. Designed for use with BPDU filtering (See“Configuring BPDU filtering” (page 150)).

`loop-guard`

Enables SNMP notification when a loop guard inconsistency is detected. Designed for use with the Loop Guard option (See “Configuring STP loop guard” (page 154)).

`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. See 149.

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.

Example

Example 117 Displaying spanning tree traps in the default configuration

```
HP 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
```

Displaying 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.

Example

Example 118 Displaying debug counters for all VLANs

```
HP 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
```

See “Field descriptions for RPVST+ debug command output” (page 174).

Displaying 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.

Example

Example 119 Displaying debug counters for a specific VLAN

```
HP 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
```

See "Field descriptions for RPVST+ debug command output" (page 174).

Displaying 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.

Example

Example 120 Displaying 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.

Field	Shows 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 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.
`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 doesn't 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`*
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`

Event	Log message
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 appendix titled "Troubleshooting" in the latest *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/or VLANs are specified.

`ports` *`port-list`*

Clears spanning tree debug counters on the specified ports.

`vlan` *`vid-list`*

Clears spanning tree debug counters for all ports on the VLAN.

Using the `vlan` and `ports` options together clears the spanning tree debug counters on the specified port(s) for the specified VLAN(s). Counters maintained on the same port(s) for other VLAN(s) are not cleared.

Syntax:

`[no] debug rpvst [event [filter 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. (For more information, see the "Troubleshooting" appendix in the latest *Management and Configuration Guide* for your switch.)

`event`

Displays RPVST+ Event Log messages.

`filter vlan` *`vid-list`*

Limits log messages to those generated on the specified VLAN(s).

`packet`

Displays RPVST+ packets sent and received.

`filter port` *`port-list`* `[vlan` *`vid-list`*`]`

Limits packets displayed to those generated on the specified port(s). If the `vlan` option is used, then packets displayed are further limited to the ports on the specified VLAN(s).

The `no` form of the command disables display of RPVST+ debug messages on the destination device.

About RPVST+

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. Note that 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 Figure 18 (page 178)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 18 RSTP forming a single spanning tree across all VLANs

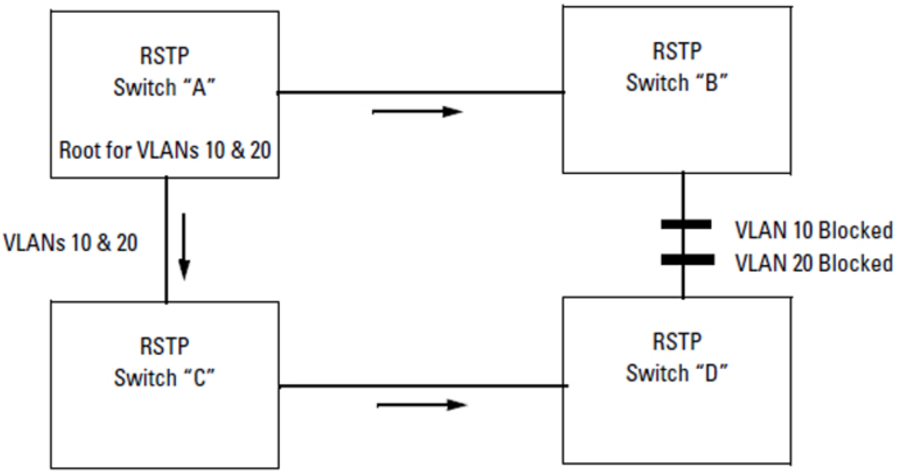

The topology has four switches running RSTP. Switch "A" is the root switch. In order to prevent a loop, RSTP blocks the link between switch "B" and switch "D". There are two VLANs in this network (VLAN 10 and VLAN 20). Since RSTP does not have VLAN intelligence, it forces all VLANs in a layer 2 domain 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 19 RPVST+ creating a spanning tree for VLAN 10

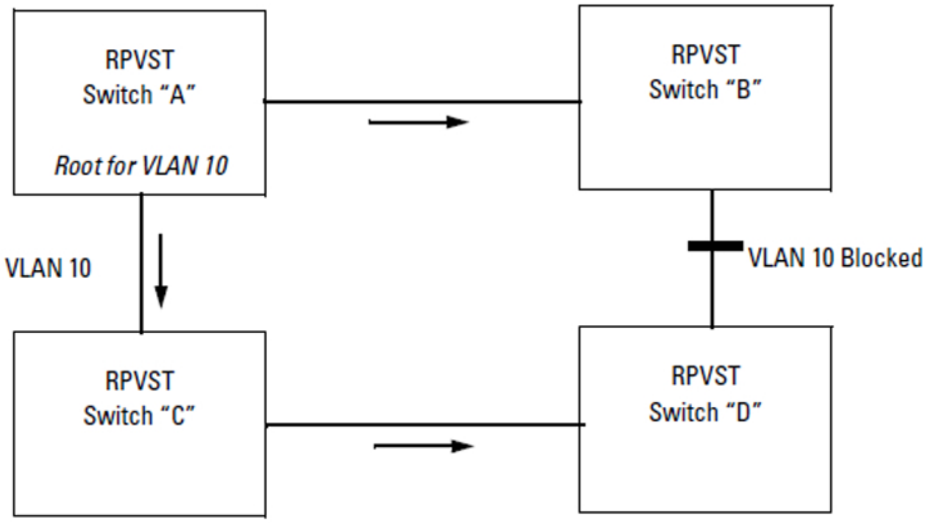

Figure 20 RPVST+ creating a spanning tree for VLAN 20

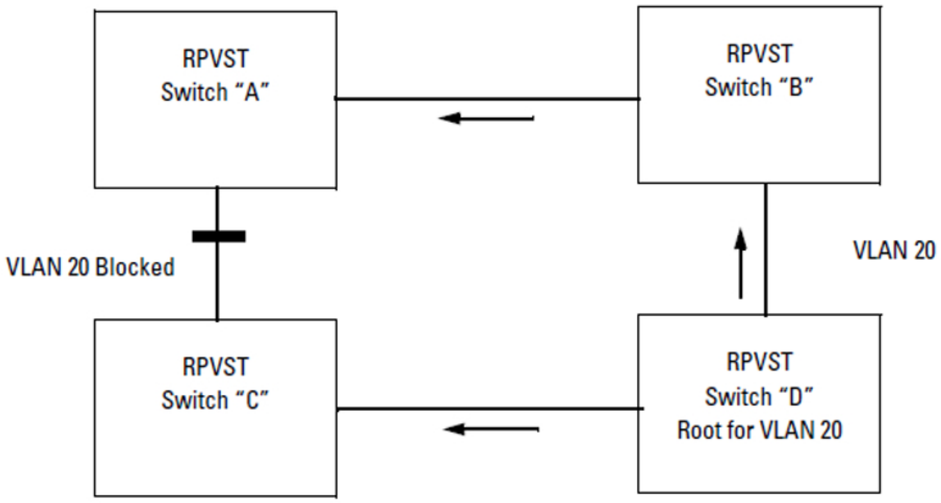

The two topologies above are the same as the first topology, but now the switches run RPVST+ and can span different trees for different VLANs. Switch "A" is the root switch for the VLAN 10 spanning tree and switch "D" is the root switch for the VLAN 20 spanning tree. The link between switch "B" and switch "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 21 Sample RPVST+ network

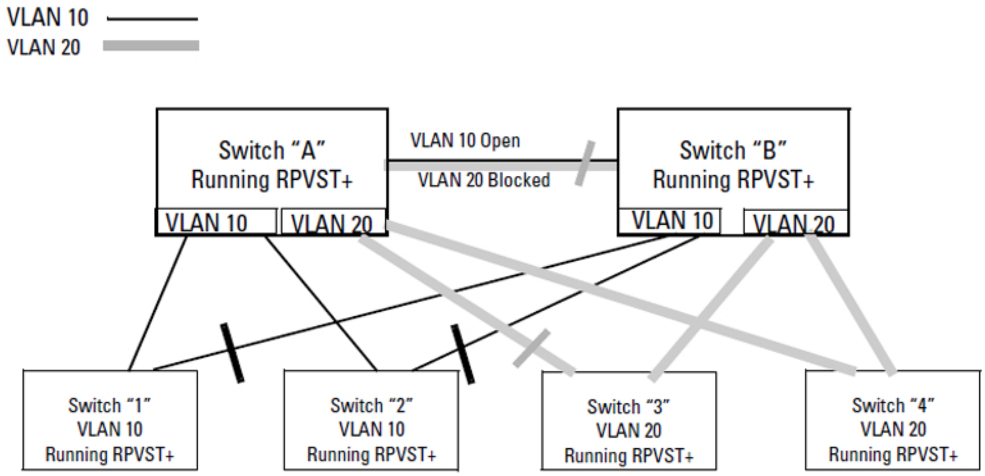

Working with the default RPVST+ configuration

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, as well as parameters that apply across the switch. *Although these parameters can be adjusted, HP 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.*

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. For more information, see "Displaying RPVST+ VLAN and vPort system limits" (page 162).
- **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.

5 Switch meshing

Command syntax	Description	Menu reference page	CLI reference page
[no] `mesh` [e] *`port-list`* `[id` *`1...31`* `]`	Configures a switch mesh	185	184
`show mesh`	Displays the mesh configuration	185	189
`show mesh mac-address`			
`show mesh trace route mac-address` *`MAC-addr`* `vlan` *`vid`*			

Introduction

Switch meshing is a load-balancing technology that enhances reliability and performance in these ways:

- Provides significantly better bandwidth utilization than either Spanning Tree Protocol (MSTP) or standard port trunking.
- Uses redundant links that remain open to carry traffic, removing any single point of failure for disabling the network, and allowing quick responses to individual link failures. This also helps to maximize investments in ports and cabling.
- Unlike trunked ports, the ports in a switch mesh can be of different types and speeds (10 and 100 Mbps, gigabit, and 10 gigabit). For example, a 10Base-FL port and a 1GB port can be included in the same switch mesh.

Example 121 Switch meshing

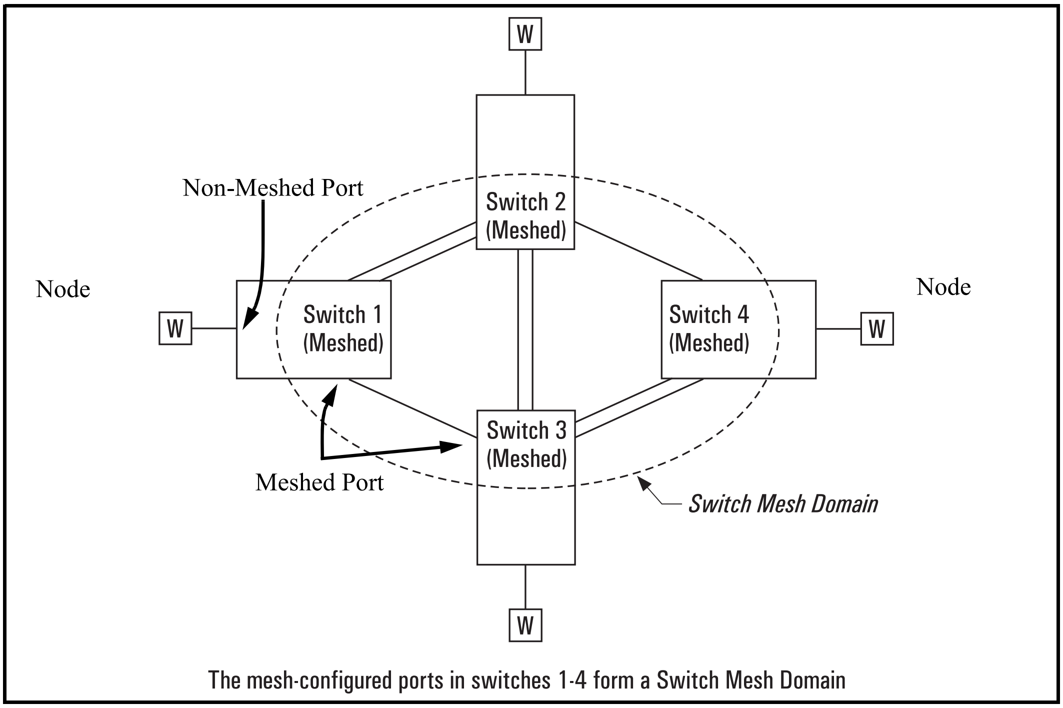

Finding the fastest path

Using multiple switches redundantly linked together to form a meshed switch domain, switch meshing dynamically distributes traffic across load-balanced switch paths by seeking the fastest paths for new traffic between nodes. In actual operation, the switch mesh periodically determines the lowest latency paths, then assigns these paths as the need arises. The path assignment remains until the related MAC address entry times out. The mesh sees later traffic between the same nodes as new traffic, and may assign a different path, depending on conditions at the time. For example, at one time the best path from node A to node B is through switch 2. However, if traffic between node A and node B ceases long enough for the path assignment to age out, then the next time node A has traffic for node B, the assigned path between these nodes may be through switch 3 if network conditions have changed significantly.

NOTE: The `mac-age-time` parameter determines how long an inactive path assignment remains in memory. See "System Information" in the chapter titled "Interface Access and System Information" in the *Management and Configuration Guide* for your switch.

Because redundant paths are active, meshing adjusts quickly to link failures. If a link in the mesh fails, the fast convergence time designed into meshing typically has an alternate route selected in less than a second for traffic that was destined for the failed link.

Meshing allows scalable responses to increasing bandwidth demand

As more bandwidth is needed in a LAN backbone, another switch and another set of links can be added. This means that bandwidth is not limited by the number of trunk ports allowed in a single switch.

Preparing to configure switch meshing

Before configuring switch meshing:

1. Review the Operating Rules (page 194), and particularly the restrictions and requirements for using switch meshing in environments that include static trunks, multiple static VLANs, GVRP, IGMP, and MSTP.
2. To avoid unnecessary system disruption, plan the mesh bring-up to minimize temporary port blocking. (See "Bringing up a switch mesh domain" 196).
3. To view the current switch mesh status on the switch, use the CLI `show mesh` command. (See"Viewing switch mesh status (CLI)" 189).

Configuring switch meshing (CLI)

Syntax:

`[no] mesh [e] port-list [id 1...31 ]`

Enables or disables meshing operation on the specified ports.

`port-list`

Specifies the ports to be added or removed from a mesh.

`id 1...31`

Specifies the mesh ID. It must be unique; if two switches in a mesh are assigned the same mesh ID, they auto-negotiate and generate unique IDs. Only configured mesh IDs can be changed by a user.

NOTE: When the mesh ID is configured, a reboot is required for this ID to take effect.

To view the mesh IDs, enter the `show mesh id` command.

The `no mesh id` command resets the configured mesh ID to zero.

Default: Mesh ID = 0

All meshed ports on a switch belong to the same mesh domain. To configure multiple meshed ports on a switch, you need to:

1. Specify the ports you want to operate in the mesh domain.
2. Use `write memory` to save the configuration to the startup-config file.
3. Reboot the switch. For switches with redundant management modules, such as 8200zl switches, you must reboot both management modules. Use the boot system … command.

Examples

Example 122 Configuring meshing

To configure meshing on ports A1-A4, B3, C1, and D1-D3:

```
HP Switch(config)# mesh a1-a4, b3, c1, d1-d3
Command will take effect after saving configuration and reboot.
HP Switch(config)# write memory
HP Switch(config)# boot
Device will be rebooted, do you want to continue [y/n]? y
```

Example 123 Removing a port from meshing

To remove a port from meshing, use the no version of `mesh`, followed by `write memory` and rebooting the switch. For example, to remove port C1 from the mesh:

```
HP Switch# config
HP Switch(config)# no mesh c1
Command will take effect after saving configuration and reboot.
HP Switch(config)# write memory
HP Switch(config)# boot
Device will be rebooted, do you want to continue [y/n]? y
```

Configuring switch meshing (Menu)

1. From the Main Menu, select: **2. Switch Configuration** —> **2. Port/Trunk Settings**
2. Press **E** (for Edit) to access the load balancing parameters.

Example 124 Configuring ports for meshing (menu)

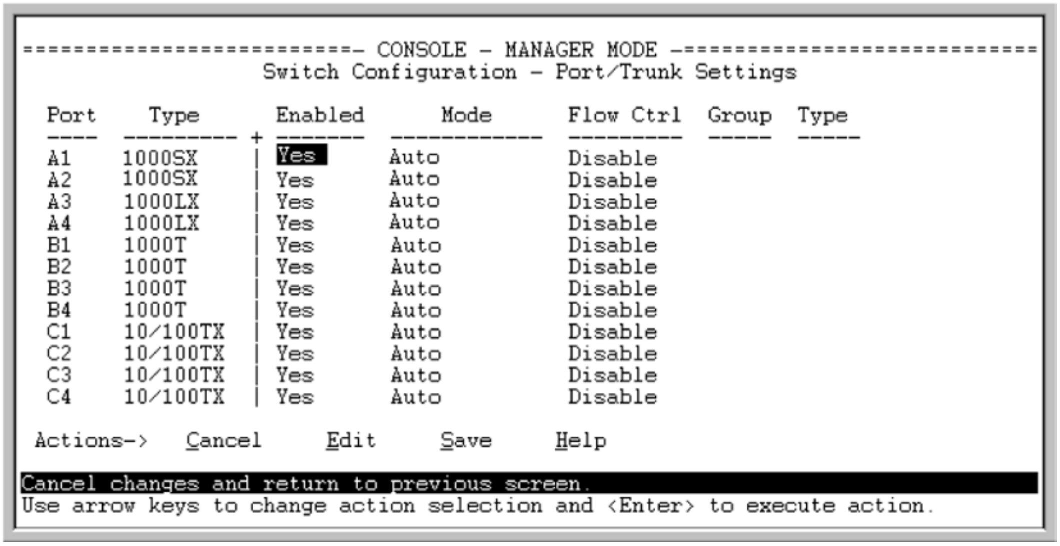

```
==========================- CONSOLE - MANAGER MODE -==========================
                    Switch Configuration - Port/Trunk Settings

  Port    Type         Enabled      Mode          Flow Ctrl  Group  Type
  ----   ---------- + -------  ------------    ---------  -----  -----
  A1     1000SX     | Yes      Auto             Disable
  A2     1000SX     | Yes      Auto             Disable
  A3     1000LX     | Yes      Auto             Disable
  A4     1000LX     | Yes      Auto             Disable
  B1     1000T      | Yes      Auto             Disable
  B2     1000T      | Yes      Auto             Disable
  B3     1000T      | Yes      Auto             Disable
  B4     1000T      | Yes      Auto             Disable
  C1     10/100TX   | Yes      Auto             Disable
  C2     10/100TX   | Yes      Auto             Disable
  C3     10/100TX   | Yes      Auto             Disable
  C4     10/100TX   | Yes      Auto             Disable

 Actions->   Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

3. In the Group column, move the cursor to the port you want to assign to the switch mesh.
4. Press **M** to choose Mesh for the selected port.

5. Use the up-arrow or down-arrow key to select the next port you want to include in your mesh domain, then press **M** again. For example, if you were adding ports A1 and A2 to your mesh domain, the screen would appear similar to Example 125 (page 186):

Example 125 Mesh group assignments for several ports

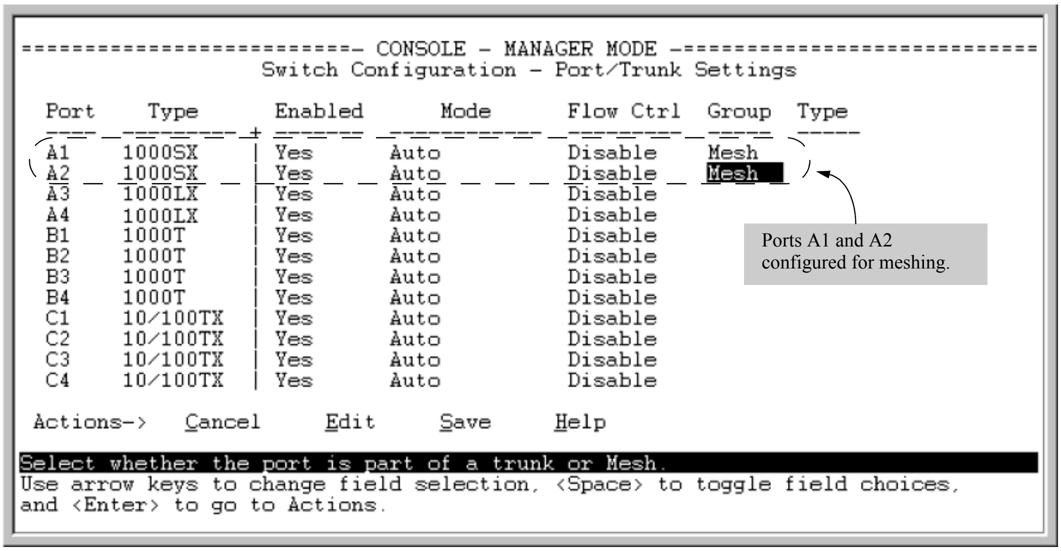

```
=========================- CONSOLE - MANAGER MODE -===========================
                   Switch Configuration - Port/Trunk Settings

  Port    Type      Enabled      Mode        Flow Ctrl  Group  Type
  ----  ---------+ -------  ------------  ---------  -----  -----
  A1    1000SX    | Yes      Auto          Disable    Mesh
  A2    1000SX    | Yes      Auto          Disable    Mesh
  A3    1000LX    | Yes      Auto          Disable
  A4    1000LX    | Yes      Auto          Disable
  B1    1000T     | Yes      Auto          Disable
  B2    1000T     | Yes      Auto          Disable
  B3    1000T     | Yes      Auto          Disable
  B4    1000T     | Yes      Auto          Disable
  C1    10/100TX  | Yes      Auto          Disable
  C2    10/100TX  | Yes      Auto          Disable
  C3    10/100TX  | Yes      Auto          Disable
  C4    10/100TX  | Yes      Auto          Disable

 Actions->   Cancel     Edit     Save     Help

Select whether the port is part of a trunk or Mesh.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Ports A1 and A2 configured for meshing.

6. Repeat Step 5 for all ports you want in the mesh domain.

NOTE: As meshed ports do not accept a `Type` setting, leave the `Type` setting blank.

All meshed ports in the switch automatically belong to the same mesh domain, see Example 133 (page 193).

7. When you finish assigning ports to the switch mesh, press **Enter**, then **S** (for Save). You will then see the following screen.

Figure 22 After saving a mesh configuration change, rebooting the switch

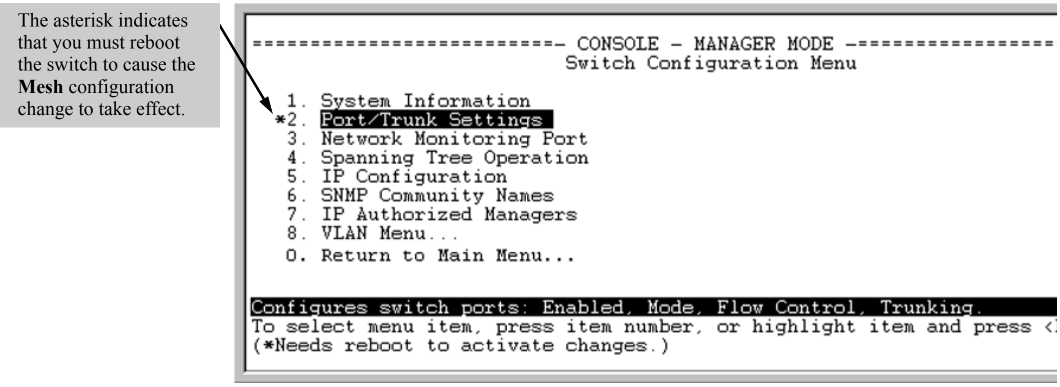

The asterisk indicates that you must reboot the switch to cause the **Mesh** configuration change to take effect.

```
=========================- CONSOLE - MANAGER MODE -=================
                     Switch Configuration Menu

     1. System Information
    *2. Port/Trunk Settings
     3. Network Monitoring Port
     4. Spanning Tree Operation
     5. IP Configuration
     6. SNMP Community Names
     7. IP Authorized Managers
     8. VLAN Menu...
     0. Return to Main Menu...

Configures switch ports: Enabled, Mode, Flow Control, Trunking.
To select menu item, press item number, or highlight item and press <
(*Needs reboot to activate changes.)
```

8. Press **0** to return to the Main menu.
9. To activate the mesh assignments from the Main menu, reboot the switch by pressing the following keys:
   a. **6** (for `Reboot Switch`)
   b. Space bar (to select `Yes`).
   c. **13** (to start the reboot process).

The switch cannot dynamically reconfigure ports to enable or disable meshing, so it is always necessary to reboot the switch after adding or deleting a port in the switch mesh.

Configuring concurrent meshing and routing

Concurrent meshing and routing is only supported on the HP 5400 series and HP 8200 series switches using these modules.

Module	Description
J9534A	24-Port Gig-T PoE+ v2 zl Module
J9535A	20-Port Gig-T PoE+ / 4-port SFP v2 zl Module
J9536A	20-Port Gig-T PoE+ / 2-port 10-GbE SFP+ v2 zl Module
J9537A	24-Port SFP v2 zl Module
J9538A	8-Port 10-GbE SFP+ v2 zl Module
J9547A	24-Port 10/100 PoE+ v2 zl Module
J9548A	20-Port Gig-T / 2-port 10-GbE SFP+ v2 zl Module
J9549A	20-Port Gig-T / 4-port SFP v2 zl Module
J9550A	24-Port Gig-T v2 zl Module
J9637A	12-Port SFP / 12-port PoE+ v2 zl Module

NOTE: Since concurrent meshing and routing is only supported on V2 modules, the `no allow-v1-modules` configuration parameter must be set on switches that are configured for meshing and routing.

Meshing and routing can be configured simultaneously. A packet can be routed into a mesh, or be switched through a mesh and then routed. Two routers can be connected by mesh links, which offers additional network topologies between routers and switches. Concurrent meshing and routing makes it possible to implement meshing throughout a broadcast domain without the need for additional switches or the use of another Layer 2 technology such as Spanning Tree to connect meshing domains with routing switches.

It is important to remember that meshing provides Layer 2 load balancing only; two traffic streams going to the same router will take the same path through the mesh as the traffic is going to the same MAC address.

NOTE: The mesh port on a switch is tagged on all VLANs on that switch. It is recommended that every meshed switch have the same VLANs, whether they are used or not.

Meshing routers and switches

When Router A has no ports belonging to VLAN Z, the packets arriving on Router A's non-mesh ports at VLAN X can be routed to VLAN Y, travel through the mesh, and arrive at Router B. After that they can be routed from VLAN Y to VLAN Z.

Example 126 A router mesh

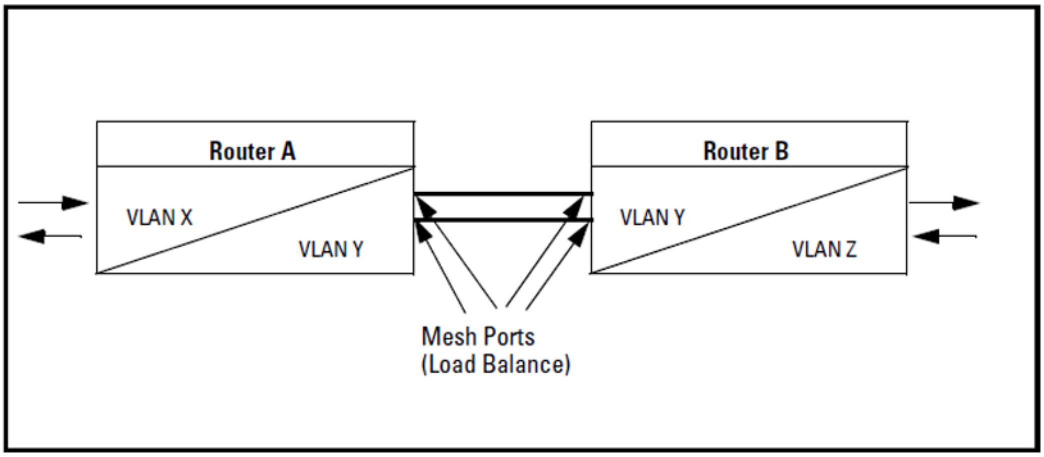

Switches and routers can be meshed together to create more complex local area networks with many redundant mesh links. Load balancing utilizes redundant links in the mesh to deliver traffic efficiently.

Packets arriving at Router A's non-mesh port at VLAN X may be routed to VLAN Y, then switched through the mesh to a host connected to Switch B. Depending on the cost of the link between Router A and Switch B, the packets may be delivered to Switch C, and then forwarded to Switch B if that mesh path is less costly.

Packets arriving at the non-mesh ports for Switch B and Switch C may need to reach Router A. These packets are delivered through the mesh with the least mesh path cost.

Example 127 Meshing with one router and two switches

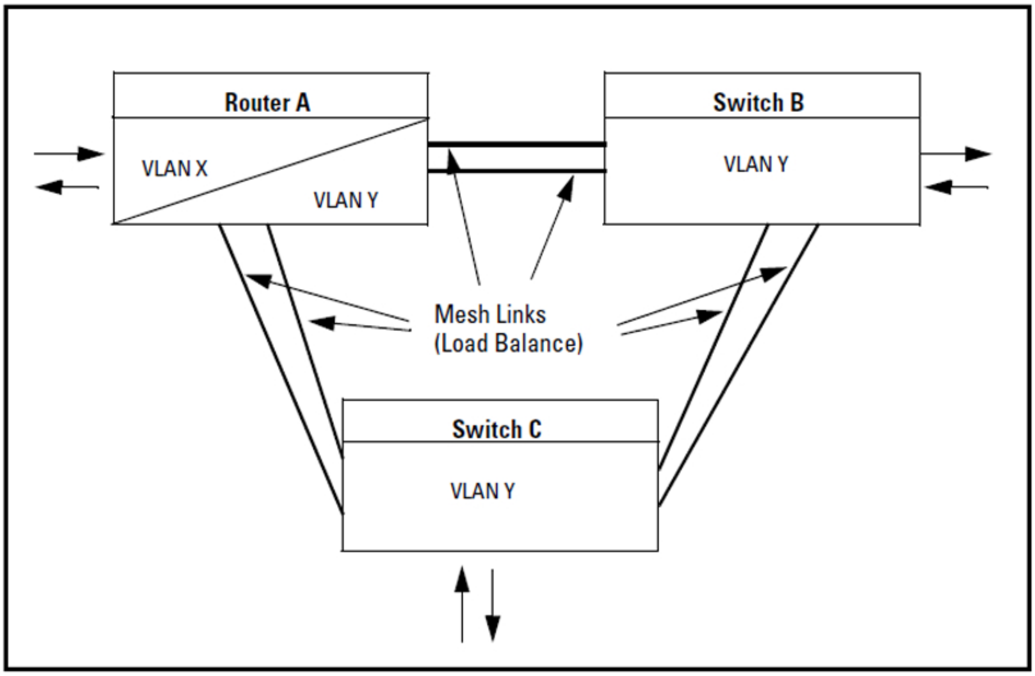

The following example shows two routers and one switch that are meshed together. Packets arriving at the non-mesh ports on VLAN X on Router A may be routed to VLAN Y. Load balancing determines which port Router A should send the packets to in order for the packets to reach Router B.

Example 128 Meshing with two routers and one switch

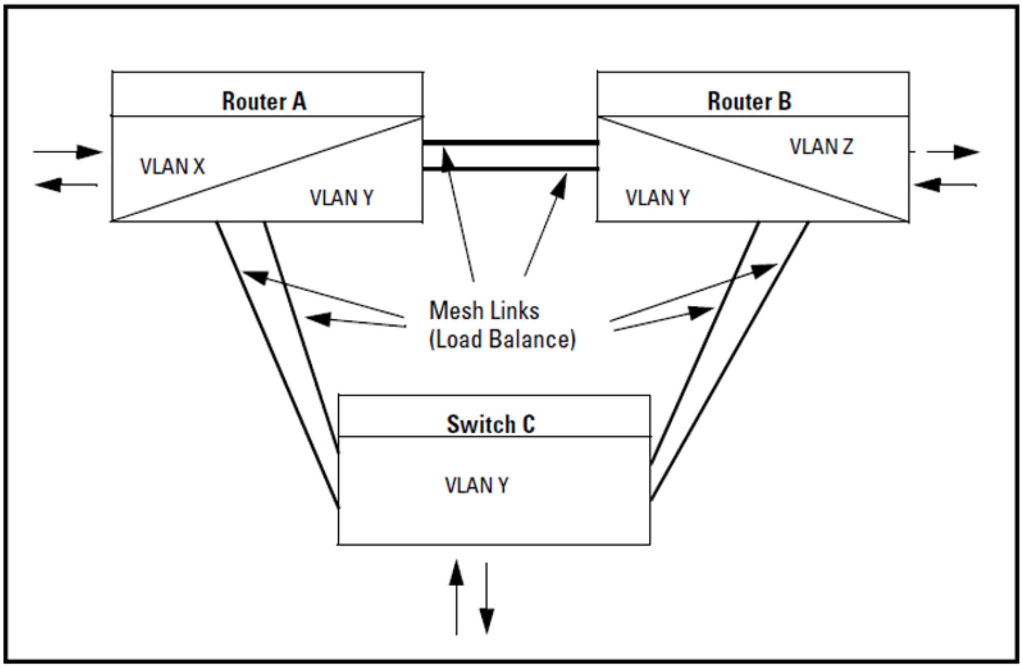

Viewing switch mesh status (CLI)

Syntax:

```
show mesh
show mesh mac-address
show mesh traceroute mac-address MAC-addr vlan vid
```

These three commands list the switch ports configured for meshing, along with the state of each mesh-configured connection, the hostname and MAC address of the switch on the opposite end of the link (Adjacent Switch), the MAC address of the port on the opposite end of the link (Peer Port), and whether or not any mesh warnings have been generated. Mesh warnings are written to the event log.

The switch presents show mesh output in the following format:

```
Adjacent Hosts

Port  State       | Hostname  Address     Peer Port   Mesh Warning
----- ----------- + -------- ----------- ----------- ------------
See the log file for details of warnings.
```

Port

Lists the ports on the switch that have been configured for meshing.

State

Shows the operating state of the port:

Established

The port is linked to a meshed port on another switch and meshing traffic is flowing across the link. The show mesh listing includes the MAC addresses of the adjacent switch and direct connection port on the adjacent switch.

Not Established

The port may be linked to a switch on a port that is not configured for meshing or has gone down.

Initial

The port has just come up as a meshed port and is trying to negotiate meshing.

Disabled

The port is configured for meshing but is not connected to another device.

Error

Error caused by external traffic that also happens to be an apparent duplicate MAC address.

Topology Error

Other sources of traffic may trigger a topology error, but may be simply dropped.

Hostname

The name of the adjacent switch (the switch at the other end of the link).

Address

The base MAC address of the adjacent switch.

Peer port

The MAC address of the port on the adjacent switch to which the mesh link connects.

Mesh Warning

Indicates whether the meshing process has generated a warning for the mesh link.

Yes

One or more warning messages have been generated and written to the event log. Use the show log command to see the warnings.

No

No warning messages have been generated by the meshing process for this port.

Examples

Example 129 Displaying `show mesh` output

```
North# show mesh

Adjacent Hosts


Port   State           | Hostname Address        Peer Port     Mesh Warning
------ --------------- + -------- ------------- ------------- ------------
13 Established         | East    001560-f9e300  001560-f9f3f3     No
14 Established         | South   001f28-244a00  001f28-245ac3     No
15 Established         | West    001279-884300  001279-8853f3     No

See the log file for details of warnings.
```

Example 130 A mesh topology

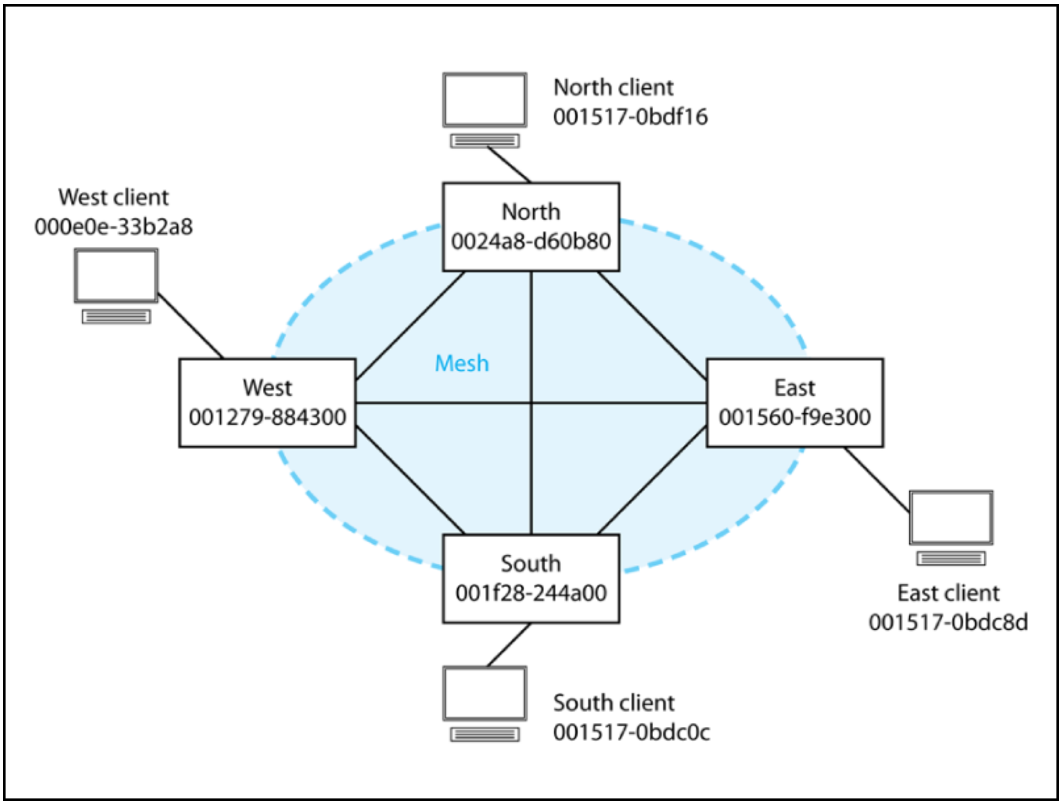

Syntax:

`show mesh mac-address`

Lists information about devices connected to the switch mesh in the following format.

```
MAC Address  VLAN  Port  Owner  Switch   Hostname
------------ ----  ----  ------ -------  ----------
```

MAC Address

The MAC address of the device connected to the switch mesh.

VLAN

The VLAN of the switch mesh.

Port

The port on the originating switch through which the device's MAC address was obtained.

Owner Switch

The base MAC address of the switch to which the device is connected.

Hostname

The hostname of the switch to which the device is connected.

NOTE: The information shown by the `show mesh mac-address` command is not static. Addresses age out of a switch's meshing address table after approximately five minutes of inactivity, and the switches in a mesh exchange meshing address information at intervals of approximately five minutes. So client devices that appear in the listing may have disappeared the next time you view the listing. If you wish to see device information for a device that is no longer shown in the MAC address listing, a simple way to activate the device's entry in the meshing address table and make its MAC address reappear in the listing is to ping the device's IP address.

Example

Example 131 Displaying device information

For the mesh topology shown in Example 130 (page 191), the `show mesh mac-address` command issued from the North switch displays the following device information. This listing shows the MAC addresses of the clients connected to the South, East, and West switches.

```
North# show mesh mac-address
MAC Address  VLAN Port Owner Switch   Hostname
------------ ---- ---- ------------- ----------
001517-0bdc0c 1   14  001f28-244a00   South
001517-0bdc8d 1   13  001560-f9e300    East
000e0c-33b2a8 1   15  001279-884300    West
There are 3 addresses in the meshing address table
```

Syntax:

`show mesh traceroute mac-address` *`MAC-addr`* `vlan` *`vid`*

Traces the route from a source switch in a mesh to a device connected to the mesh.

`MAC-addr`

MAC address of the target device.

`vid`

VLAN number of the mesh, between 1 and 4094.

Hop

The hop count of the trace, beginning at 0.

Address

The base MAC address of the switch involved in the hop.

Hostname

The hostname of the switch involved in the hop.

Inbound-Port

The number of the port through which the trace enters the switch.

Outbound-Port

The number of the port through which the trace exits the switch.

The direction of the trace is from the originating switch to the target device.

Example

Example 132 Displaying traceroute information

For the example mesh topology shown in Example 130 (page 191), the following `switch mesh traceroute` command traces the route to the client attached to the South switch.

```
 North# show mesh traceroute mac-address 001517-0bdc0c vlan 1

Traceroute to MAC Address: 001517-0bdc0c VID: 1

Hop    Address        Hostname      Inbound-Port     Outbound-Port
0     0024a8-d60b80    North                              13
1     001560-f9e300    East              A13              A14
2     001f28-244a00    South             C14              A1

Trace route is completed. Total hops to destination is 3
```

NOTE: In this case the trace traverses two links (North-to-East and East-to-South), rather than going directly from the North switch to the South switch. The mesh chooses the route based on traffic loads on the links and other factors. As a result, the mesh may choose different paths between any two points at different times. You may notice these differences when you trace routes through the mesh.

About switch meshing

Switch mesh domain

This is a group of meshed switch ports exchanging meshing protocol packets. Paths between these ports can have multiple redundant links without creating broadcast storms.

Example 133 A switch mesh domain in a network

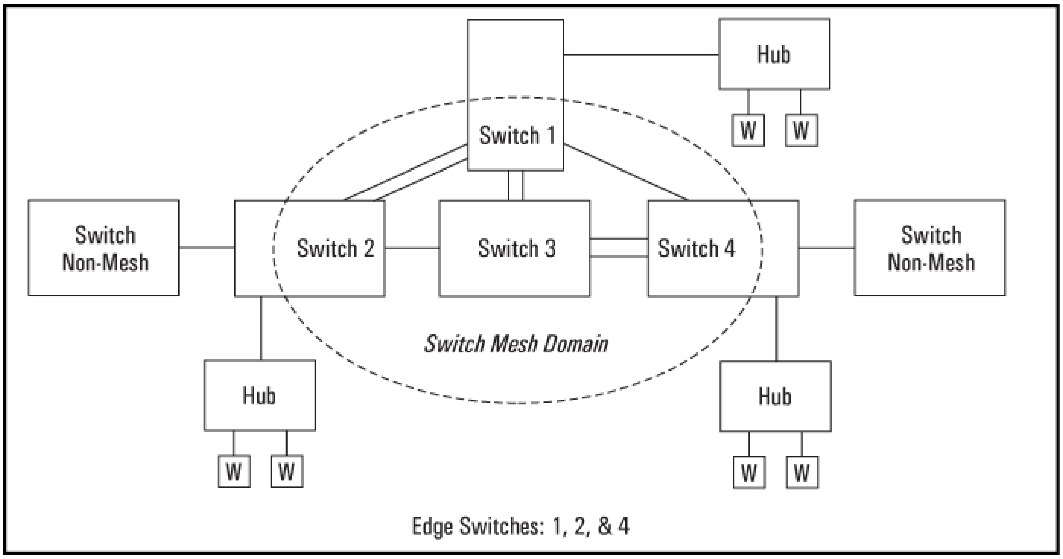

Edge switch

This is a switch that has some ports in the switch meshing domain and some ports outside of the domain. (See Example 133 (page 193).)

Operating rules

(See also “Mesh design optimization” (page 201).)

- A meshed switch can have some ports in the meshed domain and other ports outside the meshed domain. That is, ports within the meshed domain must be configured for meshing, while ports outside the meshed domain must not be configured for meshing.
- Meshed links must be point-to-point switch links.
- On any switch, all meshed ports belong to the same mesh domain.
- A switch can have up to 24 meshed ports.
- A mesh domain can include up to 12 switches.
- Up to five inter-switch, meshed hops are allowed in the path connecting two nodes through a switch mesh domain. A path of six or more meshed hops between two nodes is unusable. However, in most mesh topologies, there would normally be a shorter path available, and paths of five hops or fewer through the same mesh will continue to operate.
- Other sources of traffic between meshed switch links are not allowed.
- If the switch has multiple static VLANs and you configure a port for meshing, the port becomes a tagged member of all such VLANs . If you remove a port from meshing, it becomes an untagged member of only the default VLAN.
- A port configured as a member of a static trunk (LACP or Trunk) cannot also be configured for meshing.
- If a port belongs to a dynamic LACP trunk and you impose meshing on the port, it automatically ceases to be a member of the dynamic trunk.
- Meshing is not supported on ports configured with 802.1X access control.
- On a port configured for meshing, if you subsequently remove meshing from the port's configuration and reboot the switch, the port returns to its default configuration. (It does not revert to any non-default configuration it had before being configured for meshing).
- In a given mesh domain, switches in the same product family must run the same switch software version. For example, if you update the software version on one 8212zl switch, then you must update the software version on any other 8212zl switch in the mesh. HP recommends that you always use the most recent software version available for the switches in your network.
- The spanningtree configuration must be the same for all switches in the mesh (enabled or disabled). If spanning tree is enabled in the mesh, it must be the same version on all switches in the mesh: 802.1D, 802.1w, or 802.1s.
- If a switch in the mesh has GVRP enabled, then all switches in the mesh must have GVRP enabled. Otherwise, traffic on a dynamic VLAN may not pass through the mesh.

  If a switch in the mesh has a particular static VLAN configured, then all switches in the mesh must have that static VLAN configured.
- If a switch in the mesh has IGMP enabled, then all switches in the mesh must have IGMP enabled.
- If a switch in the mesh has LLDP enabled, then all switches in the mesh must have LLDP enabled.
- After adding or removing a port from the mesh, you must save the current configuration and reboot the switch in order for the change to take effect.
- Dynamic IP Lockdown and Virus Throttling should not be activated on mesh ports. These are security features for edge ports and mesh ports are not edge ports.

- DHCP Snooping and ARP protection are enabled through VLANs. Mesh ports belong to all VLANs, so if these security features are enabled on a switch that has mesh ports, the mesh ports must be configured as "trusted" ports because meshing may move the port of a MAC address in the mesh based on the least cost path.
- Multiple meshed domains require separation by either a non-meshed switch or a non-meshed link. For example:

Figure 23 Multiple meshed domains separated by a non-mesh switch or a non-mesh link

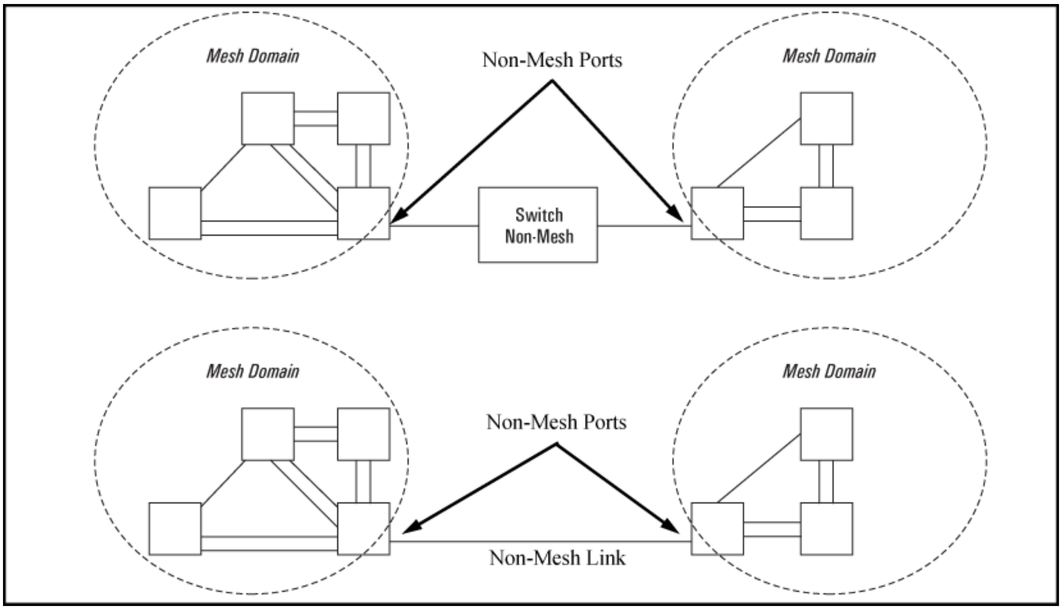

- If GVRP is enabled, meshed ports in a switch become members of any dynamic VLANs created in the switch in the same way that they would if meshing was not configured in the switch. (For more on GVRP, see "GVRP" (page 67).)

NOTE:

- A switch mesh domain (Example 121 (page 183)) cannot include either a switch that is not configured for meshing, or other sources of traffic.
- Where a given pair of switches are linked with meshed ports, you must not also link the pair together through non-meshed ports unless you have also enabled STP, RSTP, or MSTP to prevent a loop from forming.

Figure 24 A unsupported topology

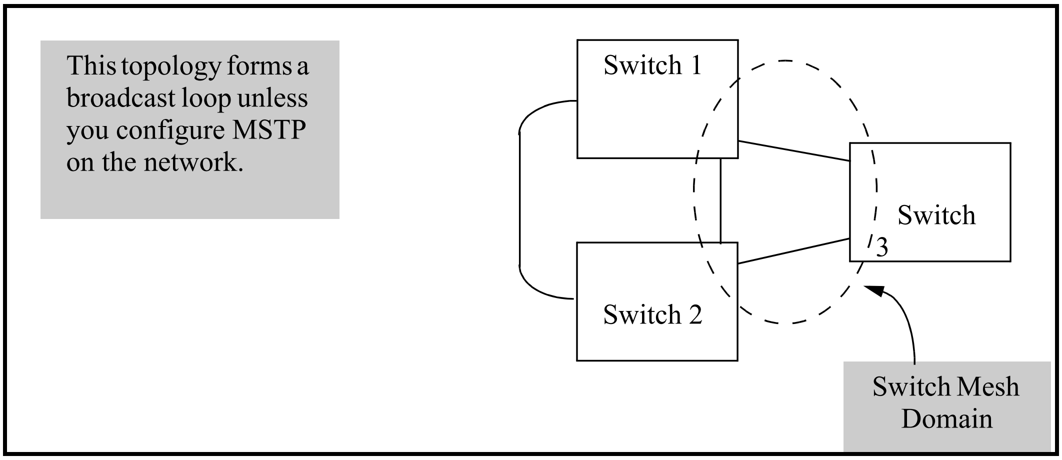
This topology forms a broadcast loop unless you configure MSTP on the network.
Switch 1
Switch 3
Switch 2
Switch Mesh Domain

- The switch blocks traffic on a meshed port connected to a non-meshed port on another switch.
- Switch meshing does not allow trunked links (LACP or Trunk) between meshed ports.

Linking a non-mesh device or port into the mesh causes the meshed switch ports connected to that device to shut down.

Using a heterogeneous switch mesh

You can use the switches covered in this guide with the HP Switch Series 5300xl in normal mode.

Example 134 A supported heterogeneous topology in normal mode

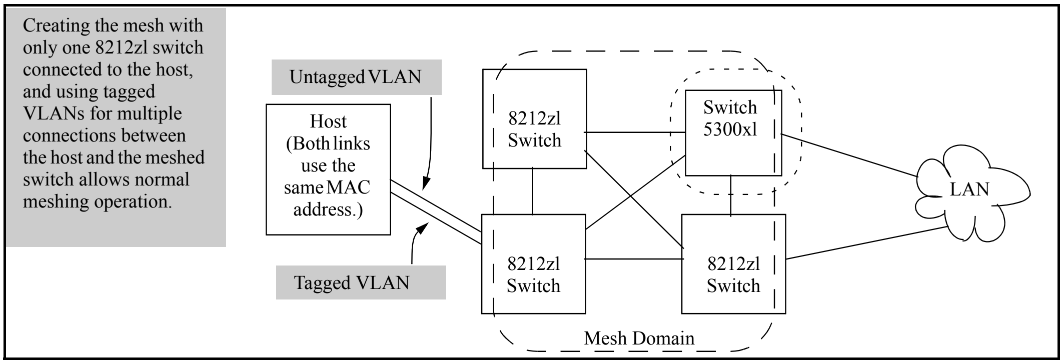
Creating the mesh with only one 8212zl switch connected to the host, and using tagged VLANs for multiple connections between the host and the meshed switch allows normal meshing operation.
Untagged VLAN
Host (Both links use the same MAC address.)
Tagged VLAN
8212zl Switch
Switch 5300xl
8212zl Switch
8212zl Switch
LAN
Mesh Domain

Bringing up a switch mesh domain

When a meshed port detects a non-meshed port on the opposite end of a point-to-point connection, the link will be blocked. Thus, as you bring up switch meshing on various switches, you may temporarily experience blocked ports where meshed links should be running. These conditions should clear themselves after all switches in the mesh have been configured for meshing and their switches rebooted. To reduce the effect of blocked ports during bring-up, configure meshing and reboot the switches before installing the meshed switches in the network. Also, since adding (or removing) a meshed port requires a switch reboot to implement, you can avoid repeated system disruptions by waiting to implement the mesh until you have finished configuring meshing on all ports in your intended mesh domain.

Further operating information

For more operating information, see “Operating notes for switch meshing” (page 197).

Operating notes for switch meshing

In a switch mesh domain traffic is distributed across the available paths with an effort to keep latency the same from path to path. The path selected at any time for a connection between a source node and a destination node is based on these latency and throughput cost factors:

- Outbound queue depth, or the current outbound load factor for any given outbound port in a possible path
- Port speed, such as 10Mbps versus 100Mbps; full-duplex or half-duplex
- Inbound queue depth, or how busy a destination switch is in a possible path
- Increased packet drops, indicating an overloaded port or switch

Pathshaving a lower cost will have more traffic added than those having a higher cost. Alternate paths and cost information is discovered periodically and communicated to the switches in the mesh domain. This information is used to assign traffic paths between devices that are newly active on the mesh. This means that after an assigned path between two devices has timed out, new traffic between the same two devices may take a different path than previously used.

To display information on the operating states of meshed ports and the identities of adjacent meshed ports and switches, see “Viewing switch mesh status (CLI)” (page 189).

Flooded traffic

Broadcast and multicast packets will always use the same path between the source and destination edge switches unless link failures create the need to select new paths. (Broadcast and multicast traffic entering the mesh from different edge switches are likely to take different paths.) When an edge switch receives a broadcast from a non-mesh port, it floods the broadcast out all its other non-mesh ports, but sends the broadcast out only those ports in the mesh that represent the path from that edge switch through the mesh domain. (Only one copy of the broadcast packet gets to each edge switch for broadcast out of its non-meshed ports. This helps to keep the latency for these packets to each switch as low as possible.)

Example 135 A broadcast path through a switch mesh domain

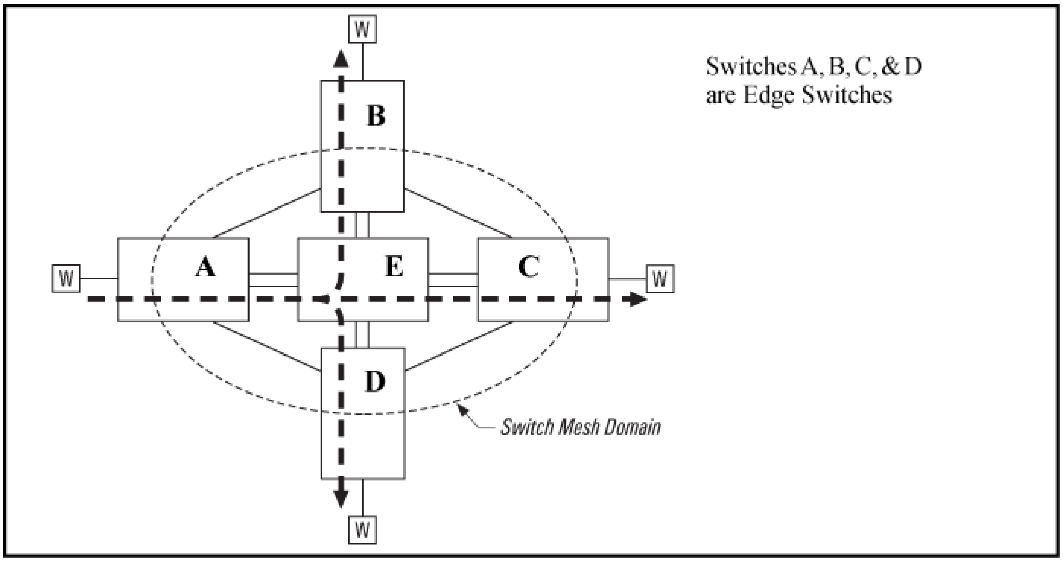

Any mesh switches that are not edge switches will flood the broadcast packets only through ports (paths) that link to separate edge switches in the controlled broadcast tree. The edge switches that receive the broadcast will flood the broadcast out all non-meshed ports. Some variations on broadcast/multicast traffic patterns, including the situation where multiple VLANs are configured and a broadcast path through the mesh domain leads only to ports that are in the same VLAN as the device originating the broadcast.

Unicast packets with unknown destinations

A meshed switch receiving a unicast packet with an unknown destination does not flood the packet onto the mesh. Instead, the switch sends a query on the mesh to learn the location of the unicast destination. The meshed switches then send 802.2 test packets through their non-meshed ports. After the unicast destination is found and learned by the mesh, subsequent packets having the same destination address will be forwarded. By increasing the `MAC Age Time` you can cause the switch address table to retain device addresses longer. (For more on `MAC Age Time`, see "System Information" in the chapter titled "Interface Access and System Information" in the *Management and Configuration Guide* for your switch.) Because the switches in a mesh exchange address information, this will help to decrease the number of unicast packets with unknown destinations, which improveslatency within the switch mesh. Also, in an IP environment, HP Networking recommends that you configure IP addresses on meshed switches. This makes the discovery mechanism more robust, which contributes to decreased latency.

Spanning tree operation with switch meshing

NOTE: Switch meshing cannot run concurrently with RPVST+.

Using MSTP with several switches and no switch meshing configured can result in unnecessarily blocking links and reducing available bandwidth. For example:

Example 136 Using STP without and with switch meshing

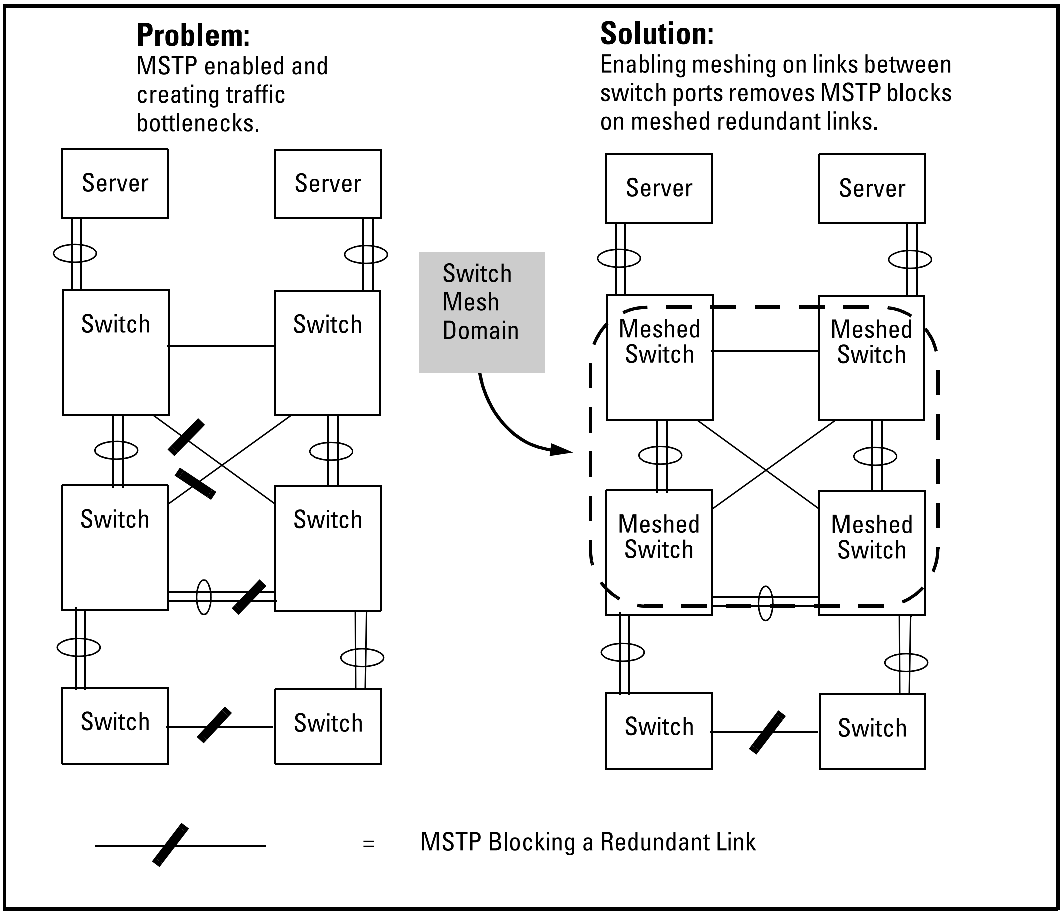
Problem:
MSTP enabled and creating traffic bottlenecks.
Solution:
Enabling meshing on links between switch ports removes MSTP blocks on meshed redundant links.
Server
Server
Switch
Switch
Switch
Switch
Switch
Switch
Switch Mesh Domain
Server
Server
Meshed Switch
Meshed Switch
Meshed Switch
Meshed Switch
Switch
Switch
= MSTP Blocking a Redundant Link

Example 137 Connecting a switch mesh domain to non-meshed devices

If you are going to use spanning tree in a switch mesh, all switches in the mesh should be configured with the same type of spanning tree: 802.1d/STP, 802.1w/RSTP, or 802.1s/MSTP. spanning tree interprets a meshed domain as a single link. However, on edge switches in the domain, MSTP will manage non-meshed redundant links from other devices. For example:

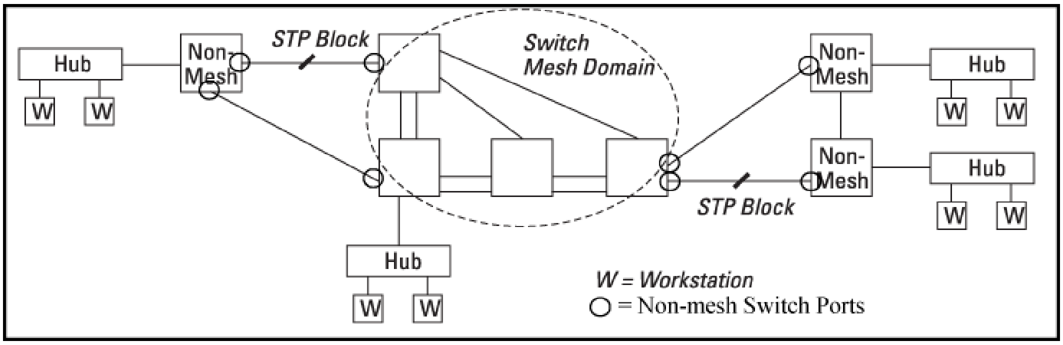

NOTE: When using MSTP and interconnecting switches covered in this guide in a mesh with switches that are not in the mesh, all the non-mesh switch ports (as indicated in the Example 137 "Connecting a switch mesh domain to non-meshed devices") should have the `edge-port` parameter disabled.

Example 138 Interconnecting switch mesh domains with redundant links

MSTP should be configured on non-mesh devices that use redundant links to interconnect with other devices or with multiple switch mesh domains. For example:

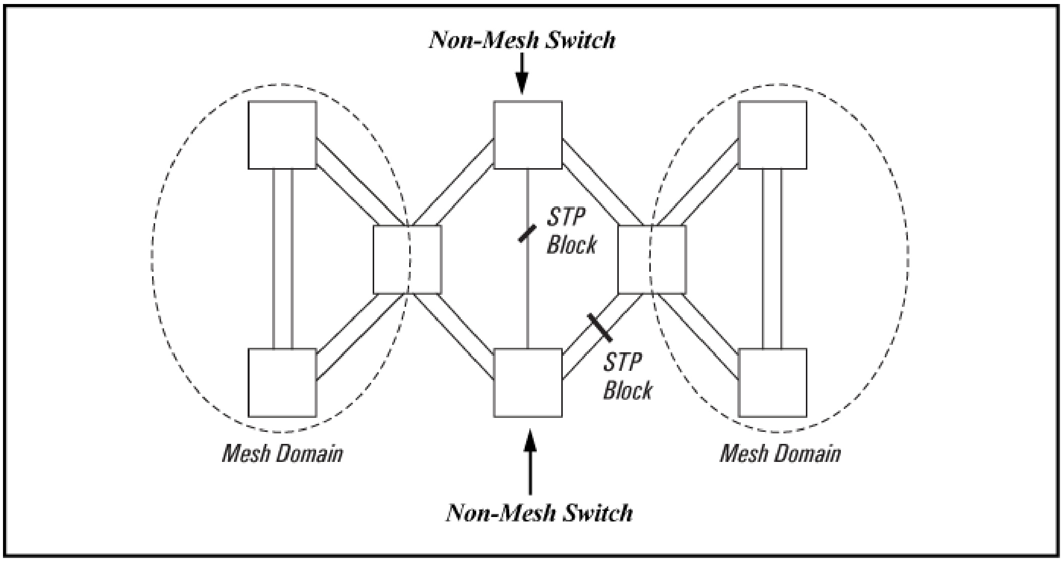

In the above case of multiple switch meshes linked with redundant trunks, there is the possibility that spanning tree will temporarily block a mesh link. This is because it is possible for spanning tree to interpret the cost on an external trunked link to be less than the cost on a meshed link. However, if this condition occurs, the meshed switch that has a blocked link will automatically increase the cost on the external (non-meshed) link to the point where spanning tree will block the external link and unblock the meshed link. This process typically resolves itself in approximately 30 seconds.

Filtering/security in meshed switches

Because paths through the mesh can vary with network conditions, configuring filters on meshed ports can create traffic problems that are difficult to predict, and is not recommended. However, configuring filters on non-meshed ports in an edge switch provides you with control and predictability.

IP Multicast (IGMP) in meshed switches

Like trunked ports, the switch mesh domain appears as a single port to IGMP. However, unlike trunked ports, IGMP protocol and multicast traffic may be sent out over several links in the mesh in the same manner as broadcast packets.

Static VLANs

In a network having a switch mesh domain and multiple static VLANs configured, all static VLANs must be configured on each meshed switch, even if no ports on the switch are assigned to any VLAN. (The switch mesh is a member of all static VLANs configured on the switches in the mesh.)

When static VLANs are configured, the mesh is seen as a single entity by each VLAN. All ports in the mesh domain are members of all VLANs and can be used to forward traffic for any VLAN. However, the non-mesh ports on edge switches that allow traffic to move between the mesh and non-meshed devices belong to specific VLANs and do not allow packets originating in a specific VLAN to enter non-meshed devices that do not belong to that same VLAN. (It is necessary to use a router to communicate between VLANs.) For example, in Figure 25 (page 201), traffic from host A entering the switch mesh can only exit the mesh at the port for hosts B and E. Traffic from host

A for any other host (such as C or D) will be dropped because only hosts B and E are in the same VLAN as host A.

Figure 25 VLAN operation with a switch mesh domain

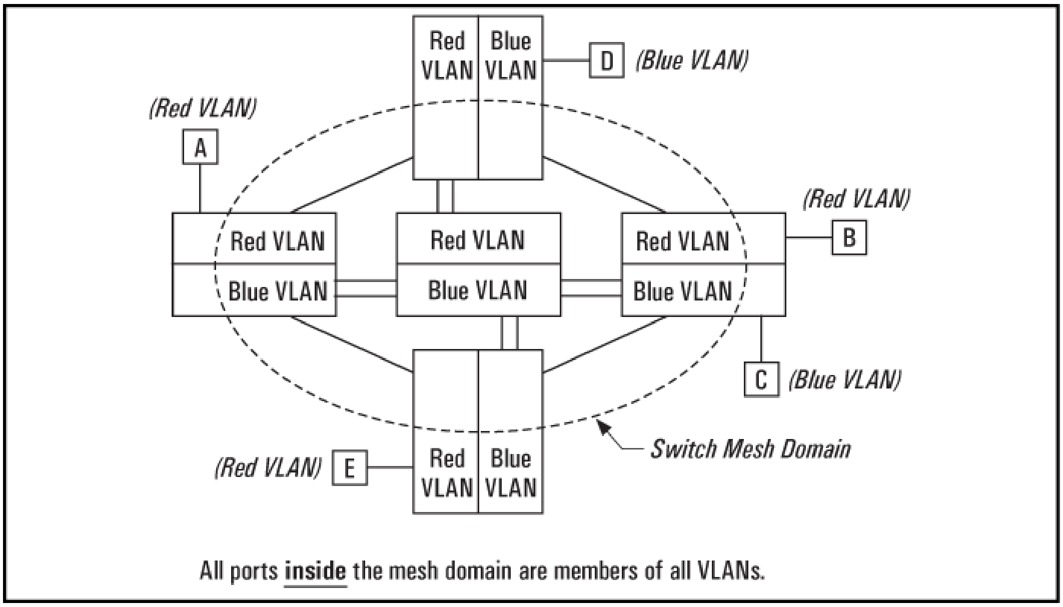

Dynamic VLANs

If GVRP is enabled, meshed ports in a switch become members of any dynamic VLANs created in the switch in the same way that they would if meshing was not configured in the switch. (For more on GVRP, see "GVRP" (page 67).)

Jumbo packets

If you enable jumbo traffic on any VLAN, then all meshed ports on the switch will be enabled to support jumbo traffic. (On a given meshed switch, every meshed port becomes a member of every VLAN configured on the switch.) If a port in a meshed domain does not belong to any VLANs configured to support jumbo traffic, then the port drops any jumbo packets it receives from other devices. In this regard, if a mesh domain includes any HP 8212zl switches, 6200yl switches, Series 5400zl switches, Series 3500yl switches, Series 3400cl or Series 6400cl switches that are configured to support jumbo traffic, only these switches can transmit and receive jumbo packets. Other switch models in the mesh will drop jumbo packets as they are not supported by those switches. See the chapter titled "Port Traffic Controls" in the *Management and Configuration Guide* for your switch.

Mesh design optimization

Mesh performance can be enhanced by using mesh designs that are as small and compact as possible while still meeting the network design requirements. The following are limits on the design of meshes and have not changed:

- Any switch in the mesh can have up to 24 meshed ports.
- A mesh domain can contain up to 12 switches.
- Up to 5 inter-switch meshed hops are allowed in the path connecting two nodes.
- A fully interconnected mesh domain can contain up to 5 switches.

Mesh performance can be optimized by keeping the number of switches and the number of possible paths between any two nodes as small as possible. As mesh complexity grows, the overhead associated with dynamically calculating and updating the cost of all of the possible paths between nodes grows exponentially. Cost discovery packets are sent out by each switch in the mesh every

30 seconds and are flooded to all mesh ports. Return packets include a cost metric based on inbound and outbound queue depth, port speed, number of dropped packets, etc. Also, as mesh complexity grows, the number of hops over which a downed link has to be reported may increase, thereby increasing the reconvergence time.

The simplest design is the two-tier design because the number of possible paths between any two nodes is kept low and any bad link would have to be communicated only to its neighbor switch.

Other factors affecting the performance of mesh networks include the number of destination addresses that have to be maintained, and the overall traffic levels and patterns. However, a conservative approach when designing new mesh implementations is to use the two-tier design and limit the mesh domain to eight switches where possible.

Example 139 A two-tier mesh design

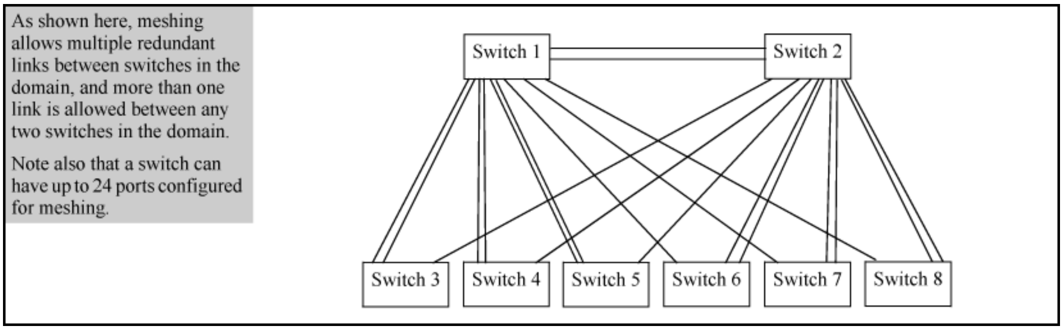

As shown here, meshing allows multiple redundant links between switches in the domain, and more than one link is allowed between any two switches in the domain.

Note also that a switch can have up to 24 ports configured for meshing.

Example 140 A fully interconnected mesh with the maximum switch count

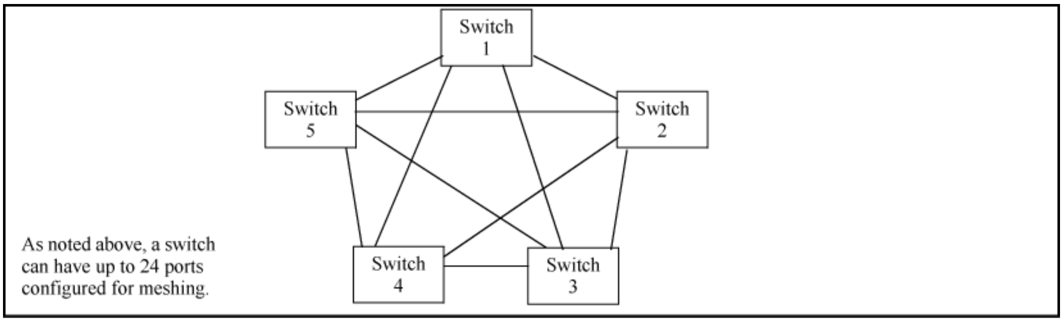

As noted above, a switch can have up to 24 ports configured for meshing.

Other factors affecting the performance of mesh networks include the number of destination addresses that have to be maintained, and the overall traffic levels and patterns. However, a conservative approach when designing new mesh implementations is to use the two-tier design and limit the mesh domain to eight switches where possible.

Configuring VRRP with concurrent meshing and routing

For more information, see the Virtual Router Redundancy Protocol (VRRP) chapter in the *Multicast and Routing Guide* for your switch.

- The VRRP configuration parameter `preempt-delay-time` should be set to at least 60 for each virtual router. This is because meshing is a distributed protocol that takes some time to stabilize when a new switch enters the mesh. A failing switch will be treated as new when it reenters the mesh. Since meshing is an L2 protocol, it must be stable before L3 VRRP can become active. The preempt-delay-time is only observed in the case where an already active

VRRP router exists. This will allow the existing router to continue serving its role until the preferred router, owner or higher priority backup, is truly ready to take over.

- The VRRP configuration parameter `nonstop` will not be configurable when VRRP is configured within a mesh. This is because the VRRP `nonstop` configuration attempts to make VRRP hitless in the case of a failed management module. Meshing, however, is not hitless in this case. Having the nonstop parameter set will cause the VRRP virtual router to ignore the `preempt-delay-time` and will have the virtual router attempt to become active before meshing is ready. This will result in a potential 30 or more second routed traffic gap while meshing becomes stable.
- For best network resiliency during a VRRP failover event, all switches in the mesh domain must be running x.15.09 or later version of the switch software. This is because changes in versions x.15.09 and later allow a VRRP virtual router MAC address to move from the master to the backup without being blocked by meshing on connected switches in the mesh.
- Using 5300 series switches in the same mesh domain that implements VRRP with concurrent meshing and routing is not recommended.

Other requirements and restrictions

- **Mesh support within the domain**

  All switches in the mesh domain, including edge switches, must support the HP witch meshing protocol.

- **Switch hop count in the mesh domain**

  A maximum of five (meshed) switch hops is allowed in the path connecting two nodes in a switch mesh domain. A path of six meshed hops is unusable. However, this does not interfere with other, shorter paths in the same domain.

- **Connecting mesh domains**

  To connect two separate switch meshing domains, you must use non-meshed ports. (The non-meshed link can be a porttrunk or a single link.) See Figure 23 (page 195).

- **Multiple links between meshed switches**

  Multiple mesh ports can be connected between the same two switches, to provide higher bandwidth. Each port that you want in the mesh domain should be configured as Mesh (and not as a trunk—Trk). If you configure a port as Mesh, there is no Type selection for that port.

- **Network monitor port**

  If a network monitor port is configured, broadcast packets may be duplicated on that port if more than one port is being monitored and switch meshing is enabled.

- **Compatibility with other switches**

  The switches covered in this guide operate with the Series 5300xl switches in normal mode.

- **Rate limiting not recommended on meshed ports**

  Rate limiting can reduce the efficiency of paths through a mesh domain.

See also “Operating rules” (page 194).

For additional information on troubleshooting meshing problems, see “Using a heterogeneous switch mesh” (page 196) and "Mesh-Related Problems" in Appendix C, "Troubleshooting" of the *Management and Configuration Guide* for your switch.

6 Quality of Service: Managing bandwidth effectively

<table>
<tr><th>Command Syntax</th><th>Description</th><th>Default value</th><th>CLI page reference</th></tr>
<tr><td>show qos <i>global-classifier</i></td><td>Displays a global QoS configuration</td><td></td><td>209</td></tr>
<tr><td>qos [<i>udp-port</i> | <i>tcp-port</i>] [ipv4 | ipv6 | ip-all] [<i>port-number</i> | range <i>start end</i>] priority <i>0 - 7</i></td><td>Assigns an 802.1p priority</td><td>Disabled</td><td>209</td></tr>
<tr><td>show qos tcp-udp-port-priority</td><td>Displays a list of all TCP and UDP QoS classifiers</td><td></td><td>210</td></tr>
<tr><td>qos dscp-map <i>codepoint</i> priority <i>0 - 7</i></td><td>Creates a DSCP policy</td><td>No override</td><td>210</td></tr>
<tr><td>show qos resources</td><td>Displays resource usage for policies</td><td></td><td>214</td></tr>
<tr><td>qos device-priority [<i>ipv4-address</i> | [ipv4] <i>ipv4-address/mask-length</i>] priority <i>0 - 7</i></td><td rowspan="2">Assigns a priority for a global IP-device classifier</td><td rowspan="2"></td><td rowspan="2">215</td></tr>
<tr><td>qos device-priority[<i>ipv6-address</i> | ipv6 <i>ipv6-address/mask-length</i> priority <i>0 - 7</i>]</td></tr>
<tr><td>qos dscp-map <i>codepoint</i> priority <i>0 - 7</i></td><td rowspan="3">Creates a policy based on IP address</td><td rowspan="3"></td><td rowspan="3">217</td></tr>
<tr><td>qos device-priority [<i>ipv4-address</i> | [ipv4] <i>ipv4-address/mask-length</i>] dscp <i>codepoint</i></td></tr>
<tr><td>qos device-priority <i>ipv6-address</i> | [ipv6] <i>ipv6-address/mask-length</i> dscp <i>codepoint</i></td></tr>
<tr><td>qos type-of-service ip-precedence</td><td>Assigns an 802.1p priority for a global IP-precedence classifier</td><td>Disabled</td><td>221</td></tr>
<tr><td>qos type-of-service diff-services <i>codepoint</i></td><td>Uses a global IP-Diffserv classifier to mark matching packets with an 801.p priority</td><td></td><td>222</td></tr>
<tr><td>qos type-of-service diff-services <i>current-codepoint</i> dscp <i>new-codepoint</i></td><td>Uses a global IP-Diffserv classifier to mark matching packets with a new DSCP policy</td><td></td><td>223</td></tr>
<tr><td>qos protocol [ip | ipx | arp | appletalk | sna | netbeui] priority <i>0 - 7</i></td><td>Assigns a priority for a global layer-3 protocol classifier</td><td>No-override</td><td>226</td></tr>
</table>

Command Syntax	Description	Default value	CLI page reference
vlan *vid* qos priority *0 - 7*	Assigns a priority for a global VLAN-ID classifier		227
qos dscp-map *codepoint* priority *0 - 7*	Creates a policy based on a VLAN-ID classifier	No-override	229
interface *port-list* qos priority *0 - 7*	Assigns a DSCP policy for a global source-port classifier	No-override	233
qos dscp-map *codepoint* priority *0 - 7*	Creates a policy based on source-port classifiers	No-override	234
Global configuration context: [no] class [ipv4 | ipv6 | *classname*]	Configures classifier-based QoS		237
Class configuration context: [no] [*seq-number*] [match | ignore]*ip-protocol* *source-address* *destination-address* [dscp *codepoint*] [precedence *precedence-value*] [tos *tos-value*] [vlan *vlan-id*]			
Global configuration context: [no] policy qos *policy-name*			
Policy configuration context: [no] [*seq-number*] class [ipv4 | ipv6 | *classname* action *qos-action*] [action *qosaction* ...]			
Global configuration context: [no] [*seq-number*]class [ipv4 | ipv6] *classname* action *qos-action* [action *qosaction* ...]	Configures QoS actions in a policy		240
Global configuration context: qos dscp-map *codepoint* priority *0 - 7*	Reconfigures the 802.1p priority value currently assigned to a DSCP codepoint		243
Policy configuration context: class ipv4 | ipv6 *classname* action dscp *codepoint* priority *0 - 7*			
show class ipv4 *classname*	Displays a classifier-based QoS configuration		244
show class ipv6 *classname*			
show class config			
qos dscp-map *codepoint* priority *0 - 7* [name *ascii-string*]	Configures Differentiated Services Codepoint (DSCP) mapping		249

Command Syntax	Description	Default value	CLI page reference
`qos queue-config [ 2-queues | 4-queues | 8-queues ]`	Configures QoS queues	8 queues	252
`show qos queue-config`	Displays the QoS queue configuration		253

Overview

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 Quality of Service (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

Quality of Service 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 Quality of Service 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 26 802.1p priority based on CoS (Class-of-Service) types and use of VLAN tags

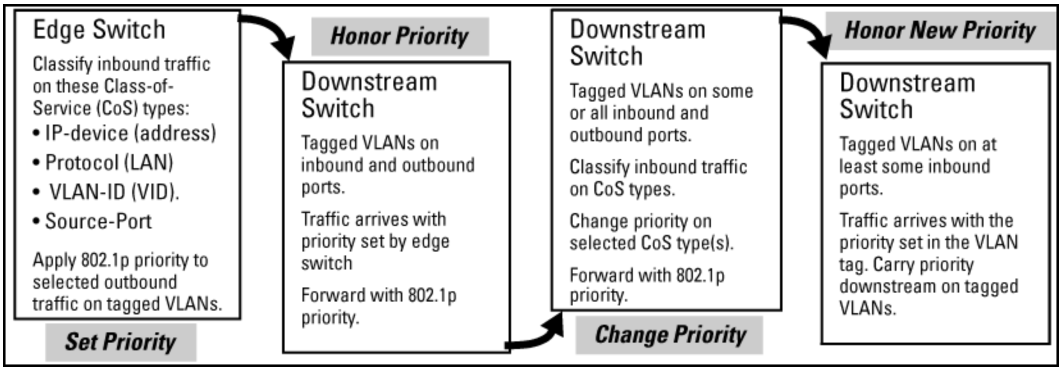

Edge Switch
Classify inbound traffic on these Class-of-Service (CoS) types:
• IP-device (address)
• Protocol (LAN)
• VLAN-ID (VID).
• Source-Port
Apply 802.1p priority to selected outbound traffic on tagged VLANs.
Set Priority
Honor Priority
Downstream Switch
Tagged VLANs on inbound and outbound ports.
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 27 Application of Differentiated Services Codepoint (DSCP) policies

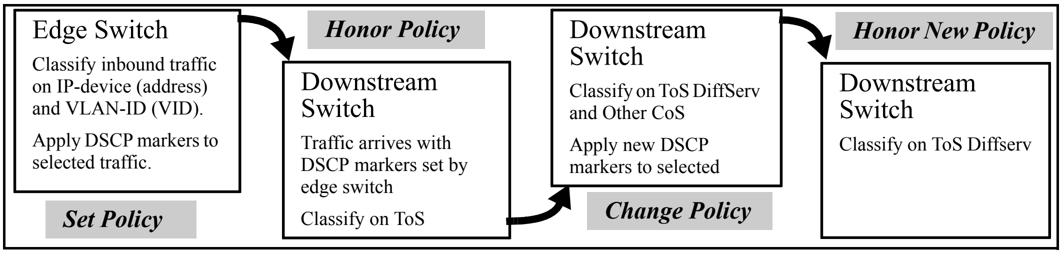
Edge Switch
Classify inbound traffic on IP-device (address) and VLAN-ID (VID).
Apply DSCP markers to selected traffic.
Set Policy
Honor Policy
Downstream Switch
Traffic arrives with DSCP markers set by edge switch
Classify on ToS
Downstream Switch
Classify on ToS DiffServ and Other CoS
Apply new DSCP 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.

Using classifier-based QoS to provide additional policy actions and aid migration in networks with legacy and OEM devices

Starting in software release K.14.01, HP Switch QoS configuration supports a classifier-based model that provides added functionality to create and manage QoS policies across a network consisting of HP switches as well as OEM and legacy devices.

The classifier-based configuration model is a single, simplified procedure and command syntax for cross-feature usage, which offers:

- Finer granularity than globally-configured QoS for classifying IPv4 and IPv6 traffic
- Additional actions for managing selected traffic, such as rate limiting and IP precedence marking
- The application of QoS policies to inbound traffic flows on specific port and VLAN interfaces (instead of using only globally-configured, switch-wide QoS settings)
- The use of configured traffic classes by different software features, such as QoS or port mirroring

Classifier-based QoS is designed to work with existing globally-configured, switch-wide QoS policies by allowing you to zoom in on a subset of port or VLAN traffic to further manage it.

Classifier-based policies take precedence over and may override globally-configured, switch-wide QoS settings.

Classifier-based QoS policies provide greater control for managing network traffic. Using multiple match criteria, you can finely select and define the classes of traffic that you want to manage. QoS-specific actions determine how you can handle the selected traffic.

Configuring QoS globally

To globally configure a QoS policy on the switch, follow these steps:

1. Determine the global QoS policy to implement on the switch by analyzing the types of traffic flowing through the network and identifying one or more traffic types to prioritize. The order of precedence in which global QoS classifiers are applied, from *a* (highest) to *h* (lowest), is as follows:
   - **a.** TCP/UDP applications.
   - **b.** Device priority—IP source or destination address. Destination has precedence over source, see Table 13 (page 208).
   - **c.** IP precedence bit set (leftmost three bits in the ToS/Traffic Class field of IP packets).
   - **d.** IP differentiated services bit set (leftmost six bits in the ToS/Traffic Class field of IP packets).
   - **e.** Layer-3 protocol.
   - **f.** VLAN ID. At least one tagged VLAN is required on the network.
   - **g.** Source port.
   - **h.** Incoming 802.1p priority (requires at least one tagged VLAN on the network).

     Default: In a tagged VLAN environment, the incoming 802.1p priority is used as the default QoS classifier if no global QoS classifier with a higher precedence matches.

2. Select the global QoS classifier to use. The following table shows the types of QoS marking (802.1p priority or DSCP codepoint) supported by each global QoS classifier.

Table 13 QoS marking supported by global 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 (see Table 16 (page 258)). 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. For 802.1p priority settings to be included in outbound packets, ensure that tagged VLANs are configured on the appropriate downstream links.
4. Determine the global QoS policy required for each QoS-capable device in the network and configure the necessary settings.

   For downstream devices to recognize and use DSCP codepoints in IP packets sent from the switch, enable ToS (Type-of-Service) Differentiated Service mode on the devices and configure the appropriate DSCP policies. Note that certain DSCP policies have a default 802.1p priority automatically assigned (see Table 24 (page 275)).

NOTE: For more information on how to use global QoS classifiers, see "Global QoS restrictions" (page 261).

Viewing a global QoS configuration

To display the existing switch-wide configurations for a global QoS classifier, use one of the following `show qos` commands.

Syntax:

`show qos` *`global-classifier`*

`tcp-udp-port-priority`

Displays the current TCP/UDP port priority configuration, see Figure 29 (page 214).

`device-priority`

Displays the current device (IP address) priority configuration, see Example 143 (page 217).

`type-of-service`

Displays the current type-of-service priority configuration. The display output differs according to the option used for IP Precedence, see Figure 32 (page 222). See also Figure 33 (page 223).

`protocol-priority`

Displays the current protocol priority configuration.

`vlan-priority`

Displays the current VLAN priority configuration.

`port-priority`

Displays the current source-port priority configuration, see Figure 40 (page 233).

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. See "Operating notes on using TCP/UDP port ranges" (page 263). If you specify a range, the minimum port number must precede the maximum port number in the range.

`priority` *`0 - 7`*

Marks the specified 802.1p priority in matching TCP or UDP packets.

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.

Displaying a list of all TCP and UDP QoS classifiers

Syntax:

```
show qos tcp-udp-port-priority
```

Displays a listing of all TCP and UDP QoS classifiers currently in the running-config file.

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:

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 Figure 50 (page 264), above).
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 "Using Differentiated Services Codepoint (DSCP) mapping" (page 249).)
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.

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 (`show qos dscp-map command)`, you must first configure a priority for the codepoint before proceeding (`qos dscp-map priority` `command)`. See "Using Differentiated Services Codepoint (DSCP) mapping" (page 249).

Syntax:

`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 (see Table 14 (page 250)).

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
`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. See "The default DSCP policy table" (page 250).

4. Configure the switch to assign the DSCP policy to packets with the specified TCP or UDP port number or range of port numbers.

   Syntax:

   ```
   qos [ udp-port | tcp-port ][ ipv4 | ipv6 | ip-all ]
   [ port-number | range  start end ] dscp codepoint
   ```

   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:

   - ipv4 marks only IPv4 packets (default).
   - ipv6 marks only IPv6 packets.
   - ip-all marks all IP traffic (both IPv4 and IPv6 packets).
   - port-number specifies a TCP/UDP port-number from 1 to 65535.
   - range start end specifies a range of TCP/UDP ports; see "Operating notes on using TCP/UDP port ranges" (page 263). 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 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 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 (see Table 22 (page 272)). 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.

Assigning DSCP policies to packets matching specified TCP and UDP port applications (Example)

Port Applications	DSCP Policies	
	DSCP	**Priority**
23-UDP	000111	7
80-TCP	000101	5
914-TCP	000010	1
1001-UDP	000010	1

1. Determine whether 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).

   NOTE: A DSCP codepoint must also have a priority configured before you can use it to mark matching packets.

   Example 141 Displaying the current DSCP-map configuration

   ```
   HP 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 trafic.

     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.

Figure 28 Assigning priorities to the selected DSCPs

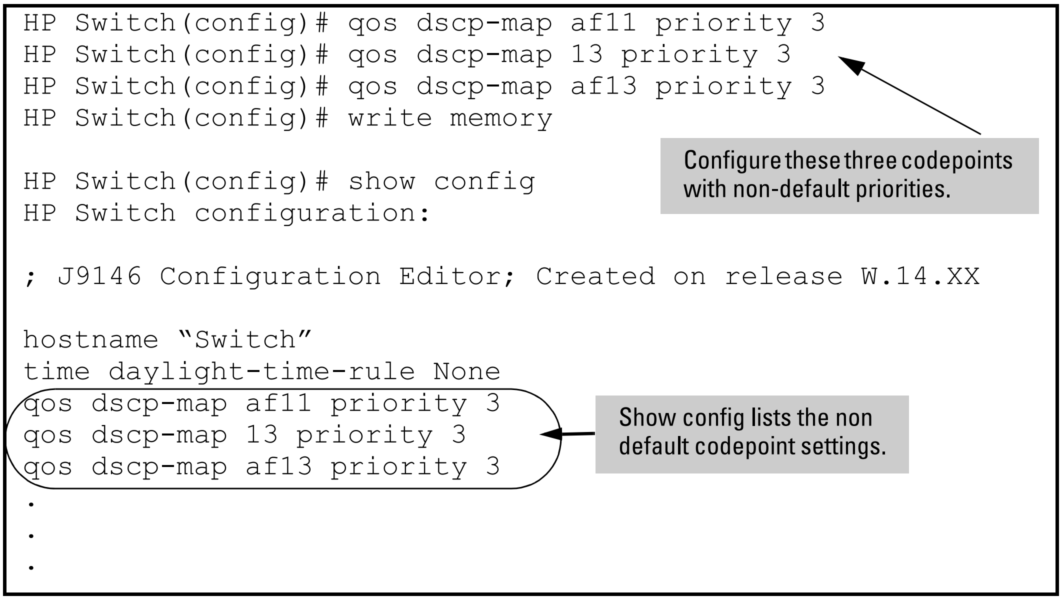

```
HP Switch(config)# qos dscp-map af11 priority 3
HP Switch(config)# qos dscp-map 13 priority 3
HP Switch(config)# qos dscp-map af13 priority 3
HP Switch(config)# write memory

HP Switch(config)# show config
HP Switch configuration:

; J9146 Configuration Editor; Created on release W.14.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
.
.
.
```

Configure these three codepoints with non-default priorities.

Show config lists the non default codepoint settings.

3. Assign the DSCP policies to the selected TCP/UDP port applications and display the result.

Figure 29 Configuring a DSCP policy for global TCP/UDP port classifiers

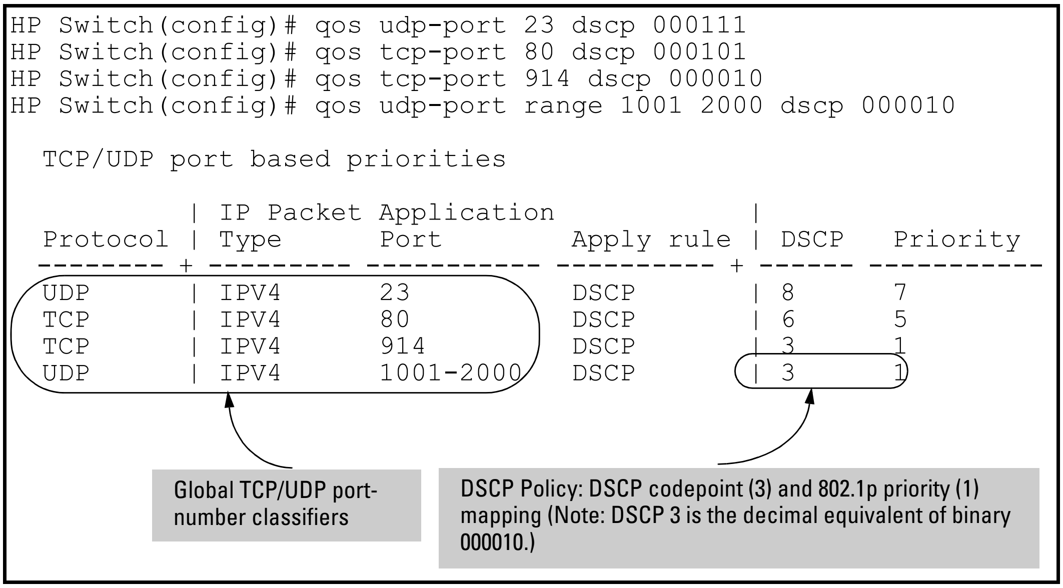

```
HP Switch(config)# qos udp-port 23 dscp 000111
HP Switch(config)# qos tcp-port 80 dscp 000101
HP Switch(config)# qos tcp-port 914 dscp 000010
HP Switch(config)# qos udp-port range 1001 2000 dscp 000010

  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
```

Global TCP/UDP port-number classifiers

DSCP Policy: DSCP codepoint (3) and 802.1p priority (1) mapping (Note: DSCP 3 is the decimal equivalent of binary 000010.)

The switch applies the DSCP policies in Figure 29 (page 214) to IP packets with the specified TCP/UDP port applications that are received in the switch. The switch manages the packets as follows:

1. Overwrites the original DSCPs in the selected packets with the new DSCPs specified in the above policies.
2. Assigns the 802.1p priorities in the above policies to the selected packets.

Displaying resource usage for QoS policies

When configurng global QoS classifiers using TCP/UDP and a Layer 4 Application port number or port range, the switch automatically assigns two QoS resources for each policy—one for traffic to the TCP/UDP destination port and one for traffic to the TCP/UDP source port.

The `show qos resources` command displays the number of hardware resources currently in use by QoS policies as well as other software features.

Example 142 Displaying the hardware resources used by currently configured QoS policies

```
HP Switch(config)# show qos resources

 Resource usage in Policy Enforcement Engine

       |    Rules     |  Rules Used
 Slots |  Available   | ACL | QoS | IDM |  VT | Mirror | PBR | Other |
 ------+--------------+-----+-----+-----+-----+--------+-----+-------|
   A   |         3014 |  15 |  11 |   0 |   1 |      0 |   0 |   3   |

       |    Meters    |  Meters Used
 Slots |  Available   | ACL | QoS | IDM |  VT | Mirror | PBR | Other |
 ------+--------------+-----+-----+-----+-----+--------+-----+-------|
   A   |          250 |     |   5 |   0 |     |        |     |     0 |

       | Application  |
       | Port Ranges  |  Application Port Ranges Used
 Slots |  Available   | ACL | QoS | IDM |  VT | Mirror | PBR | Other |
 ------+--------------+-----+-----+-----+-----+--------+-----+-------|
   A   |           14 |   2 |   0 |   0 |     |      0 |   0 |     0 |

 0 of 8 Policy Engine management resources used.

 Key:
 ACL = Access Control Lists
 QoS = Device & Application Port Priority, QoS Policies, ICMP rate limits
 IDM = Identity Driven Management
 VT  = Virus Throttling blocks
 Mirror = Mirror Policies, Remote Intelligent Mirror endpoints
 PBR = Policy Based Routing Policies
 Other = Management VLAN, DHCP Snooping, ARP Protection, Jumbo IP-MTU,
         Transparent Mode.

 Resource usage includes resources actually in use, or reserved for future
 use by the listed feature.  Internal dedicated-purpose resources, such as
 port bandwidth limits or VLAN QoS priority, are not included.
```

NOTE: ACLs and QoS policies share the same application port ranges. If a new QoS policy specifies a port range that is already configured for one or more ACLs, the QoS column increases by 1, but the **Application Port Ranges Available** column remains unchanged. Likewise, if an ACL is configured for a port range on which a QoS policy is already applied, the ACL column increases by 1, while the **Available** column remains unchanged.

Similarly, when you remove a port range, the **Application Port Ranges Available** column increases only if the port range is not configured for an existing ACL or QoS policy on the switch.

Assigning a priority for a global IP-device classifier

This global QoS packet-marking option assigns an 802.1p priority to all IP packets that have the specified IP address as either a source or destination. If both the source and destination addresses match, the priority configured for the IP destination address has precedence.

Syntax:

qos device-priority[*ipv4-address* | [ipv4]*ipv4-address/mask-length*]
priority *0 - 7*

qos device-priority[*ipv6-address* | ipv6 *ipv6-address/mask-length*]
priority *0 - 7*

Marks an 802.1p priority in outbound packets with the specified IP address or subnet mask in the source or destination field in a packet header, where:

- *`ipv4-address`* or *`ipv6-address`* is an IPv4 or IPv6 address used to match the source or destination address in packet headers.

  NOTE: An IPv6 local-link address (such as `fe80::110:252%vlan20`) that is automatically generated on a VLAN interface is not supported as an `ipv6-address` value.

- `[ipv4]` *`ipv4-address/mask-length`* is the subnet identified by the IPv4 mask for the specified address that is used to match the IPv4 in the source or destination field of packet headers.
- `ipv6` *`ipv6-address/prefix-length`* is the subnet identified by the IPv6 prefix-length for the specified address that is used to match the IPv6 address in the source or destination field of packet headers.

  Enter the IPv4 mask or IPv6 prefix length with an address in CIDR format by using the number of significant bits (for example, `2001:db8::1:262:a03:e102:127/64` or `10.28.31.1/24`).

- `priority` *`0 - 7`* marks the specified 802.1p priority in matching IP packets.

  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.

  The `no` form of the command deletes the specified IP address or subnet mask as a QoS classifier, and resets the priority for the VLAN to `No-override`.

`show qos device-priority`

Displays a listing of all IP device-priority QoS configurations currently in the running-config file.

Examples

Example 143 Configuring and displaying 802.1p priority

This example shows how to configure and display the 802.1p priority used to mark packets that match each global IP-device classifier:

IP Address / Mask or Prefix Length	802.1p Priority
10.28.31.1	7
10.28.31.130	5
10.28.31.100/24	1
2001:db8:2:1:212:79ff:fe88:a100	3
2001:db8:3:3::/64	1

```
HP Switch(config)# qos device-priority 10.28.31.1 priority 7
HP Switch(config)# qos device-priority 10.28.31.130 priority 5
HP Switch(config)# qos device-priority ipv4 10.28.32.100/24 priority 1
HP Switch(config)# qos device-priority 2001:db8:2:1:212:79ff:fe88:a100 priority
HP Switch(config)# qos device-priority ipv6 2001:db8:3:3::/64 priority 1
HP Switch(config)# show qos device-priority

  Device priorities

  Device Address                                Apply rule | DSCP  Priority
  --------------------------------------------  ---------- + ------ -----------
  10.28.31.1                                    Priority   |        7
  10.28.31.130                                  Priority   |        5
  10.28.32.100/24                               Priority   |        1
  2001:db8:2:1:212:79ff:fe88:a100               Priority   |        3
  2001:db8:3:3::/64                             Priority   |        1
```

Assigning a DSCP policy for a global IP-device classifier

This global QoS packet-marking option assigns a previously configured DSCP policy (codepoint and 802.1p priority) to outbound IP packets having the specified IP address or subnet mask in the source or destination field of their packet header. The switch:

1. Selects an incoming IPv4 or IPv6 packet on the basis of the source or destination IP address or subnet mask it carries.
2. Overwrites the DSCP in matching packets with the globally configured DSCP codepoint, and assigns the 802.1p priority associated with the new DSCP. For more information, see "Using Differentiated Services Codepoint (DSCP) mapping" (page 249).
3. Forwards the packet through the appropriate outbound port queue.

Creating a policy based on IP address

This procedure creates a DSCP policy for IP packets carrying the selected IP address (source or destination).

1. Identify the IPv4 or IPv6 address to use as a classifier for assigning a DSCP policy.
2. Determine the DSCP policy for packets carrying the selected IP address:
   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 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.

   Prerequisite: A DSCP codepoint must have a preconfigured 802.1p priority (0 - 7) to use the codepoint to mark matching packets. If a codepoint shows `No-override` in the Priority column of the DSCP Policy table (`show qos dscp-map`command), first configure a priority for the codepoint before proceeding (`qos dscp-map priority` command). See "Using Differentiated Services Codepoint (DSCP) mapping" (page 249).

   ### *Syntax:*

   `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, see Table 14 (page 250).

   When the switch applies this 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. See "The default DSCP policy table" (page 250).

4. Configure the switch to assign the DSCP policy to packets with the specified IP address or subnet mask.

   Syntax:

   ```
   qos device-priority[ ipv4-address | [ipv4 ipv4-address/mask-length]]
   dscp codepoint
   ```

   ```
   qos device-priority[ ipv6-address | ipv6 ipv6-address/mask-length ]
   dscp codepoint
   ```

   Assigns a DSCP policy in packets with the specified IP address or subnet mask in the source or destination field in a packet header, where:

   - *ipv4-address* or *ipv6-address* is an IPv4 or IPv6 address used to match the source or destination address in packet headers.

     NOTE: An IPv6 local-link address (such as `fe80::110:252%vlan20`) that is automatically generated on a VLAN interface is not supported as an `ipv6-address` value.

   - `[ipv4]` *ipv4-address/mask-length* is the subnet identified by the IPv4 mask for the specified address that is used to match the IPv4 in the source or destination field of packet headers.
   - `ipv6 ipv6-address/prefix-length` is the subnet identified by the IPv6 prefix-length for the specified address that is used to match the IPv6 address in the source or destination field of packet headers.

     Enter the IPv4 mask or IPv6 prefix length with an address in CIDR format by using the number of significant bits (for example, `2001:db8:2:1:262:a03:e102:127/64` or `10.28.31.1/24`).

   - `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 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 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 (see Table 22 (page 272)). 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 IP address or subnet mask as a QoS classifier. If you configured a subnet mask as match criteria, you must enter the entire IP address and mask-length in the `no` command.

   Syntax:

   ```
   show qos device-priority
   ```

   Displays a listing of all IP addresses and subnet masks used as QoS classifiers currently in the running-config file.

Creating a policy based on IP address 219

Assigning DSCP policies to packets matching specified global classifiers

This example shows how to assign the following DSCP policies to the packets that match the specified global IP-device classifiers:

IP address	DSCP Policy	
	DSCP codepoint	802.1p priority
10.28.31.1	000111	7
10.28.31.130	000101	5
10.28.31.100/24	000010	1
2001:db8:2:1:212:79ff:fe88:a100	000101	3
2001:db8:3:3::/64	000010	1

1. Determine whether 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). This is not a problem if the configured priorities are acceptable for all applications that use the same DSCP (see "Notes on changing priority settings" (page 274)).

   Note that a DSCP codepoint must have an associated priority before you can use it to mark matching packets.

   Figure 30 Display the current DSCP-map configuration

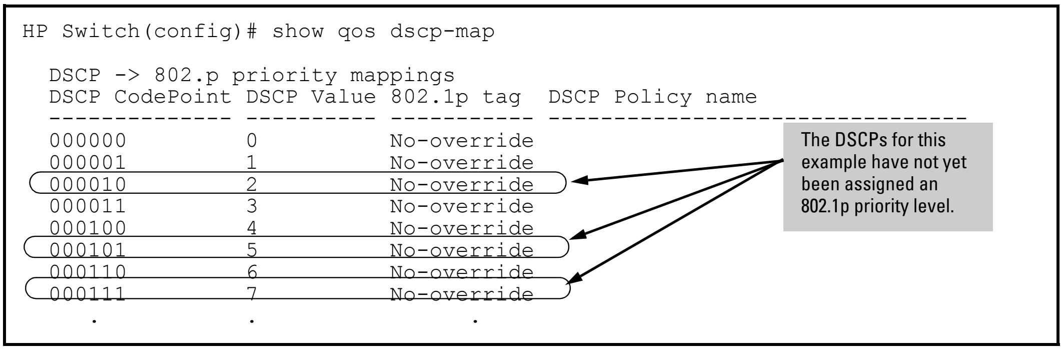

```
HP Switch(config)# show qos dscp-map

  DSCP -> 802.p priority mappings
  DSCP CodePoint DSCP Value 802.1p tag  DSCP Policy name
  -------------- ---------- ----------- --------------------------------
  000000         0          No-override
  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
     .             .                 .
```

The DSCPs for this example have not yet been assigned an 802.1p priority level.

2. Configure the priorities for the DSCPs you want to use to mark packets.

   Figure 31 Assigning 802.1p priorities to the selected DSCPs

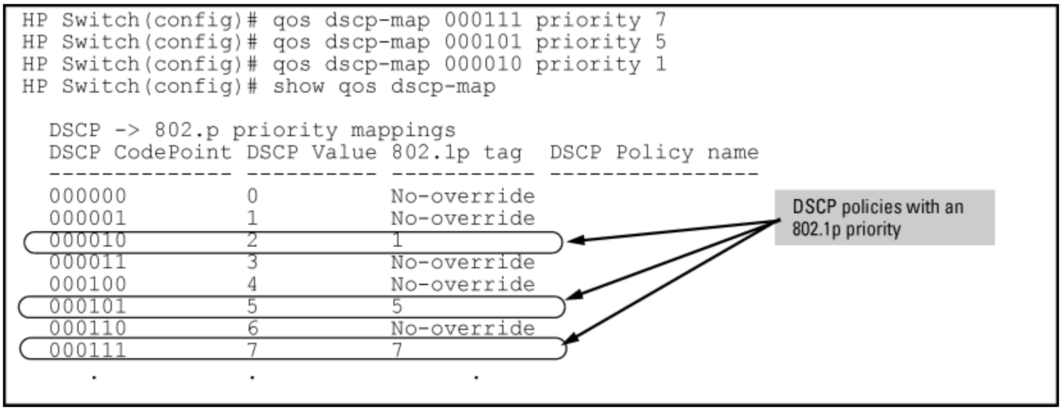

```
HP Switch(config)# qos dscp-map 000111 priority 7
HP Switch(config)# qos dscp-map 000101 priority 5
HP Switch(config)# qos dscp-map 000010 priority 1
HP Switch(config)# show qos dscp-map

  DSCP -> 802.p priority mappings
  DSCP CodePoint DSCP Value 802.1p tag  DSCP Policy name
  -------------- ---------- ----------- ----------------
  000000         0          No-override
  000001         1          No-override
  000010         2          1
  000011         3          No-override
  000100         4          No-override
  000101         5          5
  000110         6          No-override
  000111         7          7
     .             .                 .
```

DSCP policies with an 802.1p priority

3. Assign the DSCP policies to the specified IP-device addresses and display the result.

```
HP Switch(config)# qos device-priority 10.28.31.1 dscp 000111
HP Switch(config)# qos device-priority 10.28.31.130 dscp 000101
HP Switch(config)# qos device-priority ipv4 10.28.32.100/24 dscp 000010
HP Switch(config)# qos device-priority 2001:db8:2:1:212:79ff:fe88:a100 dscp 000
HP Switch(config)# qos device-priority ipv6 2001:db8:3:3/64 dscp 000010
HP Switch(config)# show qos device-priority

  Device priorities

  Device Address                                Apply rule | DSCP  Priority
  --------------------------------------------  ---------- + ------ -----------
  10.28.31.1                                     Priority  | 000111 7
  10.28.31.130                                   Priority  | 000101 5
  10.28.32.100/24                                Priority  | 000010 1
  2001:db8:2:1:212:79ff:fe88:a100                Priority  | 000101 3
  2001:db8:3:3/64                                Priority  | 000010 1
```

The switch applies the DSCP policies in Figure 31 (page 220) to IP packets with the specified IP addresses and subnet masks (source or destination) 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 appropriate packets.

Assigning an 802.1p priority for a global IP-precedence classifier

If a device or application upstream of the switch sets the precedence bits in the ToS/Traffic Class byte of IPv4/IPv6 packets, you can use this global packet-marking option to prioritize 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 IP packets (IPv4 and IPv6) by computing a 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 Type-of-Service option is configured), this command displays the ToS configuration status. If the Diff-serv mode is enabled, codepoint data is displayed as described in "Assigning a DSCP policy for a global IP-Diffserv classifier" (page 223).

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 32 Enabling ToS IP-precedence prioritization

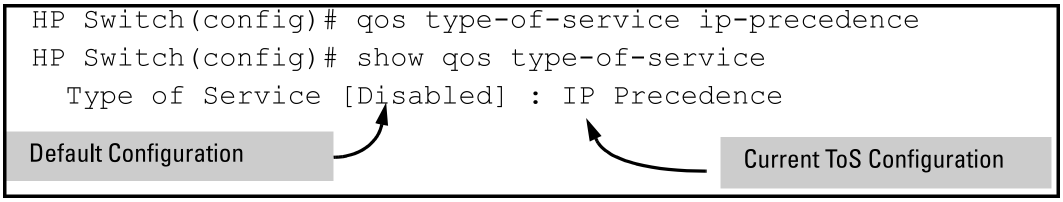

```
HP Switch(config)# qos type-of-service ip-precedence
HP Switch(config)# show qos type-of-service
  Type of Service [Disabled] : IP Precedence
```

To change from IP-precedence to IP-Diffserv mode, follow the procedure in "Assigning a priority for a global IP-device classifier" (page 215), which automatically disables IP-Precedence. To disable IP-Precedence without enabling the IP-Diffserv option, enter the `no qos type-of-service` command.

Using a global IP-Diffserv classifier to mark matching packets with an 802.1p priority

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. 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.
4. Enable IP-Diffserv mode by entering the `qos type-of-service diff-services` command.

Syntax:

`qos type-of-service diff-services` *`codepoint`*

Causes the switch to read the *`codepoint`* (DSCP) of an incoming IP packet and, when a match occurs, assign the associated 802.1p priority in the DSCP Policy table (see "The default DSCP policy table" (page 250)).

`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`. Note that 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 33 (page 223) you cannot change the priority for the 000000 codepoint until you redirect the DSCP policy for 000001 from using 000000 as a policy.

For more information see "Notes on changing priority settings" (page 274) and "Using Differentiated Services Codepoint (DSCP) mapping" (page 249).

`show qos type-of-service`

Displays the current Type-of-Service configuration. In IP-Diffserv mode it also shows the current direct 802.1p assignments and the current DSCP assignments covered later in this section.

Examples

An edge switch A in an untagged VLAN assigns a DSCP of 000110 on IP packets it receives on port A6, 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`, and then enable `diff-services`:

Figure 33 Displaying the codepoints available for 802.1p priority assignments

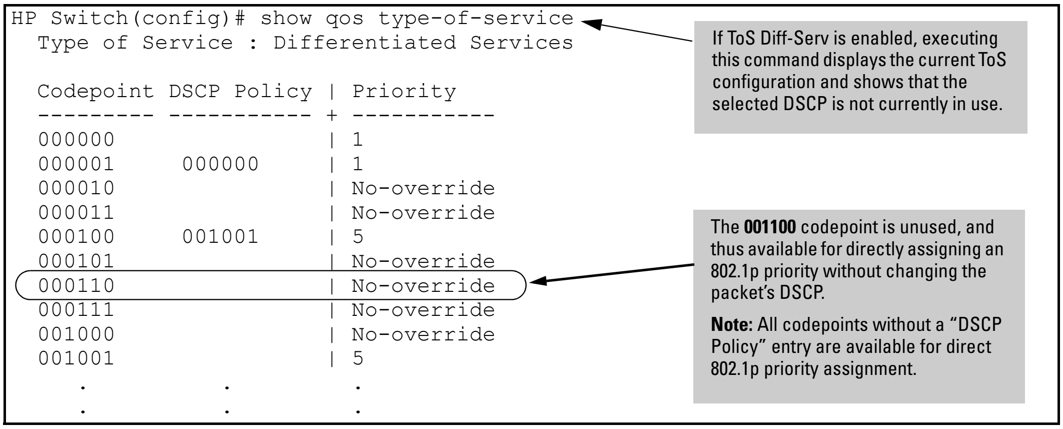

```
HP Switch(config)# show qos type-of-service
  Type of Service : Differentiated Services

  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 34 Type-of-Service configuration that enables both 802.1p priority and DSCP policy assignment

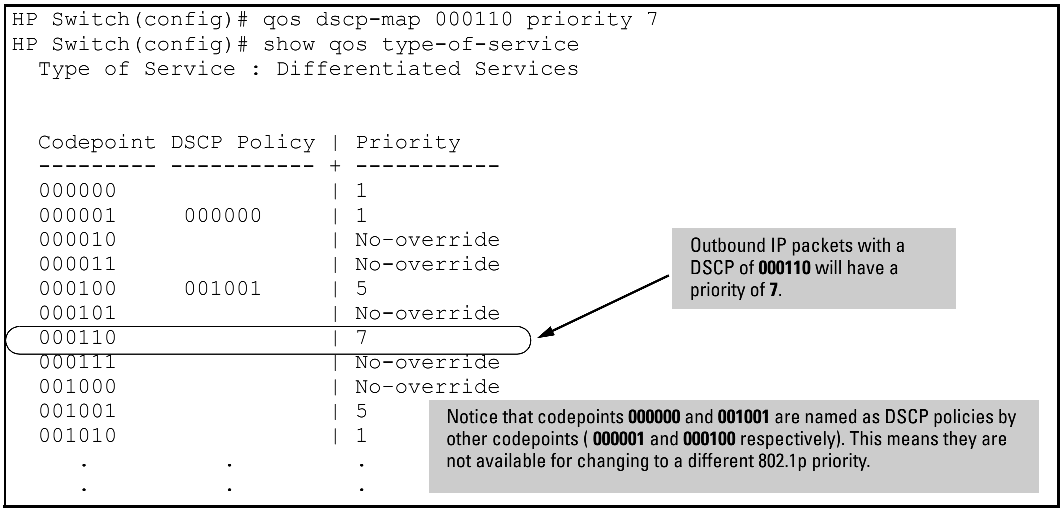

```
HP Switch(config)# qos dscp-map 000110 priority 7
HP Switch(config)# show qos type-of-service
  Type of Service : Differentiated Services


  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 for a global IP-Diffserv classifier

The preceding section describes how to forward an 802.1p priority level set by an edge (or upstream) switch. This section describes how to use a global IP-Diffserv classifier to mark matching packets with a new DSCP policy. A DSCP policy consists of a DSCP codepoint and an associated 802.1p priority.

You can use a global IP-Diffserv classifier to mark a DSCP policy at the same time with a global IP-Diffserv classifier that marks an 802.1p priority if different DSCP codepoints are configured with each classifier.

To use a global IP-Diffserv classifier to mark matching packets with a new DSCP policy, follow these steps:

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. (For more information, see "Creating a service policy" (page 377))

   .

3. Use the `qos type-of-service diff-services` *`incoming-DSCP`* `dscp` *`outgoing-DSCP`* command to change the policy on packets coming from the edge or upstream switch with the specified incoming DSCP.

   Figure 52 (page 268) illustrates this scenario.

Syntax:

`qos type-of-service diff-services`

Enables ToS Diff-services.

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.

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

  Enter `?` to display the list of valid codepoint entries.

  To reconfigure the 802.1p priority currently assigned to a DSCP codepoint, use the `qos dscp-map` command as described in "Using Differentiated Services Codepoint (DSCP) mapping" (page 249).

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 policy assigned to it.

Example 144 Configuring new DSCP policies

The following example shows how to configure new DSCP policies on matching packets with the specified DSCP codepoints.

Received DSCP	Policy DSCP	802.1p Priority	Policy Name (Optional)
001100	000010	6	Level 6
001101	000101	4	Level 4

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). This is not a problem as long as the configured priorities are acceptable for all applications using the same DSCP, see "Notes on changing priority settings" (page 274).

   Also, note that 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 (`show qos dscp-map` command), you must first configure a priority for the codepoint before proceeding (`qos dscp-map priority` command). See "Using Differentiated Services Codepoint (DSCP) mapping" (page 249).

   ### Figure 35 Displaying the current DSCP-map configuration

   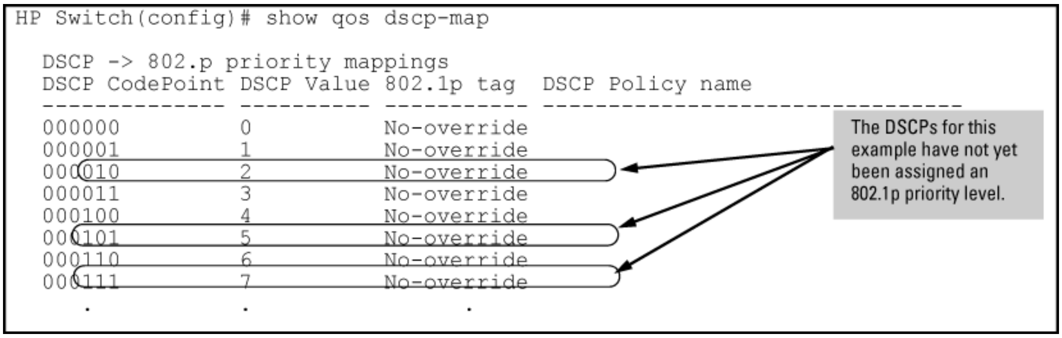

```
HP Switch(config)# show qos dscp-map

  DSCP -> 802.p priority mappings
  DSCP CodePoint DSCP Value 802.1p tag  DSCP Policy name
  -------------- ---------- ----------- --------------------------------
  000000         0          No-override
  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
     .           .                .
```

   The DSCPs for this example have not yet been assigned an 802.1p priority level.

2. Configure the desired policies (codepoint and associated 802.1p priority) in the DSCP table:

   ### Example 145 Configuring DSCP policies in the DSCP table

```
HP Switch(config)# qos dscp-map 000010 priority 6 name 'Level 6'
HP Switch(config)# qos dscp-map 000101 priority 4 name 'Level 4'
HP Switch(config)# show qos dscp-map
  DSCP -> 802.p priority mappings
  DSCP policy 802.1p tag Policy name
  ----------- ---------- -----------------
  000000      No-override
  000001      No-override
  000010      6          Level 6
  000011      No-override
  000100      No-override
  000101      4          Level 4
  000110      No-override
  000111      No-override
     .            .          .
     .            .          .
     .            .          .
```

3. Assign the new policies to mark matching packets with the specified codepoints.

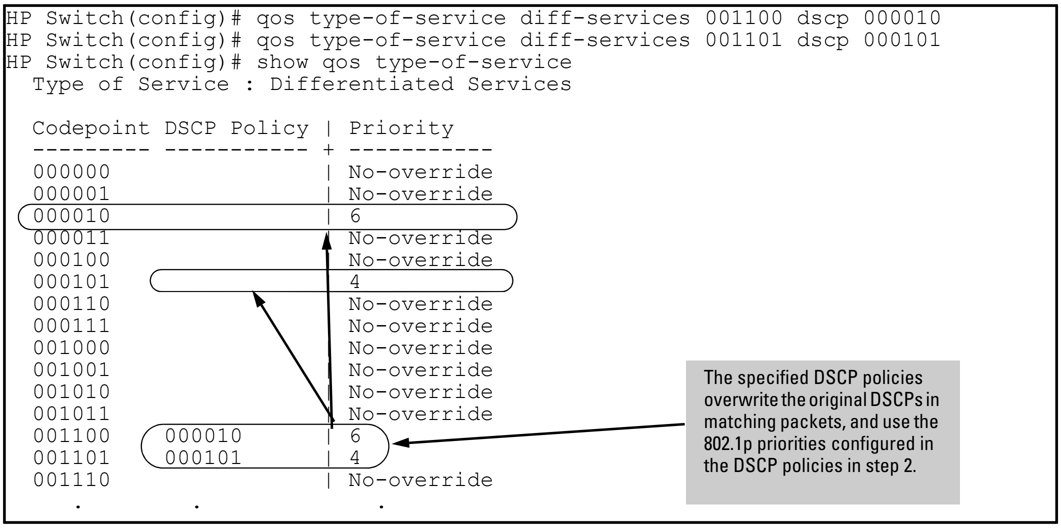

```
HP Switch(config)# qos type-of-service diff-services 001100 dscp 000010
HP Switch(config)# qos type-of-service diff-services 001101 dscp 000101
HP Switch(config)# show qos type-of-service
  Type of Service : Differentiated Services

  Codepoint DSCP Policy | Priority
  --------- ----------- + -----------
  000000                | No-override
  000001                | No-override
  000010                | 6
  000011                | No-override
  000100                | No-override
  000101                | 4
  000110                | No-override
  000111                | No-override
  001000                | No-override
  001001                | No-override
  001010                | No-override
  001011                | No-override
  001100    000010      | 6
  001101    000101      | 4
  001110                | No-override
     .         .             .
```

The specified DSCP policies overwrite the original DSCPs in matching packets, and use the 802.1p priorities configured in the DSCP policies in step 2.

Figure 36 Assigning DSCP policies to outbound packets based on the DSCP codepoint from upstream devices

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

`no qos protocol  ip | ipx | arp | appletalk | sna | netbeui priority` *`0 - 7`*

Disables use of the specified protocol as a QoS classifier and resets the protocol priority to `No-override`.

`show qos protocol`

Lists the QoS protocol classifiers with their priority settings.

Example

Example 146 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 37 Adding, displaying, removing, and changing QoS protocol classifiers

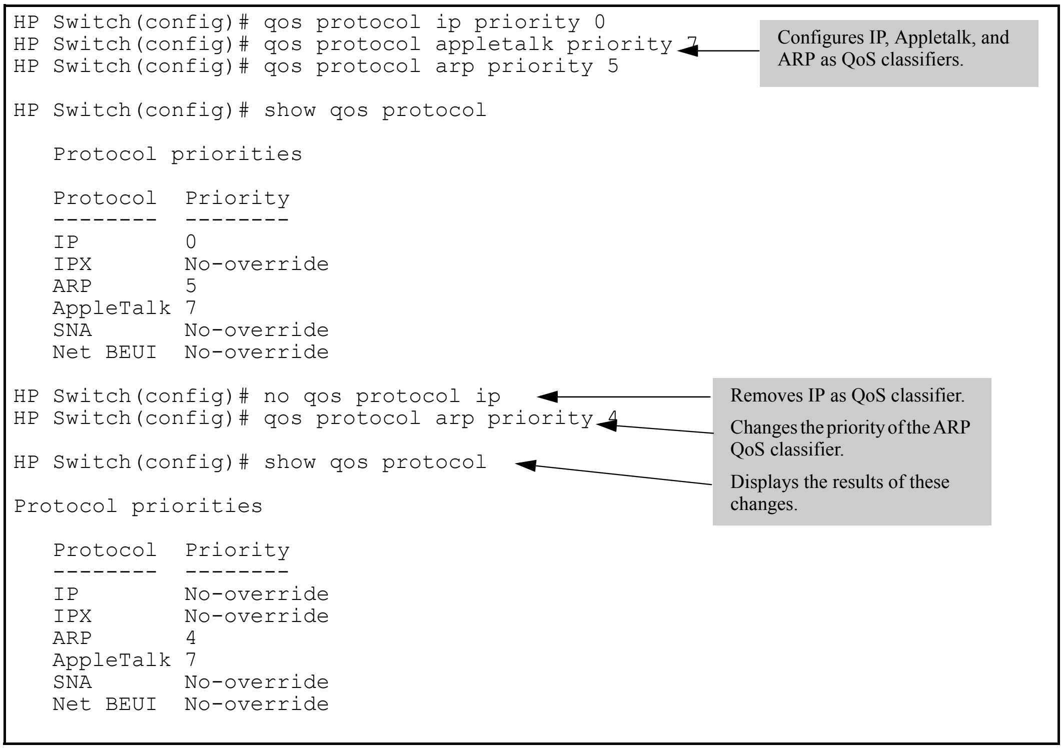

```
HP Switch(config)# qos protocol ip priority 0
HP Switch(config)# qos protocol appletalk priority 7
HP Switch(config)# qos protocol arp priority 5

HP 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

HP Switch(config)# no qos protocol ip
HP Switch(config)# qos protocol arp priority 4

HP 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.

Assigning a priority for a global VLAN-ID classifier

This global QoS packet-marking option assigns a priority to all outbound packets having the specified VLAN-ID (VID). You can configure this option by either specifying the VLAN-ID 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

`no vlan` *`vid`* `qos`

Removes the specified VLAN-ID as a QoS classifier and resets the priority for that VLAN to `No-override`.

```
show qos vlan-priority
```

Displays a listing of the QoS VLAN-ID classifiers currently in the running-config file, with their priority data.

Example

Example 147 Displaying the VLANs available for QoS prioritization

In this example, 802.1p priorities are assigned to packets received in VLANs 1, 20, 30, and 40.

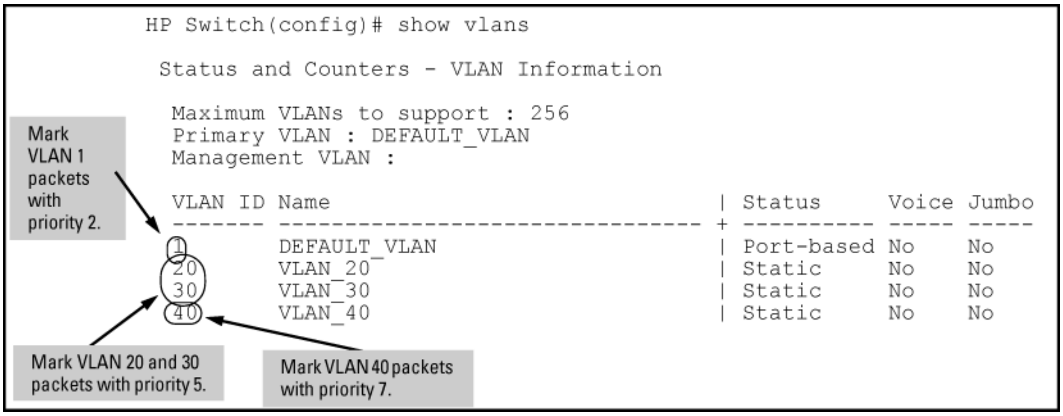

```
HP Switch(config)# show vlans

 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
  20      VLAN_20                          | Static     No    No
  30      VLAN_30                          | Static     No    No
  40      VLAN_40                          | Static     No    No
```

Mark VLAN 1 packets with priority 2.

Mark VLAN 20 and 30 packets with priority 5.

Mark VLAN 40 packets with priority 7.

Enter the following commands to mark VLAN packets that match the specified VLAN IDs with an 802.1p priority:

```
HP Switch(config)# vlan 1 qos priority 2
HP Switch(config)# vlan 20 qos priority 5
HP Switch(config)# vlan 30 qos priority 5
HP Switch(config)# vlan 40 qos priority 7
HP Switch(config)# show qos vlan

  VLAN priorities

  VLAN ID Apply rule   | DSCP   Priority
  ------- ------------+ ------ -------------
  1       Priority     |        2
  20      Priority     |        5
  30      Priority     |        5
  40      Priority     |        7
```

Example 148 Returning a QoS-prioritized VLAN to "No-override" status

If later it is necessary to remove VLAN 20 from QoS prioritization, enter the following command:

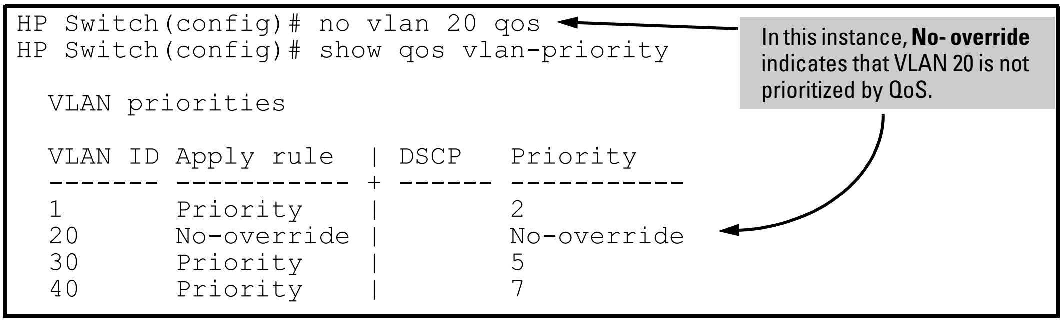

```
HP Switch(config)# no vlan 20 qos
HP Switch(config)# show qos vlan-priority

  VLAN priorities

  VLAN ID Apply rule  | DSCP   Priority
  ------- ----------- + ------ -----------
  1       Priority    |        2
  20      No-override |        No-override
  30      Priority    |        5
  40      Priority    |        7
```

In this instance, **No- override** indicates that VLAN 20 is not prioritized by QoS.

Assigning a DSCP policy for a global VLAN-ID classifier

This global QoS packet-marking option assigns a previously configured DSCP policy (codepoint and 802.1p priority) to outbound IP packets having the specified VLAN-ID (VID). The switch:

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 for matching packets.
3. Assigns the 802.1p priority associated with the new DSCP. (See "Using Differentiated Services Codepoint (DSCP) mapping" (page 249).)
4. Forwards the packet through the appropriate outbound port queue.

Creating a policy based on the 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 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.

   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 (`show qos dscp-map` command), you must first configure a priority for the codepoint before proceeding (`qos dscp-map priority` command). See "Using Differentiated Services Codepoint (DSCP) mapping" (page 249).

4. Configure the switch to assign the DSCP policy to packets with the specified VLAN-ID.

   ### Syntax:

   `qos dscp-map` *`codepoint`*`priority` *`0 - 7`*

   This command is optional if a priority has already been assigned to the *`codepoint`*. The command creates a DSCP policy by assigning an 802.1p priority to a specific DSCP.

   When the switch applies this priority 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 codepoint in the packet header is replaced by the codepoint specified in this command.

   Default: For most codepoints, No-override. See Table 14 (page 250).

Syntax:

`vlan` *`vid`* `qos dscp` *`codepoint`*

Assigns a DSCP policy to IP packets carrying the specified VLAN ID, and overwrites the DSCP in these packets with the assigned *`codepoint`* value.

- 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 name for a binary DSCP bit set.

  Enter `?` to display the list of valid codepoint entries.

  The DSCP 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 vlan` *`vid`* `qos`

Removes a global QoS classifier for the specified VLAN.

Syntax:

`show qos device-priority`

Displays a listing of all QoS VLAN-ID classifiers currently in the running-config file.

Example 149 Assigning DSCP policies to packets

This example assigns the following DSCP policies (codepoint and associated 802.1p priority) to packets with the specified VLAN IDs:

VLAN-ID	DSCP	Priority
40	000111	7
30	000101	5
20	000010	1
1	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-mapcommand)`. This is not a problem as long as the configured priorities are acceptable for all applications using the same DSCP (see "Notes on changing priority settings" (page 274).

   A DSCP codepoint must also have a priority configured before you can use it to mark matching packets.

   Figure 38 Displaying the current DSCP-priority mapping in the DSCP policy table

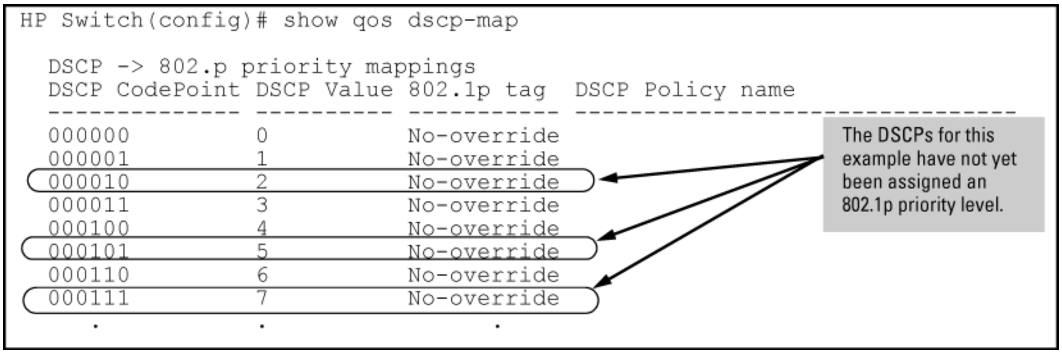

```
HP Switch(config)# show qos dscp-map

  DSCP -> 802.p priority mappings
  DSCP CodePoint DSCP Value 802.1p tag  DSCP Policy name
  -------------- ---------- ----------- -------------------------------
  000000         0          No-override
  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
        .             .               .
```

The DSCPs for this example have not yet been assigned an 802.1p priority level.

2. Configure the priorities for the DSCPs.

   Figure 39 Assign priorities to the selected DSCPs

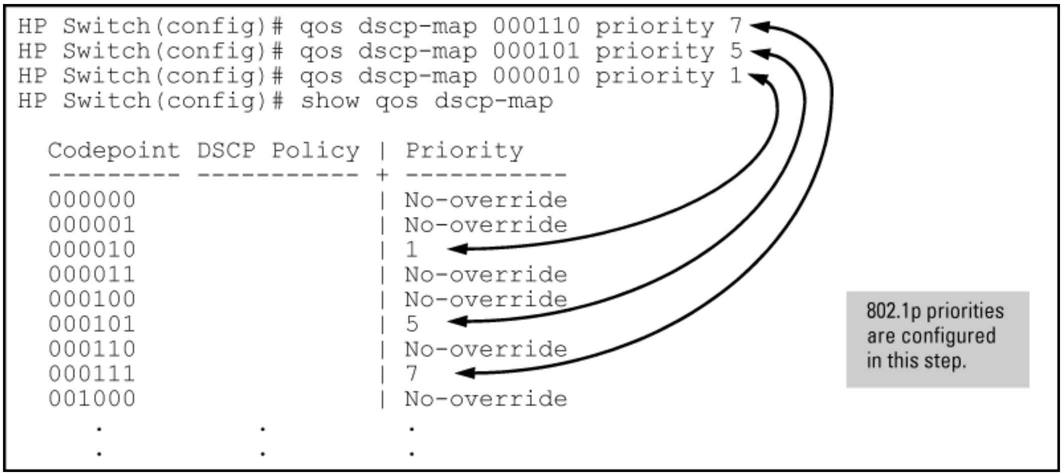

```
HP Switch(config)# qos dscp-map 000110 priority 7
HP Switch(config)# qos dscp-map 000101 priority 5
HP Switch(config)# qos dscp-map 000010 priority 1
HP Switch(config)# show qos dscp-map

  Codepoint DSCP Policy | Priority
  --------- ----------- + -----------
  000000                | No-override
  000001                | No-override
  000010                | 1
  000011                | No-override
  000100                | No-override
  000101                | 5
  000110                | No-override
  000111                | 7
  001000                | No-override
      .          .          .
      .          .          .
```

802.1p priorities are configured in this step.

Creating a policy based on the VLAN-ID classifier 231

3. Assign the DSCP policies to the selected VLAN IDs and display the result.

```
HP Switch(config)# vlan 1 qos dscp 000010
HP Switch(config)# vlan 20 qos dscp 000010
HP Switch(config)# vlan 30 qos dscp 000101
HP Switch(config)# vlan 40 qos dscp 000111

HP Switch(config)# show qos vlan-priority

  VLAN priorities

  VLAN ID Apply rule  | DSCP   Priority
  ------- ----------- + ------ -----------
  1       DSCP        | 000010 1
  20      DSCP        | 000010 1
  30      DSCP        | 000101 5
  40      DSCP        | 000111 7
```

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.

Assigning a priority for a global source-port classifier

This global QoS packet-marking option assigns a priority to all outbound packets having the specified source-port.

This option can be configured 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. For configuring multiple source-ports with the same priority, it is 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`*

Disables use of the specified source-ports for QoS classifiers and resets the priority for the specified source-ports to `No-override`.

Syntax:

`show qos port-priority`

Lists the QoS port-priority classifiers with their priority data.

Example

Example 150 Prioritizing inbound traffic on source-ports

This example shows how to prioritize inbound traffic on the following source-ports:

Source-Port	Priority
A1 - A3	2
A4	3
B1, B4	5
C1-C3	6

Enter the following commands to prioritize packets received from the specified source ports:

Figure 40 Configuring and displaying source-port QoS priorities

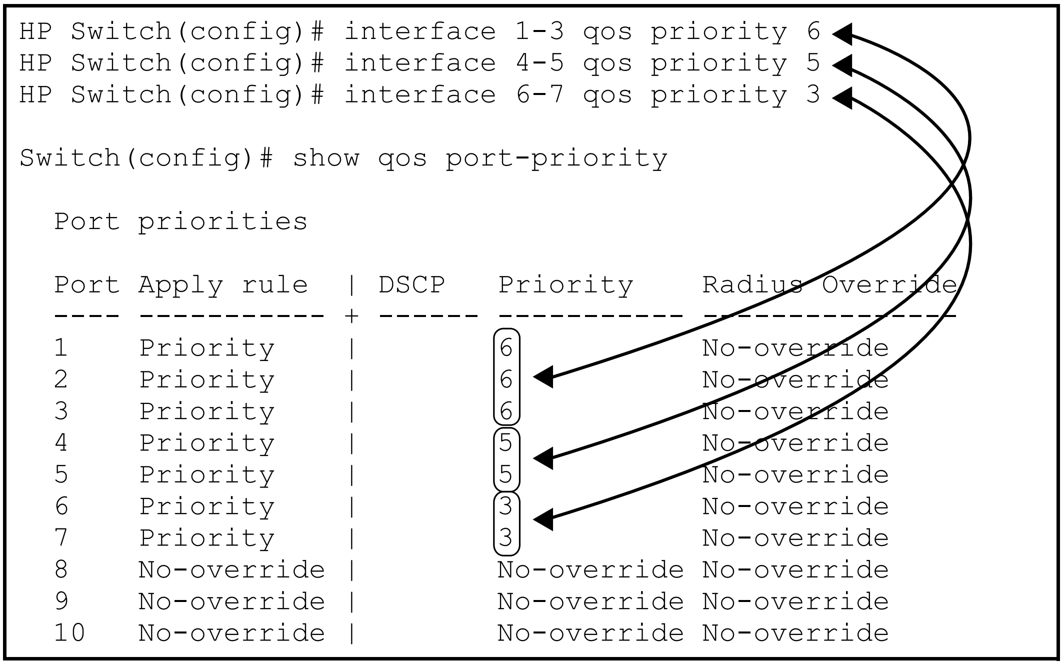

```
HP Switch(config)# interface 1-3 qos priority 6
HP Switch(config)# interface 4-5 qos priority 5
HP Switch(config)# interface 6-7 qos priority 3

Switch(config)# show qos port-priority

  Port priorities

  Port Apply rule  | DSCP   Priority    Radius Override
  ---- ----------- + ------ ----------- ---------------
  1    Priority    |        6           No-override
  2    Priority    |        6           No-override
  3    Priority    |        6           No-override
  4    Priority    |        5           No-override
  5    Priority    |        5           No-override
  6    Priority    |        3           No-override
  7    Priority    |        3           No-override
  8    No-override |        No-override No-override
  9    No-override |        No-override No-override
  10   No-override |        No-override No-override
```

If you later decided to remove source-port A1 from QoS prioritization, you would enter the following command:

Figure 41 Returning a QoS-prioritized VLAN to "No-override" status

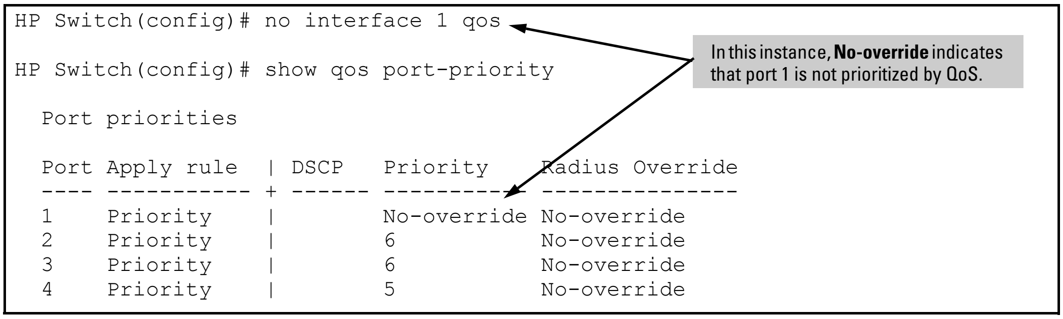

```
HP Switch(config)# no interface 1 qos

HP Switch(config)# show qos port-priority

  Port priorities

  Port Apply rule  | DSCP   Priority    Radius Override
  ---- ----------- + ------ ----------- ---------------
  1    Priority    |        No-override No-override
  2    Priority    |        6           No-override
  3    Priority    |        6           No-override
  4    Priority    |        5           No-override
```

In this instance, **No-override** indicates that port 1 is not prioritized by QoS.

Assigning a DSCP policy for a global source-port classifier

This global QoS packet-marking option assigns a previously configured DSCP policy (codepoint and 802.1p priority) to outbound IP packets received from the specified source-ports. The switch:

1. Selects an incoming IP packet on the basis of its source-port.
2. Overwrites the packet's DSCP with the DSCP configured for matching packets.
3. Assigns the 802.1p priority associated with the new DSCP. (See "Using Differentiated Services Codepoint (DSCP) mapping" (page 249).)
4. Forwards the packet through the appropriate outbound port queue.

For more on DSCP, see "About QoS" (page 254).

Creating a policy based on source-port classifiers

Only one DSCP per source-port may be used to mark matching packets.

Configuring a new DSCP for a source-port automatically overwrites any previous DSCP or 802.1p priority configuration for that source-port classifier.

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. 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 (`show qos dscp-map` command), you must first configure a priority for the codepoint before proceeding (`qos dscp-map priority` command). See "Using Differentiated Services Codepoint (DSCP) mapping" (page 249).

Syntax:

`qos dscp-map` *`codepoint`* `priority` *`0 - 7`*

This command is optional if a priority has already been assigned to the *`codepoint`*.

The command creates a DSCP policy by assigning an 802.1p priority to a specific DSCP. When the switch applies this priority 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 codepoint in the packet header is replaced by the codepoint specified in this command.

Default: For most codepoints, No-override. See Table 14 (page 250).

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 IP packets from the specified source-ports, and overwrites the DSCP in these packets with the assigned `codepoint` value.

- 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 name for a binary DSCP bit set.

  Enter `?` to display the list of valid codepoint entries.

Syntax:

`interface  port-list qos dscp codepoint`

The DSCP 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  [e]  port-list qos`

Removes a QoS classifier for the specified source-ports.

Syntax:

`show qos source-port`

Displays a listing of all source-port QoS classifiers currently in the running-config file.

Example

Example 151 Assigning DSCP policies (codepoint and associated 802.1p priority) to matching packets

In this example, the following DSCP policies (codepoint and associated 802.1p priority) are assigned to matching packets with the specified source-ports:

Source-Port	DSCP	Priority
A2	000111	7
B1-B3	000101	5
B4, C2	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). This is not a problem as long as the configured priorities are acceptable for all applications using the same DSCP.

   Also, note that a DSCP must have an 802.1p priority configured before you can use it to mark matching packets. 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.

Creating a policy based on source-port classifiers 235

Figure 42 Displaying the current DSCP-priority mapping in the DSCP policy table

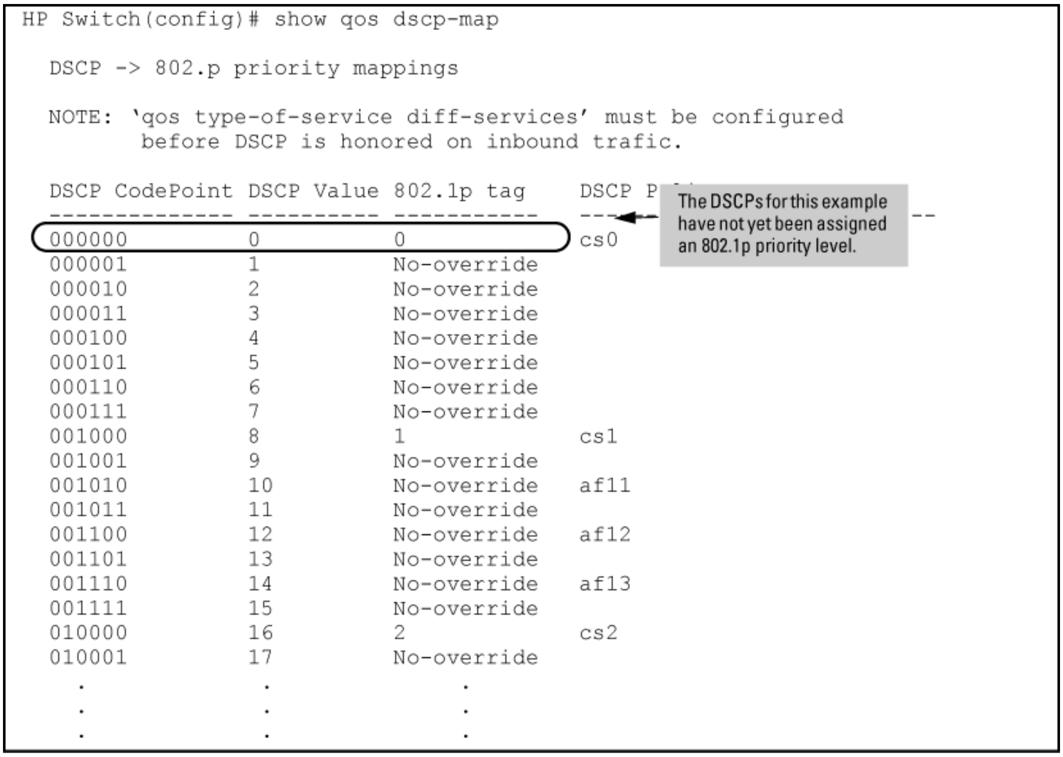

```
HP Switch(config)# show qos dscp-map

  DSCP -> 802.p priority mappings

  NOTE: 'qos type-of-service diff-services' must be configured
         before DSCP is honored on inbound trafic.

  DSCP CodePoint DSCP Value 802.1p tag     DSCP P
  -------------- ---------- -----------    ---                   --
  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
  001010         10         No-override    af11
  001011         11         No-override
  001100         12         No-override    af12
  001101         13         No-override
  001110         14         No-override    af13
  001111         15         No-override
  010000         16         2              cs2
  010001         17         No-override
     .             .             .
     .             .             .
     .             .             .
```

The DSCPs for this example have not yet been assigned an 802.1p priority level.

2. Configure the priorities for the DSCPs that you want to use to mark matching packets.

Figure 43 Assigning priorities to the specified DSCP codepoints

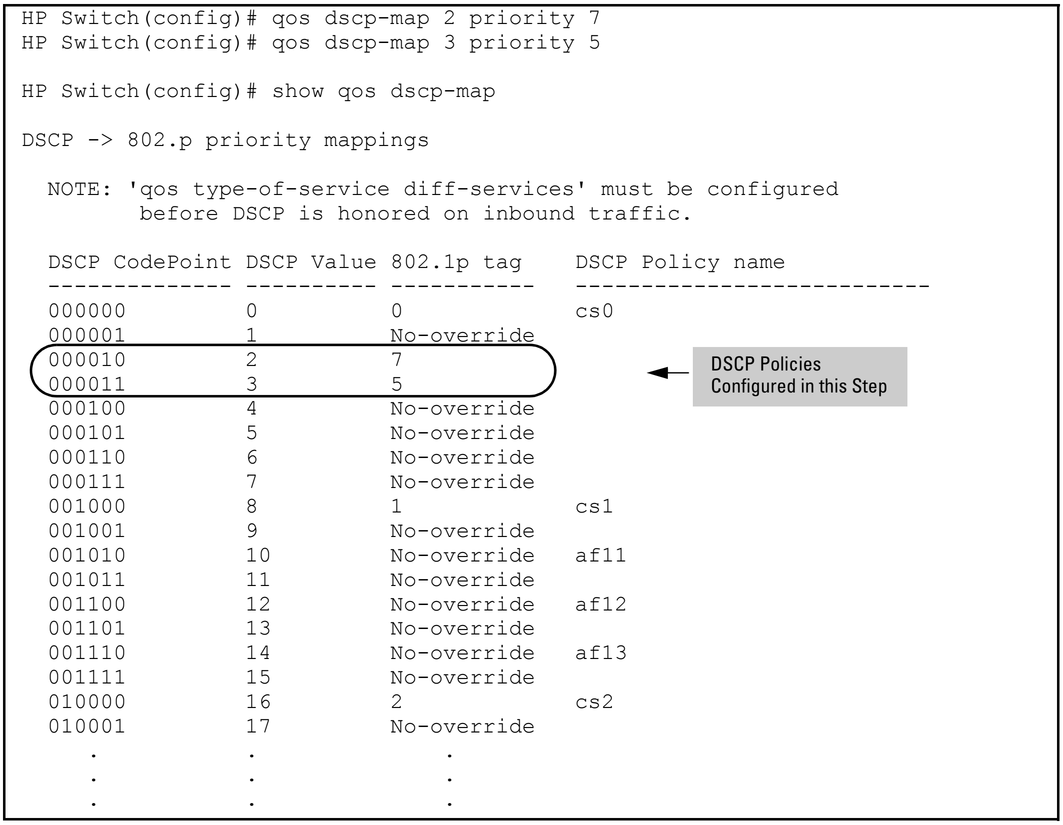

```
HP Switch(config)# qos dscp-map 2 priority 7
HP Switch(config)# qos dscp-map 3 priority 5

HP 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          7
  000011         3          5
  000100         4          No-override
  000101         5          No-override
  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
  001100         12         No-override    af12
  001101         13         No-override
  001110         14         No-override    af13
  001111         15         No-override
  010000         16         2              cs2
  010001         17         No-override
     .             .             .
     .             .             .
     .             .             .
```

DSCP Policies Configured in this Step

3. Assign the DSCP policies to the selected source-ports and display the result.

Figure 44 Displaying global source-port classifier with DSCP-priority marking

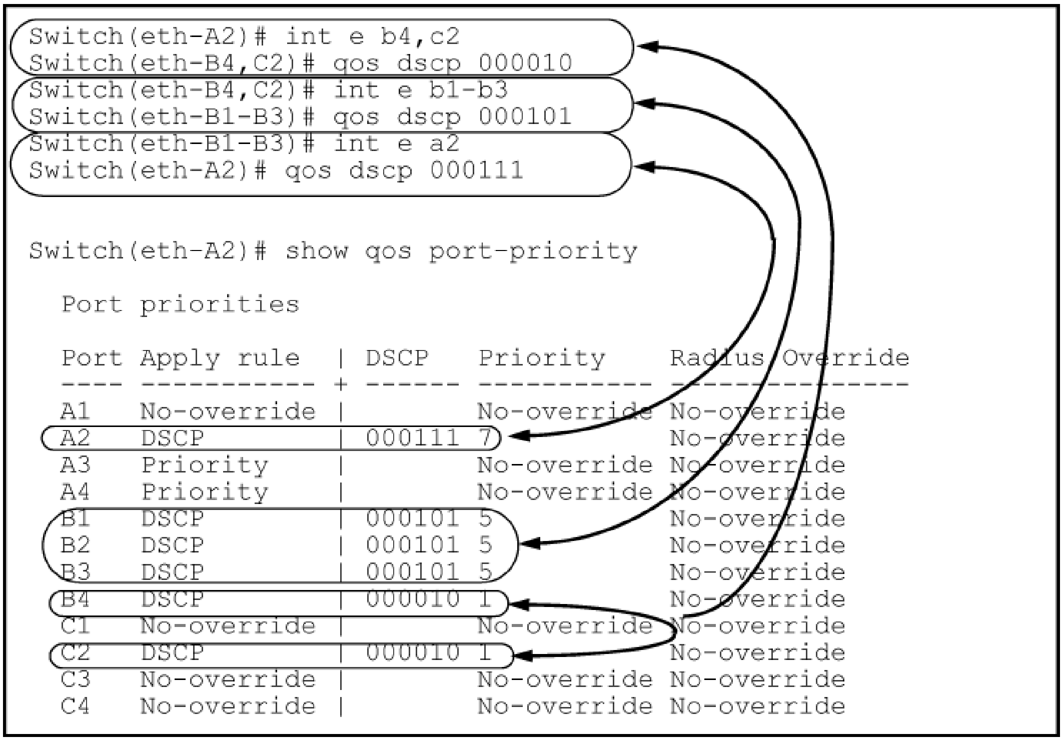

```
Switch(eth-A2)# int e b4,c2
Switch(eth-B4,C2)# qos dscp 000010
Switch(eth-B4,C2)# int e b1-b3
Switch(eth-B1-B3)# qos dscp 000101
Switch(eth-B1-B3)# int e a2
Switch(eth-A2)# qos dscp 000111


Switch(eth-A2)# show qos port-priority

  Port priorities

  Port Apply rule   | DSCP   Priority    Radius Override
  ---- ----------- + ------ ----------- ---------------
  A1   No-override  |        No-override No-override
  A2   DSCP         | 000111 7           No-override
  A3   Priority     |        No-override No-override
  A4   Priority     |        No-override No-override
  B1   DSCP         | 000101 5           No-override
  B2   DSCP         | 000101 5           No-override
  B3   DSCP         | 000101 5           No-override
  B4   DSCP         | 000010 1           No-override
  C1   No-override  |        No-override No-override
  C2   DSCP         | 000010 1           No-override
  C3   No-override  |        No-override No-override
  C4   No-override  |        No-override No-override
```

Configuring classifier-based QoS

To use the classifier-based model to configure a QoS policy and apply it to a selected class of traffic on a port or VLAN interface, follow these steps:

1. Evaluate the types of traffic in your network and identify the traffic types that you want to prioritize or rate limit.
2. Create an IPv4 or IPv6 traffic class using the `class` command to select the packets you want to manage.

   Context: Global configuration

Syntax:

`[no] class  ipv4 | ipv6  `*`classname`*

Defines the name of a traffic class and specifies whether a policy is to be applied to IPv4 or IPv6 packets, where *`classname`* is a text string (64 characters maximum). After you enter the `class` command, you enter the class configuration context to specify match criteria. A traffic class contains a series of `match` and `ignore` commands, which specify the criteria used to classify packets.

A traffic class consists of match criteria, which consist of `match` and `ignore` commands.

- The `match` commands define the values that header fields must contain for a packet to belong to the class and be managed by policy actions.
- The `ignore` commands define the values which, if contained in header fields, exclude a packet from the policy actions configured for the class.

NOTE: Enter match/ignore statements in the precise order in which you want their criteria to be used to check packets.

The following match criteria are supported in match/ignore statements for inbound IPv4/IPv6 traffic:

- IP source address (IPv4 and IPv6)
- IP destination address (IPv4 and IPv6)
- Layer 2 802.1Q VLAN ID
- Layer 3 IP protocol
- Layer 3 IP precedence bits
- Layer 3 DSCP codepoint
- Layer 4 TCP/UDP application port
- VLAN ID

3. Enter one or more `match` or `ignore` commands from the class configuration context to filter traffic and determine the packets on which policy actions will be performed.

   Context: Class configuration

   ### Syntax:

   [no] [*seq-number*] [match | ignore] *ip-protocol source-address destination-address* [dscp *codepoint*] [precedence *precedence-value*] [tos *tos-value*] [vlan *vlan-id*]

4. Create a QoS policy to perform QoS actions on selected packets by entering the `policy qos` command from the global configuration context.

   Context: Global configuration

   ### Syntax:

   [no] policy qos *policy-name*

   Defines the name of a QoS policy and enters the policy configuration context.

   A traffic policy consists of one or more classes, and one or more QoS actions configured for each class of traffic. The configured actions are executed on packets that match a `match` statement in a class. No policy action is performed on packets that match an `ignore` statement.

   NOTE: Be sure to enter each class and its associated QoS actions in the precise order in which you want packets to be checked and processed by QoS actions.

   To configure the QoS actions that you want to execute on packets that match the criteria in a specified class, enter one or more `class action` commands from the policy configuration context:

   Context: Class configuration

   ### Syntax:

   [no] [*seq-number*] class [ipv4 | ipv6] *classname* action *qos-action* [action *qos action ...*]

   Defines the QoS actions to be applied on a pre-configured IPv4 or IPv6 traffic class when a packet matches the `match` criteria in the traffic class. You can enter multiple action statements for the same traffic class.

   [no] [*seq-number*] class [ipv4 | ipv6] *classname*

   seq-number

   (Optional) Sequentially orders the QoS actions that you enter in a policy configuration. Actions are executed on matching packets in numerical order.

Default: QoS action statements are numbered in increments of 10, starting at 10.

`class [ipv4 | ipv6] classname`

Defines the preconfigured traffic class on which the QoS actions in the policy are executed, and specifies whether the QoS policy is applied to IPv4 or IPv6 traffic in the class. The `classname` is a text string (64 characters maximum).

NOTE: Multiple `class action` statements can be configured for different traffic classes in the same policy. The execution of QoS actions is performed in the order in which the actions are numerically listed in the policy.

`action qos-action [action qos-action ...]`

Configures the QoS action specified by the *qos-action* replaceable. The action is executed on any packet that matches the `match` criteria in the class. The action is not executed on packets that match `ignore` criteria.

The complete `no` form of the `class action` command or the `no seq-number` command removes a QoS action from the policy configuration.

The following QoS commands are supported by the *qos-action* replaceable:

- `rate-limit kbps`
- `priority priority-value`
- `ip-precedence precedence-value`
- `dscp dscp-value`

For information on the complete syntax of each QoS command, see "Configuring QoS actions in a policy" (page 240).

To manage packets that do not match the `match` or `ignore` criteria in any class in the policy, and therefore have no QoS actions performed on them, enter an optional default class. The default class is placed at the end of a policy configuration and specifies the QoS actions to perform on packets that are neither matched nor ignored.

5. (Optional) To configure a default class in a policy, enter the `default-class` command at the end of a policy configuration and specify one or more QoS actions to be executed on packets that are not matched and not ignored.

   Context: Class configuration

   ## Syntax:

   `[no] default-class action qos-action [action qos-action ...]`

   Configures a default class that allows one or more QoS actions to be executed on packets that are not matched or ignored by any of the class configurations in a QoS policy. The default-class supports the same QoS commands as the `class ipv4 |ipv6 action` command: `rate-limit`, `priority`, `ip-precedence`, and `dscp`.

6. Apply the QoS policy to inbound traffic on a port (`interface service-policy in` command) or VLAN (`vlan service-policy in` command) interface.

   The following restrictions apply to a QoS service policy:

   - Only one QoS policy is supported on a port or VLAN interface.
   - If you apply a QoS policy to a port or VLAN interface on which a QoS policy is already configured, the new policy replaces the existing one.
   - A QoS policy is supported only on inbound traffic.

   Because only one QoS policy is supported on a port or VLAN interface, ensure that the policy you want to apply contains all the required classes and actions for your configuration.

   To apply a QoS policy on a port or VLAN interface, enter one of the following commands from the global configuration context.

   Context: Global configuration

   ### Syntax:

   `interface` *`port-list`* `service-policy` *`policy-name`*

   Configures specified ports with a QoS policy that is applied to inbound traffic on each interface.

   Separate individual port numbers in a series with a comma; for example, `a1, b4, d3`.

   Enter a range of ports by using a dash; for example, `a1-a5`.

   The QoS policy name you enter must be the same as the policy name you configured with the `policy qos` command in Step 2.

   ### Syntax:

   `vlan` *`vlan-id`* `service-policy` *`policy-name`* `in`

   Configures a QoS policy on the specified VLAN that is applied to inbound traffic on the VLAN interface.

   Valid VLAN ID numbers range from 1 to 4094.

   The QoS policy name you enter must be the same as the policy name you configured with the `policy` command in Step 2.

7. Determine the additional QoS configurations that you need to apply to each QoS-capable device in your network and configure the appropriate policy.

   Optional: For802.1p (CoS) priority settings to be included in outbound packets, configure tagged VLANs on the appropriate downstream links.

Configuring QoS actions in a policy

In QoS policy-configuration mode, you define the actions to be applied to a pre-configured IPv4 or IPv6 traffic class when a packet matches the `match` criteria in the class. Note: Actions are not executed on packets that match `ignore` criteria. You can enter multiple action statements in a traffic class, including the default class.

The following commands are supported in a QoS policy configuration:

`rate-limit`

Configures the rate limit for matching packets.

`ip-precedence`

Configures (marks) the IP precedence bits in the ToS byte of IPv4 packet headers and Traffic Class byte of IPv6 headers.

`dscp`

Configures the DSCP bits in the IPv4 ToS byte and IPv6 Traffic Class byte of packet headers.

`priority`

Configures the 802.1p class of service (CoS) priority in Layer 2 frame headers.

For information on the difference between the DSCP bits and precedence bits in the ToS byte of an IPv4 header and the Traffic Class byte of an IPv6 header.

Context: Global configuration

Syntax:

[no] [*seq-number*]class [ipv4 | ipv6] *classname* action *qos-action* [action *qosaction ...*]

In a QoS policy configuration, the *qos-action* parameter can be any of the following commands:

`rate-limit` *kbps*

Configures the maximum transmission rate for matching packets in a specified traffic class. All packets that exceed the configured limit are dropped.

The rate limit is specified in kilobits per second, where *kbps* is a value from 0 to 10000000.

NOTE: **Rate limiting usage**:

- Rate limit values below 13 kbps may result in unpredictable rate limiting behavior.
- Configuring a rate limit of 0 (zero) kilobits on a port blocks all traffic on the port. If blocking all traffic is the desired behavior, HP recommends that you configure `deny` ACL instead configuring a rate limit of 0.
- A rate limit that you apply with a classifier-based policy overrides any globally-configured per-port rate limit on the selected packets.

  For more information on per-port rate limiting, see the Port Traffic Controls chapter in the *Management and Configuration Guide*.

Rate limiting restrictions:

- A rate limit is calculated on a per-module or per port-bank basis. If trunked ports or VLANs with a configured rate limit span multiple modules or port-banks, the configured rate limit is not guaranteed.
- A QoS policy that uses the `class action rate-limit` command is not supported on a port interface on which ICMP rate limiting has already been globally configured. To apply the QoS policy, you must first disable the ICMP rate limiting configuration.

  In cases where you want to maintain an ICMP rate limiting configuration, configure a class in which you specify the necessary match statements for ICMP traffic, and a QoS policy in which you configure the rate limit action for the class.

  For information on globally-configured ICMP, see the Configuring ICMP section in the Configuring IP Parameters for Routing Switches chapter in the *Multicast and Routing Guide*.

`priority` *priority-value*

Configures the 802.1p class of service (CoS) bits in Layer 2 frames of matching packets in a specified traffic class. Valid CoS values range from 0 to 7.

The 802.1p CoS value controls the outbound port-queue priority for traffic leaving the switch. In an 802.1Q VLAN network, downstream devices may honor or change the 802.1p priority in incoming packets. For more information, see "Layer 2802.1p prioritization" (page 257).

Table 16 (page 258) shows how the Layer 2 802.1p priority value determines to which outbound port queue a packet is sent both on the switch and on a downstream device.

The 802.1p CoS numeric value (from 0 to 7) corresponds to the hexadecimal equivalent of the three binary 0 and 1 bit settings in the Layer 2 header. For example if the CoS bit values are `1 1 1`, the numeric value is `7` (1+2+4). Similarly, if the CoS bits are `0 1 1`, the numeric value is `3` (1+2+0).

NOTE: If you want the 802.1p CoS priority settings included in outbound packets to be honored on downstream devices, configure tagged VLANs on the appropriate inbound and outbound ports.

`ip-precedence` *`precedence-value`*

Configures the IP precedence value in the IPv4 ToS byte or IPv6 Traffic Class byte of matching packets in a specified traffic class. Valid IP precedence values are either a numeric value from 0 (low priority) to 7 (high priority) or its corresponding name:

0

routine

1

priority

2

immediate

3

flash

4

flash-override

5

critical

6

internet (for internetwork control)

7

network (for network control)

Table 16 (page 258) shows how the Layer 2 802.1p priority value determines to which outbound port queue a packet is sent.

Table 21 (page 267) shows the 802.1p priority value (0 to 7) associated, by default, with each IP Precedence three-bit setting and automatically assigned by the switch to the Layer 2 header of matching packets.

`dscp` *`dscp-value`*

Configures the DSCP codepoint in the IPv4 ToS byte or IPv6 Traffic Class byte of matching packets in a specified traffic class.

Valid values for the DSCP codepoint are any of the following:

- A binary eight-bit set (such as `100110`)
- A decimal value from `0` (low priority) to `63` (high priority) that corresponds to a binary DSCP bit set
- The ASCII standard 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
`af41` (100010)	`cs7` (111000) = *precedence* 7
`default` (000000)	

Prerequisite: The DSCP value you enter must already be configured with an 802.1p priority in the DSCP Policy table ("The default DSCP policy table" (page 250)) before you can use it to mark matching packets.

NOTE: DSCP-802.1p Mapping: The 802.1p priority currently associated with each DSCP codepoint is stored in the DSCP Policy table (displayed with the `show qos dscp-map` command and shown in "The default DSCP policy table" (page 250)). Note that certain DSCP codepoints have 802.1p priorities assigned by default. The 802.1p priority mapped to a DSCP codepoint is automatically applied in matching packets whose codepoint is reset with the `class action dscp` command in a QoS policy.

Reconfiguring the 802.1p priority value currently assigned to a DSCP codepoint

To reconfigure the 802.1p priority value currently assigned to a DSCP codepoint, enter one of the following commands:

- Global configuration context:
  `qos dscp-map` *`codepoint`* `priority` *`0 - 7`*
- Policy configuration context:
  `class [ ipv4 | ipv6 ]` *`classname`* `action dscp` *`codepoint`* `priority 0 - 7`

If you do not enter a `priority` value with the `class action dscp` command in a QoS policy, one of the following occurs:

- The switch refers to the DSCP Policy table to assign the 802.1p value that is currently configured for the specified DSCP codepoint to remark matching packets.
- If the specified DSCP codepoint is not associated with an 802.1p priority in the DSCP Policy table, an error message is displayed and the `class action dscp` *`codepoint`* command is not executed. You are prompted to re-enter the command with an 802.1p priority: `class action dscp` *`codepoint`* `priority 0 - 7.`

To ensure that the desired 802.1p priority is assigned to matching packets, you may need to first re-map the priority to the new codepoint before you configure the policy, by using the `qos dscp-map` *`codepoint`* `priority` *`0 - 7`* command.

NOTE: After you reconfigure the 802.1p priority for a DSCP codepoint, the switch immediately applies the new 802.1p priority value to packets transmitted with the associated codepoint as a result of:

- Globally-configured QoS commands
- `class action dscp` commands in other QoS policies

Example 152 Applying classifier-based QoS policy to inbound traffic on VLAN

In the following example, a classifier-based QoS policy (`dscp-remap`) that assigns a new DSCP codepoint (af43) and associated 802.1p priority (5) to matching packets with a specified DSCP codepoint (af11) is applied to the inbound traffic on a VLAN.

```
HP Switch(config)# qos dscp-map af43 priority 5
HP Switch(config)# class ipv4 dscp5
HP Switch(config-class)# match ip any any dscp af11
HP Switch(config-class)# exit
HP Switch(config)# policy qos dscp-remap
HP Switch(config-policy)# class ipv4 dscp5 action dscp af43
HP Switch(config-policy)# exit
HP Switch(config)# vlan 3 service-policy dscp-remap in
```

NOTE: In this example, the desired 802.1p priority is mapped to the specified DSCP codepoint by using the `qos dscp-map` *`codepoint`* `priority` *`0 - 7`* command before the QoS policy is configured.

Viewing a classifier-based QoSconfiguration

Use the following `show` commands to display information about a classifier-based QoS configuration and statistics or resource usage on QoS policies.

Syntax:

```
show class ipv4 classname
show class ipv6 classname
show class config
```

`ipv4` *`classname`*

Lists the statements that make up the IPv4 class identified by classname.

`ipv6` *`classname`*

Lists the statements that make up the IPv6 class identified by classname.

`config`

Displays all classes, both IPv4 and IPv6, and lists the statements that make up each class.

Additional variants of the `show class` command provide information on classes that are members of policies that have been applied to ports or VLANs.

Example

Example 153 Displaying `show class` output for a QoS policy

```
HP Switch(config)# show class ipv4 gnutella
Statements for Class ipv4 "gnutella"
  10 match tcp 0.0.0.0 255.255.255.255 range 6346 6347 0.0.0.0 255.255.255.255
  20 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 range 6346 6347
  30 match udp 0.0.0.0 255.255.255.255 range 6346 6347 0.0.0.0 255.255.255.255
  40 match udp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 range 6346 6347

HP Switch(config)# show class ipv4 kazaa
Statements for Class ipv4 "kazaa"
  10 match tcp 0.0.0.0 255.255.255.255 eq 1214 0.0.0.0 255.255.255.255
  20 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 1214
  30 match udp 0.0.0.0 255.255.255.255 eq 1214 0.0.0.0 255.255.255.255
  40 match udp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 1214

HP Switch(config)# show class ipv4 http
Statements for Class ipv4 "http"
  10 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 80
  20 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 443
  50 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 8080
```

Syntax:

```
show policy policy-name
show policy config
```

`policy-name`

Lists the statements that make up the specified policy.

`config`

Displays the names of all policies defined for the switch and lists the statements that make up each policy.

Additional variants of the `show policy` command provide information on policies that have been applied to ports or VLANs.

Example

Example 154 Displaying `show policy` output for a QoS policy

```
HP Switch(config)# show policy suspect-traffic
Statements for Policy "suspect-traffic"
  10 class ipv4 "http" action rate-limit kbps 2000 action priority 3
  20 class ipv4 "kazaa" action rate-limit kbps 1000 action priority 2
  30 class ipv4 "gnutella" action rate-limit kbps 1000 action priority 2
```

Syntax:

[show | clear] statistics policy *policy-name* port *port-num*
[show | clear] statistics policy *policy-name* vlan *vid* in

`show`

Displays the statistics for a specified policy applied to a specified port or VLAN.

`clear`

Clears statistics for the specified policy and port or VLAN.

`policy-name`

Specifies the name of the policy.

port-num

Specifies the number of the port on which the policy is applied (single port only, not a range).

vid

Specifies the number or name of the vlan on which the policy is applied. VLAN ID numbers range fro 1 to 4094.

`in`

Specifies that statistics are shown for inbound traffic only.

Example 155 Displaying `show statistics policy` output for a QoS policy

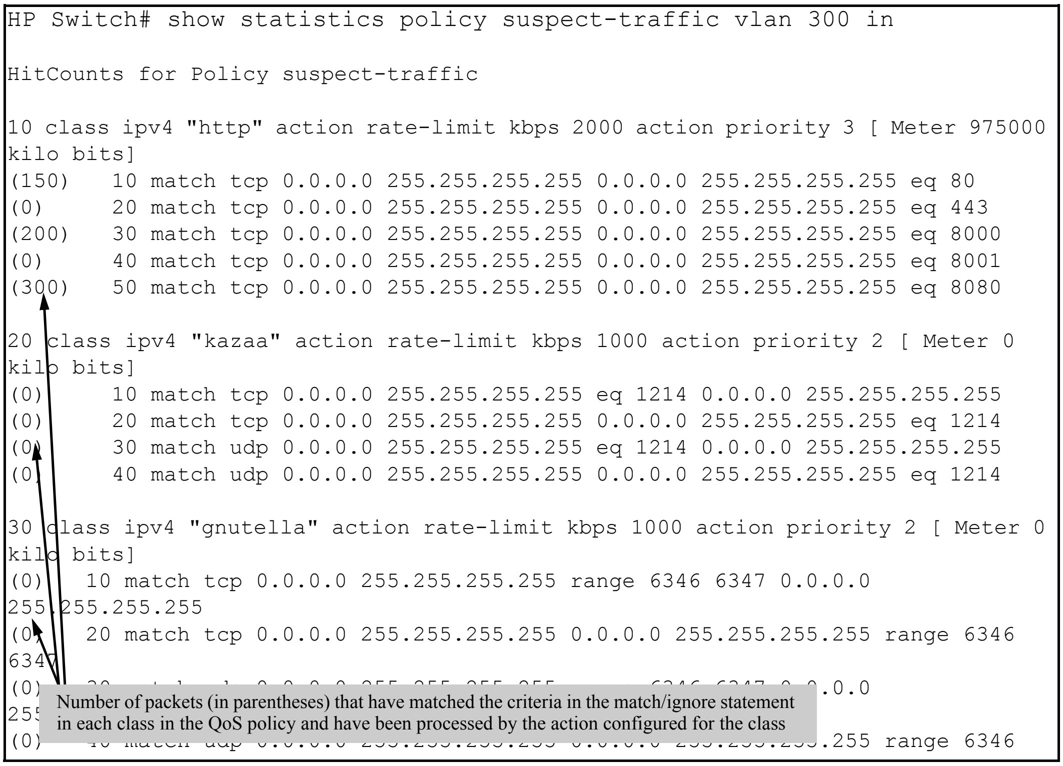

```
HP Switch# show statistics policy suspect-traffic vlan 300 in

HitCounts for Policy suspect-traffic

10 class ipv4 "http" action rate-limit kbps 2000 action priority 3 [ Meter 975000
kilo bits]
(150)   10 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 80
(0)     20 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 443
(200)   30 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 8000
(0)     40 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 8001
(300)   50 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 8080

20 class ipv4 "kazaa" action rate-limit kbps 1000 action priority 2 [ Meter 0
kilo bits]
(0)     10 match tcp 0.0.0.0 255.255.255.255 eq 1214 0.0.0.0 255.255.255.255
(0)     20 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 1214
(0)     30 match udp 0.0.0.0 255.255.255.255 eq 1214 0.0.0.0 255.255.255.255
(0)     40 match udp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 eq 1214

30 class ipv4 "gnutella" action rate-limit kbps 1000 action priority 2 [ Meter 0
kilo bits]
(0)   10 match tcp 0.0.0.0 255.255.255.255 range 6346 6347 0.0.0.0
255.255.255.255
(0)   20 match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 range 6346
6347
(0)   30 match udp 0.0.0.0 255.255.255.255 range 6346 6347 0.0.0.0
255
(0)   40 match udp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 range 6346
```

Number of packets (in parentheses) that have matched the criteria in the match/ignore statement in each class in the QoS policy and have been processed by the action configured for the class

Syntax:

`show policy resources`

Displays the number of hardware resources (rules, meters, and application port ranges) used by classifier-based QoS policies that are currently applied to interfaces on the switch, as well as mirroring policies and other software features.

NOTE: The information displayed is the same as the output of `show qos resources` (see Example 142 (page 215)) and `show access-list resources` commands. For a detailed explanation of the information displayed with the `show [qos | access-list | policy] resources` command, see the *"Monitoring Resources"* appendix of the Management and Configuration Guide.

Example 156 Displaying `show policy resources` output for all currently configured QoS policies

```
HP Switch(config)# show policy resources

 Resource usage in Policy Enforcement Engine

        |    Rules     |  Rules Used
  Slots |  Available  | ACL | QoS | IDM |  VT | Mirror | PBR | Other |
  ------+-------------+-----+-----+-----+-----+--------+-----+-------|
    A   |         3014 |  15 |  11 |    0 |    1 |       0 |   0 |    3  |

        |    Meters    |  Meters Used
  Slots |  Available  | ACL | QoS | IDM |  VT | Mirror | PBR | Other |
  ------+-------------+-----+-----+-----+-----+--------+-----+-------|
    A   |          250 |     |   5 |    0 |     |        |     |     0 |

        | Application |
        | Port Ranges |  Application Port Ranges Used
  Slots |  Available  | ACL | QoS | IDM |  VT | Mirror | PBR | Other |
  ------+-------------+-----+-----+-----+-----+--------+-----+-------|
    A   |           14 |   2 |   0 |    0 |     |       0 |   0 |     0 |

 0 of 8 Policy Engine management resources used.

 Key:
 ACL = Access Control Lists
 QoS = Device & Application Port Priority, QoS Policies, ICMP rate limits
 IDM = Identity Driven Management
 VT  = Virus Throttling blocks
 Mirror = Mirror Policies, Remote Intelligent Mirror endpoints
 PBR = Policy Based Routing Policies
 Other = Management VLAN, DHCP Snooping, ARP Protection, Jumbo IP-MTU,
         Transparent Mode.

 Resource usage includes resources actually in use, or reserved for future
 use by the listed feature.  Internal dedicated-purpose resources, such as
 port bandwidth limits or VLAN QoS priority, are not included.
```

Configuring a QoS policy for Voice over IP and Data traffic (Example)

In this example, an administrator would like to configure the following Layer 2 802.1p CoS and Layer 3 DSCP values to prioritize how VoIP traffic from different phones is handled compared to data traffic:

Softphone traffic

  DSCP 46; 802.1p CoS priority 6

Avaya phone traffic

  DSCP 34; 802.1p CoS priority 3

Miscellaneous phone traffic

  DSCP 26; 802.1p CoS priority 3

Data traffic

  DSCP 000000; 802.1p CoS priority 0

The following QoS configuration creates and assigns a QoS policy to VLAN 1 that prioritizes VoIP and data traffic in this way:

Figure 45 A QoS policy for voice over IP and data traffic

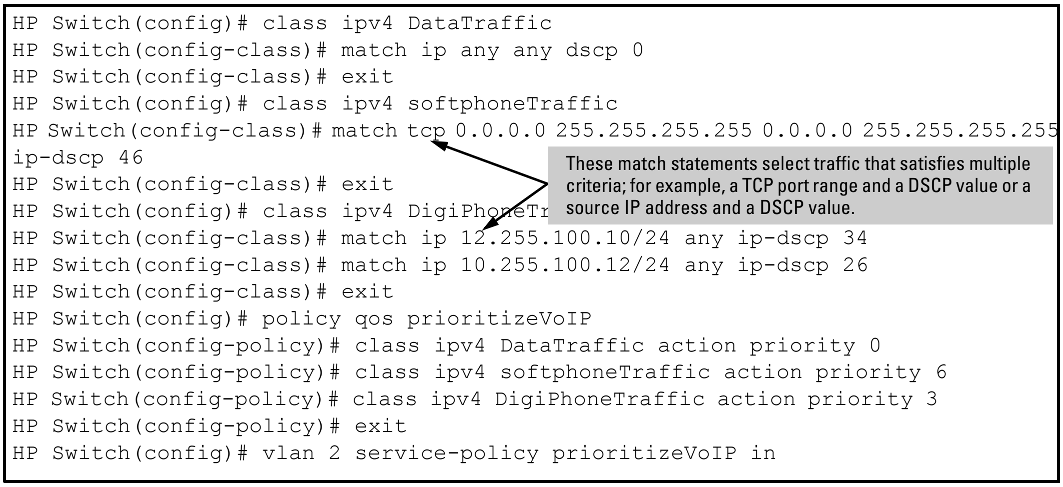

```
HP Switch(config)# class ipv4 DataTraffic
HP Switch(config-class)# match ip any any dscp 0
HP Switch(config-class)# exit
HP Switch(config)# class ipv4 softphoneTraffic
HP Switch(config-class)# match tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
ip-dscp 46
HP Switch(config-class)# exit
HP Switch(config)# class ipv4 DigiPhoneTr
HP Switch(config-class)# match ip 12.255.100.10/24 any ip-dscp 34
HP Switch(config-class)# match ip 10.255.100.12/24 any ip-dscp 26
HP Switch(config-class)# exit
HP Switch(config)# policy qos prioritizeVoIP
HP Switch(config-policy)# class ipv4 DataTraffic action priority 0
HP Switch(config-policy)# class ipv4 softphoneTraffic action priority 6
HP Switch(config-policy)# class ipv4 DigiPhoneTraffic action priority 3
HP Switch(config-policy)# exit
HP Switch(config)# vlan 2 service-policy prioritizeVoIP in
```

These match statements select traffic that satisfies multiple criteria; for example, a TCP port range and a DSCP value or a source IP address and a DSCP value.

Configuring a QoS policy for layer 4 TCP/UDP traffic (Example)

The following example shows how to configure a rate limiting policy for TCP/UDP application streams and apply the policy on all inbound switch ports.

```
HP Switch(config)# class ipv4 http
HP Switch(config-class)# match tcp any any eq 80
HP Switch(config-class)# match tcp any any eq 443
HP Switch(config-class)# match tcp any any eq 8080
HP Switch(config-class)# exit
HP Switch(config)# class ipv4 kazaa
HP Switch(config-class)# match tcp any eq 1214 any
HP Switch(config-class)# match tcp any any eq 1214
HP Switch(config-class)# exit
HP Switch(config)# class ipv4 gnutella
HP Switch(config-class)# match tcp any range 6346 6347 any
HP Switch(config-class)# match tcp any any range 6346 6347
HP Switch(config-class)# match udp any range 6346 6347 any
HP Switch(config-class)# match udp any any range 6346 6347
HP Switch(config-class)# exit
HP Switch(config)# policy qos PrioritizeSuspectTraffic
HP Switch(config-policy)# class ipv4 http action rate-limit kbps 7000
HP Switch(config-policy)# class ipv4 kazaa action rate-limit kbps 3500
HP Switch(config-policy)# class ipv4 gnutella action rate-limit kbps 3500
HP Switch(config-policy)# exit
HP Switch(config)# interface all service-policy PrioritizeSuspectTraffic in
```

Configuring a QoS policy for subnet traffic (Example)

The next example shows how to configure a QoS policy that prioritizes inbound traffic sent to and received from a specified subnet (15.29.16.0/10) and TCP port range on VLAN 5.

Figure 46 A QoS policy for IPv4 and IPv6 subnet traffic on a VLAN interface

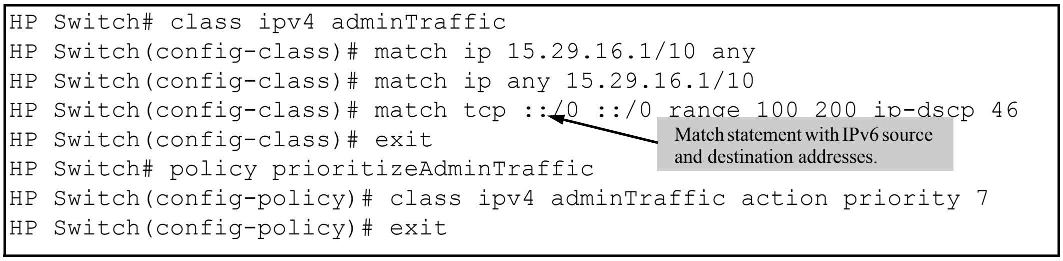

```
HP Switch# class ipv4 adminTraffic
HP Switch(config-class)# match ip 15.29.16.1/10 any
HP Switch(config-class)# match ip any 15.29.16.1/10
HP Switch(config-class)# match tcp ::/0 ::/0 range 100 200 ip-dscp 46
HP Switch(config-class)# exit
HP Switch# policy prioritizeAdminTraffic
HP Switch(config-policy)# class ipv4 adminTraffic action priority 7
HP Switch(config-policy)# exit
```

Match statement with IPv6 source and destination addresses.

Using Differentiated Services Codepoint (DSCP) mapping

The DSCP Policy Table associates an 802.1p priority with a DSCP codepoint in an IPv4/IPv6 packet. Using DSCP codepoints in your network allows 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 Table 14 (page 250). However, some codepoints, such as Assured Forwarding and Expedited Forwarding, have a default 802.1p priority setting.

Use the following commands to display the DSCP Policy table, configure the codepoint-priority assignments, and assign optional names to the codepoints.

Syntax:

`show qos dscp-map`

Displays the DSCP Policy table.

`qos dscp-map` *`codepoint`* `priority`*`0`*` - 7`[`name` *`ascii-string`*]

Configures an 802.1p priority for the specified codepoint and an optional (DSCP policy)name.

`no qos dscp-map` *`codepoint`*

Removes the currently configured 802.1p priority that is associated with the specified *`codepoint`* and displays the `No-override` setting. The codepoint policy name, if configured, is also removed.

`no qos dscp-map` *`codepoint`*`name`

Deletes only the policy name, if configured, for the specified `codepoint`.

Table 14 The default DSCP policy table

DSCP Policy	802.1p Priority	DSCP Policy	802.1p Priority	DSCP Policy	802.1p Priority
000000	No-override	010110	3[1]	101011	No-override
000001	No-override	010111	No-override	101100	No-override
000010	No-override	011000	No-override	101101	No-override
000011	No-override	011001	No-override	101110	7[2]
000100	No-override	011010	4[1]	101111	No-override
000101	No-override	011011	No-override	110000	No-override
000110	No-override	011100	4[1]	110001	No-override
000111	No-override	011101	No-override	110010	No-override
001000	No-override	011110	5[1]	110011	No-override
001001	No-override	011111	No-override	110100	No-override
001010	1[1]	100000	No-override	110101	No-override
001011	No-override	100001	No-override	110110	No-override
001100	1[1]	100010	6[1]	110111	No-override
001101	No-override	100011	No-override	111000	No-override
001110	2[1]	100100	6[1]	111001	No-override
001111	No-override	100101	No-override	111010	No-override
010000	No-override	100110	7[1]	111011	No-override
010001	No-override	100111	No-override	111100	No-override
010010	0[1]	101000	No-override	111101	No-override
010011	No-override	101001	No-override	111110	No-override
010100	0[1]	101010	No-override	111111	No-override
010101	No-override				

[1] Assured Forwarding codepoints; configured by default on the switches covered in this guide.

[2] Expedited Forwarding codepoint configured by default.

Displaying non-default codepoint settings (Example)

Default priority settings for selected codepoints

In a few cases, such as 001010 and 001100, a default DSCP policy (implied by the DSCP standards for Assured-Forwarding and Expedited-Forwarding) is used. You can change the priorities for the default policies by using the `qos dscp-map` *`codepoint`*`priority` *`0 - 7`* command.

The currently configured DSCP policies (codepoint and associated 802.1p priority) are not enabled until you configure a global or classifier-based QoS policy to mark matching packets or configure a global IP-Diffserv classifier.

Table 14 (page 250) displays the switch's default codepoint-priority assignments. If you change the priority of any codepoint to a non-default value and then enter the `write memory` command, the switch will list the non-default setting in the `show config` display.

The default configuration has the following DSCP-priority settings:

Codepoint	Default Priority
001100	1
001101	No-override
001110	2

If you reconfigure these three codepoints to a priority of 3 and then enter the `write memory` command, the switch displays the changes in the `show config` listing:

Figure 47 Displaying non-default priority settings in the DSCP table

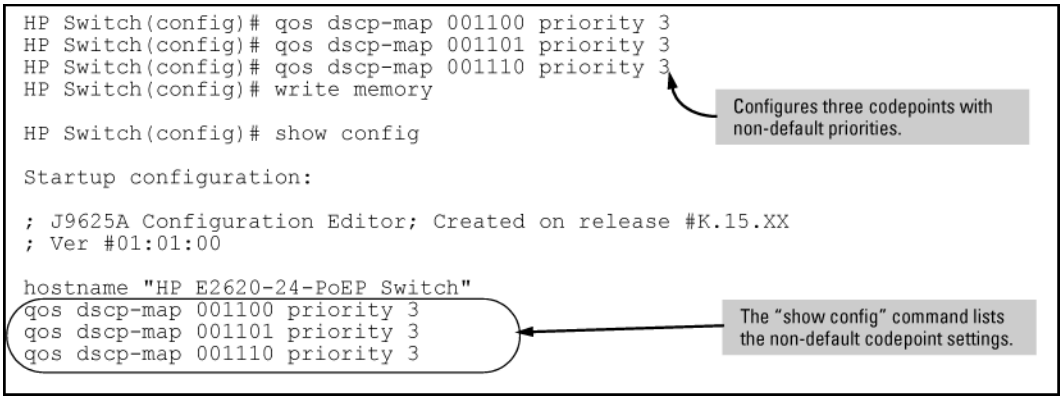

```
HP Switch(config)# qos dscp-map 001100 priority 3
HP Switch(config)# qos dscp-map 001101 priority 3
HP Switch(config)# qos dscp-map 001110 priority 3
HP Switch(config)# write memory

HP Switch(config)# show config

Startup configuration:

; J9625A Configuration Editor; Created on release #K.15.XX
; Ver #01:01:00

hostname "HP E2620-24-PoEP Switch"
qos dscp-map 001100 priority 3
qos dscp-map 001101 priority 3
qos dscp-map 001110 priority 3
```

Configures three codepoints with non-default priorities.

The “show config” command lists the non-default codepoint settings.

Changing the priority setting on a policy when classifiers are currently using the policy (Example)

In this example, the codepoint 000001 is in use by one or more global QoS policies. If you try to modify the priority currently associated with the codepoint, an error message similar to the following is displayed:

```
HP Switch(config)# qos dscp-map 1 priority 2
Cannot modify DSCP Policy 1 - in use by other qos rules.
```

In this case, follow these steps to change the priority:

1. Identify the global and classifier-based QoS policies that use the codepoint whose DSCP-priority mapping you want to change.

Figure 48 Identifying the QoS policies that use a codepoint

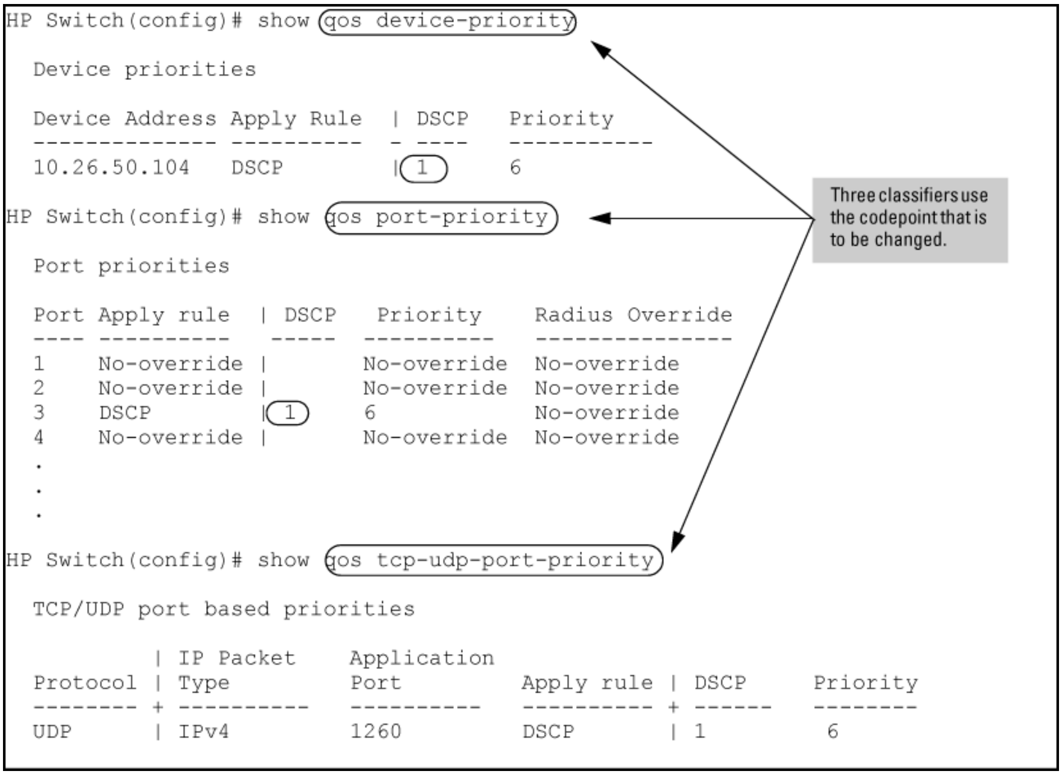

```
HP Switch(config)# show qos device-priority

  Device priorities

  Device Address Apply Rule  | DSCP   Priority
  -------------- ----------  - ----   -----------
  10.26.50.104   DSCP        | 1      6

HP Switch(config)# show qos port-priority

  Port priorities

  Port Apply rule  | DSCP   Priority     Radius Override
  ---- ----------   -----  ----------   ---------------
  1    No-override |       No-override  No-override
  2    No-override |       No-override  No-override
  3    DSCP        | 1     6            No-override
  4    No-override |       No-override  No-override
  .
  .
  .

HP Switch(config)# show qos tcp-udp-port-priority

  TCP/UDP port based priorities

           | IP Packet   Application
  Protocol | Type        Port         Apply rule | DSCP      Priority
  -------- + ---------- ----------   ---------- + ------    --------
  UDP      | IPv4        1260         DSCP       | 1          6
```

Three classifiers use the codepoint that is to be changed.

2. Change each QoS configuration by assigning a different DSCP policy or a different 802.1p priority, or by removing the currently configured DSCP policy and restore the default `No-override` setting; for example:
   a. Delete the current DSCP policy used to mark matching packets for a global IP-device policy (`no qos device-priority` command) and reset the default priority mapping to `No-override`.
   b. Create a new DSCP policy to use when you reconfigure QoS policies to use the new codepoint-priority mapping.
   c. Configure a global QoS source-port policy to mark matching packets with the new DSCP policy.
   d. Assign the global QoS policy that matches `udp-port 1260` packets to a different 802.1p priority.

   ```
   HP Switch(config)# no qos device-priority 10.26.50.104

   HP Switch(config)# qos dscp-map 000100 priority 6
   HP Switch(config)# int 3 qos dscp 000100
   HP Switch(config)# qos udp-port 1260 priority 2
   ```

3. Reconfigure the desired priority for the 000001 codepoint.

   ```
   HP Switch(config)# qos dscp-map 000001 priority 4
   ```

4. Reconfigure QoS policies with the original codepoint (000001) to mark packets with the new DSCP-priority mapping, or leave QoS policies as currently configured from Step 2.

Configuring QoS queues

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 eight 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 | 8-queues
```

Configures the number of outbound priority queues for all ports on the switch using one of the following options: 2-queues, 4-queues, or 8-queues.

Default: 8-queues

The new configuration will:

- Remove any previously configured bandwidth-min output settings
- 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.

△ **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.

Changing the number of priority queues (Example)

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. In addition to setting the number of outbound port queues, the new configuration will remove any previously configured `bandwidth-min output` settings.

To change the number of outbound priority queues for all ports on the switch from eight queues (the default) to four:

1. Configure the number of outbound priority queues by using the `qos queue-config` command.

   ```
   HP Switch(config)# qos queue-config 4-queues
   ```

   A caution message is displayed (see the Caution note above) concluding with the following prompt.

   ```
   Do you wish to proceed? [Proceed/Cancel]
   ```

2. Type **`Proceed`** to continue.

   A second confirmation prompt appears:

   ```
   Please confirm reset. [Yes/Cancel]
   ```

3. Type **`Yes`** to initiate a write memory followed by an immediate reboot. (If you enter **`Cancel`** at either of the two prompts, the command is aborted and the current queue configuration is maintained on the switch).

   The changes will be committed to the startup configuration and the switch will reboot automatically with the new priority queue changes in effect. See Table 25 (page 276) for a listing of the default GMB percentages that are allocated per queue.

Viewing the QoS queue configuration

Syntax:

`show qos queue-config`

Displays the current priority queue configuration and memory allocations per queue. For example:

```
HP Switch# show qos queue-config

        802.1p
Queue   Priority   Memory %
-----   --------   --------
  1       1-2        10
  2       0,3        70
  3       4-5        10
  4       6-7        10
```

Using the outbound queue monitor

NOTE: Outbound queue monitoring is not supported on HP 3800 switches.

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:

```
[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:

```
HP 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.

Example 157 Monitoring egress queues on a port

```
HP 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
```

About QoS

QoS operation

On the switches covered in this guide, QoS operation may be configured through a combination of the following methods:

- Globally-configured, switch-wide QoS settings
- Classifier-based per-port and per-VLAN QoS policies.

Classifier-based QoS policies are designed to work with existing globally-configured, switch-wide QoS settings by allowing you to zoom in on a subset of port or VLAN traffic to further manage it. You can use multiple match criteria to more finely select and define the classes of traffic that you want to manage. QoS policy actions determine how you can handle the selected traffic.

NOTE: While providing greater control for implementing QoS policies, classifier-based QoS policies may override globally-configured QoS settings. For more information, see "Viewing a classifier-based QoSconfiguration" (page 244).

Carefully plan your QoS strategies in advance, identifying the network traffic that you can globally configure and the traffic on which you want to execute customized, classifier-based QoS actions.

Globally-configured QoS

Globally-configured QoS operation supports the following types of packet classification and traffic marking on outbound port and VLAN traffic. For information on how to configure and use global QoS settings, see "Configuring QoS globally" (page 208).

- Globally configured packet classification criteria include:
  - IPv4 device: source and destination address
  - Layer 2 802.1p priority (VLAN header)
  - Layer 3 protocol (such as ARP, IP, IPX, RIP)
  - Layer 3 IPv4 Type of Service (ToS) byte: IP precedence or DSCP bits
  - Layer 3 IPv6 Traffic Class byte: IP precedence or DSCP bits
  - Layer 4 UDP/TCP application port
  - Source port on the switch
  - VLAN ID
- Traffic marking options are as follows:
  - Setting the Layer 2 802.1p priority value in VLAN-tagged and untagged packet headers
  - Setting the Layer 3 Differentiated Services Codepoint (DSCP) bits in the ToS byte of IPv4 packet headers and Traffic Class byte of IPv6 headers.

Classifier-based QoS

Classifier-based QoS feature	Default	Page reference
Classifier-Based QoS Configuration Procedure		237
Override of Global QoS Settings		272
Viewing a Classifier-Based QoS Configuration		244
Classifier-Based QoS Restrictions		273
Classifier-Based QoS Configuration Examples		247
DSCP Policy Table	Various	249
Queue Configuration	8 Queues	252

Starting in release K.14.01, classifier-based QoS operation provides additional QoS actions on a per-port and per-VLAN basis.

- Classifier-based match criteria on inbound IPv4/IPv6 traffic include:
  - IP source address (IPv4 and IPv6)
  - IP destination address (IPv4 and IPv6)
  - IP protocol (such as ICMP or SNMP)
  - Layer 3 IP precedence bits
  - Layer 3 DSCP codepoint

  - Layer 4 UDP/TCP application port
  - VLAN ID
- Classifier-based QoS policy actions on matching IPv4/IPv6 packets are as follows:
  - Setting Layer 2 802.1p priority value (class of service) in VLAN-tagged and untagged packet headers
  - Setting the Layer 3 IP precedence bits
  - Setting the Layer 3 Differentiated-Services Codepoint (DSCP) bits
  - Rate limiting inbound traffic on port and VLAN interfaces

    For information on operation with globally-configured QoS settings, see "Advanced classifier-based QoS" (page 271).

QoS packet classification

To manage network traffic using QoS features, you must first classify (select) the packets you want to manage. You can use any combination of the following packet classification methods to select packets for QoS management:

- Globally configured, switch-wide classification criteria
- Classifier-based match criteria applied to inbound traffic on specific port and VLAN interfaces

NOTE: Starting in software release K.14.01, global and classifier-based QoS policies support IPv6 and IPv4 packet classification.

Using multiple global criteria

NOTE: HP recommends that you configure a minimum number of global QoS classifiers to prioritize a specific packet type. Increasing the number of enabled global QoS classifiers increases the complexity of possible outcomes and consumes switch resources.

The switches covered in this guide provide six types of globally-configured QoS classifiers (match criteria) to select packets for QoS traffic marking.

When multiple, global QoS classifiers are configured, a switch uses the highest-to-lowest search order shown in the following table to identify the highest-precedence classifier to apply to any given packet. When a match between a packet and a classifier is found, the switch applies the QoS policy configured for the classifier and the packet is handled accordingly.

Table 15 Globally-configured packet classification: 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: precedence and DSCP bit sets (IP packets only)
4	4	IP protocol (IP, IPX, ARP, AppleTalk, SNA, and NetBeui)
5	5	VLAN ID
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.

NOTE: On the switches covered in this guide, if the switch is configured with multiple global classifiers that match the same packet, the switch only applies the QoS marking configured for the QoS classifier with the highest precedence. In this case, the QoS configuration for another, lower-precedence classifier that matches is ignored.

For example, if QoS assigns a high priority to packets belonging to VLAN 100 and normal priority to all IP protocol packets, because the IP protocol priority (4) has precedence over the VLAN priority (5), IP protocol packets on VLAN 100 are set to normal priority.

Classifier-based match criteria

In classifier-based packet classification, *match criteria* provide a way to select the packets on which you want to execute QoS actions, such as rate limiting or 802.1p prioritization.

Match criteria are configured by creating a *class* of IPv4 orIPv6 traffic, which contains one or more match or ignore statements. A traffic class may be used by any classifier-based software feature, such as QoS or port mirroring.

By using classifier-based QoS, you can configure multiple match criteria that search multiple fields in packet headers to select the exact traffic you want to rate limit or prioritize for a port or VLAN interface. A classifier-based QoS policy is especially useful when you want to manage different types of traffic in the same way (for example, to prioritize both IP subnet and voice traffic).

For information on how to use match criteria to configure a traffic class, see "Classifier-based software configuration" (page 363).

QoS traffic marking

As described in "QoS operation" (page 254), when you apply or reconfigure QoS actions for selected packets, QoS supports different types of traffic marking in globally-configured QoS settings and classifier-based per-port or per-VLAN QoS policies.

Globally-configured traffic marking

If a packet matches one of the globally-configured packet classifiers, QoS applies one of the following types of traffic marking to the outbound packet:

Layer 2 802.1p prioritization

Controls the outbound port-queue priority for traffic leaving the switch, and (if traffic exits through a VLAN-tagged port) sends the priority setting with the individual packets to downstream devices.

Layer 3 DSCP marking

Enables the switch to set, change, and honor prioritization policies by using the Differentiated Services (diff-serv) bits in the IPv4 ToS byte and IPv6 Traffic Class byte of packet headers.

Layer 2802.1p prioritization

By setting a new 802.1p priority value, QoS allows you to control the priority of outbound packets moving through the switch. The Layer 2 802.1p priority setting in a packet header determines the outbound port queue to which the packet is sent.

By default, the switches covered in this guide have eight outbound traffic queues (0 through 7). A lower-numbered queue has a lower outbound priority; a higher-numbered queue has a higher outbound priority. Packets are transmitted from the switch port on the basis of their queue assignment and whether any higher queues are empty. (To increase bandwidth, you can reconfigure the switch to use four or two outbound queues. See "Configuring QoS queues" (page 252).)

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 queue
1 and 2	Low priority (1, 2)
0 or 3	Normal priority (3, 4)
4 and 5	Medium priority (5, 6)
6 and 7	High priority (7, 8)

If a packet is transmitted in an untagged-VLAN environment, the 802.1p priority settings in the preceding table control only the outbound queue to which the packet is sent on the local switch. Because no VLAN tag is used, an 802.1p priority value is not added to the 802.1Q field in the packet header for use by downstream devices.

However, if your network uses only one VLAN and does not require VLAN-tagged ports, you can preserve 802.1p priority settings in outbound traffic by configuring the ports on links between devices on which you want 802.1p priorities to be honored as tagged VLAN members.

If a packet is transmitted in an 802.1Q VLAN-tagged environment, the QoS-configured 802.1p setting is also added to the VLAN packet header as an 802.1p priority for use by downstream devices and applications.

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 outbound port queue on which packets are forwarded.

Table 17 Mapping 802.1p priorities to outbound port queues on the switch and downstream devices

Configured 802.1p priority	Outbound port queue in the switch	802.1p priority added to tagged VLAN packets exiting the switch	Queue assignment in downstream devices with:		
			8 queues	4 queues	2 queues
1	Queue 1	1 (low priority)	Queue 1	Queue 1	Queue 1
2	Queue 2	2	Queue 2		
0	Queue 3	0 (normal priority)	Queue 3	Queue 2	
3	Queue 4	3	Queue 4		
4	Queue 5	4 (medium priority)	Queue 5	Queue 3	Queue 2
5	Queue 6	5	Queue 6		
6	Queue 7	6 (high priority)	Queue 7	Queue 4	
7	Queue 8	7	Queue 8		

NOTE: You can reconfigure the QoS queue setting to change the number of outbound port queues in the switch from eight (default) to four or two queues. For more information, see "Configuring QoS queues" (page 252).

Layer 3 DSCP marking

By changing or honoring the settings of the DSCP codepoint in IP packet headers, QoS allows you to control the DSCP and associated 802.1p priority values in outbound IP packets that are sent to downstream devices.

You can later configure downstream devices to read and use the DSCP policy that QoS sets. When marking the DSCP bits in IP packets, a QoS policy is not dependent on VLAN-tagged ports to carry

802.1p packet priorities to downstream devices (as shown in "QoS traffic marking supported in tagged and untagged VLANs" (page 259)).

When configuring a Layer 3 DSCP policy, specify:

- Bit values for the DSCP codepoint (the upper six bits in the ToS/Traffic Class byte in IP packet headers), entered in either binary format, the decimal equivalent, or an ASCII standard (hexadecimal) name
- An 802.1p priority value that is associated with the new DSCP bit values

  Certain DSCP codepoints (such as Assured Forwarding and Expedited Forwarding) have default 802.1p priorities as shown in "The default DSCP policy table" (page 250).

A DSCP policy assigns a DSCP codepoint and 802.1p priority value to IPv4 and IPv6 packets. As shown in "Application of Differentiated Services Codepoint (DSCP) policies" (page 207), you can classify traffic on an edge switch and use Layer 3 DSCP-marking (instead of only 802.1p priority) to assign and preserve QoS policies on downstream devices. In this case, if you reconfigure the 802.1p priority associated with the DSCP codepoint, the new 802.1p assignment takes effect starting on the switch on which it is configured and is used in packets sent to downstream devices.

If you configure a different 802.1p priority for a DSCP codepoint, the new DSCP policy overrides the 802.1p priority value in packets which enter the switch with the specified codepoint. The Layer 2 802.1p priority setting (0 through 7) determines the outbound port queue to which a packet is sent (as shown in Table 16 (page 258)).

VLAN and untagged VLAN environments

QoS operates in VLAN-tagged and untagged environments. If your network does not use multiple VLANs, you can still implement the 802.1Q VLAN capability to allow packets to carry an 802.1p priority to the next downstream device. To do so, configure the ports on links to other network devices as VLAN-tagged members.

In a tagged or untagged VLAN, you can also ensure that IPv4/IPv6 packets carry an 802.1p priority to downstream devices by configuring DSCP marking in the ToS/Traffic Class byte.

The following table summarizes the QoS options for traffic-marking in VLAN-tagged and untagged environments.

Table 18 QoS traffic marking supported in tagged and untagged VLANs

QoS marking supported on outbound packets	Port membership in VLANs	
	Tagged	Untagged
Assign an 802.1p priority that determines the outbound port queue to which a packet is sent	Supported	Supported
Carry the 802.1p priority to the next downstream device	Supported	Not Supported
Carry a DSCP policy (DSCP codepoint[1] and associated 802.1p priority[2]) to downstream devices	Supported	Supported

[1] DSCP marking (DSCP codepoint and associated 802.1p priority) are not supported on non-IP packets and packets selected using the following global QOS classifiers: Layer 3 Protocol and IP-Precedence. Also, in order for DSCP policy marking to be honored on a downstream device, the device must be configured to use the DSCP policy in IP packet headers.

[2] The 802.1p priority associated with a DSCP codepoint (see "The default DSCP policy table" (page 250)) is used to determine the packet's outbound port queue. When used in a VLAN-tagged environment, an 802.1p priority is also carried in the 802.1Q field of outbound packet headers.

Classifier-based traffic marking

Classifier-based per-port or per-VLAN QoS policies support the following traffic-marking actions. Note that in addition to globally-configured QoS traffic marking (802.1p and DSCP prioritization), classifier-based QoS policies also support IP precedence and rate limiting.

Layer 2802.1p prioritization

Controls the outbound port queue priority for traffic leaving the switch, and (if traffic exits through a VLAN-tagged port) sends the priority setting in packet headers to downstream devices.

Layer 3 IP precedence-bitmarking

Enables the switch to set, change, and honor prioritization policies by using the IP precedence bits in the ToS byte of IPv4 packet headers and Traffic Class byte of IPv6 headers.

Layer 3 DSCP marking

Enables the switch to set, change, and honor prioritization policies by using the Differentiated Services (Diffserv) bits in the ToS byte of IPv4 headers and Traffic Class byte of IPv6 headers.

Rate limiting

Enables a port or VLAN interface to allow only the specified amount of bandwidth to be used for inbound traffic. When traffic exceeds the configured limit, it is dropped.

For information on how to configure and use classifier-based QoS policies, see "Configuring classifier-based QoS" (page 237).

NOTE: After you apply a classifier-based QoS policy on a port or VLAN interface:

- The 802.1p (CoS) priority and DSCP codepoint marking applied to classified packets override any 802.1p and DSCP codepoint values that are globally-configured using the QoS commands, described in "Configuring QoS globally" (page 208).
- The rate limit applied to classified packets overrides any globally configured rate limit globally-configured with the commands described in the Port Traffic Controls chapter in the *Management and Configuration Guide*.

For more information on how classifier-based traffic marking overrides globally-configured traffic marketing, see "Override of global QoS settings" (page 272).

No override

By default, the `show qos` output for following global QoS classifiers may display `No-override` for QoS marking: IP Precedence, IP Diffserv, Layer-3 Protocol, VLAN ID, and Source-port (see "Displaying `show qos` output" (page 261)). `No-override` means that the global QoS policy used to mark matching packets does not assign an 802.1p value.

- 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.

Example 158 Show QoS command output

"Displaying `show qos` output" (page 261)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 49 Displaying `show qos` output

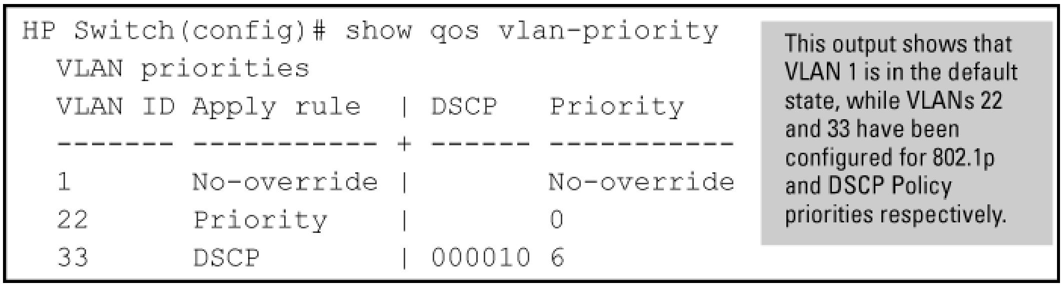

```
HP Switch(config)# show qos vlan-priority
  VLAN priorities
  VLAN ID Apply rule  | DSCP   Priority
  ------- ----------- + ------ -----------
  1       No-override |        No-override
  22      Priority    |        0
  33      DSCP        | 000010 6
```

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 QoS restrictions

This table shows the packet types supported by different global QoS classifiers and DSCP marking.

Table 19 Restrictions for global QoS support

Type of packets supported	Global QoS classifiers							DSCP overwrite (re-marking)
	TCP/UDP	IP Device	IP Type-of-Service	Layer 3 Protocol	VLAN ID	Source Port	Incoming 802.1p	
IP packets (IPv4 andIPv6 [1]) only	Yes	Yes	Yes	No	No	No	No	Yes
Layer-2 SAP encapsulation	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

[1] Globally-configured QoS supports IPv6 packets starting in release K.14.01.

All switches

For explicit QoS support of IP subnets, HP 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. For more on VLANs, see "Static Virtual LANs (VLANs)" (page 14).

Maximum global QoS remarking entries

The switches covered in this guide accept the maximum number of configured outbound 802.1p priority and DSCP entries shown in the following table.

Table 20 Maximum number of QoS entries.

Switch	Maximum QoS remarking	Notes
3800 Switches		• Each IP Device (IP address) QoS configuration uses two entries. • Each TCP/UDP Port QoS configuration uses two entries. • All other global QoS classifier configurations use one entry each.
Switch 8212zl Series 5400zl	250[1] configured entries	
Series 5300yl		

[1] Configuring IP Device (IP address) and TCP/UDP global QoS classifiers reduces this maximum. For more information, see the Notes column.

If the global QoS configurations on a switch exceed the maximum number of entries shown in Table 20 (page 262), the following error message is displayed:

```
Unable to add this QoS rule. Maximum number (entry-#) already
reached.
```

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.

Fragmented packets and TCP/UDP

QoS is not performed on fragmented packets under TCP/UDP.

Monitoring shared resources

The QoS 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, additional QoS provisions 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 Management and Configuration Guide for your switch.

Global QoS classifiers

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.

NOTE: Starting in software release K.14.01, global TCP/UDP classifiers are supported on IPv4, IPv6, or both IPv4 and IPv6 packets. In previous releases, only IPv4 packets were supported.

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 an associated 802.1p priority; inbound packets can be IPv4 or IPv6.)

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:

www.iana.org

Then click on:

Protocol Number Assignment Services

P under **Directory of General Assigned Numbers**)

Port Numbers

Operating notes on using TCP/UDP port ranges

- Only 6 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:

```
HP Switch(config)# qos udp-port range 1300 1399 dscp 001110
HP Switch(config)# no qos range 1300 1399
```

Example 159 Configuration for TCP and UDP port prioritization

The following example displays the following 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
TCP Port 80 (World Wide Web HTTP)	2	2
UDP Port 80 (World Wide Web HTTP)	1	1

Figure 50 Configuring 802.1p priority assignments on TCP/UDP ports

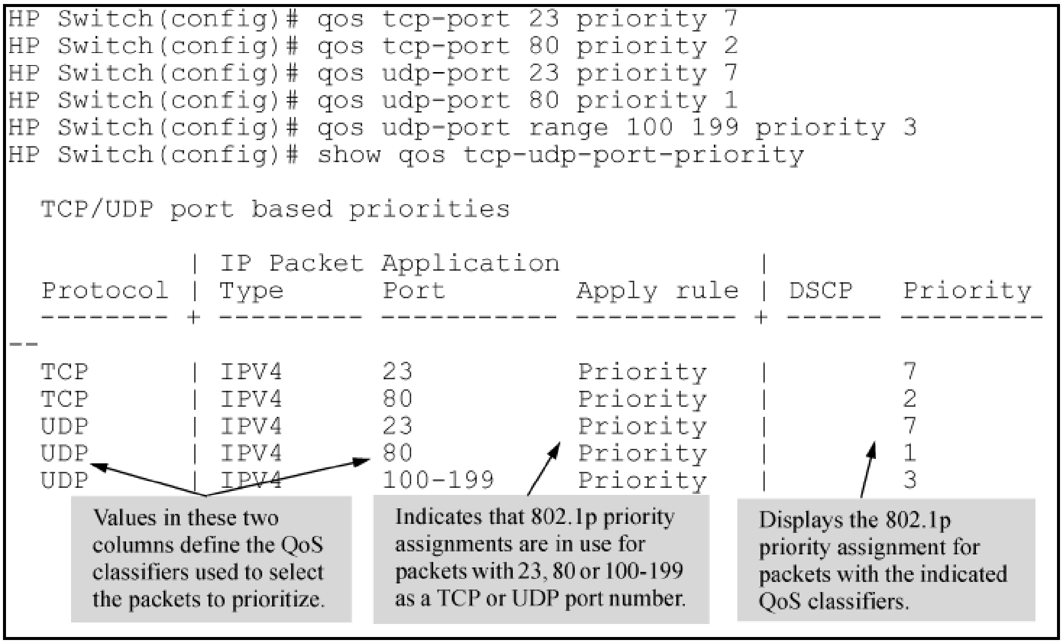

```
HP Switch(config)# qos tcp-port 23 priority 7
HP Switch(config)# qos tcp-port 80 priority 2
HP Switch(config)# qos udp-port 23 priority 7
HP Switch(config)# qos udp-port 80 priority 1
HP Switch(config)# qos udp-port range 100 199 priority 3
HP Switch(config)# show qos tcp-udp-port-priority

  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.

About global IP-device classifier

Global QoS classifier precedence: 2

The global IP-device classifier enables you to configure up to 300 IP addresses to select IP packets according to source *or* destination address.

NOTE: IPv6 Support: Starting in software release K.14.01, IP device classifiers are supported on IPv4,IPv6, and IPv4/IPv6 subnets. In previous releases, only IPv4 packets are supported.

When a globally-configured IP-device address has the highest precedence in the switch for traffic addressed to or from the device, traffic received on the switch with the configured IP address is marked with the specified priority level. You can configure different IP-device classifiers with different priority levels.

NOTE: QoS IP-Device Restriction: The configuration of a QoS IP-device priority on the Management VLAN IP address (if configured) is not supported. If no Management VLAN is configured, the configuration of a QoS IP-device priority on the default VLAN IP address is not supported.

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 information on global QoS operation when other global classifiers apply to the same traffic, see to "Using multiple global criteria" (page 256).

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.

Global IP type-of-service classifier

Global QoS classifier precedence: 3

The global IP Type-of-Service classifier enables you to classify and mark IP packets according to the following modes:

IP-precedence mode

All IP packets generated by upstream devices and applications include a precedence bit set in the ToS/Traffic Class byte. In IP-precedence mode, the switch uses the precedence bits to compute and assign the corresponding 802.1p priority.

IP Differentiated Services (Diffserv) Mode

The Diffserv mode uses the codepoints set in IP packets by upstream devices and applications to assign an 802.1p priority to packets. You can use Diffserv mode to mark packets in the following ways:

Assign a new DSCP policy: A policy includes both a codepoint and a corresponding 802.1p priority. This option selects an incoming IP 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 "Using Differentiated Services Codepoint (DSCP) mapping" (page 249).)

Assign an 802.1p priority: This option reads the DSCP of an incoming IP packet and, without changing this codepoint, assigns the 802.1p priority to the packet, as configured in the DSCP Policy Table (see "Using Differentiated Services Codepoint (DSCP) mapping" (page 249)). This means that a priority value of 0 - 7 must be configured for a DSCP before the switch can perform a QoS match on the packet's DSCP bits.

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 (`show qos dscp-map` command), you must first configure a priority for the codepoint before proceeding (`qos dscp-map priority`command). See "Using Differentiated Services Codepoint (DSCP) mapping" (page 249). Note that some 802.1p priorities are assigned by default to well-known DSCP codepoints, such as the "Assured Forwarding" and "Expedited Forwarding" codepoints (see "The default DSCP policy table" (page 250)).

Unless IP-Precedence mode and Diffserv mode are both disabled (the default setting), enabling one automatically disables the other. For more information on Type-of-Service operation, see "IPv4 ToS/IPv6 traffic class byte" (page 267).

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. (For information on QoS operation when other global QoS classifiers match the same traffic, see "Using multiple global criteria" (page 256).)

Global VLAN-ID classifier

Global QoS Classifier Precedence: 5

The global VLAN-ID (VID) classifier allows you to use up to 4094 VLAN IDs to match packets. When a particular VLAN-ID classifier has the highest precedence in the switch, traffic received in the VLAN is marked with the configured priority level. You can configure different global VLAN-ID classifiers to mark packets with different priority levels.

Options for assigning priority

The global QoS packet-marking options for packets that carry a specified VLAN-ID include:

- 802.1p priority
- DSCP policy (Assigning a new DSCP and an associated 802.1p priority; inbound packets must be IPv4.)

For information on QoS operation when other global QoS classifiers match the same traffic, see to "Using multiple global criteria" (page 256).

NOTE: A global VLAN-ID classifier marks priority levels only in packets received on static VLANs. Packets received in a dynamic VLAN created byGVRP operation are not marked by a global VLAN-ID classifier.

The VLAN ID used as a global QoS classifier must currently exist on the switch. If you remove a VLAN from the switch, all global QoS configurations that use the VLAN ID for packet marking are also removed.

Global source-port classifier

Global QoS Classifier Precedence: 6

The global QoS source-port classifier allows you to use a packet's source-port on the switch to mark packets. When a global source-port classifier has the highest precedence in the switch for traffic entering through a port, traffic received on the port is marked with the configured priority level. Different source-port classifiers can have different priority levels.

Options for assigning priority on the switch

The global QoS packet-marking options for matching packets from a specified source-port include:

- 802.1p priority
- DSCP policy: Assigning a new DSCP and an associated 802.1p priority

For information on QoS operation when other global QoS classifiers match the same traffic, see to "Using multiple global criteria" (page 256).

Options for assigning priority from a RADIUS server

You can use a RADIUS server to assign a QoS source-port priority during an 802.1X port-access authentication session. See the RADIUS chapter in the *Access Security Guide* for your switch.

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 *Access Security Guide* for your switch.

IPv4 ToS/IPv6 traffic class byte

IPv4 packet headers contain a Type of Service (ToS) byte; IPv6 packet headers contain a Traffic Class byte. In an IPv6 packet, the Traffic Class byte is used in the same way as the ToS byte in an IPv4 packet. A ToS/Traffic Class byte includes a DSCP codepoint and precedence bits:

- Differentiated Services Codepoint (DSCP)

  Consists of the upper six bits of the ToS/Traffic Class byte. There are 64 possible codepoints.

  In the switches covered in this guide, the default QoS configuration includes some codepoints, such as Assured Forwarding and Expedited Forwarding, that are preconfigured with an 802.1p priority setting. All other codepoints are not configured with an 802.1p priority and display `No-override` as shown in the default DSCP Policy table ("The default DSCP policy table" (page 250)).

  Use the `qos dscp map` command to configure the switch to assign different 802.1p priorities to IP packets with different codepoints. Also, you can configure the switch to assign a new codepoint with its associated priority level (0-7) to matching packets as follows:

  1. Configure a DSCP codepoint with the desired priority in an edge switch.
  2. Configure the local switch to mark specified inbound traffic with the DSCP (and thus create a policy for that traffic type).
  3. Configure the internal switches in your LAN to honor the policy.

  For example, you could configure an edge switch to assign a codepoint of 000001 to all packets received from a specific VLAN, and then handle all traffic with that codepoint at high priority.

  For a codepoint listing and the commands for displaying and changing the DSCP Policy table, see "Using Differentiated Services Codepoint (DSCP) mapping" (page 249).

- Precedence Bits

  A subset of the DSCP codepoint, consisting of the upper three bits of the ToS/Traffic Class byte. When a global IP-Precedence classifier is configured, the switch uses the precedence bit set to determine the priority for selected packets as shown in the following table. (The switch does not change the setting of the precedence bits.)

Table 21 IP precedence-to-802.1p priority mapping

ToS/Traffic Class Byte: IP Precedence Bits	Corresponding 802.1p Priority	Service Priority Level
000	1	Lowest
001	2	Low
002	0	Normal
003	3	
004	4	
005	5	
006	6	
007	7	Highest

NOTE: Using a global IP-Precedence classifier to prioritize IP packets relies on priorities set in upstream devices and applications.

Figure 51 (page 268) shows the difference between the diffserv bits and precedence bits in an IPv4 ToS byte and an IPv6 Traffic Class byte. Note that:

- Precedence bits are a subset of the Differentiated Services bits.
- The right-most two bits are reserved.

Figure 51 IPv4 ToS/IPv6 traffic class byte with DSCP codepoint and precedence bits

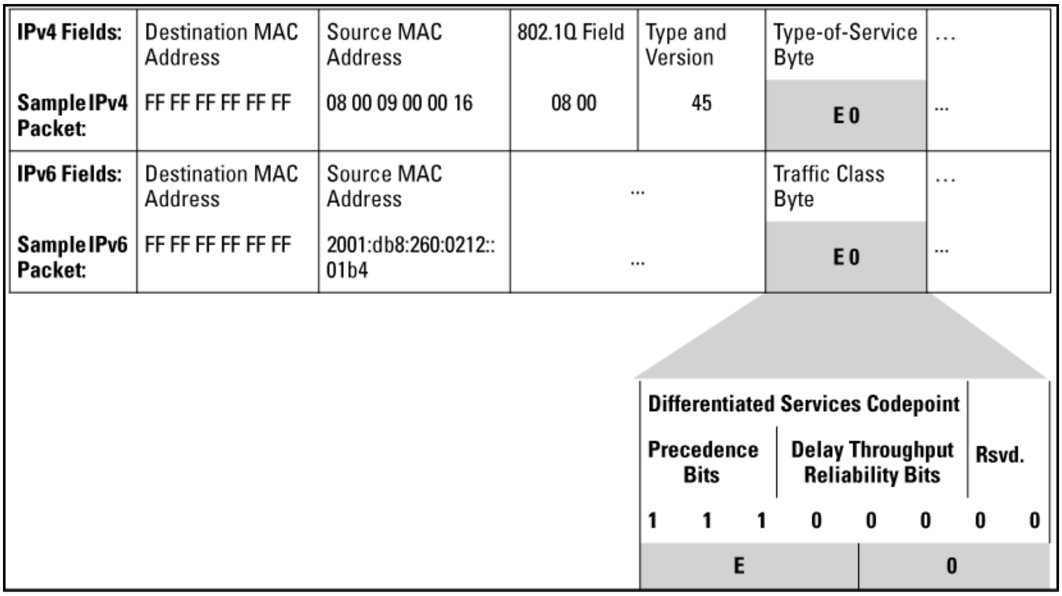

IPv4 Fields:	Destination MAC Address	Source MAC Address	802.1Q Field	Type and Version	Type-of-Service Byte	...
Sample IPv4 Packet:	FF FF FF FF FF FF	08 00 09 00 00 16	08 00	45	E 0	...
IPv6 Fields:	Destination MAC Address	Source MAC Address	...		Traffic Class Byte	...
Sample IPv6 Packet:	FF FF FF FF FF FF	2001:db8:260:0212::01b4	...		E 0	...

Differentiated Services Codepoint						**Rsvd.**	
Precedence Bits			**Delay Throughput Reliability Bits**				
1	1	1	0	0	0	0	0
E				0			

Assigning an 802.1p priority for a global IP-diffserv classifier

One of the best uses for this global QoS packet-marking option is on an interior switch to honor (continue) a policy set on an edge switch. The IP-diffserv classifier enables selecting incoming packets having a specific DSCP and forwards these packets with the desired 802.1p priority. For example, if an edge switch A marks all packets received on port A5 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 52 Interior switch B honors the policy established in edge switch A

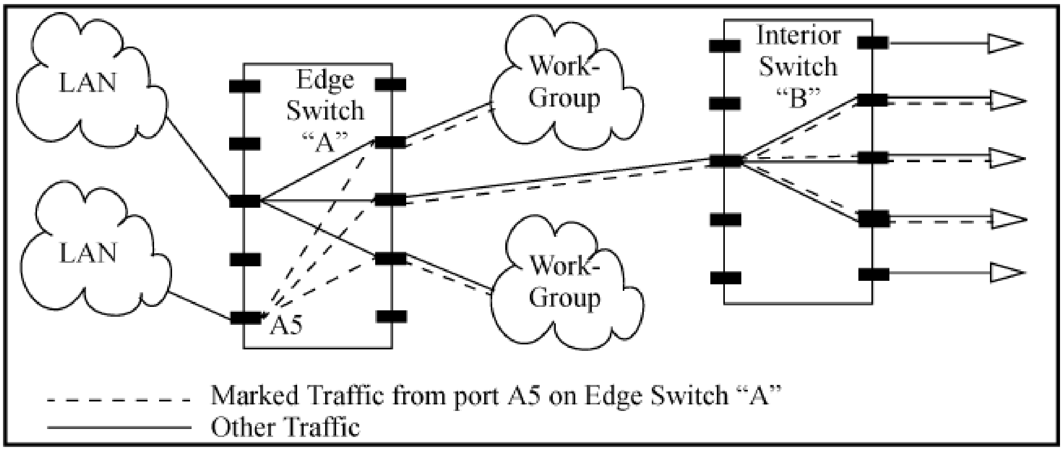

To do so, assign the desired 802.1p priority to the same codepoint that the upstream or edge switch assigns to the selected packets. When the downstream switch receives an IP packet carrying one of these codepoints, it assigns the configured priority to the packet and sends it out the appropriate outbound port 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, as long as the DSCPs specified in the two options do not match.

NOTE: 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.

Comparing global IP type-of-service classifiers

The next table shows the difference in how global IP-Precedence and IP-Diffserv classifiers are implemented in the switch.

Outbound port	IP Type-of-Service classifiers	
	IP-Precedence mode	**IP differentiated services mode**
IP Packet Sent Out an Untagged Port in a VLAN	Based on the IP Precedence bit set in a packet's ToS/Traffic Class field, the packet is sent to one of eight outbound port queues in the switch: • 1 - 2 = low priority (queue 1, 2) • 0 - 3 = normal priority (queue 3, 4) • 4 - 5 = medium priority (queue 5, 6) • 6 - 7 = high priority (queue 7, 8)	Based on the DSCP codepoint that the switch has been configured to detect, one of the following actions is taken: • The codepoint is re-marked according to the configured DSCP policy and the 802.1p priority currently configured for the codepoint in the DSCP Policy see (Table 21 (page 267)). • The codepoint is not changed, but the 802.1p priority is marked with the currently configured value for the codepoint in the DSCP Policy table. Based on the new 802.1p priority marking, the packet leaves the switch through one of the following queues: • 1 - 2 = low priority (queue 1, 2) • 0 - 3 = normal priority (queue 3, 4) • 4 - 5 = medium priority (queue 5, 6) • 6 - 7 = high priority (queue 7, 8) If `No-override` (the default) is configured for the 802.1p priority associated with a codepoint, the priority in the packet header is not re-marked by the global IP-Diffserv classifier and, by default, is sent to the "normal priority" outbound port queue.
IP Packet Sent Out a Tagged Port in a VLAN	Based on the IP Precedence bit set in a packet's ToS/Traffic Class field: • The packet is sent to one of eight outbound port queues in the switch as described above. • The IP Precedence value (0 - 7) is used to set the corresponding 802.1p priority in the VLAN tag carried by the packet to the next downstream device (see Table 21 (page 267)).	Based on the DSCP codepoint that the switch has been configured to detect, one of the following actions is taken: • The codepoint is re-marked according to the configured DSCP policy and the 802.1p priority currently configured for the codepoint in the DSCP Policy Table (Table 21 (page 267)). • The codepoint is not changed, but the 802.1p priority is marked with the currently configured value for the codepoint in the DSCP Policy Table (Table 21 (page 267)). Based on the new 802.1p priority marking, the packet leaves the switch through one of the outbound port queues described above. In addition, the priority value (0 - 7) is used to set the 802.1p priority in the VLAN tag carried by the packet to the next downstream device. If the priority is configured as `No-override` in the DSCP Policy table, the VLAN tag carries a "0" (normal priority) 802.1p setting if not prioritized by other global QoS classifiers.

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.

Advanced classifier-based QoS

Starting in software release K.14.01, in addition to the packet classification and prioritization methods described in “Configuring QoS globally” (page 208), QoS configuration also supports advanced classifier-based functions. Advanced classifier-basedQoS introduces:

- A finer granularity than globally-configured QoS for classifying IPv4 andIPv6 traffic
- Additional actions for managing selected traffic, such as rate limiting and IP precedence marking
- The application of QoS policies to inbound traffic flows on specific port and VLAN interfaces (instead of using only globally-configured, switch-wide QoS settings)
- The ability to re-use traffic classes in different software-feature configurations, such as QoS and port mirroring

Classifier-based QoS is designed to work with existing globally-configured, switch-wide QoS policies by allowing you to zoom in on a subset of port or VLAN traffic to further manage it. Classifier-based policies take precedence over, and may override, globally-configured QoS settings that apply to all traffic on the switch.

Classifier-based QoS policies provide greater control for managing network traffic. Using multiple match criteria, you can finely select and define the classes of traffic that you want to manage. QoS-specific policy actions determine how you can handle the selected traffic.

For more information, see the "Classifier-Based Software Configuration" chapter in the *Advanced Traffic Management Guide*.

Classifier-based QoS model

Classifier-based QoS configuration consists of the following general steps:

1. Classify the traffic that you want to manage by configuring a class.
2. Configure a QoS policy in which you specify the QoS actions to execute on each class of traffic.
3. Assign the QoS policy to a port or VLAN (inbound only) interface.

NOTE: Classifier-based QoS operation supports all globally-configured packet classification criteria (except for Source-port and Layer-3 protocol) and traffic marking functions, and provides additional QoS actions on a per-port and per-VLAN basis.

- Classifier-based match criteria on inbound IPv4/IPv6 traffic include:
  - IP source address (IPv4 and IPv6)
  - IP destination address (IPv4 and IPv6)
  - IP protocol (such as ICMP or SNMP)
  - Layer 3 IP precedence bits
  - Layer 3 DSCP codepoint
  - Layer 4 TCP/UDP application port (including TCP flags)
  - VLAN ID
- Classifier-based QoS policy actions on matching IPv4/IPv6 packets are as follows:
  - Setting the Layer 2 802.1p priority value (class of service) in VLAN-tagged and untagged packet headers
  - Setting the Layer 3 IP precedence bits
  - Setting the Layer 3 Differentiated Services Codepoint (DSCP) bits
  - Rate limiting inbound traffic on port and VLAN interfaces

Override of global QoS settings

After you apply a QoS policy to an interface, the classifier-based settings configured by QoS actions in the policy override any 802.1p CoS or DSCP codepoint values that were globally-configured on the switch to mark packets using the QoS commands described in "Configuring QoS globally" (page 208).

If you use a classifier-based QoS configuration along with globally-configured QoS commands, the order of precedence in which 802.1p priority, IP precedence, and DSCP settings mark selected packets is as follows, from highest (1) to lowest (9):

Table 22 Order of precedence for classifier-based QoS over global QoS

Precedence order	QoS feature	Reference
1	Classifier-based port-specific policy	(page 237)
2	Classifier-based VLAN-specific policy	(page 237)
3	Globally-configured TCP/UDP priority	(page 262)
4	Globally-configured IP-device priority	(page 264)
5	Globally-configured IP Type-of-Service priority	(page 265)
6	Globally-configured Layer 3-Protocol priority	(page 266)
7	Globally-configured VLAN-ID priority	(page 227)
8	Globally-configured Source-Port priority	(page 266)
9	802.1p CoS in Layer 2 VLAN header[1]	(page 256)

[1] In a tagged VLAN environment, the incoming 802.1p priority is used as the default QoS classifier to determine how a packet is handled if no global or classifier-based QoS match criterion with a higher precedence matches.

Effect of No-override

If you configure a global IP-Diffserv classifier and `No-override` is displayed for the 802.1p priority associated with a codepoint, DSCP marking cannot be performed on matching outbound packets. However, QoS does not affect the packet-queuing 802.1p priority or VLAN tagging carried in the packet.

In this case, the packets are handled as follows (as long as no other QoS classifier marks a new 802.1p priority on the matching packets):

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

Classifier-based QoS restrictions

The following restrictions apply to QoS policies configured with the classifier-based model:

- A classifier-based QoS policy cannot be applied on a port or VLAN interface on which a classifier-based QoS policy is already configured. It is possible to apply a classifier-based policy of a different type, such as port mirroring.
- A QoS policy that uses the `rate-limit` command is not supported on a port interface on which ICMP rate limiting has already been globally configured. To apply the QoS policy, you must first disable the ICMP rate limiting configuration. For more information, see the Configuring ICMP section in the Configuring IP Parameters for Routing Switches chapter in the *Multicast and Routing Guide*.

  In cases where an ICMP rate limiting configuration is to be maintained, configure a QoS policy by adding the necessary `match` statements for the ICMP traffic in a class configuration, then configure a `rate-limit` action for the class in the policy configuration.
- In a QoS policy that uses the `class action rate-limit` command, the rate limit is calculated on a per-module or per port-bank basis. If trunked ports or VLANs with a configured rate limit span multiple modules or port-banks, the configured rate limit is not guaranteed.
- In a QoS policy that uses the `class action dscp` command, the DSCP value entered must be already configured with an 802.1p priority in the DSCP Policy table (see “The default DSCP policy table” (page 250)).

Interaction with other software features

After applying a QoS policy to an interface, an error message appears if there are not sufficient hardware resources to support the policy. In this case, use the `show resources` command to verify the amount of resources that are currently in use and the resources available on the switch. QoS policies share the same hardware resources with other software features, such as mirroring policies, ACLs, virus throttling, the management VLAN, and so on.

Use the displayed information to decide whether to re-prioritize current resource usage by reconfiguring or disabling software features to free the resources reserved for less important features.

For more information, see “Displaying resource usage for QoS policies” (page 214) and the Monitoring Resources chapter in the *Management and Configuration Guide*.

Notes on changing priority settings

If you try to modify the priority associated with a DSCP codepoint in a DSCP policy using the `qos dscp-map priority` command, and if the DSCP policy is currently used by one or more global QoS or classifier-based QoS policies, the following error message is displayed:

```
Cannot modify DSCP Policy codepoint - in use by other qos rules.
```

In this case, enter the following QoS `show` commands to identify in which global and classifier-based QoS configurations the DSCP policy is currently used:

```
show policy qos-policy
show qos tcp-udp-port-priority
show qos device-priority
show qos type-of-service
show qos protocol
show qos vlan
show qos port-priority
```

After determining the QoS configurations in which the DSCP-priority mapping is used, you can either delete a QoS configuration and reset the DSCP-priority mapping to `No-override`, or change either the 802.1p priority or the codepoint used in the QoS configuration.

Example

Example 160 Changing the priority of a codepoint

If codepoint 000001 is currently mapped to priority 6, and several global QoS policies use this codepoint to assign a priority to their respective types of matching traffic, you can change the priority associated with the codepoint using the following procedure.

1. Identify the global and classifier-based QoS policies that use the codepoint.
2. Do one of the following:
   a. Reconfigure each QoS policy by re-entering a different DSCP codepoint or a different 802.1p priority associated with the codepoint.
   b. Enter the `no qos` *`classifier`* or `no policy` *`qos-policy`* command to remove the current DSCP policy with codepoint 000001 and reset the priority to `No-override`.
3. Use the `qos dscp-map 000001 priority 0 - 7` command to remap DSCP 000001 to the desired priority.
4. Do one of the following:
   a. Reconfigure codepoint 000001 in the QoS policies in which you want to use the new DSCP-priority mapping to mark matching packets.
   b. Leave a QoS policy in which you use DSCP 000001 to mark matching packets with the associated `No-override` priority mapping.

Error messages for DSCP policy changes

See the error messages in the following table to troubleshoot an error condition that results from reconfiguring a DSCP policy.

Table 23 Error messages generated by DSCP policy changes

Error message	Description
DSCP Policy *`decimal-codepoint`* not configured	You have tried to configure a codepoint in a global or classifier-based QoS policy for which there is no associated priority (No-override). Use the `qos dscp-map` command to configure a priority for the codepoint, then re-enter the codepoint in the QoS configuration.
Cannot modify DSCP Policy *`codepoint`* - in use by other qos rules.	You have tried to configure a codepoint in a global or classifier-based QoS policy that is already in use by other QoS policies. Before remapping the codepoint to a new priority, you must first reconfigure the other QoS policies so that they do not use the current codepoint-priority mapping. You can have multiple QoS policies that use the same codepoint to mark packets as long as it is acceptable for all such policies to use the same 802.1p priority.

Mapping of outbound port queues

This table shows the mapping of 802.1p priorities to outbound port queues.

Table 24 Mapping 802.1p priorities to outbound port queues

802.1p priority	8 Queues (default)	4 Queues	2 Queues
1 (lowest)	1	1	1
2	2		
0 (normal)	3	2	
3	4		
4	5	3	2
5	6		
6	7	4	
7 (highest)	8		

Impact of QoS queue configuration on guaranteed minimum bandwidth (GMB)

Changing the number of queues removes any `bandwidth-min output` settings in the startup configuration, and automatically re-allocates the GMB per queue as shown in the following table.

Table 25 Default GMB percentage allocations per QoS queue configuration

802.1p priority	8 Queues (default)		4 Queues		2 Queues	
	Queue	**GMB**	**Queue**	**GMB**	**Queue**	**GMB**
1 (lowest)	1	2%	1	8%	1	20%
2	2	3%				
0 (normal)	3	30%	2	17%		
3	4	10%				
4	5	10%	3	30%	2	80%
5	6	10%				
6	7	15%	4	45%		
7 (highest)	8	20%				

NOTE: For more information on configuring GMB, see the chapter titled *"Port Traffic Controls"* in the *Management and Configuration Guide.*

Setting minimum guaranteed bandwidth with 8 queues

When 10 Mbps ports on an 8200zl or 5400zl switch are configured in QoS for eight outbound queues (the default), and the guaranteed minimum bandwidth is set for 5 Mbps or less for a given queue, then packets in the lower-priority queues may be discarded on ports that are oversubscribed for extended periods of time. If the oversubscription cannot be corrected, HP recommends reconfiguring the switch to operate with four outbound queues. The command to do this is:

```
HP Switch(config)# qos queue-config 4-queues
```

This issue applies to 8200zl and 5400zl switches operating with any of the following modules installed.

HP device	Product number	Minimum supported software version
HP Switch 24-port 10/100/1000 PoE+v2 zl Module	J9534A	K.15.02.0004
HP Switch 20-port 10/100/1000 PoE+ / 4-port SFP v2 zl Module	J9535A	K.15.02.0004
HP Switch 20-port 10/100/1000 PoE+ / 2-port 10-GbE SFP+ v2 zl Module	J9536A	K.15.02.0004
HP Switch 24-port SFP v2 zl Module	J9537A	K.15.02.0004
HP Switch 8-port 10-GbE SFP+ v2 zl Module	J9538A	K.15.02.0004
HP 24-port 10/100 PoE+ v2 zl Module	J9547A	K.15.02.0004
HP 20-port Gig-T / 2-port 10-GbE SFP+ v2 zl Module	J9548A	K.15.02.0004
HP 20-port Gig-T / 4-port SFP v2 zl Module	J9549A	K.15.02.0004

HP device	Product number	Minimum supported software version
HP 24-port Gig-T v2 zl Module	J9550A	K.15.02.0004
HP 12-port Gig-T / 12-port SFP v2 zl Module	J9637A	K.15.02.0004

7 Stack management for the 3500, 3500yl, 6200yl and 6600 switches

Command syntax	Description	Default value	CLI page reference	Menu page reference
`show stack [ candidates | view | all ]`	Displays stack status		282	283 303 303
`stack commander` *`name-str`*	Makes a switch a Commander		285	
`no stack stack commander` *`stack name`*	Converts a Member to be Commander of new stack		286	
`stack member` *`switch-number mac-address mac-addr`* `[ password` *`password-str`* `]`	Adds to a stack, and moves switches between stacks	Commander's auto-grab = No	287	289
`[no] stack auto-join`	Uses auto-join on a Candidate	Candidate's auto-join = Yes	292	
`[no] stack join` *`mac-addr`*	Pushes a Candidate into a stack		292	
`stack member` *`switch-number mac-address mac-addr`* `[ password` *`password-str`*`]`	Pulls a Member from a stack		293	
`stack join` *`mac-addr`*	Pushes a Member into a stack		294	
`no stack name` *`stack name`* `stack join` *`mac-address`*	Converts a Commander to Member of another stack		294	
`[no] stack member` *`switch-num`* `mac-address` *`mac-addr`*	Removes a Member from a stack		295	297
`telnet` *`switch-number`*	Accesses Member switches for configuration changes and traffic monitoring		298	298
`no stack`	Disables or re-enables stacking		300	
`stack transmission-interval` *`seconds`*	Sets the transmission	60	300	

Command syntax	Description	Default value	CLI page reference	Menu page reference
	interval in seconds			

Introduction

This feature is available on the 3500, 3500yl, 6200yl and 6600 switches, but not on the 5400zl and 8200zl switches.

HP Switch Stack 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:

- Stacking and meshing cannot both be enabled at the same time.
- In the default configuration, stacking in the "candidate" state is enabled.
- For additional rules and restrictions, see “Operating rules for stacking” (page 306).

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.
- If you are including an HP Switch 8000M, 4000M, 2424M, 2400M, or 1600M in a stack, you must first update all such devices to software versionC.08.03 or later. Copies of the latest software version are available from the HP Switch Networking web site or can be copied from one switch to another. For downloading instructions, see appendix A, File Transfers, in the *Management and Configuration Guide* for your switch.

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 reconfigured 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 26 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)[1]
Manually add Candidate to Stack Prevent 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:

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 a Tagged VLAN" on "Stacking operation with multiple VLANs configured" (page 308).)
- No Manager password
- 14 or fewer stack members at the moment

Creating a stack (Overview)

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 53 Using the system name to help identify individual switches

   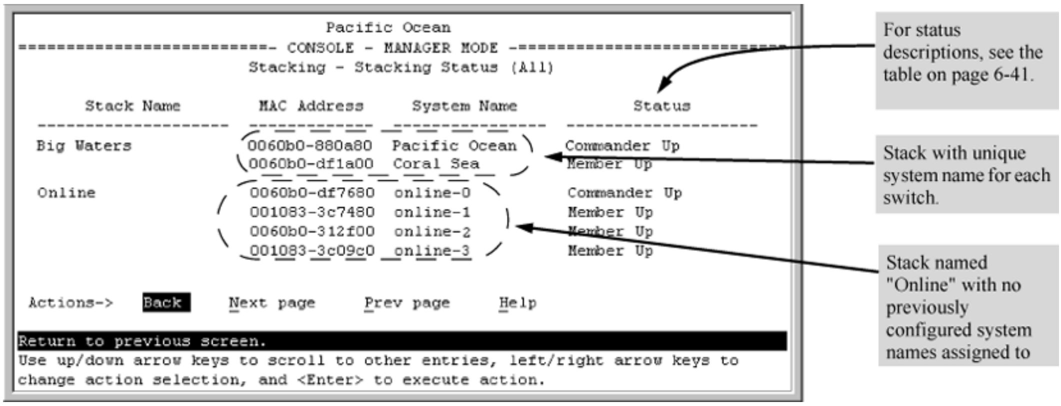

   ```
                                   Pacific Ocean
   ==========================- CONSOLE - MANAGER MODE -============================
                              Stacking - Stacking Status (All)

         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 to 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, the easiest way to avoid unintentionally adding a Candidate to the wrong stack is to 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, you can leave such switches in their default stacking configuration. If you need 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" (page 308).

6. Ensure that all switches intended for the stack are connected to the same subnet. As soon as you connect the Commander, it will begin discovering the available Candidates in the subnet.
   - If the Commander is configured to automatically add Members (`Auto Grab=Yes`), the first fifteen discovered Candidates meeting both of the following criteria will automatically join the stack:
     - `Auto Join` parameter set to `Yes` (the default)
     - Manager password not configured
   - 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.

Example

Example 161 Displaying `show stack` output

```
HP 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  HP Switch            Commander Up
 1 000883-08f980  Big_Waters-1  HP Switch            Member Up
```

Viewing the status of candidates the Commander has detected (CLI)

Syntax:

`show stack candidates`

Example

Example 162 Using the `show stack candidates` command to list candidates

```
HP Switch(config)# show stack candidates
 Stack Candidates

  Candidate MAC System Name            Device Type
  ------------- ---------------------- ----------------------
  0060b0-889e00 DEFAULT_CONFIG         3500yl
```

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.

Example

Example 163 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.

```
HP 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.

Example

Example 164 Using the `show stack view` command to list the stack assigned to the selected commander

```
HP Switch(config)# show stack view
Stack Members

 SN MAC Address    System Name    Device Type           Status
 -- -------------  -------------  --------------------  ------------
 0  1cc1de-cfbc80  Big_Waters-0   HP Switch 3500yl      Commander Up
 1  000883-08f980  Big Waters-1   HP Switch 3500yl      Member Up
```

Viewing stack status and configuring a Commander switch (Menu)

1. Configure an IP address and subnet mask on the Commander switch. (See the *Management and Configuration Guide* for your switch.)

2. Display the Stacking Menu by selecting `Stacking` in the Main Menu.

Figure 54 The default stacking menu

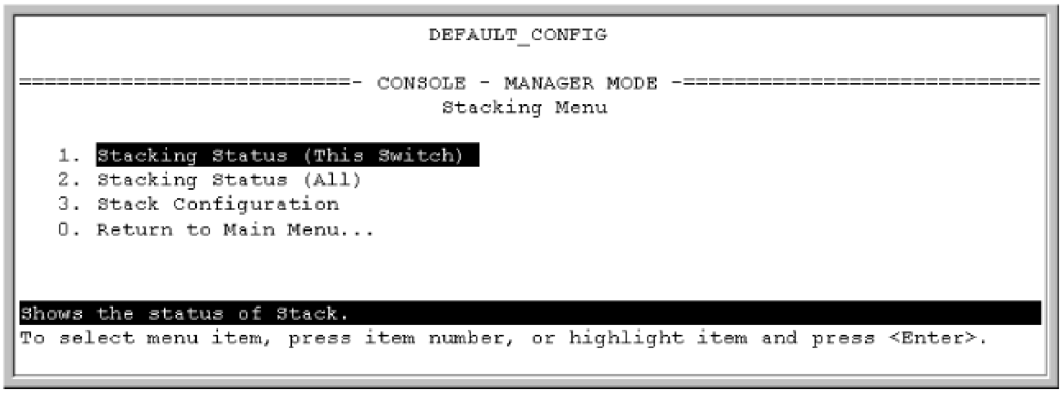

```
                          DEFAULT_CONFIG

==========================- CONSOLE - MANAGER MODE -============================
                           Stacking Menu

   1. Stacking Status (This Switch)
   2. Stacking Status (All)
   3. Stack Configuration
   0. Return to Main Menu...


Shows the status of Stack.
To select menu item, press item number, or highlight item and press <Enter>.
```

3. Display the Stack Configuration screen by pressing **3** to select `Stack Configuration`.

Figure 55 The default Stack Configuration screen

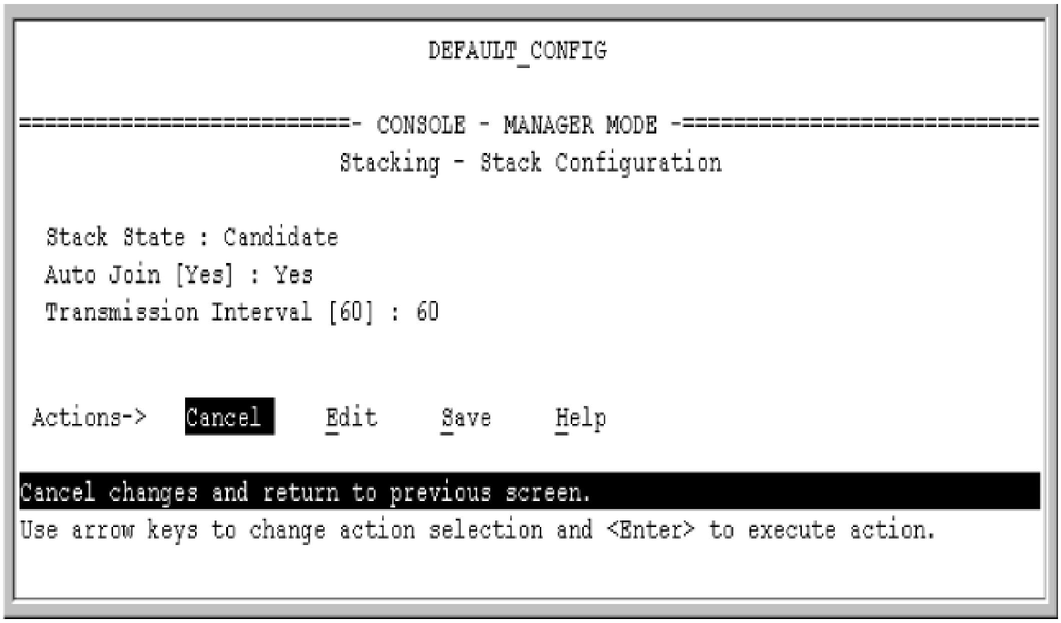

```
                          DEFAULT_CONFIG

==========================- CONSOLE - MANAGER MODE -============================
                  Stacking - Stack Configuration

  Stack State : Candidate
  Auto Join [Yes] : Yes
  Transmission Interval [60] : 60


 Actions->   Cancel    Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

4. Move the cursor to the Stack State field by pressing **E** (for Edit). Then use the Space bar to select the `Commander` option.
5. Press the down arrow key to display the Commander configuration fields in the Stack Configuration screen.

Figure 56 The default Commander configuration on the Stack Configuration screen

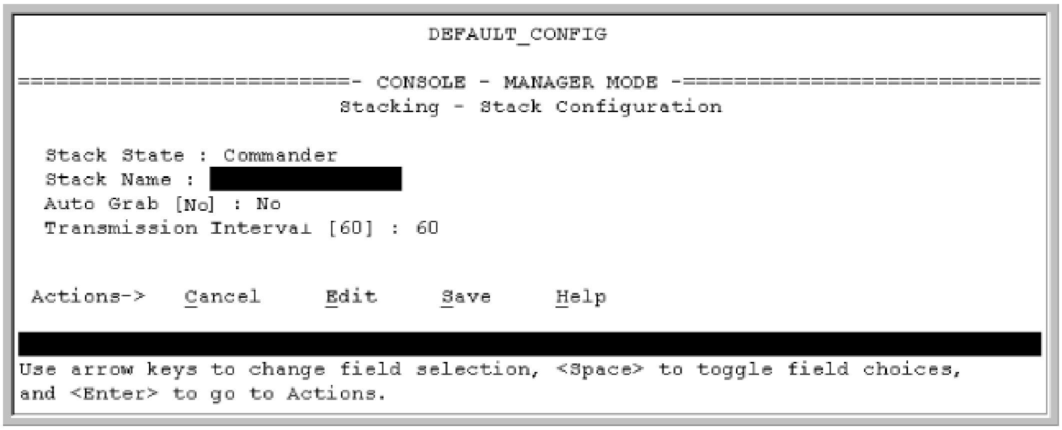

```
                              DEFAULT_CONFIG

==========================- CONSOLE - MANAGER MODE -============================
                          Stacking - Stack Configuration

  Stack State : Commander
  Stack Name : [redacted]
  Auto Grab [No] : No
  Transmission Interval [60] : 60


 Actions->    Cancel     Edit     Save     Help


Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

6. Enter a unique stack name (up to 15 characters; no spaces) and press the down arrow key.
7. Ensure that the Commander has the desired `Auto Grab` setting, then press the down arrow key:
   - No (the default) prevents automatic joining of Candidates that have their `Auto Join` set to `Yes`.
   - Yes enables the Commander to automatically take a Candidate into the stack as a Member if the Candidate has `Auto Join` set to `Yes` (the default Candidate setting) and does not have a previously configured password.
8. Accept or change the transmission interval (default: 60 seconds), then press **Enter** to return the cursor to the `Actions` line.
9. Press **S** (for Save) to save your configuration changes and return to the Stacking menu.

Your Commander switch should now be ready to automatically or manually acquire Member switches from the list of discovered Candidates, depending on your configuration choices.

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:

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 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 "The pimary VLAN" (page 58).

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 *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.

Example

Example 165 Creating a Command switch

To create a Commander switch with a stack name of `Big_Waters`:

```
HP 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.

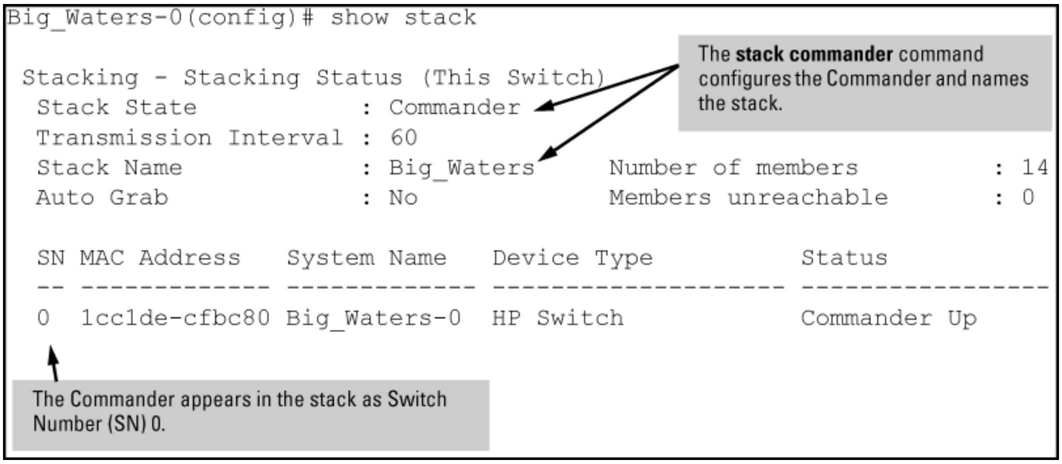

```
Big_Waters-0(config)# show stack

 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  HP 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 convert the Member to be the 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
stack commander stack name
```

Example

Example 166 Using a member's CLI to convert the member to the commander of a new stack

Suppose an HP switch named `Bering Sea` is a Member of a stack named `Big_Waters`. To use the switch's CLI to convert it from a stack Member to the Commander of a new stack named "Lakes", use the following commands:

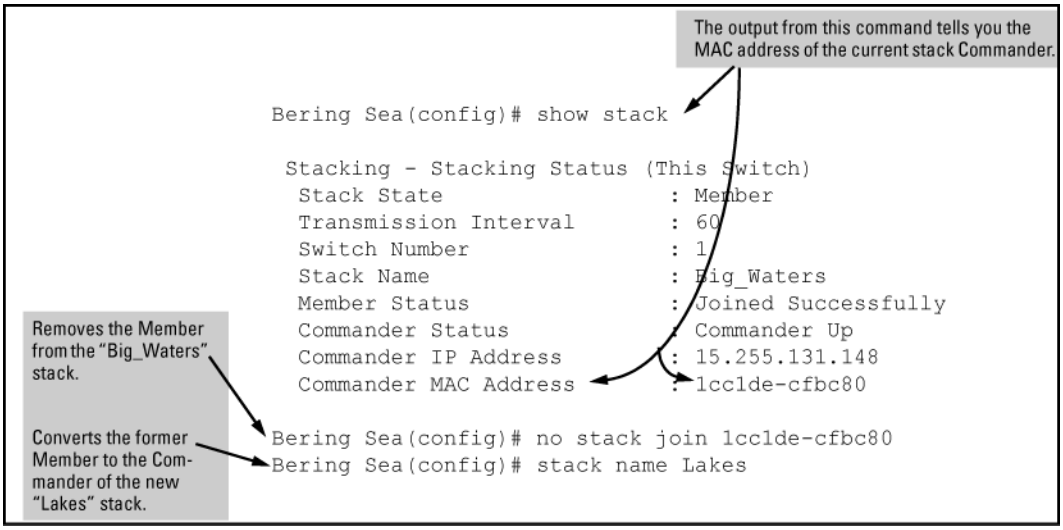

The output from this command tells you the MAC address of the current stack Commander.

```
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 this data, use the `show stack candidates` listing.

Example

Example 167 Determining available switch numbers (SNs)

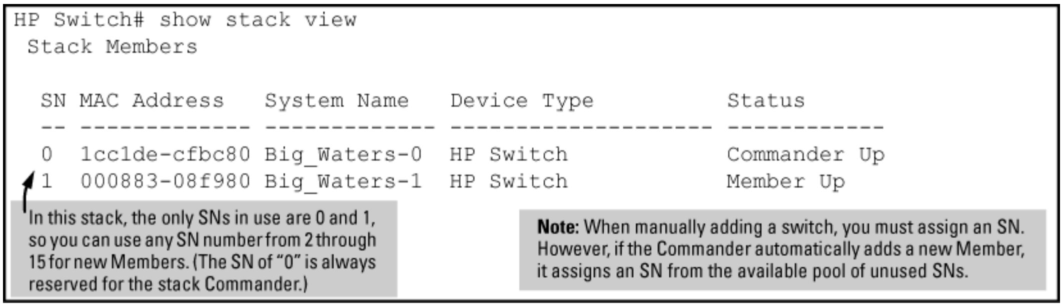

```
HP Switch# show stack view
 Stack Members

  SN MAC Address   System Name   Device Type          Status
  -- ------------- ------------- -------------------- ------------
  0  1cc1de-cfbc80 Big_Waters-0  HP Switch            Commander Up
  1  000883-08f980 Big_Waters-1  HP 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, execute the `show stack candidates` command from the Commander's CLI. For example, to list the discovered candidates for the above Commander:

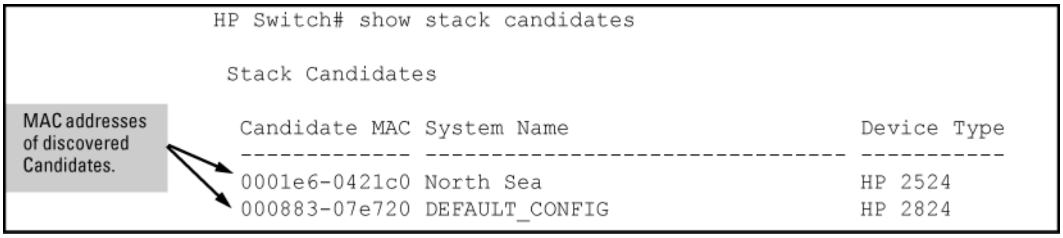

```
HP Switch# show stack candidates

 Stack Candidates

  Candidate MAC System Name                      Device Type
  ------------- -------------------------------- -----------
  0001e6-0421c0 North Sea                        HP 2524
  000883-07e720 DEFAULT_CONFIG                   HP 2824
```

MAC addresses of discovered Candidates.

Knowing the available switch numbers (`SN`s) and Candidate MAC addresses, proceed to 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 listing did not have a Manager password and you want to make it a stack Member with an `SN` of `2`, execute the following command:

```
HP Switch(config)# stack member 2 mac-address 0060b0-dfla00
```

The `show stack view` command then lists the Member added by the above command:

Example 168 Displaying the stack after adding a new member

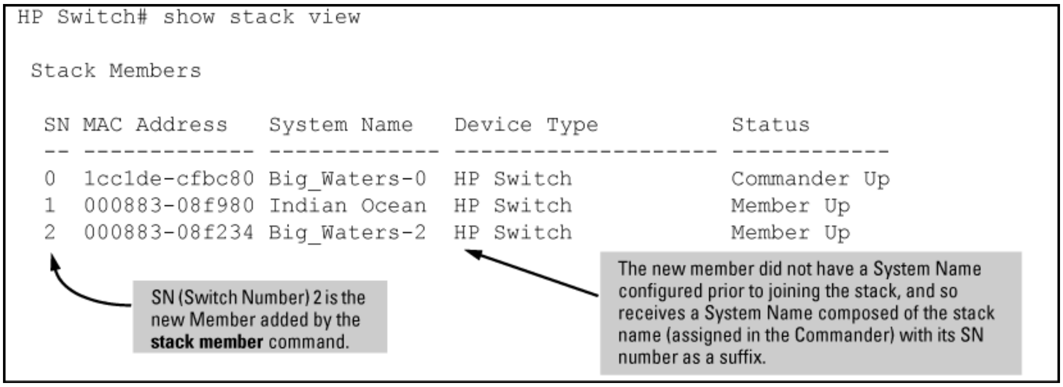

```
HP Switch# show stack view

 Stack Members

  SN MAC Address   System Name   Device Type          Status
  -- ------------- ------------- -------------------- ------------
  0  1cc1de-cfbc80 Big_Waters-0  HP Switch            Commander Up
  1  000883-08f980 Indian Ocean  HP Switch            Member Up
  2  000883-08f234 Big_Waters-2  HP 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.

Manually adding a Candidate to a stack (Menu)

In the default configuration, you must manually add stack Members from the Candidate pool. Reasons for a switch remaining a Candidate instead of becoming a Member include any of the following:

- `Auto Grab` in the Commander is set to `No` (the default).
- `Auto Join` in the Candidate is set to `No`.

NOTE: When a switch leaves a stack and returns to Candidate status, its `Auto Join` parameter resets to `No` so that it will not immediately rejoin a stack from which it has just departed.

- A Manager password is set in the Candidate.
- The stack is full.

Unless the stack is already full, use the Stack Management screen to manually convert a Candidate to a Member. If the Candidate has a Manager password, you will need to use it to make the Candidate a Member of the stack.

1. To add a Member, start at the Main Menu and select: **9. Stacking...** —> **4. Stack Management**

   You will then see the Stack Management screen. For status descriptions, see "Status messages" (page 308).

Example 169 The Stack Management screen

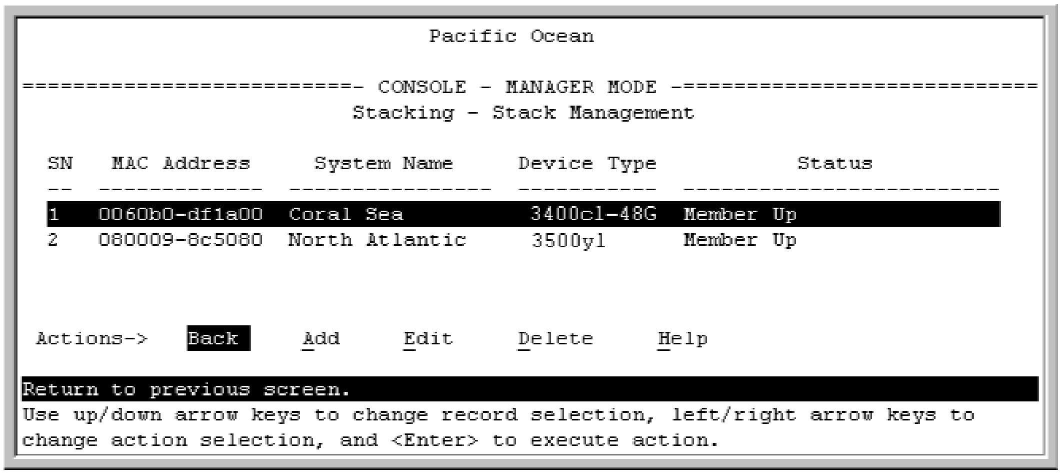

```
                              Pacific Ocean

==========================- CONSOLE - MANAGER MODE -============================
                        Stacking - Stack Management

  SN   MAC Address     System Name      Device Type          Status
  --  -------------  ---------------  -----------  -------------------------
  1   0060b0-df1a00  Coral Sea          3400cl-48G  Member Up
  2   080009-8c5080  North Atlantic     3500yl      Member Up



 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.
```

2. Press **A** (for Add) to add a Candidate. You will then see this screen listing the available Candidates:

Example 170 A Candidate list on the Stack Management screen

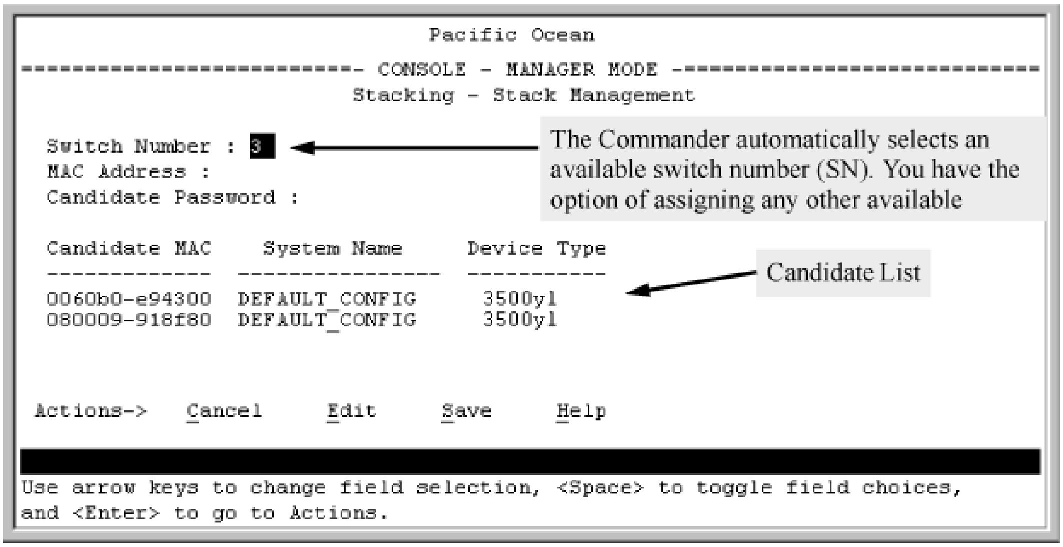

```
                              Pacific Ocean
==========================- CONSOLE - MANAGER MODE -============================
                        Stacking - Stack Management

  Switch Number : 3
  MAC Address :
  Candidate Password :

  Candidate MAC     System Name      Device Type
  -------------  ----------------   -----------
  0060b0-e94300  DEFAULT_CONFIG       3500yl
  080009-918f80  DEFAULT_CONFIG       3500yl


 Actions->    Cancel     Edit     Save     Help

Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

The Commander automatically selects an available switch number (SN). You have the option of assigning any other available

Candidate List

3. Either accept the displayed switch number or enter another available number. (The range is 0 - 15, with 0 reserved for the Commander.)
4. Use the down arrow key to move the cursor to the MAC Address field, then enter the MAC address of the desired Candidate from the Candidate list in the lower part of the screen.
5. Do one of the following:
    - If the desired Candidate has a Manager password, press the down arrow key to move the cursor to the Candidate Password field, then enter the password.
    - If the desired Candidate does not have a password, go to step 6.
6. Press **Enter** to return to the Actions line, then press **S** (for Save) to complete the Add process for the selected Candidate. You will then see a screen similar to the one in the next figure, with the newly added Member listed. For a description of the Status column entries, see "Status messages" (page 308).

NOTE: If the message `Unable to add stack member: Invalid Password` appears in the console menu's Help line, then you either omitted the Candidate's Manager password or incorrectly entered the Manager password.

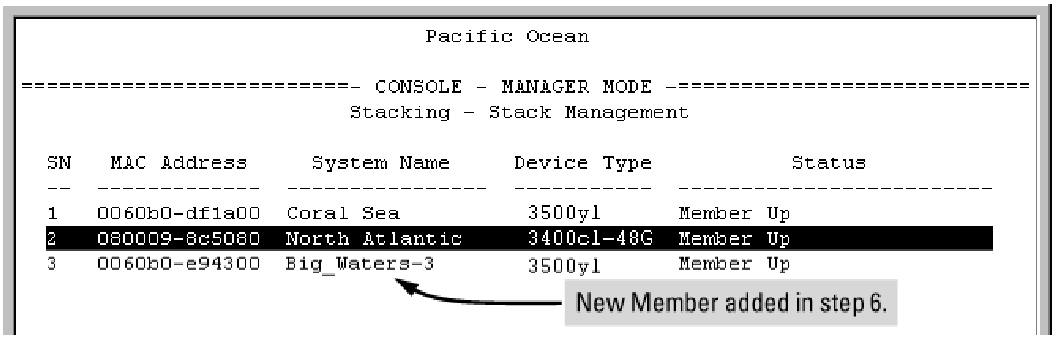

```
                              Pacific Ocean

==========================- CONSOLE - MANAGER MODE -============================
                        Stacking - Stack Management

  SN   MAC Address      System Name      Device Type           Status
  --  -------------  -----------------  -----------  -------------------------
  1   0060b0-df1a00  Coral Sea            3500yl     Member  Up
  2   080009-8c5080  North Atlantic       3400cl-48G Member  Up
  3   0060b0-e94300  Big_Waters-3         3500yl     Member  Up
```

New Member added in step 6.

Moving a Member from one stack to another (Menu)

Where two or more stacks exist in the same subnet, it is easy to move a Member of one stack to another stack if the destination stack is not full. This procedure is nearly identical to manually adding a Candidate to a stack. If the stack from which you want to move the Member has a Manager password, you will need to know it to make the move.

1. To move a Member from one stack to another, go to the Main Menu of the Commander in the destination stack and display the Stacking Menu by selecting **9. Stacking...**
2. To learn or verify the MAC address of the Member you want to move, display a listing of all Commanders, Members, and Candidates in the subnet by selecting **2. Stacking Status (All)**.

   You will then see the Stacking Status (All) screen. For a description of the Status column entries, see "Status messages" (page 308).

Figure 57 How the Stacking Status (All) screen helps you find member MAC addresses

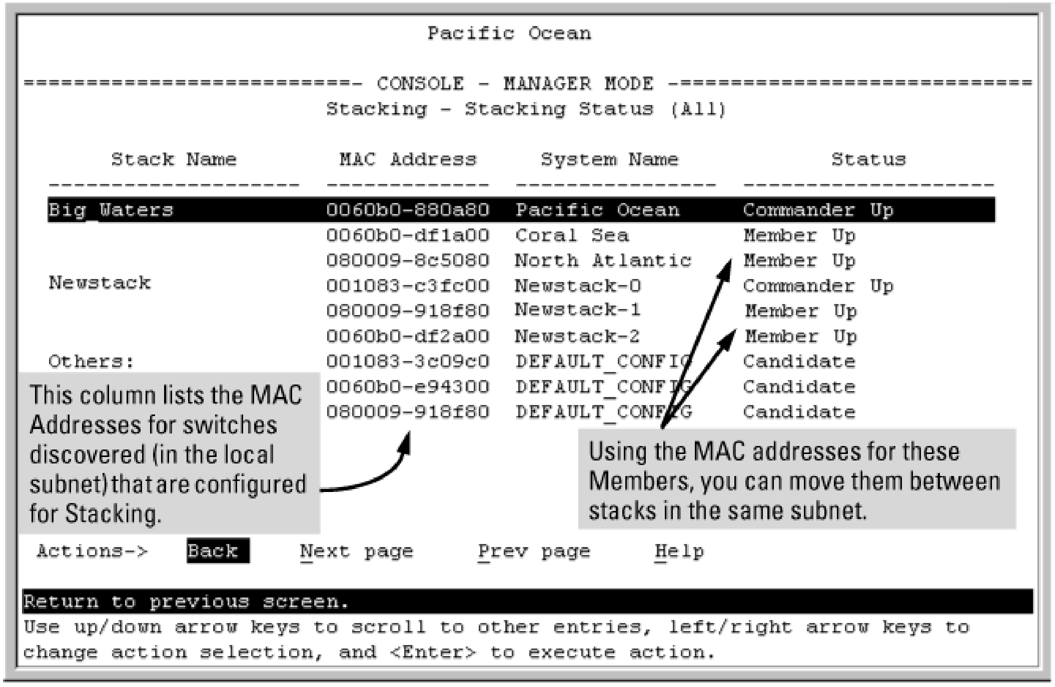

```
                           Pacific Ocean

=========================- 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
                       080009-8c5080  North Atlantic   Member Up
  Newstack             001083-c3fc00  Newstack-0       Commander Up
                       080009-918f80  Newstack-1       Member Up
                       0060b0-df2a00  Newstack-2       Member Up
  Others:              001083-3c09c0  DEFAULT_CONFIG   Candidate
                       0060b0-e94300  DEFAULT_CONFIG   Candidate
                       080009-918f80  DEFAULT_CONFIG   Candidate

 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.
```

This column lists the MAC Addresses for switches discovered (in the local subnet) that are configured for Stacking.

Using the MAC addresses for these Members, you can move them between stacks in the same subnet.

3. In the Stacking Status (All) screen, find the Member switch that you want to move and note its MAC address, then press **B** (for `Back`) to return to the Stacking Menu.
4. Display the Commander's Stack Management screen by selecting
   4. Stack Management

   (For an example of this screen, see Example 169 (page 289) and Example 172 (page 294).)
5. Press **A** (for Add) to add the Member. You will then see a screen listing any available candidates. (See Example 170 (page 290) and Example 172 (page 294).) You will not see the switch you want to add because it is a Member of another stack and not a Candidate.)
6. Either accept the displayed switch number or enter another available number. (The range is 0 - 15, with 0 reserved for the Commander.)
7. Use the down arrow key to move the cursor to the MAC Address field, then enter the MAC address of the desired Member you want to move from another stack.
8. Do one of the following:
   - If the stack containing the Member you are moving has a Manager password, press the down arrow key to select the Candidate Password field, then enter the password.
   - If the stack containing the Member you want to move does not have a password, go to step 9.

9. Press **Enter** to return to the Actions line, then press **S** (for Save) to complete the Add process for the selected Member. You will then see a screen similar to the one in Example 169 (page 289) and Example 172 (page 294), with the newly added Member listed.

   NOTE: If the message `Unable to add stack member: Invalid Password` appears in the console menu's Help line, then you either omitted the Manager password for the stack containing the Member or incorrectly entered the Manager password.

   You can "push" a Member from one stack to another by going to the Member's interface and entering the MAC address of the destination stack Commander in the Member's `Commander MAC Address` field. Using this method moves the Member to another stack without a need for knowing the Manager password in that stack, but also blocks access to the Member from the original Commander.

   If you are using VLANs in your stack environment, see "Stacking Operation with a Tagged VLAN" ("Status messages" (page 308)).

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:

```
[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.

Example

Example 171 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.

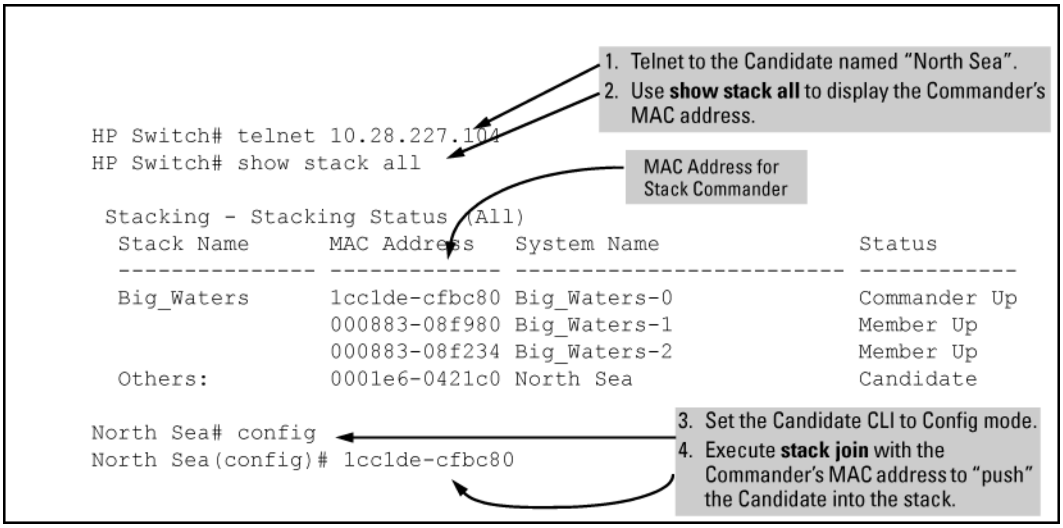

```
HP Switch# telnet 10.28.227.104
HP Switch# show stack all

 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.

Example

Example 172 Sstack listing with two stacks in the subnet

Suppose 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:

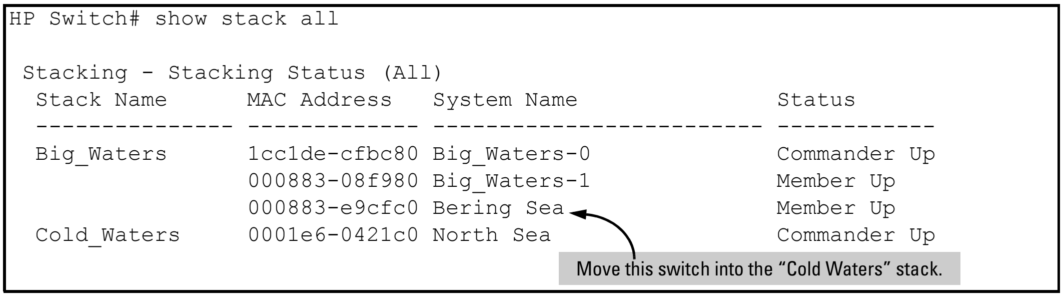

```
HP Switch# show stack all

 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:

```
HP 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.

To identify the MAC address of the destination Commander, use the `show stack all` command.

Example

Example 173 Converting a Commander to a Member

Suppose you have a switch operating as the Commander for a temporary stack named Test. When it is time 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:

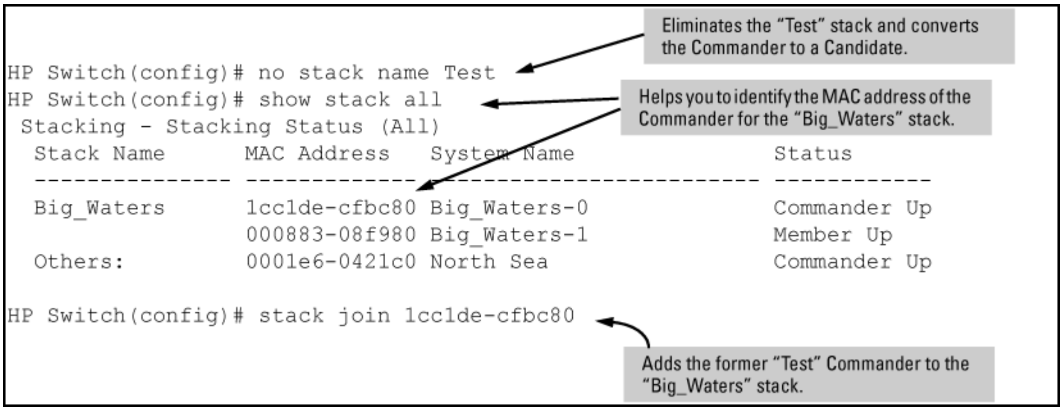

```
HP Switch(config)# no stack name Test
HP Switch(config)# show stack all
 Stacking - Stacking Status (All)
  Stack Name      MAC Address   System Name               Status
  --------------- ------------- ------------------------- ------------
  Big_Waters      1cc1de-cfbc80 Big_Waters-0              Commander Up
                  000883-08f980 Big_Waters-1              Member Up
  Others:         0001e6-0421c0 North Sea                 Commander Up

HP 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.

Converting a Commander or Member to a Member of another stack (Commander Menu)

When moving a Commander, the following procedure returns the stack members to Candidate status (with Auto-Join set to `No`), and converts the stack Commander to a Member of another stack. When moving a member, the procedure simply pulls a Member out of one stack and pushes it into another.

1. From the Main Menu of the switch you want to move, select **9. Stacking**.
2. To determine the MAC address of the destination Commander, select **2. Stacking Status (All)**.
3. Press **B** (for Back) to return to the Stacking Menu.
4. To display Stack Configuration menu for the switch you are moving, select **3. Stack Configuration**.
5. Press **E** (for Edit) to select the Stack State parameter.
6. Use the Space bar to select `Member`, then press **v** to move to the **Commander MAC Address** field.
7. Enter the MAC address of the destination Commander and press **Enter**.
8. Press **S** (for Save).

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.

Example 174 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:

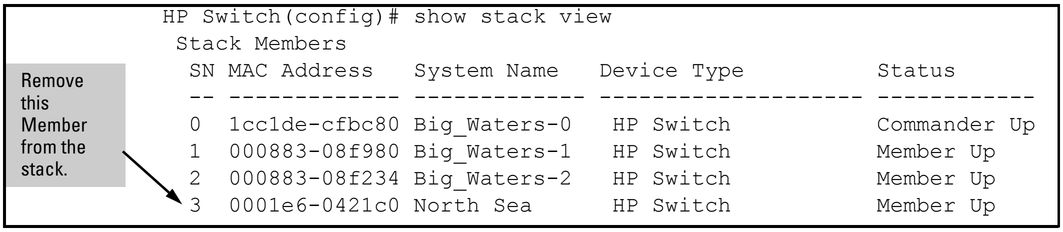

```
HP Switch(config)# show stack view
 Stack Members
  SN MAC Address   System Name   Device Type          Status
  -- ------------- ------------- -------------------- ------------
  0  1cc1de-cfbc80 Big_Waters-0   HP Switch           Commander Up
  1  000883-08f980 Big_Waters-1   HP Switch           Member Up
  2  000883-08f234 Big_Waters-2   HP Switch           Member Up
  3  0001e6-0421c0 North Sea      HP Switch           Member Up
```

Execute this command to remove the `North Sea` switch from the stack:

```
HP 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.

Example

Example 175 Identifying the commander's MAC address from a member switch

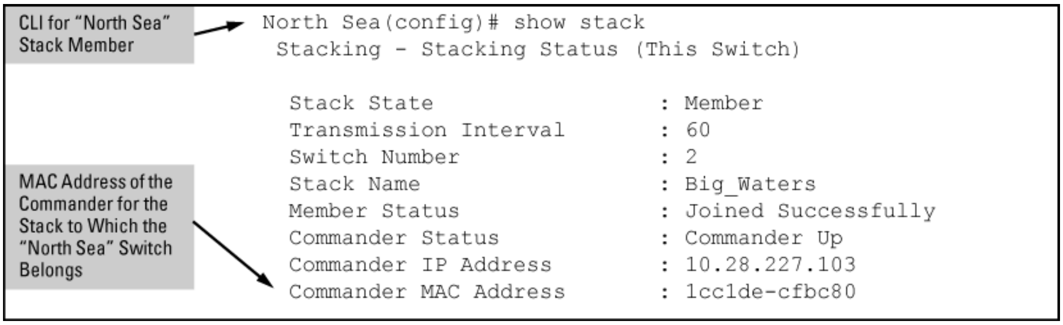

CLI for "North Sea" Stack Member

MAC Address of the Commander for the Stack to Which the "North Sea" Switch Belongs

```
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
```

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
```

Removing a stack Member (Menu)

These rules affect removals from a stack:

- When a Candidate becomes a Member, its `Auto Join` parameter is automatically set to `No`. This prevents the switch from automatically rejoining a stack as soon as you remove it from the stack.
- When you use the Commander to remove a switch from a stack, the switch rejoins the Candidate pool for your IP subnet (broadcast domain), with `Auto Join` set to `No`.
- When you remove a Member from a stack, it frees the previously assigned switch number (`SN`), which then becomes available for assignment to another switch that you may subsequently add to the stack. The default switch number used for an add is the lowest unassigned number in the Member range (1 - 15; 0 is reserved for the Commander).

To remove a Member from a stack, use the Stack Management screen.

1. From the Main Menu, select **9. Stacking...** —> **4. Stack Management**

   You will then see the Stack Management screen. See "Status messages" (page 308).

   Figure 58 The stack management screen with stack members listed

   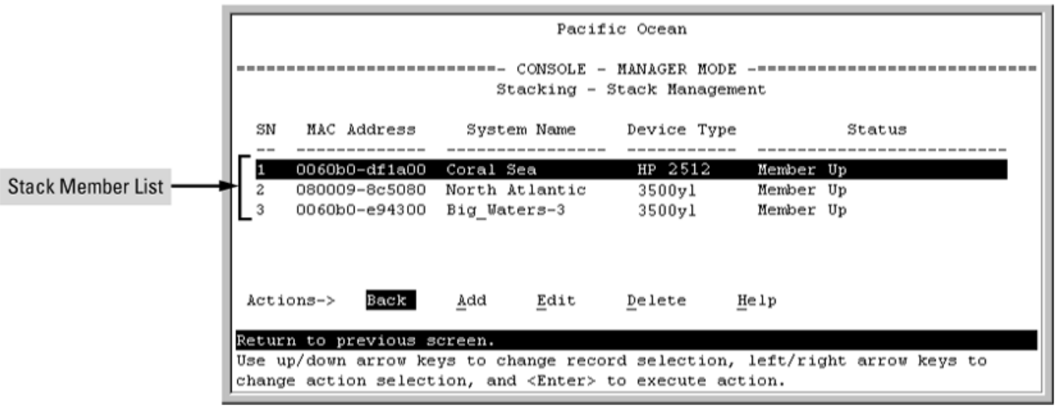

   Stack Member List

   ```
                                Pacific Ocean

   ==========================- CONSOLE - MANAGER MODE -===========================
                         Stacking - Stack Management

     SN   MAC Address     System Name       Device Type          Status
     --  -------------  ----------------  -----------  -------------------------
     1   0060b0-df1a00  Coral Sea          HP 2512     Member Up
     2   080009-8c5080  North Atlantic     3500yl      Member Up
     3   0060b0-e94300  Big_Waters-3       3500yl      Member Up



    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.
   ```

2. Use the down arrow key to select the Member to remove from the stack.
3. Type **D** (for Delete) to remove the selected Member from the stack.

   You will see the following prompt:

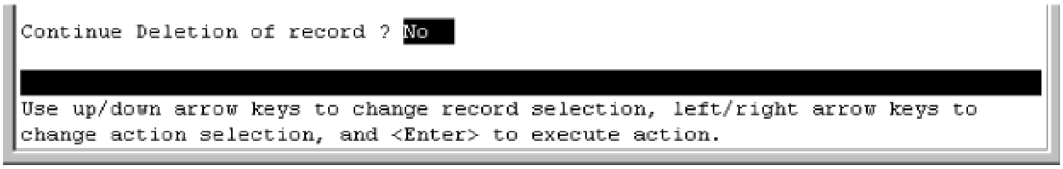

```
Continue Deletion of record ? No

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

4. To continue deleting the selected Member, press the Space bar once to select **Yes** for the prompt, then press **Enter** to complete the deletion. The Stack Management screen updates to show the new stack Member list.

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.

Example

Example 176 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:

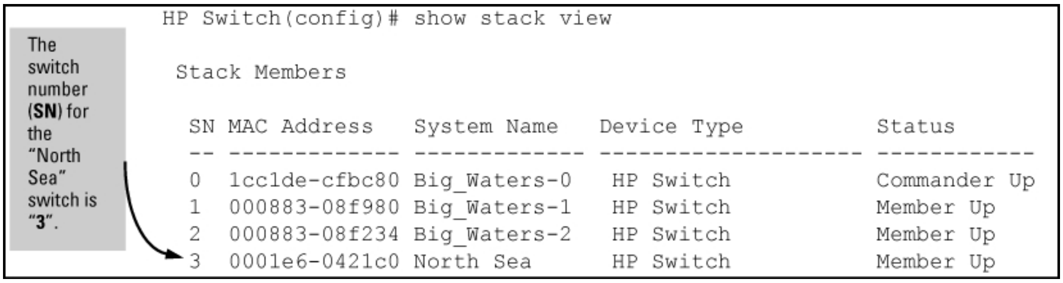

The switch number (**SN**) for the "North Sea" switch is "3".

```
HP Switch(config)# show stack view

 Stack Members

  SN MAC Address   System Name   Device Type          Status
  -- ------------- ------------- -------------------- ------------
  0  1cc1de-cfbc80 Big_Waters-0   HP Switch           Commander Up
  1  000883-08f980 Big_Waters-1   HP Switch           Member Up
  2  000883-08f234 Big_Waters-2   HP Switch           Member Up
  3  0001e6-0421c0 North Sea      HP Switch           Member Up
```

To access the `North Sea` console, execute the following Telnet command:

```
HP 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.

Accessing Member switches for configuring changes and monitoring traffic (Commander Menu)

After a Candidate becomes a stack Member, you can use that stack's Commander to access the Member's console interface for the same configuration and monitoring that you would do through a Telnet or direct-connect access.

1. From the Main Menu, select **9. Stacking...** —> **5. Stack Access**

   You will then see the Stack Access screen. See "Status messages" (page 308).

   Figure 59 The Stack Access screen

   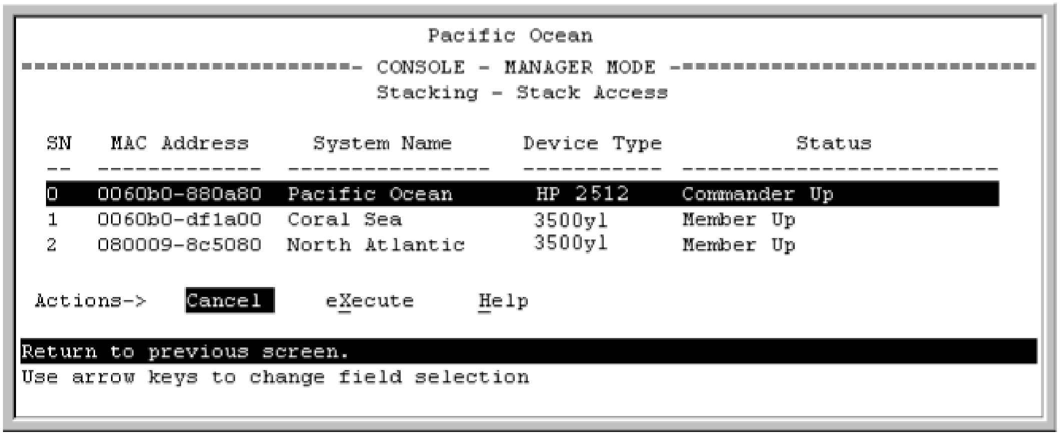

```
                               Pacific Ocean
=========================- CONSOLE - MANAGER MODE -==========================
                           Stacking - Stack Access

  SN   MAC Address      System Name      Device Type          Status
  --   -------------   ----------------  -----------  -------------------------
  0    0060b0-880a80   Pacific Ocean       HP 2512    Commander Up
  1    0060b0-df1a00   Coral Sea          3500yl      Member Up
  2    080009-8c5080   North Atlantic     3500yl      Member Up

 Actions->    Cancel      eXecute     Help

Return to previous screen.
Use arrow keys to change field selection
```

   Use the down arrow key to select the stack Member you want to access, then press **X** (for eXecute) to display the console interface for the selected Member. For example, if you selected switch number 1 (system name: `Coral Sea`) in Figure 59 (page 299) and then pressed **X**, you would see the Main Menu for the switch named `Coral Sea`.

   Figure 60 The eXecute command displaying the console Main Menu for the selected stack Member

   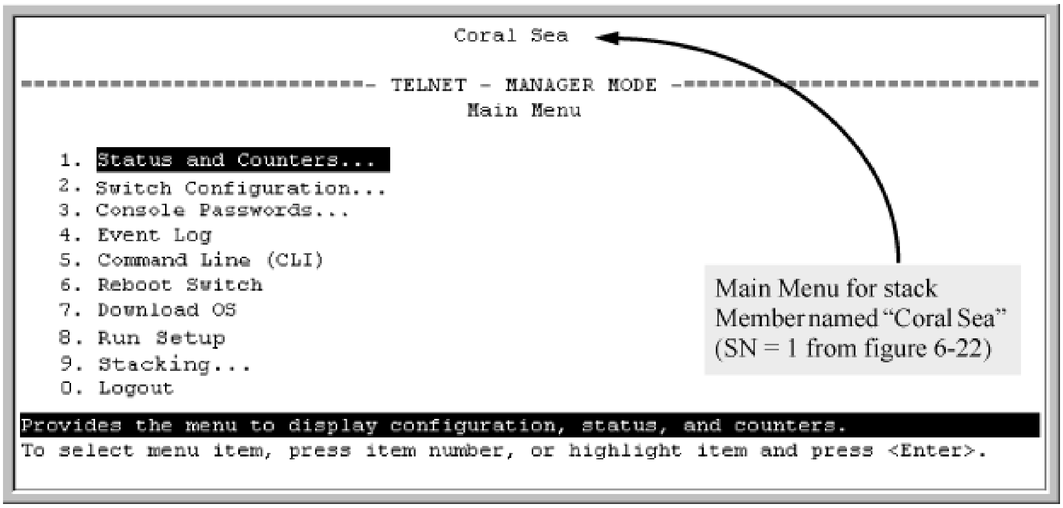

```
                                 Coral Sea

==========================- TELNET - MANAGER MODE -==========================
                                 Main Menu

    1. Status and Counters...
    2. Switch Configuration...
    3. Console Passwords...
    4. Event Log
    5. Command Line (CLI)
    6. Reboot Switch
    7. Download OS
    8. Run Setup
    9. Stacking...
    0. Logout

Provides the menu to display configuration, status, and counters.
To select menu item, press item number, or highlight item and press <Enter>.
```

   Main Menu for stack Member named "Coral Sea" (SN = 1 from figure 6-22)

2. You can now make configuration changes and view status data for the selected Member in the same way t you would if you were directly connected or telnetted into the switch.
3. When you finish accessing the selected Member, do the following to return to the Commander's Stack Access screen:
   a. Return to the Member's Main Menu.
   b. Press **0** (for Logout), then **Y** (for Yes).
   c. Press **Return**.

   You should now see the Commander's Stack Access screen. (For an example, see Figure 59 (page 299) on Figure 63 (page 303).)

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. HP recommends that you leave this parameter set to the default 60 seconds.

Syntax:

```
stack transmission-interval seconds
```

Managing a Candidate switch (Menu)

Perform these actions on a Candidate switch:

- Add (push) the Candidate into an existing stack.
- Modify the Candidate's stacking configuration (`Auto Join` and `Transmission Interval`).
- Convert the Candidate to a Commander.
- Disable stacking on the Candidate so that it operates as a standalone switch.

In its default stacking configuration, a Candidate switch can either automatically join a stack or be manually added (pulled) into a stack by a Commander, depending on the Commander's `Auto Grab` setting. The following table lists the Candidate's configuration options:

Table 27 Candidate configuration options in the menu interface

Parameter	Default setting	Other settings
`Stack State`	Candidate	Commander, Member, or Disabled
`Auto Join`	Yes	No
`Transmission Interval`	60 Seconds	Range: 1 to 300 seconds

Pushing a switch into a stack, modifying the switch's configuration, or disabling stacking on the switch (Menu)

Use Telnet or the WebAgent to access the Candidate if it has an IP address. Otherwise, use a direct connection from a terminal device to the switch's console port.

1. Display the Stacking Menu by selecting `Stacking` in the console Main Menu.
2. Display the Stack Configuration menu by pressing **3** to select `Stack Configuration`.

Figure 61 The default stack configuration screen

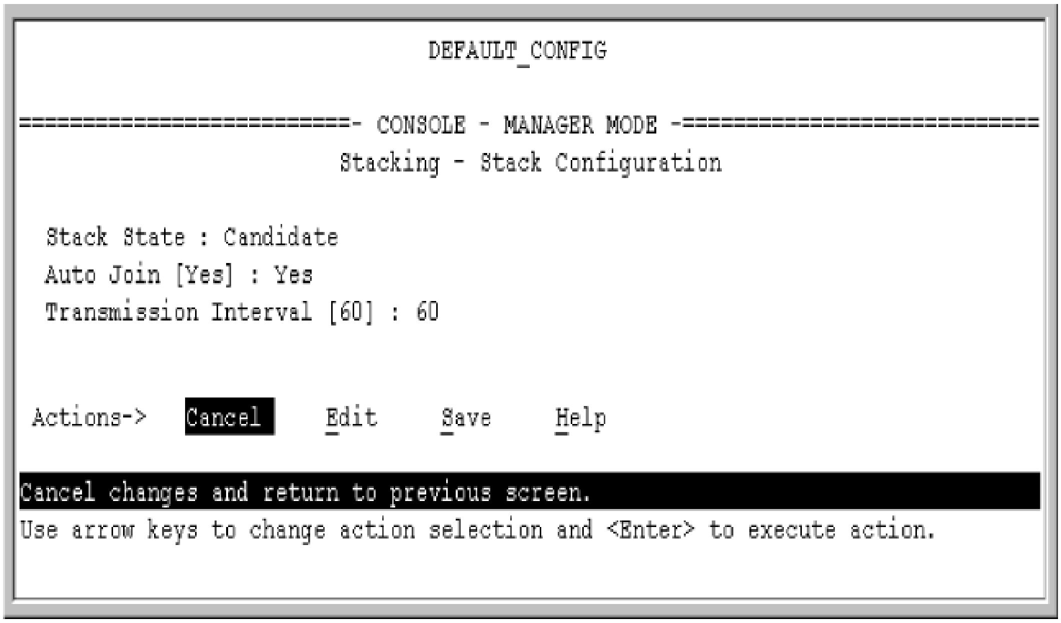

```
                          DEFAULT_CONFIG

==========================- CONSOLE - MANAGER MODE -============================
                    Stacking - Stack Configuration

  Stack State : Candidate
  Auto Join [Yes] : Yes
  Transmission Interval [60] : 60


 Actions->   Cancel     Edit     Save     Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

3. Move the cursor to the Stack State field by pressing **E** (for Edit).
4. Do one of the following:
   - To disable stacking on the Candidate, use the Space bar to select the `Disabled` option, then go to step 5.

     NOTE: Using the menu interface to disable stacking on a Candidate removes the Candidate from all stacking menus.

   - To insert the Candidate into a specific Commander's stack:
     a. Use the space bar to select Member.
     b. Press **Tab** once to display the `Commander MAC Address` parameter, then enter the MAC address of the desired Commander.
   - To change `Auto Join` or `Transmission Interval`, use **Tab** to select the desired parameter, and:
     - To change `Auto Join`, use the Space bar.
     - To change `Transmission Interval`, enter the new value in the range of 1 to 300 seconds.

       NOTE: All switches in the stack must be set to the same transmission interval to help ensure proper stacking operation. HP recommends that you leave this parameter set to the default 60 seconds.

   Then go to Step 5.
5. Press **Enter** to return the cursor to the `Actions` line.

6. Press **S** (for Save) to save your configuration changes and return to the Stacking menu.

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.

Monitoring stack status (Menu)

Using the stacking options in the menu interface for any switch in a stack, you can view stacking data for that switch or for all stacks in the subnet. This helps in determining the stacking configuration for individual switches, identifying stack Members and Candidates, and determining the status of individual switches in a stack.

Table 28 Stack status environments

Screen name	Commander	Member	Candidate
Stack Status (This Switch)	• Commander's stacking configuration • Data on stack Members: • `Switch Number` • `MAC Address` • `System Name` • `Device Type` • `Status`	• Member's stacking configuration • Member Status • Data identifying Member's Commander: • Commander Status • Commander IP Address • Commander MAC Address	Candidate's stacking configuration
Stack Status (All)	Lists devices by stack name or Candidate status (if device is not a stack Member). Includes: • Stack Name • MAC Address • System Name • Status	Same as for Commander.	Same as for Commander.

If you are using VLANs in your stack environment, see "Stacking operation with multiple VLANs configured" (page 308).

Using a stacked switch to view status for all switches with stacking enabled (Menu)

This procedure displays the general status of all switches in the IP subnet that have stacking enabled.

Go to the console Main Menu for any switch configured for stacking and select **9. Stacking ...** —> **2. Stacking Status (All)**.

You will then see a Stacking Status screen similar to the following. For a description of the Status column entries, see "Status messages" (page 308).

Figure 62 Stacking status for all detected switches configured for stacking

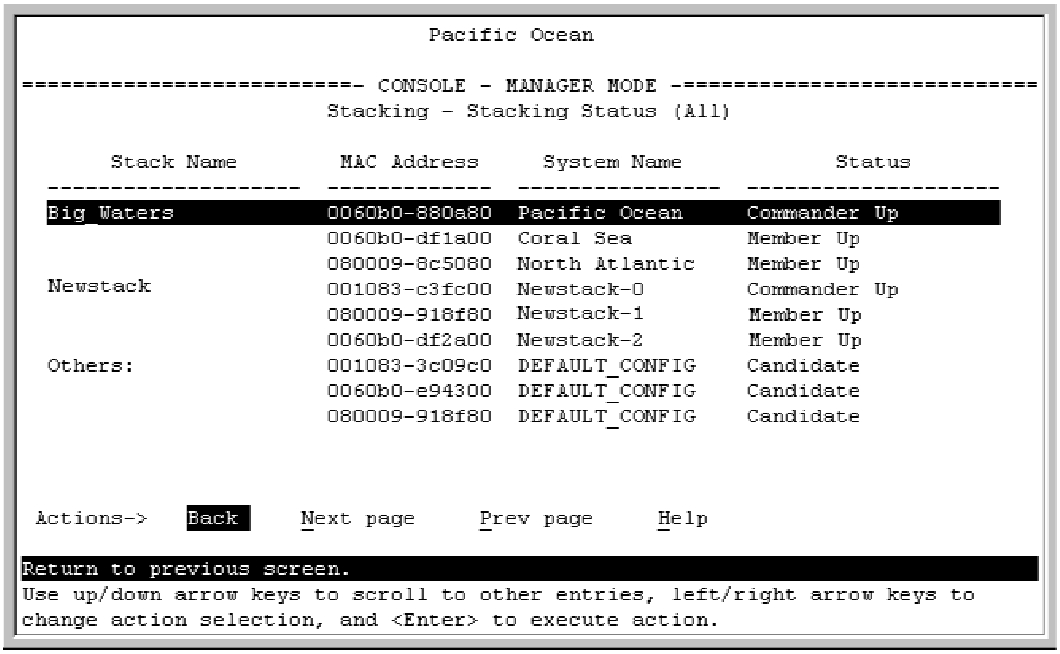

```
                                Pacific Ocean

==========================- 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
                        080009-8c5080  North Atlantic    Member Up
  Newstack              001083-c3fc00  Newstack-0        Commander Up
                        080009-918f80  Newstack-1        Member Up
                        0060b0-df2a00  Newstack-2        Member Up
  Others:               001083-3c09c0  DEFAULT_CONFIG    Candidate
                        0060b0-e94300  DEFAULT_CONFIG    Candidate
                        080009-918f80  DEFAULT_CONFIG    Candidate



 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.
```

Viewing Commander status (Menu)

This procedure displays the Commander and stack configuration, plus information identifying each stack member.

To display the status for a Commander, go to the console Main Menu for the switch and select **9. Stacking ...** —> **1. Stacking Status (This Switch)**.

You will then see the Commander's Stacking Status screen:

Figure 63 The Commander’s Stacking Status screen

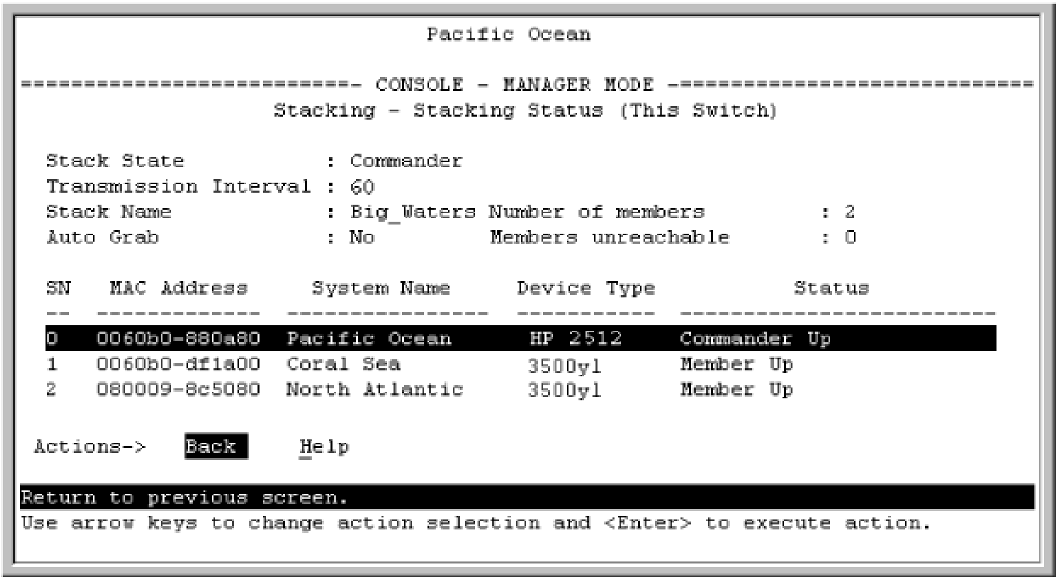

```
                                Pacific Ocean

==========================- CONSOLE - MANAGER MODE -============================
                   Stacking - Stacking Status (This Switch)

  Stack State            : Commander
  Transmission Interval  : 60
  Stack Name             : Big_Waters Number of members         : 2
  Auto Grab              : No         Members unreachable       : 0

  SN   MAC Address     System Name      Device Type          Status
  --  -------------  ---------------- -----------  -------------------------
  0   0060b0-880a80  Pacific Ocean     HP 2512     Commander Up
  1   0060b0-df1a00  Coral Sea         3500yl      Member Up
  2   080009-8c5080  North Atlantic    3500yl      Member Up

 Actions->   Back     Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Viewing Member status, and a Commander’s IP and MAC addresses and status (Menu)

This procedure displays the Member's stacking information plus the Commander's status, IP address, and MAC address.

To display the status for a Member:

1. Go to the console Main Menu of the Commander switch and select **9. Stacking ...** —> **5. Stack Access**.
2. Use the downarrow key to select the Member switch whose status you want to view, then press **X** (for eXecute). You will then see the Main Menu for the selected Member switch.
3. In the Member's Main Menu screen, select **9. Stacking ...** —> **1. Stacking Status (This Switch)**.

   You will see the Member's Stacking Status screen:

Figure 64 A Member's stacking status screen

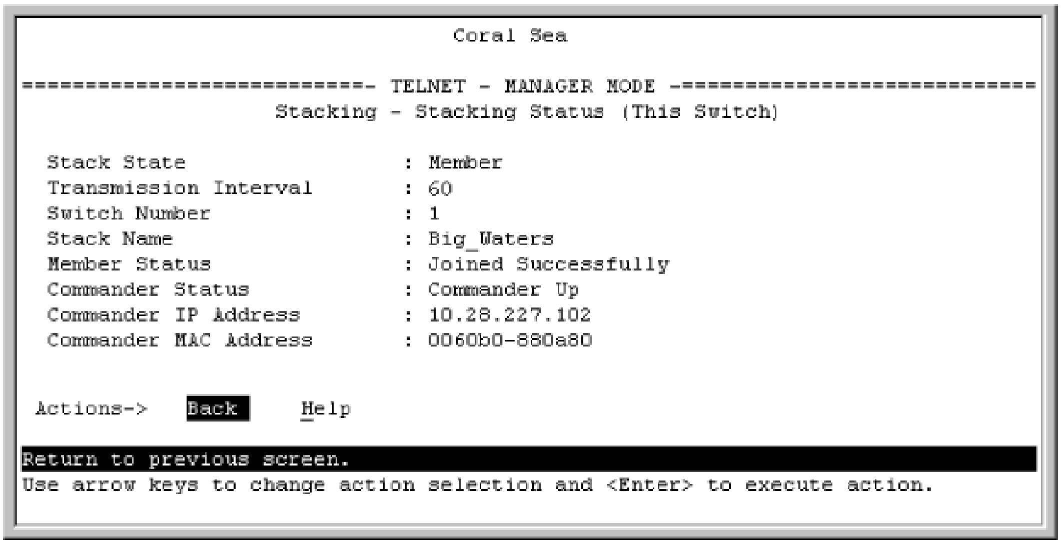

```
                                   Coral Sea

===========================- TELNET - MANAGER MODE -============================
                      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         : 10.28.227.102
  Commander MAC Address        : 0060b0-880a80


 Actions->   Back     Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

Viewing Candidate status (Menu)

This procedure displays the Candidate's stacking configuration.

To display the status for a Candidate:

Use Telnet (if the Candidate has a valid IP address for your network) or a direct serial port connection to access the menu interface Main Menu for the Candidate switch and select **9. Stacking ...** —> **1. Stacking Status (This Switch)**.

You will then see the Candidate's Stacking Status screen:

Figure 65 A Candidate's Stacking screen

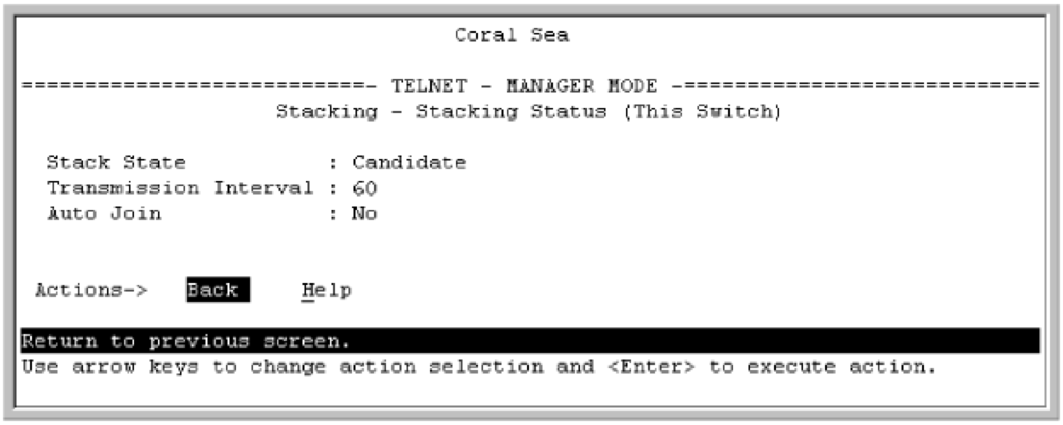

```
                                   Coral Sea

===========================- TELNET - MANAGER MODE -============================
                      Stacking - Stacking Status (This Switch)

  Stack State           : Candidate
  Transmission Interval : 60
  Auto Join             : No


 Actions->   Back     Help

Return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
```

About stack management

Components of HP Switch stack management

Table 29 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 66 A switch moving from Candidate to Member

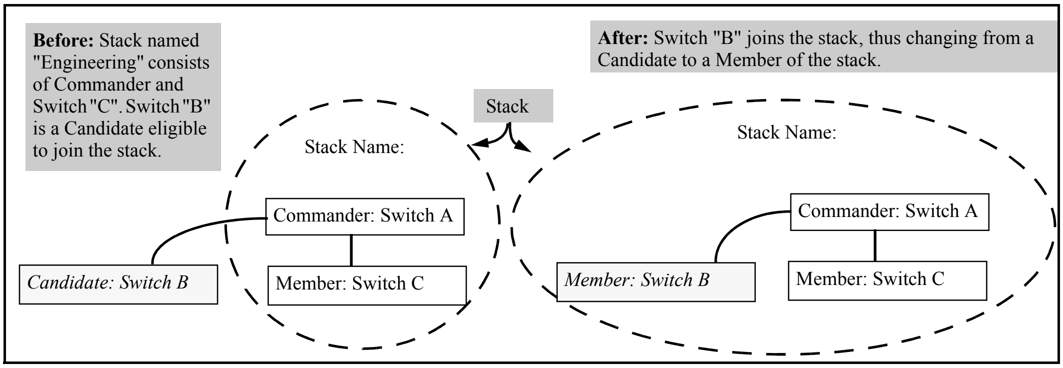
Before: Stack named "Engineering" consists of Commander and Switch "C". Switch "B" is a Candidate eligible to join the stack.
After: Switch "B" joins the stack, thus changing from a Candidate to a Member of the stack.
Stack
Stack Name:
Commander: Switch A
Member: Switch C
Candidate: Switch B
Stack Name:
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 67 Stacking with one commander controlling access to wiring closet switches

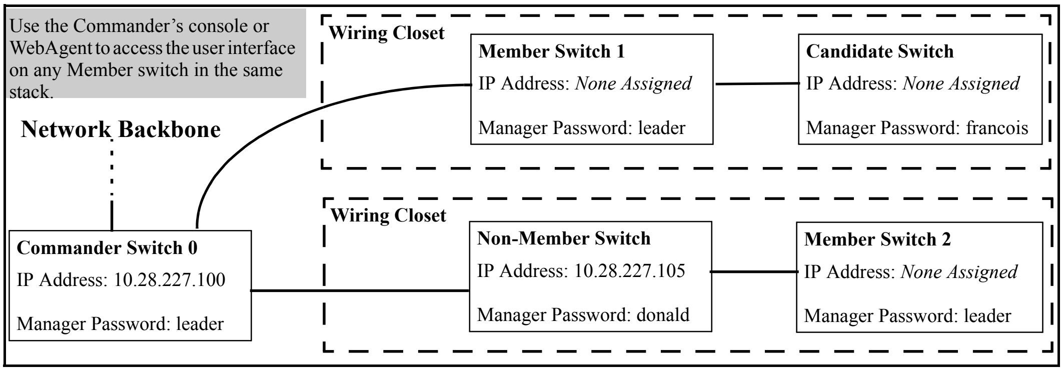
Use the Commander's console or WebAgent to access the user interface on any Member switch in the same stack.
Network Backbone
Wiring Closet
Member Switch 1
IP Address: None Assigned
Manager Password: leader
Candidate Switch
IP Address: None Assigned
Manager Password: francois
Commander Switch 0
IP Address: 10.28.227.100
Manager Password: leader
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" (page 308) and "The pimary VLAN" (page 58).)
- Stacking allows intermediate devices that do not support stacking. This enables you to include switches that are distant from the Commander.

Figure 68 A non-stacking device used in a stacking environment

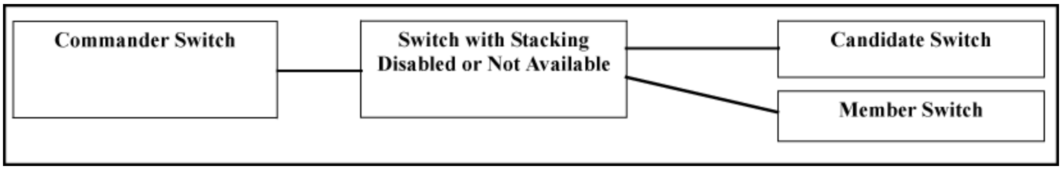

Specific rules

Table 30 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 fails and you need to convert a Member switch to operate as a replacement Commander. `Stack Name:` N/A	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 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 the same SNMP communities as the Commander (which serves as an SNMP proxy to the Member for communities to which the Commander 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. See "Components of HP Switch stack management" (page 305).

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, HP 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 pimary VLAN" (page 58).)

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, reconfigure 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 69 An SNMP community operation with stacking

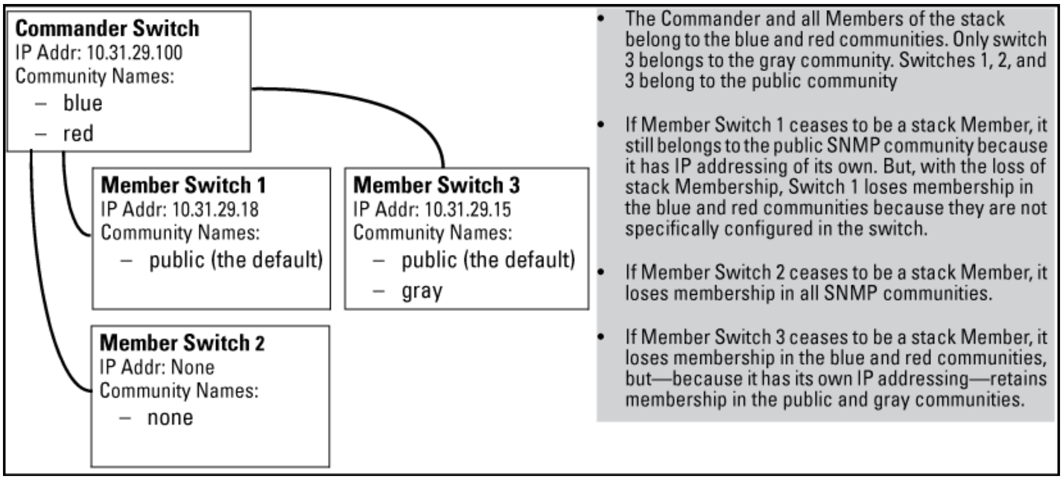

- The Commander and all Members of the stack belong to the blue and red communities. Only switch 3 belongs to the gray community. Switches 1, 2, and 3 belong to the public community
- If Member Switch 1 ceases to be a stack Member, it still belongs to the public SNMP community because it has IP addressing 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 Figure 69 (page 309), 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 above example (Figure 69 (page 309)) 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
```

8 Stack management for the 3800 switches

Command syntax	Description	Default value	CLI page reference
`stacking enable`	Enables stacking on the switch	Enabled	279
`stacking set-stack`	Generates a stack ID		314
`stacking member` *N* `pri` *X* `type` *JxxxxA* [`mac` *MAC-Addr*]	Adds a switch to the stack		314
`stacking member` *N* `remove`	Removes a switch from the stack		315
`redundancy switchover`	Forces Commander status over to the Standby switch		315
`stack member 2 type` *type* `mac-address` *MAC address*	Clears the MAC address of a switch member to allow for renumbering of the stack		315
`member-context` *stack-member*	After downloading new software to the stack, sets the CLI context so that subsequent commands apply to the stack member that is specified		318
`show stacking`	Shows a summary of the stack status		318
`show stacking detail`	Shows detailed stack status		320
`show stacking member` *stack-member-list*	Shows detailed information about switches in the *stack-member-list* only		320
`show stacking stack-ports member` *stack-member*	Shows the current state of the stacking ports of the specified member, or for all physically		321

Command syntax	Description	Default value	CLI page reference
	present devices in the stack		

Introduction

This feature is available on the HP 3800 switches only. See the *HP 3800 Switch Installation and Getting Started Guide* for information on supported stacking topologies.

NOTE: This feature is different from the stacking feature that is implemented on some other HP Networking switches. The other feature is implemented via the front-panel networking cables, and it does not have the high bandwidth and redundancy features of the HP 3800 stacking.

The stacking feature for the HP 3800 switches allows you to connect up to ten switches and have them act as a single high-bandwidth switch for both data and management.

One switch in the stack is designated as “Commander” and one switch is elected to be the “Standby.” The other switches are designated “Member.” The Commander is responsible for the overall management of the stack. The Standby provides redundancy for the stack and takes over stack management operations should the Commander fail, or if a Commander failover is forced by an administrator. The Members are not part of the overall stack management, however, they must manage their local subsystems and ports to operate correctly as part of the stack. The Commander and Standby are also responsible for their own local subsystems and ports.

HP Switch Stack 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 for the HP 3800 switches enables you to:

- Simplify management of small workgroups or wiring closets while scaling your network to handle increased bandwidth demand.
- Add switches to your network without having to first perform IP addressing tasks.
- Reduce the number of IP addresses needed in your network.
- Reduce downtime with high availability in the event of a failure.

NOTE: In the default configuration, stacking is enabled on HP 3800 switches. However, if an HP 3800 switch is powered on and it does not have a Stacking Module installed, stacking is disabled. If a Stacking Module is subsequently installed in the switch, stacking must be enabled from the switch CLI (in the config context) by entering the following command:

```
HP Switch(config)# stacking enable
```

Configuring a stack

Creating a stack

Ways to create a stack include:

- By the sequence in which the switches are booted. You choose which member becomes Commander.
- Plug-and-go method. Ensure that stacking is enabled on all the switches, and then connect them together in the desired stacking topology. The plug-and-go method lets stacking decide which member is the Commander.

Using a deterministic method

1. Install a Stacking Module into an HP 3800 switch and then boot the switch, as described in the *HP 3800 Switch Installation and Getting Started Guide*.
2. Make sure that stacking is enabled for the switch:
   a. Enter the `show stacking` command.
   b. If stacking is disabled, enter `stacking enable` (in global config context). This command causes the switch to reboot.
3. When the switch finishes booting, enter the `show stacking` command again. The switch now has the status of Commander. It has a Member ID of 1 (one) and a default priority of 128.

Example 177 Displaying `show stacking` output

```
HP Stack 3800# show stacking
Stack ID : NO ID - will merge upon connectivity
MAC Address : 1cc1de-4d48c0
Stack Topology : No Stack Formed
Stack Status : No Stack Formed
Uptime : 0d 0h 5m
Software Version : KA.15.03
Mbr
ID Mac Address Model Pri Status
--- ------------- ------------------------------------- --- -------
1 1cc1de-4d48c0 HP J9574A 3800-48G-PoE+-4SFP+ Switch 128 Commander
```

4. To generate a stack ID, enter the following command:

   ```
   HP Switch(config)#stacking set-stack
   ```

5. (Optional) To have this switch retain its function as Commander through stack boots and other situations, you can increase its priority. The switch with the highest priority becomes Commander when all the switches are booted simultaneously. The default priority is 128. The priority can be set to any value between 1 and 255. To increase the switch's stacking priority, enter the following command:

   ```
   HP Switch(config)# stacking member 1 priority 255
   ```

6. (Optional) Pre-configure (provision) the stack for the other switches that become members of the stack. You can assign a member number and a priority by entering the following command for each switch:

   ```
   HP Switch(config)# stacking member N pri X type JxxxxA [mac
   MAC-Addr]
   ```

   where:

   - *N* is the stacking member number for the switch
   - (Optional) *X* is the priority (1 - 255, but it should be less than the priority assigned to the Commander. The priority for the Standby should be the second highest in the stack. The member switches can be left at the default priority value of 128.)
   - *JxxxxA* is the product number of the switch (required). Any of the HP 3800 models can be installed and assume this provisioned configuration. If you specify a value for this parameter, then only a switch of this specific model type can assume this provisioned configuration
   - (Optional) *MAC-Addr* can be specified if you want a specific switch to assume this provisioned configuration. If this value is entered, then the type value for

the switch that has this MAC address must be correct, or a configuration error is logged and the switch is not allowed to join the stack.

7. Connect the stacking cables to the module ports for the desired stacking topology. For example, plug port 1 and 2 in a ring.
8. Install Stacking Modules into the other switches that will be members of the stack, but do not boot them yet.

   NOTE: It is highly recommended that you create a mesh topology for maximum throughput and resiliency of the stack. At a minimum, a ring topology should be created. A chain topology is not recommended because any hardware or software failure in the stack results in lost ports, which increases the amount of time for the recovery of full stack operation due to multiple reboots. See the *HP 3800 Installation and Getting Started Guide* for supported topologies.

9. Boot the Standby and Member switches. The second switch that is booted becomes the Standby. The remaining switches become Members when booted.
10. When all of the switches are booted, enter the show stacking command to confirm that the stack is operating correctly. The following example shows four switches in a meshed topology.

Example 178 Displaying output for four switches in a mesh topology

```
HP Stack 3800# show stacking
 Stack ID         : 00031cc1-de4d48c0
 MAC Address      : 1cc1de-4d48c9
 Stack Topology   : Mesh
 Stack Status     : Active
 Uptime           : 1d 2h 35m
 Software Version : KA.15.05

  Mbr
  ID  Mac Address   Model                                  Pri Status
  --- ------------- -------------------------------------- --- ---------------
  1   1cc1de-4d48c0 HP J9574A 3800-48G-PoE+-4SFP+ Switch   250 Commander
  2   1cc1de-4d8680 HP J9573A 3800-24G-PoE+-2SFP+ Switch   230 Standby
  3   1cc1de-4e3180 HP J9574A 3800-48G-PoE+-4SFP+ Switch   128 Member
  4   78acc0-3c2180 HP J9576A 3800-48G-4SFP+ Switch        128 Member
```

Using the plug-and-go method

1. Install a Stacking Module into an each HP 3800 switch that will be in the stack, as described in the *HP 3800 Switch Installation and Getting Started Guide*, but do not connect them together with stacking cables yet
2. Make sure that stacking is enabled for each switch.
   a. You can determine this by connecting a console to each switch and entering the command `show stacking` from the switch CLI.
   b. If stacking is disabled, enter the command `stacking enable`. This command causes the switch to reboot.

   NOTE: By default, stacking is enabled on all the HP 3800 switches when a Stacking Module is installed before the switch is powered up for the first time, but if the switches were powered up without a Stacking Module installed, then stacking is disabled

   If you are connecting stacking cables during/after switch boot, then multiple stacks can form (plug-and-go method).

3. Connect the stacking cables between the switches to form the desired stacking topology, then power on all switches.

When the switches that are stacked together complete booting up, one of the switches is elected as the Commander, one of the switches is elected as the Standby, the remaining switches become Members of the stack, and the stack becomes fully operational.

To find out the roles of the switches in the stack, connect a console to any of the switches and enter the `show stacking` command. You can use the MAC address and other information in the display to determine the roles of each of the switches.

Adding a switch to a stack as a new member

HP Networking stacking allows for switches to be added to the stack while the stack is operational (as long as the maximum number of ten switches in the stack is not exceeded).

1. Provision the stack for the new switch by entering the following command:

   ```
   HP Switch(config)# stacking member N pri X type JxxxxA [mac MAC-Addr]
   ```

   where:

   - *N* is the stacking member number for the switch
   - (Optional) *X* is the priority (1 - 255, but it should be less than the priority assigned to the Commander. The priority for the Standby should be the second highest in the stack. The member switches can be left at the default priority value of 128.)
   - *JxxxxA* is the product number of the switch (required). Any of the HP 3800 models can be installed and assume this provisioned configuration. If you specify a value for this parameter, then only a switch of this specific model type can assume this provisioned configuration
   - (Optional) *MAC-Addr* can be specified if you want a specific switch to assume this provisioned configuration. If this value is entered, then the type value for the switch that has this MAC address must be correct, or a configuration error is logged and the switch is not allowed to join the stack.

   NOTE: When the new switch has been provisioned in the stack, a complete configuration can be applied to the switch even before it is physically connected to the stack, connected to the network, and powered up.

2. Power on the new switch. The new switch does not become a member of the stack unless stacking has been enabled on the switch.
3. Install a Stacking Module into the new switch, connect the switch into the stack via the stacking cables, and form the desired stacking topology.
4. When the switch has finished booting, establish a console session with it and, from the config context, issue the command to enable stacking:

   ```
   HP Switch(config)# stacking enable
   ```

   This causes the switch to reboot. When the reboot is complete, the switch is a member of the stack with the attributes that you provisioned for it.

5. Confirm that the switch is now a member of the stack by issuing a `show stacking` command via a console session with any of the switches in the stack. The command output should show that the new switch is a Standby or Member of the stack with the member number and priority that you assigned to it.

When you add the switch to the stack, the following occurs:

- The Stack Revision Number is incremented by one.
- The Commander verifies that the new switch has the same switch software as the other switches in the stack, and downloads the software to the new switch if it does not. When downloading new software, there will be an automatic reboot during this process.
- A stack ID is assigned, even if the switch is later disconnected from the stack.

- The member's console is automatically redirected to the Commander.
- The OOBM IP address for that member is no longer reachable.

Removing a switch from the stack

You can remove a switch from the stack to be redeployed in another part of the network. The procedures vary depending on whether the switch is the Commander of the stack or not.

Removing a Member or the Standby

1. Establish a console session with the stack via direct console cable connection or telnet. If using the console cable, connect it to the Standby.
2. Enter the following command to remove the switch from the stack:

   ```
   HP Stack 3800(config)# stacking member N remove
   ```

   This causes the switch to lose its complete configuration and to be removed from the stack configuration. A subsequent show stacking command issued to the stack will show that the removed switch no longer exists in the stack.
3. Power down the removed switch.
4. Disconnect the stacking cables from the removed switch and from the other switches in the stack.

Removing the Commander

1. Establish a console session with the stack via direct console cable connection or telnet. If using the console cable, connect it to a switch other than the Commander
2. Enter the following command to force the Commander status over to the Standby switch:

   ```
   HP Stack 3800(config)# redundancy switchover
   ```

   This results in the Standby switch taking the role of the Commander and a new Standby being selected from the remaining member switches. The former Commander becomes a Member of the stack.
3. Enter the following command to remove the former Commander from the stack:

   ```
   HP Stack 3800(config)# stacking member N remove
   ```

   where *N* is the member number of the former Commander
4. Power down the removed switch.
5. Disconnect the stacking cables from the removed switch and from the other switches in the stack.

Renumbering stack members

If you did not provision the stack for the switches when you first created the stack, it is possible that members did not acquire the desired member numbers. The stack members can be renumbered.

A five-member stack is used in the following example with switches A, B, C, D, and E. These switches are members 1, 2, 3, 4, and 5, respectively. Switch B acquired member number 3 and switch C acquired member number 2.

1. In the global config context, enter the `remove` command option for switch B (member 3) and switch C (member 2):

   ```
   HP Stack 3800(config)# stack member 3 remove
   HP Stack 3800(config)# stack member 2 remove
   ```

   All configurations on the removed member switch are deleted, not just the stacking configuration.
2. Enter the following command:

```
HP Stack 3800(config)# stack member 2 type B's type
mac-address B's MAC address
```

This command clears the MAC address of the member 2 configuration to allow switch C's MAC address to be entered in the next command, without a duplicate MAC address occurring in the stack.

3. Reboot switch B (new member 2) and then switch C (new member 3).
4. To confirm that each switch now has the desired member number, enter the `show stacking` command.

Restoring the operation of a stack

Restoring operation after disconnecting a power cord

If a power cord becomes disconnected from one of the switches in the stack, the stack operation is affected. The stacking status of the switch that lost power is "Missing." Its record is retained in the stack configuration. The effect of the power loss depends on the role of the switch in the stack.

- If the Commander loses power, the Standby switch takes over as the Commander and one of the member switches in the stack is elected as the new Standby.
- If the Standby loses power, one of the member switches in the stack is elected as the new Standby.
- For any switch that loses power, all the network ports and stacking ports are non-operational until power is restored to the switch and it rejoins the stack. This affects the stacking topology.
- Reconnecting the power cord restores the operation of the switch, however, if the switch was either the Commander or the Standby, then it returns in a different role if the topology has 3 or more members. In a 2-member stack, a Standby that reboots will rejoin as Standby.

Restoring operation after disconnecting a stacking cable

If a stacking cable becomes disconnected from one of the switches in the stack, the effect depends on the stacking topology that is being used:

- Mesh—The stack topology is temporarily changed to a ring. To recover, simply reconnect the stacking cable; the mesh topology and the previous stack configuration is restored.
- Ring—There is little effect. The stack topology is temporarily changed to a chain topology. To recover, simply reconnect the stacking cable; the ring topology and the previous stack configuration is restored.
- Chain—The following occurs:
  - The smaller section (fragment) of the stack that results from the disconnection becomes `Inactive` (the `Stack Status` value shown in the output of the show stacking command is `Inactive`.
  - If the two resulting fragments are the same size, the fragment that contains the Commander will be `Active`, and the other fragment becomes `Inactive`.
  - Both fragments will have a Commander and a Standby selected (if there is more than one switch in each fragment).
  - When the stacking cable is reconnected to reform the chain:
    - The Commander and Standby of the Active fragment retain those roles for the resulting stack. If the original Commander was not in that fragment, then the stack will have a new Commander when the stack is reformed.
    - The switches in the Inactive fragment reboot and assume their new roles in the reformed chain.

Replacing a failed stack member

If a Stack Member fails, the effect on the stack depends on which member failed.

- If the Commander fails, the Standby switch takes over as the Commander and one of the Member switches in the stack is elected as the new Standby. All network ports and stacking ports on the failed switch become non-operational.
- If the Standby fails, one of the Member switches in the stack is elected as the new Standby. All network ports and stacking ports on the failed switch become non-operational.
- If a Member fails, all network ports and stacking ports on that switch become non-operational.

If a Stack Member fails:

1. Physically remove the Stack Member from the stack.
2. Replace the failed Stack Member.

NOTE: HP recommends using the same type (product or "J" number) switch as a replacement since all configuration information is retained.

- If you are using the same type switch as a replacement:
  1. Provision the new switch using the `stacking member N` command.
  2. Reconnect all Ethernet ports as they were on the failed switch.
- If you are using a different type switch as a replacement:
  1. Remove the failed switch from the stack configuration using the `stacking member N remove` command.
  2. Provision the new switch using the `stacking member N` command.
  3. Reconnect Ethernet ports and create a new stack configuration on the new switch.

If the replacement switch uses a different version of software, it will be updated automatically to match the software version running on the stack.

Replacing a failed stacking module

Replacing a failed stacking module is simpler than replacing a Stack Member since the switch configuration itself does not change. In this case, there is no need to re-provision the switch as a member of the stack. After you replace the stacking module, if the switch that experienced the module failure was Commander or Standby, the election of a new Commander and Standby is the same as described in “Replacing a failed stack member” (page 317).

Merging stack fragments

When two fragments have the same stack-id, the merge of the fragments is almost always allowed regardless of the merge policy. The Commander and Standby of the merged stack are selected based on the election rules. All of the switches in the previously inactive fragment or fragments reboot, and then join the Active fragment as Members.

If both fragments are Inactive, then an election process occurs. The two (or more) Commanders in the fragments are compared. The Commander is selected using the following criteria:

1. Highest Stack Rev
2. If the stack rev is the same for both, then choose the switch with the highest configured priority
3. If the priorities are the same for both, then choose the switch with the highest OS revision
4. If the OS revisions are the same, then choose the switch with the longest uptime
5. If the uptimes are the same, then choose the switch with the lowest MAC address

Modifying the stack topology

You can increase the efficiency and redundancy of the stack by adding stacking cables to create a stacking mesh instead of a ring. This modification can be performed while the switches are powered on and the stack is operating. After connecting the cables, enter the `show stack` command. The `Stack Topology field` value displays the new topology.

NOTE: Transitioning between mesh and ring topologies temporarily leaves the stack in an unsupported state (for example, when the stack is only connected to half of the additional ports). If a failure occurs during this time, redundancy is not guaranteed.

Downloading new software to the stack

The stack is essentially a single switch with the Commander unit controlling the management functionality, so the process of loading new software is identical to the process for a standalone HP 3800. (See the Appendix A, "File Transfers" in the *HP 3800 Management and Configuration Guide* for information on downloading software.)

After new software is loaded on the Commander, the Commander installs the software on all the stack members. The loading process can take some time.

To load the new software:

1. Load the new software onto the Commander via TFTP, USB, or Xmodem.
2. Once the new software is loaded, establish a console session with the stack and enter the following command:

   ```
   HP Stack 3800# boot system
   ```

   This causes the entire stack to be rebooted. Each unit is booted from its image unless you specify otherwise with options to this command. Make sure that you boot from the image to which you downloaded (that is, primary or secondary). If you add a new member to an existing stack, the Commander updates the new switch's software to match the current stack software. Multiple versions of software are not supported across stack members.

3. Confirm that the new software has been loaded on each stack member by entering the `member-context` command for each member. From the stack member context, you can see the switch software version that is running on that switch by entering the `show flash` or `show version` command.

   Syntax:

   `member-context` *`stack-member`*

   Sets the CLI context so that subsequent commands apply to the stack member that is specified.

Monitoring stacking

Use the following commands to monitor the status and configuration of the stack.

Syntax:

`show stacking`

Shows the current state of the stack.

Example 179 Displaying `show stacking` summary output

```
HP Stack 3800# show stacking
 Stack ID         : 00031cc1-de4d48c0
 MAC Address      : 1cc1de-4d48c9
 Stack Topology   : Mesh
 Stack Status     : Active
 Uptime           : 0d 2h 30m
 Software Version : KA.15.05

  Mbr
  ID  Mac Address   Model                                 Pri Status
  --- ------------- ------------------------------------- --- ------
  1   1cc1de-4d48c0 HP J9574A 3800-48G-PoE+-4SFP+ Switch   250 Commander
  2   1cc1de-4d8680 HP J9573A 3800-24G-PoE+-2SFP+ Switch   230 Standby
  3   1cc1de-4e3180 HP J9574A 3800-48G-PoE+-4SFP+ Switch   128 Member
  4   78acc0-3c2180 HP J9576A 3800-48G-4SFP+ Switch        128 Member
```

If stacking is disabled on the switch, the `show stacking` command displays this message:

```
Stacking is disabled.
```

Possible values for the various parameters are:

- Stack Topology: Chain, Ring, Mesh, Unknown
- Stack Status: Active, Fragment Active, Fragment Inactive
- Pri (Priority): 1 - 255>
- Status: Commander, Standby, Member, Standby Booting, Booting, Missing, Not Joined, Failed

Syntax:

`show stacking detail`

Shows the same output as the `show stacking` command, but with much more information about each device's stack port and connectivity, CPU state, uptime, and so on.

Example 180 Displaying `show stacking detail` output

```
HP Stack 3800(config)# show stacking detail
Stack ID          : 00031cc1-de4d48c0
MAC Address       : 1cc1de-4d48c9
Stack Topology    : Mesh
Stack Status      : Active
Uptime            : 4d 8h 50m
Software Version : KA.15.05.0000x
Name              : HP Stack 3800
Contact           : HP Admin
Location          : Cupertino

Member ID         : 1
Mac Address       : 1cc1de-4d48c0
Type              : J9574A
Model             : HP J9574A 3800-48G-PoE+-4SFP+ Switch
Priority          : 250
Status            : Commander
ROM Version       : KA.15.03
Serial Number     : SG3BDLW047
Uptime            : 0d 8h 44m
CPU Utilization   : 0%
Memory - Total    : 756,760,576 bytes
Free              : 647,056,384 bytes
Stack Ports -
#1 : Active, Peer member 2
#2 : Active, Peer member 3
#3 : Active, Peer member 4
#4 : Inactive
```

Syntax:

`show stacking member` *`stack-member-list`*

Shows detailed information about switches in the *`stack-member-list`* only.

Syntax:

`show stacking stack-ports member` *`stack-member`*

Shows the current state of the stacking ports of the specified member. If a member is not specified, the command shows the state of the ports of all physically present devices in the stack.

Example 181 Displaying `show stacking stack-ports` output

```
HP Stack 3800# show stacking stack-ports

       Stack       Peer   Peer
Member Port  State Member Port
------ ----- ----- ------ ----
1      1     Up    2      1
1      2     Up    3      2
1      3     Up    4      3
1      4     Down  0      0
2      1     Up    1      1
2      2     Up    3      1
2      3     Down  0      0
2      4     Up    4      4
3      1     Up    2      2
3      2     Up    1      2
3      3     Up    4      1
3      4     Down  0      0
4      1     Up    3      3
4      2     Down  0      0
4      3     Up    1      3
4      4     Up    2      4
```

If you specify specific stack members in the command, then the stacking port information for those members displays.

Troubleshooting stacking

Troubleshooting OOBM and split stack issues

If all the OOBM ports in the stack are in the same VLAN, you can use the `show oobm` commands to view the current state of all the switches. For example, if you have a five-member chain and member 4 fails or has the power removed, a stack split will occur with an active fragment on members 1-2-3 and an inactive fragment on member 5.

There is one IP address for the active fragment. This can be statically set by assigning an IP address to the global OOBM port.

If the stack splits, you can connect to the Active Fragment using the global OOBM IP address and then enter the `show oobm discovery` command to see if this active fragment has discovered any other members that are connected using the OOBM LAN.

In the following five member chain example, connect using the global IP address of 10.0.11.49. Once logged on, enter the `show stacking` command.

Example 182 Displaying stacking member status

```
HP Stack 3800# show stacking

Stack  ID          : 00011cc1-de4d87c0

MAC  Address       : 1cc1de -4d87e5
Stack  Topology    : Chain
Stack  Status      : Fragment
Active Uptime      : 0d 0h 5m
Software  Version  : KA.15.03.0000x

Mbr
ID  Mac Address    Model                                 Pri  Status
--- ------------- -------------------------------------- ---  ---------------
1   1cc1de-4d87c0  HP J9573A 3800-24G-PoE+-2SFP+ Switch  128  Commander
2   1cc1de-4dc740  HP J9573A 3800-24G-PoE+-2SFP+ Switch  128  Member
3   1cc1de-4dbd40  HP J9575A 3800-24G-2SFP+ Switch       128  Standby
4   1cc1de-4d79c0  HP J9576A 3800-48G-4SFP+ Switch       128  Missing
5   1cc1de-4da900  HP J9576A 3800-48G-4SFP+ Switch       200  Missing
```

Enter `show oobm discovery` to see if the members have been discovered using OOBM.

Example 183 Displaying oobm discovery

```
HP Stack 3800#  show oobm discovery

Active Stack Fragment(discovered) IP Address : 10.0.11.49

Mbr
ID   Mac Address    Status
--- -------------- ----------
1   1cc1de-4d87c0  Commander
2   1cc1de-4dc740  Member
3   1cc1de-4dbd40  Member

Inactive Stack Fragment(discovered) IP Address : 10.0.10.98

Mbr
ID  Mac Address    Status
--- -------------- ----------
5   1cc1de-4da900  Commander
```

Member 5 is up, but is an inactive fragment. It has an addressable IP address, which can be used to connect to this fragment.

Example 184 Connecting to a stack fragment

```
HP Stack 3800# telnet 10.0.10.98 oobm
HP J9576A 3800-48G-4SFP+ Switch
Software revision KA.15.03.0000x

Copyright (C) 1991-2011 Hewlett-Packard Development Company, L.P.

RESTRICTED RIGHTS LEGEND
Confidential computer software. Valid license from HP required for possession,
use or copying. Consistent with FAR 12.211 and 12.212, Commercial
Computer Software, Computer Software Documentation, and Technical
Data for Commercial Items are licensed to the U.S. Government under
vendor's standard commercial license.
HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P.
20555 State Highway 249, Houston, TX 77070
```

Enter the `show stacking` command.

Example 185 Displaying missing stack members

```
HP Stack 3800# show stacking

Stack  Topology      : Chain
Stack  Status        : Fragment Inactive
Uptime               : 0d 0h 7m
Software  Version    : KA.15.03.0000x

Mbr
ID  Mac Address   Model                                 Pri  Status
--- ------------- ------------------------------------- --- ---------------
1   1cc1de-4d87c0 HP J9573A 3800-24G-PoE+-2SFP+ Switch   128  Missing
2   1cc1de-4dc740 HP J9573A 3800-24G-PoE+-2SFP+ Switch   128  Missing
3   1cc1de-4dbd40 HP J9575A 3800-24G-2SFP+ Switch        128  Missing
4   1cc1de-4d79c0 HP J9576A 3800-48G-4SFP+ Switch        128  Missing
5   1cc1de-4da900 HP J9576A 3800-48G-4SFP+ Switch        200  Commander
```

Confirm by entering the `oobm discovery` command. Member 4 is down.

Example 186 Confirming stack member 4 is down

```
HP Stack 3800# show oobm discovery

Inactive Stack Fragment(discovered) IP Address: 10.0.10.98

Mbr
ID   Mac Address   Status
--- -------------- ----------
5   1cc1de-4da900 Commander

Active Stack Fragment(discovered) IP Address: 10.0.11.49

Mbr
ID   Mac Address   Status
--- -------------- ----------
1   1cc1de-4d87c0  Commander
2   1cc1de-4dc740  Member
3   1cc1de-4dbd40  Member
```

Using fault recovery/troubleshooting tools

Stacking provides tools and logging information to aid in troubleshooting problems specific to stacking. Problems may include:

- Installation/deployment issues
- Problems with initial stack creation
  - Problems with adding or removing members
  - Booting an existing stack
- Stacking failures encountered while running an existing stack

The tools used in troubleshooting problems are:

- Event Log
- Show commands
  - Show stacking
  - Show system
  - Show boot history
- Show tech
- LEDs

Troubleshooting installation and deployment issues

Installation and deployment issues include the initial deployment or creation of a stack, as well as adding additional members or removing members from a stack.

Problem:

When using the Deterministic method, one or more of the statically provisioned members did not join the stack.

Possible reasons a switch does not join a stack are:

- The switch being added is already a member of another stack and has a different stack ID.
- The maximum number of switches is already configured.

- The switch being added has been statically provisioned. The MAC address matches, but the switch type does not.
- There is a problem with the stack cable.
- The stack cables are connected in a way that creates an unsupported topology.
- Stack module failure.

Solution:

Perform a diagnostic fingerprint.

Troubleshooting issues with adding or removing members in the stack

Problem:

Cannot add a new switch to an existing stack.

Solution:

Identify root cause. Possible reasons for a member not joining an existing stack are:

- The switch being added has already been a member of another stack and has a different stack ID.
- The maximum number of switches is already configured.
- The switch being added has been statically provisioned, but switch type and MAC address in the configuration do not match the switch being added.
- There is a problem with the stack cable.
- There is a problem with the stack physical cabling. (illegal topology).

Problem:

The entire stack does not come up after a boot.

Solution:

There are several reasons why all members do not join the stack:

- There is a problem with the stack cable.
- Physical cabling was changed.
- Stack booted on incorrect configuration.
- One or more of the switches has a hardware problem (for example, bad power supply, back stacking module, corrupt flash).

Problem:

One or more of the members keeps rebooting and does not join the stack.

Possible reasons:

- An unresponsive member.
- Heartbeat loss—a stack that has a member no longer in the stack or a member failing after joining the stack.
- Illegal topology.

Problem:

After initial boot sequence, the activity and Link LEDs of an interface are not on and the ports are not passing traffic.

Solutions:

- Identify the “inactive fragment” and provide alternatives for recovery.
- Verify that all OOBMs are connected so that there is uninterrupted access.

Problem:

After a reboot, the selected Command or Standby are not the expected switches.

Solutions:

Check to see if the log files provide a reason why the Commander and Standby were chosen and which rule they matched.

Troubleshooting a strictly provisioned, mismatched MAC address

When switches are strictly provisioned, it is possible to enter an incorrect type or incorrect MAC address. If this occurs, the switch does not match the intended configuration entry and stacking attempts to add this switch as a new “plug-and-go” switch. If the stacking configuration already has 10 switches, then the “plug-and-go” fails.

The following example shows a stack with 9 members. There is a new J9576A switch that is supposed to be member 4; however, the MAC address was mis-typed, therefore, there is an “opening” for a plug-and-go at member 10. It will join as member 10.

Example 187 Displaying a stack with 9 members

This shows the stack before boot.

```
HP Stack 3800(stacking)# show stacking

Stack ID         : 00031cc1-de4d87c0
MAC Address      : 1cc1de-4dc765
Stack Topology   : Ring021560
Stack Status     : Active
Uptime           : 0d 0h 56m
Software Version : KA.15.05.0000x

Mbr
ID  Mac Address   Model                                Pri  Status
--- ------------- ------------------------------------- --- --------------
1   1cc1de-4d87c0 HP J9573A 3800-24G-PoE+-2SFP+ Switch  200 Standby
2   1cc1de-4dc740 HP J9573A 3800-24G-PoE+-2SFP+ Switch  128 Commander
3   1cc1de-4dbd40 HP J9575A 3800-24G-2SFP+ Switch       128 Member
4   1cc1de-444444 HP J9576A 3800-48G-4SFP+ Switch       175 Not Joined
5   1cc1de-000005 HP J9576A 3800-48G-4SFP+ Switch       128 Not Joined
6   1cc1de-000006 HP J9576A 3800-48G-4SFP+ Switch       128 Not Joined
7   1cc1de-000007 HP J9573A 3800-24G-PoE+-2SFP+ Switch  128 Not Joined
8   1cc1de-000008 HP J9573A 3800-24G-PoE+-2SFP+ Switch  128 Not Joined
9   1cc1de-000009 HP J9574A 3800-48G-PoE+-4SFP+ Switch  128 Not Joined
```

Example 188 Displaying a member joining the stack

This shows that, after booting, the switch is joined as member 10.

```
HP Stack 3800(config)# show stacking

Stack ID : 00031cc1-de4d87c0
MAC Address : 1cc1de-4dc765
Stack Topology : Mesh
Stack Status : Active
Uptime : 0d 1h 11m
Software Version : KA.15.05.0000x

Mbr
ID  Mac Address   Model                                 Pri  Status
--- ------------- -------------------------------------- --- -------------
1   1cc1de-4d87c0 HP J9573A 3800-24G-PoE+-2SFP+ Switch   200 Standby
2   1cc1de-4dc740 HP J9573A 3800-24G-PoE+-2SFP+ Switch   128 Commander
3   1cc1de-4dbd40 HP J9575A 3800-24G-2SFP+ Switch        128 Member
4   1cc1de-444444 HP J9576A 3800-48G-4SFP+ Switch        175 Not Joined
5   1cc1de-000005 HP J9576A 3800-48G-4SFP+ Switch        128 Not Joined
6   1cc1de-000006 HP J9576A 3800-48G-4SFP+ Switch        128 Not Joined
7   1cc1de-000007 HP J9573A 3800-24G-PoE+-2SFP+ Switch   128 Not Joined
8   1cc1de-000008 HP J9573A 3800-24G-PoE+-2SFP+ Switch   128 Not Joined
9   1cc1de-000009 HP J9574A 3800-48G-PoE+-4SFP+ Switch   128 Not Joined
10  1cc1de-4d79c0 HP J9576A 3800-48G-4SFP+ Switch        128 Member
```

To correct this issue:

1. Write down the correct MAC address.
2. Remove the member that was added using plug-and-go with the strictly provisioned, mismatched MAC address as shown in the following example.
3. Update the strictly provisioned entry with the correct MAC address.
4. Boot the switch.

Example 189 Removing a member and updating the entry with a MAC address

```
HP Stack 3800(config)# stacking member 10 remove reboot

The specified stack member will be removed from the stack and
its configuration will be erased. The resulting configuration
will be saved. The stack member will be rebooted and join as
a new member. Continue [y/n]? y

HP Stack 3800(config)# stacking member 4 type J9576A mac-address
```

Example 190 Displaying that member 4 joined the stack

This shows that member 4 has joined the stack.

```
HP Stack 3800(config)# show stacking

Stack ID         : 00031cc1-de4d87c0
MAC Address      : 1cc1de-4dc765
Stack Topology   : Mesh
Stack Status     : Active
Uptime           : 0d 1h 19m
Software Version : KA.15.05.0000x

Mbr
ID  Mac Address   Model                                  Pri Status
--- ------------- -------------------------------------- --- ---------------
1   1cc1de-4d87c0 HP J9573A 3800-24G-PoE+-2SFP+ Switch   200  Standby
2   1cc1de-4dc740 HP J9573A 3800-24G-PoE+-2SFP+ Switch   128  Commander
3   1cc1de-4dbd40 HP J9575A 3800-24G-2SFP+ Switch        128  Member
4   1cc1de-4d79c0 HP J9576A 3800-48G-4SFP+ Switch        175  Member
5   1cc1de-000005 HP J9576A 3800-48G-4SFP+ Switch        128  Not Joined
6   1cc1de-000006 HP J9576A 3800-48G-4SFP+ Switch        128  Not Joined
7   1cc1de-000007 HP J9573A 3800-24G-PoE+-2SFP+ Switch   128  Not Joined
8   1cc1de-000008 HP J9573A 3800-24G-PoE+-2SFP+ Switch   128  Not Joined
9   1cc1de-000009 HP J9574A 3800-48G-PoE+-4SFP+ Switch   128  Not Joined
```

Troubleshooting a mismatched stack-ID

Example 191 Displaying a stack with 3 unjoined switches

This is an example of a stack that has two members with three more members that have been strictly provisioned, following the deterministic method of initial installation.

```
HP Stack 3800# show stack

Stack ID         : 00031cc1-de4d87c0
MAC Address      : 1cc1de-4dc765
Stack Topology   : Chain
Stack Status     : Active
Uptime           : 0d 0h 2m
Software Version : KA.15.05.0000x

Mbr
ID  Mac Address   Model                                 Pri Status
--- -----------   ------------------------------------- --- -----------
1   1cc1de-4d87c0 HP J9573A 3800-24G-PoE+-2SFP+ Switch  200 Standby
2   1cc1de-4dc740 HP J9573A 3800-24G-PoE+-2SFP+ Switch  128 Commander
3   1cc1de-4dbd40 HP J9575A 3800-24G-2SFP+ Switch       128 Not Joined
4   1cc1de-4d79c0 HP J9576A 3800-48G-4SFP+ Switch       175 Not Joined
5   1cc1de-4da900 HP J9576A 3800-48G-4SFP+ Switch       128 Not Joined
```

When powering on switch #3, it does not join the stack.

Example 192 Logging output

```
HP Stack 3800 # show logging -r -s
I 10/02/00 00:46:56 02558 chassis: ST1-STBY: Stack port 3 is now on-line.
I 10/02/00 00:46:56 02558 chassis: ST2-CMDR: Stack port 2 is now on-line.
```

The stack ports for the new switch appear online, however, the `show stacking` command shows that the switch has not been recognized.

Example 193 Displaying the switch is not recognized

```
HP Stack 3800(config)# show stacking stack-ports member 1,2

Member 1

Member Stacking Port  State Peer Member   Peer Port
------ -------------- ----- ------------- -------
1      1              Down  0             0
1      2              Up    2             1
1      3              Down  0             0
1      4              Down  0             0

Member 2

Member Stacking Port  State Peer Member   Peer Port
------ -------------- ----- ------------- -------
2      1              Up    1             2
2      2              Up    0             0
2      3              Down  0             0
2      4              Down  0             0
```

The `show stacking` command does not show that the member is “Not Joined.”

A log file indicates that a “topo /hello” was seen from a switch that was not part of the current stack ID. The console of the switch that should have been member 3 shows the following example output.

Example 194 Displaying output from the “not joined” switch

```
HP Stack 3800# show stacking

Stack ID         : 00011cc1-de4dbd40
MAC Address      : 1cc1de-4dbd64
Stack Topology   : Unknown
Stack Status     : Active
Uptime           : 0d 0h 1m
Software Version : KA.15.05.0000x

Mbr
ID  Mac Address   Model                              Pri Status
--- ------------- ---------------------------------- --- -------------------
1   1cc1de-4dbd40 HP J9575A 3800-24G-2SFP+ Switch    128 Commander
```

The output is different if you have an inactive fragment, since this switch can have the configuration from an old stack. In this case, it might be inactive and show ‘missing’ switches from the old configuration. The stack-id value does not match the stack ID of the HP Stack 3800 stacking factory reset.

```
HP Stack 3800# stacking factory-reset
Configuration will be deleted and device rebooted,continue [y/n]?
 y
```

To join this switch to the other stack, execute the `stacking factory-reset` command to erase all of the stale stacking configuration information. This command automatically reboots the switch and on its subsequent boot, the switch is able to join the new stack.

Troubleshooting a strictly provisioned, mismatched type

When the MAC address matches a strictly provisioned configuration, it either matches the configured type and succeeds, or it does not match the type and fails. This is because the MAC address is unique and you cannot have duplicate MAC addresses.

The log messages indicate that this was the type of failure. The information in the log message helps you correct the configuration.

The switch that fails to join automatically reboots. Execute the `show stacking` command to view the mis-configured entry.

Example 195 Displaying the mis-configured entry

```
HP Stack 3800(config)# show stacking

Stack ID         : 00011cc1-de4d87c0
MAC Address      : 1cc1de-4d87e5
Stack Topology   : Mesh
Stack Status     : Active
Uptime           : 4d 0h 2m
Software Version : KA.15.05.0000x

Mbr
ID  Mac Address   Model                                 Pri Status
--- ------------- ------------------------------------- --- ----------------
1   1cc1de-4d87c0 HP J9573A 3800-24G-PoE+-2SFP+ Switch  128 Commander
2   1cc1de-4dc740 HP J9573A 3800-24G-PoE+-2SFP+ Switch  128 Standby
3   1cc1de-4dbd40 HP J9575A 3800-24G-2SFP+ Switch       128 Member
4   1cc1de-4d79c0 HP J9576A 3800-48G-4SFP+ Switch       175 Member
5   1cc1de-4da900 HP J9575A 3800-24G-2SFP+ Switch       128 Not Joined
```

The configuration entry for member 5 matches a J9576A switch that will be added, however, it will fail because it is configured as a J9575A switch.

The following example shows the log entries with the failure to join the stack.

Example 196 Log entries displaying stacking failures

```
W 10/06/00 03:24:37 03255 stacking: ST2-STBY: Provisioned switch with Member ID
5 removed due to loss of communication
I 10/06/00 03:24:37 02558 chassis: ST2-STBY: Stack port 4 is now on-line.
I 10/06/00 03:24:35 02558 chassis: ST4-MMBR: Stack port 2 is now on-line.
W 10/06/00 03:24:35 03274 stacking: ST1-CMDR: Member 5 (1cc1de-4da900) cannot
join stack due to incorrect product id: J9576A
```

You cannot re-type the configuration command with the same MAC address, member ID, and a different J-number. You need to remove the configuration and then reconfigure this switch member entry.

Example 197 Removing a stack member and reconfiguring

```
HP Stack 3800(config)# stacking member 5 remove

The specified stack member configuration will be erased. The
resulting configuration will be saved. Continue [y/n]? y

HP Stack 3800(config)# stacking member 5 type J9576A mac 1cc1de-4da900
This will save the current configuration. Continue [y/n]? y

Stack ID : 00011cc1-de4d87c0

tty=ansi HP Stack 3800(config)# show stacking
Strictly provisioned: Mis-Matched Type

Stack ID         : 00011cc1-de4d87c0
MAC Address      : 1cc1de-4d87e5
Stack Topology   : Mesh
Stack Status     : Active
Uptime           : 4d 0h 35m
Software Version : KA.15.05.0000x

Mbr
ID  Mac Address   Model                                 Pri Status
--- ------------- ------------------------------------- --- ---------------
1   1cc1de-4d87c0 HP J9573A 3800-24G-PoE+-2SFP+ Switch  128 Commander
2   1cc1de-4dc740 HP J9573A 3800-24G-PoE+-2SFP+ Switch  128 Standby
3   1cc1de-4dbd40 HP J9575A 3800-24G-2SFP+ Switch       128 Member
4   1cc1de-4d79c0 HP J9576A 3800-48G-4SFP+ Switch       175 Member
5   1cc1de-4da900 HP J9576A 3800-48G-4SFP+ Switch       128 Not Joined
```

Boot the switch with the matching MAC/Type.

Example 198 Displaying joined stack members

```
HP Stack 3800(config)# show stacking

Stack ID         : 00011cc1-de4d87c0
MAC Address      : 1cc1de-4d87e5
Stack Topology   : Mesh
Stack Status     : Active
Uptime           : 4d 0h 50m
Software Version : KA.15.05.0000x

Mbr
ID  Mac Address   Model                                 Pri Status
--- ------------- ------------------------------------- --- ----------------
1   1cc1de-4d87c0 HP J9573A 3800-24G-PoE+-2SFP+ Switch  128 Commander
2   1cc1de-4dc740 HP J9573A 3800-24G-PoE+-2SFP+ Switch  128 Standby
3   1cc1de-4dbd40 HP J9575A 3800-24G-2SFP+ Switch       128 Member
4   1cc1de-4d79c0 HP J9576A 3800-48G-4SFP+ Switch       175 Member
5   1cc1de-4da900 HP J9576A 3800-48G-4SFP+ Switch       128 Member
```

Troubleshooting maximum stack members exceeded

This failure can happen if you have an active stack that has already reached its maximum number of members. It can also happen when the maximum number of switches is reached with a combination of active members and strictly provisioned members.

Since one of the suggested initial deployment techniques is a deterministic method using strictly provisioned entries, this failure example demonstrates what occurs if the maximum number of members is reached by strictly provisioning ten members. At least one of these configuration entries has an incorrect MAC addresses. Similar to the mismatched MAC address example, the stack attempts to “plug-and-go” to add the switch, however, since the maximum number of membership has already been reached, the switch cannot join the stack.

The following example shows the `show stacking` output before the switch attempts to join.

Example 199 Displaying stack members before the join

```
HP Stack 3800(config)# show stacking

Stack ID         : 00031cc1-de4d87c0
MAC Address      : 1cc1de-4dc765
Stack Topology   : Mesh
Stack Status     : Active
Uptime           : 0d 1h 27m
Software Version : KA.15.05.0000x

Mbr
ID  Mac Address   Model                                   Pri Status
--- ------------- --------------------------------------  --- ---------------
1   1cc1de-4d87c0 HP J9573A 3800-24G-PoE+-2SFP+ Switch    200 Standby
2   1cc1de-4dc740 HP J9573A 3800-24G-PoE+-2SFP+ Switch    128 Commander
3   1cc1de-4dbd40 HP J9575A 3800-24G-2SFP+ Switch         128 Member
4   1cc1de-4d79c0 HP J9576A 3800-48G-4SFP+ Switch         175 Member
5   1cc1de-000005 HP J9576A 3800-48G-4SFP+ Switch         128 Not Joined
6   1cc1de-000006 HP J9576A 3800-48G-4SFP+ Switch         128 Not Joined
7   1cc1de-000007 HP J9573A 3800-24G-PoE+-2SFP+ Switch    128 Not Joined
8   1cc1de-000008 HP J9573A 3800-24G-PoE+-2SFP+ Switch    128 Not Joined
9   1cc1de-000009 HP J9574A 3800-48G-PoE+-4SFP+ Switch    128 Not Joined
10  1cc1de-00000a HP J9574A 3800-48G-PoE+-4SFP+ Switch    128 Not Joined
```

When a switch that does not match the MAC addresses attempts to join, that switch reboots when the maximum configuration is detected. The active stack logs the following:

```
W 10/07/00 06:01:11 03253 stacking: ST3-CMDR: Maximum number of switches in the stack has been
 reached.
Cannot add 1cc1de-4da900 type J9576A
```

The failure can be due to one of the strictly provisioned entries being incorrect. To correct this entry, reboot the switch. If there are already 10 switches in the stack, you cannot add additional switches at this time.

Troubleshooting a bad cable

Bad cables can cause the stack port to flap or go down completely. If there are an excessive number of port flaps, the port is disabled and the following log message appears:

```
W 10/06:00 23:23:16 03260 chassis: ST4-CMDR:
Stack port 1 disabled due to excessive errors. Check cable.
To reenable use 'stacking member 4 stack-port 1 enable'.
```

When this occurs, the `show stacking stack-ports` command shows the port with a status of “Disabled”.

Example 200 Displaying a disabled stack port

```
HP Stack 3800$ show stacking stack-ports

Member   Stacking  Port State  Peer Member Peer Port
------------------------------------------------------
1          1         Up           5            2
1          2         Up           2            1
1          3         Up           3            3
1          4         Up           4            3
2          1         Up           1            2
2          2         Up           3            1
2          3         Up           4            4
2          4         Up           5            3
3          1         Up           2            2
3          2         Down         0            0
3          3         Up           1            3
3          4         Up           5            4
4          1         Disabled     0            0
4          2         Up           5            1
4          3         Up           1            4
4          4         Up           2            3
5          1         Up           4            2
5          2         Up           1            1
5          3         Up           2            4
5          4         Up           3            4
```

If the cable failure is more solid, the port is in the DOWN state. The logs show any transition.

```
I 10/07/00 06:01:15 02559 chassis: ST4-STBY: Stack port 3 is now off-line.
I 10/07/00 06:01:16 02559 chassis: ST3-CMDR: Stack port 3 is now off-line.
I 10/07/00 06:01:16 02559 chassis: ST2-MMBR: Stack port 1 is now off-line.
I 10/07/00 06:01:15 02559 chassis: ST5-MMBR: Stack port 2 is now off-line.
I 10/07/00 06:01:15 02558 chassis: ST2-MMBR: Stack port 1 is now on-line.
I 10/07/00 06:01:12 02558 chassis: ST5-MMBR: Stack port 2 is now on-line.
I 10/07/00 06:01:10 02558 chassis: ST4-STBY: Stack port 3 is now on-line.
```

The following example shows member 3, port 2, which should be connected to member 4, port 1. The ports are down because the cable is bad or disconnected.

Example 201 Displaying that two ports are down due to a bad connection

```
HP Stack 3800# show stacking stack-ports

Member   Stacking  Port State  Peer Member Peer Port
------------------------------------------------------
1          1         Up           5            2
1          2         Up           2            1
1          3         Up           3            3
1          4         Up           4            3
2          1         Up           1            2
2          2         Up           3            1
2          3         Up           4            4
2          4         Up           5            3
3          1         Up           2            2
3          2         Down         0            0
3          3         Up           1            3
3          4         Up           5            4
4          1         Down         0            0
4          2         Up           5            1
4          3         Up           1            4
4          4         Up           2            3
5          1         Up           4            2
5          2         Up           1            1
5          3         Up           2            4
5          4         Up           3            4
```

The solution in both cases is to ensure that the cable is firmly connected at both ends. If the problem continues, replace the cable. It is possible that there could be a problem with the stack port itself. In this case, validation of this issue requires the installation of a known good cable to see if that cable also fails.

The port state is not UP until both ends of the cable are connected and the cable has been validated as a genuine HP cable.

To view the statistics on the physical port, execute the `show tech` command in member-context 4. The following examples show the types of information displayed.

Example 202 Displaying `show tech` output

```
Port Number : 1                    State : Available
Last Event : Available             Start Req : 1
NE Present : 1                     HPID Good : 1
HPID Fails : 0                     FE Present : 1
Rem Dev Rdy : 1
ESSI Link : 1                      ESSI Good : 1
ESSI Fails : 0                     ESSI TX En : 1
ICL Good : 1                       ICL Enabled : 1
LP Local RDY: 1                    LP Rem RDY : 1
LP DONE : 1                        ICL FailCnt : 0 (10 second interval)
ICL FailCnt : 0 (10 minute interval)
NE Presence HW : 1
FE Presence HW : 1
Rem Dev Rdy HW : 1
Local Dev Rdy HW : 1
Asserted NE Presence HW : 1
Asserted FE Presence HW : 1
Asserted Rem Dev Rdy HW : 1
Phy Frame Errors : 0
Invalid Status Errors : 0
Invalid Packet Type Errors : 0
Incomplete Packet Errors : 0
Checksum Errors : 0
ESSI Flow Out This Port (HW) : 0x2
```

Example 203 Displaying trace information for a port

```
Trace for Port 1
[ 0] [Info ] Start Request Received (Empty) [0]
[ 1] [Info ] Waiting for Stack Module Good (Empty) [0]
[ 2] [Info ] Stack Module Good Received (Empty) [0]
[ 3] [Info ] Cable Insertion Detected (Empty) [1]
[ 4] [Info ] Re-enable NE Present Int [487]
[ 5] [Info ] Starting Cable HPID Validation (Inserted) [488]
[ 6] [Info ] Skipping Cable HPID Validation (Inserted) [488]
[ 7] [Info ] Far End Insertion Detected (Valid) [988]
[ 8] [Info ] Polling for ESSI phy link up (Valid) [988]
[ 9] [Info ] ESSI Link Up [9988]
[10] [Info ] ESSI Link Good (Valid) [9988]
[11] [Info ] ESSI Linked at 9988 ms [9988]
[12] [Info ] Remote Device Ready Detected (Valid) [10898]
[13] [Info ] ICL Change Request Enable (Cable Ready) [10898]
[14] [Info ] Detected Remote Ready Drop. (Cable Ready) [12651]
[15] [Info ] ICL Good. Behind. Partner ready. (Cable Ready) [12988]
[16] [Info ] ICL GOOD received at 2091 ms [12988]
[17] [Info ] Partner LP ready. (Cable Ready) [13980]
[18] [Info ] Set Device Ready. (Cable Ready) [13987]
[19] [Info ] ESSI Link Verfied [13988]
[20] [Info ] ESSI Able to Transmit (Cable Ready) [13988]
[21] [Info ] ESSI Verified at 3091 ms [13988]
[22] [Info ] Cable Available (Available) [13988]
```

Troubleshooting when a switch crashes and reboots

Although the switch software is highly reliable, a switch in the stack can experience a software issue that results in the crash and reboot of that switch. This crash can happen in the software running on the CPU in the management CPU or on the software running on the CPUs in the interfaces. In either case, crash information is generated and the switch is rebooted.

The resiliency of the stack is determined by the stacking topology, however, in all cases, the interfaces/ports on the switch that crashes are brought down and a reboot of that switch occurs.

The following table describes how the stack reacts to the crashing switch, depending on what role the switch had when the crash occurred. The assumption in this table is that the topology is a resilient topology (that is, a mesh or ring).

Stacking role	**Description**
Commander	• The standby takes over as the new Commander • A new standby is elected • Crashing switch writes core file to local stable storage • Crashing switch reboots and joins the stack • Core file and crash information for this switch is available from the Commander
Standby	• A new standby is elected • Crashing switch writes core file to local stable storage • Crashing switch reboots and joins the stack • Core file and crash information for this switch is available from the Commander
Member	• Crashing switch writes core file to local stable storage • Crashing switch reboots and joins the stack • Core file and crash information for this switch is available from the Commander

After a switch crashes, you can collect data to help HP understand why the crash occurred. The information is a combination of crash data, crash log, and core-dump files. The `show tech` command displays logs of events that happened right before the crash.

Troubleshooting an unresponsive reboot

An unresponsive reboot occurs when a member does not respond to an update packet.

Example 204 Reboot output

```
 /* SSM_SWITCH_LOST_EVENT */
// 30 States -> Initial !Discovery !Bid for Cmdr !Become cmdr !
/* Switch Lost */{ssmIgnore ,ssmMbrRmvMbr ,ssmMbrRmvMbr ,ssmMbrRmvCmdr ,
// States -> Cmdr chas wait!Cmdr RFS start! Commander ! Cmdr Merge !
/* Switch Lost */ ssmMbrRmvCmdr ,ssmMbrRmvCmdr ,ssmMbrRmvCmdr ,ssmMbrRmvCmdr ,
// States -> Wait for chas !wait for type !stby RFS start!stby RFS sync !
/* Switch lost */ ssmMbrRmvMbr ,ssmMbrRmvMbr ,ssmMbrRmvMbr ,ssmMbrRmvMbr ,
// States -> Standby !Mbr RFS wait ! Member !Pass through !
/* Switch Lost */ ssmMbrRmvStby ,ssmMbrRmvMbr ,ssmMbrRmvMbr ,ssmIllegal },
```

Troubleshooting an unexpected Commander or Standby switch selection

When a switch stack is established and a boot/reboot of the stack is performed, the Commander and Standby are selected based on the configured switch priority. There are other rules in the election process that can override this priority.

Example 205 Displaying the running configuration with priority

```
HP Switch(config)# show running-config

; hpStack Configuration Editor; Created on release #KA.15.05.0000x
; Ver #01:00:01

hostname "HP Stack 3800"
stacking
member 1 type "J9573A" mac-address 1cc1de-4d87c0
member 2 type "J9573A" mac-address 1cc1de-4dc740
member 3 type "J9575A" mac-address 1cc1de-4dbd40
member 3 priority 200
member 4 type "J9576A" mac-address 1cc1de-4d79c0
member 4 priority 175
member 5 type "J9576A" mac-address 1cc1de-4da900exit
```

On a boot of the stack, member 3 becomes a Commander and member 4 becomes a Standby, based on priority. If this were a chain with member 1 at one end of the chain and member 5 at the other end, the number of hops between switches will be part of the election process.

Managing interactions with other switch features

Managing SSH or Telnet sessions

Switches in a non-stacking configuration support up to six sessions running SSH or Telnet concurrently. However, if stacking is configured, each stacking connection reduces the number of sessions available. For example, five connections into the stack leaves only one session available for SSH or Telnet.

Managing switch-level configuration interactions

In a stack, the Commander functions as a single switch and the Standby and Members function as additional network ports for that switch. Switch configuration is performed in the same manner as for any other switch, as described in these manuals for the HP 3800 switches:

- *Basic Operation Guide*
- *Management and Configuration Guide*
- *Advanced Traffic Management Guide*
- *Multicast and Routing Guide*
- *Access Security Guide*
- *IPv6 Configuration Guide*

Managing port-level configuration interactions

For features that are configured on specific switch ports in a stack, the configuration procedures are the same as for stand-alone switches, but the port designations for the ports in the stack are modified. Each port is identified by the stack member ID of its switch, followed by a slash and then the port number as it is shown on the switch. For example, for a switch with stack member ID 3, port 10 on that switch would be identified as port 3/10.

Example 206 Displaying `show interfaces brief` output for port 3/10

```
HP Switch(config)# show interfaces brief 3/10

 Status and Counters - Port Status

                         | Intrusion                           MDI  Flow Bcast
  Port         Type      | Alert     Enabled Status Mode       Mode Ctrl Limit
  ------------ --------- + --------- ------- ------ ---------- ---- ---- -----
  3/10         100/1000T | No        Yes     Down   1000FDx         off  0
```

Similarly, CLI commands requiring specific port (interface) numbers on an HP 3800 switch configured for stacking require the modified port designations. For example, to enter the port context for port 10 on stack member 2, enter this command:

```
HP Switch(config)# interface 2/10
 HP Switch(eth-2/10)#_
```

In the output containing designated port numbers for an HP 3800 switch configured for stacking, the port numbers are likewise listed in the modified format.

Example 207 Displaying `show interfaces config` output

```
HP Switch(config)# show interfaces config

 Port Settings

  Port   Type      | Enabled Mode         Flow Ctrl MDI
  ------ --------- + ------- ------------ --------- ----
  1/1    100/1000T | Yes     Auto         Disable   Auto
  1/2    100/1000T | Yes     Auto         Disable   Auto
  1/3    100/1000T | Yes     Auto         Disable   Auto
   .
   .
   .
  2/1    100/1000T | Yes     Auto         Disable   Auto
  2/2    100/1000T | Yes     Auto         Disable   Auto
  2/3    100/1000T | Yes     Auto         Disable   Auto
  2/4    100/1000T | Yes     Auto         Disable   Auto
.
.
.
```

Attempting to enter a CLI command for a port on a stack member without using the modified port number format generates a “Module not present...” message such as the following:

```
HP Switch(config)# interface 10
Module not present for port or invalid port: 10
HP Switch(config)#
```

LACP support

LACP trunking can support up to 144 trunks in a stacking configuration, each with up to eight links (ports) per trunk.

Managing OOBM ports

Each OOBM port of a member is assigned one MAC address from that member's manufacturing allocated range. The OOBM port also can be assigned an IP address (IPv4 /v6/DHCP/Manual/Auto-Config/LinkLocal). The commander's OOBM IP address (called the Global IP address) is used for managing the commander through the OOBM port.

After switchover/failover of control from the Commander to the Standby, the OOBM port IP address of the new Commander is the Global IP address. This change in address causes some undesirable behavior (after failover):

- When using DHCP or DHCPv6, the new Commander requests a new lease and typically receives a new network address (IPv4 or IPv6). With OOBM high availability (HA), it will seem as if a new link has come up requesting a network address.
- IPV6 link-local or auto-config addresses will change.

Using a static IP address avoids these issues. During failover, it will be as if the IP address is reconfigured. All ARP entries are updated automatically.

For more information on OOBM operation, see appendix J, "Network Out-of-Band for the HP 3800 Switches," in the latest *Management and Configuration Guide* for your HP 3800 switch.

Understanding stacking election

Electing a Commander

In those cases in which the Commander of the stack is not identified as described in "Creating a stack" (page 311), the stack undergoes a Commander Election process. This occurs when the entire stack is rebooted simultaneously, such as during a building power failure recovery, or when the stack becomes split and the Commander is isolated in the Inactive fragment, requiring the Active fragment to elect a new Commander.

All of the switches go through discovery and election at the same time. There is an election timer that is set for 60 seconds, and if there are no new switches discovered during that timeout period, the switches in the stack enter the election phase.

During the election process, for each group of switches that has the same STACK-ID (they should all be the same), these steps occur:

1. The switches with the highest Stack Revision are discovered.
2. The switch with the highest configured priority is selected as the Commander.
3. If there are switches with the same "highest" priority, the switch that was the previous Commander is selected.
4. If no switches were previous commanders, the switch that was the previous Standby is selected.
5. If none of the above conditions apply, the switch with the lowest MAC Address is selected as the Commander.

Electing a Standby

The Standby switch is selected by the Commander following the same rules used to elect the Commander. Like the Commander, the Standby switch is not changed unless a failure occurs (for example, the Standby switch fails, or the Commander fails and the Standby becomes the Commander).

NOTE: Since the Commander will update the revision software and set the stack IDs of all the switches, this information will be the same for all Standby switch contenders.

The criteria used by the Commander to select the Standby is in this priority order:

1. A switch with the same system revision software as the commander is available. This speeds up the initial boot since the stack will not have to wait for the standby to be updated. (If this were not the case, then the selected Standby would need to be reloaded with the new system and rebooted, resulting in the selection of a new Standby. This process would continue until either the original Standby was rebooted or a Standby was chosen that already had the correct system revision).
2. For all switches with the Commander's revision software, the switch with the highest priority that is not the current Commander will become the Standby switch.

NOTE: It is possible for the Standby to have a higher priority than the Commander, if the priority of the Standby was increased after the Commander becomes the Commander. (The Commander is not changed unless it fails or is on the Inactive fragment side of a stack that becomes split).

3. If there are two or more switches whose priority is equally high, then the Commander will look at the topology of the stack and pick a switch that is the most hops away from the Commander. This will increase the probability that the Commander and Standby switch will be in different stack fragments should a failure occur.
4. If the priorities and hop counts of the contenders are the same, then the switch with the lowest MAC address is selected as the Standby switch.

9 QinQ (Provider bridging)

Command syntax	Description	Default value	CLI reference page
qinq mixedvlan *tag-type* [*tpid*]	Enables QinQ	Disabled	347
qinq svlan *tag-type* [*tpid*]			
svlan [*vid* | *ascii-name-string*] [no] svlan *vid*	Sets up S-VLANs		347
svlan *vid*	Configures per-port S-VLAN membership		348
[no] interface [*port-list*] | Trkx qinq port-type { *customer-network* | *provider-network* }	Configures port types	port-type: provider	349
no qinq	Disables QinQ	Standard VLAN operations apply	350
show qinq	Displays QinQ configuration and status		351
show vlans	Displays a switch VLAN configuration		351
show vlans *vlan-id*	Displays the configuration for a VLAN		352
show vlans *vlan-id*	Displays the VLAN membership of ports		353
show spanning tree	Displays spanning tree status		354

Introduction

This chapter describes how to enable QinQ operations on the switch and how to configure provider bridge S-VLANs and port assignments.

The IEEE 802.1ad specification, commonly known as QinQ or provider bridging, extends the IEEE 802.1Q standard by providing for a second tier of VLANs in a bridged network. The general purpose of QinQ is to allow frames from multiple customers to be forwarded (or tunneled) through another topology (provider network) using service VLANs or S-VLANs. The provider bridge, which may comprise multiple devices in the service provider domain, looks like a simple bridge port to the customer's traffic and maintains the customer's VLANs.

Figure 70 (page 342) shows a sample QinQ topology and use model. Customer A has LANs spread across multiple site locations and may want to link them together in a single logical LAN. To do this, the customer could have a cable laid out for the entire distance interconnecting the three sites. A more cost-effective and scalable alternative, however, would be to tunnel frames through the provider's network to interconnect all the sites subscribing to the service. This solution can be delivered using QinQ.

Figure 70 QinQ network diagram

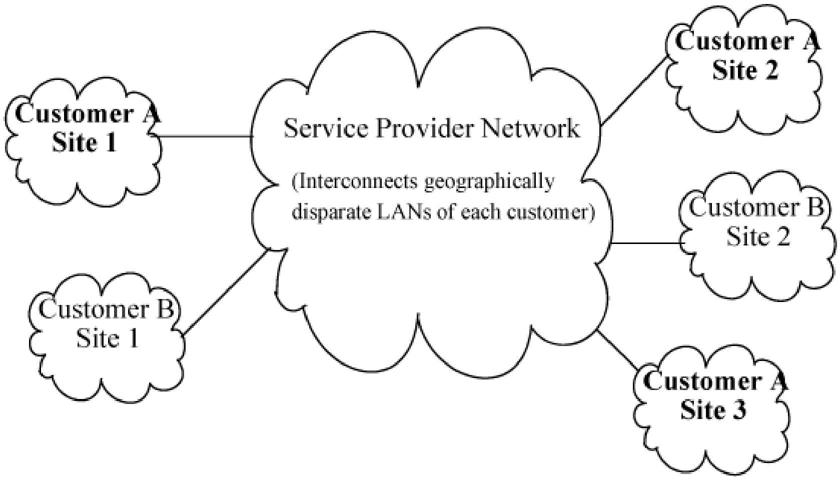

NOTE: The Service Provider and customers may belong to the same business entity, as in the case where a single enterprise uses QinQ to help segregate local networks and increase the scalability of their backbone infrastructure.

How QinQ works

Under QinQ, the provider network operates on a different VLAN space, independent of the VLANs that are used in the customer network as shown in Figure 71 (page 342).

Figure 71 VLANs in a QinQ configuration

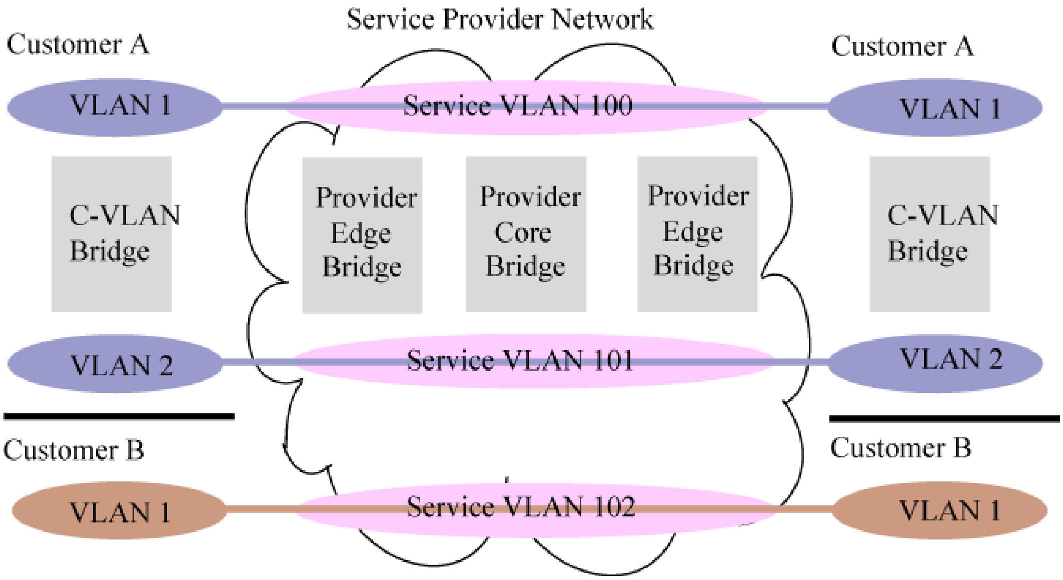

Customer VLANs (referred to as *C-VLANs* by the IEEE 802.1ad specification) are not used to make any forwarding decisions inside the provider network where customer frames get assigned to service VLANs (S-VLANs). Inside the provider cloud, frames are forwarded based on the S-VLAN tag only, while the C-VLAN tag remains shielded during data transmission. The S-VLAN tag is removed when the frame exits the provider network, restoring the original customer frame.

Features and benefits

- Increases the VLAN space in a provider network or enterprise backbone.
- Reduces the number of VLANs that a provider needs to support within the provider network for the same number of customers.
- Enables customers to plan their own VLAN IDs, without running into conflicts with service provider VLAN IDs.
- Provides a simple Layer 2VPN solution for small-sized MANs (Metropolitan Area Networks) or intranets.
- Provides for customer traffic isolation at Layer 2 within a Service Provider network.

Configuring QinQ (Overview)

QinQ must be configured on all the devices and ports participating in the provider bridge. Typically, customer facing ports are configured as untagged members of S-VLANs and provider facing ports are configured as tagged members of S-VLANs. Per the IEEE 802.1ad specification, there is no condition binding port types (customer or provider) to untagged or tagged S-VLAN memberships. Therefore, when configuring QinQ tunnelling on the switch, you would first configure per-port S-VLAN membership (tagged or untagged), and then configure the port type as `customer-network` or `provider-network`, depending on the device to which the switch port is connected.

NOTE: A customer-network port can receive S-VLAN tagged frames if the customer and provider agree on the S-VID association for that customer and the customer device is capable of sending S-VLAN tagged frames.

To configure QinQ take the following steps on all participating provider switches:

1. Enable QinQ on the device, selecting the appropriate QinQ mode (S-VLAN or mixed VLAN mode).
2. Save the configuration and reboot the switch.
3. Configure S-VLANs and assign per port VLAN membership.
4. Configure port-types for all of the switch ports that carry QinQ traffic across the network.
5. (Optional) Assign priorities to traffic passing through the provider network (See “Displaying QinQ configuration and status” (page 351)).

△ **CAUTION:** A reboot is required to enable/disable QinQ operations on the switch. When moving between QinQ modes (`qinq mixedvlan` to `qinq svlan` or vice versa), the switch boots up with a default configuration for the new qinq mode and the configuration parameters of the current mode will be erased. See “Disabling QinQ” (page 350) for details.

Configuration example

This configuration example uses four HP switches to establish a QinQ tunnel through the provider network.

Figure 72 QinQ configuration example

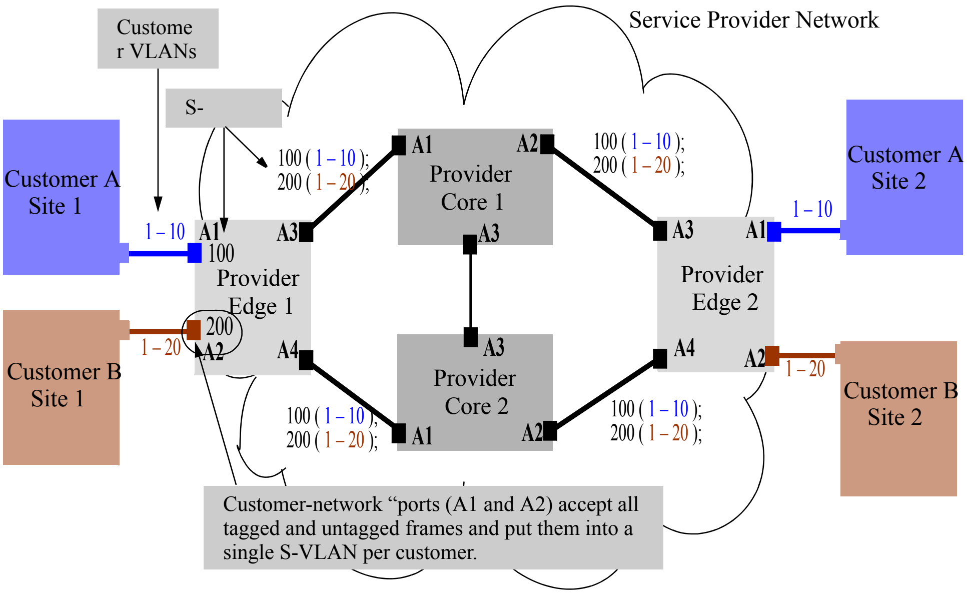

The design parameters for this example are as follows:

- The provider edge bridge and the provider core bridge are configured in svlan mode.
- Each customer is associated with a single S-VLAN connecting two separate sites: customer A's VLANs (C-VLANs 1-10) are associated with S-VLAN 100; and customer B's VLANs (C-VLANs 1-20) are associated with S-VLAN 200.

NOTE:

- The VLANs of customers A and B can overlap: this will not result in intermixing of customer frames in the provider cloud because the S-VLANs associated with each customer are different.
- Core devices are not mandatory to establish a QinQ tunnel. For example, two edge-bridges can be connected directly to create a provider bridge network.
- The relationship between S-VLANs and C-VIDs is typically one to many. An alternative configuration might associate a single customer's C-VIDs with more than one S-VLAN. Such a configuration would most likely be used to tunnel distinct C-VIDs through various S-VLANs, but seldom be used to send the same C-VID through multiple S-VLANs.

Figure 73 Configuration example: Edge Switch 1

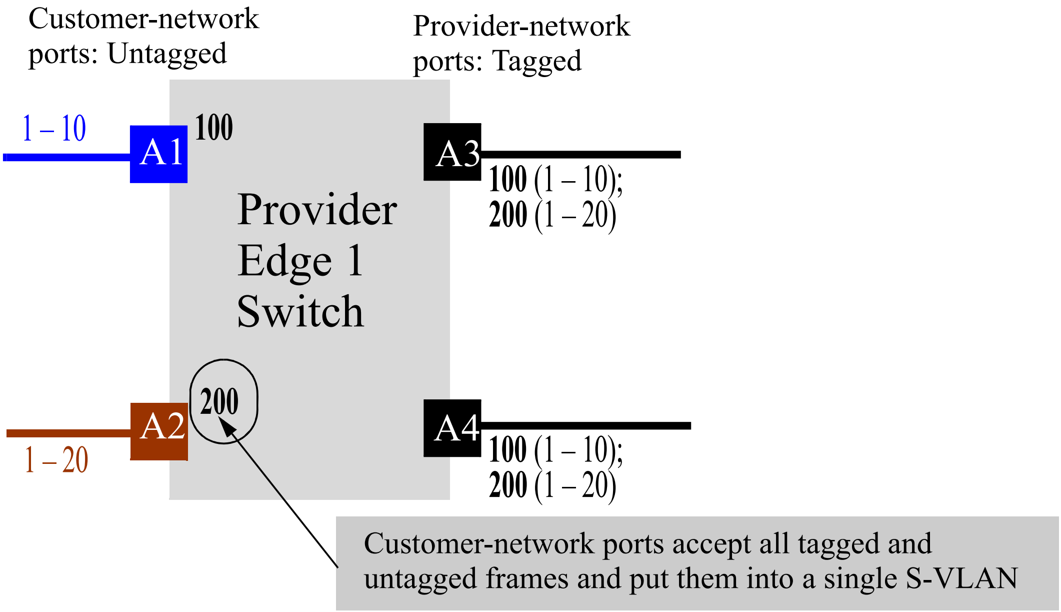
Customer-network ports: Untagged
Provider-network ports: Tagged
1 – 10
A1
100
A3
100 (1 – 10);
200 (1 – 20)
Provider Edge 1 Switch
200
A2
1 – 20
A4
100 (1 – 10);
200 (1 – 20)
Customer-network ports accept all tagged and untagged frames and put them into a single S-VLAN

At the end of the configuration, the following settings will apply:

- All customer A site traffic received on port A1 will be associated with S-VLAN 100. This is independent of the C-VLAN tag information that the customer frames may carry.
- All customer B Site 1 traffic will be associated with S-VLAN 200 and be switched out to the core (uplinks A3, A4) with the S-VLAN tag-id of 200.
- The frame size will increase by 4 since ports A3 and A4 are tagged members of S-VLAN 100 and 200.

To configure the switch, follow these steps:

1. Enable QinQ:

```
Edge 1(config)# qinq svlan tag-type 88a8
```

2. Reboot the box with the configuration saved to transfer into svlan bridge mode.

NOTE: A reboot is required for the QinQ enable command to take effect.

3. Configure S-VLANs and ports connected to the customer network.

```
Edge1(config)# svlan 100
Edge1(svlan-100)# untagged A1
Edge1(svlan-100)# exit
Edge1(config)# int A1 qinq port-type customer-network
Edge1(config)# svlan 200
Edge1(svlan-200)# untagged A2
Edge1(svlan-200)# exit
Edge1(config)# int A2 qinq port-type customer-network
```

NOTE: In this example, customer A is assigned S-VLAN 100 and customer B is assigned S-VLAN 200. However, the same customer can be associated with more that one SVLAN. Also, interfaces A1 and A2 are configured as customer network ports because they are linked to customer bridges.

4. Configure the provider ports leading to the core of the provider network.

```
Edge1(config)# svlan 100 tagged A3, A4
Edge1(config)# svlan 200 tagged A3, A4
Edge1(config)# interface A3,A4 qinq port-type provider-network
```

NOTE: As recommended by IEEE 802.1ad specification, uplink ports should generally be configured as tagged ports for S-VLANs that are used to carry customer traffic. However, this is not a mandatory requirement on HP switches—S-VLANs that are used for internal provider network use (not carrying customer traffic but for management of the provider network devices) can have untagged port memberships.

Configure provider Edge 2 switch

The configuration details for the Edge 2 switch mirrors the configuration for the Edge 1 switch. All customer traffic received on port A1 from customer A's site 2 will be associated with S-VLAN 100. Similarly, all customer B's site 2 traffic will be associated with S-VLAN 200.

To configure the switch, follow these steps:

1. Enable QinQ:

```
Edge 2(config)# qinq svlan tag-type 88a8
```

2. Reboot the box with the configuration saved to transfer into S-VLAN bridge mode.
3. Configure S-VLANs and customer ports connected to the customer network.

```
Edge2(config)# svlan 100
Edge2(svlan-100)# untagged A1
Edge2(svlan-100)# exit
Edge2(config)# int A1 qinq port-type customer-network
Edge2(config)# svlan 200
Edge2(svlan-200)# untagged A2
Edge2(svlan-200)# exit
Edge2(config)# int A2 qinq port-type customer-network
```

4. Configure the provider ports leading to the core of the provider network.

```
Edge1(config)# svlan 100 tagged A3, A4
Edge1(config)# svlan 200 tagged A3, A4
Edge1(config)# interface A3,A4 qinq port-type provider-network
```

Configuring provider Core 1 switch

Figure 74 Configuration example: Core 1 Switch

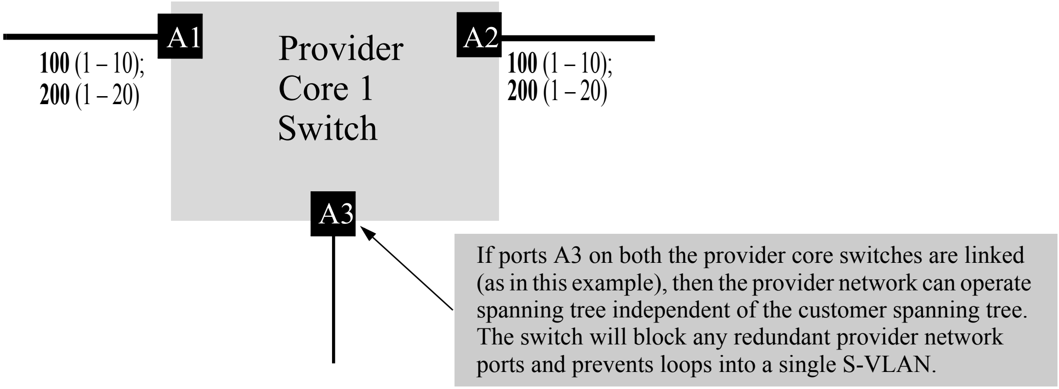

If ports A3 on both the provider core switches are linked (as in this example), then the provider network can operate spanning tree independent of the customer spanning tree. The switch will block any redundant provider network ports and prevents loops into a single S-VLAN.

To configure the Core 1 switch:

1. Enable QinQ:

```
Core 1(config)# qinq svlan tag-type 88a8
```

2. Reboot the box with the configuration saved to transfer into svlan bridge mode.

3. Configure S-VLANs and port assignments.

   ```
   Core 1(config)# svlan 100
   Core 1(svlan-100)# tagged A1, A2
   Core 1(svlan-100)# exit
   Core 1(config)# svlan 200
   Core 1(svlan-200)# tagged A1, A2
   Core 1(svlan-200)# exit
   Core 1(config)# interface A1,A2 qinq port-type provider-network
   ```

NOTE: The S-VLAN configuration for the core devices is based on what VLANs the edge devices (Edge 1 and 2) can send. Per the 802.1ad specification, all ports carrying customer traffic will be tagged on the VLAN that the port carries customer frames on.

To configure the Core 2 switch:

1. Enable QinQ:

   ```
   Core 2(config)# qinq svlan tag-type 88a8
   ```

2. Reboot the box with the configuration saved to transfer into svlan bridge mode.
3. Configure S-VLANs and port assignments.

   ```
   Core 2(config)# svlan 100
   Core 2(svlan-100)# tagged A1, A2
   Core 2(svlan-100)# exit
   Core 2(config)# svlan 200
   Core 2(svlan-100)# tagged A1, A2
   Core 2(svlan-100)# exit
   Core 2(config)# interface A1,A2 qinq port-type provider-network
   ```

Verify the configuration

After the edge and core switch configurations are completed, QinQ operations can begin. To verify operations, it should be possible to assign IP-addresses to customer A or B devices in site 1 and site 2 and ping them. If everything has been configured correctly, traffic will flow through the provider network cloud and reach the other site seamlessly. To verify the configuration, see also "Displaying QinQ configuration and status" (page 351).

Enabling QinQ

By default, QinQ is disabled on the switch. To enable QinQ, the switch must be put into either in mixed VLAN mode or QinQ SVLAN mode by issuing one of the following commands from configuration mode on the CLI.

Syntax:

`qinq mixedvlan` *`tag-type`* [*`tpid`*]

From config mode, globally enables QinQ mixed mode, an environment that supports both S-VLAN and C-VLAN traffic on the same device. This command requires a reboot to take effect. Default: Disabled.

Syntax:

`qinq svlan` *`tag-type`* [*`tpid`*]

From config mode, globally enables QinQ SVLAN mode, an S-VLAN only environment that supports supports port-based or s-tagged interfaces of the standard. Requires a reboot to take effect. Default: Disabled.

Setting up S-VLANs

S-VLANs are created via the CLI using the `svlan vid` command.

Syntax:

`svlan` *`vid`* | *`ascii-name-string`*

`[no] svlan` *`vid`*

If *`vid`* does not exist in the switch, this command creates a port-based S-VLAN with the specified *`vid`*. If the command does not include options, the CLI moves to the newly created S-VLAN context. If you do not specify an optional name, the switch assigns a name in the default format: `svlan` where `n` is the *`vid`* assigned to the S-VLAN. If the S-VLAN already exists and you enter either the *`vid`* or the `ascii-name-string`, the CLI moves to the specified S-VLAN's context.

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

- If one or more ports belong only to the S-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 S-VLAN, there is no `move` prompt.

NOTE: When QinQ is disabled, all VLANs must be C-VLANs. When QinQ is enabled in SVLAN mode, all VLANs must be S-VLANs. When QinQ is enabled in mixed VLAN mode, VLANs can be configured as either C-VLANs or S-VLANs. For more on S-VLAN configuration constraints, see the restrictions on “Operating notes and restrictions” (page 357).

Configuring per-port S-VLAN membership

The `svlan` *`vid`* command supports tagged and untagged options to configure per-port S-VLAN memberships. Use these options from the configuration level by beginning the command with `svlan` *`vid`* , or from the context level of the specific VLAN by entering the command option.

Syntax:

`svlan` *`vid`*

`tagged` *`port-list`*

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

`untagged` *`port-list`*

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

`forbid` *`port-list`*

QinQ S-VLAN mode only. Used in port-based S-VLANs to configure *`port-list`* as forbidden to become a member of the specified VLAN, as well as other actions. The `no` version sets the ports to either `No` or (if GVRP is enabled) to `Auto`. See “GVRP” (page 67).

`auto` *`port-list`*

QinQ S-VLAN mode only. Available if GVRP is enabled on the switch. Returns the per-port settings for the specified S-VLAN to `Auto` operation. `Auto` is the default per-port setting for a static VLAN if GVRP is running on the switch. See “GVRP” (page 67).

NOTE: Since provider-gvrp is not supported in a QinQ mixed VLAN mode environment, the `forbid` and `auto` configurations are available only in QinQ S-VLAN mode. For more information on dynamic VLAN and GVRP operation, see “GVRP” (page 67).

In QinQ mixed VLAN mode

An interface (port or trunk) must be explicitly GVRP-disabled before it can be assigned to the S-VLAN space. When you first attempt to configure a port as tagged for an S-VLAN, the CLI will issue a message disallowing the configuration.

```
config# svlan 200 tagged a1,a2 GVRP enabled ports cannot be members of svlans.
Disable the interface level gvrp configuration.
```

To disable GVRP at the interface, issue the following command:

```
config# interface a1,a2 unknown-vlans disable
```

When you configure the port, the CLI will issue a warning prompt:

```
config# svlan 200 tagged a1,a2 Ports a1, a2 will lose their cvlan memberships if any
Do you want to continue? [y/n]
```

Press **Y** to continue and automatically configure both ports as port-type `provider-network` (the default for all S-VLAN member ports).

Configuring port-types

When QinQ is enabled on the switch all S-VLAN member ports must be categorized as either `port-type customer-network` or `provider-network` (See Figure 75 (page 349)).

Figure 75 Customer or provider ports in the provider network

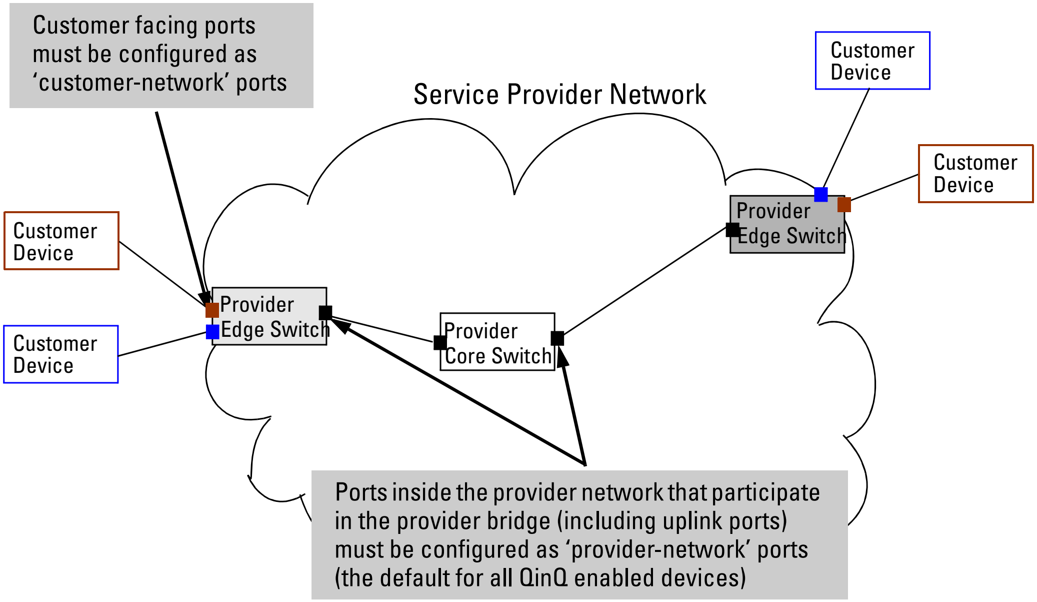
Customer facing ports must be configured as 'customer-network' ports
Service Provider Network
Customer Device
Customer Device
Customer Device
Customer Device
Provider Edge Switch
Provider Core Switch
Provider Edge Switch
Ports inside the provider network that participate in the provider bridge (including uplink ports) must be configured as 'provider-network' ports (the default for all QinQ enabled devices)

All ports of a QinQ-enabled device default to `provider-network`. Any ports participating in the provider bridge that are used to connect to customer equipment, must be manually configured as port-type `customer-network`. In a mixed mode device, ports that are members of C-VLANs and that do not participate in the provider-bridge cannot be configured to any port-type.

The following command allows you to configure the appropriate port-type.

Syntax:

[no] interface *[port-list]* | Trkx
qinq port-type { *customer-network* | *provider-network* }

Configures the specified ports/trunks as a customer network port or provider network port.

Default: port-type provider (for QinQ S-VLAN mode)

Disabling QinQ

To disable QinQ once it has been enabled, issue the following commands from configuration mode on the CLI.

Syntax:

```
no qinq
```

This is the default mode when QinQ is disabled on the switch.

Moving into this configuration from another QinQ configuration requires a reboot to take effect. Upon reboot, all configuration information for the prior QinQ mode will be lost.

Default setting. Standard VLAN operations apply.

Changing VLAN port memberships (mixed VLAN mode)

On mixed VLAN mode devices, certain per-port features are not supported on S-VLANs that are supported on C-VLANs. Ports that are currently members of a regular VLAN can move to an S-VLAN only if there is no conflicting configuration.

NOTE: To avoid a misconfiguration, HP recommends that you use a default interface configuration when moving ports between C-VLANs and S-VLANs.

When configuring S-VLAN port memberships using the `svlan` command, the CLI issues a warning and prompt if any of the ports listed already belong to a regular VLAN. For example:

```
HP Switch(config)# svlan 200 tagged a1,a2
Ports a1, a2 will lose their cvlan memberships if any.
Do you want to continue: y/n?
```

The warning prompt is displayed only when there is at least one port in the port list that needs to be moved out from the C-VLAN space to the S-VLAN domain. Similarly, if ports being added to the C-VLAN are already members of an S-VLAN, the CLI issues a warning that the port's membership with its existing VLANs will be removed and will prompt for a confirmation before continuing.

If all ports are just being added or removed from within the same VLAN type domain, no prompt will appear. For example, moving ports from S-VLAN 200 to S-VLAN 300, will not result in any warning as the ports are already part of the S-VLAN domain.

Moving ports between C-VLANs and S-VLANs (mixed VLAN mode)

A port (or trunk) that is a member of C-VLANs cannot be moved into the S-VLAN space with conflicting configurations for the S-VLAN mode. The following is a list of conflicting protocols/features. If a port has any of these enabled, the feature must be disabled before the port can be moved in to the S-VLAN space.

- An interface has to beGVRP-disabled to move it from the C-VLAN to the S-VLAN space. This is because S-VLANs of mixed VLAN mode do not support provider-GVRP, and also because a GVRP-enabled configuration (when the port is a C-VLAN member) is in the context of customer-GVRP which must be disabled before the port can operate in the S-VLAN space.
- Interface should not have any mirroring or monitoring sessions when moving between C-VLANs and S-VLANs. All mirror/monitor sessions that involve the port must be unconfigured.
- An interface that has auth-vid or unauth-vid configuration cannot move into the S-VLAN space. They have to be unset.
- Interfaces cannot have LACP enabled (active or passive modes) when moving into the S-VLAN space. They have be disabled.

Displaying QinQ configuration and status

This section outlines changes and additions to existing `show` command outputs to display QinQ configuration and status.

The `show qinq` command displays QinQ configuration information.

Syntax:

`show qinq`

Shows QinQ global and port configurations on the switch, including:

Bridge-mode

- **cvlan bridge**: QinQ is disabled, normal VLANs apply.
- **mixedvlan bridge**: Both S-VLANs and regular C-VLANs are available in a mixed VLAN mode environment.
- **svlan**: No regular VLAN commands are available. All VLANs configured on the switch are S-VLANs only.

Tag-id: Displays only if QinQ is enabled on the switch.

port-type: Displays only if QinQ is enabled on the switch. On a mixed mode device, port type is shown only for S-VLAN ports.

Example

Example 208 Displaying `show qinq` output (QinQ S-VLAN mode)

```
HP Switch(config)# show qinq

QinQ Global Configuration:
----------------------------------------------
Bridge-mode        : svlan bridge

QinQ Interface Configuration:
-------------------------------------------------
interface   port-type
---------   ----------
   A1       provider-network
   A2       provider-network
   Trk1     customer-network
```

Displaying a switch VLAN configuration

The following `show` commands are a subset of those listed in the chapter on Static Virtual LANs (VLANs) highlighting the changes made to show the additional QinQ VLAN types (C-VLANs and S-VLANs). For a full listing of all command parameters, see "Static Virtual LANs (VLANs)" (page 14).

The `show vlans` command lists the VLANs currently running in the switch, including the VID, VLAN name, and VLAN status. Once QinQ is enabled in mixed VLAN mode, an additional field showing the VLAN type is added to the display output.

Syntax:

`show vlans`

Changes to parameters when QinQ is enabled:

VLAN ID

Field name changes from 802.1Q VLAN ID to VLAN ID only.

Type

In a QinQ mixed mode environment, the VLAN type can be either a regular customer VLAN CVLAN, or it can be a tunnel VLAN in the provider network S-VLAN.

Example

Example 209 Displaying `show vlans` command output with QinQ disabled

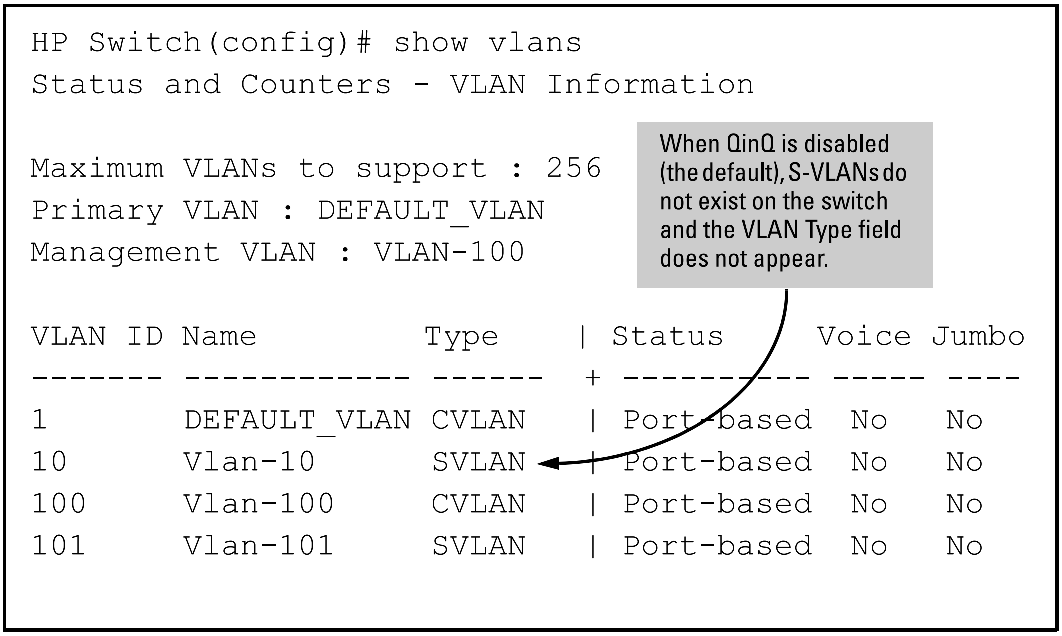

```
HP Switch(config)# show vlans
Status and Counters - VLAN Information

Maximum VLANs to support : 256
Primary VLAN : DEFAULT_VLAN
Management VLAN : VLAN-100

VLAN ID Name          Type     | Status       Voice Jumbo
------- ------------ ------   + ---------- ----- ----
1       DEFAULT_VLAN CVLAN    | Port-based  No    No
10      Vlan-10      SVLAN    | Port-based  No    No
100     Vlan-100     CVLAN    | Port-based  No    No
101     Vlan-101     SVLAN    | Port-based  No    No
```

When QinQ is disabled (the default), S-VLANs do not exist on the switch and the VLAN Type field does not appear.

Displaying the configuration for a particular VLAN

This command uses the VID to identify and display the data for a specific VLAN. Once QinQ is enabled in mixed VLAN mode, an additional field showing the VLAN type is added to the display output.

Syntax:

`show vlans vlan-id`

Changes to parameters when QinQ is enabled:

VLAN ID

Field name changes from 802.1Q VLAN ID to VLAN ID only.

Type

In a QinQ enabled environment, the VLAN type can be either a regular customer VLAN CVLAN, or it can be a tunnel VLAN in the provider network S-VLAN.

Example

Example 210 Displaying `show vlan` output with QinQ enabled

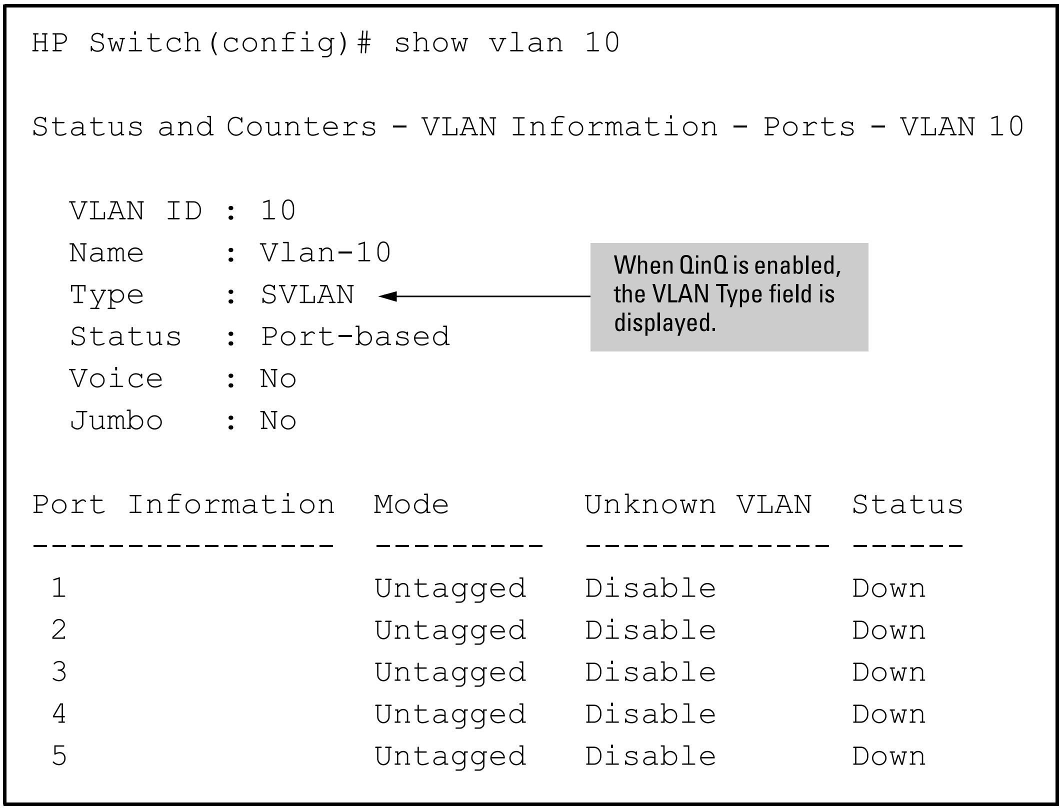

```
HP Switch(config)# show vlan 10

Status and Counters - VLAN Information - Ports - VLAN 10

  VLAN ID : 10
  Name    : Vlan-10
  Type    : SVLAN
  Status  : Port-based
  Voice   : No
  Jumbo   : No

Port Information  Mode       Unknown VLAN  Status
----------------  ---------  ------------- ------
 1                Untagged   Disable       Down
 2                Untagged   Disable       Down
 3                Untagged   Disable       Down
 4                Untagged   Disable       Down
 5                Untagged   Disable       Down
```

When QinQ is enabled, the VLAN Type field is displayed.

Displaying the VLAN membership of one or more ports

This command shows to which VLAN a port belongs. Once QinQ is enabled, an additional field showing the VLAN Type is added to the display output.

Syntax:

`show vlans` *`vlan-id`*

Changes to parameters when QinQ is enabled:

VLAN ID

Field name changes from 802.1Q VLAN ID to VLAN ID only.

Type

In a QinQ enabled environment, the VLAN type can be either a regular customer VLAN CVLAN, or it can be a tunnel VLAN in the provider network S-VLAN.

Example

Example 211 Displaying VLAN membership

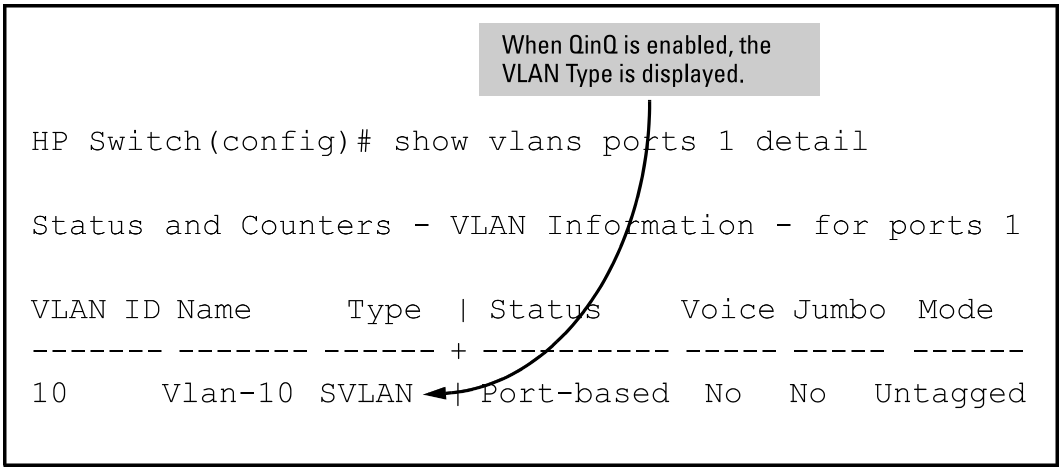

When QinQ is enabled, the VLAN Type is displayed.

```
HP Switch(config)# show vlans ports 1 detail

Status and Counters - VLAN Information - for ports 1

VLAN ID Name      Type  | Status     Voice Jumbo  Mode
------- ------- ------ + ---------- ----- -----  ------
10      Vlan-10  SVLAN  | Port-based  No    No   Untagged
```

Displaying spanning tree status

In QinQ mixed mode, only ports that are members of C-VLANs will be displayed in `show spanning tree` output. This is due to the fact that ports that are members of S-VLANs do not participate in C-VLAN spanning tree and will always be in forwarding state (treated as edge ports).

About QinQ

Operating rules and guidelines

This section provides an overview of QinQ operations and restrictions on the switch. For details of CLI commands and configuration procedures, see "Configuring QinQ (Overview)" (page 343).

Enabling QinQ and configuring QinQ modes

By default, QinQ is disabled. WhenQinQ is enabled via the CLI, an operating mode is globally configured on the switch. Two QinQ modes are supported:

qinq mixedvlan

C-VLANs and S-VLANs are both supported, with regular switching/routing based on C-VLAN tags in the C-VLAN domain, while S-VLANs are used for QinQ tunneling through the provider network.

qinq svlan

C-VLANs are not supported on the device. All configured VLANs on the switch must be S-VLANs.

The following table shows how the various QinQ modes and operations impact VLAN configuration options on the switch.

Table 31 Relationship of QinQ operating modes to VLAN environments

QinQ Operation	CLI Command	VLAN Options
QinQ disabled		
No QinQ support (Default)	`no qinq`	Only regular VLAN commands are available. If QinQ is disabled, S-VLAN commands are not available.

Table 31 Relationship of QinQ operating modes to VLAN environments *(continued)*

QinQ Operation	CLI Command	VLAN Options
QinQ enabled		
QinQ mixed VLANmode	`qinq mixedvlan`	Both S-VLAN and regular VLAN commands (known as C-VLANs in a mixed vlan environment) are available.
QinQ S-VLAN mode	`qinq svlan`	No regular VLAN commands are available. All VLANs configured on the switch are S-VLANs only.

QinQ mixed VLAN mode

The QinQ mixed VLAN mode configuration supports both C-VLAN and S-VLAN operations on the same device. This allows the use of S-VLAN member ports for QinQ tunneling, while regular ports can still do switching or routing within the C-VLAN space. To tunnel customer frames through the provider network, you can externally connect a regular port to a customer-network port, eliminating the need for a separate S-VLAN bridge device to perform such operations. When configuring VLANs on a mixed VLAN mode device, a separate `svlan` *`vid`* command is used to distinguish the S-VLAN type from regular VLANs.

The main advantage for QinQ mixed VLAN mode is that users do not have to dedicate the entire switch as a QinQ access switch. For a high density chassis switch such as the 5400zl or 8200zl series, customers can use regular ports for normal LAN switching, while S-VLAN member ports can be configured to access the QinQ provider network (see Figure 72 (page 344)). There are some additional restrictions in mixed-VLAN mode (see "Operating notes and restrictions" (page 357) for details).

Figure 76 HP Switch in mixed-VLAN mode

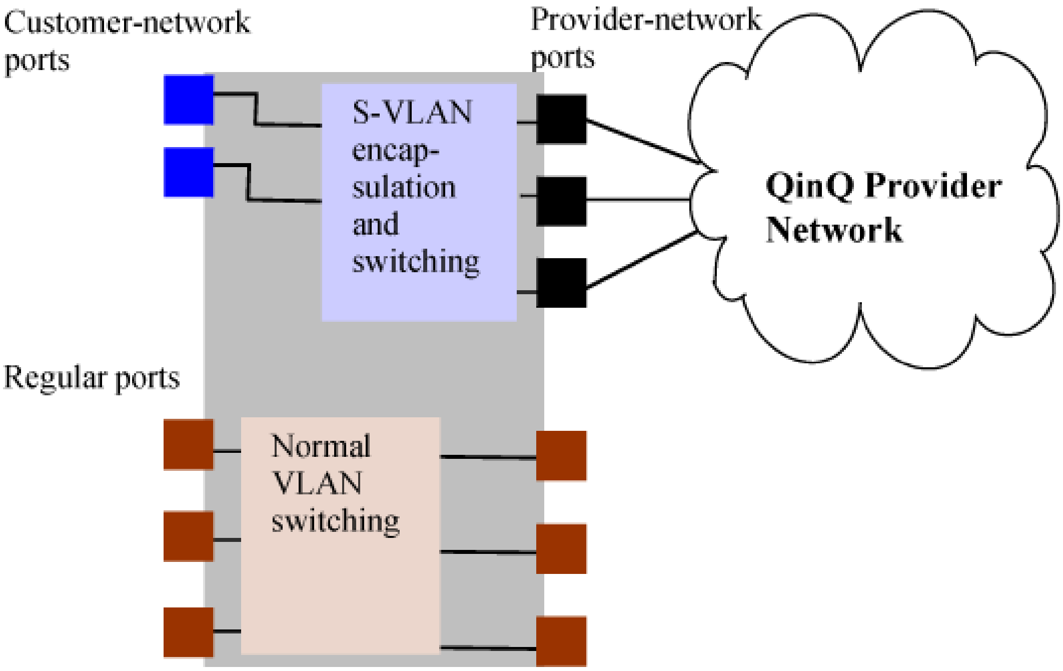

Configuring VLANs

- A VLAN created on a QinQ mixed VLAN mode device can be either a regular VLAN (C-VLAN) or a tunnel VLAN (S-VLAN). C-VLANs have no mapping/relation whatsoever to the S-VLANs on the device.
- VLANs created on a QinQ S-VLAN mode device can be S-VLANs only. S-VLANs provide QinQ tunneling of customer frames and behave like a port-based/s-tagged interface (see "Setting up S-VLANs" (page 347) for configuration details).

QinQ and duplicate VIDs

Duplicate VID's for c-tagged and s-tagged VLANs (for example, C-VID=100; S-VID=100) are allowed in certain cases and disallowed in others. Customer-network ports are essentially S-VLAN ports: they simply read the C-tags in the customer frame to insert them into the appropriate untagged S-VLAN for that port. Once this double-tagging occurs, frames are forwarded based on the S-VLAN tag only, while the C-VLAN tag remains shielded during data transmission. See Figure 77 (page 356) for examples of allowed configurations.

Figure 77 QinQ and duplicate VIDs: examples of allowed configurations

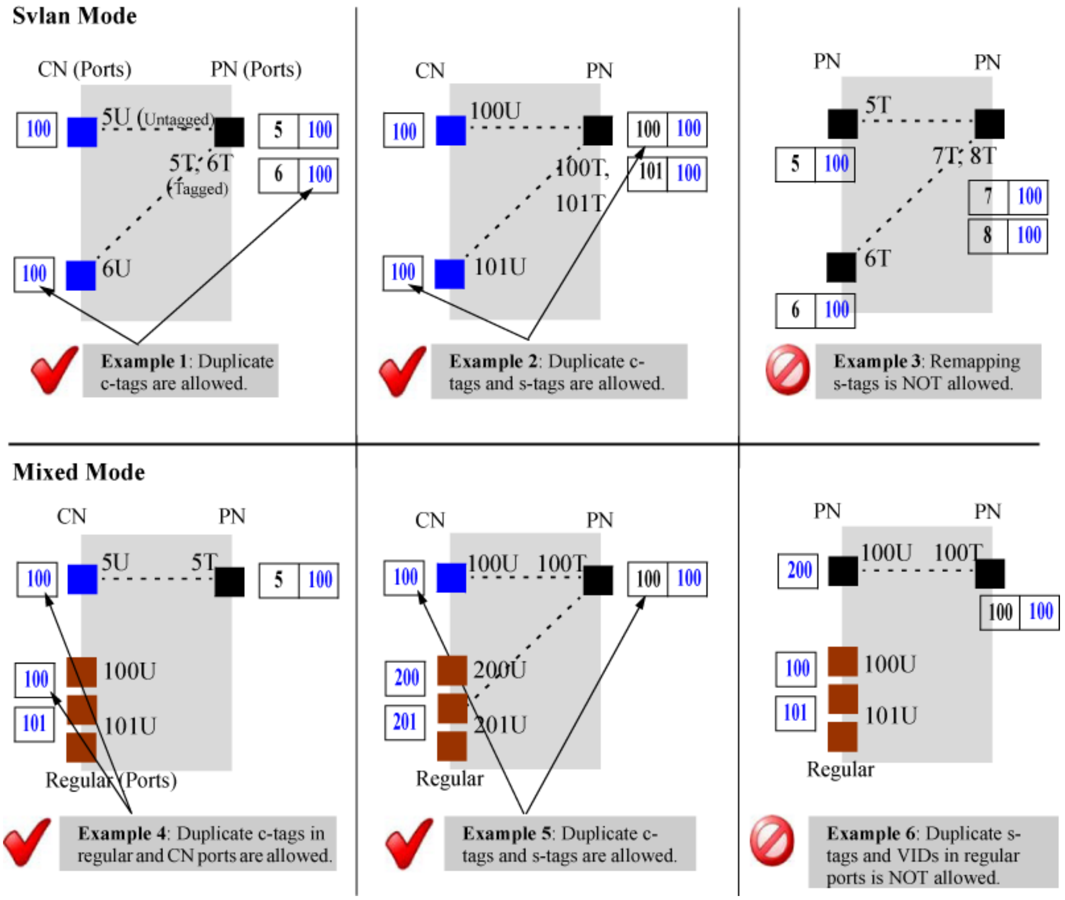
Svlan Mode
CN (Ports)
PN (Ports)
100
5U (Untagged)
5 100
6 100
5T, 6T (Tagged)
100
6U
Example 1: Duplicate c-tags are allowed.
CN
PN
100
100U
100 100
100T, 101T
101 100
100
101U
Example 2: Duplicate c-tags and s-tags are allowed.
PN
PN
5T
5 100
7T, 8T
7 100
8 100
6T
6 100
Example 3: Remapping s-tags is NOT allowed.
Mixed Mode
CN
PN
100
5U
5T
5 100
100
100U
101
101U
Regular (Ports)
Example 4: Duplicate c-tags in regular and CN ports are allowed.
CN
PN
100
100U
100T
100 100
200
200U
201
201U
Regular
Example 5: Duplicate c-tags and s-tags are allowed.
PN
PN
200
100U
100T
100 100
100
100U
101
101U
Regular
Example 6: Duplicate s-tags and VIDs in regular ports is NOT allowed.

Assigning ports to VLANs

In mixed VLAN mode, a port can be a member of a C-VLAN or of an S-VLAN but not both. For details, on assigning membership to provider-based VLANs, see "Configuring per-port S-VLAN membership" (page 348).

Configuring port types

The IEEE 802.1ad standard requires that every S-VLAN member port be configured as either a provider-network or as a customer-network port. In a typical deployment scenario, customer-network ports will be configured as untagged members of S-VLANs while provider-network ports will be configured as tagged members of S-VLANs. Note the following configuration rules and guidelines:

- All ports of a device that is QinQ enabled (in S-VLAN mode or mixed VLAN mode) are provider-network ports by default—if there are any ports that connect to a customer device, they must be manually configured as customer-network ports.
- Configuring a port-type is applicable only if the device is QinQ enabled and the port is a member of an S-VLAN. In QinQ mixed mode, ports that are members of C-VLANs cannot be configured to any port-type.

For more information, see "Configuring port types" (page 356).

NOTE: If a device running in QinQ S-VLAN mode has one or more customer-network ports, it is considered to be a provider edge and not a provider core bridge. This may affect certain operations, such as meshing, UDLD, and stacking. This is because at the edge of the provider network such proprietary protocol are filtered out at customer network ports. This prevents the intermix of stacking meshing/UDLD protocols in the customer and provider domains (since they use the same `dst-mac` address in either domain).

Operating notes and restrictions

Cannot run concurrently with RPVST+

QinQ cannot run concurrently with RPVST+

Changing bridge modes requires a reboot

When changing the operating mode (to/from: QinQ S-VLAN mode, QinQ mixed VLAN mode, or QinQ disabled), you will prompted to restart the system before the changes can take effect. Upon reboot, all configuration information for the prior QinQ mode will be lost. Any configurations created will be erased, and the device will boot up with a default configuration for the new QinQ mode.

Provider edge devices at Layer 2 only

QinQ does not provide Layer 3 capabilities of complete network isolation between customers. In a mixed VLAN configuration, there is no switching/routing between C-VLANs and S-VLANs. S-VLANs are essentially Layer 2 VLANs that switch packets based on S-VIDs.

IP support

Regular VLANs support IP and can be routing enabled. S-VLANs of mixed VLAN mode devices cannot be ip enabled. S-VLANs of S-VLAN mode devices can be ip-enabled, though routing related features (such as ip routing) are not supported.

Double-tagging causes frame size increases

Since there is both a provider VLAN tag and customer VLAN tag in each QinQ frame, the size of each double-tagged frame increases by 4 bytes. To accommodate the frame size increase, HP recommends that you configure all port-based S-VLANs to accept jumbo frames. See the section on Jumbo Frames in the Management and Configuration Guide for details.

S-VLAN configuration restrictions

S-VLAN commands are not available when QinQ is disabled on the switch.

VLAN configuration restrictions in mixed VLAN mode

- Both C-VLANs and S-VLANs can be configured on the switch. In a mixed mode device, the default VLAN is always a C-VLAN.
- VLAN types cannot be updated dynamically. A VLAN can be classified only as an S-VLAN or a C-VLAN at the time its created. Once created, the VLAN cannot be moved between being a C-VLAN and an S-VLAN. If a VID that was initially created as a regular VLAN needs to be used for an S-VLAN, the VID must be deleted and re-created as an S-VLAN.

- If a VLAN being configured as an S-VLAN already exists as a GVRP C-VLAN or a static C-VLAN on the switch, the S-VLAN creation is blocked. Similarly, a C-VLAN creation is blocked if the same VID exists as a static S-VLAN on the device.
- S-VLANs in a mixed vlan device cannot be configured as avoice-VLAN, primary-VLAN, or management-VLAN.
- S-VLANs cannot be configured with ip-layer functionality, except for ip-acls.

VLAN configuration restrictions in S-VLAN mode

- Only S-VLANs are supported—the keyword on all vlan-related command syntax changes from `vlan` to `svlan`.
- Routing related features such as ip-routing, RIP, OSPF, PIM, and VRRP are not supported in S-VLAN mode.

Port-based restrictions

- In QinQ mixed VLAN mode, a port must be explicitly GVRP-disabled before it can be assigned to the S-VLAN space (see "In QinQ mixed VLAN mode" (page 349) for details).
- In QinQ mixed VLAN mode, only ports that are members of S-VLANs can be configured as customer network or provider network ports; ports that are members of C-VLANs cannot be configured to any port-type.
- QinQ mixed VLAN mode devices cannot be connected in an S-VLAN mesh topology. This is because STP cannot be run in the S-VLAN space, and so a mesh topology (or the presence of any redundant links) would result in loops.
- A port can either be a member of S-VLANs or C-VLANs only, but not a combination of both.
- A port cannot be configured as a Customer-Edge as specified in Section 12.13.3 of the IEEE 802.1ad specification. In the current software release, such C-tagged interfaces are not supported—only port-based/S-tagged interfaces are supported.
- Moving ports between C-VLANs and S-VLANs may cause conflicts. For example, if a port has any mirroring/monitoring sessions set up, they will not be allowed to change VLAN domains until these sessions are unconfigured. See "Changing VLAN port memberships (mixed VLAN mode)" (page 350) for additional details.

Interoperating with other vendor devices

When enabling QinQ, you can configure a unique tpid value, such as 0x8100, to allow the device to interoperate with devices that require this value for the inner and outer VLAN-tag. If the provider tag-type is configured as 0x8100, then:

- Customer-network ports cannot be configured as tagged-S-VLAN members
- Tagged-S-VLAN members cannot be configured as customer-network ports.

Configuring QinQ with other network protocols

The networks for both the customer and provider can be complex. For information on how QinQ may impact other network protocols (such as spanning tree, LLDP, and GVRP), see Figure 76 (page 355)

Changing QinQ modes

Changing QinQ modes (and/or disabling QinQ operations) will result in the current configuration being erased. See the following Caution for details.

△ **CAUTION:** Configuring the switch to operate in a different bridge mode requires a reboot to take effect. Upon reboot, all configuration information for the prior QinQ mode will be lost. Any configurations created under the existing QinQ mode will be erased, and the device will boot up with a default configuration for the new QinQ mode.

For information on the effect of the different QinQ modes on switch protocols and operations,see Table Table 32 (page 359).

Effects of QinQ on other switch features

Per the IEEE standards, protocols such as STP and GVRP are assigned separate addresses for customer networks and provider networks, ensuring that QinQ has no impact on their operations. Bridge Protocol Data Units (BPDUs) that need to be tunneled through the provider network are treated as normal multicast frames at the provider bridge and forwarded out.

However, other protocols use common addresses for both customer and provider networks, and so are not supported when QinQ is enabled on the switch. Similarly, proprietary features such as meshing, discovery, UDLD, and loop-protect do not provide tunneling support. In such cases, where provider networks could run an instance of the same protocol as a customer could run local to their site, these frames are dropped at the customer-network ports of the provider bridge.

NOTE: The IEEE standards group is devising new addressing schemes that may support additional QinQ tunneling operations. Check the latest product release notes for implementation updates as they apply to HP switches.

When QinQ is not enabled (the default setting), there are no impacts to the switch's normal operations. The following table shows the impacts of QinQ on the operation of switch protocols and features based on the QinQ mode that is configured as QinQ mixed VLAN mode (C-VLANs and S-VLANs are allowed) or QinQ S-VLAN mode (S-VLANs only).

Table 32 Impacts of QinQ configurations on other switch features

Switch feature	Impacts of QinQ configurations and allowed operations
ACLs	In QinQ mixed VLAN or S-VLAN modes: • On double-tagged frames , the VID applicable when applying ACLs will be the S-VLAN tag and not the C-VLAN tag.
aaa	In QinQ mixed VLAN mode: • auth-vid/unauth-vid configuration is not supported on S-VLAN ports; the auth-vid/unauth-vid cannot be an S-VLAN id. • If a port that is a member of C-VLANs is configured with auth-vid or unauth-vid and it needs to be added to the S-VLAN domain, the auth/unauth configuration must first be undone.
arp-protect	In QinQ mixed VLAN mode: • ARP-protect is not supported on S-VLANs, nor on S-VLAN ports.
CDP	In QinQ VLAN or S-VLAN modes: • CDP frames are consumed at customer network ports, if CDP is enabled on the device port, and the customer device shows up as a CDP neighbor on the customer-network port. If not, the frames are dropped.
DHCP	In QinQ mixed VLAN or S-VLAN modes: • DHCP relay applies only to C-VLANs. • DHCP snooping is not supported on S-VLANs.
directed-broadcast	In QinQ S-VLAN mode: • directed-broadcast is not supported on provider core devices.
GVRP	In QinQ mixed VLAN mode: • S-VLAN ports cannot be GVRP enabled. • Regular VLANs will participate in C-VLAN GVRP if enabled to do so. S-VLANs will tunnel all C-VLAN GVRP frames through. • An explicit GVRP disable on a port is a prerequisite for moving the port to an S-VLAN domain.

Table 32 Impacts of QinQ configurations on other switch features *(continued)*

Switch feature	Impacts of QinQ configurations and allowed operations
	• Port-based interfaces do not have support for provider-GVRP protocols. Provider GVRP frames received at S-VLAN interfaces will be dropped. • If a VLAN being configured as an S-VLAN is already a GVRP VLAN on the switch, this S-VLAN creation would be blocked. In QinQ S-VLAN mode: • GVRP is supported on S-VLAN ports if the qinq mode is S-VLAN.
igmp-proxy	In QinQ mixed VLAN mode: • IGMP-proxy cannot be configured on S-VLANs. In QinQ S-VLAN mode: • IGMP-proxy is not supported.
IPv6	In QinQ mixed VLAN mode: • IPv6 features are not supported on S-VLANs.
ip-recv-mac	In QinQ mixed VLAN mode: • `ip-recv-mac` cannot be configured on S-VLANs. In QinQ S-VLANmode: • `ip-recv-mac` is not supported.
Jumbo	In QinQ mixed VLAN or S-VLAN modes: • No change in operations. HP recommends to jumbo-enable all S-VLANs used for customer data tunneling to support the addition of the extra S-tag in each frame.
LACP/ Port Trunks	In QinQ mixed VLAN mode: • Dynamic-LACP is not supported on S-VLAN ports: LACP manual trunks alone are supported. The new trunk will be a member of C-VLANs (port types are not applicable). • If two ports are added to a trunk, the resultant trunk will be a member of the default-vlan (vid-1) which is always a C-VLAN. The trunk can subsequently be manually assigned to an S-VLAN. • Port-type and VLAN configurations are not mapped. If the port-type is updated through CLI or SNMP and the port is subsequently moved from the C-VLAN space to the S-VLAN space then back again, the last configured port-type is retained through each move. In QinQ S-VLAN mode: • On S-VLAN bridges, both manual and dynamic LACP trunks are supported. HP does not recommend that you configure dynamic trunks on customer ports because they cannot become dynamic members of S-VLANs (there is no provider-gvrp for a dynamic trunk to become a member of S-VLANs.) • A newly formed trunk will by default be of type provider-network. When the trunk is manually assigned to an S-VLAN for the first time after being created, the port-type is provider-network.
Layer 3 Protocols (IP, IP+, DHCP, ARP, IGMP Layer 3, Layer 3 ACLs)	In QinQ mixed VLAN mode: • There is no IP layer functionality on S-VLANs. • No change in IP layer functionality on regular C-VLANs. • S-VLANs cannot be configured as RIP, OSPF, PIM, or VRRP interfaces. In QinQ S-VLAN mode: • S-VLANs can be ip enabled. • IP routing is not supported.
LLDP	In QinQ mixed VLAN or S-VLAN modes: • LLDP is supported on the device (in both qinq modes). However, there is no provision for tunneling customer LLDP BPDUs through the provider-network. • LLDP BPDUs received from a customer's network will be consumed at the customer-network ports of a provider device and the customer device will be displayed as an LLDP neighbor. Similarly

Table 32 Impacts of QinQ configurations on other switch features *(continued)*

Switch feature	Impacts of QinQ configurations and allowed operations
	the provider network device will show up as a neighbor on the customer's network if the customer-network ports send out LLDP advertisements.
load-sharing	In QinQ S-VLAN mode: • Equal cost multi-path (ECMP) is not supported on provider core devices.
management VLAN	In QinQ mixed VLAN mode: • The management VLAN cannot be an S-VLAN.
Meshing	In QinQ mixed VLAN mode: • Meshing is not supported on the device. In QinQ S-VLAN mode: • On an all provider-network ports of an S-VLAN bridge, meshing is supported. • Meshing cannot be enabled on customer-network ports.
Mirroring/Monitoring	In QinQ mixed VLAN mode: • Remote mirroring is not supported on S-VLANs. • Cannot monitor a VLAN with mirror ports in the other VLAN domain. That is, an S-VLAN or an S-VLAN port cannot be monitored using a C-VLAN port as its mirror, and vice-versa. • When a port is moved from the S-VLAN space to the C-VLAN space (or vice versa), all mirror/monitor sessions on the port must be unconfigured before the move will be allowed.
multicast-routing	In QinQ S-VLAN mode: • Multicast routing is not supported on provider core devices.
QoS	In QinQ mixed VLAN or S-VLAN modes: • HP does not recommend that you enable DSCP on S-VLANs used for tunneling as the customer IP-pkt will be modified in the S-VLAN space.
Routing	In QinQ S-VLAN mode: • Routing is not supported on provider core devices.
source-binding	In QinQ mixed VLAN or S-VLAN modes: • source-binding cannot be configured on S-VLANs.
source-route	In QinQ S-VLAN mode: • source-route is not supported on provider core devices.
Spanning Tree	In QinQ mixed VLAN mode: • Customer (C-VLAN) spanning tree is supported. All C-VLAN ports will receive/transmit customer STP BPDUs and participate in regular VLAN spanning tree as usual. • When customer STP BPDUs are received at S-VLAN ports on the switch, they will be flooded out of the other ports on the S-VLAN. All such frames will be tunneled through the S-VLAN tunnel unscathed. • Provider (S-VLAN) spanning tree is not supported on the switch. If S-VLAN STP frames are received on any S-VLAN enabled ports, they will be re-forwarded out of the other ports on the S-VLAN. • STP configuration on S-VLAN ports is not supported. • If a port that is a member of C-VLANs is moved into being a member of S-VLANs, the port would, by default, tunnel customer STP BPDUs. • If a C-VLAN port has been configured with any non-default STP parameters (such as `admin-edge, auto-edge`, and `bpdu-protect`) and is then moved into an S-VLAN, the port will be put into a forwarding state regardless of the STP configurations done when the port was a member of the C-VLAN. • MSTP instances cannot include S-VLANs.

Table 32 Impacts of QinQ configurations on other switch features *(continued)*

Switch feature	Impacts of QinQ configurations and allowed operations
	In QinQ S-VLAN mode: • Provider (S-VLAN) spanning tree is supported—both provider-network ports and customer-network ports will receive/transmit provider STP BPDUs. • Customer (VLAN) spanning tree tunneling is supported on S-VLAN interfaces—customer-network or provider-network ports will tunnel customer STP BPDUs through the appropriate S-VLAN.
Stacking	In QinQ mixed VLAN mode: • Stacking is supported only on C-VLANs. The device does not advertise itself (using the stack discovery protocol) in the S-VLAN space. In QinQ S-VLAN mode: • Stacking discovery protocol frames will not be sent out of customer-network ports; similarly, any stacking discovery protocol frames received on customer-network ports will be dropped.
UDLD	In QinQ mixed vlan or S-VLAN modes: • UDLD frames received on udld-disabled customer network ports will be dropped. However, if the customer-network port is udld-enabled, it can peer with a customer device. • UDLD frames received on udld-disabled provider network ports will be re-forwarded out of other udld-disabled provider network ports on the same VLAN. • UDLD re-forwarding in the C-VLAN space (QinQ disabled or mixed VLAN mode) will remain unaltered.
udp-bcast-forward	In QinQ S-VLAN mode: • `udp-bcast-forward` is not supported on provider core devices.
unknown-vlans	In QinQ mixed VLAN mode: • GVRP (learn and disabled modes) not supported on S-VLAN ports. • A C-VLAN port that has GVRP enabled will need to disable it before it can be added to S-VLANs.
Voice VLANs	In QinQ mixed VLAN mode: • S-VLANs cannot be configured as voice-VLANs.
VRRP	In QinQ mixed VLAN or S-VLAN modes: • VRRP is not supported on S-VLANs.

Event log messages

See the *Event Log Message Reference Guide* for information about Event Log messages.

10 Classifier-based software configuration

Command syntax	Description	Default value	CLI page reference
[no] class [ipv4 | ipv6] *classname*	Configures a traffic class		364
[no] [*seqnumber*] [match | ignore] icmp *sourceaddress destinationaddress* [*icmptypenumber* | *icmpv4typename* | *icmpv6typename*] [ipdscp *codepoint*] [precedence *precedencevalue*] [tos *tosvalue*] [vlan *vlanid*]	Defines the ICMP match criteria		369
[no] [*seqnumber*] [match | ignore] igmp *sourceaddress destinationaddress* [*igmptype*] [ipdscp *codepoint*] [precedence *precedencevalue*] [tos *tosvalue*] [vlan *vlanid*]	Defines the IGMP match criteria		371
[no] [seqnumber] [match | ignore] [tcp | udp] *sourceaddress* [*operator tcpsrcport* | *udpsrcport*] *destinationaddress* [*operator tcpdestport* [established] [tcpflag tcpflag ...] udpdestport] [ipdscp *codepoint*] [precedence *precedencevalue*] [tos *tosvalue*] [vlan *vlanid*]	Defines the TCP and UDP match criteria		372
class resequence [ipv4 | ipv6] *name seq-number interval*	Resequences match/ignore statements		376
[no] policy *feature-name policy-name*	Creates a service policy		377
policy resequence *name seq-number interval*	Resequences classes in a policy	10	382
interface *port-list* service-policy *policy-namein*	Applies a service policy to an interface		382
show policy resources	Checks resource usage		384
[no] class zone *zone name*	Creates a zone class		387
[no] policy zone *policy-name*	Creates a zone policy		388v
[no] zone-service-policy *policy-name* zone [enable | disable | update] bind [Ethernet] *logical port* appname *application name* appinstance *instance description* [fail-action [bypass | block]] [expire [app-down | permanent | slot-down]]	Applies a zone policy to a ONE application		390

Introduction

Classifier-based service policies are designed to work with existing globally configured switch-wide and port-wide settings by allowing you to select a subset of:

- Traffic sent to or from certain ports
- VLAN traffic

Once the traffic is selected, you can further manage it.

Classifier-based service policies take precedence over, and may override, globally configured settings. These policies provide greater control for managing network traffic. Using multiple match criteria, you can finely select and define the classes of traffic that you want to manage. You can then use policy actions to determine how the selected traffic is handled.

Classes can be based on IPv4 or IPv6 addresses (which you specify in the policy). For information about traffic classes, see "Introduction" (page 363).

When using an HP AllianceONE Extended Services zl Module that supports Transparent Mode, you can also classify traffic based on zones. For information about zone classes, see "Configuring class-based zones" (page 386).

Configuring a traffic class

To configure a traffic class to be used in one or more policies, follow these steps:

1. Enter the `class` command from the global configuration context.

   Context: Global configuration

   ### *Syntax:*

   [no] `class [ ipv4 | ipv6 ]` *`classname`*

   Defines a traffic class and specifies whether a policy is to be applied to IPv4 or IPv6 packets, where *`classname`* is a text string (64 characters maximum).

   After you enter the `class` command, you enter the class configuration context to specify match criteria. A traffic class contains a series of `match` and `ignore` commands, which specify the criteria used to classify packets.

   The `no` form of the command removes the existing class

2. Enter one or more `match` or `ignore` commands from the traffic class configuration context to filter traffic and determine the packets on which policy actions will be performed.

   Context: Class configuration

   ### *Syntax:*

   [no] [*`seq–number`*] `[ match | ignore ] igmp` *`source–address destination–address`* `[` *`igmp–type`* `] [ ip–dscp` *`codepoint`* `] [ precedence` *`precedence–value`* `] [ tos` *`tos–value`* `] [ vlan` *`vlan–id`* `]`

   `seq-number`

   (Optional) Sequentially orders the match/ignore statements that you enter in a traffic class configuration. Packets are checked by the statements in numerical order.

   Default: Match/ignore statements are numbered in increments of 10, starting at 10. To re-number the match/ignore statements in a class configuration, use the `resequence` command. (See "Resequencing match/ignore statements" (page 376)).

   `match | ignore`

   Defines the classifier criteria used to determine which packets belong to the traffic class.

   If a packet matches a `match` criterion, it becomes a member of the traffic class and is forwarded according to the actions configured with the `policy` command. If a packet matches an `ignore` criterion, no policy action is performed on the packet. You can enter one or more match/ignore statements in a traffic class.

   To remove a match/ignore statement from a class configuration, enter the `no` *`seq-number`* command or the complete form of a `no match` or `no ignore` command.

   `ip-protocol`

   Specifies an IP protocol to be matched in packet fields of IPv4 or IPv6 traffic, where *`ip-protocol`* is one of the values described below.

When entering a match/ignore command in an IPv4 or IPv6 class, enter `?` to display a list of valid *ip-protocol* entries.

- In an IPv4 class, you can enter any of the following IPv4 protocol match criteria:

  ahesp
  gre
  icmp[1]
  igmp[1]
  ipip-in-ip
  ipv6–in-ip
  ospf
  pim
  sctptcp[1]
  udp[1]
  vrrp

  To specify an IPv4 protocol as match criteria, you can also enter its protocol number. Valid values are from 0 to 255.

  For example, 8 means Exterior Gateway Protocol; 121 means Simple Message Protocol. For a list of IPv4 protocol numbers and corresponding protocol names, see the IANA "Protocol Number Assignment Services" at www.iana.com.

- In an IPv6 class, you can enter any of the following IPv6 protocol match criteria:

  ahesp
  icmp[2]
  ipv6
  sctptcp[2]
  udp[2]

source-address destination-address

Defines the source IP address (SA) and destination IP address (DA) that a packet must contain to match a match/ignore statement in an IPv4 or IPv6 traffic class. Both the source and destination address parameters are required entries in a match/ignore statement.

Valid values for *source-address* and *destination-address* are as follows:

- `any`: Matches IPv4 or IPv6 packets from, or destined to, any SA or DA.
- `host [ SA | DA ]`: Matches only packets from a specified IPv4 or IPv6 host address. Use this match criterion when you want to match IP packets from only one SA/DA.
- *SAv4 mask | DAv4 mask*: Matches packets received from, or destined to, a subnet or a group of IP4 addresses defined by the IPv4 mask. Enter an IPv4 mask in dotted-decimal format for an IPv4 address (for example, 10.28.31.1 0.0.0.255).

  See “How IPv4 mask bit settings define a match (Example)” (page 374).

  NOTE: An IPv6 address and mask are not supported as `SAv6 mask` and `DAv6` mask match criteria.

- *SAv4/mask-length | DAv4/mask-length*: Matches packets received from, or destined to, an IPv4 subnet or a group of IPv4 addresses defined by the mask length.

[1] For IPv4 ICMP, IGMP, TCP, and UDP packets, you can enter additional match criteria; see:
"Defining the ICMP match criteria"
"Defining the IGMP match criteria"
"Defining TCP and UDP match criteria"

[2] For IPv6 ICMP, TCP, and UDP packets, you can enter additional match criteria; see:“Defining the ICMP match criteria” (page 369), “Defining the IGMP match criteria” (page 371), “Defining TCP and UDP match criteria” (page 372).

Enter the mask length for an IPv4 SA or DA mask in CIDR format by using the number of significant bits. (for example, 10.28.31.3/24).

An IPv4 mask-length is applied to an SA or DA in a match/ignore statement to define which bits in a packet's SA/DA must exactly match the specified SA/DA and which bits need not match. For example, 10.28.31.3/24 means that the leftmost 24 bits in an IPv4 source or destination address in a packet header must match the same bit set in the specified IPv4 address (in this case, 10.28.3.3). For more information, see "How IPv4 mask bit settings define a match (Example)" (page 374).

An IPv4 mask-length is applied from right to left, starting from the rightmost bits. For example, 10.10.10.1/24 and 10.10.10.1 0.0.0.255 both match IPv4 addresses in the range 10.10.10.(1 to 255).

NOTE: Specifying a group of non-contiguous IP source addresses may require more than one match/ignore statement.

- `SAv6/prefix-length` | `DAv6/prefix-length`: Matches packets received from, or destined to, an IPv6 subnet or a group of IPv6 addresses defined by the prefix length. Enter the prefix length for an IPv6 SA/DA in CIDR format by using the number of significant bits; for example: 2001:db8:2620:212::01b4/64.

  An IPv6 prefix-length is applied to an SA/DA in a match/ignore statement to define which bits in a packet's SA/DA must exactly match the specified SA/DA and which bits need not match. For example, 2001:db8:2620:212::01b4/64 means that the leftmost 64 bits in a 128-bit IPv6 source or destination address in a packet header must match the same bit set in the specified IPv6 address (in this case, 2001:db8:2620:212::01b4).

  For more information, see "How IPv4 mask bit settings define a match (Example)" (page 374).

  An IPv6 prefix-length is applied from left to right, starting from the leftmost bits. For example, 2001:db8::0001: 2620:a03:e102:127/64 and 2001:db8::1: 244:17ff:feb6:d37d/64 both match IPv6 addresses with a network prefix of 2001:db8:0000:0001.

`ip-dscp codepoint`

(Optional) Matches the six-bit DSCP codepoint DSCP codepoint in IPv4 or IPv6 packets to further define match criteria. Valid values for `codepoint` are one of the following:

- Numeric equivalent of a binary DSCP bit set from `0` (low priority) to `63` (high priority)
- ASCII standard 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
af41 (100010) cs7 (111000) = precedence 7
default (000000)
```

To display a list of valid `codepoint` entries when you enter `ip-dscp` in a match/ignore statement, enter **?**.

The DSCP codepoints are the leftmost six bits of the ToS/Traffic Class byte (See "A ToS/traffic class field" (page 368)).

`precedence` *`precedence-value`*

(Optional) Matches the three-bit IP precedence value in IPv4 or IPv6 packets to further define match criteria. Valid values for `precedence-value` are either the numeric value (0 to 7) or corresponding name of an IP precedence bit set:

0 routine
1 priority
2 immediate
3 flash
4 flash-override
5 critical
6 internet (for internetwork control)
7 network (for network control)

To display a list of valid `precedence-value` entries when you enter `precedence` in a match/ignore statement, enter **`?`**.

NOTE: When used as a match criteria, the IP precedence value is applied in addition to all other criteria configured in the match/ignore statement. You can enter a match/ignore statement either with or without a `precedence-value`.

The IP precedence bits are the leftmost three bits of the ToS/Traffic Class byte (see "A ToS/traffic class field" (page 368)). The numeric value (0 to 7) of the IP precedence bits corresponds to the hexadecimal equivalent of the three binary 0 and 1 bits in the IP precedence field. For example if the IP precedence-bit binary values are `1 1 1`, the numeric value is `7` (1+2+4). Similarly, if the IP precedence bits are `0 1 0`, the numeric value is `2` (0+2+0).

`tos` *`tos-value`*

(Optional) Matches the Delay Throughput Reliability (DTR) bit set in the IPv4 Type-of-Service or IPv6 Traffic Class byte to further define match criteria.

Valid values are the numeric value or corresponding name of the DTR bit set. Some useful values are as follows:

0 — normal
2 — max-reliability
4 — max-throughput
8 — minimize-delay

Default: 0 or `normal`.

To display a list of valid `tos-value` entries when you enter `tos` in a match/ignore statement, enter **`?`**.

NOTE: When used as a match criteria, the ToS/Traffic Class byte entry is applied in addition to all other criteria configured in the match/ignore statement. You can enter a match/ignore statement either with or without a `tos-value`.

"A ToS/traffic class field" (page 368) shows the DTR bit set in a ToS/Traffic Class byte in an IPv4/IPv6 packet header and the difference between the DSCP, DTR, and precedence bits.

`vlan` *`vlan-id`*

(Optional) Matches the VLAN ID number in the Layer 2 header of 802.1Q VLAN packets to further define match criteria. Valid VLAN IDs are from 1 to 4094.

"A ToS/traffic class field" (page 368) uses a sample ToS/Traffic Class field of `10101000` to show the differences between the IP precedence (`101`), DSCP (`101010`), and ToS/Traffic Class (`10101000`) bits. The rightmost two bits are reserved as `00`.

Figure 78 A ToS/traffic class field

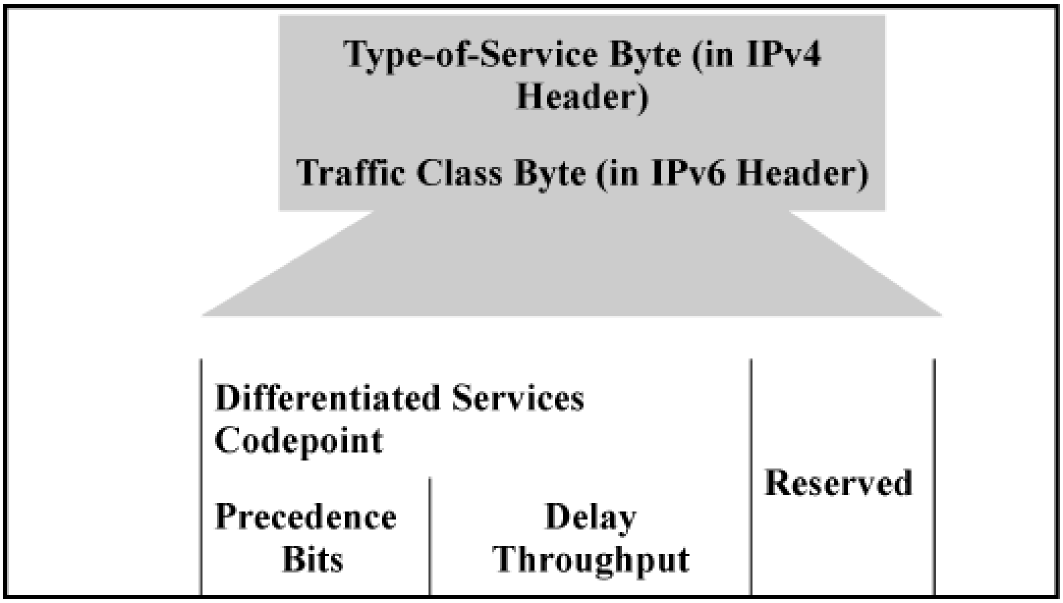
Type-of-Service Byte (in IPv4 Header)
Traffic Class Byte (in IPv6 Header)
Differentiated Services Codepoint
Precedence Bits
Delay Throughput
Reserved

3. A ToS/traffic class field.

To display a class configuration, enter the following command (see "Resequencing a class configuration" (page 377)):

```
show class  [ ipv4 | ipv6 ][classname]
```

To edit a class configuration, re-enter the class configuration context (`class` command) and enter new match/ignore statements as follows:

- If you do not enter a sequence number, a new statement is inserted at the end of the class configuration.
- To remove a match/ignore statement from a class configuration, enter the `no` *`sequence-number`* command or the complete form of the `no match` or `no ignore` command.
- To `resequence` the order in which match/ignore statements are listed, include the `resequence` option in the class command. (See "Resequencing match/ignore statements" (page 376)).
- To replace an existing match/ignore statement, enter the `no` *`sequence-number`* command to delete the entry and re-enter a complete *`sequence-number`* `match` or *`sequence-number`* `ignore` command.

When exiting the class configuration context, the changes are automatically saved and applied to existing policy configurations on the switch that use the class if the policies have not been applied to an interface. If a policy has already been applied to an interface, the editing changes are not accepted and an error message is displayed.

Example 212 Class configurations

The following example shows two class configurations:

- `AdminTraffic` selects the administrative traffic sent to, and received from, the IPv4 address of an administrator's PC.
- `http` selects HTTP traffic sent to TCP ports 80, 443, and 8080, and excludes HTTP traffic sent to, and received from, TCP port 1214.

```
HP Switch(config)# class ipv4 AdminTraffic
HP Switch(config-class)# match ip 15.29.16.1/10 any
HP Switch(config-class)# match ip any 15.29.16.1/10
HP Switch(config-class)# exit
HP Switch(config)# class ipv4 http
HP Switch(config-class)# match tcp any any eq 80
HP Switch(config-class)# match tcp any any eq 443
HP Switch(config-class)# match tcp any any eq 8080
HP Switch(config-class)# ignore tcp any eq 1214 any
HP Switch(config-class)# ignore tcp any any eq 1214
HP Switch(config-class)# exit
```

Defining the ICMP match criteria

To more precisely define the ICMP packets to match in an IPv4 or IPv6 traffic class, use the optional parameter settings below. For example, instead of matching or ignoring all ICMP traffic, you can configure a class that matches only a specific ICMP packet type by entering its numeric value.

Context*:* Class configuration

Syntax:

[no] [*seq–number*] [`match` | `ignore`] [icmp] *`source–address destination–address`* [*`icmp–type–number`* | *`icmpv4–type–name`* | *`icmpv6–type–name`*] [`ip–dscp` *`codepoint`*] [`precedence` *`precedence–value`*] [`tos` *`tos–value`*] [*`vlan–id`*]

If you enter `icmp` as the IP protocol type in a match/ignore statement, you can optionally specify an ICMP packet type to more precisely define match criteria for a traffic class. Enter the optional ICMP match criteria immediately after the destination address (DA) value in the command syntax; for example:

```
HP Switch(config-class)# match icmp any any host-unknown
HP Switch(config-class)# match icmp any any 3 7
```

`icmp-type-number`

Configures an ICMP packet type as match criteria in a class configuration by entering its numeric identifier. Valid values are from `0` to `255`.

For information on ICMP packet-type names and numeric identifiers, go to the Internet Assigned Numbers Authority (IANA) website at www.iana.com, click **Protocol Number Assignment Services**, and then go to the selections under **Internet Control Message Protocol (ICMP) Parameters**.

`icmpv4-type-name`

Enter any of the following ICMPv4 packet-type names to configure more precise match criteria for ICMP packets in an IPv4 class configuration.

To display a list of valid `icmpv4-type-name` entries when entering `icmp` as the IP protocol type in a match/ignore statement, enter **?**. Some of the valid values are:

- `administratively-prohibitednet-tos-unreachable`
- `alternate-addressnet-unreachable`
- `conversion-errornetwork-unknown`
- `dod-host-prohibitedno-room-for-option`
- `dod-net-prohibitedoption-missing`
- `echopacket-too-big`
- `echo-replyparameter-problem`
- `general-parameter-problemport-unreachable`
- `host-isolatedprecedence-unreachable`
- `host-precedence-unreachableprotocol-unreachable`
- `host-redirectreassembly-timeout`
- `host-tos-redirectredirect`
- `host-tos-unreachablerouter-advertisement`
- `host-unknownrouter-solicitation`
- `host-unreachablesource-quench`
- `information-replysource-route-failed`
- `information-requesttime-exceeded`
- `mask-replytimestamp-reply`
- `mask-requesttimestamp-request`
- `mobile-redirecttraceroute`

- `net-redirectttl-exceeded`
- `net-tos-redirectunreachable`

`icmpv6-type-name`

You can also enter any of the following ICMPv6 packet-type names to configure more precise match criteria for ICMP packets in an IPv6 class configuration.

To display a list of valid `icmpv6-type-name` entries when you enter `icmp` as the IP protocol type in a match/ignore statement, enter **?**. Some of the valid values are as follows:

- `cert-path-advertisemobile-advertise`
- `cert-path-solicitmobile-solicit`
- `destination-unreachablend-na`
- `echo-replynd-ns`
- `echo-requestnode-info`
- `home-agent-replynode-query`
- `home-agent-requestpacket-too-big`
- `inv-nd-naparameter-problem`
- `inv-nd-nsredirect`
- `mcast-router-advertiserouter-advertisement`
- `mcast-router-solicitrouter-renum`
- `mcast-router-terminate router-solicitation`
- `mld-done time-exceeded`
- `mld-query ver2-mld-report`
- `mld-report`

Defining the IGMP match criteria

To more precisely define the IGMP packets to match in an IPv4 traffic class, use the optional parameter settings described in this section. For example, instead of matching all IGMP traffic, configure a class that matches only a specific IGMP packet type.

Context: Class configuration

Syntax:

[no] [*seq–number*] [match | ignore] igmp *source–address destination–address* [*igmp–type*] [ip–dscp *codepoint*] [precedence *precedence–value*] [tos *tos–value*] [vlan *vlan–id*]

If you enter `igmp` as the IP protocol type in a match/ignore statement, you can optionally specify an IGMP packet type to more precisely define match criteria for a traffic class. Enter the optional IGMP match criteria immediately after the destination IP address (DA) value in the command syntax; for example:

```
HP Switch(config-class)# match igmp any any host-query
```

`igmp-type`

Configures an IGMP packet type as match criteria in a class configuration. Some of the valid values for IGMP packet-type names are as follows:

- `dvmrpmtrace-requesttrace`
- `host-querymtrace-replyv2-host-leave`
- `host-reportpimv2-host-report`
- `v3-host-report`

To display a list of valid `igmp-type` entries when you enter `igmp` as the IP protocol type in a match/ignore statement, enter **`?`**.

Defining TCP and UDP match criteria

In a class configuration, you can enter match/ignore statements that more precisely define the TCP or UDP traffic to match in an IPv4 or IPv6 traffic class. For example, enter a port number as a match criterion that specifies one or more TCP source ports, destination ports, or both.

Context*:* Class configuration

Syntax:

[no] [`seq–number`] [`match` | `ignore`] `tcp` | `udp` *`source–address`* [*`operator tcp–src–port`* | *`udp–src–port`*] *`destination–address`* [*`operator tcp–dest–port`* `[established] [tcp–flag tcp–flag ... ] udp–dest–port`] [`ip–dscp` *`codepoint`*] [`precedence` *`precedence–value`*] [`tos` *`tos–value`*] [`vlan` *`vlan–id`*]

If you use TCP or UDP as the IP protocol type in a match/ignore statement, you can optionally configure TCP or UDP source and destination port numbers or ranges of numbers to more precisely define match criteria for a traffic class. Enter the optional TCP/UDP match criteria immediately after the source and destination address in the command syntax; for example:

```
HP Switch(config-class)# match tcp host 10.20.10.17 eq 23 host 10.20.10.155
established
HP Switch(config-class)# match tcp host 10.10.10.100 host 10.20.10.17  eq telnet
HP Switch(config-class)# ignore udp 10.30.10.1/24 host 10.20.10.17  range 161 162
```

`operator` *`tcp-src-port`* | *`udp-src-port`*

To specify a TCP or UDP source port number as a match criteria, enter a comparison operator from the following list with a TDP/UDP port number or well-known port name immediately after the source-address value in the command.

Comparison Operators:

- `eq` *`tcp/udp-port-number`*

Equal To matches a packet with the same TCP or UDP source port number as *`tcp/udp-port-number`*.

- `gt` *`tcp/udp-port-number`*

Greater Than matches any packet with a TCP or UDP source port number greater than *`tcp/udp-port-number`*.

- `lt` *`tcp/udp-port-number`*

Less Than matches any packet with a TCP or UDP source port number less than *`tcp/udp-port-number`*.

- `neq` *`tcp/udp-port-number`*

Not Equal matches any packet with a TCP or UDP source port number that is not equal to *`tcp/udp-port-number`* .

- `range` *`start-port-number end-port-number`*

Matches any packet with a TCP or UDP source port number in the range*`start-port-number`* to*`end-port-number`*.

TCP/UDP well-known source-port names and numbers

Enter a comparison operator with the source TCP or UDP port number used by the applications you want to match. Valid port numbers are from `0` to `255`. You can also enter well-known TCP or UDP port names as an alternative to the corresponding port number; for example:

- TCP: `bgp, dns, ftp, http, imap4, ldap, nntp, pop2, pop3, smtp, ssl, telnet`
- UDP: `bootpc, bootps, dns, ntp, radius, radius-old, rip, snmp, snmp-trap, tftp`

To display a list of valid TCP/UDP source ports, enter **?** after you enter an operator.

`operator tcp-dest-port` `established [tcp-flag` *`tcp-flag`* `... ] |` *`udp-destport`*

To specify a TCP or UDP destination port number as a match criteria, enter a comparison operator with a TDP/UDP port number or well-known port name immediately after the destination-address value in the command.

NOTE: The optional `established` and *`tcp-flag`* values apply only to TCP destination-port criteria.

TCP/UDP well-known destination-port names and numbers

The same operators, port numbers and well-known names are supported for TCP/UDP destination-port match criteria as for TCP/UDP source-port criteria. To display a list of valid TCP/UDP destination ports, enter **?** after you enter an operator.

`established`

(Optional) Applies only to TCP destination-port match criteria and matches only on the TCP Acknowledge (ACK) or Reset (RST) flags. The `established` keyword ignores the synchronizing packet associated with the establishment of a TCP connection in one direction on a port or VLAN, and matches all other IP traffic in the opposite direction.

For example, a Telnet connection requires TCP traffic to move both ways between a host and the target device. If you configure a match statement for inbound Telnet traffic, policy actions are normally applied to Telnet traffic in both directions because responses to outbound requests are also matched. However, if you enter the `established` option, inbound Telnet traffic arriving in response to outbound Telnet requests is matched, but inbound Telnet traffic trying to establish a connection is not matched.

`tcp-flag` *`tcp-flag`* ...

(Optional) Applies only to TCP bit settings in packets destined to a TCP destination port configured as match criteria (with the *`tcp-dest-port`* parameter) and can be one or more of the following values:

ack

Acknowledge matches TCP packets with the ACK flag.

fin

Finish matches TCP packets with the FIN flag.

rst

Reset matches TCP packets with the RST bit set.

syn

Synchronized matches TCP packets with the SYN flag.

How IPv4 mask bit settings define a match (Example)

For this example, the following configuration exists:

- A match statement in a class configuration uses an IPv4 source-address/mask-length of 10.38.31.125/21. The mask-length of 21 results in an IPv4 mask of 0.0.7.255. In the second octet of the mask, 7 means that the rightmost three bits are on or 1 (see the Mask for SA row in "How IPv4 mask defines a match" (page 374)).
- The second octet of the corresponding source address is 31, which means that the rightmost five bits are on or 1 (see the SA in Match Statement row in "How IPv4 mask defines a match" (page 374)).

In this example, a match occurs when the second octet of the SA in a packet being classified has a value in the range of 24 (binary 00011000) to 31 (binary 00001111), as shown in the last row in the following table.

Table 33 How IPv4 mask defines a match

Location of octet	**Bit position in the octet**							
	128	64	32	16	8	4	2	1
SA in match statement	0	0	0	1	1	1	1	1
Mask for SA	0	0	0	0	0	1	1	1
Bits in the corresponding octet of a packet's SA that must exactly match	0	0	0	1	1	0/1	0/1	0/1

The shaded area indicates the bits in the packet that must exactly match the bits in the source IPv4 address in the match/ignore statement.

- If a mask bit is 1 (wildcard value), the corresponding bits in a source/destination address in an IPv4 packet header can be any value.
- If a mask bit is 0, the corresponding bits in a source/destination address must be the same value as in the IPv4 address in the match/ignore statement.

NOTE: This example covers only one octet in an IPv4 address used as a match criterion. The mask in a match/ignore statement may apply a packet filter to all four octets of a source/destination address in IPv4 packet headers.

How IPv6 mask bit settings define a match (Example)

For an example in which an IPv6 prefix-length of 126 is used to select four IPv6 addresses in a match statement, see Figure 79.

The specified source IPv6 address is: `2001:DB8:0000:0000:244:17FF:FEB6:D37D`.

The IPv6 prefix-length (/126) results in the IPv6 mask:

`FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFC`.

Figure 79 Mask for matching four IPv6 devices

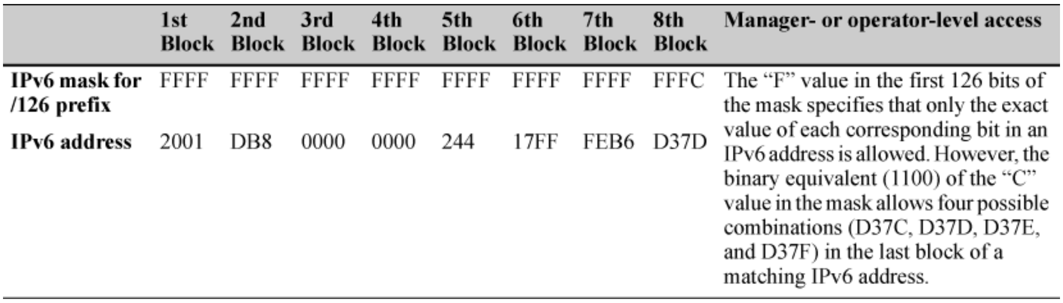

	1st Block	2nd Block	3rd Block	4th Block	5th Block	6th Block	7th Block	8th Block	Manager- or operator-level access
IPv6 mask for /126 prefix	FFFF	FFFF	FFFF	FFFF	FFFF	FFFF	FFFF	FFFC	The "F" value in the first 126 bits of the mask specifies that only the exact value of each corresponding bit in an IPv6 address is allowed. However, the binary equivalent (1100) of the "C" value in the mask allows four possible combinations (D37C, D37D, D37E, and D37F) in the last block of a matching IPv6 address.
IPv6 address	2001	DB8	0000	0000	244	17FF	FEB6	D37D	

To see the on and off settings in the last block of the resulting IPv6 mask that determine the matching IPv6 addresses, see Figure 79. In this mask, all bits except the last two are set to 1 (on) and must be the same in an IPv6 address. The binary equivalent of hexadecimal `C` is 1100, which allows the last two bits to differ.

Figure 80 How a mask determines four authorized IPv6 manager addresses

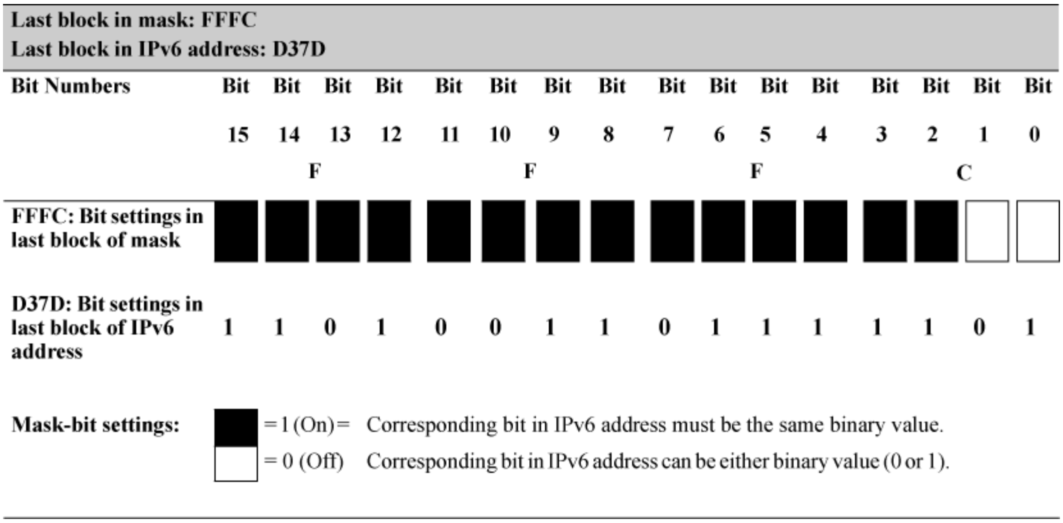

Last block in mask: FFFC
Last block in IPv6 address: D37D

Bit Numbers	Bit 15	Bit 14	Bit 13	Bit 12	Bit 11	Bit 10	Bit 9	Bit 8	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
		F				F				F				C		
FFFC: Bit settings in last block of mask	■	■	■	■	■	■	■	■	■	■	■	■	■	■	□	□
D37D: Bit settings in last block of IPv6 address	1	1	0	1	0	0	1	1	0	1	1	1	1	1	0	1

Mask-bit settings:

■ = 1 (On) = Corresponding bit in IPv6 address must be the same binary value.

□ = 0 (Off) Corresponding bit in IPv6 address can be either binary value (0 or 1).

To see how the binary equivalent (1100) of the C value in the last block of the resulting IPv6 mask supports four possible combinations (D37C, D37D, D37E, and D37F) in the last block of a matching IPv6 address, see Figure 81. Therefore, the IPv6 mask that results from a /126 prefix-length matches inbound traffic from four IPv6-based devices.

Figure 81 How hexadecimal C in an IPv6 mask matches four IPv6 addresses

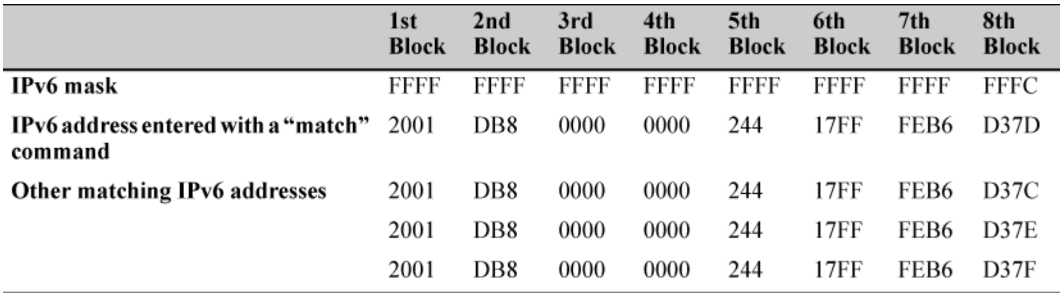

	1st Block	2nd Block	3rd Block	4th Block	5th Block	6th Block	7th Block	8th Block
IPv6 mask	FFFF	FFFF	FFFF	FFFF	FFFF	FFFF	FFFF	FFFC
IPv6 address entered with a "match" command	2001	DB8	0000	0000	244	17FF	FEB6	D37D
Other matching IPv6 addresses	2001	DB8	0000	0000	244	17FF	FEB6	D37C
	2001	DB8	0000	0000	244	17FF	FEB6	D37E
	2001	DB8	0000	0000	244	17FF	FEB6	D37F

CIDR notation

For more detailed information on how to use CIDR notation to specify masks in match criteria, see the "How an ACE Uses a Mask To Screen Packets for Matches" section in the *Access Control Lists (ACLs)* chapter in the *Access Security Guide*.

Resequencing match/ignore statements

Use the `class` command with the `resequence` option to reconfigure the number at which the first match/ignore statement in the class starts, and reset the interval used to number other match/ignore statements.

Resequencing match/ignore statements is useful when you want to insert a new match/ignore statement between two numbered entries (see “Resequencing a class configuration” (page 377)).

Context: Global configuration

Syntax:

```
class resequence [ ipv4 | ipv6 ] nameseq-number interval
```

`resequence`

Resets the sequence numbers for all match/ignore statements in the class.

`name`

Specifies the name of the class that contains the match/ignore statements that you want to resequence.

`seq-number`

Specifies the sequence number of the first match/ignore statement in the class. Default: 10.

`interval`

Specifies the interval between sequence numbers of match/ignore statements in the class to allow additional match/ignore statements to be inserted. Default: 10.

To view the current sequence numbering in a particular class, enter the following command:

```
show class [ ipv4 | ipv6 ] classname
```

Example 213 Resequencing a class configuration

The following example shows how to resequence a class configuration so that you can insert new match/ignore statements between sequentially numbered statements. In this example, the resequenced class contains two additional match/ignore statements and renumbers the criteria with an interval of 10.

Figure 82 Resequencing a class configuration

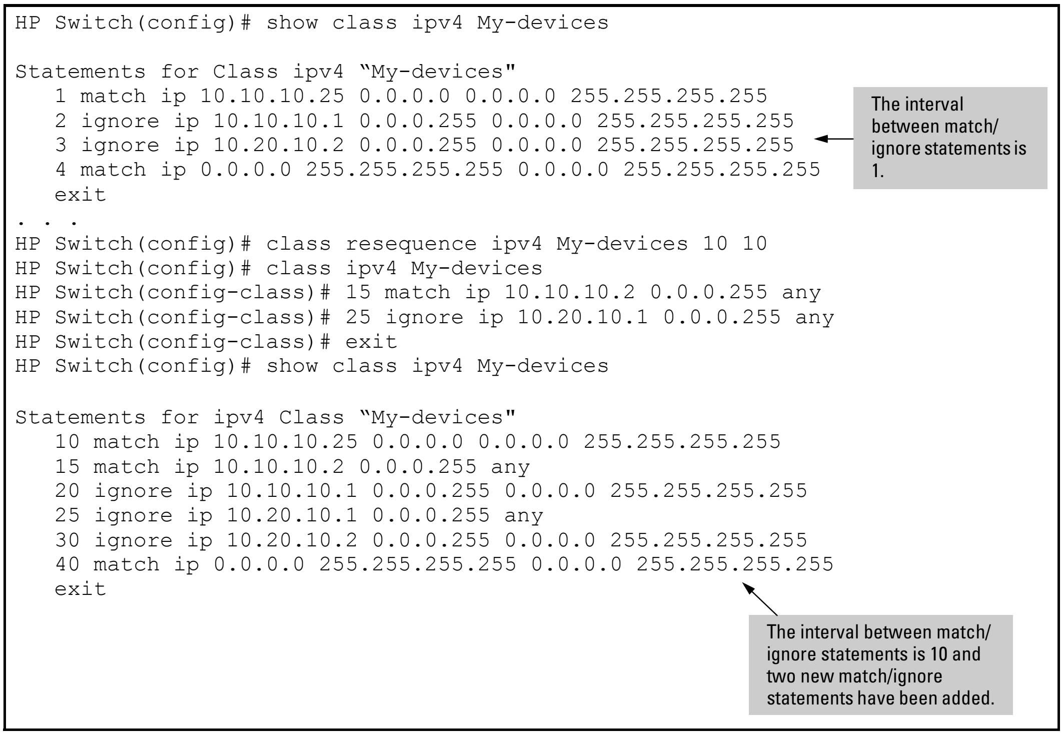

```
HP Switch(config)# show class ipv4 My-devices

Statements for Class ipv4 “My-devices"
   1 match ip 10.10.10.25 0.0.0.0 0.0.0.0 255.255.255.255
   2 ignore ip 10.10.10.1 0.0.0.255 0.0.0.0 255.255.255.255
   3 ignore ip 10.20.10.2 0.0.0.255 0.0.0.0 255.255.255.255
   4 match ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
   exit
. . .
HP Switch(config)# class resequence ipv4 My-devices 10 10
HP Switch(config)# class ipv4 My-devices
HP Switch(config-class)# 15 match ip 10.10.10.2 0.0.0.255 any
HP Switch(config-class)# 25 ignore ip 10.20.10.1 0.0.0.255 any
HP Switch(config-class)# exit
HP Switch(config)# show class ipv4 My-devices

Statements for ipv4 Class “My-devices"
   10 match ip 10.10.10.25 0.0.0.0 0.0.0.0 255.255.255.255
   15 match ip 10.10.10.2 0.0.0.255 any
   20 ignore ip 10.10.10.1 0.0.0.255 0.0.0.0 255.255.255.255
   25 ignore ip 10.20.10.1 0.0.0.255 any
   30 ignore ip 10.20.10.2 0.0.0.255 0.0.0.0 255.255.255.255
   40 match ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
   exit
```

The interval between match/ ignore statements is 1.

The interval between match/ ignore statements is 10 and two new match/ignore statements have been added.

Creating a service policy

In the classifier-based configuration model, the service policy you create for one or more traffic classes is always relative to a software feature, such as QoS, port and VLAN mirroring, or PBR. The software feature must support class and policy configuration. Each feature supports different actions for managing selected packets.

For example, QoS policies support QoS-specificactions, such as rate limiting, 802.1p-priority, IP-precedence, and DSCP-codepoint assignment. Port and VLAN mirroring policies support mirror-destination assignment for matching packets. PBR policies support specifying the IP next-hop and IP default next-hop, tunnel ID, or null for matching packets.

1. To create a service policy that performs feature-specific actions on selected packets, enter the `policy` *`feature-name`* command from the global configuration context.

   Context: Global configuration

   ### *Syntax:*

   [no] `policy` [`qos` | `mirror` | `pbr`] [*`policy-name`*]

   Defines the name of a service policy and enters the policy configuration context, where *`policy-name`* is a text string (64 characters maximum).

   A traffic policy consists of one or more actions that are configured for each class of traffic. The configured actions are executed on packets that match a `match` statement in a class. No policy action is performed on packets that match an `ignore` statement. You can configure multiple classes in a policy.

2. To configure the actions that you want to execute on packets that match the `match` criteria in a specified class, enter one or more `class action` commands from the policy configuration context.

   Context: Policy configuration

   Syntax:

   `[no] [seq-number] class [ ipv4 | ipv6` *`classname`* `action` *`action-name`* `]`
   `[action` *`action-name`* `...]`

   Defines the actions to be applied on a pre-configured IPv4 or IPv6 traffic class when a packet matches the `match` criteria in the class.

   You can enter multiple class-action statements for the same class. The actions supported for a class command differ according to the feature-specific policy (for example, QoS or mirroring) configured with the `policy` command in Step 1.

   `seq-number`

   (Optional) Sequentially orders the class-action statements in a policy configuration. Actions are executed on matching packets in numerical order.

   Default: Class-action statements are numbered in increments of 10, starting at 10.

   `class ipv4|ipv6` *`classname`*

   Defines the preconfigured class on which the actions in a class-action statement are executed, and specifies whether the class consists of IPv4 or IPv6 traffic. The class name is a text string (64 characters maximum).

   NOTE: You can configure multiple class-action statements to include different classes in a policy. The execution of actions is performed in the order in which the class-actions are numerically listed.

   `action` *`action-name`* `[action` *`action-name`* `...]`
   The `action` keyword configures the action specified by the `action-name` parameter. The action is executed on any packet that matches the `match` criteria in the class. The action is not executed on packets that match `ignore` criteria. You can configure more than one action for a class. The complete `no` form of the `class action` command or the `no` *`seq-number`* command removes an action from the policy configuration. For information on the exact actions supported by each classifier-based software feature, see the appropriate chapter in the HP Switch documentation set, as described in "Configuring class-based zones" (page 386)

   Be sure to enter a class and its associated actions in the precise order in which you want packets to be checked and handled by `class action` commands.

3. (Optional) To configure adefault class, enter the `default-class` command and specify one or more actions to be executed on packets that are not matched and not ignored.

   Context: Policy configuration

   Syntax:

   `[no] default-class action` *`action-name`* `[action` *`action-name`* `...]`

   Configures a default class to be used to execute one or more actions on packets that are not matched nor ignored in any of the class configurations in a policy. The `default-class action` command supports only the feature-specific commands supported in the `class action`command.

   The default class manages packets that do not match the `match` or `ignore` criteria in all classes in a policy, and otherwise would have no actions performed on them.

The default class differs from other classes because it contains no match/ignore statements and uses implicit `match ipv4 any any` and `match ipv6 any any` statements to manage all unmatched packets. If you do not configure a default class, unmatched and unignored packets are transmitted without an action performed on them.

4. Enter the `exit` command to exit the policy configuration context.

   To display a policy configuration, enter the `show policy` *`policy-name feature-name`* command (see "Resequencing a policy configuration" (page 383)), where *`feature-name`* is a software feature (such as `qos`, `mirror`, or `pbr`) that supports classifier-based configuration.

   To edit a policy configuration, re-enter the policy context (`policy` command) and modify class-action statements as described in "Modifying classes in a policy" (page 382).

   To resequence the order in which class-action statements are listed, enter the `resequence` command (see "Resequencing classes in a policy" (page 382)).

Example 214 A policy configuration

In the following QoS policy configuration, matching HTTP packets are rate limited to 10000 kbps. All unmatched packets are managed by the default class, which assigns a slightly higher 802.1p priority (4) and a new DSCP codepoint (5).

```
HP Switch(config)# class ipv4 http
HP Switch(config-class)# match tcp any any eq 80
HP Switch(config-class)# match tcp any any eq 8080
HP Switch(config-class)# exit
HP Switch(config)# policy qos RateLimitPrioritizeSuspectTraffic
HP Switch(policy-qos)# class ipv4 http action rate-limit kbps 10000
HP Switch(policy-qos)# default-class action priority 4 action dscp 5
HP Switch(policy-qos)# exit
```

A policy configuration requires a feature-specific `policy` command to identify the software feature used to manage one or more traffic classes:

- To configure a QoS policy, use the `policy qos` command as described in the "Quality of Service" chapter in the *Advanced Traffic Management Guide*.
- To configure a mirroring policy, use the `policy mirror` command as described in the *Monitoring and Analyzing Switch Operation* appendix in the *Management and Configuration Guide*.

Creating a PBR policy

PBR provides the ability to manipulate a packet's path based on attributes of the packet. Traffic with the same destination can be routed over different paths, so that different types of traffic, such as VOIP or traffic with special security requirements, can be better managed.

The supported actions for PBR are:

- Setting the next hop for routing the packet (`[ipv4 | ipv6] next-hop [`*`ip-addr`*`]`).
- Setting the next hop for routing the packet if there is no explicit route for this destination (`[ipv4 | ipv6] ip default-next-hop [`*`ip-addr`*`]`).
- Setting the outbound tunnel interface for the packet (`interface tunnel [`*`tunnel-ID`*`]`). See the chapter "Tunneling IPv6 over IPv4" in the *IPv6 Configuration Guide* for your switch for information on configuring tunnels.
- Setting `interface null`, which specifies that the packets are dropped if no other actions have occurred.

Operating notes for PBR

- Multiple actions can be configured for a class, up to 8 actions per class.
- If you configure an action of interface null, no more actions for that class may be configured.
- Only one of the 8 possible actions can be active at one time.
- The precedence of actions is indicated by the order in which they are added to the policy.
- Actions can only be added to a class, and they are added to the end of the action list for the class.
- To remove actions from a class, the entire class must be removed from the policy.
- When an action becomes inactive, for example, if the configured address becomes unreachable (for `next-hop` and `default-next-hop)` or the interface goes down (for a tunnel), the policy is configured with the next action for that class, if possible. If that action is not active, the next action is tried, and so on, until an interface null or the end of the list of configured actions is encountered. If the end of the list is reached, the policy action for that class behaves as if no PBR policy is applied.
- The maximum combined number of unique IP next-hops and default-next-hops supported is 256.

Example 215 TCP and UDP traffic routing

The following example shows TCP and UDP traffic routed on different network paths. First, the traffic classes are created, then the PBR policy is created, and lastly the PBR policy is applied to an interface.

```
HP Switch(config)# class ipv4 TCP
HP Switch(config-class)# match tcp 10.0.8.1/24 15.29.16.104/24 eq 80
HP Switch(config-class)# match tcp 10.0.8.1/24 15.29.16.104/24 eq 22
HP Switch(config-class)# match tcp 10.0.8.1/24 15.29.16.104/24 eq 23
HP Switch(config-class)# exit
HP Switch(config)# class ipv4 UDP
HP Switch(config-class)# match udp 10.0.8.1/24 15.29.16.104/24 eq 80
HP Switch(config-class)# match udp 10.0.8.1/24 15.29.16.104/24 eq 22
HP Switch(config-class)# match upd 10.0.8.1/24 15.29.16.104/24 eq 23
HP Switch(config-class)# exit
HP Switch(config)# class ipv6 TCP
HP Switch(config-class)# match tcp 2001::1/64 3001::1/64 eq 80
HP Switch(config-class)# match tcp 2001::1/64 3001::1/64 eq 22
HP Switch(config-class)# match tcp 2001::1/64 3001::1/64 eq 23
HP Switch(config-class)# exit
HP Switch(config)# class ipv6 UDP
HP Switch(config-class)# match udp 2001::1/64 3001::1/64 eq 80
HP Switch(config-class)# match udp 2001::1/64 3001::1/64 eq 22
HP Switch(config-class)# match udp 2001::1/64 3001::1/64 eq 23
HP Switch(config-class)# exit
HP Switch(config)# policy pbr TCP_UDP
HP Switch(policy-pbr)# class ipv4 TCP
HP Switch(policy-pbr-class)# action ip next-hop 20.0.0.1
HP Switch(policy-pbr-class)# action interface null
HP Switch(policy-pbr-class)# exit
HP Switch(policy-pbr)# class ipv4 UDP
HP Switch(policy-pbr-class)# action ip default-next-hop 30.0.0.1
HP Switch(policy-pbr-class)# action interface tunnel 3
HP Switch(policy-pbr-class)# exit
HP Switch(policy-pbr)# class ipv6 TCP
HP Switch(policy-pbr-class)# action ip next-hop 20.0.0.1
HP Switch(policy-pbr-class)# exit
HP Switch(policy-pbr)# class ipv6 UDP
HP Switch(policy-pbr-class)# action ip next-hop 30.0.0.1
HP Switch(policy-pbr-class)# exit
HP Switch(policy-pbr)# exit
HP Switch(config)# vlan 10
HP Switch(vlan-10)# service-policy TCP_UDP in
```

Troubleshooting PBR

Use the `show statistics policy` command to display information about which PBR action for an applied policy is active. Hit counts for each entry in the class and policy with the active action are displayed.

```
HP Switch(vlan-111)# show statistics policy TCP_UDP vlan 111 in
HitCounts for Policy TCP_UDP
Total
100 class ipv4 TCP action
( 5 ) 10 match tcp 10.0.8.1 0.0.0.255 15.29.16.104 0.0.0.255 eq 80
( 2 ) 20 match tcp 10.0.8.1 0.0.0.255 15.29.16.104 0.0.0.255 eq 22
( 2 ) 30 match tcp 10.0.8.1 0.0.0.255 15.29.16.104 0.0.0.255 eq 23
110 class ipv4 voice action
( 4 ) 10 match tcp 10.0.8.1 0.0.0.255 15.29.16.104 0.0.0.255 eq 80
```

To enable debug logging for PBR, enter the `debug ip pbr` command. A message will be logged when a PBR policy is applied, when the action in a class becomes inactive, and when an action in a class becomes active. (See “Troubleshooting” in the *Management and Configuration Guide* for your switch for general information on debugging.)

Modifying classes in a policy

You can modify the classes and class-action statements in a policy configuration without removing them from the policy:

- To modify the match/ignore statements in a class, enter the class-configuration context with the command, and make the necessary changes by removing or replacing existing statements. To display a class configuration, enter the following command as shown in "Resequencing a class configuration" (page 377):

  `show class [ ipv4 | ipv6 ] classname`

  When you exit class configuration context, the changes are automatically saved and applied to existing policy configurations on the switch that use the class if the policies have not been applied to an interface. If a policy has already been applied to an interface, the editing changes are not accepted and an error message is displayed.

- To modify the class-action statements in a policy, enter the policy-configuration context with the `policy` *`feature-name policy-name`* command. To display a policy configuration, enter the following command as shown in "Resequencing a policy configuration" (page 383):

  `show policy feature-name policy-name`

  Then do one of the following:

  - You can enter a new class-action statement. If you do not enter a sequence number, the new class-action statement is inserted at the end of the policy configuration.
  - To remove a class-action statement from a policy configuration, enter the `no` *`sequence-number`* command or the complete form of the `no class ... action` command.
  - To resequence the order in which class-action statements are listed, enter the `resequence`command (see "Resequencing classes in a policy" (page 382)).
  - To replace an existing class-action statement, enter the `no` *`sequence-number`* command to delete the entry, and re-enter the following complete command:

    ```
    class [  ipv4 | ipv6 ] classname
    action action-name  or default-class action action-name
    ```

When exiting the policy-configuration context, the changes are automatically saved and applied to the policy configuration if the policy has not been applied to an interface. If the policy has already been applied to an interface, the editing changes are not accepted and an error message is displayed.

Resequencing classes in a policy

You can use the `policy` command with the `resequence` option to reconfigure the number at which the first class-action statement starts, and reset the interval used to number other class-actions.

Resequencing class-actions is useful when you want to insert a new class-action between two numbered entries.

Context: Global configuration

Syntax:

```
policy resequence name seq-number interval
```

`resequence`

Resets the sequence numbers for all class-action statements in the policy.

`name`

Specifies the name of the policy that contains the class-action statements that you want to resequence.

seq-number

Specifies the sequence number of the first class-action-statement in the policy. Default: 10.

interval

Specifies the interval between sequence numbers of class-action statements in the policy to allow additional statements to be inserted. Default: 10.

NOTE: When resequencing class-action statements in a policy, the default class-action-statement always remains as the last class-action statement.

To view the current class-action numbering in a policy, enter the following command:
`show policy` *`feature-name policy-name`*

Example 216 Resequencing a policy configuration

The following example shows how to resequence a policy configuration after displaying its contents. The resequenced policy allows you to add a new class-action statement between entries 100 and 200.

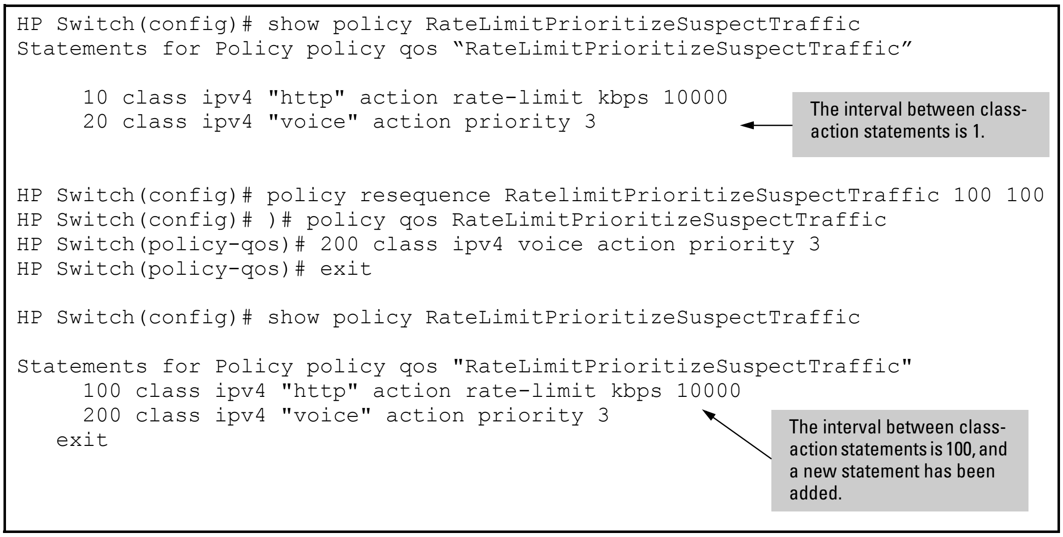

```
HP Switch(config)# show policy RateLimitPrioritizeSuspectTraffic
Statements for Policy policy qos “RateLimitPrioritizeSuspectTraffic”

     10 class ipv4 "http" action rate-limit kbps 10000
     20 class ipv4 "voice" action priority 3


HP Switch(config)# policy resequence RatelimitPrioritizeSuspectTraffic 100 100
HP Switch(config)# )# policy qos RateLimitPrioritizeSuspectTraffic
HP Switch(policy-qos)# 200 class ipv4 voice action priority 3
HP Switch(policy-qos)# exit

HP Switch(config)# show policy RateLimitPrioritizeSuspectTraffic

Statements for Policy policy qos "RateLimitPrioritizeSuspectTraffic"
     100 class ipv4 "http" action rate-limit kbps 10000
     200 class ipv4 "voice" action priority 3
   exit
```

The interval between class-action statements is 1.

The interval between class-action statements is 100, and a new statement has been added.

Applying a service policy to an interface

To apply feature-specific service policies to inbound port or VLAN interfaces, use the `interface service-policy in` or `vlan service-policy in` command.

The following service-policy restrictions apply to all software features:

- A service policy is supported only on inbound traffic.
- Only one feature-specific policy (for example, QoS or mirroring) is supported on a port or VLAN interface.
- PBR is only supported on a VLAN interface.
- If you apply a policy to a port or VLAN interface on which a policy of the same type (for example, QoS) is already configured, an error message is displayed. The new policy does not overwrite the existing one.

  Before you can apply a new policy, you must first remove the existing policy with the `no interface service-policy in` or `no vlan service-policy in` command.

Because only one policy of each type is supported on a port or VLAN interface, ensure that the policy you want to apply contains all the required classes and actions for your configuration.

NOTE: If ICMP rate limiting is already configured on a port, a service policy cannot be applied to the port until you disable the ICMP rate limiting configuration.

To apply a service policy to the port, maintain ICMP rate limiting by configuring a QoS policy in which you add the necessary `match` statements for ICMP packets to a class configuration and configure a `rate-limit` action for the class in the policy configuration.

For information on globally configured ICMP, see the "Configuring ICMP" section in the "Configuring IP Parameters for Routing Switches" chapter in the Multicast and Routing Guide.

To apply a service policy on a port or VLAN interface, enter one of the following commands from the global configuration context.

Context: Global configuration

Syntax:

`interface` *`port-list`* `service-policy` *`policy-name`* `in`

> Configures the specified ports with a policy that is applied to inbound traffic on each interface. Separate individual port numbers in a series with a comma; for example, `a1, b4, d3`. Enter a range of ports by using a dash; for example, `a1-a5`.
>
> The policy name you enter must be the same as the policy name you configured with the `policy` command (see "Creating a service policy" (page 377)).

Context: Global configuration

Syntax:

`vlan` *`vlan-id`* `service-policy` *`policy-name`* `in`

> Configures a policy on the specified VLAN that is applied to inbound traffic on the VLAN interface. Valid VLAN ID numbers range from `1` to `4094`.
>
> The policy name you enter must be the same as the policy name you configured with the `policy` command (see "Creating a service policy" (page 377)).

Example

Example 217 Applying a QoS policy to a port range and a VLAN interface

The following example shows how to apply a QoS policy to a port range and a VLAN interface:

```
HP Switch(config)# interface a4 service-policy RateLimitPrioritizeSuspectTraffic in
HP Switch(config)# vlan 10 service-policy RateLimitPrioritizeSuspectTraffic in
```

Checking resource usage

Syntax:

`show policy resources`

> After applying service policies to an interface, use the `show policy resources` command to verify the amount of additional resources used and the amount of resources that are still available on the switch. Classifier-based service policies (such as QoS or mirroring) share the same hardware resources with other software features, such as ACLs, virus throttling, management VLAN, globally configured QoS policies, MAC-based mirroring policies, and so on.
>
> Use the displayed information to decide whether to re-prioritize current resource usage by reconfiguring or disabling software features to free the resources reserved for less important features. For a detailed explanation of the information displayed

with the `show policy resources` command, see the "Monitoring Resources" appendix in the *Management and Configuration Guide*.

Example 218 Displaying policy resources

In this example, the `show policy resources` command output displays the number of hardware resources (rules, meters, and application port ranges) used by classifier-based QoS and mirroring policies that are currently applied to interfaces on the switch, as well as other software features.

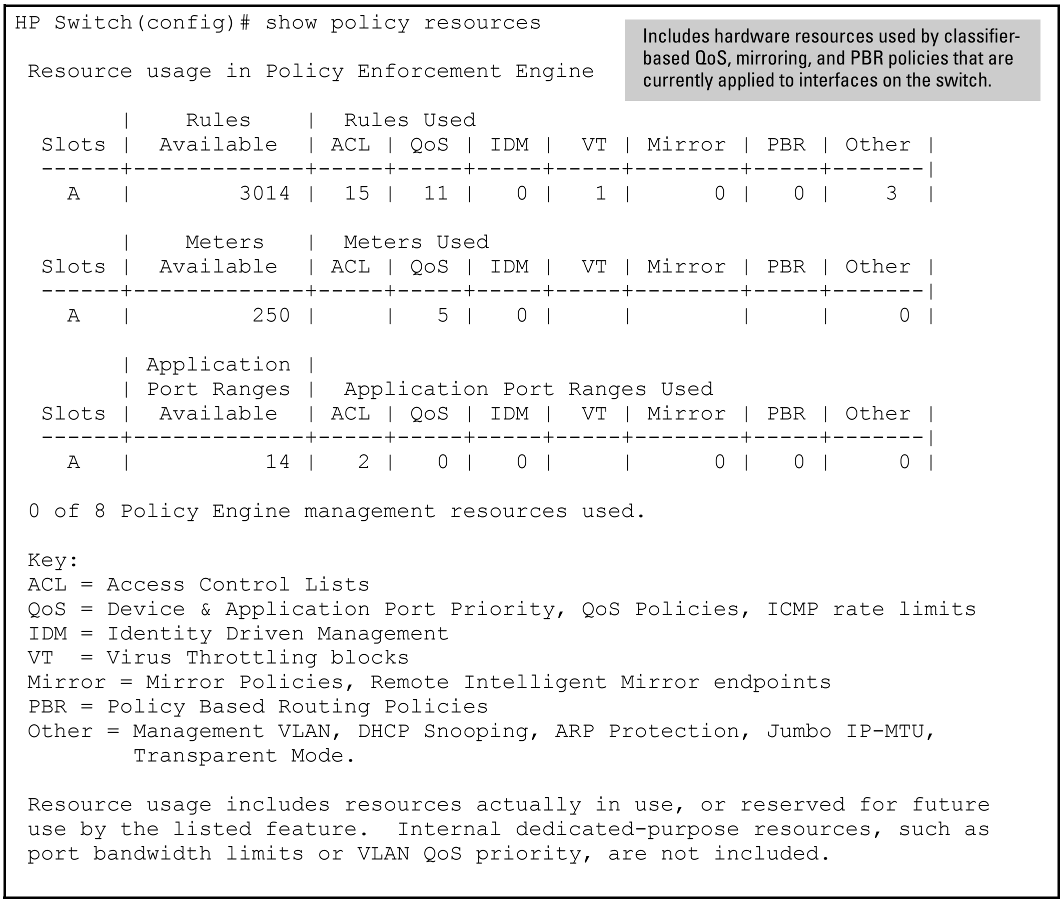

```
HP Switch(config)# show policy resources

 Resource usage in Policy Enforcement Engine

       |    Rules     |  Rules Used
  Slots |  Available   | ACL | QoS | IDM |  VT | Mirror | PBR | Other |
  ------+--------------+-----+-----+-----+-----+--------+-----+-------|
    A   |         3014 |  15 |  11 |   0 |   1 |      0 |   0 |    3  |

       |    Meters    |  Meters Used
  Slots |  Available   | ACL | QoS | IDM |  VT | Mirror | PBR | Other |
  ------+--------------+-----+-----+-----+-----+--------+-----+-------|
    A   |          250 |     |   5 |   0 |     |        |     |     0 |

       | Application |
       | Port Ranges |  Application Port Ranges Used
  Slots |  Available   | ACL | QoS | IDM |  VT | Mirror | PBR | Other |
  ------+--------------+-----+-----+-----+-----+--------+-----+-------|
    A   |           14 |   2 |   0 |   0 |     |      0 |   0 |     0 |

 0 of 8 Policy Engine management resources used.

 Key:
 ACL = Access Control Lists
 QoS = Device & Application Port Priority, QoS Policies, ICMP rate limits
 IDM = Identity Driven Management
 VT  = Virus Throttling blocks
 Mirror = Mirror Policies, Remote Intelligent Mirror endpoints
 PBR = Policy Based Routing Policies
 Other = Management VLAN, DHCP Snooping, ARP Protection, Jumbo IP-MTU,
         Transparent Mode.

 Resource usage includes resources actually in use, or reserved for future
 use by the listed feature.  Internal dedicated-purpose resources, such as
 port bandwidth limits or VLAN QoS priority, are not included.
```

Includes hardware resources used by classifier-based QoS, mirroring, and PBR policies that are currently applied to interfaces on the switch.

Displaying statistics for a policy

Only the active redirects (matches, ignores, etc.) are displayed when executing the `show statistics` command.

Example 219 Statistical output for a policy with active redirects

```
HP Switch(vlan-111)# show statistics policy TCP_UDP vlan 111 in

HitCounts for Policy TCP_UDP

  Total

 100 class ipv4 TCP action
(      0 )     10 match tcp 10.0.8.1 0.0.0.255 15.29.16.104 0.0.0.255 eq 80
(      0 )     20 match tcp 10.0.8.1 0.0.0.255 15.29.16.104 0.0.0.255 eq 22
(      0 )     30 match tcp 10.0.8.1 0.0.0.255 15.29.16.104 0.0.0.255 eq 23

  110 class ipv4 voice action
(      0 )     10 match tcp 10.0.8.1 0.0.0.255 15.29.16.104 0.0.0.255 eq 80
```

Configuring class-based zones

Zone class-based software configuration consists of the following general steps:

1. Identify the traffic that you want the HP 8200zl or 5400zl switch to intercept and forward to the ONE application. Specifically, you need to know:
   - Source
     - Which users or devices are sending the traffic that you want the switch to intercept?
     - Which switch ports are connected to these users' workstations and devices?
   - Destination
     - What is the destination of the traffic?

   Based on the answers to these questions, you can begin to plan your zones and zone policies.

2. Create a zone class.
   A zone class is a logical group of switch ports. For example, you might create an internal zone and assign all the ports on the internal network to this zone. You might then create an external zone and assign the switch port that connects to the firewall or router to this zone

3. Configure a zone policy for one or more zone classes, including an optional, default zone class.

   A policy specifies the traffic that should be intercepted based on the source and destination zones. Specifically, you define a zone policy by specifying:

   Source zone

   Where the traffic you want to accelerate originates.

   Destination zone

   Where the traffic will be sent.

   Intercept rules

   Define the action the switch will take—intercepting the traffic—and the direction of the traffic. For Transparent Mode, traffic is unidirectional, or one way.

   A zone policy can contain one or more intercept rules.

4. Associate the policy with the ONE application.

Creating a zone class

To use Transparent Mode, you create a zone class and use the `port-list` command to specify the ports that belong to a zone class. By default, the HP 5400zl or 8200zl switch supports a maximum of ten zones. Two are created automatically—BYPASS and SWITCH_SELF.

- BYPASS—contains the ports that should not be included in your Transparent Mode configuration. That is, the switch will not intercept traffic sent to or from the ports in the BYPASS zone. By default, all ports on the switch become part of the BYPASS zone, until you assign the ports to a different zone.
- SWITCH_SELF—contains only the switch. You cannot add any ports to this zone.

In addition to these two zones, you can create eight zones for a total of 10 zones.

Use the following guidelines when assigning ports to zones.

- The switch ports you add to a zone can be in different VLANs.
- Each switch port can belong to only one zone.
- If a port is already associated with a zone, adding the port to another zone removes that port from its existing zone and adds it to the new zone.
- Ports cannot be added to the SWITCH_SELF zone.

You may also need to create a zone for the ONE Application. Consult the HP Installation and Getting Started Guide for your ONE application.

The zone classes you configure will be used later in the zone policies you create.

To configure a zone class to be used in one or more policies, follow these steps:

1. Enter the `class zone` command from the global configuration context.

   Context: Global configuration

   ### *Syntax:*

   [no] `class zone` *`zone name`*

   Defines a zone class, where *zone name* is a text string (64 characters maximum).

   After you enter the `class zone` command, you enter the class zone configuration context. For transparent mode, you can then enter the port-list you want to define for this zone.

2. Enter the `port-list` command from the class configuration context to determine the ports on which policy actions will be performed.

   Context*:* Class configuration

   ### *Syntax:*

   [no] `port-list` [`ethernet`] *`port`*

   Defines the port or ports that are assigned to this zone class. A port may belong to only one zone; if a port is already associated with a zone, adding the port to another zone removes that port from its exiting zone and moves it to the new zone.

3. Enter the `exit` command to exit the class configuration context.
4. To display all the classes configured, including the zone classes, enter the `show class` command. To display a specific zone class configuration, enter the `show class zone`*`zone name`* command.
5. To edit a class configuration, re-enter the class configuration context (`class` command) and enter new port-list statements as follows:
6. To remove a port from the zone class, enter the `no port-list` *`port`* command.

Zone class configuration examples

The following example shows several class configurations:

- Ports A10-A24 belong to the internal zone class.
- Port A1 belongs to the external zone class.

Example 220 A zone class configuration

```
HP Switch(config)# class zone internal
HP Switch(config-class)# port-list a10-a24
HP Switch(config-class)# exit
HP Switch(config)# class zone external
HP Switch(config-class)# port-list a1
HP Switch(config-class)# exit
```

Creating a zone policy

1. To create a zone policy that performs zone-specific actions on selected packets, enter the `policy zone` *`policy-name`* command from the global configuration context.

   Context: Global configuration

   ### *Syntax:*

   [no] `policy zone` *`policy-name`*

   Defines the name of a service policy and enters the policy configuration context, where `policy-name` is a text string (64 characters maximum). This name should not be the same as a zone name.

   A zone policy consists of one or more actions that are configured for specific zones.

   No action is performed on packets sent to or from ports in the BYPASS zone. By default, ports are assigned to the BYPASS zone unless you explicitly assign them to a different zone.

2. To configure the actions that you want to execute on ports associated with a zone, enter one or more `class` commands from the policy configuration context.

   Context: Policy configuration

   ### *Syntax:*

   [no] [`seq-number`] `class zone`*`source zone name destination zone name`* `action intercept unidirectional`

   Defines the source and destination zones for packets that must be intercepted and forwarded to the ONE application.

   `seq-number`

   (Optional) Sequentially orders the class-action statements in a policy configuration. Actions are executed on matching packets in numerical order.

   `source zone name`

   Defines the source zone for packets that must be intercepted and forwarded to the ONE application.

   `destination zone name`

   Defines the destination zone for packets that must be intercepted and forwarded to the ONE application.

   `action intercept unidirectional`

   Defines the action as intercept and the flow of traffic as unidirectional (one-way).

   Default: Class-action statements are numbered in increments of 10, starting at 10.

The configured actions are executed on packets that arrive on the ports associated with the source zone and are destined for ports associated with the destination zone.

You cannot configure intercept rules for the BYPASS zone class. As such, traffic to and from the BYPASS zone cannot be intercepted.

3. Enter the `exit` command to exit the policy configuration context.
4. To display a policy configuration, enter the `show policy` *`policy-name`* command.

   To edit a policy configuration, re-enter the policy context (`policy` command) and modify class-action statements as described in "Modifying zones and policies" (page 389).

Example

Example 221 Forwarding zone traffic

In the following policy configuration, traffic being sent from the internal zone to the external zone is intercepted, so that it can be forwarded to an application that is running on an HP AllianceONE Extended Services zl Module.

```
HP Switch(config)# class zone internal
HP Switch(config-class)# port-list a10-a24
HP Switch(config-class)# exit
HP Switch(config)# class zone external
HP Switch(config-class)# port-list a1-a4
HP Switch(config-class)# exit
HP Switch(config)# policy zone Firewall
HP Switch(policy-config)# class zone internal external action intercept unidirectional
HP Switch(policy-config)# exit
```

Modifying zones and policies

You can modify the zones and class-action statements in a zone policy configuration without removing them from the policy:

- To modify the ports associated with a zone, enter the `class zone` *`classname`* command. Remember that the classname you entered is case sensitive. From the class-configuration context, make the necessary changes by removing or adding ports. (To display a class configuration, enter the `show class zone` *`classname`* command.)

  When you exit class configuration context, the changes are automatically saved and applied to existing policy configurations on the switch that use the class if the policies have not been applied to a ONE application. If a policy has already been applied, the editing changes are not accepted, and an error message is displayed.

- To modify the class-action statements in a policy, enter the `policy` *`policy-name`* command. (To display a policy configuration, enter the `show policy` *`policy-name`* command as shown.) From the policy-configuration context, complete one of the following:
  - Enter a new class-action statement. If you do not include a sequence number, the new class-action statement is inserted at the end of the policy configuration.
  - Remove a class-action statement by entering the `no` *`sequence-number`* command.
  - Replace an existing class-action statement by:
    - Entering the `no` *sequence-number* command to delete the entry.
    - Entering a new `class zone` *`source zone name destination zone name`* `action intercept unidirectional` command.

When you exit the policy-configuration context, the changes are automatically applied to the policy configuration if the policy has not been applied to an interface. If the policy has already been applied to an interface, the editing changes are not accepted and an error message is displayed.

Applying a zone policy to a ONE application

To apply a zone policy to a ONE Application, you can complete one of the following steps:

- Enter the `zone-service-policy` command on the HP 8200zl or 5400zl switch
- Use the ONE application's management interface to apply the zone policies

To apply zone policies through the ONE application, consult the *HP Installation and Getting Started Guide* for that application.

To apply a zone service policy from the switch CLI, enter the following command from the global configuration context.

Syntax:

```
[no] zone-service-policy policy-name zone[ enable | disable | update ]
bind  [Ethernet] logical port appname application name  appinstance
instance description[ fail-action  [ bypass |  block ]][ expire [
app-down |  permanent  |  slot-down ]]
```

`policy-name`

Specifies the name of the policy you configured with the `policy` command (see "Creating a service policy" (page 377)).

`enable|disable|update`

Makes the policy active, inactive, or updates options that have been assigned to a policy.

`logical port`

Specifies the HP AllianceONE Extended Services zl Module's port 1, using the format `slot1`.

`application name`

Specifies the name of the ONE application. See the *Installation and Getting Started Guide* for the ONE application.

`instance description`

Specifies the name of the application and the slot in which it resides. See the *Installation and Getting Started Guide* for the ONE application.

`fail-action bypass | block`

Specifies the action the switch will take if the ONE application is unavailable). Enter `bypass` if you want the switch to ignore the policies and not intercept traffic if the ONE application is unavailable. Specify `block` if you want the switch to drop traffic that matches your policy criteria if the ONE application is unavailable.

`expire`

Determines if the policy persists if the ONE application is down or unavailable.

`app-down`

Specifies if you want the policy to expire if the ONE application is unavailable. Specify `permanent` if you do not want the policy to ever expire.

`slot-down`

Specifies if you want the policy to expire if the slot in which the ONE application is installed is unavailable.

Example

Example 222 Applying a zone policy

The following example shows how to apply a zone policy:

```
HP Switch(config)# zone-service-policy Firewall zone enable bind F1 appname
```

Verify that the zone policy is associated with the ONE application by entering the following command:

```
hostswitch# show ONE_app slot_ID
```

Replace *slot_ID* with the slot in which the AllianceONE Extended Services zl Module is installed.

Figure 83 Displaying output for the show ONE_app command

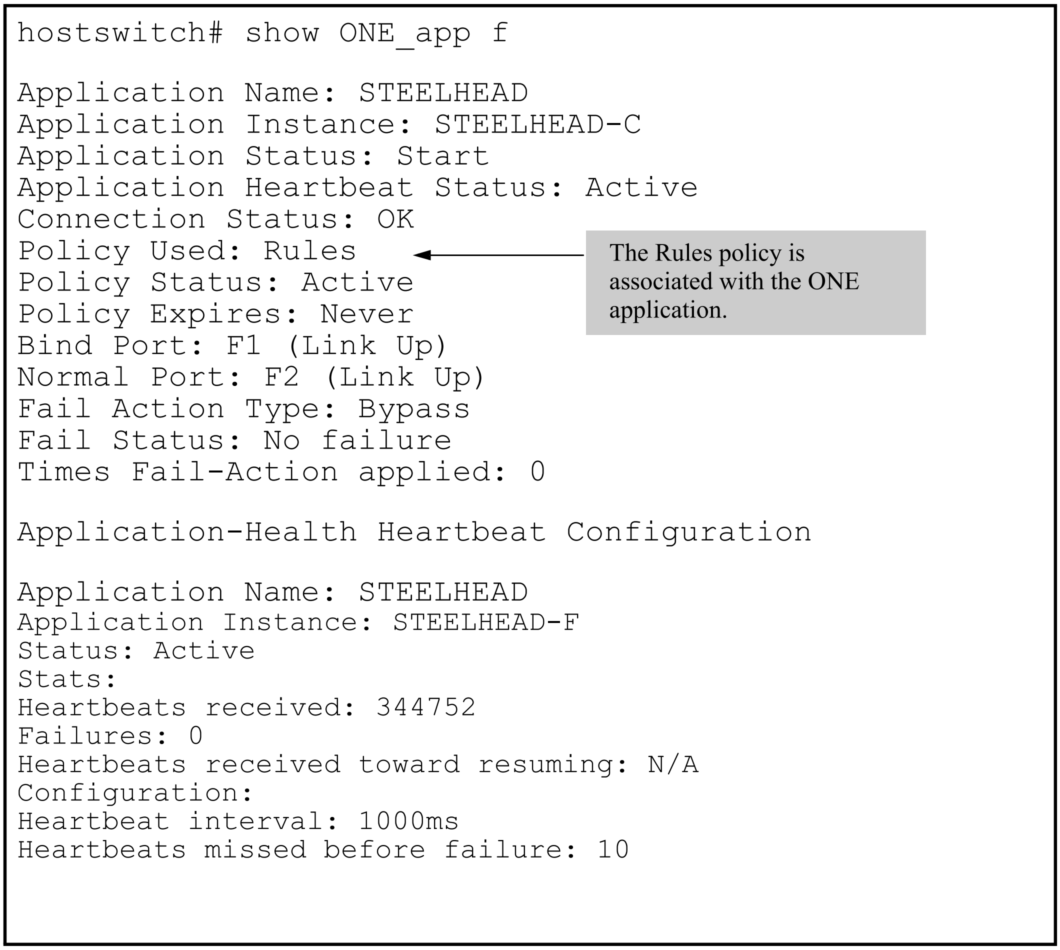

```
hostswitch# show ONE_app f

Application Name: STEELHEAD
Application Instance: STEELHEAD-C
Application Status: Start
Application Heartbeat Status: Active
Connection Status: OK
Policy Used: Rules
Policy Status: Active
Policy Expires: Never
Bind Port: F1 (Link Up)
Normal Port: F2 (Link Up)
Fail Action Type: Bypass
Fail Status: No failure
Times Fail-Action applied: 0

Application-Health Heartbeat Configuration

Application Name: STEELHEAD
Application Instance: STEELHEAD-F
Status: Active
Stats:
Heartbeats received: 344752
Failures: 0
Heartbeats received toward resuming: N/A
Configuration:
Heartbeat interval: 1000ms
Heartbeats missed before failure: 10
```

The Rules policy is associated with the ONE application.

About Classifier-based configuration

Traffic classes and software releases

The Classifier feature introduces:

- A finer granularity than globally configured features for placing network traffic (IPv4 or IPv6) into classes that can be used in cross-feature software configurations
- Additional policy actions, such as rate limiting and IP precedence marking, to manage selected traffic

- The configuration of service policies for classified traffic with the following software features:
  - Quality of Service (QoS)
  - Traffic mirroring
  - Policy Based Routing (PBR)
- The application of service policies to specific inbound traffic flows on individual port and VLAN interfaces (rather than only on switch-wide or port-wide traffic).

Using CIDR notation for IPv4/IPv6 addresses

You can use CIDR (Classless Inter-Domain Routing) notation to enter an IPv4 mask-length or an IPv6 prefix-length with a source and destination address that are used as match criteria in a match/ignore statement. The switch interprets the IP address with CIDR notation to compute the range of corresponding IP source or destination addresses in packet headers that are considered to be a match for the traffic class.

When the switch uses a match/ignore statement to compare an IP address and corresponding mask/prefix length to the IP source/destination address carried in a packet, the IPv4 mask-bit settings and IPv6 prefix-bit settings select packets in different ways.

- An IPv4 mask length creates a mask in which:
  - A mask-bit setting set to 0 (off) requires the corresponding bit in a packet's IPv4 source/destination address to be the same binary value as the mask-bit in the matching IPv4 source/destination address.
  - A mask-bit setting set to 1 (on) is used as a wildcard and allows the corresponding bit in a packet's IPv4 source/destination address to be either binary value (0 or 1).

Table 34 How CIDR notation is used with IPv4 SA/DA match criteria

IPv4 Source/Destination address used with CIDR notation in a Match/Ignore statement	Resulting mask	Range of IPv4 addresses selected by the match criteria
10.38.240.125/15	0.1.255.255	The leftmost 15 bits must match; the remaining bits are wildcards.
10.38.240.125/20	0.0.15.255	The leftmost 20 bits must match; the remaining bits are wildcards.
10.38.240.125/21	0.0.7.255	The leftmost 21 bits must match; the remaining bits are wildcards.
10.38.240.125/24	0.0.0.255	The leftmost 24 bits must match; the remaining bits are wildcards.
18.38.240.125/32	0.0.0.0	All bits must match.

- An IPv6 prefix-length creates a mask in which:
  - A mask-bit setting set to 1 (on) requires the corresponding bit in a packet's IPv6 source/destination address to be the same binary value as the mask-bit in the matching IPv6 source/destination address.
  - A mask-bit setting set to 0 (off) is used as a wildcard and allows the corresponding bit in a packet's IPv6 source/destination address to be either binary value (0 or 1).

Table 35 How CIDR notation is used with IPv6 SA/DA match criteria

IPv6 source/destination address used with CIDR notation in a Match/Ignore statement	Resulting mask	Range of IPv6 addresses selected by the match criteria
2001:db8:0:7::5/64	FFFF:FFFF:FFFF:FFFF::	The leftmost 64 bits must match; the remaining bits are wildcards.
2001:db8:0:7::5/72	FFFF:FFFF:FFFF:FFFF:FF00::	The leftmost 72 bits must match; the remaining bits are wildcards.
2001:db8::244:17ff:feb6:d37d/126	FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFC	The first 126 bits mst match; the C value in the mask allows four possible combinations (D37C, D37D, D37E, and D37F) in the last block of a matching IPv6 address.
2001:db8:0:7:af:e2:c1:5/128	FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF	All bits must match.

NOTE: Although IPv4 and IPv6 masks are applied in opposite directions:

- An IPv4 mask-length is applied from right to left, starting from the rightmost bits.
- An IPv6 prefix-length is applied from left to right, starting from the leftmost bits.

The behavior of IPv4 and IPv6 masks as match criteria and wildcards is the same.

Where to go from here

Classifier-based service policies are designed to work with your existing globally-configured software settings. While existing software features allow you to globally manage *all* network traffic on a switch or port, classifier-based service policies allow you to zoom in on subsets of network traffic to further manage it on a per-port or per-VLAN basis.

You can use the match criteria described in this chapter across software features to configure classes of traffic for use in feature-specific service policies.

After you decide on the IPv4 and IPv6 network traffic you want to manage, see the following chapters for more information about how to configure and use classifier-based quality-of-service and mirroring policies:

- *Quality of Service (QoS)* chapter in the *Advanced Configuration Guide*
- "Traffic Mirroring" section in the *Monitoring and Analyzing Switch Operation* appendix in the *Management and Configuration Guide*.

Traffic class-based configuration model

Traffic class-based software configuration consists of the following general steps:

1. Determine the inbound traffic you want to manage and how you want to manage it. For example, you may want to rate limit certain traffic, prioritize it, mirror it, and so on.
2. Classify the traffic that you want to manage by configuring a class, using `match` and `ignore` commands. A traffic class is configured separately from service policies and can be used in various policies.
3. Configure a service policy for one or more address classes, including an optional, default class. A policy consists of configuration commands executed on specified traffic classes for one of the following software features:
   - Quality of Service (`policy qos` command)
   - Port and VLAN mirroring (`policy mirror` command)
   - Policy Based Routing (`policy pbr` command)

4. Assign the policy to an inbound port or VLAN interface using the `interface service-policy in` or `vlan service-policy in` command.

The following figure shows an overview of traffic class-based software configuration.

Figure 84 Traffic class-based configuration model

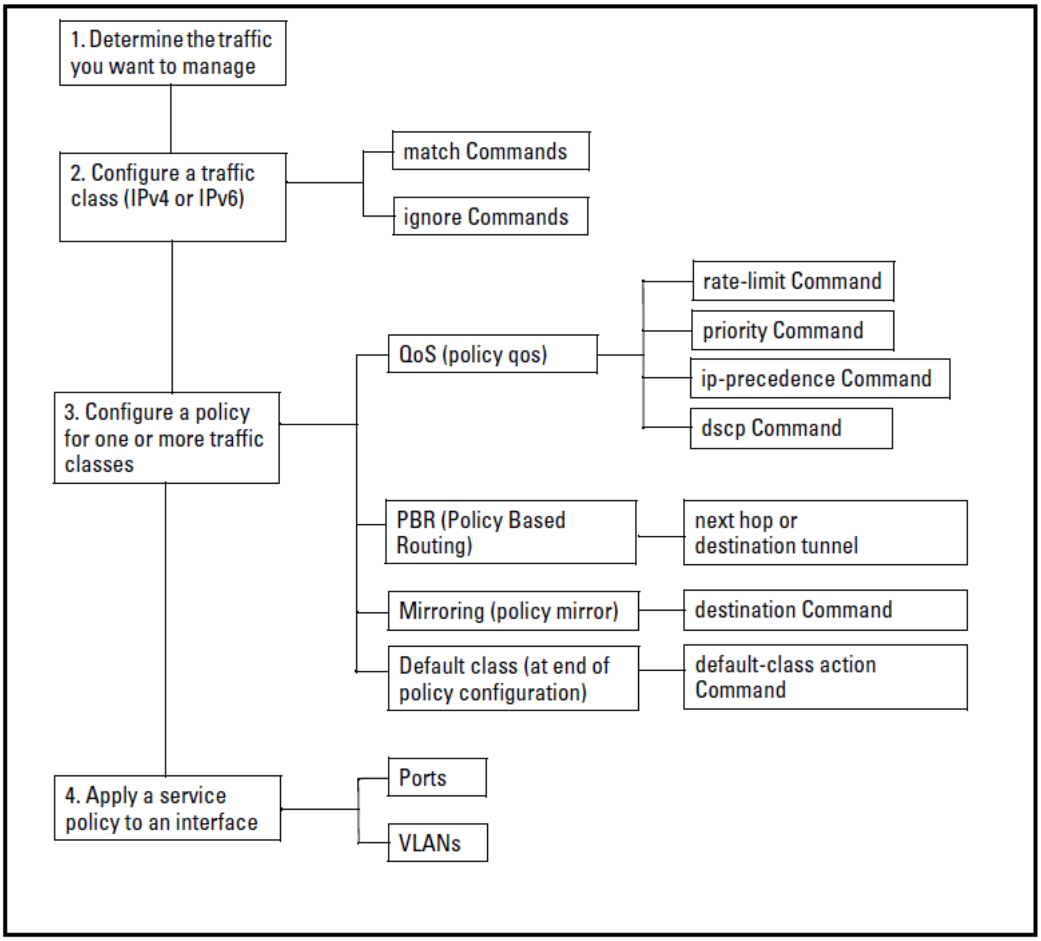
1. Determine the traffic you want to manage
2. Configure a traffic class (IPv4 or IPv6)
match Commands
ignore Commands
rate-limit Command
priority Command
QoS (policy qos)
ip-precedence Command
3. Configure a policy for one or more traffic classes
dscp Command
PBR (Policy Based Routing)
next hop or destination tunnel
Mirroring (policy mirror)
destination Command
Default class (at end of policy configuration)
default-class action Command
Ports
4. Apply a service policy to an interface
VLANs

Creating a traffic class

In the traffic class-based configuration model, you use match criteria to create a class of IPv4 or IPv6 traffic and select the packets you want to manage. In a traffic class configuration, match criteria consist of `match` and `ignore` commands. These commands determine the packets that belong to a class. (Match/ignore criteria are modelled on the permit/deny criteria used in ACLs.)

The traffic classes you configure can be used later in the service policies you create for different software features, such as QoS and port mirroring. The match criteria used in match/ignore statements are the same across software features.

Using match criteria

To identify the packets that belong to a traffic class for further processing by policy actions, use `match` and `ignore` commands in a class configuration:

`match` commands

Define the values that header fields must contain for a packet to belong to the class and be managed by policy actions.

`ignore` commands

Define the values which, if contained in header fields, exclude a packet from the policy actions configured for the class. An ignored packet is transmitted without having a policy action performed on it.

The switch compares match/ignore statements to the values in packet fields. It compares the specified criteria in the sequential order in which the statements are entered in the class, until a match is found. Be sure to enter match/ignore statements in the precise order in which you want their criteria to be used to check packets.

- As soon as a field in a packet header matches the criteria in a `match` statement, the sequential comparison of match criteria in the class stops, and the policy actions configured for the class are executed on the packet (see "Creating a service policy" (page 377)).
- If a packet matches the criteria in an `ignore` statement, the sequential comparison of match criteria in the class stops, and no policy action is performed on the packet.

If a packet does not match the criteria in any match/ignore statement in a traffic class configuration, one of the following actions is taken:

- The packet is transmitted without a policy action performed on it.
- If a default class is configured in the policy, the actions specified in the `default-class` command are performed on packets that do not match the criteria in preceding classes in the policy (see Step 3 in "Creating a service policy" (page 377)).

The following match criteria are supported in match/ignore statements for inbound IPv4/IPv6 traffic:

- IP source address (IPv4 and IPv6)
- IP destination address (IPv4 and IPv6)
- Layer 2 802.1Q VLAN ID
- Layer 3 IP protocol
- Layer 3 IP precedence bits
- Layer 3 DSCP bits
- Layer 4 TCP/UDP application port (including TCP flags)
- VLAN ID

Using zone classes

Zone classes are used in conjunction with a technology called Transparent Mode. Both zone classes and Transparent Mode are supported on the HP 5400zl and 8200zl Switch Series when:

- The switch is running software version K.14.58 and above
- An HP AllianceONE Extended Services zl Module is installed in the switch

Transparent Mode enables the HP 5400zl or 8200zl switch to intercept packets that match certain criteria and redirect them to an application that is running on an HP AllianceONE Extended Services zl Module. In effect, Transparent Mode allows that application to be in the path of packet flow. (Applications that run on AllianceONE Extended Services zl Modules are referred to as *ONE applications*.)

Packet interception differs from port mirroring. With port mirroring, the switch copies the packets being sent from a particular source port to a particular destination port and sends these packets to a mirror port. The original packets continue to be sent from the source port to the destination port without interruption.

With packet interception, on the other hand, the switch does not copy packets to another port. Instead the switch actually intercepts the packets and forwards them to the ONE application. The ONE application can then make decisions based upon those packets and control or modify the packets before they are delivered to their final destination.

After intercepting and redirecting the packets to the ONE application, the HP 5400zl or 8200zl switch does not perform any further action on the intercepted packets.

Because a ONE application is required to act on the intercepted packets, Transparent Mode is available only when an HP AllianceONE Extended Services zl Module is installed in the zl switch. In addition, the switch must also be running software version K.14.58 or above.)

NOTE: Check the release notes for the switch software you are using to ensure it supports the ONE application that is running on your AllianceONE Extended Services zl Module.

You will configure Transparent Mode commands only when your ONE application supports this functionality.

Troubleshooting problems

If you experience problems with your Transparent Mode configuration but cannot pinpoint the cause, you can use the following command to gather detailed information about your Transparent Mode configuration:

```
HP Switch# show tech transparentmode
```

The output from this command is displayed on your terminal emulator. However, using your terminal emulator's text capture features, you can save `show tech transparentmode` data to a text file for viewing, printing, or sending to an associate or even HP Support. For example, if your terminal emulator is the HyperTerminal application available with Microsoft® Windows® software, you can copy the `show tech transparentmode` to a file and then use either Microsoft Word or Notepad to display the data. (In this case, Microsoft Word provides the data in an easier-to-read format.)

Where to go from here

Zone service policies are designed to work with ONE applications that support Transparent Mode. (ONE applications run on HP AllianceONE Extended Services zl Modules.) See the ONE application's documentation for detailed information on how to configure zone classes and zone policies.

11 Support and other resources

Intended audience

This guide is intended for network administrators with intermediate-to-advanced knowledge of computer networking.

Related documentation

The following sources provide related information:

- *HP Basic Operation Guide*
- *HP Management and Configuration Guide*
- *HP Access Security Guide*
- *HP Multicast and Routing Guide*
- *HP IPv6 Configuration Guide*
- *Power over Ethernet (PoE/PoE+) Planning and Implementation Guide*
- *HP Switch 620 Redundant and External Power Suppy Installation and Getting Started Guide*
- *HP Switch 630 Redundant and/or External Power Supply Installation and Getting Started Guide*

You can also find the documents referenced in this guide on the Manuals page of the HP Business Support Center website: http://www.hp.com/support/manuals.

Contacting HP

HP technical support

For worldwide technical support information, see the HP support website:

http://www.hp.com/networking/support.

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Subscription service

HP recommends that you register your product at the Subscriber's choice for business website:

http://www.hp.com/go/e-updates

After registering, you will receive email notifications of product enhancements, new driver versions, firmware updates, and other product resources.

Related information

HP websites

- HP: http://www.hp.com
- HP Networking: http://www/hp.com/go/networking
- HP Partner Locator: http://www.hp.com/service_locator
- HP Software Downloads: http://www.hp.com/support/downloads

Typographical conventions

Table 36 Document conventions

Convention	Element
Blue text: Table 26	Cross-reference links and email addresses
Blue underlined text: http://www.hp.com	Website addresses
Bold text	• Keys that are pressed • Text entered into a GUI element, such as a box • GUI elements that are clicked or selected, such as menu▯ and list items, buttons, tabs, and check boxes
Italic text	Text emphasis
Monospace text	• File and directory names • System output • Code • Commands, their arguments, and argument values
Monospace italic text	• Code variables • Command variables
Monospace bold text	Emphasized monospace text
. . .	Indication that example continues

Product warranties

For information about HP Networking product warranties, see the warranty information website:

http://www.hp.com/networking/support

"Applicable Products" (page 1) lists related products and their part numbers.

HP customer support services

If you are having trouble with your switch, Hewlett-Packard offers support 24 hours a day, seven days a week through the use of a number of automated electronic services. See the Customer Support/Warranty booklet that came with your switch for information on how to use these services to get technical support. HP provides up-to-date customer care, support and warranty information at http://www.hp.com/networking/support.

Your HP authorized network reseller can also provide assistance, both with services that they offer and with services offered by HP.

Before calling support

Before calling your networking dealer or HP Support, to make the support process most efficient, first retrieve the following information:

Information item	Information location
• Product identification, including mini-GBICs	The front of the switch and on labels on the mini-GBICs
• Details about the switch's status including the software (OS) version, a copy of the switch configuration, a copy of the switch Event Log, and a copy of the switch status and counters information	Switch console: show tech command
• Copy of your network topology map, including network addresses assigned to the relevant devices	Your network records

Glossary

This glossary defines acronyms and terms used in this guide or related to this product and is not a comprehensive glossary of computer terms.

802.1p priority A traffic priority setting carried by a VLAN-tagged packet moving from one device to another through ports that are tagged members of the VLAN to which the packet belongs.

802.1Q field A four-byte field contained in the header of Ethernet packets entering or leaving the switch through a port that is a tagged member of a VLAN. This field includes an 802.1p priority setting, a VLAN tag, or ID number (VID), and other data.

ACL Access Control List. A list of one or more Access Control Entries (ACEs) specifying the criteria the switch uses to either permit (forward) or deny (drop) IP packets traversing the switch's interfaces.

Active Stack A stack that has one or more Missing members which have not been removed. All switches in the stack actively switch packets, and the stack has a Commander. Note: In the stand-alone case a single switch can be an active stack; however, it will not pick a stack ID unless one existed from a previous boot.

Active stack fragment When a stack becomes fragmented, only one fragment remains Active; the other fragments become Inactive (all network ports are disabled). The active stack fragment inherits the MAC address and IP addressing of the stack for management. The fragment that has more switches in it will be the Active fragment. This allows more of the network ports to remain operational. If the fragments have the same number of switches in them, then the fragment that has the original Commander will be the Active fragment.

Boot time In a stacking configuration, the time that the switch waits for all members of the stack to come up and for the stack protocol to elect a Commander. The value can be reset each time a new switch is added before the previous timeout has expired. Once the timeout has expired, an election for the Commander occurs. Note: Establishment of a ring topology or the discovery of all the switches in the stack, as determined from a previous boot of the stack, cancels this timeout.

BPDU Bridge Protocol Data Units. BPDUs are data messages that are exchanged between the switches within an extended LAN that use a spanning tree protocol topology. BPDUs contain information on ports, addresses, priorities and costs and ensure that the data ends up where it was intended to go. BPDU messages are exchanged across bridges to detect loops in a network topology. The loops are then removed by placing redundant switch ports in a backup, or blocked, state.

BPDU filtering Spanning tree configuration mode that prevents the switch from receiving and transmitting BPDU frames on a specific port.

BPDU protection spanning tree configuration mode that disables a port where BPDU frames are received.

C-VLAN Customer network VLAN that can exist across multiple locations.

C-VLAN bridge A customer-owned device operating regular 802.1Q VLANs.

Chain topology A stack topology in which all the switches are connected in a ring with a single stacking link removed. The switches on the ends of the chain only have a single link back to the peer switch. This topology is considered an error condition and is only expected when a Ring topology has been broken between two peer switches.

CIST Common and Internal Spanning Tree. Identifies the regions in a network and administers the CIST root bridge for the network, the root bridge for each region, and the root bridge for each spanning tree instance in each region.

Class Defines a set of different types of traffic, using match and ignore commands, to be managed in a specific way, for example by prioritizing all High Availability traffic.

Classifier Method of matching or ignoring certain traffic attributes that allows you to classify packets using more than one traffic attribute at a time. This method can be used by multiple features, such as QoS and mirroring.

Classifier model Applied to various software features, this configuration model requires you to first classify the traffic that you want to manage, and then configure a policy containing the actions to be executed on the class.

Commander In a stacking configuration, the physical switch that operates as the controller of the stack. The Commander responds to management requests for the stack. The Commander runs all the routing and switching protocols for the entire stack. There is only one Commander switch elected at any given time in the stack. The Commander's MAC address is used for management, in addition to the Commander's OOBM IP address.

CoS Class of Service: Priority level (0-7) used to transmit a packet on an outbound queue.

CST Common Spanning Tree. Administers the connectivity among the MST regions, STP LANs, and RSTP LANs in a bridged network.

Customer The consumer of network services delivered by a service provider.

Customer-network port Customer-facing port on a provider edge device.

Downstream device A device linked directly or indirectly to an outbound switch port. The switch sends traffic to downstream devices.

DSCP Differentiated Services Codepoint. Also known as *DiffServe Codepoint*. A DSCP consists of the upper six bits of the:

- Type of Service (ToS) byte in an IPv4 packet
- Traffic Class byte in an IPv6 packet

There are 64 possible DSCP codepoints. In the default QoS configuration for the switches covered in this guide, some codepoints (such as Assured Forwarding and Expedited Forwarding) are configured by default with an 802.1p priority. All other codepoints have no 802.1p priority assigned and are listed as `No-override`.

DSCP policy A DSCP codepoint that is configured with an 802.1p priority (0 to 7).

Dynamic VLAN An 802.1Q VLAN membership temporarily created on a port linked to another device, where both devices are running GVRP.

Edge switch In QoS, this is a switch that receives traffic from the edge of the LAN or from outside the LAN, and forwards it to devices within the LAN.

GVRP GARP VLAN Registration Protocol. An application of the Generic Attribute Registration Protocol—GARP. GVRP is defined in the IEEE 802.1Q standard, and GARP is defined in the IEEE 802.1D-1998 standard.

GVRP BPDU In this manual, a GVRP BPDU is termed an *advertisement*. Advertisements are sent outbound from ports on a switch to the devices directly connected to those ports.

IEEE 802.1ad Specification that allows a service provider to assign a unique VLAN identifier (called the Service VLAN ID or S-VID) to customers using multiple VLANs, thereby extending the total number of VLANs that can be supported within the provider network.

IGMP The Internet Group Management Protocol (IGMP) is a communications protocol used by hosts and adjacent routers on IP networks to establish multicast group memberships.

Inactive stack fragment The switches in this fragment do not actively switch packets. They are powered on, however, the network ceases to carry traffic. All user ports are disabled. Only the OOBM and stack ports remain active.

Inbound port Any port on the switch through which traffic enters the switch.

IP options Optional fields supported in an IPv4 packet header.

IP precedence bits The upper three bits in the Type of Service (ToS) byte of an IPv4 packet.

IPv4 Version 4 of the IP protocol.

IPv6 Version 6 of the IP protocol.

IST Internal Spanning Tree. Administers the topology within a given MST region. When you configure a switch for MSTP operation, the switch automatically includes all of the static VLANs configured on the switch in a single, active spanning tree topology (instance) within the IST. This is termed the “IST instance”. Any VLANs you subsequently configure on the switch are added to this IST instance. To create separate forwarding paths within a region, group specific VLANs into different Multiple Spanning Tree Instances (MSTIs). (Refer to “Multiple Spanning Tree Instance (MSTI)”.)

IST instance 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.

Member In a stacking configuration, a switch in the stack that is neither a Commander nor a Standby. It serves as set of additional ports in the stack. There can be up to eight stack Members, for a maximum total of ten switches in the stack.

Member ID In a stacking configuration, he identification number of this switch in the stack. It can be set by the Commander or configured by the user. This ID is retained by the switch unless the switch is put into a different stack.

Mesh topology A stack topology in which every member of the stack has a direct connection to every other member. For example, you can have a four-member mesh where only three stacking ports are used on each switch. This topology supports up to five members.

Missing switch A switch that was previously a member of a stack, but which has been disconnected from the stack, or is still connected to the stack but is shut down. The missing switch retains its stacking configuration and can rejoin the stack when it is reconnected or rebooted.

Mixed VLAN mode device Device that supports both C-VLANs and S-VLANs. A device configured in `qinq mixedvlan` mode can do regular CVLAN switching/routing (standard bridge behavior) and can also serve as a provider edge device tunneling frames into and out of the provider network.

MST region A region that forms a multiple spanning tree domain and is a component of a single spanning tree domain within a network. It comprises the VLANs configured on physically connected MSTP switches. All switches in a given region must be configured with the same VLANs, the same Multiple Spanning Tree Instances (MSTIs), and the same MST configuration identifiers.

MSTI Multiple Spanning Tree Instance. A configurable spanning tree instance that comprises all static VLANs you specifically assign to it, and must include at least one VLAN.

MSTP Multiple Spanning Tree Protocol. A network supporting MSTP allows multiple spanning tree instances within configured regions, and a single spanning tree among regions, STP bridges, and RSTP bridges.

MSTP BPDU (MSTP Bridge Protocol Data Units) BPDUs that carry region-specific information, such as the region identifier (region name and revision number). If a switch receives an MSTP BPDU with a region identifier that differs from its own, then the port on which that BPDU was received is on the boundary of the region in which the switch resides.

MSTP bridge In this manual an MSTP bridge is a switch or other 802.1s compatible device, configured for MSTP operation.

NIC Network interface card.

Outbound packet A packet leaving the switch through any LAN port.

Outbound port Any port on the switch through which traffic leaves the switch.

Outbound port queue On any port, a buffer that holds outbound traffic until it is transmitted from the switch through the port.

Outbound port queue On any port, a buffer that holds outbound traffic until it is transmitted from the switch through the port.

Port-based interface Untagged customer-network ports or trunks on a QinQ enabled device.

Provider-network port Port on an S-VLAN bridge that connects to the provider network. This equates to PN ports of the IEEE 802.1ad standard.

Provisioned switch A switch for which the stack has been configured before it is connected into the stack. Once such a switch configuration has been created in the stack, a complete network configuration can be created for it, even if the switch is not physically present in the stack.

PVST Per VLAN Spanning Tree. HP switches that offer PVST protection and PVST filtering features, enabled per-port.

QinQ A feature that enables service providers to use a single VLAN-ID to support multiple customer VLANs by encapsulating the 802.1Q VLAN tag within another 802.1Q frame.

QoS policy Defines a policy configured in one of the following ways:

- Classifier-based: rate limiting and/or prioritizing packets in a specified traffic class, defined by the class command).
- Globally-configured: 802.1p priority and/or DSCP codepoint, with global QoS commands.

Re-marking Assigning a new QoS policy to an outbound packet by changing the:

- Class-of-Service (CoS) 802.1p bit setting in Layer 2 VLAN headers
- DSCP bit setting in the Layer 3 IPv4 ToS byte or IPv6 Traffic Class byte.

Ring topology A stack topology in which all the switches are connected via two connections on each Stacking Module installed in each switch.

RPVST+ BPDU RPVST+ Bridge Protocol Data Unit. These BPDUs carry region-specific information, such as the region identifier (region name and revision number). If a switch receives an RPVST+ BPDU with a region identifier that differs from its own, then the port on which that BPDU was received is on the boundary of the region in which the switch resides.

RPVST+ bridge A switch configured for RPVST+ operation.

RSTP Rapid Spanning Tree Protocol, defined in IEEE 802.1w and ratified in IEEE 802.1D-2004.

S-tagged interface Tagged customer-network ports or trunks on a QinQ enabled device.

S-VLAN Service VLANs that are used to tunnel customer frames through the provider network to customer sites. These are managed by the service provider who can assign each customer a unique S-VLAN ID.

S-VLAN bridge Provider-owned device configured in QinQ VLAN mode that uses S-VLANs only to forward frames in the provider network.

S-VLAN ID Unique S-VLAN identifier.

Service provider The provider of the network that provides one or more service instances to a customer.

SNMP Simple Network Management Protocol, used to remotely manage network devices.

Spanning tree Generic term to refer to the many spanning tree flavors: now deprecated STP, MSTP, RSTP and VLAN-aware RPVST+.

Stack The composite switch made of two or more switches in a stack configuration. The stack responds as a single entity and is managed as a single device through the Commander. It operates as if there is a single switch with all the additional switches operating to provide additional ports. The stack, at a minimum, consists of a Commander and a Standby switch.

Stack fragment A stack that previously had more members (that is, some of its previous members are now missing). The fragment can be Active or Inactive based on the rules described.

Stack revision A value is stored in the configuration to ensure that the Commander retains knowledge of the most current stack topology. This value is incremented whenever the stack configuration changes. For example, if a member is rebooted, a new member is added. The stack revision is determined by starting with the highest value among all members in the stack, and then increase it by 1 and then distribute that new value to all switches in the stack.

Standby In a stacking configuration, the physical switch that serves as a backup for the Commander. If the Commander fails over, the Standby takes control and becomes the Commander for the stack. There is only one Standby chosen at any given time in the stack.

Static VLAN A port-based or protocol-based VLAN configured in switch memory.

STP Spanning Tree Protocol. Part of the original IEEE 802.1D specification. The 2004 edition completely deprecates STP. Both RSTP and RPVST+ have fallback modes to handle STP. RPVST+ has a fallback mode to support PVST and RSTP has a fallback mode to support STP.

Switch mesh domain A group of meshed switch ports exchanging meshing protocol packets. Paths between these ports can have multiple redundant links without creating broadcast storms.

Switch priority A user-configured value used to influence which switches becomes the Command and Standby when a stack is initially booted. Note: Regardless of its priority configuration, when a new switch is added to a stack, the Commander and Standby switches are not changed unless forced by the user. This prevents unnecessary disruption to the network. The switch priority also influences which device will be chosen as the new Standby after a takeover

Tagged packet A packet that carries an IEEE 802.1Q VLAN ID (VID), which is a two-byte extension that precedes the source MAC address field of an Ethernet frame. A VLAN tag is layer 2 data and is transparent to higher layers.

Tagged port membership Identifies a port as belonging to a specific VLAN and enables VLAN-tagged packets to carry an 802.1p priority when sent from the port in outbound traffic. When a port is an untagged member of a VLAN, outbound packets belonging to the VLAN do not carry an 802.1p priority setting.

Tagged VLAN A VLAN that complies with the 802.1Q standard, including priority settings, and allows a port to join multiple VLANs. (*See also* Untagged VLAN.)

ToS Type-of-Service. One of the bytes in an IPv4 packet header. IPv4 packets may be classified according to two type-of service modes:

- IP-precedence mode, using the upper three bits of the ToS byte
- Differentiated Services (Diff-Serv) mode, using the upper six bits of the ToS byte

Traffic class byte In IPv6 packets, the byte that corresponds to the IPv4 Type-of-Service field, consisting of a six-bit (high order) Differentiated Services (Diff-Serv) field and a two-bit (low-order) reserved field. The three-bit (high-order) precedence field in the Traffic Class byte corresponds to the IP precedence bit set in IPv4 packets.

Tunnel VLAN *See* S-VLAN.

Untagged packet A packet that does not carry an IEEE 802.1Q VLAN ID (VID).

Untagged VLAN A VLAN that does not use or forward 802.1Q VLAN tagging, including priority settings. A port can be a member of only one untagged VLAN of a given type (port-based and the various protocol-based types).

Upstream device A device linked directly or indirectly to an inbound switch port. The switch receives traffic from upstream devices.

VID VLAN Identification Number. Each 802.1Q-compliant VLAN must have its own unique VID number, and that VLAN must be given the same VID in every device in which it is configured.

VOIP Voice-Over-IP.

VRRP Virtual Router Redundancy Protocol.

Index