Chapter 3

Switching

Introduction

This chapter gives an overview of Layer 1 (the physical layer), 2 (the data link layer), and 3 (the network layer) switching, and describes the support for switching and how to configure and operate the switching functions.

The switch, also referred to as a MAC (media access control) bridge, a data link relay or a level 2 relay, can connect multiple Local Area Network (LAN) segments together to form an extended LAN. Stations connected to different LANs can be configured to communicate with one another as if they were on the same LAN. It can also divide one physical LAN into multiple Virtual LANs (VLANs). Stations connected to each other on the same extended LAN can be grouped in separate VLANs, so that a station in one VLAN can communicate directly with other stations in the same VLAN, but must go through higher layer routing protocols to communicate with stations in other VLANs.

The switch operates at the data link layer, transparent to higher layer protocols, transferring frames between the data link layers of the networks to which it is attached. A bridge accesses each physical link according to the rules for that particular network. Access may not always be instant, so a bridge must be capable of storing and forwarding frames. Since the switch can store and forward frames, it can examine and discard or admit frames according to their VLAN tag fields. The switch can also examine the address fields of the frames and forward the frames based on knowledge of which network contains the station with an address matching the frame's destination address. In this way, the switch can act as an intelligent filtering device, redirecting or blocking the movement of frames between networks.

Because the switch may receive frames faster than it can forward them, the switch has Quality of Service queues in which frames await transmission according to their priority.

The switch can be used to:

- Increase the physical extent and/or the maximum number of stations on a LAN.

  LANs are limited in their physical extent by the signal distortion and propagation delay characteristics of the media. The switch overcomes this limitation by receiving a frame on one LAN and then retransmitting the frame on another LAN, using the normal access methods for each LAN. The physical characteristics of the LAN media also place a practical limit on the number of stations that can be connected to a single LAN segment. The switch overcomes this limitation by joining LAN segments together to form an extended LAN capable of supporting more stations than either of the individual LANs.

- Connect LANs that have a common data link layer protocol but different physical media, for example, Ethernet 10BASET, 100BASET and 10BASEF.

- Increase the availability of LANs by allowing multiple redundant paths to be physically configured, and selected dynamically, using the Spanning Tree algorithm.

- Reduce the load on a LAN or increase the effective bandwidth of a LAN, by filtering traffic.

- Prioritise the transmission of data with high Quality of Service requirements.

By using Virtual LANs (VLANs), a single physical LAN can be separated into multiple Virtual LANs. VLANs can be used to:

- Further improve LAN performance, as broadcast traffic is limited to LAN segments serving members of the VLAN to which the sender belongs.
- Provide security, as frames are forwarded to those stations belonging to the sender's VLAN, and not to stations in other VLANs on the same physical LAN.
- Reduce the cost of moving or adding stations to function or security based LANs, as this generally requires only a change in the VLAN configuration.

Switch Ports

The term *port* is used frequently in switch terminology. Each port in a switch is associated with one of the physical interfaces on the switch. Each port is uniquely identified by a port number. The switch supports a number of features at the physical level that allows it to be connected in a variety of physical networks. This physical layer (Layer 1) versatility includes:

- Enabling and disabling Ethernet ports.
- Autonegotiation of port speed and duplex mode for all 10/100 Ethernet ports and copper gigabit ports.
- Manual setting of port speed and duplex mode for all 10/100 Ethernet ports and copper gigabit ports.
- Port trunking.
- Packet storm protection.
- Port mirroring.
- Support for SNMP management.
- Link triggers for fibre ports.

Enabling and disabling switch ports

A switch port that is enabled is available for packet reception and transmission. Its administrative status in the Interfaces MIB is UP. Conversely, a port that is disabled is not available for packet reception and transmission. It does not send or receive frames and its administrative status in the Interfaces MIB is DOWN. Every port on the switch is enabled by default. A switch port that has been disabled by the Port Security feature cannot be enabled using the ENABLE SWITCH PORT command.

To enable or disable a switch port, use the commands:

```
ENABLE SWITCH PORT={port-list|ALL}
```

```
DISABLE SWITCH PORT={port-list|ALL}
```

Resetting ports at the hardware level discards all frames queued for reception or transmission on the port, and restarts autonegotiation of port speed and duplex mode. Ports are reset using the command:

```
RESET SWITCH PORT={port-list|ALL} [COUNTER]
```

To display information about switch ports, use the command:

```
SHOW SWITCH PORT[={port-list|ALL}]
```

Autonegotiation of port speed and duplex mode

Each of the switch ports can operate at either 10 Mbps or 100 Mbps, in either full duplex or half duplex mode. In full duplex mode, a port transmits and receives data simultaneously. In half duplex mode, the port either transmits or receives, but not at the same time. This versatility makes it possible to connect devices with different speeds and duplex modes to different ports on the switch. This versatility also requires that each port on the switch know which speed and mode to use.

Autonegotiation allows the ports to adjust their speed and duplex mode to accommodate devices connected to them. Each switch port can be either configured with a fixed speed and duplex mode, or configured to autonegotiate speed and duplex mode with a device connected to it to determine a speed and mode that allows successful transmission. An autonegotiating port adopts the speed and duplex mode required by devices connected to it. If another autonegotiating device is connected to the switch, they negotiate the highest possible common speed and duplex mode. Setting the port to a fixed speed and duplex mode allows it to support equipment that cannot autonegotiate.

If you override a port's autonegotiation by setting it to a fixed speed/duplex setting, automatic MDI/MDI-X detection is also overridden. The port defaults to MDI-X.

It is also possible to require a port to operate at a single speed without disabling autonegotiation by allowing the port to autonegotiate but constrain the speed/duplex options to the desired combination. For example, if one end of a link is set to AUTO and the other to 100MFULL, then the AUTO end selects 100MHALF operation because without the other end autonegotiating, the AUTO end has no way of knowing that the fixed end is full duplex capable. If a particular speed is required, it is better to fix the speed/duplex combination using one of the autonegotiating speed values. Therefore, using 100MFAUTO at one end of a link allows the AUTO end to autonegotiate 100MFULL.

Switch ports autonegotiate by default when they are connected to a new device. To change this setting, use the command:

```
SET SWITCH PORT={port-list|ALL} SPEED={AUTONEGOTIATE|10MHALF|
   10MFULL|10MHAUTO|10MFAUTO|100MHALF|100MFULL|100MHAUTO|
   100MFAUTO|1000MHALF|1000MFULL|1000MHAUTO|1000MFAUTO}
```

Settings available on different models are shown in Table 3-1 on page 3-7. Autonegotiation can also be activated at any time after this, on any port that is set to autonegotiate by using the command:

```
ACTIVATE SWITCH PORT={port-list|ALL} AUTONEGOTIATE
```

The SHOW SWITCH PORT command displays the port speed and duplex mode settings.

Table 3-1: Port speed and duplex settings for switch ports .

Speed	AT-8624T/2M 10/100
10MHALF	Yes
10MFULL	Yes
100MHALF	Yes
100MFULL	Yes
1000MHALF	No
1000MFULL	No
10MHAUTO	Yes
10MFAUTO	Yes
100MHAUTO	Yes
100MFAUTO	Yes
1000MHAUTO	No
1000MFAUTO	No
AUTONEGOTIATE	Yes

Port trunking

Port trunking, also known as port bundling or link aggregation, allows a number of ports to be configured to join together to make a single logical connection of higher bandwidth. This can be used where a higher performance link is required, and makes links even more reliable.

The switch supports static 802.3ad link aggregation, and is also compatible with third party devices that do not support static 802.3ad link aggregation.

The switch supports up to 6 trunk groups, of up to 8 switch ports each. The two gigabit Ethernet ports can also be grouped together to form a trunk group. For trunking to work properly, avoid having a trunk group that spans multiple switch instances. It is not possible for a trunk group to include both 10/100 Ethernet and gigabit Ethernet ports. Ports in the trunk group do not have to be contiguous.

Port trunk groups are created and destroyed on the switch using the commands:

```
CREATE SWITCH TRUNK=trunk [PORT=port-list] [SELECT={MACSRC|
   MACDEST|MACBOTH|IPSRC|IPDEST|IPBOTH}] [SPEED={10M|100M|
   1000M}]

DESTROY SWITCH TRUNK=trunk
```

Port trunk groups can be destroyed on the switch only when no ports belong to them.

All the ports in a trunk group must have the same VLAN configuration: they must belong to the same VLANs and have the same tagging status. All ports in a trunk group must be added to VLANs together, and can only be deleted from a VLAN as a group. Similarly, if the tagged or untagged status of the ports is changed, it must be changed for all ports in the trunk group at the same time.

The members of a trunk group can be specified when it is created, and ports can be added to or removed from a trunk group using the commands:

```
ADD SWITCH TRUNK=trunk PORT=port-list

DELETE SWITCH TRUNK=trunk PORT={port-list|ALL}
```

Ports in a trunk group are set to autonegotiate at the trunk speed at full duplex. When a port is added to a trunk group, the speed setting for the group overrides the speed setting previously configured for the port. When a port is removed from a trunk group, the port returns to its previously configured speed and duplex mode settings.

The speed of the trunk group can either be specified when it is created, or set using the command:

```
SET SWITCH TRUNK=trunk [SELECT={MACSRC|MACDEST|MACBOTH|IPSRC|
   IPDEST|IPBOTH}] [SPEED={10M|100M|1000M}]
```

To display information about trunks on the switch, use the command:

```
SHOW SWITCH TRUNK[=trunk]
```

To display the VLANs to which the ports in the trunk groups belong, use the command:

```
SHOW VLAN[=ALL]
```

Port trunking must be configured on both ends of the link, or network loops may result.

Packet storm protection

The packet storm protection feature allows the user to set limits on the reception rate of broadcast, multicast and destination lookup failure packets. The software allows separate limits to be set for each port, beyond which each of the different packet types are discarded. The software also allows separate limits to be set for each of the packet types. Which of these options can be implemented depends on the model of switch hardware.

By default, packet storm protection is set to NONE, that is, disabled. It can be enabled, and each of the limits can be set using the command:

```
SET SWITCH PORT=port-list [BCLIMIT={NONE|limit}]
   [DLFLIMIT={NONE|limit}] [MCLIMIT={NONE|limit}]
```

Packet storm protection limits cannot be set for each individual port on the switch, but can be set for each processing block of ports. The processing blocks are sets of 8 ports (e.g. as many as are applicable of ports 1-8, 9-16 and 17-24) and each uplink port is a further processing block. Therefore, a 16-port switch has four processing blocks and a 24-port switch has five. The two uplink ports are numbered sequentially after the last port, and therefore are 17 and 18 for a 16-port switch, and 25 and 26 for a 24-port switch. Only one limit can be set per processing block, and then applies to all three packet types. Thus each of the packet types are either limited to this value or unlimited (NONE).

The SHOW SWITCH PORT command displays the packet storm protection settings.

Port mirroring

Port mirroring allows traffic being received and transmitted on a switch port to be sent to another switch port, the mirror port, usually for the purposes of capturing the data with a protocol analyser. This mirror port is the only switch port that belongs to no VLANs, and therefore does not participate in any other switching. Before the mirror port can be set, it must be removed from all VLANs except the default VLAN. The port cannot be part of a trunk group.

To set the mirror port (and remove it from the default VLAN) use the command:

```
SET SWITCH MIRROR={NONE|port}
```

If another port was previously set as the mirror port, this command returns the previous mirror port to the default VLAN as an untagged port. Return this port to any VLANs to which it should belong, by using the ADD VLAN PORT command, or set it as a tagged port using the SET VLAN PORT command if required.

Either traffic received on a port or traffic transmitted by the port, or both, can be mirrored. This setting and the source port(s) from which traffic is sent to the mirror port are specified using the command:

```
SET SWITCH PORT={port-list|ALL} MIRROR={NONE|RX|TX|BOTH}
```

Mirroring four or more ports may significantly reduce switch performance.

To send packets that match particular criteria to the mirror port, first create a filter match using the command:

```
ADD SWITCH L3FILTER MATCH
```

Then create a filter entry with the ACTION parameter set to SENDMIRROR, using the command:

```
ADD SWITCH L3FILTER=filter-id ENTRY ACTION=SENDMIRROR.
```

By default, when mirroring is disabled, no mirror port is set and no source ports are set to be mirrored. Mirroring functions when a switch mirror port is set to a valid port. When mirroring is enabled and the switch mirror port is set to NONE, then mirroring can be disabled using the commands:

```
ENABLE SWITCH MIRROR

DISABLE SWITCH MIRROR
```

The SHOW SWITCH PORT and SHOW SWITCH commands display the switch and port mirroring settings.

Port security

The port security feature allows control over the stations connected to each switch port, by MAC address. If enabled on a port, the switch learns MAC addresses up to a user-defined limit from 1 to 256, then locks out all other MAC addresses. One of the following options can be specified for the action taken when an unknown MAC address is detected on a locked port:

- Discard the packet and take no further action,
- Discard the packet and notify management with an SNMP trap,
- Discard the packet, notify management with an SNMP trap and disable the port.

To enable port security on a port, set the limit for learned MAC addresses to a value greater than zero, and specify the action to take for unknown MAC addresses on a locked port. To disable port security on a port, set the limit for learned MAC addresses to zero or NONE. Port security can be enabled or disabled on a port using the command:

```
SET SWITCH PORT={port-list|ALL} LEARN={NONE|0|1..256}
   [INTRUSIONACTION={DISCARD|TRAP|DISABLE}]
```

If INTRUSIONACTION is set to TRAP or DISABLE, a list of MAC addresses for devices that are active on a port, but which are not allowed or learned for the port, can be displayed (Figure 3-25 on page 3-147) using the command:

```
SHOW SWITCH PORT={port-list|ALL} INTRUSION
```

A switch port can be manually locked before it reaches the learning limit by using the command:

```
ACTIVATE SWITCH PORT={port-list|ALL} LOCK
```

Addresses can be manually added to a port locked list up to a total of 256 MAC addresses, and the learning limit can be extended to accommodate them. Use the command:

```
ADD SWITCH FILTER ACTION={FORWARD|DISCARD} DESTADDRESS=macadd
   PORT=port [ENTRY=entry] [LEARN] [VLAN={vlan-name|1..4094}]
```

Learned addresses on locked ports can be saved as part of the switch configuration, so that they become part of the configuration after a power cycle. Use the command:

```
CREATE CONFIG=filename
```

If the configuration is not saved when there is a locked list for a port, the learning process begins again after the router is restarted.

Virtual Local Area Networks (VLANs)

A Virtual LAN (VLAN) is a logical, software-defined subnetwork. It allows similar devices on the network to be grouped together into one broadcast domain, irrespective of their physical position in the network. Multiple VLANs can be used to group workstations, servers, stacks, and other network equipment connected to the switch, according to similar data and security requirements.

Decoupling logical broadcast domains from the physical wiring topology offers several advantages, including the ability to:

- Move devices and people with minimal, or no, reconfiguration
- Change a device's broadcast domain and access to resources without physically moving the device, by software reconfiguration or by moving its cable from one switch port to another
- Isolate parts of the network from other parts by placing them in different VLANs
- Share servers and other network resources without losing data isolation or security
- Direct broadcast traffic to only those devices that need to receive it thereby reducing traffic across the network
- Connect 802.1q-compatible switches together through one port on each switch

Devices that are members of the same VLAN exchange data with each other through the switch's switching capabilities. To exchange data between devices in separate VLANs, the switch's routing capabilities are used. The switch passes VLAN status information, indicating whether a VLAN is up or down, to the Internet Protocol (IP) module. IP uses this information to determine route availability.

The switch has a maximum of 255 VLANs, ranging from a VLAN identifier (VID) of 1 to 4094.

When the switch is first powered up, a "default" VLAN is created and all ports are added to it. In this initial unconfigured state, the switch broadcasts all the packets it receives to the default VLAN. This VLAN has a VID of 1 and an interface name of vlan1. It cannot be deleted, and ports can be removed from it only when they also belong to at least one other VLAN. When all devices on the physical LAN belong to the same logical LAN (same broadcast domain), the default settings are acceptable and no additional VLAN configuration is necessary.

VLAN Tagging

An Ethernet packet can contain a *VLAN tag* with fields that specify VLAN membership and user priority. The VLAN tag is described in IEEE Standard 802.3ac, and is four octets that can be inserted between the Source Address and the Type/Length fields in the Ethernet packet (Figure 3-1 on page 3-12). To accommodate the tag, IEEE 802.3ac also increased the maximum allowable length for an Ethernet frame to 1522 octets (the minimum size is 64 octets). IEEE 802.1q specifies how the data in the VLAN tag switches frames. VLAN-aware devices are able to add the VLAN tag to the packet header. VLAN-unaware devices cannot set or read the VLAN tag.

Table 3-2 on page 3-12 lists the meaning and use of the fields in the Ethernet frame. Figure 3-1 on page 3-12 shows the format of VLAN data in an Ethernet frame. Twelve bits of the tag are the VLAN Identifier (VID), which indicates the VLAN to which the packet belongs. Table 3-3 on page 3-12 lists the VLAN Identifier values that have specific meaning.

Table 3-2: Fields in the Ethernet frame for QoS and VLAN switching .

Field	Length	Meaning and use
TPID	2 octets	The Tag Protocol Identifier (TPID) is defined by IEEE Standard 802.1q as 0x81-00.
User Priority	3 bits	The User Priority field is the priority tag for the frame, which can be used by the switch to determine the Quality of Service to apply to the frame. The three bit binary number represents eight priority levels, 0 to 7.
CFI	1 bit	The Canonical Format Indicator (CFI flag) indicates whether all MAC address information that may be present in the MAC data carried by the frame is in canonical format.
VID	12 bits	The VLAN Identifier (VID) field uniquely identifies the VLAN to which the frame belongs.

Figure 3-1: Format of user priority and VLAN data in an Ethernet frame.

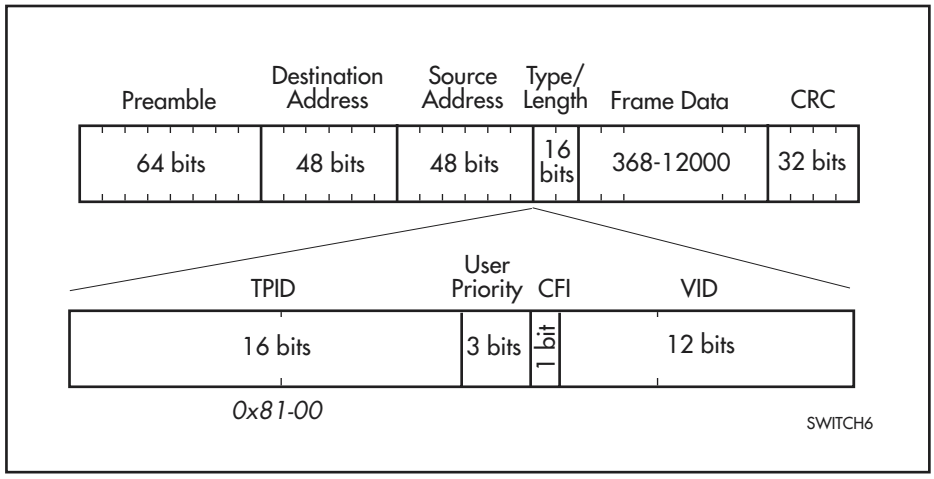

Table 3-3: Reserved VID values .

VID value (hexadecimal)	Meaning and use of reserved VID values
0	The null VLAN ID. Indicates that the tag header contains only user priority information; no VLAN Identifier is present in the frame. This VID value must not be configured in any Forwarding Database entry, or used in any management operation. Frames that contain the null VLAN ID are also known as priority-tagged frames.
1	The default VID value used for classifying frames on ingress through an untagged switch port.
FFF	Reserved for implementation use. This VID value must not be configured in any Forwarding Database entry, used in any management operation, or transmitted in a tag header.

Ethernet packets that contain a VLAN tag are referred to as *tagged* frames, and switch ports that transmit tagged frames are referred to as *tagged ports*. Ethernet packets that do not contain a VLAN tag are referred to as *untagged* frames, and switch ports that transmit untagged frames are referred to as *untagged ports*. VLANs can consist of simple logical groupings of untagged ports in which the ports receive and transmit untagged packets. Alternatively, VLANs can contain only tagged ports or a mixture of tagged and untagged ports.

The switch is VLAN aware. It can accept VLAN tagged frames, and supports the VLAN switching required by such tags. A network can contain a mixture of VLAN aware devices, for example, other 802.1q-compatible switches, and VLAN unaware devices, for example, workstations and legacy switches that do not support VLAN tagging. The switch can be configured to send VLAN tagged or untagged frames on each port, depending on whether the devices connected to the port are VLAN aware. By assigning a port to two different VLANs, to one as an untagged port and to another as a tagged port, it is possible for the port to transmit both VLAN-tagged and untagged frames. A port must belong to a VLAN at all times unless the port has been set as the mirror port for the switch.

Every frame admitted by the switch has a VID associated with it. When a frame arrives on a tagged port, the associated VID is determined from the VLAN tag the frame had when it arrived. When a frame arrives on an untagged port, it is associated with the VID of the VLAN for which the incoming port is untagged. When the switch forwards a frame over a tagged port, it adds a VLAN tag to the frame. When the switch forwards the frame over an untagged port, it transmits the frame as a VLAN-untagged frame, not including the VID in the frame.

The VLAN tag that the switch adds to a frame on egress depends on whether the frame is switched in Layer 2 or Layer 3. In Layer 3 switching, the switch determines the destination VLAN from its routing tables. The VID of the destination VLAN is added to the frame on egress. In Layer 2 switching, the frame's source and destination VLANs are the same. The VID that was associated with the frame on ingress is associated with it on egress.

VLAN Membership using VLAN Tags

Ports can belong to many VLANs as tagged ports. Because VLAN tags determine to which VLAN a packet belongs, it is easy to:

- Share network resources, such as servers and printers, across several VLANs
- Configure VLANs that span several switches

For tagged ports, the switch uses the VID of incoming frames, and the frame's destination field to switch traffic through a VLAN aware network. Frames are transmitted only on ports belonging to the required VLAN. Other vendors' VLAN-aware devices on the network can be configured to accept traffic from one or more VLANs. A VLAN-aware server can be configured to accept traffic from many different VLANs, and then return data to each VLAN without mixing or leaking data into the wrong VLANs.

Figure 3-2 on page 3-14 shows a network configured with VLAN tagging. Table 3-4 on page 3-14 shows the VLAN membership. The server on port 2 on Switch A belongs to both the *admin* and *marketing* VLANs. The two switches are connected through uplink port 26 on Switch A and uplink port 25 on

Switch B, which belong to both the *marketing* VLAN and the *training* VLAN, so devices on both VLANs can use this link.

Figure 3-2: VLANs with tagged ports.

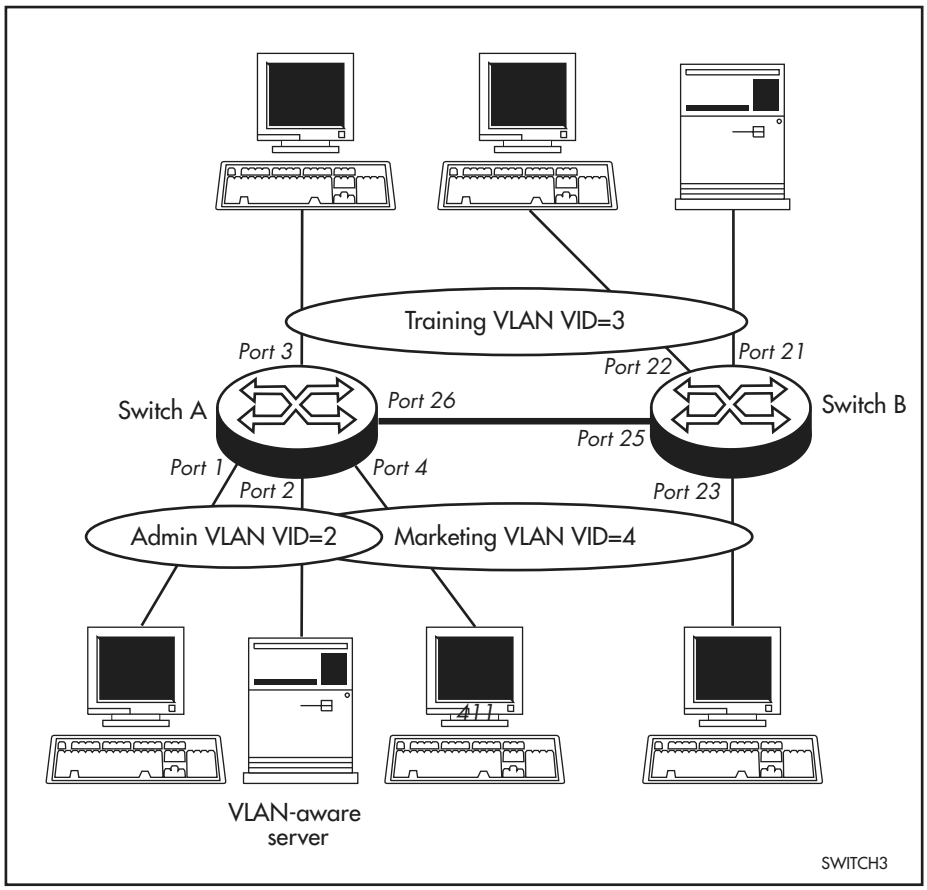

Table 3-4: VLAN membership of example of a network using tagged ports .

VLAN	Member ports
Training	3, 26 on Switch A 21, 22, 25 on Switch B
Marketing	2, 4, 26 on Switch A 23, 25 on Switch B
Admin	1, 2 on Switch A

VLAN Membership of Untagged Packets

A VLAN that does not send VLAN-tagged frames is a logical grouping of ports. All untagged traffic arriving at those ports belongs to that VLAN.

VLANs based on untagged ports are limited because each port can belong only to one VLAN as an untagged port. Limitations include:

- It is difficult to share network resources, such as servers and printers, across several VLANs. The routing functions in the switch must be configured to interconnect using untagged ports only.
- A VLAN that spans several switches requires a port on each switch for the interconnection of the various parts of the VLAN. When there are several VLANs in the switch that span more than one switch, then many ports are occupied with connecting the VLANs, and so are unavailable for other devices.

If the network includes VLANs that do not need to share network resources or span several switches, VLAN membership can usefully be based on untagged ports. Otherwise, VLAN membership should be determined by tagging (see *VLAN Tagging on page 3-11*).

Figure 3-3 on page 3-15 shows two port-based VLANs with untagged ports. Ports 1-3 belong to the *marketing* VLAN, and ports 14-16 belong to the *training* VLAN. The switch acts as two separate bridges: one that forwards traffic between the ports belonging to the *marketing* VLAN, and a second one that forwards traffic between the ports belonging to the *training* VLAN. Devices in the *marketing* VLAN can communicate with devices in the *training* VLAN only by using the switch's routing functions.

Figure 3-3: VLANS with untagged ports.

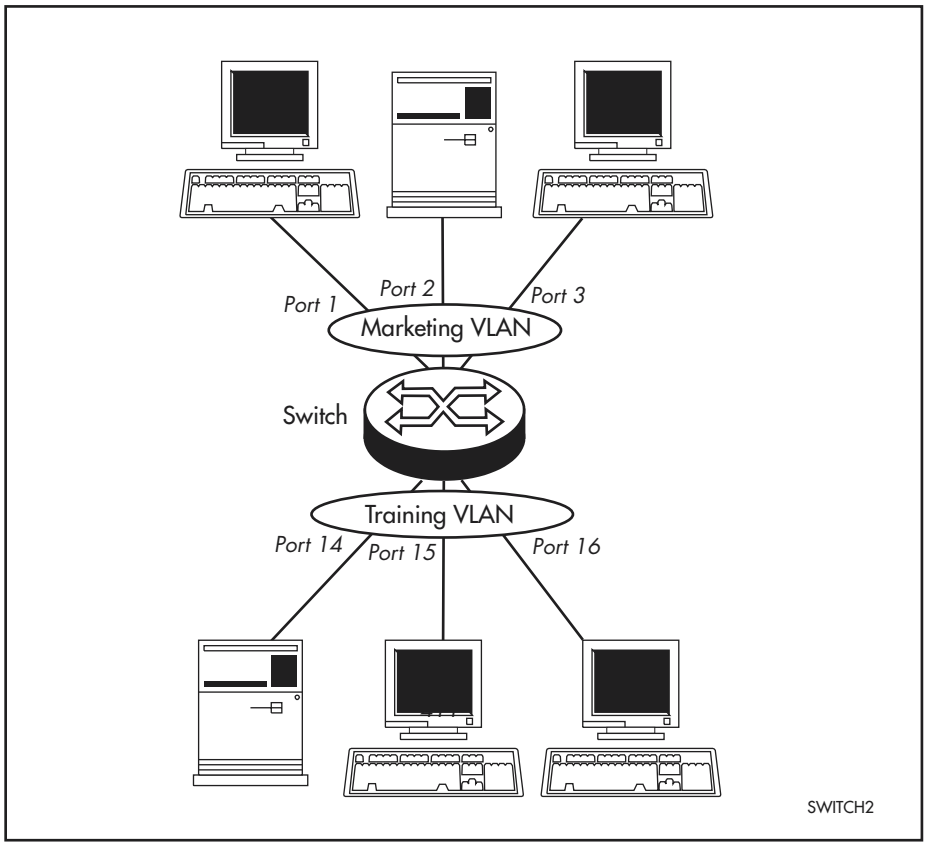

Creating VLANs

To summarise the process of creating a VLAN:

1. Create the VLAN.
2. Add tagged ports to the VLAN, if required.
3. Add untagged ports to the VLAN, if required.

To create a VLAN, use the command:

```
CREATE VLAN=vlan-name VID=2..4094
```

Every port must belong to a VLAN unless it is the mirror port. By default, all ports belong to the default VLAN as untagged ports.

To add tagged ports to a VLAN, use the command:

```
ADD VLAN={vlan-name|1..4094} PORT={port-list|ALL}
    FRAME=TAGGED
```

A port can be tagged for any number of VLANs.

To add untagged ports to a VLAN, use the command:

```
ADD VLAN={vlan-name|1..4094} PORT={port-list|ALL}
    [FRAME=UNTAGGED]
```

A port can be untagged for zero or one VLAN. A port can be added only to the default VLAN as an untagged port when it is not untagged for another VLAN. A port cannot transmit both tagged and untagged frames for the same VLAN (that is, it cannot be added to a VLAN as both a tagged and an untagged port).

To remove ports from a VLAN, use the command:

```
DELETE VLAN={vlan-name|1..4094} PORT={port-list|ALL}
```

Removing an untagged port from a VLAN returns it to the default VLAN unless it is a tagged port for another static VLAN. An untagged port can be deleted from the default VLAN only when the port is a tagged port for another static VLAN.

Ports tagged for some VLANs and left in the default VLAN as untagged ports transmit broadcast traffic for the default VLAN. If this is not required, the unnecessary traffic in the switch can be reduced by deleting those ports from the default VLAN.

To change the tagging status of a port in a VLAN, use the command:

```
SET VLAN={vlan-name|1..4094} PORT={port-list|ALL}
    FRAME=TAGGED
```

To destroy a VLAN, use the command:

```
DESTROY VLAN={vlan-name|2..4094|ALL}
```

VLANs can be destroyed only when no ports belong to them.

To display the VLANs configured on the switch, use the command:

```
SHOW VLAN[={vlan-name|1..4094|ALL}]
```

Information that may be useful for troubleshooting a network can be displayed with the VLAN debugging mode. This is disabled by default, and can be enabled for a specified time, disabled, and displayed using the commands:

```
ENABLE VLAN={vlan-name|1..4094|ALL} DEBUG={PKT|ALL}
   [OUTPUT=CONSOLE] [TIMEOUT={1..4000000000|NONE}]

DISABLE VLAN={vlan-name|1..4094|ALL} DEBUG={PKT|ALL}

SHOW VLAN DEBUG
```

To view packet reception and transmission counters for a VLAN, use the command (see the *Interfaces* chapter of the switch's Software Reference):

```
SHOW INTERFACE=VLANn COUNTER
```

Summary of VLAN tagging rules

When designing a VLAN and adding ports to VLANs, consider the following rules:

1. Except for the mirror port, each port must belong to at least one static VLAN. By default, a port is an untagged member of the default VLAN.
2. A port can be untagged for zero or one VLAN. A port that is untagged for a VLAN transmits frames destined for that VLAN without a VLAN tag in the Ethernet frame.
3. A port can be tagged for zero or more VLANs. A port that is tagged for a VLAN transmits frames destined for that VLAN with a VLAN tag, including the numerical VLAN Identifier of the VLAN.
4. A port cannot be untagged and tagged for the same VLAN.
5. The mirror port, if present, is not a member of any VLAN.

VLAN Interaction with Trunk Groups

All the ports in a trunk group must have the same VLAN configuration. They must belong to the same VLANs and have the same tagging status; and they must be operated on as a group.

Static and dynamic VLANs

All VLANs created by the user on the command line are static VLANs. The default VLAN is also a static VLAN. A port must belong to at least one static VLAN.

Dynamic VLANs are created by GVRP, a GARP application whose purpose is to propagate VLAN information between VLAN aware switches (see the *Generic Attribute Registration Protocol (GARP)* chapter). These dynamic VLANs are entitled gvrpxxx, where xxx is the VLAN's VLAN Identifier. Dynamic VLANs are created only when GVRP is enabled on the switch. GVRP is disabled by default.

All static VLANs except for the default VLAN can be destroyed by the user. Dynamic VLANs cannot be directly destroyed by the user, but may be destroyed according to the operations of GVRP by using the RESET GARP command on page 5-12 of *Chapter 5, Generic Attribute Registration Protocol (GARP)* or by disabling the GVRP instance.

A user can add, delete, or modify ports for a static VLAN, but not for a dynamic VLAN. Dynamic VLANs created by GVRP include only tagged ports.

Protected VLANs

If a VLAN is protected, Layer 2 traffic between ports that are members of a protected VLAN is blocked. Traffic can be Layer 3 switched to another VLAN. This feature prevents members of a protected VLAN from communicating with each other yet still allows members to access another network. Layer 3 Routing between ports in a protected VLAN can be prevented by adding a Layer 3 filter. The protected VLAN feature also allows all of the members of the protected VLAN to be in the same subnet.

A typical application is a hotel installation where each room has a port that can be used to access the Internet. In this situation it is undesirable to allow communication between rooms.

To create a protected VLAN, use the command:

```
CREATE VLAN=vlan-name VID=2..4094 [PROTECTED]
```

VLAN Relaying

VLAN relaying allows the passage of traffic between the VLANs on one switch, for protocols that are not processed by the switch's routing functions. Particular protocols or protocol groups can be specified, and filtering occurs on the basis of protocol identification number. VLAN relaying is similar to the bridging function of an Allied Telesyn router.

Protocol names have been predefined for many protocol types. Those protocols that are transferred by VLAN relay and that have predefined names are given in Table 3-5 on page 3-18, with their associated protocol identification numbers. Other protocols can be specified by entering their protocol identification numbers. Protocols that are routed by the switch, including IP, IPX, AppleTalk, STP and GARP, cannot be VLAN relayed.

Table 3-5: Predefined protocol types implemented by VLAN relay .

Protocol Name	Protocol Number	Encapsulation
All802	all SAP protocols	SAP
Netbeui	F0	SAP
SNA Path Control	04	SAP
PROWAY-LAN	0E	SAP
EIA-RS	4E	SAP
PROWAY	8E	SAP
ISO CLNS IS	FE	SAP
AllEthII	all EthII protocols	EthII
XEROX PUP	0200	EthII
PUP Addr Trans	0201	EthII
XEROX NS IDP	0600	EthII

Table 3-5: Predefined protocol types implemented by VLAN relay (Continued).

Protocol Name	Protocol Number	Encapsulation
X.75 Internet	0801	EthII
NBS Internet	0802	EthII
ECMA Internet	0803	EthII
Chaosnet	0804	EthII
X.25 Level 3	0805	EthII
XNS Compat	0807	EthII
Banyan Systems	0BAD	EthII
BBN Simnet	5208	EthII
DEC MOP Dump/Ld	6001	EthII
DEC MOP Rem Cons	6002	EthII
DEC LAT	6004	EthII
DEC Diagnostic	6005	EthII
DEC Customer	6006	EthII
DEC LAVC	6007	EthII
RARP	8035	EthII
DEC LANBridge	8038	EthII
DEC Encryption	803D	EthII
IBM SNA	80D5	EthII
SNMP	814C	EthII
AllSNAP	all SNAP protocols	SNAP

VLAN relaying operates in three stages:

1. The user creates one or more VLAN relay entities and adds the required VLANs and protocols to each entity.
2. The VLAN relay entity attaches to each specified VLAN and receives traffic. If more than one VLAN relay entity is attached to the same VLAN for the same protocol type, an intermediate attachment level receives the packet, duplicates it, and sends it to separate VLAN relay entities as required.
3. The VLAN relay entity sends the packet to the appropriate destination VLAN. Destination addresses are determined from the switch's learned address tables. If the destination address cannot be found, the packet is sent to all ports on all VLANs that are part of the VLAN relay entity. If the packet is destined for the VLAN on which it was received, the relaying entity does not send it to that VLAN because the packet causes a destination lookup failure, and the switch itself sends the packet to all ports in the VLAN.

Configuring VLAN relaying

To configure VLAN relaying on the switch, first create a VLAN relay entity and give it a unique name, using the command:

```
CREATE VLANRELAY=name
```

An existing VLAN relay entity can be disabled or destroyed using the commands:

```
DISABLE VLANRELAY=name

DESTROY VLANRELAY=name
```

In many networks, only one VLAN relay entity is required. The following configurations are examples of situations when more than one VLAN relay entity is used.

- If a number of protocols and VLANs are part of VLAN relaying but not all protocols on all VLANs, then setting up a number of VLAN relay entities allows only relevant protocols and VLANs to be part of relaying.
- If traffic is to be relayed between certain VLANs but not others (for example, between VLAN 1 and VLAN 2, and between VLAN 1 and VLAN 3, but not between VLAN 2 and VLAN 3), then separate VLAN relay entities are required.

To initiate relaying, add the VLANs which packets are to be sent between, and the desired protocols, to the VLAN relay entity, using the command:

```
ADD VLANRELAY=name [PROTOCOL=protocoltype] [VLAN={vlan-name|
   1..4094}]
```

Protocols are specified by protocol type and number, or by allowing all protocols of a certain type. A predefined list of common protocols is provided in Table 3-5 on page 3-18.

VLANs and/or protocols can be removed from an existing VLAN relay entity using the command:

```
DELETE VLANRELAY=name [PROTOCOL=protocoltype] VLAN=[{vlan-
   name|1..4094}]
```

A count of the packets relayed by the VLAN relay entity or entities, which shows the packets relayed from and to each VLAN, can be displayed using the command:

```
SHOW VLANRELAY[=name]
```

The traffic being relayed, including the source and destination VLANs and the relevant VLAN relay entity, can be displayed using the command:

```
ENABLE VLANRELAY DEBUG
```

VLAN relay debugging can be disabled using the command:

```
DISABLE VLANRELAY DEBUG
```

Debugging is disabled by default. It can be enabled for one specified VLAN relay entity, and can be disabled for all entities or for a specified entity.

The Layer 2 Switching Process

The Layer 2 switching process comprises related but separate processes. The *Ingress Rules* admit or discard frames based on their VLAN tagging. The *Learning Process* learns the MAC addresses and VLAN membership of frames admitted on each port. The *Forwarding Process* determines to which ports the frames are forwarded, and the *Quality of Service* priority with which they are transmitted. Finally, the *Egress Rules* determine for each frame whether VLAN tags are included in the Ethernet frames that are transmitted. These processes assume that each station on the extended LAN has a unique data link layer address, and that all data link layer frames have a header that includes the source (sender's) MAC address and destination (recipient's) MAC address.

The Ingress Rules

When a frame first arrives at a port, the Ingress Rules for the port check the VLAN tagging in the frame to determine whether to discard it or forward it to the Learning Process.

The first check depends on whether the *Acceptable Frame Types* parameter is set to *Admit All Frames* or to *Admit Only VLAN Tagged Frames*. A port that transmits only VLAN tagged frames is automatically set to Admit Only VLAN Tagged Frames regardless of the VLAN to which the port belongs. The user cannot change this setting. Frames with a null numerical VLAN Identifier (VID) are VLAN-untagged frames or frames with priority tagging only.

Every frame received by the switch must be associated with a VLAN. When a frame is admitted by the Acceptable Frame Types parameter, the second part of the Ingress Rules associates each untagged frame admitted with the VID of the VLAN for which the port is untagged.

Every port belongs to one or more VLANs so every incoming frame has a VID that shows to which VLAN it belongs. The final part of the Ingress Rules depends on whether *Ingress Filtering* is enabled for the port. If Ingress Filtering is disabled, all frames are passed on to the Learning Process, regardless of which VLAN they belong to. If Ingress Filtering is enabled, frames are admitted only when they have the VID of a VLAN to which the port belongs. Otherwise, they are discarded.

The default settings for the Ingress Rules are to Admit All Frames, and for Ingress Filtering to be OFF. This means that if no VLAN configuration has been done, all incoming frames pass on to the Learning Process, regardless of whether not they are VLAN tagged. The parameters for each port's Ingress Rules can be configured using the command:

```
SET SWITCH PORT={port-list|ALL} [ACCEPTABLE={VLAN|ALL}]
   [INFILTERING={ON|OFF}]
```

The Learning Process

The Learning Process uses an *adaptive learning* algorithm, sometimes called *backward learning*, to discover the location of each station on the extended LAN.

All frames admitted by the Ingress Rules on any port are passed on to the Forwarding Process if they are for destinations within the same VLAN. Frames destined for other VLANs are passed to the layer three protocol, for instance IP. For every frame admitted, the frame's source MAC address and numerical VLAN Identifier (VID) are compared with entries in the Forwarding Database for the VLAN (also known as a MAC address table, or a forwarding table) maintained by the switch. The Forwarding Database contains one entry for every unique station MAC address the switch knows in each VLAN.

If the frame's source address is not already in the Forwarding Database for the VLAN, the address is added and an ageing timer for that entry is started. If the frame's source address is already in the Forwarding Database, the ageing timer for that entry is restarted. By default, switch learning is enabled, and it can be disabled or enabled using the commands:

```
DISABLE SWITCH LEARNING

ENABLE SWITCH LEARNING
```

If the ageing timer for an entry in the Forwarding Database expires before another frame with the same source address is received, the entry is removed from the Forwarding Database. This prevents the Forwarding Database from being filled up with information about stations that are inactive or have been disconnected from the network, while ensuring that entries for active stations are kept alive in the Forwarding Database. By default, the ageing timer is enabled, and it can be disabled or enabled using the commands:

```
ENABLE SWITCH AGEINGTIMER

DISABLE SWITCH AGEINGTIMER
```

If switch learning is disabled and the ageing timer has aged out all dynamically learned filter entries, only statically entered MAC source addresses are used to decide which packets to forward or discard. If the switch finds no matching entries in the Forwarding Database during the Forwarding Process, then all switch ports in the VLAN are flooded with the packet, except the port on which the packet was received.

The default of the ageing timer is 300 seconds (5 minutes), and this can be modified using the command:

```
SET SWITCH AGEINGTIMER=10..1000000
```

The Forwarding Database relates a station's (source) address to a port on the switch, and is used by the switch to determine from which port to transmit frames with a destination MAC address matching the entry in the station map.

To display the contents of the Forwarding Database, use the command:

```
SHOW SWITCH FDB [ADDRESS=macadd] [DISCARD={SOURCE|
    DESTINATION}] [HIT={YES|NO}] [L3={YES|NO}]
    [PORT={portlist|ALL}] [STATUS={STATIC|DYNAMIC}]
    [VLAN={vlan-name|1..4094}]
```

To display general switch settings, including settings for switch learning and the switch ageing timer, use the command:

```
SHOW SWITCH
```

The Forwarding Process

The Forwarding Process forwards received frames that are to be relayed to other ports in the same VLAN, filtering out frames on the basis of information contained in the station map and on the state of the ports. When a frame is received on the port for a destination in a different VLAN, it is either Layer 3 switched if it is an IP packet, or looked up in the Layer 3 routing tables.

Forwarding occurs only when the port on which the frame was received is in the Spanning Tree 'Forwarding' or 'Disabled' states. The destination address is then looked up in the Forwarding Database for the VLAN. If the destination address is not found, the switch floods the frame on all ports in the VLAN except the port on which the frame was received. If the destination address is found, the switch discards the frame if the port is not in the STP 'Forwarding' or 'Disabled' states, if the destination address is on the same port as the source address, or if there is a static filter entry for the destination address set to DISCARD (see *Layer 2 Filtering on page 3-24*). Otherwise, the frame is forwarded on the indicated port.

This whole process can further be modified by the action of static switch filters. These are configurable filters that allow switched frames to be checked against a number of entries.

The Forwarding Process provides storage for queued frames to be transmitted over a particular port or ports. More than one transmission queue may be provided for a given port. The transmission queue where a frame is sent is determined by the user priority tag in the Ethernet frame and the Quality of Service mapping (see *Quality of Service on page 3-23*).

Quality of Service

The switch hardware has a number of Quality of Service (QOS) *egress queues* that can be used to give priority to the transmission of some frames over other frames on the basis of their user priority tagging. The user priority field in an incoming frame (with value 0 to 7) determines which of the eight priority levels the frame is allocated. When a frame is forwarded, it is sent to a QOS egress queue on the port determined by the mapping of priority levels to QOS egress queues. All frames in the first QOS queue are sent before frames in the second QOS egress queue, and so on, until frames in the last QOS egress queue, which are sent when there are no frames waiting to be sent in any of the higher QOS egress queues.

The mapping between user priority and a QOS egress queue can be configured using the command:

```
SET SWITCH QOS=P0,P1,P2,P3,P4,P5,P6,P7
```

The switch has four QOS egress queues. It has a default mapping of priority levels to QOS egress queues as defined in *IEEE 802.1q* (Table 3-25 on page 3-114).

Table 3-6: Default priority level to queue mapping for four QOS egress queues.

Priority level	QOS Egress Queue
0	1
1	0
2	0
3	1
4	2
5	2
6	3
7	3

To display the mapping of user priority to QOS egress queues, use the command:

```
SHOW SWITCH QOS
```

The Egress Rules

After the Forwarding Process determines the ports and transmission queues from which a frame is forwarded, the Egress Rules for each port determine whether the outgoing frame is VLAN-tagged with its numerical VLAN Identifier (VID).

When a port is added to a VLAN, it is configured to transmit either untagged or VLAN tagged packets, using the command:

```
ADD VLAN={vlan-name|1..4094} PORT={port-list|ALL}
    [FRAME={TAGGED|UNTAGGED}]
```

This setting can be changed for a port that is already part of a VLAN, using the command:

```
SET VLAN={vlan-name|1..4094} PORT={port-list|ALL}
    FRAME={UNTAGGED|TAGGED}
```

Layer 2 Filtering

The switch has a Forwarding Database, entries that determine whether frames are forwarded or discarded over each port. Entries in this Forwarding Database are created dynamically by the Learning Process. A dynamic entry is automatically deleted from the Forwarding Database when its ageing timer expires. Filtering is specified in the IEEE 802.1d.

The user can configure static switch filter entries using the command line interface. Static switch filter entries associate a MAC address with a VLAN and a port in the VLAN. When the switch receives a frame with a destination address and VLAN Identifier that match those of a static filter entry, the frame can be either forwarded to the port specified in the static filter entry, or discarded.

The Forwarding Database supports queries by the Forwarding Process as to whether frames with given values of the destination MAC address field should be forwarded to a given port.

To add or delete a static switch filter entry, use the command:

```
ADD SWITCH FILTER ACTION={FORWARD|DISCARD} DESTADDRESS=macadd
   PORT=port [ENTRY=entry] [LEARN] [VLAN={vlan-name|1..4094}]

DELETE SWITCH FILTER PORT=port ENTRY=entry-list
```

To display current static and learned switch filter entries, use the command:

```
SHOW SWITCH FILTER [PORT={port-list|ALL}]
   [DESTADDRESS=macadd] [ENTRY=entrylist] [VLAN={vlan-name|
   1..4094}]
```

For each VLAN, the destination MAC address of a frame to be forwarded is checked against the Forwarding Database. If there is no entry for the destination address and VLAN, the frame is transmitted on all ports in the VLAN that are in the 'Forwarding' or 'Disabled' states, except the port on which the frame was received. This process is referred to as *flooding*. If an entry is found in the Forwarding Database, but the entry is not marked as 'Forwarding' or the entry points to the same port the frame was received on, the frame is discarded. Otherwise, the frame is transmitted on the port specified by the entry in the Forwarding Database.

Spanning Tree Protocol

The Spanning Tree Protocol (STP) makes it possible to automatically disable redundant paths in a network to avoid loops, and enable them when a fault in the network means they are needed to keep traffic flowing. A sequence of LANs and switches may be connected together in an arbitrary physical topology resulting in more than one path between any two switches. If a loop exists, frames transmitted onto the extended LAN would circulate around the loop indefinitely, decreasing the performance of the extended LAN. On the other hand, multiple paths through the extended LAN provide the opportunity for redundancy and backup in the event of a bridge experiencing a fatal error condition. The spanning tree is created through the exchange of Bridge Protocol Data Units (BPDUs) between the bridges in the LAN when they start up, or when a change in the configuration of the network is detected.

The spanning tree algorithm ensures that the extended LAN contains no loops and that all LANs are connected by:

- Detecting the presence of loops and automatically computing a logical loop-free portion of the topology, called a *spanning tree*. The topology is dynamically pruned to a spanning tree by declaring the ports on a switch redundant, and placing the ports into a 'Blocking' state.
- Automatically recovering from a switch failure that would partition the extended LAN by reconfiguring the spanning tree to use redundant paths, if available.

The logical tree computed by the spanning tree algorithm has the following properties:

- A single switch, called the *root bridge*, forms a unique root to the tree. The root bridge is the bridge with the lowest Bridge ID. Each switch in an extended LAN is uniquely identified by its Bridge ID, which comprises the switch's root priority (a spanning tree parameter) and its MAC address.
- Each switch or LAN in the tree, except the root bridge, has a unique parent, known as the *designated bridge.* Each LAN has a single switch, called the designated bridge, that logically connects the LAN to which the switch is attached, to the next LAN closer to the root bridge.
- Each port connecting a switch to a LAN has an associated *cost*. The *root path cost* is the sum of the costs for each port between the switch and the root bridge. The designated bridge for a LAN is the switch on the LAN with the lowest root path cost, and therefore logically closer to the root bridge. If two switches on the same LAN have the same lowest root path cost, the switch with the lowest bridge ID is elected the designated bridge.

The spanning tree computation is a continuous, distributed process. The algorithm uses the following steps to establish the spanning tree:

1. A unique *root bridge* is elected by the switches in the LAN.
2. A *designated bridge* is elected for each LAN in the extended LAN by the switches in the LAN.
3. The logical spanning tree is computed and redundant paths are removed.

Once the spanning tree is established, it is maintained by:

1. Replacing a failed path with a redundant backup path, if one is available.
2. Detecting and removing loops by declaring ports redundant and removing them from the logical spanning tree.
3. Maintaining timers that control the ageing of the Forwarding Database entries.

The logical spanning tree, sometimes called the active topology, includes the root bridge and all designated bridges, i.e. all the ports that are to be used for communication within the STP. These ports are in the Forwarding state. Ports removed from the logical spanning tree are not in the Forwarding state. To implement the spanning tree algorithm, switches communicate with one another using the Spanning Tree Protocol. The primary protocol data unit (PDU) is the *Hello message* or *Configuration Bridge Protocol Data Unit* (BPDU), which includes the following information:

- The bridge ID of the root bridge.
- The distance (or cost) from this switch to the root bridge.
- The bridge ID of the designated bridge on this LAN.

Hello messages are initiated at regular intervals by the root bridge and propagate through the extended LAN.

Electing the Root Bridge and Designated Bridge

Each spanning tree has a *root bridge*, which initiates the propagation of Hello messages through the extended LAN, and sets the values of parameters that control the spanning tree computation process. The root bridge is the switch with the lowest bridge ID and is elected by the exchange of Hello packets. When a switch receives a Hello packet it compares the value of the root bridge ID in the message to the value of the root bridge ID parameter in its own spanning tree database. If the value in the message is better, the switch stores the new value in its database and sends Hello messages with the new value out on its other ports. Otherwise, the switch continues to send Hello messages with the value currently stored in its spanning tree database. By this process, all switches in the extended LAN eventually learn the bridge ID of the root bridge.

Each LAN has a single switch, called the *designated bridge*, that logically connects the LAN to the next LAN closer to the root bridge. The designated bridge for a LAN is the switch on the LAN with the lowest root path cost and bridge ID. The designated bridge is elected by the exchange of Hello messages, in the same way that the root bridge is elected. The election of a new root bridge or a switch becoming unavailable due to a fatal error condition, typically results in the election of a new designated bridge in the next few rounds of Hello messages.

Spanning Tree modes

STP can run in *standard* mode or *rapid* mode. Rapid mode allows for rapid configuration of the spanning tree. The Rapid Spanning Tree Protocol (RSTP) is specified in IEEE 802.1w.

A spanning tree running in standard mode can take up to one minute to rebuild after a topology or configuration change. The Rapid Spanning Tree algorithm provides for a more rapid recovery of connectivity following the failure of a bridge, bridge port, or a LAN. RSTP provides rapid recovery by including port roles in the computation of port states, and by allowing neighbouring bridges to explicitly acknowledge signals on a point-to-point link that indicate that a port wants to enter the forwarding mode.

In rapid mode, the rapid transition of a port to the Forwarding state is possible when the port is considered to be part of a Point-to-Point link, or when the port is considered to be an *Edge* port. An edge port is a port that attaches to a LAN that is known to have no other bridges attached.

In order to ensure that rapid transitions take place on an edge port, the port must be explicitly configured as an edge port using the SET STP PORT= {port-list | ALL} EDGEPORT=TRUE command.

Rapid Mode Spanning Tree Types

The RSTP algorithm has two types of operation: *normal* and *stp compatible.* If normal is specified as the type, then the algorithm uses rapid port role transitions and transmits and receives RST BPDUs. If STP compatible is specified, then rapid transitions are disabled and RST BPDUs are discarded. The default is normal. Setting the RSTP type to be STP compatible allows RSTP to support applications and protocols that may be sensitive to frame duplication and misordering, for example NetBeui.

Setting RSTPTYPE to NORMAL, when normal has already been set, sets all ports to the "sending RSTP" state. This is referred to in IEEE802.1w as mCheck, and is useful for restoring full rapid mode operation when one or more ports on the switch has entered the "sending STP" state. RSTP capable devices operating with RSTP set to NORMAL that receive the RST BPDUs enter the "sending RSTP" state. After the mCheck operation, if an STP BPDU is received, either as a result of a device operating in rapid mode with RSTPTYPE set to STPCOMPATIBLE, or as a result of a device operating in standard mode, the ports that received the STP BPDUs reverts to the "sending STP" state.

mCheck is most effective on switches acting as designated bridges for LANs because they regularly propagate BPDUs. Other bridges in the LAN do not transmit BPDUs as frequently.

Spanning Tree and Rapid Spanning Tree port states

If STP is running in STANDARD mode, then each port can be in one of five Spanning Tree states, and one of two switch states. If STP is running in RAPID mode, then each port can be in one of four states. The state of a switch port is taken into account by STP. To be involved in STP negotiations, STP must be enabled on the switch, the port must be enabled on the switch, and enabled for the STP it belongs to.

The Spanning Tree port states (Table 3-7 on page 3-28 and Table 3-8 on page 3-29) affect the behaviour of ports whose switch state is enabled.

Table 3-7: Spanning Tree port states.

State	Meaning
DISABLED	STP operations are disabled on the port. The port does not participate in frame relay or the operation of the Spanning Tree Algorithm and Protocol. The port can still switch if its switch state is enabled.
BLOCKING	The port does not participate in frame relay. The forwarding process discards received frames and does not submit forwarded frames for transmission. This is the "standby" mode.
LISTENING	The port is enabled for receiving frames only. The port is preparing to participate in frame relay. The forwarding process discards received frames and does not submit forwarded frames for transmission.

Table 3-7: Spanning Tree port states. (Continued)

State	Meaning
LEARNING	The port is enabled for receiving frames only, and the Learning Process can add new source address information to the Forwarding Database.
FORWARDING	The normal state for a switch port. The Forwarding Process and the Spanning Tree entity are enabled for transmit and receive operations on the port.

Table 3-8: Rapid Spanning Tree port states.

State	Meaning
DISABLED	STP operations are disabled on the port.
DISCARDING	The port does not participate in frame relay. The forwarding process discards received frames and does not submit forwarded frames for transmission.
LEARNING	The port is enabled for receiving frames only, and the Learning Process can add new source address information to the Forwarding Database. The port does not forward any frames.
FORWARDING	The normal state for a switch port. The Forwarding Process and the Spanning Tree entity are enabled for transmit and receive operations on the port.

Multiple Spanning Trees and STP interaction with VLANs

In a legacy network that has no VLANs configured, and has STP enabled, switches in the LAN run a distributed Spanning Tree Algorithm to create a single Spanning Tree.

In a network of switches with VLANs configured, all VLANs belong to a default Spanning Tree called *default*. Multiple Spanning Trees can be created with each Spanning Tree encompassing multiple VLANs. Spanning Tree Protocol entities, called STPs here, operate independently of each other; each STP has its own root bridge and active path. Once an STP is created, one or more VLANs can be assigned to it. In operation, additional STPs in the switch place no significant burden on the CPU.

If creating multiple STPs in a network, consider the following:

- A VLAN can only belong to a single STP.
- A port can belong to multiple STPs when the port is a member of more than one VLAN.

Overlapping VLANs belonging to multiple Spanning Tree instances

The AT-8600 series switches support the situation where a port is contained in more than one Spanning Tree instance when the port is a member of more than one VLAN and those VLANs belong to different STPs (See Figure 3-4 on page 3-30). You can configure up to 255 STPs.

Figure 3-4: Port membership of VLANs which belong to different spanning tree instances.

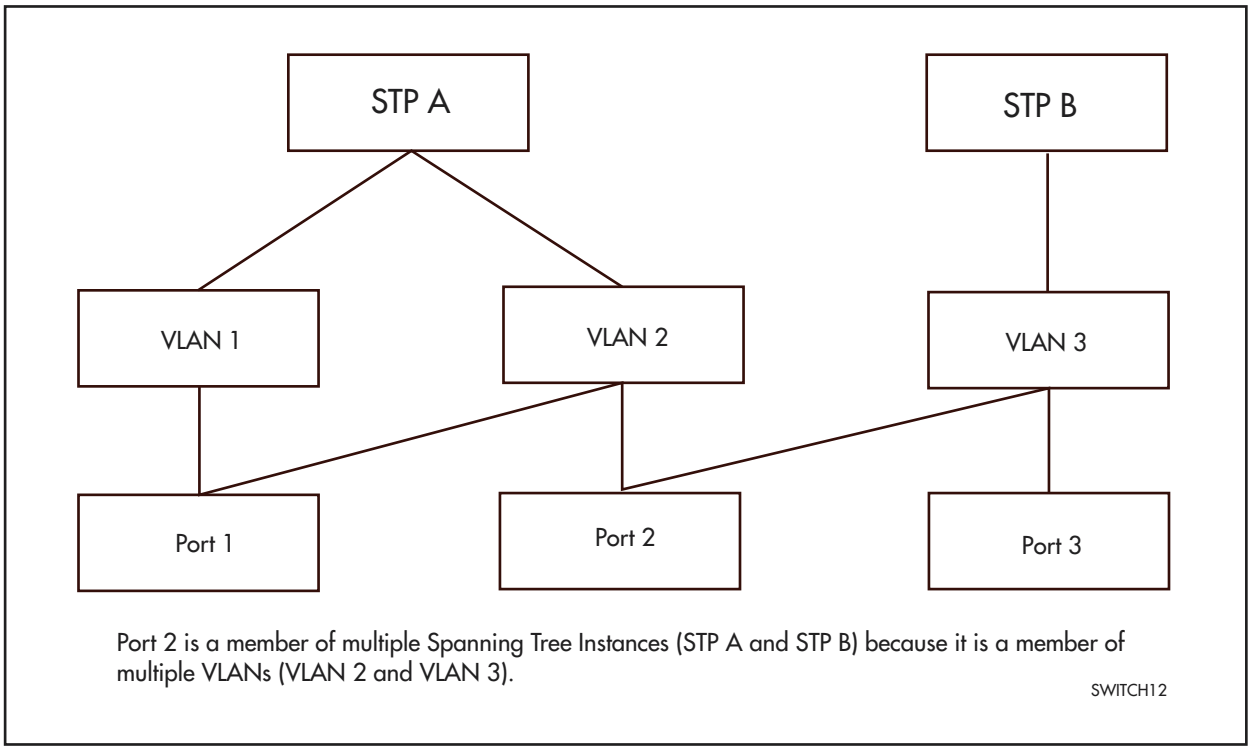

Configuring STP

By default, the switch has one *default* STP that cannot be destroyed. This default is sufficient in most situations. However, further instances of the Spanning Tree Protocol (STPs) can be created and destroyed using the commands:

```
CREATE STP=stp-name

DESTROY STP={stp-name|ALL}
```

By default, all VLANs, and therefore all ports, belong to the *default* STP. To add or delete a VLAN and all the ports belonging to it from any other STP, use the commands:

```
ADD STP=stp-name VLAN={vlan-name|2..4094}

DELETE STP=stp-name VLAN={vlan-name|2..4094|ALL}
```

The default STP is disabled by default at switch start up, and STPs created by a user are disabled by default when they are created. To enable or disable STPs, use the commands:

```
ENABLE STP={stp-name|ALL}

DISABLE STP={stp-name|ALL}
```

The Spanning Tree Protocol uses three configurable parameters for the time intervals that control the flow of STP information on which the dynamic STP topology depends: the HELLOTIME, FORWARDDELAY, and MAXAGE parameters. All switches in the same spanning tree topology must use the same values for these parameters, but can themselves be configured with different, and potentially incompatible time intervals. The parameter values actually used by each switch are those sent by the root bridge, and forwarded to all other switches by the designated bridges.

The HELLOTIME parameter, with a default of 2 seconds, determines how often the switch sends Hello messages containing spanning tree configuration information if it is the *root bridge*, or is trying to become the root bridge in the network. Setting a shorter value for HELLOTIME than the default of 2 seconds makes the network more robust; setting a longer time uses less processing overhead.

The MAXAGE parameter, with a default of 20 seconds, determines the maximum time that dynamic STP configuration information is stored in the switch, before it is considered too old, and discarded. The value can be set at approximately two seconds for every hop across the network. If this value is too small, the STP may sometimes configure unnecessarily. If it is too long, there can be delays in adapting to a change in the topology, for instance when a fault occurs.

The FORWARDDELAY parameter prevents temporary loops in the network occurring in the briefly unstable topology while a topology change is propagated through the network. When STP is running in standard mode and a port that has been in the Blocking state is to move into the Forwarding state, it must first pass through the Listening and Learning states. The FORWARDDELAY parameter determines how long the port remains in each of these intermediate states before moving on to the Forwarding state in the active topology; that is, half the time between when it is decided that the port will become part of the spanning tree and when it is allowed to forward traffic. When STP is running in rapid mode, a port only has to pass from the Discarding state through the Learning state to reach the Forwarding State. In this case, the FORWARDDELAY parameter should be at least half the time it takes for a topology change message to reach the whole network. A value that is too short risks the temporary creation of loops, which can seriously degrade switch performance. A longer value can result in delays in the network after topology changes. The default FORWARDDELAY value is 15 seconds.

The FORWARDDELAY, MAXAGE and HELLOTIME parameters should be set according to the following formulae, as specified in IEEE 802.1d:
2 x (FORWARDDELAY - 1.0 seconds) >= MAXAGE
MAXAGE >= 2 x (HELLOTIME + 1.0 seconds)

To modify the parameters controlling these time intervals, use the command:

```
SET STP={stp-name|ALL} [FORWARDDELAY=4..30] [HELLOTIME=1..10]
    [MAXAGE=6..40]
```

The value of the PRIORITY parameter sets the writable portion of the bridge ID, i.e. the first two octets of the (8-octet long) Bridge Identifier. The remaining 6 octets of the bridge ID are given by the MAC address of the switches. The Bridge Identifier parameter is used in all configuration Spanning Tree Protocol packets transmitted by the switch. The first two octets, specified by the PRIORITY parameter, determine the switch's priority for becoming the *root bridge* or a *designated bridge* in the network, with a lower number indicating a higher priority. In a fairly simple network with a small number of switches in a meshed topology, it may make little difference which switch is selected as the

root bridge, and no modifications may be needed to the default PRIORITY parameter, which has a default of 32768. In more complex networks, one or more switches are likely to be more suitable candidates for the root bridge role, for instance by virtue of being more central in the physical topology of the network. In these cases the STP PRIORITY parameters for at least one of the switches should be modified.

To change the STP priority value, use the command:

```
SET STP={stp-name|ALL} [PRIORITY=0..65535]
```

To restore STP timer and priority defaults, use the command:

```
SET STP={stp-name|ALL} DEFAULT
```

Changing the STP PRIORITY using either of the previous commands initialises the STP, so that elections for the root bridge and designated bridges begin again, without resetting STP counters. To display general information about STPs on the switch, use the command:

```
SHOW STP={stp-name|ALL}
```

Each port has a port priority, with a default of 128, used to determine which port should be the root port for the STP when two ports are connected in a loop. A lower number indicates the higher priority.

```
SET STP={stp-name|ALL} PORT={port-list|ALL}
   PORTPRIORITY=0..255
```

Each port also has a path cost, which is used when the port is the root port for the STP on the switch. The path cost is added to the root path cost field in configuration messages received on the port to determine the total cost of the path to the root bridge. The default PATHCOST values and the range of recommended PATHCOST values depend on the port speed, see Table 3-9 on page 3-32, and Table 3-10 on page 3-32. If the path cost for a port is not explicitly set, it varies as the speed of the port varies.

Table 3-9: Path cost values and port speed for STANDARD mode.

Port speed	Default PATHCOST	Recommended PATHCOST range
10Mbps	100	50-600
100Mbps	19	10-60
1Gbps	4	3-10

Table 3-10: Path cost values and port speed for RAPID mode.

Port Speed	Default PATHCOST	Recommended PATHCOST range
Less than 100 Kb/s	200,000,000	20,000,000-200,000,000
1Mbps	20,000,000	2,000,000-20,000,000
10Mbps	2,000,000	200,000-2,000,000
100 Mbps	200,000	20,000-200,000
1 Gbps	20,000	2,000-20,000
10 Gbps	2,000	200-2,000
100 Gbps	200	20-200
1Tbps	20	2-200
10 Tbps	2	2-20

Setting the path cost to a larger value on a particular port is likely to reduce the traffic over the LAN connected to it. This may be appropriate if the LAN has lower bandwidth, or if there are reasons for limiting the traffic across it. To modify the STP port path cost, for the '*Admin*' STP use the command:

```
SET STP=ADMIN PORT=ALL PATHCOST=1..200000000
```

If the path cost of a port in the '*Admin*' STP has been explicitly set to a particular value, it can be returned to its self-adjusting default path cost and priority, using the command:

```
SET STP=ADMIN PORT=4 DEFAULT
```

When an STP is enabled in a looped or meshed network, it disables and enables particular ports belonging to it dynamically, to eliminate redundant links. All ports in a VLAN belong to the same STP, and their participation in STP configuration, and hence the possibility of them being elected to the STP's active topology is enabled by default. To enable or disable particular ports, use the commands:

```
ENABLE STP={stp-name|ALL} PORT={port-list|ALL}

DISABLE STP={stp-name|ALL} PORT={port-list|ALL}
```

STP treats a trunk group configured on both ends of a link as a single path.

To display STP port information, use the command:

```
SHOW STP[={stpname|ALL}] PORT={port-list|ALL}
```

The spanning tree algorithm can be recalculated at any time, and all timers and counters be initialised, using the command:

```
RESET STP={stp-name|ALL}
```

To display STP counters, use the command:

```
SHOW STP={stp-name|ALL} COUNTER
```

Enabling one or more STP debugging modes for a period of time displays information for STP troubleshooting, see Table 3-11 on page 3-33, to the port on which the switch received the command, or to the console.

Table 3-11: STP debugging options.

Option	Debug Mode	Description
MSG	Message	Decoded display of received and transmitted STP packets
PKT	Packet	Raw ASCII display of received and transmitted STP packets
STATE	State	Port state transitions.
ALL	All	All debug options

To enable, disable or show the debug modes, use the commands:

```
ENABLE STP={stp-name|ALL} DEBUG={MSG|PKT|STATE|ALL}
    [OUTPUT=CONSOLE] [TIMEOUT={1..4,000,000,000|NONE}]

ENABLE STP DEBUG={MSG|PKT|STATE|ALL} PORT={port-list|ALL}
    [OUTPUT=CONSOLE] [TIMEOUT={1..4,000,000,000|NONE}]

DISABLE STP={stp-name|ALL} DEBUG={MSG|PKT|STATE|ALL}

DISABLE STP DEBUG={MSG|PKT|STATE|ALL} PORT={port-list|ALL}

SHOW STP DEBUG
```

STP debugging can be enabled or disabled for either a particular port(s) or a particular STP(s). Use of one of these commands overrides the other.

Set OUTPUT to CONSOLE if using this command in a script. Each of the debug modes can be enabled or disabled independently. Use the TIMEOUT parameter to prevent the switch or display from being overloaded with debugging data.

If necessary, all the STP configuration that users have created on the switch can be removed, so that all STPs except the default STP are destroyed, and all other defaults are restored, using the command:

```
PURGE STP
```

The PURGE STP command should be used with caution, and generally only before major reconfiguration of the switch, as it removes all STP configuration entered on the switch.

Hardware Packet Filters

The switch hardware can be configured to discard, forward, mirror, or change the priority of packets matching specified criteria at wirespeed. Filters can also be configured to provide a range of Quality of Service (QoS) controls, including changing the DSCP byte, and actions can be specified for packets that match the ingress and egress ports of the filter (if set), but do not match the filter's other parameters.

Two sets of commands are available, one based on the Packet Classifier (see *Chapter 6, Generic Packet Classifier*), and one based on Layer 3 filter matches and entries. These two filter types cannot be used together.

When Internet Group Management Protocol (IGMP) snooping is enabled, it uses a hardware filter, so the number of available filters is reduced. IGMP snooping is enabled by default, but can be disabled to make this filter available, using the command (see *IGMP Snooping on page 11-6 of Chapter 11, IP Multicasting*):

```
DISABLE IGMPSNOOPING
```

When IGMP snooping is disabled, multicast packets flood the VLAN.

IGMP snooping cannot be enabled unless a filter is available. To enable IGMP snooping, use the command:

```
ENABLE IGMPSNOOPING
```

Classifier-based Packet Filters

The switch hardware can be configured through entries in the Packet Classifier to copy, drop, forward, and associate QOS attributes to Layer 3 packets that match the criteria set using the classifier (see *Chapter 7, Quality of Service (QoS)* and *Chapter 6, Generic Packet Classifier*).

Every packet passing through the switch is matched against a series of classification tables by the Packet Classifier. Packets can be classified according to:

- Packet type
- Physical source/destination port
- Layer 3 protocol
- Source/destination IP address
- Destination IPX address
- Layer 4 protocol (for example: TCP/UDP/Socket number)
- Layer 4 source/destination ports
- Any 16-bit word in the first 64 bytes of a packet

See *Chapter 6, Generic Packet Classifier* for information on configuring classifiers.

Hardware-based packet filters can be configured by the user to take action upon the results of the classification tables. These actions are:

- Discard the packet
- Forward the packet
- Send the packet to the mirror port
- Forward the packet to a specified egress port, for unicast packets
- Send the packet to a Class of Service queue
- Replace the packet's 802.1p priority

The filter can also perform the following Quality of Service actions:

- Replace the packet's IP TOS value and/or the IP DSCP value.
- Direct non-unicast packets that were scheduled to be dropped or sent to the CPU to a specified port.
- Forward packets that were marked to be dropped. This option allows bandwidth limiting to be overridden for particular packets.

All actions are also available on packets that match the ingress and egress ports of the classifier (if either or both are set), but do not match the classifier's other parameters.

For more information about the circumstances when hardware filters are useful for performing QoS, see Table 7-1 on page 7-6 in *Chapter 7, Quality of Service (QoS)*.

A classifier-based packet filter comprises a single classifier entry. A number of filters can be created at one time with the same action by specifying a list of classifiers, but each classifier is contained in a single filter. The number of packet filters supported by the switch is determined by the switch model and how different each filter is.

To enable and disable classifier-based hardware filtering, use the commands:

```
ENABLE SWITCH HWFILTER

DISABLE SWITCH HWFILTER
```

This command can be useful for testing filter functionality.

When Internet Group Management Protocol (IGMP) Snooping is enabled, hardware filtering is also enabled. Hardware filtering cannot be disabled unless IGMP snooping is first disabled, using the command DISABLE IGMPSNOOPING (see IGMP Snooping on page 11-6 of Chapter 11, IP Multicasting). IGMP snooping is enabled by default.

To add hardware-based packet filters to the switch, use the command:

```
ADD SWITCH HWFILTER CLASSIFIER=classifier-list
    [ACTION={SETPRIORITY|SENDCOS|SETTOS|DENY|SENDEPORT|
    SENDMIRROR|MOVEPRIOTOTOS|MOVETOSTOPRIO|SETIPDSCP|
    SENDNONUNICASTTOPORT|NODROP|FORWARD}[,...]]
    [NEWIPDSCP=dscp-value] [NEWTOS=0..7]
    [NOMATCHACTION={SETPRIORITY|SENDCOS|SETTOS|DENY|
    SENDEPORT|SENDMIRROR|MOVEPRIOTOTOS|MOVETOSTOPRIO|
    SETIPDSCP|SENDNONUNICASTTOPORT|FORWARD}[,...]]
    [NOMATCHDSCP=dscp-value] [NOMATCHPORT=port-number]
    [NOMATCHPRIORITY=0..7] [NOMATCHTOS=0..7]
    [PORT=port-number] [PRIORITY=0..7]
```

To delete one or more hardware-based packet filters from the switch, use the command:

```
DELETE SWITCH HWFILTER CLASSIFIER=classifier-list
```

To display information about hardware-based packet filters, use the command:

```
SHOW SWITCH HWFILTER [CLASSIFIER=classifier-list]
```

Layer 3 Filter Matches

As an alternative to classifier-based filters, Layer 3 filter matches can be configured to determine which fields in each packet are matched, whether ingress or egress ports are to be matched, and the source and destination class of IP masks to apply to the packets. An entry added to a filter specifies the values to be matched for each field and the action to be taken on packets matching the filter entry. Layer 3 filter matches can perform the same actions as classifier-based hardware filters, but classifiers match a wider range of packet types.

Filters can be configured while Layer 3 filtering is disabled or enabled, but it must be enabled for any of the existing filters to take effect. To enable or disable the Layer 3 filter function, use the commands:

```
ENABLE SWITCH L3FILTER

DISABLE SWITCH L3FILTER
```

When Internet Group Management Protocol (IGMP) Snooping is enabled, Layer 3 filtering is also enabled. Layer 3 filtering cannot be disabled unless IGMP snooping is first disabled, using the command DISABLE IGMPSNOOPING (see IGMP Snooping on page 11-6 of Chapter 11, IP Multicasting). IGMP snooping is enabled by default.

To add Layer 3 filter match criteria, use the command:

```
ADD SWITCH L3FILTER MATCH={DIPADDR|IPDSCP|PROTOCOL|SIPADDR|
    TCPACK|TCPFIN|TCPDPORT|TCPSPORT|TCPSYN|TOS|TTL|UDPDPORT|
    UDPSPORT}[,...] [DCLASS={A|B|C|HOST}] [EMPORT={YES|NO|ON|
    OFF|TRUE|FALSE}] [IMPORT={YES|NO|ON|OFF|TRUE|FALSE}]
    [NOMATCHACTION={SETPRIORITY|SENDCOS|SETTOS|DENY|
    SENDEPORT|SENDMIRROR|MOVEPRIOTOTOS|MOVETOSTOPRIO|
    SETIPDSCP|SENDNONUNICASTTOPORT|FORWARD}[,...]]
    [NOMATCHDSCP=dscp-value] [NOMATCHPORT=port-number]
    [NOMATCHPRIORITY=0..7] [NOMATCHTOS=0..7] [SCLASS={A|B|C|
    HOST}] [TYPE={802|ETHII|SNAP}]
```

To display any hardware-based Layer 3 filtering match criteria configured on the switch, and their filter entries, use the command:

```
SHOW SWITCH L3FILTER[=filter-id [ENTRY=entry-id]]
```

Filter match criteria can be changed only when no filter entries belong to them. To change filter match criteria, delete any entries associated with them, then use the command:

```
SET SWITCH L3FILTER=filter-id MATCH={DIPADDR|IPDSCP|PROTOCOL|
    SIPADDR|TCPACK|TCPFIN|TCPDPORT|TCPSPORT|TCPSYN|TOS|TTL|
    UDPDPORT|UDPSPORT}[,...] [DCLASS={A|B|C|HOST}]
    [EMPORT={YES|NO|ON|OFF|TRUE|FALSE}] [IMPORT={YES|NO|ON|
    OFF|TRUE|FALSE}] [NOMATCHACTION={SETPRIORITY|SENDCOS|
    SETTOS|DENY|SENDEPORT|SENDMIRROR|MOVEPRIOTOTOS|
    MOVETOSTOPRIO|SETIPDSCP|SENDNONUNICASTTOPORT|
    FORWARD}[,...]] [NOMATCHDSCP=dscp-value]
    [NOMATCHPORT=port-number] [NOMATCHPRIORITY=0..7]
    [NOMATCHTOS=0..7] [SCLASS={A|B|C|HOST}] [TYPE={802|ETHII|
    SNAP}]
```

To delete the Layer 3 filter match criteria, first delete any entries belonging to it, then use the command:

```
DELETE SWITCH L3FILTER=filter-id
```

To configure a Layer 3 filter entry, first add the filter match criteria, then add a filter entry.

Layer 3 Filter Entries

Filter matches specify the aspect of the packet that the filter checks. Filter entries specify what that aspect must be set to in order for the traffic to be filtered by the filter. To add a Layer 3 switch filter entry to the match criteria described above, use the command:

```
ADD SWITCH L3FILTER=filter-id ENTRY [ACTION={SETPRIORITY|
    SENDCOS|SETTOS|DENY|SENDEPORT|SENDMIRROR|MOVEPRIOTOTOS|
    MOVETOSTOPRIO|SETIPDSCP|SENDNONUNICASTTOPORT|NODROP|
    FORWARD}[,...]] [DIPADDR=ipadd] [EPORT=port-number]
    [IPDSCP=number] [IPORT=port-number] [NEWIPDSCP=dscp-value]
    [NEWTOS=tos-number] [PORT=port-number] [PRIORITY=0..7]
    [PROTOCOL={TCP|UDP|ICMP|IGMP|protocol}] [SIPADDR=ipadd]
    [TCPACK={TRUE|FALSE}] [TCPDPORT=port-id] [TCPFIN={TRUE|
    FALSE}] [TCPSPORT=port-id] [TCPSYN={TRUE|FALSE}]
    [TOS=number] [TTL=number] [TYPE=protocol-type]
    [UDPSPORT=port-id] [UDPDPORT=port-id]
```

All criteria specified in the filter match should also be set in the filter entry. Criteria not in the filter match are not valid in the filter entry. The L3FILTER parameter specifies the number of the filter match to be modified. Filter match numbers are displayed in the output of the SHOW SWITCH L3FILTER command.

To change the parameters for a filter entry, use the command:

```
SET SWITCH L3FILTER=filter-id ENTRY=entry-id
    [ACTION={SETPRIORITY|SENDCOS|SETTOS|DENY|SENDEPORT|
    SENDMIRROR|MOVEPTIOTOTOS|MOVETOSTOPRIO|SETIPDSCP|
    SENDNONUNICASTTOPORT|FORWARD}[,...]] [DIPADDR=ipadd]
    [EPORT=port-number] [IPORT=port-number] [NEWIPDSCP=dscp-
    value] [NEWTOS=tos-number] [PORT=port-number]
    [PRIORITY=0..7] [PROTOCOL={TCP|UDP|ICMP|IGMP|protocol}]
    [SIPADDR=ipadd] [TCPACK={TRUE|FALSE}] [TCPDPORT=port-id]
    [TCPFIN={TRUE|FALSE}] [TCPSPORT=port-id] [TCPSYN={TRUE|
    FALSE}] [TOS=number] [TTL=number] [TYPE=protocol-type]
    [UDPSPORT=port-id] [UDPDPORT=port-id]
```

To delete a Layer 3 filter entry, use the command:

```
DELETE SWITCH L3FILTER=filter-id ENTRY=entry-id
```

Access Control Lists (ACLs)

Classifiers and hardware packet filters can be configured to provide Access Control List functionality.

For example, to allow WWW servers in the 192.168.10.0 subnet to be accessed only from the 192.168.20.0 subnet:

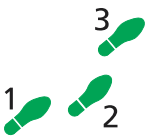

1. **Create a classifier to match all WWW traffic to the subnet**

   Create a classifier to match all WWW traffic to the 192.168.10.0 subnet.

   ```
   CREATE CLASSIFIER=1 IPDADDR=192.168.10.0/24 TCPDPORT=80
   ```

2. **Create a hardware packet filter to deny this traffic**

   ```
   ADD SWITCH HWFILTER CLASSIFIER=1 ACTION=DENY
   ```

3. **Create a classifier to match the subset of this traffic that is to be allowed**

   Create a classifier to match WWW traffic from the 192.168.20.0 subnet to the 192.168.10.0 subnet.

   ```
   CREATE CLASSIFIER=2 IPDADDR=192.168.10.0/24
      IPSADDR=192.168.20.0/24 TCPDPORT=80
   ```

4. **Create a hardware packet filter to allow this traffic**

   This filter must be created last so that it is the first filter that the switch processes.

   ```
   ADD SWITCH HWFILTER CLASSIFIER=2 ACTION=NODROP
   ```

The parameter NOMATCHACTION can be used to create a hardware filter that acts upon traffic that does not match the classifier or any other hardware filters. For example, to allow traffic destined for TCP ports 25 and 80 and UDP port 5151, and block all other traffic, create the following set of classifiers and filters:

```
CREATE CLASSIFIER=1 TCPDPORT=80

ADD SWITCH HWFILTER CLASSIFIER=1 ACTION=FORWARD
   NOMATCHACTION=DENY

CREATE CLASSIFIER=2 TCPDPORT=25

ADD SWITCH HWFILTER CLASSIFIER=2 ACTION=FORWARD
   NOMATCHACTION=DENY

CREATE CLASSIFIER=3 UDPDPORT=5151

ADD SWITCH HWFILTER CLASSIFIER=3 ACTION=FORWARD
   NOMATCHACTION=DENY
```

If the NOMATCHACTION was not specified in these filters, all traffic would be forwarded, including traffic that matched the classifiers.

Triggers

The Trigger Facility can be used to automatically run specified command scripts when particular triggers are activated. When a trigger is activated by an event, global parameters and parameters specific to the event are passed to the script that runs. For a full description of the Trigger Facility, see *Chapter 17, Trigger Facility*.

The switch can generate triggers to activate scripts when a switch port goes up or down.

The following section lists the events that may be specified for the Switching module for the EVENT parameter, the parameters that may be specified as *module-specific-parameters* for the Switching module, and the arguments passed to the script activated by the trigger.

Module Layer 3 Switching module: MODULE=SWI

Event LINKDOWN

Description The port link specified by the PORT parameter has just gone down.

Parameters The following command parameter(s) must be specified in the CREATE/SET TRIGGER commands:

Parameter	Description
PORT=*port*	The port where the event activates the trigger.

Script Parameters The trigger passes the following parameter(s) to the script:

Argument	Description
%1	The port number of the port that has just gone down.

Event LINKUP

Description The port link specified by the PORT parameter has just come up.

Parameters The following command parameter(s) must be specified in the CREATE/SET TRIGGER commands:

Parameter	Description
PORT=*port*	The port where the event activates the trigger.

Script Parameters The trigger passes the following parameter to the script:

Argument	Description
%1	The port number of the port that has just come up.

To create or modify a switch trigger, use the commands:

```
CREATE TRIGGER=trigger-id MODULE=SWITCH EVENT={LINKDOWN|
    LINKUP} PORT=port [AFTER=hh:mm] [BEFORE=hh:mm] [DATE=date|
    DAYS=day-list] [NAME=name] [REPEAT={YES|NO|ONCE|FOREVER|
    count}] [SCRIPT=filename...] [STATE={ENABLED|DISABLED}]
    [TEST={YES|NO|ON|OFF|TRUE|FALSE}]

SET TRIGGER=trigger-id [PORT=port] [AFTER=hh:mm]
    [BEFORE=hh:mm] [DATE=date|DAYS=day-list] [NAME=name]
    [REPEAT={YES|NO|ONCE|FOREVER|count}] [TEST={YES|NO|ON|
    OFF|TRUE|FALSE}]
```

Configuration Examples

This section shows examples of configuring the Layer two switch functions on the switch. All examples assume that the switch configuration begins from factory default settings.

Note that routing, required for communication between the VLANs, is not shown in these examples.

Example using one switch to extend a local LAN

The example in Figure 3-5 on page 3-41 uses a single switch to connect two (or more) physical LANs and a server. All the devices connected belong to the same broadcast domain, and separate collision domains. The Learning and Forwarding Processes in the switch give this topology better performance than a single LAN would give, and allow more devices to be attached than would a single physical LAN.

Figure 3-5: Example of switch with default configuration

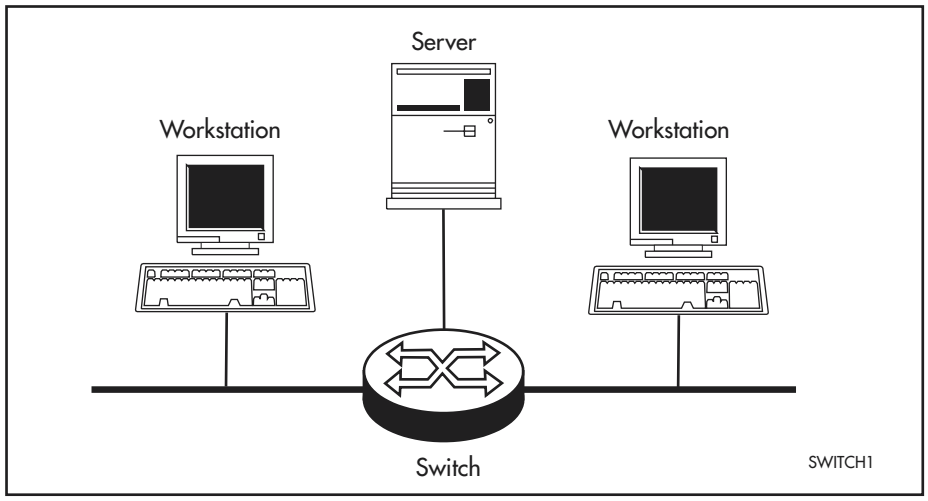

No software configuration is required. The default switch settings lets the switch learn source addresses and forward frames to correct ports as soon as it is physically connected and powered up.

Example of a meshed network without VLANs

The example in Figure 3-6 on page 3-41 has redundant links between the switches, and all ports belong only to the default VLAN. STP is needed because of the loop in the physical topology.

Figure 3-6: Example of switch with default configuration

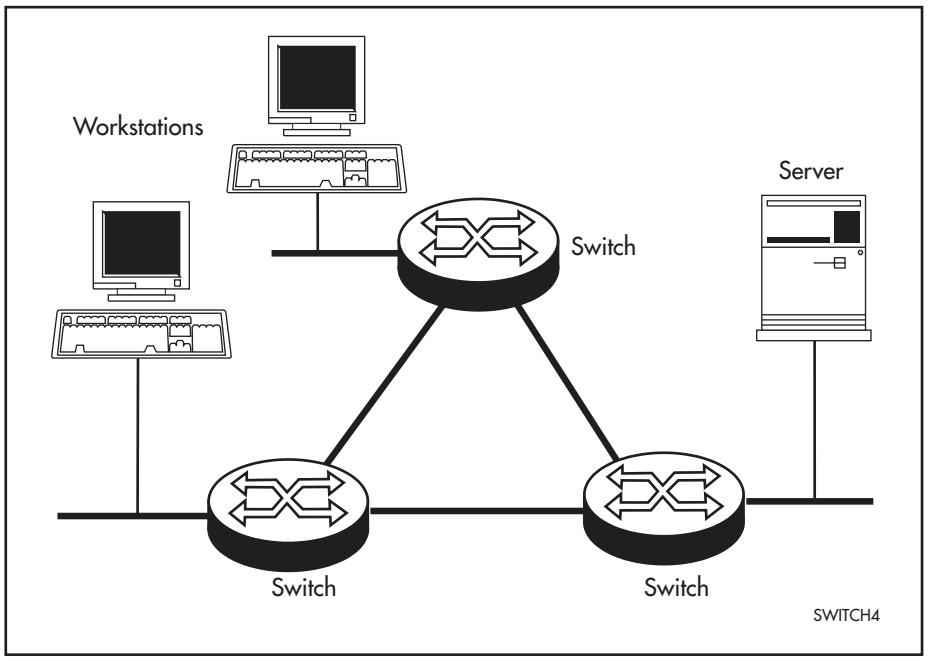

The only software configuration required is to enable the default STP on each of the switches, to eliminate loops in the network. The switches begin switching as soon as they are physically connected and powered up.

Table 3-12: Parameters for meshed network without VLANs.

All switches		
STP	default STP	Enabled

Configure all switches

1. Enable STP

The default VLAN to which all ports belong by default, is a member of the default STP. Enable the default STP on each switch using the command:

```
ENABLE STP=default
```

VLAN example using untagged ports

The example in Figure 3-7 on page 3-42 has two VLANs using untagged ports. Ports 1-3 belong to one broadcast domain, the *marketing* VLAN, and ports 14-16 belong to another broadcast domain, the *training* VLAN. The switch acts as two separate bridges: one that forwards between the ports belonging to the *marketing* VLAN, and a second one that forwards between the ports belonging to the *training* VLAN. Devices on ports 2 and 14 can only communicate with each other by using the switch's IP routing functions.

Figure 3-7: VLANS with untagged ports

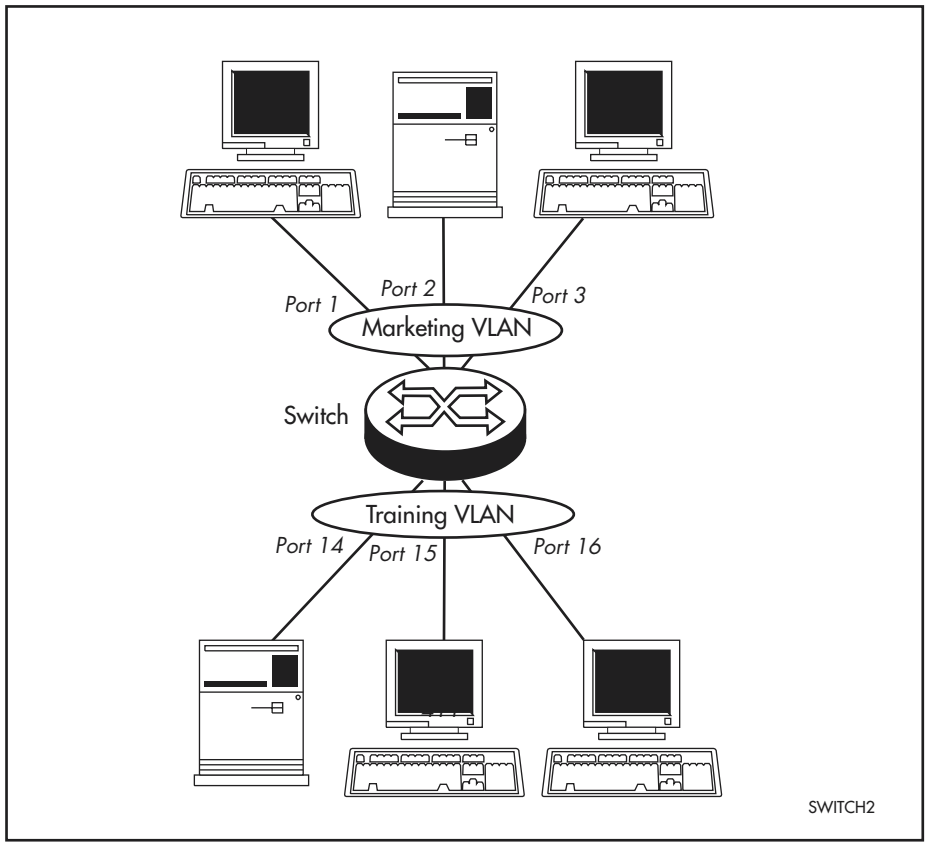

Table 3-13 on page 3-43 shows the parameters used to configure this example. Since there is only one switch and no loops in this topology, the Spanning Tree Protocol (STP) is not needed. This example assumes that the switch has factory default settings.

Table 3-13: Parameters for port-based VLAN example.

VLAN name	VLAN ID	Ports
Marketing	VID=2	PORT 1-3
Training	VID=3	PORT 14-16

Configure the switch

1. Create VLANs

Create the two VLANs using the following commands on the switch:

```
CREATE VLAN=Marketing VID=2
CREATE VLAN=Training VID=3
```

2. Add ports to VLANs

Add the ports to these VLANs on the switch by using the following commands:

```
ADD VLAN=Marketing PORT=1-3
ADD VLAN=Training PORT=14-16
```

Check the VLAN configuration by using the command:

```
SHOW VLAN
```

Check

Check that the switch is switching across the ports. Traffic on the switch can be monitored using the command:

```
SHOW SWITCH PORT=1-3,14-16 COUNTER
```

VLAN example using tagged ports

Figure 3-8 on page 3-44 shows a network that must be configured with VLAN tagging, since the VLAN aware server on port 2 on Switch A belongs to both the *admin* VLAN and the *marketing* VLAN. Using VLAN tags, port 26 on Switch A and port 25 on Switch B belong to both the *marketing* VLAN and the *training* VLAN, so that devices on both VLANs can use this uplink to communicate with other devices in the same VLAN on the other switch. There are no loops in this topology, so STP is not needed.

Figure 3-8: VLANs with tagged ports

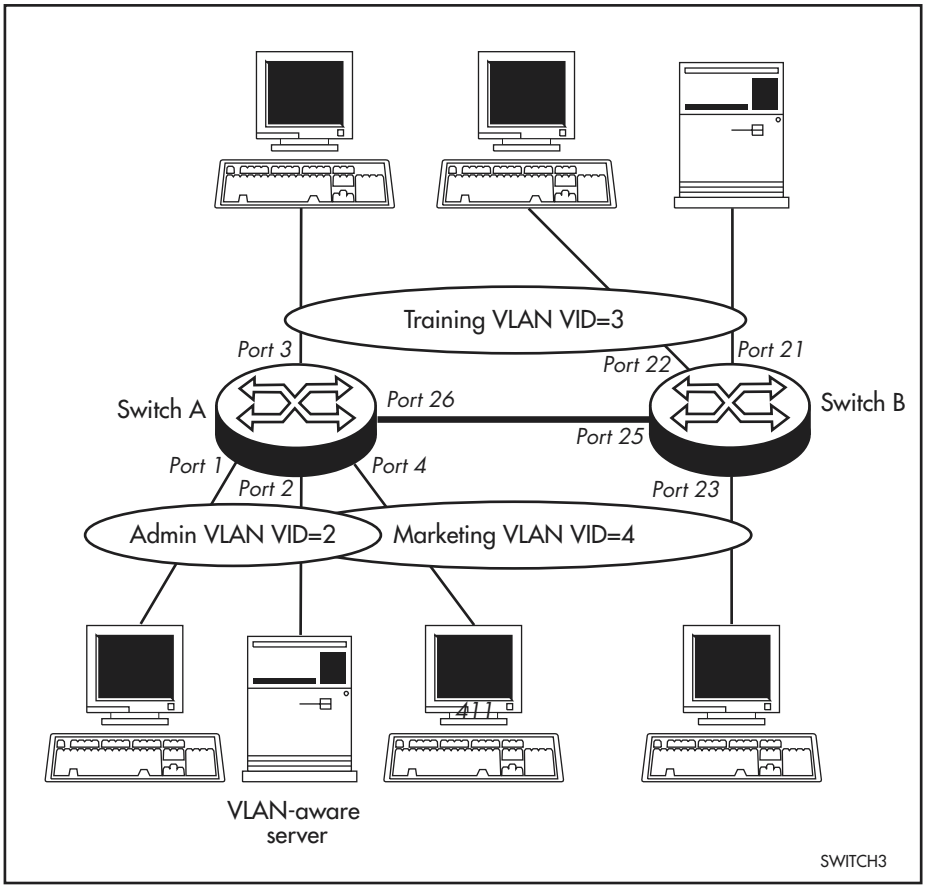

Table 3-14: Configuration example parameters for VLANs with tagged ports.

		Switch A		Switch B	
VLAN name	**VID**	**Tagged ports**	**Untagged ports**	**Tagged ports**	**Untagged ports**
Admin	VID=2	PORT 2	PORT 1		
Training	VID=3	PORT 26	PORT 3	PORT 25	PORT 21,22
Marketing	VID=4	PORT 2,26	PORT 4	PORT 25	PORT 23

Configure Switch A

1. Create VLANs

Create the three VLANs using the following commands on the switch:

```
CREATE VLAN=Admin VID=2
CREATE VLAN=Training VID=3
CREATE VLAN=Marketing VID=4
```

2. Add ports to VLANs

Add the ports to these VLANs on the switch by using the following commands:

```
ADD VLAN=Admin PORT=2 FRAME=TAGGED
ADD VLAN=Admin PORT=1
ADD VLAN=Training PORT=26 FRAME=TAGGED
ADD VLAN=Training PORT=3
ADD VLAN=Marketing PORT=2,26 FRAME=TAGGED
ADD VLAN=Marketing PORT=4
```

Check the VLAN configuration by using the command:

```
SHOW VLAN
```

Configure Switch B

1. Create VLANs

Create the two VLANs using the following commands on the switch:

```
CREATE VLAN=Training VID=3
CREATE VLAN=Marketing VID=4
```

2. Add ports to VLANs

Add the ports to these VLANs on the switch by using the following commands:

```
ADD VLAN=Training PORT=25 FRAME=TAGGED
ADD VLAN=Training PORT=21,22
ADD VLAN=Marketing PORT=25 FRAME=TAGGED
ADD VLAN=Marketing PORT=23
```

Check the VLAN configuration by using the command:

```
SHOW VLAN
```

Check

Check that the switch is switching across the ports. Traffic on Switch A can be monitored using the command:

```
SHOW SWITCH PORT=1-4,26 COUNTER
```

Traffic on Switch B can be monitored using the command:

```
SHOW SWITCH PORT=21-23,25 COUNTER
```

Example of meshed network with VLAN tagged ports

In this example, the uplink ports on all three switches connect the VLANs. Server S on Switch B is VLAN aware, and is shared between all three VLANs. The other devices shown are VLAN-unaware end stations, connected to untagged ports. Because both uplink ports on all three switches belong to the *marketing* VLAN, the Spanning Tree Protocol eliminates the loop in this VLAN, and provides redundancy in case links fail. Because the VLAN-aware shared server on Switch B, and the uplink ports belong to all three VLANs, these VLANs must all belong to the same STP.

Figure 3-9: Example of meshed network with VLAN tagged ports

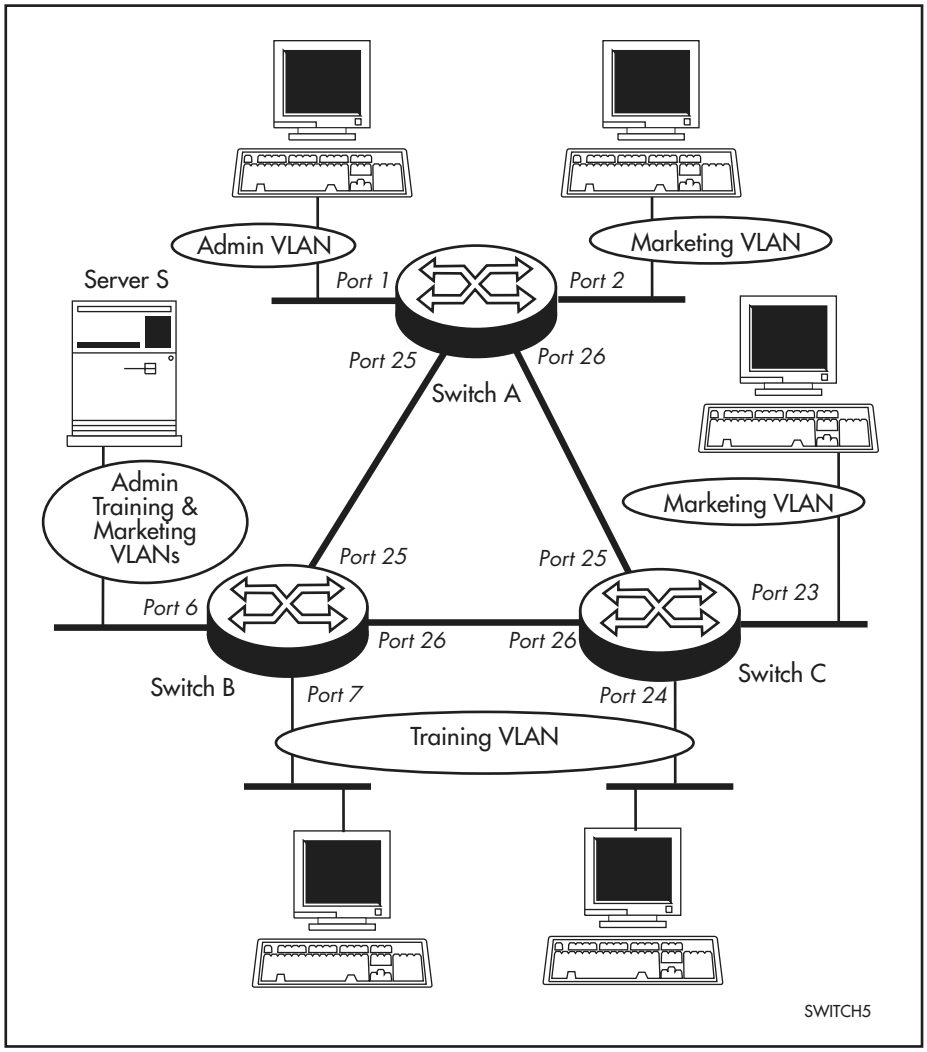

Table 3-15 on page 3-47 shows the parameters for creating the VLANs on the switches and adding ports to the VLANs. Note that by default all VLANs belong to the default STP, which is disabled at switch start-up.

Note that all three VLANs are created on all three switches, and all uplink ports belong to all three VLANs. This should be done even though the training VLAN has no devices on Switch A that need to communicate with Switch B or C, and Switch C has no devices belonging to the admin VLAN requiring links to Switch A or B. This is because STP is enabled, and inevitably blocks ports on one of the three links to prevent a loop in the marketing VLAN. This also blocks traffic over these ports for the other VLANs. Therefore the training and admin VLANs must be able to communicate over either of the links on each switch to ensure full VLAN operation. Failing to include the switches and uplink ports in the VLANs for which they have no devices attached is likely to block either the admin or training VLANs access to some of their members.

Table 3-15: Parameters for meshed VLAN network with tagged ports.

		Switch A		Switch B		Switch C	
VLAN name	**VID**	**Tagged ports**	**Untagged ports**	**Tagged ports**	**Tagged ports**	**Tagged ports**	**Tagged ports**
Admin	VID=2	25,26	1	6,25,26	-	25,26	-
Training	VID=3	25,26	-	6,26,25	7	26,25	24
Marketing	VID=4	25,26	2	6,25,26	-	25,26	23
STP		**Default STP**		**Default STP**		**Default STP**	
		Enabled		Enabled		Enabled	

To configure the uplink ports in the above example, use the following commands:

Configure Switch A

1. Create VLANs

Create the three VLANs using the following commands on the switch:

```
CREATE VLAN=Admin VID=2
CREATE VLAN=Training VID=3
CREATE VLAN=Marketing VID=4
```

2. Add ports to VLANs

Add the ports to these VLANs on the switch by using the following commands:

```
ADD VLAN=Admin PORT=25-26 FRAME=TAGGED
ADD VLAN=Admin PORT=1
ADD VLAN=Training PORT=25-26 FRAME=TAGGED
ADD VLAN=Marketing PORT=25-26 FRAME=TAGGED
ADD VLAN=Marketing PORT=2
```

Check the VLAN configuration by using the command:

```
SHOW VLAN
```

3. Enable STP

All VLANs belong to the default STP, which must be enabled to eliminate loops in the network. Use the command:

```
ENABLE STP=default
```

Configure Switch B

1. Create VLANs

Create the three VLANs using the following commands on the switch:

```
CREATE VLAN=Admin VID=2
CREATE VLAN=Training VID=3
CREATE VLAN=Marketing VID=4
```

2. Add ports to VLANs

Add the ports to these VLANs on the switch by using the following commands:

```
ADD VLAN=Admin PORT=6,25-26 FRAME=TAGGED
ADD VLAN=Training PORT=6,25-26 FRAME=TAGGED
ADD VLAN=Training PORT=7
ADD VLAN=Marketing PORT=6,25-26 FRAME=TAGGED
```

Check the VLAN configuration by using the command:

```
SHOW VLAN
```

3. Enable STP

All VLANs belong to the default STP, which must be enabled to eliminate loops in the network. Use the command:

```
ENABLE STP=default
```

Configure Switch C

1. Create VLANs

Create the three VLANs using the following commands on the switch:

```
CREATE VLAN=Admin VID=2
CREATE VLAN=Training VID=3
CREATE VLAN=Marketing VID=4
```

2. Add ports to VLANs

Add the ports to these VLANs on the switch by using the following commands:

```
ADD VLAN=Admin PORT=25-26 FRAME=TAGGED
ADD VLAN=Training PORT=25-26 FRAME=TAGGED
ADD VLAN=Training PORT=24
ADD VLAN=Marketing PORT=25-26 FRAME=TAGGED
ADD VLAN=Marketing PORT=23
```

Check the VLAN configuration by using the command:

```
SHOW VLAN
```

3. Enable STP

All VLANs belong to the default STP, which must be enabled to eliminate loops in the network. Use the command:

```
ENABLE STP=default
```

Check

Check that the switch is switching across the ports.

1. Check the traffic on Switch A.

```
SHOW SWITCH PORT=1,2,25,26 COUNTER
```

2. Check the traffic on Switch B.

```
SHOW SWITCH PORT=6,7,25,26 COUNTER
```

3. Check the traffic on Switch C.

```
SHOW SWITCH PORT=23-26 COUNTER
```

Command Reference

This section describes the commands available to configure and manage the switching functions on the switch.

See *Conventions on page xlvii of Preface* in the front of this manual for details of the conventions used to describe command syntax. See *Appendix A, Messages* for a complete list of messages and their meanings.

ACTIVATE SWITCH PORT

Syntax `ACTIVATE SWITCH PORT={`*`port-list`*`|ALL} {AUTONEGOTIATE} {LOCK}`

where:

- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.

Description This command activates autonegotiation of port speed and duplex mode for a port or a group of ports.

The PORT parameter specifies the port or ports for which autonegotiation is to be activated. Only ports in the list that are set to autonegotiate are actually affected by this command. Ports with a fixed speed setting or that belong to a trunk group are not modified.

The AUTONEGOTIATE parameter specifies that the port is to activate the autonegotiation process. The port begins to autonegotiate link speed and duplex mode.

The LOCK parameter manually locks the switch port before it reaches its learning limit so that no new addresses are automatically learned. The LEARN parameter for the port is set to the current number of learned MAC addresses.

Examples To activate autonegotiation on ports 1-8 and port 10, use the command:

```
ACTIVATE SWITCH PORT=1-8,10 AUTONEGOTIATE
```

Related Commands SET SWITCH PORT
SHOW SWITCH PORT

ADD STP VLAN

Syntax `ADD STP=`*`stp-name`* `VLAN={`*`vlan-name`*`|2..4094}`

where:

- *stp-name* is a character string 1 to 15 characters long.Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character (“_”), and the hyphen (-). The *stp-name* cannot be ALL.
- *vlan-name* is a unique name for the VLAN 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character (“_”), and the hyphen (-). The *vlan-name* cannot be a number or ALL.

Description This command adds a VLAN to the specified STP. If as a result of the VLAN addition, ports are moved from one STP to another STP, the two affected STPs are initialised if they are currently enabled. Any previously disabled ports in the STPs are enabled.

The default VLAN cannot be added to an STP. The default VLAN always belongs to the default STP. A VLAN cannot be explicitly added to the default STP. A VLAN is implicitly added to the default STP when it is deleted from any other STP. Only a VLAN belonging to the default STP can be added to another STP. If the VLAN already belongs to another STP, it must first be deleted from its current STP (and so be returned to the default STP), and then added to the new STP.

Within any given STP, all VLANs belonging to it use the same Spanning Tree.

If a port is a member of multiple VLANs, then all these VLANs must belong to the same STP.

A port can belong to more than one STP if the port is a member of two or more VLANs that belong to different STPs.

The VLAN parameter specifies the name or the numerical VLAN Identifier of the VLAN to be added to the STP. The name is not case sensitive, although the case is preserved for display purposes. The VLAN specified must exist.

When a VLAN is added to an STP, the ports in the VLAN will have default STP parameter values. The ports do not retain non-default STP configurations made when the VLAN was associated with any other STP.

Examples To add the *research* VLAN to the *company* STP, use the command:

```
ADD STP=company VLAN=research
```

Related Commands DELETE STP VLAN
SHOW STP

ADD SWITCH FILTER

Syntax

```
ADD SWITCH FILTER ACTION={FORWARD|DISCARD}
   DESTADDRESS=macadd PORT=port [ENTRY=entry] [LEARN]
   [VLAN={vlan-name|1..4094}]
```

where:

- *entry* is a filter entry number, in the range 0 to n+1, where n is the highest filter entry currently defined in the Permanent Forwarding Database. The Permanent Forwarding Database has a maximum of 320 entries, ranging from 0 to 319. Each port has its own Permanent Forwarding Database.
- *vlan-name* is a unique name for the VLAN 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *vlan-name* cannot be a number or ALL.
- *port* is the number of the switch port or uplink port to which this filter applies.
- *macadd* is an Ethernet six-octet MAC address, expressed as six pairs of hexadecimal digits delimited by hyphens.

Description This command adds a single static filter entry to the Permanent Forwarding Database for a specified port. If the static entry matches an existing dynamic entry that was learnt by the switch (a match means that the DESTADDRESS and VLAN parameters are the same for both entries), the static filter overwrites the existing dynamic learnt entry. All the received frames that match the static filter entry are forwarded to the specified port with an action of FORWARD or DISCARD.

The ACTION parameter specifies the outcome of the Forwarding Process for the frame. When FORWARD is specified, the frame is transmitted on the given port or ports. When DISCARD is specified, the frame is discarded.

The DESTADDRESS parameter specifies the value to be matched against the destination MAC address from frames being filtered. The destination MAC address must be an individual MAC address.

The PORT parameter specifies the outbound port over which a frame matching this filter entry is discarded or forwarded. Whether the ports are tagged ports or untagged ports is determined by the VLAN parameter. When the PORT parameter specifies tagged ports, then the VLAN parameter is required.

The ENTRY parameter specifies where in the Permanent Forwarding Database the new entry is added for the specified port. ENTRY cannot be set greater than n+1 where n is the highest filter entry currently defined. When ENTRY is not specified, the new entry is appended to the bottom of the Permanent Forwarding Database: the default is n+1 where n is the highest filter entry

currently defined. Static and dynamic entries in the Forwarding Database are kept in sorted order determined by their VLAN Identifier and MAC address. Therefore the ENTRY parameter does not affect the order of the filters in the Forwarding Database. The order in which filter entries are displayed by the SHOW SWITCH FILTER command is dependent upon the ENTRY parameter.

The LEARN parameter specifies if the filter being added should be counted and used as a learned MAC address for intrusion detection. Learned filters are not totally static, and can be lost if the learning process is stopped by setting the LEARN parameter to zero (see the SET SWITCH PORT command).

The VLAN parameter specifies the VLAN Identifier to which the filter entry is associated. The VLAN parameter is required when the PORT parameter specifies tagged ports. When the PORT parameter specifies untagged ports, the VLAN parameter is not required, and defaults to the VLAN Identifier of the VLAN for which the ports are untagged. Therefore, when the VLAN parameter is not specified, the ports are treated as untagged ports.

The switch automatically deletes static filter entries for a port if the port is deleted from the specified VLAN.

Examples To forward all frames destined for MAC address 00-00-cd-12-34-56 on the VLAN to which port 3 is an untagged port, use the command:

```
ADD SWITCH FILTER DESTADDRESS=00-00-cd-12-34-56
   ACTION=FORWARD PORT=3
```

To discard all frames destined for MAC address 00-00-cd-12-34-56 on port 4 in VLAN 4, use the command:

```
ADD SWITCH FILTER DESTADDRESS=00-00-cd-12-34-56 PORT=4
   ACTION=DISCARD VLAN=4
```

Related Commands DELETE SWITCH FILTER
SHOW SWITCH FILTER

ADD SWITCH HWFILTER CLASSIFIER

Syntax

```
ADD SWITCH HWFILTER CLASSIFIER=classifier-list
    [ACTION={SETPRIORITY|SENDCOS|SETTOS|DENY|SENDEPORT|
    SENDMIRRORMOVEPRIOTOTOS|MOVETOSTOPRIO|SETIPDSCP|
    SENDNONUNICASTTOPORT|NODROP|FORWARD}[,...]]
    [NEWIPDSCP=0..63] [NEWTOS=0..7]
    [NOMATCHACTION={SETPRIORITY|SENDCOS|SETTOS|DENY|
    SENDEPORT|SENDMIRROR|MOVEPRIOTOTOS|MOVETOSTOPRIO|
    SETIPDSCP|SENDNONUNICASTTOPORT|FORWARD}[,...]]
    [NOMATCHDSCP=0..63] [NOMATCHPORT=port-number]
    [NOMATCHPRIORITY=0..7] [NOMATCHTOS=0..7]
    [PORT=port-number] [PRIORITY=0..7]
```

where:

- *classifier-list* is either an integer in the range 1 to 9999; a range of integers (specified as 0-4) or a comma separated list of classifier numbers and/or ranges (0, 3, 4-9).
- *port-number* is the switch port number, in the range 1 to m, where m is the highest numbered Ethernet switch port, including uplink ports.

Description

This command adds hardware based filters based on the specified classifier(s). The classifiers in the list must exist, and they must not already be specified as part of an existing filter entry, neither may they be a duplicate of another classifier that is already used by a filter entry. The SWITCH HWFILTER CLASSIFIER commands may not be used with the SWITCH L3FILTER commands.

The ACTION parameter specifies a comma-separated list of actions to take when a packet matches the filter criteria specified in this entry. If SETPRIORITY is specified, the packet's 802.1p priority is set to the value specified by the PRIORITY parameter. If SENDCOS is specified, the packet is sent to the priority queue specified by the PRIORITY parameter. If SETTOS is specified, the packet's TOS (Type of Service) field is set to the value specified by the NEWTOS parameter. When DENY is specified, the packet is discarded. If SENDEPORT is specified and the frame is not a broadcast, multicast, DLF frame or marked for dropping or to be sent to the CPU, the new output port is set to the value of the PORT parameter. If SENDMIRROR is specified, the packet is sent to the mirror port. If FORWARD is specified, the packet is forwarded using the default Class of Service (priority). The default is FORWARD. If MOVEPRIOTOTOS is specified, the IP TOS field in the frame is replaced with the 802.1 priority value. If MOVETOSTOPRIO is specified, the 802.1 priority field in the frame is replaced with the IP TOS value, this also determines the egress priority queue. If SETIPDSCP is specified and the frame is an IPv4 frame, the Diffserv Codepoint field in the frame is set to the value specified by the NEWIPDSCP parameter. Actions that modify both the IP TOS and IP DSCP values in the frame are mutually exclusive. If SENDNONUNICASTTOPORT is specified, matching frames that are broadcast or DLF, multicast, marked for dropping or to be sent to the CPU are sent to the egress port specified by the PORT parameter. If NODROP is specified, matching frames previously marked for dropping are not dropped.

If the SENDEPORT action directs packets to a particular egress port, then the packet is transmitted from the mirror port with a VLAN tag.

The NEWIPDSCP parameter indicates the value to set in an IPv4 packet Diffserv CodePoint field when the ACTION parameter is set to SETIPDSCP. The range of values for this parameter is from 0 to 63.

The NEWTOS parameter specifies the new type of service value, assigning a new value to the TOS precedence field in the IP Header. When this parameter is used, only when the ACTION parameter is set to SETTOS.

The NOMATCHACTION parameter specifies a comma-separated list of actions to take when a frame matches both the IPORT and EPORT values (if they are specified in the match) on an associated entry but there is no match for the frame contents. When SETPRIORITY is specified, the packet's 802.1p priority is set to the value specified by the PRIORITY parameter. When SENDCOS is specified, the packet is sent to the priority queue specified by the PRIORITY parameter. When SETTOS is specified, the packet's TOS (Type of Service) field is set to the value specified by the NEWTOS parameter. If DENY is specified, the packet is discarded. When SENDEPORT is specified and the frame is not a broadcast, multicast, DLF frame or marked for dropping or to be sent to the CPU, the new output port is set to the value of the PORT parameter. When SENDMIRROR is specified, the packet is sent to the mirror port. When FORWARD is specified, the packet is forwarded using the default Class of Service (priority). When MOVEPRIOTOTOS is specified, the IP TOS field in the frame is replaced with the 802.1 priority value. When MOVETOSTOPRIO is specified, the 802.1 priority field in the frame is replaced with the IP TOS value, this also determines the egress priority queue. When SETIPDSCP is specified and the frame is an IPv4 frame, the Diffserv Codepoint field in the frame is set to the value specified by the NEWIPDSCP parameter. Actions that modify both the IP TOS and IP DSCP values in the frame are mutually exclusive. When SENDNONUNICASTTOPORT is specified, matching frames that are broadcast or DLF, multicast, marked for dropping or to be sent to the CPU are sent to the egress port specified by the PORT parameter. The default is FORWARD.

The NOMATCHDSCP parameter indicates the value to set in an IPv4 packet Diffserv CodePoint field when the NOMATCHACTION parameter is set to SETIPDSCP. The range of values for this parameter is from 0 to 63.

The NOMATCHPORT parameter specifies the new output port number. This port overrides the egress port selected by the Forwarding Database.

The NOMATCHPRIORITY parameter specifies the packet priority. There are eight levels of priority, from 0 to 7. This parameter is used when the NOMATCHACTION parameter is set to SETPRIORITY or SENDCOS.

The NOMATCHTOS parameter specifies the new type of service value, assigning a new value to the TOS precedence field in the IP Header. This parameter is used when the NOMATCHACTION parameter is set to SETTOS.

The PORT parameter specifies the new output port number. This port overrides the egress port selected by the Forwarding Database.

The PRIORITY parameter specifies the packet priority. There are eight levels of priority, from 0 to 7. This parameter is used when the ACTION parameter is set to SETPRIORITY or SENDCOS.

Examples To add hardware filtering entries to the switch based on classifier entries 1 to 5 that drop all matching packets, use the command:

```
ADD SWITCH HWFILTER CLASSIFIER=1-5 ACTION=DENY
```

Related Commands DELETE SWITCH HWFILTER CLASSIFIER
SET SWITCH HWFILTER CLASSIFIER
SHOW SWITCH HWFILTER

ADD SWITCH L3FILTER ENTRY

Syntax

```
ADD SWITCH L3FILTER=filter-id ENTRY [ACTION={DENY|FORWARD|
   SENDCOS|SENDEPORT|SENDMIRROR|SETPRIORITY|SETTOS|
   MOVEPRIOTOTOS|MOVETOSTOPRIO|NODROP|
   SENDNONUNICASTTOPORT|SETIPDSCP}[,...]] [DIPADDR=ipadd]
   [EPORT=port-number] [IPDSCP=number] [IPORT=port-number]
   [NEWIPDSCP=0..63] [NEWTOS=0..7] [PORT=port-number]
   [PRIORITY=0...7] [PROTOCOL={TCP|UDP|ICMP|IGMP|protocol}]
   [SIPADDR=ipadd] [TCPACK={TRUE|FALSE}]
   [TCPDPORT=port-id] [TCPFIN={TRUE|FALSE}]
   [TCPSPORT=port-id] [TCPSYN={TRUE|FALSE}] [TOS=0..7]
   [TTL=0..255] [TYPE=protocol-type] [UDPSPORT=port-id]
   [UDPDPORT=port-id]
```

where:

- *filter-id* is a decimal number in the range 1 to the number of filters defined.
- *ipadd* is an IP address in dotted decimal notation.
- *port-number* is the switch port number, in the range 1 to m, where m is the highest numbered Ethernet switch port, including uplink ports.
- *protocol* is an IP protocol number in the range 1 to 255.
- *port-id* is a TCP/UDP port number with a maximum value less than 65535.
- *protocol-type* is a valid protocol-type number. A protocol type number is 2 bytes for Ethernet type II and 802.3 (DSAP/SSAP) encapsulation, or 5 bytes for SNAP encapsulation, and is specified in hexadecimal.

Description This command adds a filter entry to an existing filter match criteria. All criteria specified in the filter match should also be set in the filter entry, and criteria not specified in the filter match are not valid in the filter entry. Up to 126 filter entries may be created.

The SWITCH HWFILTER CLASSIFIER commands may not be used with the SWITCH L3FILTER commands.

The L3FILTER parameter specifies the number of the filter match (*filter-id*) for which the entry is being created. Each filter entry is automatically assigned an *entry-id* number. Filter and filter entry numbers are displayed in the output of the SHOW SWITCH L3FILTER command on page 3-138.

The ACTION parameter specifies a comma-separated list of actions to take when a packet matches the filter criteria specified in this entry. If DENY is specified, the packet is discarded. If FORWARD is specified, the packet is forwarded using the default Class of Service (priority). If SENDCOS is specified, the packet is sent to the priority queue specified by the PRIORITY parameter. If SENDEPORT is specified, the new output port is set to the value of the PORT parameter. If SENDMIRROR is specified, the packet is sent to the mirror port. If SETPRIORITY is specified, the packet's 802.1p priority is set to

the value specified by the PRIORITY parameter. If SETTOS is specified, the packet's TOS (Type of Service) field is set to the value specified by the NEWTOS parameter. The default is FORWARD. If MOVEPRIOTOTOS is specified, the IP TOS field in the frame is replaced with the 802.1 priority value. If MOVETOSTOPRIO is specified, the 802.1 priority field in the frame is replaced with the IP TOS value, this also determines the egress priority queue. If NODROP is specified, matching frames previously marked for dropping are not dropped. If SENDEPORT is specified and the frame is not a broadcast, multicast, DLF frame or marked for dropping or to be sent to the CPU, the new output port is set to the value of the PORT parameter. If SENDNONUNICASTTOPORT is specified, matching frames that are broadcast or DLF, multicast, marked for dropping or to be sent to the CPU are sent to the egress port specified by the PORT parameter. If SETIPDSCP is specified and the frame is an IPv4 frame, the Diffserv Codepoint field in the frame is set to the value specified by the NEWIPDSCP parameter. Actions that modify both the IPTOS and IP DSCP values in the frame are mutually exclusive. The default is FORWARD.

If the SENDEPORT action directs packets to a particular egress port, then the packet is transmitted from the mirror port with a VLAN tag.

The DIPADDR parameter specifies the destination IP addresses to match.

The EPORT parameter specifies the egress port number to be matched by this filter entry, if the EMPORT parameter in the filter match is set to TRUE. The default is no port, that is, the filter entry does not apply to any egress ports. If the EMPORT parameter in the filter match is set to FALSE, the EPORT parameter is ignored, and the filter entry applies to all egress ports.

The IPDSCP parameter indicates the value to match to the IPv4 packet Diffserv Codepoint field for this entry. The range of values for this parameter is from 0 to 63.

The IPORT parameter specifies the ingress port number to be matched by this filter entry, if the IMPORT parameter in the filter match is set to TRUE. The default is no port, that is, the filter entry does not apply to any ingress ports. If the IMPORT parameter in the filter match is set to FALSE, the IPORT parameter is ignored, and the filter entry applies to all ingress ports.

The NEWIPDSCP parameter indicates the value to set in an IPv4 packet Diffserv Codepoint field when the ACTION parameter is set to SETIPDSCP. The range of values for this parameter is from 0 to 63.

The NEWTOS parameter specifies the new type of service value, assigning a new value to the TOS precedence field in the IP Header. This parameter is used when the ACTION parameter is set to SETTOS.

The PORT parameter specifies the new output port number. This port overrides the egress port selected by the Forwarding Database.

The PRIORITY parameter specifies the packet priority. There are eight levels of priority, from 0 to 7. This parameter is used when the ACTION parameter is set to SETPRIORITY or SENDCOS.

The PROTOCOL parameter specifies the IP protocol to match.

The PROTOCOL parameter specifies the IP protocol to match if the SWITCH L3FILTER MATCH value is set to PROTOCOL.

The SIPADDR parameter specifies the source IP address to match.

The TCPACK parameter specifies the ACK (acknowledgement) flag in the TCP header to match, if the protocol is TCP. This parameter is required if TCPACK is specified in the ADD or SET SWITCH L3FILTER MATCH parameter, otherwise it is invalid.

The TCPDPORT parameter specifies the destination TCP port to match, if the protocol is TCP.

The TCPFIN parameter specifies the FIN flag in the TCP header to match, if the protocol is TCP. This parameter is required if TCPFIN is specified in the ADD or SET SWITCH L3FILTER MATCH parameter, otherwise it is invalid.

The TCPSPORT parameter specifies the source TCP port to match, if the protocol is TCP.

The TCPSYN parameter specifies the SYN flag in the TCP header to match, if the protocol is TCP. This parameter is required if TCPSYN is specified in the ADD or SET SWITCH L3FILTER MATCH parameter, otherwise it is invalid.

The TOS parameter specifies the type of service to match.

The TTL parameter specifies the *Time to Live* to match.

The TYPE parameter specifies a protocol-type number to match. The number is entered in hexadecimal, e.g. 0800 for an Ethernet type II IP packet. This parameter may not be used with any other packet field matching criteria, nor may it be used with the SETTOS action. With all other packet matching criteria there is an implicit match to an IP protocol Ethernet type II packet.

The UDPDPORT parameter specifies the UDP destination port to match, if the protocol is UDP.

The UDPSPORT parameter specifies the UDP source port to match, if the protocol is UDP.

Example To add a filter to block Telnet sessions, use the commands:

```
ADD SWITCH L3FILTER MATCH=tcpdport,prot
ADD SWITCH L3FILTER=1 ENTRY ACTION=deny PROT=tcp TCPDPORT=23
```

Related Commands DELETE SWITCH L3FILTER ENTRY
SET SWITCH L3FILTER ENTRY
SHOW SWITCH L3FILTER

ADD SWITCH L3FILTER MATCH

Syntax

```
ADD SWITCH L3FILTER MATCH={DIPADDR|IPDSCP|PROTOCOL|
   SIPADDR|TCPACK|TCPFIN|TCPDPORT|TCPSPORT|TCPSYN|TOS|TTL|
   UDPDPORT|UDPSPORT}[,...] [DCLASS={A|B|C|HOST}]
   [EMPORT={YES|NO|ON|OFF|TRUE|FALSE}] [IMPORT={YES|NO|ON|
   OFF|TRUE|FALSE}] [NOMATCHACTION={SETPRIORITY|SENDCOS|
   SETTOS|DENY|SENDEPORT|SENDMIRROR|MOVEPRIOTOTOS|
   MOVETOSTOPRIO|SETIPDSCP|SENDNONUNICASTTOPORT|
   FORWARD}[,...]] [NOMATCHDSCP=1..63] [NOMATCHPORT=port-
   number] [NOMATCHPRIORITY=0..7] [NOMATCHTOS=0..7]
   [SCLASS={A|B|C|HOST}] [TYPE={802|ETHII|SNAP}]
```

where:

- *port-number* is the switch port number, in the range 1 to m.

Description

This command adds a filter that specifies the matching filter criteria used for the hardware-based packet filtering mechanism.

Up to 16 filters matches may be created.

Each filter is automatically assigned a *filter-id* number, which is displayed in the output of the SHOW SWITCH L3FILTER command on page 3-138. Once the filter has been created, entries must be added using the ADD SWITCH L3FILTER ENTRY command on page 3-55.

Enabling the Internet Group Management Protocol (IGMP) with the ENABLE IP IGMP command also enables Layer 3 filtering. IGMP uses two Layer 3 filters, so the number of available filters is reduced by two. IGMP cannot be enabled unless two filters are still available.

The SWITCH HWFILTER CLASSIFIER commands may not be used with the SWITCH L3FILTER commands.

The MATCH parameter specifies a comma-separated list of packet fields and/or types to match. There is no default.

The DCLASS parameter specifies the IP destination address mask to apply to the destination IP address field in packets when matching destination IP addresses. If A is specified, a Class A mask of 255.0.0.0 is used. If B is specified, a Class B mask of 255.255.0.0 is used. If C is specified, a Class C mask of 255.255.255.0 is used. If HOST is specified, a host mask of 255.255.255.255 is used. The default is for no mask to be used (a value of 0). The DCLASS parameter is required if DIPADDR is specified by the MATCH parameter.

The EMPORT parameter specifies whether the filter applies to all egress ports or to a specific one. If NO, OFF, or FALSE is specified, the filter is applied to all egress ports. If YES, ON, or TRUE is specified, the filter is applied to the egress port specified by the EPORT parameter in the ADD or SET SWITCH L3FILTER ENTRY command. The default is FALSE, meaning the filter is applied to all egress ports.

The IMPORT parameter specifies whether the filter applies to all ingress ports or to a specific one. If NO, OFF, or FALSE is specified, the filter is applied to all ingress ports. If YES, ON, or TRUE is specified, the filter is applied to the ingress port specified by the IPORT parameter in the ADD or SET SWITCH

L3FILTER ENTRY command. The default is FALSE, meaning the filter is applied to all ingress ports.

The NOMATCHACTION parameter specifies a comma-separated list of actions to take when a frame matches both the IPORT and EPORT values (if they are specified in the match) on an associated entry but there is no match for the frame contents. If SETPRIORITY is specified, the packet's 802.1p priority is set to the value specified by the PRIORITY parameter. If SENDCOS is specified, the packet is sent to the priority queue specified by the PRIORITY parameter. If SETTOS is specified, the packet's TOS (Type of Service) field is set to the value specified by the NEWTOS parameter. If DENY is specified, the packet is discarded. If SENDEPORT is specified and the frame is not a broadcast, multicast, DLF frame or marked for dropping or to be sent to the CPU, the new output port is set to the value of the PORT parameter. If SENDMIRROR is specified, the packet is sent to the mirror port. If FORWARD is specified, the packet is forwarded using the default Class of Service (priority). If MOVEPRIOTOTOS is specified, the IP TOS field in the frame is replaced with the 802.1p priority value. This also determines the egress priority queue. if SETIPDSCP is specified and the frame is an IPv4 frame, the Diffserv Codepoint field in the frame is set to the value specified by the NEWIPDSCP parameter. Actions that modify both the IP TOS and the IP DSCP values in the frame are mutually exclusive. If SENDNONUNICASTTOPORT is specified, matching frames that are broadcast or DLF, multicast, marked for dropping or to be sent to the CPU are sent to the egress port specified by the PORT parameter. The default is FORWARD.

The NOMATCHDSCP parameter indicates the value to set in an IPv4 packet DiffServe CodePoint field if the NOMATCHACTION parameter is set to SETIPDSCP. The range of values for this parameter is from 0 to 63.

The NOMATCHPORT parameter specifies the new output port number. This port overrides the egress port selected by the forwarding database.

The NOMATCHPRIORITY parameter specifies the packet priority. There are eight levels of priority, from 0 to 7. This parameter is used either if the NOMATCHACTION parameter is set to SETPRIORITY or SENDCOS.

The NOMATCHTOS parameter specifies the new Type of Service value, assigning a new value to the TOS precedence field in the IP header. This parameter is used when the NOMATCHACTION parameter is set to SETTOS.

The SCLASS parameter specifies the IP source address mask to apply to the source IP address field in packets when matching source IP addresses. If A is specified, a Class A mask of 255.0.0.0 is used. If B is specified, a Class B mask of 255.255.0.0 is used. If C is specified, a Class C mask of 255.255.255.0 is used. If HOST is specified, a host mask of 255.255.255.255 is used. The default is to use no mask (a value of 0). The SCLASS parameter is required if SIPADDR is specified by the MATCH parameter.

The TYPE parameter specifies the format of the protocol-type. This parameter may be used with the EMPORT and IMPORT parameters, but not with the other packet matching criteria. When other criteria are used, there is an implicit match to an IP protocol Ethernet type II packet. If 802 is specified, then the match is on the 2-byte DSAP/SSAP field of an 802.3 packet. If ETHII is specified, then the match is on the 2-byte type field of an Ethernet type II packet. If SNAP is specified, then the match is on the 5-byte variable part of the identifier field of a SNAP packet (SNAP identifiers have the format *aa-aa-03-xx-xx-xx-xx-xx*).

Example To add a filter to block Telnet sessions, use the commands:

```
ADD SWITCH L3FILTER MATCH=tcpdport,prot

ADD SWITCH L3FILTER=1 ENTRY ACTION=deny PROT=tcp TCPDPORT=23
```

Related Commands
ADD SWITCH L3FILTER ENTRY
DELETE SWITCH L3FILTER
SET SWITCH L3FILTER MATCH
SHOW SWITCH L3FILTER

ADD SWITCH TRUNK

Syntax `ADD SWITCH TRUNK=`*`trunk`* `PORT=`*`port-list`*

where:

- *trunk* is a character string 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_") and the hyphen (-).
- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.

Description This command adds ports to an existing trunk group on the switch.

The TRUNK parameter specifies the name of the trunk group. The name is not case sensitive. The name uniquely identifies the trunk group. The specified trunk group must already exist.

The PORT parameter specifies the switch ports to be added to the trunk group. Ports specified must not be in another trunk group, and must have the same VLAN configuration. They cannot include the switch's mirroring port. A trunk group can consist of a maximum of 8 fixed or uplink ports but not a mixture of both types.

When a port is added to a trunk group, its current speed and duplex mode settings are ignored and the port is set to autonegotiate to the speed of the trunk group and full duplex mode.

Port trunking must be configured on both ends of the link, or network loops may result.

Example To add ports 5 and 6 to trunk group Trunk1, use the command:

```
ADD SWITCH TRUNK=Trunk1 PORT=5,6
```

Related Commands
CREATE SWITCH TRUNK
DELETE SWITCH TRUNK
DESTROY SWITCH TRUNK
SET SWITCH TRUNK
SHOW SWITCH TRUNK

ADD VLAN PORT

Syntax

```
ADD VLAN={vlan-name|1..4094} PORT={port-list|ALL}
    [FRAME={TAGGED|UNTAGGED}]
```

where:

- *vlan-name* is a unique name for the VLAN 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *vlan-name* cannot be a number or ALL.
- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.

Description This command adds ports to the specified VLAN.

A port can belong to multiple STPs when the port is a member of more than one VLAN. If the port being added to the VLAN also belongs to another STP through concurrent membership of another VLAN, it is not removed from that VLAN or STP.

If as a result of the port addition, ports are moved from one STP to another STP, the two affected STPs are initialised if they are currently enabled. Any previously disabled ports in the STPs are enabled.

The VLAN parameter specifies the name or numerical VLAN Identifier of the VLAN. The name is not case sensitive, although the case is preserved for display purposes. The VLAN must already exist. By default, all ports belong to the default VLAN, with a numerical VLAN Identifier (VID) of 1.

The PORT parameter specifies the ports. All the ports in a trunk group must have the same VLAN configuration. If the command requires that ports be implicitly deleted from the default VLAN and these ports belong to a trunk group, then the command fails. The ports must belong to only one STP after being added to the VLAN. When the command succeeds on a subset of the specified ports but causes errors on the others, then the command as a whole fails and has no effect. The mirror port cannot be added to a VLAN.

The FRAME parameter specifies whether a VLAN tag header is included in each frame transmitted on the specified ports. If TAGGED is specified, a VLAN tag is added to frames prior to transmission. The port is then called a *tagged* port for this VLAN. If UNTAGGED is specified, the frame is transmitted without a VLAN tag. The port is then called an *untagged* port for this VLAN. A port can be untagged for one and only one of the VLANs to which it belongs, or for none of the VLANs to which it belongs. A port can have the FRAME parameter set to TAGGED for zero or more VLANs to which it belongs. It is not possible to add an untagged port to a VLAN when the port is already present in another port-based VLAN, except the default VLAN. When the port is an untagged member of the default VLAN, adding it untagged to another VLAN deletes it from the default VLAN. The default setting is UNTAGGED.

Examples To add port 4 to the port-based *marketing* VLAN, use the command:

```
ADD VLAN=Marketing PORT=4
```

To add port 25 to the *training* VLAN as a tagged port, use the command:

```
ADD VLAN=Training PORT=25 FRAME=TAGGED
```

Related Commands DELETE VLAN PORT
SHOW VLAN

ADD VLANRELAY

Syntax

```
ADD VLANRELAY=name [PROTOCOL=protocoltype] [VLAN={vlan-
name|1..4094}]
```

where:

- *name* is a unique name for the VLAN relay entity 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character (“_”), and the hyphen (-).
- *protocoltype* is either a valid protocol number in hexadecimal notation, or a recognised protocol name. A protocol number is 1 byte for SAP, 2 bytes for ETHII, or 5 bytes for an 802.2 SNAP type packet.
- *vlan-name* is a unique name for the VLAN 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character (“_”), and the hyphen (-). The *vlan-name* cannot be a number or ALL.

Description This command adds a protocol number and/or a VLAN to a VLAN relay entity. At least one protocol and two VLANs must be added to a VLAN relay entity before the entity can begin relaying packets.

The VLANRELAY parameter specifies the unique identifier for the VLAN relay entity. A VLAN relay entity with this name must already exist.

The PROTOCOL parameter specifies an Ethernet protocol number for packets that are to be relayed. A predefined list of common protocols is provided in Table 3-5 on page 3-18. To relay one of these protocols, specify the protocol name as the value for the PROTOCOL parameter. There is also the option of relaying all protocols of a given encapsulation type by use of the keywords “ALL802”, “ALLETHII” and “ALLSNAP”.

Use of the “ALL802”, “ALLETHII” and “ALLSNAP” protocols can cause traffic to be unexpectedly relayed where it is not desired. It is more desirable to explicitly enter the identification numbers of the protocols to be relayed.

The VLAN parameter specifies the name or VLAN identifier of a VLAN to add to the VLAN relay entity. Adding a VLAN allows packets from that VLAN to be received and relayed, and packets from other VLANs to be relayed to that VLAN. The VLAN must already exist, and must be a static VLAN.

Example To add the VLAN whose ID is 2, and all SAP protocols, to VLAN relay entity SNARelay, use the command:

```
ADD VLANRELAY=SNARelay VLAN=2 PROTOCOL=ALL802
```

Related Commands
CREATE VLANRELAY
DELETE VLANRELAY
DESTROY VLANRELAY
SHOW VLANRELAY

CREATE STP

Syntax
```
CREATE STP=stp-name
```

where *stp-name* is a character string 1 to 15 characters long.Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *stp-name* cannot be ALL or DEFAULT.

Description This command creates a Spanning Tree Protocol entity with a unique name. The specified STP must not already exist. The name is not case sensitive, although the case is preserved for display purposes. The STP created is disabled by default. The maximum number of STPs that can be configured is 255.

Example To create a new STP named *company*, use the command:

```
CREATE STP=company
```

Related Commands
DESTROY STP
ENABLE STP
SET STP
SHOW STP

CREATE SWITCH TRUNK

Syntax
```
CREATE SWITCH TRUNK=trunk [PORT=port-list]
   [SELECT={MACSRC|MACDEST|MACBOTH|IPSRC|IPDEST|IPBOTH}]
   [SPEED={10M|100M|1000M}]
```

where:

- *trunk* is a character string 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_") and the hyphen (-).
- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.

Description This command creates a trunk group on the switch and optionally adds ports to it and sets its speed. The maximum number of trunk groups that can be created depends on the particular switch model due to the capabilities of the switch hardware. The switch supports static 802.3ad link aggregation.

The TRUNK parameter specifies the name of the trunk group. The name is not case sensitive, although the case entered is preserved for display purposes. The name uniquely identifies the trunk group. The specified trunk group must not already exist.

The PORT parameter specifies the switch ports to be added to the trunk group. Ports specified must not be in another trunk group, and must have the same VLAN configuration. They cannot include the switch's mirroring port. A trunk group can consist of a maximum of 8 fixed or uplink ports but not a mixture of both types.

The SELECT parameter specifies the port selection criterion for the trunk group. Each packet to be sent on the trunk group is checked, using the selection criterion, and a port in the trunk group chosen down which to send the packet. If MACSRC is specified, the source MAC address is used. If MACDEST is specified, the destination MAC address is used. If MACBOTH is specified, both source and destination MAC addresses are used. If IPSRC is specified, the source IP address is used. If IPDEST is specified, the destination IP address is used. If IPBOTH is specified, both the source and destination IP addresses are used. The user of the switch should choose the value of this parameter to try to spread the load as evenly as possible on the trunk group. The default is MACBOTH.

The SPEED parameter specifies the speed of the ports in the trunk group. For gigabit ports, only the 1000M value is allowed. For switch ports, 10M and 100M values are allowed. The default is 100M. When a port is added to a trunk group, its current speed and duplex mode settings are ignored and the port is set to autonegotiate to the speed of the trunk group and full duplex mode.

Port trunking must be configured on both ends of the link, or network loops may result.

Example To create a trunk group called Trunk1 containing ports 1 to 4, use the command:

```
CREATE SWITCH TRUNK=Trunk1 PORT=1-4
```

Related Commands

ADD SWITCH TRUNK
DELETE SWITCH TRUNK
DESTROY SWITCH TRUNK
SET SWITCH TRUNK
SHOW SWITCH TRUNK

CREATE VLAN

Syntax `CREATE VLAN=`*`vlan-name`* `VID=2..4094 [PROTECTED]`

where:

- *vlan-name* is a unique name for the VLAN 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *vlan-name* cannot be a number or ALL.

Description This command creates a VLAN with a unique name and VLAN identifier (VID), and assigns it to the default STP. To change the VID of an existing VLAN, that VLAN must be destroyed and created again with a modified VID.

A maximum of 254 VLANs can be created with a VID in the range 2 to 4094.

The VLAN parameter specifies a unique name for the VLAN. This name can be more meaningful than the VID and makes administration easier. The VLAN name is used within the switch; it is not transmitted to other VLAN-aware devices, or used in the Forwarding Process or stored in the Forwarding Database. If the VLAN name begins with "vlan" and ends with a number, for instance "vlan1" or "vlan234", then the number must be the same as the VID specified. This avoids confusion when identifying to which VLAN subsequent commands refer.

The VID parameter specifies a unique VLAN identifier for the VLAN. If tagged ports are added to this VLAN, the specified VID is used in the VID field of the tag in outgoing frames. If untagged ports are added to this VLAN, the specified VID acts as an identifier for the VLAN in the Forwarding Database. The default port based VLAN has a VID of 1.

The PROTECTED parameter specifies that the VLAN is a protected VLAN. If a VLAN is protected, Layer 2 traffic is blocked between its ports.

Examples To create a VLAN named *marketing* with a VLAN Identifier of 2, use the command:

```
CREATE VLAN=marketing VID=2
```

To create a VLAN named *vlan42*, which must have a VID of 42, use the command:

```
CREATE VLAN=vlan42 VID=42
```

To create a protected VLAN named *protvlan* with a VLAN Identifier of 3, use the command:

```
CREATE VLAN=protvlan VID=3 PROTECTED
```

Related Commands ADD VLAN PORT
DESTROY VLAN
SHOW VLAN

CREATE VLANRELAY

Syntax `CREATE VLANRELAY=`*`name`*

where *name* is a unique name for the VLAN relay entity 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-).

Description This command creates a VLAN relay entity, which can be used to relay packets of a given protocol type between VLANs. The VLAN relay entity is enabled by default.

For packet relaying to commence, VLANs and protocol types must be added to this entry, using the ADD VLANRELAY command on page 3-62.

The VLANRELAY parameter specifies the unique identifier for the VLAN relay entity. No VLAN relay entity with this name may already exist. Comparisons of VLAN relay entity names are done without regard to the case of letters, although the case of letters is preserved in order to improve readability. For example, "relaying" and "RelayOne" are treated as the same VLAN relay entity name.

Example To create a VLAN relay entity called SNARelay, use the command:

```
CREATE VLANRELAY=SNARelay
```

Related Commands
ADD VLANRELAY
DELETE VLANRELAY
DESTROY VLANRELAY
SHOW VLANRELAY

DELETE STP VLAN

Syntax `DELETE STP=`*`stp-name`* `VLAN={`*`vlan-name`*`|2..4094|ALL}`

where:

- *stp-name* is a character string 1 to 15 characters long.Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *stp-name* cannot be ALL.
- *vlan-name* is a unique name for the VLAN 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *vlan-name* cannot be a number or ALL.

Description This command deletes one or all VLANs from the specified STP, and returns the VLANs to the default STP. A VLAN cannot be explicitly deleted from the default STP. The default VLAN cannot be deleted.

A port can belong to more than one STP after deletion. When a port belongs to multiple VLANs in the same STP, the port remains a member of this STP when a VLAN it was a member of is returned to the default STP.

If as a result of the VLAN deletion, ports are moved from one STP to another STP, the two affected STPs are initialised when they are currently enabled. Any previously disabled ports in the STPs are enabled.

When returned to the default STP, the ports of the VLAN have the default STP parameter values. The ports do not retain any non-default STP configuration that was made when the VLAN was associated with any other STP.

The VLAN parameter specifies the name or numerical VLAN Identifier (VID) of the VLAN to be deleted. If ALL is specified, then all VLANs are deleted from the STP.

Example To delete the Research VLAN from the *company* STP, use the command:

```
DELETE STP=company VLAN=research
```

Related Commands ADD STP VLAN
SHOW STP

DELETE SWITCH FILTER

Syntax `DELETE SWITCH FILTER PORT=port ENTRY=entry-list`

where:

- *entry-list* is an entry number, a range of entry numbers (specified as n-m), or a comma separated list of entry numbers and/or ranges. Entry numbers start at 0 and end at m, where m is the highest filter entry currently defined in the Permanent Forwarding Database. Each port has its own Permanent Forwarding Database.
- *port* is the number of one of the switch ports or an uplink port.

Description This command deletes the specified static filter entry port from the Permanent Forwarding Database. The static filter is deleted on the port specified by the PORT parameter. The ENTRY parameter must specify an existing filter entry in the Permanent Forwarding Database.

Example To delete filter entry 9 on port 2, use the command:

```
DELETE SWITCH FILTER PORT=2 ENTRY=9
```

Related Commands ADD SWITCH FILTER
SHOW SWITCH FILTER

DELETE SWITCH HWFILTER CLASSIFIER

Syntax `DELETE SWITCH HWFILTER CLASSIFIER=`*`classifier-list`*

where:

- *classifier-list* is either an integer in the range 1 to 9999; a range of integers (specified as 0-4) or a comma separated list of classifier numbers and/or ranges (0, 3, 4-9).

Description This command deletes any hardware-based filters associated with the specified classifier(s). All of the specified classifiers must exist and must already be incorporated into a filter entry. The SWITCH HWFILTER CLASSIFIER commands may not be used with the SWITCH L3FILTER commands.

The CLASSIFIER parameter specifies a list of classifiers for which hardware filter entries are to be deleted.

Examples To delete hardware filtering entries based on classifiers 1 to 5 from the switch, use the command:

```
DELETE SWITCH HWFILTER CLASSIFIER=1-5
```

Related Commands ADD SWITCH HWFILTER CLASSIFIER
SET SWITCH HWFILTER CLASSIFIER
SHOW SWITCH HWFILTER

DELETE SWITCH L3FILTER

Syntax `DELETE SWITCH L3FILTER=`*`filter-id`*

where:

- *filter-id* is a decimal number in the range 1 to the number of filters defined.

Description This command deletes the specified filter match criteria. A filter match criteria cannot be deleted if it contains a filter entry. Delete the filter entries and then delete the filter.

Example To delete filter 1, use the command:

```
DELETE SWITCH L3FILTER=1
```

Related Commands ADD SWITCH L3FILTER MATCH
SET SWITCH L3FILTER MATCH
SHOW SWITCH L3FILTER

DELETE SWITCH L3FILTER ENTRY

Syntax `DELETE SWITCH L3FILT=`*`filter-id`* `ENTRY=`*`entry-id`*

where:

- *filter-id* is a decimal number in the range 1 to the number of filters defined.
- *entry-id* is a decimal number in the range 1 to the number of entries defined.

Description This command deletes the specified entry from the specified filter. Both the entry and the filter must already exist. The L3FILTER parameter specifies the number of the filter. The ENTRY parameter specifies the number of the entry to delete. Filter and entry numbers are displayed in the output of the SHOW SWITCH L3FILTER command on page 3-138.

Example To delete entry 3 from filter 1, use the command:

```
DELETE SWITCH L3FILTER=1 ENTRY=3
```

Related Commands ADD SWITCH L3FILTER ENTRY
SET SWITCH L3FILTER ENTRY
SHOW SWITCH L3FILTER

DELETE SWITCH TRUNK

Syntax `DELETE SWITCH TRUNK=`*`trunk`* `PORT={`*`port-list`*`|ALL}`

where:

- *trunk* is a character string 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_") and the hyphen (-).
- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered switch Ethernet port, including uplink ports.

Description This command deletes ports from an existing trunk group on the switch.

The TRUNK parameter specifies the name of the trunk group. The name is not case sensitive. The name uniquely identifies the trunk group. The specified trunk group must already exist.

The PORT parameter specifies switch ports to be deleted from the trunk group. Ports specified must be in the specified trunk group. If ALL is specified, then all ports in the trunk group are deleted.

Example To delete port 3 from trunk group Trunk1, use the command:

```
DELETE SWITCH TRUNK=Trunk1 PORT=3
```

Related Commands

ADD SWITCH TRUNK
CREATE SWITCH TRUNK
DESTROY SWITCH TRUNK
SET SWITCH TRUNK
SHOW SWITCH TRUNK

DELETE VLAN PORT

Syntax

```
DELETE VLAN={vlan-name|1..4094} PORT={port-list|ALL}
```

where:

- *vlan-name* is a unique name for the VLAN 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *vlan-name* cannot be a number or ALL.
- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered switch Ethernet port (including uplink ports).

Description

This command deletes ports from the specified VLAN. An untagged port can be deleted from a VLAN when the port is still a member of a VLAN after the deletion has occurred. If the port does not belong to a VLAN as a tagged port, then the port is implicitly added to the default VLAN as an untagged port. It is not possible to delete a port that belongs only to the default VLAN as an untagged port.

If the port becomes a tagged port as a result of the deletion; that is, the port does not belong to any VLAN as an untagged port, then the ACCEPTABLE switch parameter for the port is set to VLAN. The user is not able to change the ACCEPTABLE parameter for the port.

A tagged port can be deleted from a VLAN if the port is still a member of a VLAN after the deletion has occurred.

If as a result of the port deletion, ports are moved from one STP to another STP, the two affected STPs are initialised when they are presently enabled. Previously disabled ports in the STPs are enabled.

The VLAN parameter specifies the name or numerical VLAN Identifier of the VLAN. The name is *not* case sensitive. The VLAN must already exist.

The PORT parameter specifies the ports to be deleted from the VLAN. If ALL is specified, then all ports belonging to the VLAN are deleted. When the command succeeds on a subset of the specified ports but causes errors on the others, then the command as a whole fails and has no effect.

A port can belong to multiple STPs when the port is a member of more than one VLAN. If the port being deleted from the VLAN also belongs to another STP through concurrent membership of another VLAN, it is not removed from that VLAN or STP.

If a port belongs to a trunk group, all the ports in the trunk group must be specified. A subset of the ports in a trunk group cannot be deleted from the VLAN unless they are first removed from the trunk group.

Example To delete port 3 from the *marketing* VLAN, use the command:

```
DELETE VLAN=marketing PORT=3
```

Related Commands ADD VLAN PORT
SHOW VLAN

DELETE VLANRELAY

Syntax `DELETE VLANRELAY=`*`name`* `[PROTOCOL=`*`protocoltype`*`] [VLAN={`*`vlan-name`*`|1..4094}]`

where:

- *name* is a unique name for the VLAN relay entity 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-).
- *protocoltype* is either a valid protocol number in hexadecimal notation, or a recognised protocol name. A protocol number is 1 byte for SAP, 2 bytes for ETHII, or 5 bytes for an 802.2 SNAP type packet.
- *vlan-name* is a unique name for the VLAN 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *vlan-name* cannot be a number or ALL.

Description This command deletes a protocol number and/or a VLAN from a VLAN relay entity. The relay entity must still contain at least one protocol and two VLANs in order to relay packets.

The VLANRELAY parameter specifies the unique identifier for the VLAN relay entity. A VLAN relay entity with this name must already exist.

The PROTOCOL parameter specifies an Ethernet protocol number for packets that are no longer to be relayed. The protocol number must be currently being relayed. Table 3-5 on page 3-18 lists predefined protocol types.

The VLAN parameter specifies the static VLAN to remove from the VLAN relay entity. The VLAN can be referenced by name or VLAN ID. The VLAN must already exist and must currently be part of the VLAN relay entity.

Example To delete VLAN 2 from VLAN relay entity SNARelay, use the command:

```
DELETE VLANRELAY=SNARelay VLAN=2
```

Related Commands ADD VLANRELAY
CREATE VLANRELAY
DESTROY VLANRELAY
SHOW VLANRELAY

DESTROY STP

Syntax `DESTROY STP={`*`stp-name`*`|ALL}`

where *stp-name* is a character string 1 to 15 characters long.Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *stp-name* cannot be ALL.

Description This command destroys the specified Spanning Tree Protocol entity, or all STPs except the default STP. An STP cannot be destroyed if VLANs still belong to the STP.

The STP parameter specifies the name of the STP. The name is not case sensitive, although the case is preserved for display purposes. The STP specified must exist. The default STP cannot be destroyed. If ALL is specified, then all STPs except the default STP is destroyed. When ALL is specified and the command succeeds on a subset of STPs but causes errors on the others, then the command as a whole fails and has no effect.

Examples To destroy the *company* STP, use the command:

```
DESTROY STP=company
```

To remove all user created STPs from the switch, none of which have VLANs belonging to them, use the command:

```
DESTROY STP=ALL
```

Related Commands
CREATE STP
DELETE STP VLAN
DISABLE STP
ENABLE STP
SET STP
SHOW STP

DESTROY SWITCH TRUNK

Syntax `DESTROY SWITCH TRUNK=`*`trunk`*

where *trunk* is a character string 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_") and the hyphen (-).

Description This command destroys a trunk group on the switch. The trunk group must be empty, that is, it must not contain any ports.

The TRUNK parameter specifies the name of the trunk group. The name is not case sensitive. The name uniquely identifies the trunk group. The specified trunk group must already exist.

Example To destroy a trunk group called Trunk1, use the command:

```
DESTROY SWITCH TRUNK=Trunk1
```

Related Commands
ADD SWITCH TRUNK
CREATE SWITCH TRUNK
DELETE SWITCH TRUNK
SET SWITCH TRUNK
SHOW SWITCH TRUNK

DESTROY VLAN

Syntax
`DESTROY VLAN={`*`vlan-name`*`|2..4094|ALL}`

where *vlan-name* is a unique name for the VLAN 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *vlan-name* cannot be a number or ALL.

Description
This command destroys the specified static VLAN or all static VLANs in the switch. The default VLAN, which has a numerical VLAN Identifier (VID) of 1, cannot be destroyed. If ALL is specified, then all VLANs except the default VLAN are destroyed. A VLAN cannot be destroyed when ports still belong to it or other modules are attached to it.

The RESET GARP command on page 5-12 of *Chapter 5, Generic Attribute Registration Protocol (GARP)* can be used to destroy dynamic VLANs. However, the dynamic VLANs may be recreated if the switch receives GARP packets after the RESET GARP command has been executed. Disabling a GVRP instance destroys all dynamic VLANs created by the GVRP instance. Dynamic VLANs exist only when GVRP is enabled.

Examples
To destroy the VLAN with the VLAN Identifier of 1234, use the command:

```
DESTROY VLAN=1234
```

To remove all user created VLANs from the switch, none of which have any member ports, use the command:

```
DESTROY VLAN=ALL
```

Related Commands
CREATE VLAN
SHOW VLAN

DESTROY VLANRELAY

Syntax `DESTROY VLANRELAY=`*name*

where *name* is a unique name for the VLAN relay entity 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-).

Description This command destroys a VLAN relay entity. Packet relaying as configured in this VLAN relay entity immediately stops.

The VLANRELAY parameter specifies the unique identifier for the VLAN relay entity. A VLAN relay entity with this name must already exist.

Example To destroy the VLAN relay entity called SNARelay, use the command:

```
DESTROY VLANRELAY=SNARelay
```

Related Commands
ADD VLANRELAY
CREATE VLANRELAY
DELETE VLANRELAY
SHOW VLANRELAY

DISABLE STP

Syntax `DISABLE STP={`*stp-name*`|ALL}`

where *stp-name* is a character string 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *stp-name* cannot be ALL.

Description This command disables operation of the Spanning Tree Algorithm for the specified STP or for the entire switch. User created STPs are disabled by default. The default STP is disabled on switch start-up. An STP should be disabled only when its part of the LAN topology is free of loops. When there is a loop in the topology, the performance of the LAN can be significantly reduced.

This command overrides the DISABLE STP PORT and ENABLE STP PORT commands. Once an STP has been disabled by this command, no port belonging to that STP can be enabled or disabled. The STP must be enabled before ports belonging to the STP are enabled or disabled.

Disabling an STP does not affect the debug status of that STP set by the ENABLE STP DEBUG command. However, because the STP is disabled, STP debugging produces no information.

Disabling STP operation on a port may affect the operation of GARP. Each GARP application has a GIP component whose actions depend on whether the port is in the STP Forwarding state.

Examples To disable the *company* STP, use the command:

```
DISABLE STP=company
```

To disable all STPs on the switch, use the command:

```
DISABLE STP=ALL
```

Related Commands
CREATE STP
DESTROY STP
ENABLE STP
SET STP
SHOW STP

DISABLE STP DEBUG

Syntax

```
DISABLE STP[={stp-name|ALL}] DEBUG={MSG|PKT|STATE|ALL}
   PORT={port-list|ALL}

DISABLE STP DEBUG={MSG|PKT|STATE|ALL} PORT={port-list|ALL}
```

where:

- *stp-name* is a character string 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *stp-name* cannot be ALL.
- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered switch Ethernet port (including uplink ports).

Description This command disables STP debugging options for the specified STP or ports. The DEBUG parameter specifies the debugging modes that are to be disabled (Table 3-16 on page 3-75).

If a port is a member of multiple VLANs, then all these VLANs must belong to the same STP.

A port can belong to more than one STP when the port is a member of two or more VLANs that belong to different STPs.

If ALL is specified, all debugging is disabled.

Table 3-16: STP debugging options.

Option	Debug Mode	Description
MSG	Message	Decoded display of received and transmitted STP packets
PKT	Packet	Raw ASCII display of received and transmitted STP packets
STATE	State	Port state transitions.
ALL	All	All debug options

The PORT parameter specifies the ports where the debug mode is disabled.

The PORT parameter can be supplied with the STP name. If no STP name is provided, it assumes ALL. On the port parameter, the port list does not have to perfectly match all the STP port members so the command still succeeds as a whole.

The STP parameter specifies the STP for which the debugging mode is disabled. If an STP is specified, then the PORT parameter is invalid and all ports in the STP have the debug mode disabled.

The debug status of a port is not changed if the port is moved out of its current STP by one of the following commands: the ADD VLAN PORT, DELETE VLAN PORT, ADD STP VLAN, DELETE STP VLAN. This command is effective on disabled ports or disabled STPs, but produces no debugging information until the ports and the STP are enabled.

Examples To disable the STATE debugging mode for the *company* STP, use the command:

```
DISABLE STP=company DEBUG=STATE
```

To disable all debug modes for all STPs, use the command:

```
DISABLE STP=ALL DEBUG=ALL
```

To disable the MSG debugging mode on ports 5 to 8, use the command:

```
DISABLE STP DEBUG=MSG PORT=5-8
```

Related Commands ENABLE STP DEBUG
SHOW STP DEBUG

DISABLE STP PORT

Syntax `DISABLE STP[={stp-name|ALL}] PORT={port-list|ALL}`

where:

- *stp-name* is a character string 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen(-).
- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered switch Ethernet port (including uplink ports).

Description This command disables operation of the Spanning Tree Algorithm on the specified ports. This command is effective when the STP that the port belongs to is currently enabled. Disabling the operation of STP on a port does not affect the port's ability to receive and transmit frames.

A port can belong to multiple STPs when the port is a member of more than one VLAN.

A port can belong to a single STP. This means that when a port is member of multiple VLANs, all these VLANs must belong to the same STP.

This command only disables STP operation - normal switch processing continues. Disabled ports that are part of an enabled STP can still forward packets.

Disabling the Spanning Tree Algorithm on one or more ports puts those ports in the Disabled state; all BPDUs received on these ports are discarded.

Disabling an STP port does not affect the debug status of the port as set by the ENABLE STP DEBUG command. However, no STP debugging information is produced on a disabled port.

Disabling STP operation on a port may affect the operation of GARP. Each GARP application has a GIP component whose actions depend upon whether the port is in the STP Forwarding state.

The STP parameter specifies the STP instance for which the port is disabled. If no value is provided, the default is ALL.

The PORT parameter specifies the ports. If ALL is specified, all ports in the switch are disabled. When the command succeeds on a subset of the specified ports but causes errors on the others, then the command as a whole fails and has no effect.

Examples To disable the Spanning Tree Algorithm from using port 4, use the command:

```
DISABLE STP PORT=4
```

To disable STP on all ports, use the command:

```
DISABLE STP PORT=ALL
```

To disable STP on just the administration network and only on port 4, use the command:

```
DISABLE STP=ADMIN PORT=4
```

Related Commands
ENABLE STP PORT
SET STP PORT
SHOW STP PORT

DISABLE SWITCH AGEINGTIMER

Syntax `DISABLE SWITCH AGEINGTIMER`

Description This command disables the ageing timer from ageing out dynamically learned entries in the Forwarding Database. The default setting for the ageing timer is enabled.

Example To disable the ageing of learned MAC addresses, use the command:

```
DISABLE SWITCH AGEINGTIMER
```

Related Commands
ENABLE SWITCH AGEINGTIMER
SET SWITCH AGEINGTIMER
SHOW SWITCH

DISABLE SWITCH DEBUG

Syntax `DISABLE SWITCH DEBUG={ARL|CMIC|DMA|QOS|S5600|PHY|ALL}`

Description This command disables the specified switch debug mode or all switch debugging. The DEBUG parameter specifies the switch debug mode to be disabled (Table 3-17 on page 3-78).

Table 3-17: Switch debugging options.

Debug Options	Description
ARL	Operations related to the Forwarding Database.
CMIC	Operations at the CMIC layer
DMA	Operations related to Direct Memory Access requests.
QOS	Operations related to Quality of Service
S5600	Operations related to the switching hardware.
PHY	Operations related to the PHY port interfaces.
ALL	All debug options

Example To disable all switch debugging, use the command:

```
DISABLE SWITCH DEBUG=ALL
```

Related Commands ENABLE SWITCH DEBUG
SHOW SWITCH

DISABLE SWITCH HWFILTER

Syntax `DISABLE SWITCH HWFILTER`

Description This command disables classifier-based packet filtering.

Hardware filtering is automatically disabled when the last filter match is removed, however this command may be used to manually disable filtering if this is required.

Some other modules and processes (e.g. IGMP snooping) require filtering to be enabled at all times. If any of these are active when the DISABLE SWITCH HWFILTER command is entered, it has no effect and an error message results.

Example To disable existing classifier-based packet filters, use the command:

```
DISABLE SWITCH HWFILTER
```

Related Commands ENABLE SWITCH HWFILTER
DISABLE SWITCH HWFILTER

DISABLE SWITCH L3FILTER

Syntax `DISABLE SWITCH L3FILTER`

Description This command disables hardware-based Layer 3 packet filtering.

Hardware filtering is automatically disabled when the last filter match is removed; however, this command may be used to manually disable filtering if this is required. Some other modules and processes (e.g. IGMP snooping) require filtering to be enabled at all times. If any of these are active when the DISABLE SWITCH L3FILTER command is entered, it has no effect and an error message results.

Example To disable existing hardware-based Layer 3 packet filters, use the command:

```
DISABLE SWITCH L3FILTER
```

Related Commands ENABLE SWITCH L3FILTER
SHOW SWITCH L3FILTER

DISABLE SWITCH LEARNING

Syntax `DISABLE SWITCH LEARNING`

Description This command disables the dynamic learning and updating of the Forwarding Database. The default setting for the learning function is enabled.

If switch learning is disabled and the ageing timer has aged out all dynamically learned filter entries, only statically entered MAC source addresses are used to decide which packets to forward or discard. If the switch finds no matching entries in the Forwarding Database during the Forwarding Process, then all switch ports in the VLAN are flooded with the packet, except the port on which the packet was received.

Example To disable the switch learning function, use the command:

```
DISABLE SWITCH LEARNING
```

Related Commands ENABLE SWITCH LEARNING
SHOW SWITCH

DISABLE SWITCH MIRROR

Syntax `DISABLE SWITCH MIRROR`

Description This command disables traffic mirroring on the switch. Mirrored traffic is stopped from being sent on the switch's mirror port. The mirror port and mirror settings for the sources of mirror traffic remain configured. The default state of switch mirroring is disabled.

Example To disable traffic mirroring, use the command:

```
DISABLE SWITCH MIRROR
```

Related Commands ENABLE SWITCH MIRROR
SET SWITCH MIRROR
SET SWITCH PORT
SHOW SWITCH
SHOW SWITCH PORT

DISABLE SWITCH PORT

Syntax `DISABLE SWITCH PORT={`*`port-list`*`|ALL} [FLOW=PAUSE]`

where:

- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered switch Ethernet port, including uplink ports.

Description This command disables a port or group of ports on the switch, or disables the flow control mechanism. When a port is disabled, it no longer sends or receives packets. Disabling a switch does not disable STP operation on the port. Ports should be disabled when faulty wiring or equipment is attached to them or as a security measure to stop access from intruders. Switch ports are enabled by default.

The PORT parameter specifies the port to be disabled or which are to have flow control methods disabled.

The FLOW parameter specifies the type of flow control to be disabled for the port. If PAUSE is specified, flow control is disabled for full duplex ports by sending PAUSE frames. PAUSE is enabled by default.

Example To disable ports 2, 3, 4 and 6, use the command:

```
DISABLE SWITCH PORT=2-4,6
```

Related Commands ENABLE SWITCH PORT
SHOW SWITCH PORT

DISABLE VLAN DEBUG

Syntax `DISABLE VLAN={`*`vlan-name`*`|1..4094|ALL} DEBUG={PKT|ALL}`

where *vlan-name* is a unique name for the VLAN 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *vlan-name* cannot be a number or ALL.

Description This command disables packet debugging or all debugging for the specified VLAN or all VLANs. The default is for all VLAN debugging to be disabled.

The DEBUG parameter specifies the VLAN debugging mode to be disabled. If PKT is specified, the packet debug mode (displaying raw ASCII packets) is disabled. If ALL is specified, all debugging is disabled.

Example To disable packet debugging on the *marketing* VLAN, use the command:

```
DISABLE VLAN=marketing DEBUG=PKT
```

Related Commands ENABLE VLAN DEBUG
SHOW VLAN DEBUG

DISABLE VLANRELAY

Syntax `DISABLE VLANRELAY=`*`name`*

where *name* is a unique name for the VLAN relay entity 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-).

Description This command disables packet relaying by the VLAN relay entity. The entity must exist and must be currently enabled. VLAN relay entities are enabled by default upon creation.

Example To disable packet relaying by the VLAN relay entity SNARelay, use the command:

```
DISABLE VLANRELAY=SNARelay
```

Related Commands ADD VLANRELAY
DELETE VLANRELAY
ENABLE VLANRELAY

DISABLE VLANRELAY DEBUG

Syntax `DISABLE VLANRELAY=`*`name`* `DEBUG`

where *name* is a unique name for the VLAN relay entity 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-).

Description This command disables the output of debugging information about packets relayed by a VLAN relay entity. The relay entity must already exist and VLAN relay debugging must currently be enabled. Debugging of VLAN relay entities is disabled by default.

Example To disable the display of packets relayed by the VLAN relay entity SNARelay, use the command:

```
DISABLE VLANRELAY=SNARelay DEBUG
```

Related Commands ADD VLANRELAY
DELETE VLANRELAY
ENABLE VLANRELAY
ENABLE VLANRELAY DEBUG

ENABLE STP

Syntax `ENABLE STP{=`*`stp-name`*`|ALL}`

where *stp-name* is a character string 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *stp-name* cannot be ALL.

Description This enables operation of the Spanning Tree Algorithm for the specified STP or for the entire switch. If the Spanning Tree Algorithm is to be run on a VLAN, the VLAN must be added to an STP that is enabled. User created STPs are disabled by default. The default STP is disabled on switch start-up.

This command is required before the DISABLE STP PORT and ENABLE STP PORT commands can be used. Once an STP has been enabled by this command it is then possible to enable or disable any port belonging to that STP.

Enabling STP operation on a port may affect the operation of GARP. Each GARP application has a GIP component whose actions depend upon whether the port is in the STP Forwarding state.

Examples To enable the *company* STP, use the command:

```
ENABLE STP=company
```

To enable all STPs, use the following command:

```
ENABLE STP=ALL
```

Related Commands

CREATE STP
DESTROY STP
DISABLE STP
SET STP
SHOW STP

ENABLE STP DEBUG

Syntax

```
ENABLE STP={stp-name|ALL} DEBUG={MSG|PKT|STATE|ALL}
   [OUTPUT=CONSOLE] [TIMEOUT={1..4000000000|NONE}]

ENABLE STP={stp-name|ALL} DEBUG={MSG|PKT|STATE|ALL}
   PORT={port-list|ALL} [OUTPUT=CONSOLE]
   [TIMEOUT={1..4000000000|NONE}]
```

where:

- *stp-name* is a character string 1 to 15 characters long.Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character (“_”), and the hyphen (-). The *stp-name* cannot be ALL.
- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.

Description

This command enables STP debugging for the specified STP, or ports.

Enabling debug may flood the receiving Telnet session or asynchronous port with raw data.

The STP parameter specifies the STP for which the debugging mode is enabled. If an STP is specified and ports are specified with the PORT parameter, then debug is enabled on the specified port on the specified STP. If an STP is not specified or ALL is specified with the STP parameter, and ports are specified with the PORT parameter, then debug mode for the listed ports is enabled on the STPs with the listed port as a member.

The DEBUG parameter specifies the debugging modes that are to be enabled. If ALL is specified, all debugging modes for the STP or ports are enabled. The other modes can be enabled independently of each other. The DEBUG parameter must be specified before the PORT parameter. The debugging modes enabled by each option are shown in Table 3-18 on page 3-83.

Table 3-18: STP debugging options .

Option	Description
MSG	Decoded display of received and transmitted STP packets
PKT	Raw ASCII display of received and transmitted STP packets
STATE	Port state transitions. For RSTP, states for all state machines are displayed as well the current role of the port.
ALL	All debug options

The OUTPUT parameter set to CONSOLE specifies that the debugging information produced is sent to the console. The debugging data is by default sent to the port on which it received the ENABLE STP DEBUG command. Use this option if the ENABLE STP DEBUG command is used in a script, since a script is not received on a port.

The PORT parameter specifies the ports where the debug mode is enabled, or all ports on the switch. The DEBUG parameter must be specified before the PORT parameter.

The TIMEOUT parameter specifies the time in seconds that debugging is enabled on the specified ports. This reduces the risk of the switch and the display being overloaded with too much debugging information. This value overrides previous STP debugging timeout values for these ports, even if they were specified for other debugging modes. If TIMEOUT is not specified, the time out is the most recent TIMEOUT value set in an ENABLE STP DEBUG command, or NONE if none had been set.

The debug status of a port is not changed if the port is moved out of its current STP by one of the following commands: the ADD VLAN PORT, DELETE VLAN PORT, ADD STP VLAN, DELETE STP VLAN. This command is effective on disabled ports or disabled STPs, but produces no debugging information until the ports and the STP are enabled.

Examples To view STATE debugging information for the *company* STP for the next 25 seconds, use the command:

```
ENABLE STP=company DEBUG=STATE TIMEOUT=25
```

To enable all debug modes for all STPs with output to the console and no timeout value, use this command:

```
ENABLE STP=ALL DEBUG=ALL OUTPUT=CONSOLE
```

To enable the message debug mode on ports 5 to 8 indefinitely, use the command:

```
ENABLE STP DEBUG=MSG PORT=5-8 TIMEOUT=NONE
```

Related Commands DISABLE STP DEBUG
SHOW STP DEBUG

ENABLE STP PORT

Syntax

```
ENABLE STP[={stp-name|ALL}] PORT={port-list|ALL}
```

where:

- *stp-name* is a character string 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-).
- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.

Description

This command enables operation of the Spanning Tree Algorithm on the specified ports.

The STP parameter specifies the STP that is to have ports enabled. If no value is entered, the default is ALL

If the PORT parameter specified is ALL, then all ports within the matching STP instance are enabled. This command is effective when the Spanning Tree Algorithm is enabled for the STP to which the port belongs.

Enabling an STP port may cause reconfiguration of the Spanning Tree to which the port belongs because STP messages (BPDUs) are generated on the port.

Enabling STP operation on a port may affect the operation of GARP. Each GARP application has a GIP component whose actions depend upon whether the port is in the STP Forwarding state.

The DISABLE STP command overrides the results of the DISABLE STP PORT and ENABLE STP PORT commands. Once a STP has been disabled by this command it is not possible to enable or disable any port belonging to that STP. The STP must be enabled first before any port belonging to that STP can be enabled or disabled.

Examples

To enable the Spanning Tree Algorithm to use port 4, use the command:

```
ENABLE STP PORT=4
```

To enable STP on all ports, use the command:

```
ENABLE STP PORT=ALL
```

To enable STP on just the administration network and only on port 4, use the command:

```
ENABLE STP=ADMIN PORT=4
```

Related Commands

DISABLE STP PORT
SET STP PORT
SHOW STP PORT

ENABLE SWITCH AGEINGTIMER

Syntax `ENABLE SWITCH AGEINGTIMER`

Description This command enables the ageing timer to age out dynamically learned entries in the Forwarding Database. The default setting for the ageing timer is enabled.

If the ageing timer ages out all dynamically learned filter entries, and switch learning is disabled, only statically entered MAC source addresses are used to decide which packets to forward or discard. If the switch finds no matching entries in the Forwarding Database during the Forwarding Process, then all switch ports in the VLAN are flooded with the packet, except the port on which the packet was received.

Example To enable the ageing of learned MAC addresses, use the command:

```
ENABLE SWITCH AGEINGTIMER
```

Related Commands
DISABLE SWITCH AGEINGTIMER
SET SWITCH AGEINGTIMER
SHOW SWITCH

ENABLE SWITCH BIST

Syntax

```
ENABLE SWITCH BIST

ENABLE SWITCH BIST INSTANCE=instance
```

where:

- *bist* is a single integer number.
- *instance* is 0 or 1 and specifies a switch instance on 48 port switches.

Description This command runs a set of built in self tests on the external packet buffer memory and internal memories of a switch chip (or instance). The INSTANCE parameter must be specified *only* for switches with 48 ports.

This procedure may only be performed by authorised service personnel. Network and switch performance are affected by the use of this command. After using this command the switch must be rebooted. The switch ports should be disconnected from any live networks before enabling the test.

Examples To enable the BIST test, use the command:

```
ENABLE SWITCH BIST=0
```

ENABLE SWITCH DEBUG

Syntax

```
ENABLE SWITCH DEBUG={ARL|CMIC|DMA|QOS|S5600|PHY|ALL}
    [OUTPUT=CONSOLE] [TIMEOUT={1..4000000000|NONE}]
```

Description

This command enables the specified switch debug mode or all switch debugging.

Enabling debug may flood the receiving Telnet session or asynchronous port with raw data.

The DEBUG parameter specifies the switch debug mode to be disabled (Table 3-17 on page 3-78). If ALL is specified, all switch debugging modes are enabled.

Table 3-19: Switch debugging options.

Debug Options	Description
ARL	Operations related to the Forwarding Database.
CMIC	Operations at the CMIC layer.
DMA	Operations related to Direct Memory Access requests.
QOS	Operations related to Quality of Service.
S5600	Operations related to the switching hardware.
PHY	Operations related t the PHY port interfaces.
ALL	All debug options.

The OUTPUT parameter set to CONSOLE specifies that the debugging information produced is sent to the console. The debugging data is by default sent to the port on which it received the ENABLE SWITCH DEBUG command. Use this option if the command is used in a script, since a script is not received on a port.

The TIMEOUT parameter specifies the time in seconds that switch debugging is enabled. This reduces the risk of the switch and the display being overloaded with too much debugging information. This value overrides any previous switch debugging timeout values, even if they were specified for other debugging modes. If TIMEOUT is not specified, the time out is the most recent TIMEOUT value previously used in an ENABLE VLAN DEBUG command, or NONE if it has not been previously set.

Example

To enable the ARL switch debugging mode, use the command:

```
ENABLE SWITCH DEBUG=ARL
```

Related Commands

DISABLE SWITCH DEBUG
SHOW SWITCH

ENABLE SWITCH HWFILTER

Syntax `ENABLE SWITCH HWFILTER`

Description This command enables hardware-based Layer 3 packet filtering.

Hardware filtering is automatically enabled when the first filter match is added. This command may be used to re-enable filtering if it has been temporarily disabled by the DISABLE SWITCH HWFILTER command, or to enable the filtering mechanism prior to the addition of the first filter match.

Example To enable existing hardware-based Layer 3 packet filters, use the command:

```
ENABLE SWITCH HWFILTER
```

Related Commands DISABLE SWITCH HWFILTER
SHOW SWITCH HWFILTER

ENABLE SWITCH L3FILTER

Syntax `ENABLE SWITCH L3FILTER`

Description This command enables hardware-based Layer 3 packet filtering.

Hardware filtering is automatically enabled when the first filter match is added. However this command may be used to re-enable filtering if it has been temporarily disabled by the DISABLE SWITCH L3FILTER command, or to enable the filtering mechanism prior to the addition of the first filter match.

Example To enable existing hardware-based Layer 3 packet filters, use the command:

```
ENABLE SWITCH L3FILTER
```

Related Commands DISABLE SWITCH L3FILTER
SHOW SWITCH L3FILTER

ENABLE SWITCH LEARNING

Syntax `ENABLE SWITCH LEARNING`

Description This command enables the dynamic learning and updating of the Forwarding Database. The default setting for the learning function is enabled.

Example To enable the switch learning function, use the command:

```
ENABLE SWITCH LEARNING
```

Related Commands DISABLE SWITCH LEARNING
SHOW SWITCH

ENABLE SWITCH MIRROR

Syntax `ENABLE SWITCH MIRROR`

Description This command enables traffic mirroring on the switch. Mirrored traffic is sent on the switch's mirror port as long as a valid one is defined and sources of mirror traffic have been configured. The default state of mirroring is disabled.

Four or more ports set to mirror traffic to the mirror port may significantly reduce switch performance.
If a packet is Layer 3 switched and mirrored, then the packet is always transmitted from the mirror port with a VLAN tag.

Example To enable traffic mirroring, use the command:

```
ENABLE SWITCH MIRROR
```

Related Commands DISABLE SWITCH MIRROR
SET SWITCH MIRROR
SET SWITCH PORT
SHOW SWITCH
SHOW SWITCH PORT

ENABLE SWITCH PORT

Syntax `ENABLE SWITCH PORT={`*`port-list`*`|ALL} [FLOW=PAUSE]`

where:

- *port-list* is a single port number or a group of port numbers, either a comma separated list, a range (specified as n-m) or a combination of the two. Port numbers start at 1 and end at m, where m is the highest numbered switch Ethernet port.

Description This command enables a port or group of ports on the switch, or enables the flow control mechanism. When the port is enabled, it sends and receives packets subject to the operation of STP. Enabling the switch port does not affect STP on the port. Switch ports are enabled by default.

Use the SET SWITCH PORT command to enable a port that has been disabled by the Port Security function, rather than this command.

The PORT parameter specifies the port to be enabled, or which are to have flow control methods enabled.

The FLOW parameter specifies the type of flow control to be enabled for the port. If PAUSE is specified, flow control for full duplex ports by sending PAUSE frames is enabled. PAUSE flow control is enabled by default.

Example To enable ports 2, 4 and 6, use the command:

```
ENABLE SWITCH PORT=2,4,6
```

Related Commands DISABLE SWITCH PORT
SHOW SWITCH PORT

ENABLE VLAN DEBUG

Syntax

```
ENABLE VLAN={vlan-name|1..4094|ALL} DEBUG={PKT|ALL}
   [OUTPUT=CONSOLE] [TIMEOUT={1..4000000000|NONE}]
```

where:

- *vlan-name* is a unique name for the VLAN 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *vlan-name* cannot be a number or ALL.

Description This command enables debugging options for the specified VLAN or all VLANs. The default is for all VLAN debugging to be disabled.

Caution: Enabling debug may flood the receiving Telnet session or asynchronous port with raw data.

The DEBUG parameter specifies the debugging mode that is enabled. If PKT is specified, packet debug mode (displaying raw ASCII packets) is enabled. If ALL is specified, all debugging is enabled.

The OUTPUT parameter set to CONSOLE specifies that the debugging information produced is sent to the console. The debugging data is by default sent to the port on which it received the ENABLE VLAN DEBUG command. Use this option if the command is used in a script, since a script is not received on a port.

The TIMEOUT parameter specifies the time in seconds when debugging is enabled on the specified VLAN. This reduces the risk of the switch and the display being overloaded with too much debugging information. This value overrides any previous VLAN debugging timeout values for the VLAN, even if they were specified for other debugging modes. If TIMEOUT is not specified, the time out is the most recent TIMEOUT value used in an ENABLE VLAN DEBUG command or NONE if none had been set.

Example To enable all debugging on the *marketing* VLAN, use the command:

```
ENABLE VLAN=marketing DEBUG=ALL
```

Related Commands DISABLE VLAN DEBUG
SHOW VLAN DEBUG

ENABLE VLANRELAY

Syntax `ENABLE VLANRELAY=`*`name`*

where *name* is a unique name for the VLAN relay entity 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-).

Description This command enables the relaying of packets by the VLAN relay entity. The relay entity must already exist and must be currently disabled. VLAN relay entities are enabled by default upon creation.

Example To enable packet relaying by the VLAN relay entity SNARelay, use the command:

```
ENABLE VLANRELAY=SNARelay
```

Related Commands
ADD VLANRELAY
DELETE VLANRELAY
DISABLE VLANRELAY

ENABLE VLANRELAY DEBUG

Syntax `ENABLE VLANRELAY=`*`name`* `DEBUG`

where *name* is a unique name for the VLAN relay entity 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-).

Description This command enables the output of debugging information about packets relayed by the VLAN relay entity. The relay entity must already exist, and VLAN relay debugging must be currently disabled. Debugging of VLAN relay entities is disabled by default.

The format of the output messages from packet debugging is as follows:

```
VR: 2->3: 0000cd001234 0000cd004321 040403060708090560403
```

The first part of the output shows which VLANs the packet is being relayed between. The second part shows the packet, with destination and source MAC addresses separated from the payload of the packet.

Example To enable the display of packets relayed by the VLAN relay entity SNARelay, use the command:

```
ENABLE VLANRELAY=SNARelay DEBUG
```

Related Commands
ADD VLANRELAY
DELETE VLANRELAY
DISABLE VLANRELAY DEBUG
ENABLE VLANRELAY

PURGE STP

Syntax `PURGE STP`

Description This command destroys all user created STPs, and restores the defaults to all the configurable parameters (FORWARDDELAY, HELLOTIME, MAXAGE and PRIORITY) in the remaining default STP. The debug parameters for all ports are reset to their defaults. This command returns the STP module to its status when it is first powered on.

Example To purge all STPs, use the command:

```
PURGE STP
```

Related Commands
RESET STP
SET STP
SET STP PORT
SHOW STP
SHOW STP COUNTER

RESET STP

Syntax `RESET STP={`*`stp-name`*`|ALL}`

where *stp-name* is a character string 1 to 15 characters long.Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *stp-name* cannot be ALL.

Description This command resets operation of the Spanning Tree Algorithm for the specified STP, initialises all counters for the specified STP, and initialises all timers on all ports that are members of the STP. Ports remain in the state they were before the reset command was issued, for example, ports that were enabled remain enabled, ports that were disabled remain disabled.

Example To reset the *company* STP, use the command:

```
RESET STP=company
```

Related Commands
PURGE STP
SET STP
SHOW STP
SHOW STP COUNTER

RESET SWITCH

Syntax `RESET SWITCH`

Description This command resets the switch module. All dynamic switch information is cleared. All ports are reset. All counters and timers are reset to zero.

Example To reset the switch module, use the command:

```
RESET SWITCH
```

Related Commands SHOW SWITCH
SHOW SWITCH FDB

RESET SWITCH PORT

Syntax `RESET SWITCH PORT={`*`port-list`*`|ALL} [COUNTER]`

where:

- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.

Description This command resets a port or group of ports on the switch. All packets queued for reception or transmission on the port are discarded, the port is reset at the hardware level and reconfigured to the current speed, and duplex mode is activated. Switch port counters are reset to zero. This command can be used to try to ensure that packets stuck in a queue are cleared, perhaps after a packet storm of some nature.

The PORT parameter specifies the ports to be reset.

The COUNTER parameter specifies that switch port counters be reset only. If the COUNTER parameter is not used, the switch port is fully reset.

Example To reset port 3, use the command:

```
RESET SWITCH PORT=3
```

Related Commands DISABLE SWITCH PORT
ENABLE SWITCH PORT
SHOW SWITCH PORT

SET STP

Syntax

```
SET STP={stp-name|ALL} [FORWARDDELAY=4..30]
   [HELLOTIME=1..10] [MAXAGE=6..40] [MODE={STANDARD|
   RAPID}] [PRIORITY=0..65535] [RSTPTYPE={NORMAL|
   STPCOMPATIBLE}]

SET STP={stp-name|ALL} DEFAULT
```

where:

- *stp-name* is a character string 1 to 15 characters long.Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *stp-name* cannot be ALL.

Description

This command sets parameters used by the Spanning Tree Algorithm for the specified STP. If ALL is specified, then parameters for all STPs on the switch are set. When ALL is specified and the command succeeds on a subset of STPs but causes errors on the others, then the command as a whole fails and has no effect. Each STP has its own independent FORWARDDELAY, HELLOTIME, MAXAGE, and PRIORITY parameters.

The DEFAULT parameter sets the FORWARDDELAY, HELLOTIME, MAXAGE and PRIORITY parameters back to their defaults. This parameter cannot be specified with either of the FORWARDDELAY, HELLOTIME, MAXAGE or PRIORITY parameters.

The FORWARDDELAY parameter sets the time in seconds to control how fast a port changes its spanning tree state when moving towards the Forwarding state. If the mode is set to standard, the value determines how long the port stays in each of the Listening and Learning states which precede the Forwarding state. If the mode is set to rapid, this value determines the maximum time taken to transition from Discarding to Learning and from Learning to Forwarding. This value is used only when the switch is acting as the Root Bridge. Switches not acting as the Root Bridge use a dynamic value for the FORWARDDELAY set by the Root Bridge. The FORWARDDELAY, MAXAGE, and HELLOTIME parameters are interrelated. See the note and formulas below. The default for FORWARDDELAY is 15 seconds.

The HELLOTIME parameter sets the time in seconds between the transmission of switch spanning tree configuration information when the switch is the Root Bridge of the spanning tree or is trying to become the Root Bridge. The default is 2 seconds.

The MAXAGE parameter sets the maximum time in seconds that dynamic STP configuration information is stored in the switch before it is discarded. The default is 20 seconds.

The FORWARDDELAY, MAXAGE and HELLOTIME parameters should be set according to the following formulae, as specified in IEEE 802.1d:
2 x (FORWARDDELAY - 1.0 seconds) >= MAXAGE
MAXAGE >= 2 x (HELLOTIME + 1.0 seconds)

The MODE parameter specifies whether the STP operates in STANDARD mode or RAPID mode. In STANDARD mode, the Spanning Tree Algorithm is run. In RAPID mode, the Rapid Spanning Tree Algorithm is run. The default is

STANDARD. If the mode is changed while the algorithm is running, the STP is reinitialised.

The PRIORITY parameter sets the priority of the switch to become the Root Bridge. The lower the value of the Bridge Identifier, the higher the priority. If the PRIORITY parameter is set by specifying the PRIORTY or DEFAULT parameters, the specified STP is initialised. Counters for the STP are not affected. The default for PRIORITY is 32768.

If the MODE parameter has been set to RAPID, values specified for the PRIORITY parameter must be multiples of 4096. If a value is specified that is not a multiple of 4096, the value is rounded down to the nearest multiple of 4096. The rounding scheme is defined in Table 3-20.

Table 3-20: Rounding scheme for ranges of PRIORITY parameter values when the MODE parameter is set to RAPID.

Lower boundary	Upper boundary	Rounded RSTP Bridge Priority Value
0	4095	0
4096	8191	4096
8192	12287	8192
12288	16383	12288
16384	20479	16384
20480	24575	20480
24576	28671	24576
28672	32767	28672
32768	36863	32768
36864	40959	36864
40960	45055	40960
45056	49151	45056
49152	53247	49152
53248	57343	53248
57344	61439	57344
61440	65535	61440

The RSTPTYPE parameter specifies how the RSTP algorithm operates. If NORMAL is specified, then the algorithm uses rapid port role transitions and transmits and receives RST BPDUs. If STPCOMPATIBLE is specified, then rapid transitions are disabled, standard BPDUs are transmitted and RST BPDUs are discarded. Setting RSTPTYPE to STPCOMPATIBLE allows RSTP to support applications and protocols that may be sensitive to frame duplication and misordering, for example NetBeui. The default is NORMAL.

Setting RSTPTYPE to NORMAL when normal has already been set, sets all ports to the “sending RSTP” state. This is referred to in the IEEE802.1w standard as *mCheck* and is useful for restoring full rapid mode operation when one or more ports on the switch has entered the “sending STP” state. RSTP-capable devices with RSTP set to NORMAL that receive the RST BPDUs

enter the "sending RSTP" state. When an STP BPDU is received after the mCheck operation, either as a result of a device being in rapid mode with RSTPTYPE set to STPCOMPATIBLE or as a result of a device in standard mode, the ports that received the STP BPDUs revert to the "sending STP" state.

Examples To set the forward delay to 22 seconds for the *company* STP, use the command:

```
SET STP=company FORWARDDELAY=2
```

To set the hello time to 3 seconds for the *company* STP, use the command:

```
SET STP=company HELLOTIME=3
```

To set the maximum age to 19 seconds for the *company* STP, use the command:

```
SET STP=company MAXAGE=19
```

To set the priority of the switch becoming the Root Bridge to 100 for the *company* STP, use the command:

```
SET STP=company PRIORITY=100
```

To set the Forward Delay to 12 seconds for all STPs, assuming the FORWARDDELAY-MAXAGE criterion is met for all STPs, use the command:

```
SET STP=ALL FORWARDDELAY=12
```

To set the parameters for the *company* STP to their defaults, use the command:

```
SET STP=company DEFAULT
```

Related Commands PURGE STP
RESET STP
SET STP PORT
SHOW STP

SET STP PORT

Syntax

```
SET STP={stp-name|ALL} PORT={port-list|ALL}
   [PATHCOST=pathcost] [PORTPRIORITY=0..255]
   [EDGEPORT={YES|NO|ON|OFF|TRUE|FALSE}] [PTP={AUTO|ON|
   OFF|YES|NO|TRUE|FALSE}]

SET STP[={stp-name|ALL}] PORT={port-list|ALL} DEFAULT
```

where:

- *stp-name* is a character string 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-).
- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.
- *pathcost* is a value in the range 1 to 1,000,000 if STP is running in standard mode, and 1 to 200,000,000 if STP is running in rapid mode.

Description This command sets various parameters used by the Spanning Tree Algorithm for the specified ports, or all ports within the specified STP, or all STPs.

A port can belong to multiple STPs when the port is a member of more than one VLAN.

The STP parameter specifies an STP name. If no parameter is entered, the default is ALL.

Non-default STP parameter values configured for a port are not retained when the VLAN to which the port belongs is moved to another STP by using the ADD STP VLAN or DELETE STP VLAN commands.

The PORT parameter specifies a list of ports that can belong to any STP. The default is ALL.

The DEFAULT parameter sets the PATHCOST and PORTPRIORITY parameters back to their defaults. This parameter cannot be specified with either of the PATHCOST and PORTPRIORITY parameters. The EDGEPORT and PTP parameters are not affected by this command.

The PATHCOST parameter sets the path cost for each port. The PATHCOST for a LAN port should be set to a maximum of 1,000,000 in standard mode and 200,000,000 in rapid mode. If the port is to be the root port then this value determines the total cost from the switch to the Root Bridge. Each STP has its own independent PATHCOST parameter for each member port. The default PATHCOST values and the range of recommended PATHCOST values depend on the port speed and mode (see Table 3-21 on page 3-97 and Table 3-22 on page 3-97).

Table 3-21: Path cost values and port speed for STANDARD mode.

Port Speed	Default PATHCOST	Recommended PATHCOST range
10Mbps	100	50 - 600
100Mbs	19	10 -60
1Gbps	4	3 -10

Table 3-22: Path cost values and port speed for RAPID mode.

Port Speed	Default PATHCOST	Recommended PATHCOST range
Less than 100 Kb/s	200000000	20000000-200000000
1Mbps	20000000	2000000-20000000
10Mbps	2000000	200000-2000000
100 Mbps	200000	20000-200000
1 Gbps	20000	2000-20000
10 Gbps	2000	200-2000
100 Gbps	200	20-200
1Tbps	20	2-200
10 Tbps	2	2-20

When the MODE of an STP is changed from standard to rapid, or rapid to standard, the PATHCOST parameter is mapped from one range to the other based on relative deviation from the nearest default. We recommend that the PATHCOST values be checked when changing mode to confirm that they are appropriate for the network configuration.

If the PATHCOST of a port has not been explicitly set by the user or the defaults have been restored to the port, then the default PATHCOST for the port varies as the speed of the port varies.

IEEE 802.1d, limited the range of the path cost parameter to a 16 bit unsigned integer value. The recommended path cost values for rapid mode, IEEE 802.1w, make use of the full 32 bit range available in BPDUs. The recommended values for an intermediate link speed can be calculated as 20000000000/(Link Speed in KB/s). This means that the accumulated Path Cost values cannot exceed 32 bits over a concatenation of 20 hops. In LANs where the recommended values defined in IEEE 802.1d and IEEE 802.1w are required to interwork, one set of path cost values must be reconfigured so that they are the same. The range of path costs that can be configured in an older bridge is insufficient to accommodate the range of data rates available.

The PORTPRIORITY parameter sets the value of the priority field contained in the port identifier. The Spanning Tree Algorithm uses the port priority when determining the root port for each switch. The port with the lowest value is considered to have the highest priority. The default is 128. Each STP has its own independent PORTPRIORITY parameter for each member port.

If the MODE parameter is set to RAPID, then the values specified for the PORTPRIORITY parameter must be multiples of 16. If a user specifies a value which is not a multiple of 16, then the value is rounded down to the nearest multiple of 16. The rounding scheme is identified in Table 3-23 on page 3-98.

Table 3-23: Rounding scheme for PORTPRIORITY parameter values when the MODE parameter is set to RAPID.

Lower boundary	Upper boundary	Rounded RSTP Port Priority Value.
0	15	0
16	31	16
32	47	32
48	63	48
64	69	64
80	95	80
96	111	96
112	127	112
128	143	128
144	159	144
160	175	160
176	191	176

Table 3-23: Rounding scheme for PORTPRIORITY parameter values when the MODE parameter is set to RAPID. (Continued)

Lower boundary	Upper boundary	Rounded RSTP Port Priority Value.
192	207	192
208	223	208
224	239	224
240	255	240

The EDGEPORT parameter specifies whether the port is an edge port. An edge port is a port that attaches to a LAN that is known to have no other bridges attached. If NO is specified, then the port is not considered to be an edge port. The values NO, OFF, and FALSE are equivalent. If YES is specified, then the port is considered to be an edgeport. The values YES, ON, and TRUE are equivalent. If EDGEPORT is set to YES and an RST BPDU is received on the port, which indicates that another bridge is connected to the LAN, then the port is no longer treated as an edge port. The default is NO. If STP is running in RAPID mode, then the rapid transition of a port to the Forwarding state depends on the port being considered an edgeport or part of a Point-to-Point link.

The PTP parameter specifies whether the port has a point-to-point connection with another bridge. If AUTO is specified, then the point-to-point status of the port is determined automatically by the switch. If YES is specified, then the port is treated as a point-to-point LAN segment. The values YES, ON, and TRUE are equivalent. If NO is specified, then the port is not treated as a point-to-point LAN segment. The values NO, OFF, and FALSE are equivalent. If STP is running in RAPID mode, then the rapid transition of a port to the Forwarding state depends on the port being considered an edgeport or part of a Point-to-Point link. The default is AUTO.

Examples

To set a port priority of 42 for port 10 in STP1, use the command:

```
SET STP=1 PORT=10 PORTPRIORITY=42
```

To set a path cost of 120 for all ports on all STPs, use the command:

```
SET STP=ALL PORT=ALL PATHCOST=120
```

To set the port parameters for ports 1 to 10 in STP3 to their standard defaults, use the command:

```
SET STP=3 PORT=1-10 DEFAULT
```

To set port 10 in STP3 as an edgeport, use the command:

```
SET STP=3 PORT=10 EDGEPORT=YES
```

To force port 10 in STP3 to be treated as if it were part of a point to point LAN segment, use the command:

```
SET STP=3 PORT=10 PTP=YES
```

Related Commands

PURGE STP
RESET STP
SET STP
SHOW STP

SET SWITCH AGEINGTIMER

Syntax

```
SET SWITCH AGEINGTIMER=10..1000000
```

Description This command sets the threshold value, in seconds, of the ageing timer, after which a dynamic entry in the Layer 2 Forwarding Database is automatically removed. (The maximum setting of 1 000 000 seconds is approximately 11 days 13 hours.) The default is 300 seconds (5 minutes).

Example To set the ageing timer to 180 seconds (3 minutes), use the command:

```
SET SWITCH AGEINGTIMER=180
```

Related Commands

DISABLE SWITCH AGEINGTIMER
ENABLE SWITCH AGEINGTIMER
SHOW SWITCH

SET SWITCH HWFILTER CLASSIFIER

Syntax

```
SET SWITCH HWFILTER CLASSIFIER=classifier-list
   [ACTION={SETPRIORITY|SENDCOS|SETTOS|DENY|SENDEPORT|
   SENDMIRROR|MOVEPRIOTOTOS|MOVETOSTOPRIO|SETIPDSCP|
   SENDNONUNICASTTOPORT|NODROP|FORWARD}[,...]]
   [NEWIPDSCP=0..63] [NEWTOS=0..7]
   [NOMATCHACTION={SETPRIORITY|SENDCOS|SETTOS|DENY|
   SENDEPORT|SENDMIRROR|MOVEPRIOTOTOS|MOVETOSTOPRIO|
   SETIPDSCP|SENDNONUNICASTTOPORT|FORWARD}[,...]]
   [NOMATCHDSCP=dscp-value] [NOMATCHPORT=port-number]
   [NOMATCHPRIORITY=0..7] [NOMATCHTOS=0..7]
   [PORT=port-number] [PRIORITY=0..7]
```

where:

- *classifier-list* is either an integer in the range 1 to 9999; a range of integers (specified as 0-4) or a comma separated list of classifier numbers and/or ranges (0, 3, 4-9).
- *port-number* is the switch port number, in the range 1 to m, where m is the highest numbered Ethernet switch port, including uplink ports.

Description This command sets the properties of hardware-based filters based on the specified classifier(s). All of the specified classifiers must exist and must already be incorporated into a filter entry. The SWITCH HWFILTER CLASSIFIER commands may not be used with the SWITCH L3FILTER commands.

A port can belong to multiple STPs when the port is a member of more than one VLAN. A port can belong to a single STP. This means that when the port is member of multiple VLANs, all these VLANs must belong to the same STP.

The ACTION parameter specifies a comma-separated list of actions to take when a packet matches the filter criteria specified in this entry. If SETPRIORITY is specified, the packet's 802.1p priority is set to the value specified by the PRIORITY parameter. If SENDCOS is specified, the packet is

sent to the priority queue specified by the PRIORITY parameter. If SETTOS is specified, the packet's TOS (Type of Service) field is set to the value specified by the NEWTOS parameter. If DENY is specified, the packet is discarded. If SENDEPORT is specified and the frame is not a broadcast, multicast, DLF frame or marked for dropping or to be sent to the CPU, the new output port is set to the value of the PORT parameter. If SENDMIRROR is specified, the packet is sent to the mirror port. If FORWARD is specified, the packet is forwarded using the default Class of Service (priority). The default is FORWARD. If MOVEPRIOTOTOS is specified, the IP TOS field in the frame is replaced with the 802.1 priority value. If MOVETOSTOPRIO is specified, the 802.1 priority field in the frame is replaced with the IP TOS value, this also determines the egress priority queue. If SETIPDSCP is specified and the frame is an IPv4 frame, the Diffserv Codepoint field in the frame is set to the value specified by the NEWIPDSCP parameter. Actions that modify both the IP TOS and IP DSCP values in the frame are mutually exclusive. If SENDNONUNICASTTOPORT is specified, matching frames that are broadcast or DLF, multicast, marked for dropping or to be sent to the CPU are sent to the egress port specified by the PORT parameter. If NODROP is specified, matching frames previously marked for dropping are not dropped.

If the SENDEPORT action directs packets to a particular egress port, then the packet is transmitted from the mirror port with a VLAN tag.

The NEWIPDSCP parameter indicates the value to set in an IPv4 packet Diffserv CodePoint field when the ACTION parameter is set to SETIPDSCP. The range of values for this parameter is from 0 to 63.

The NEWTOS parameter specifies the new type of service value, assigning a new value to the TOS precedence field in the IP Header. This parameter is used only when the ACTION parameter is set to SETTOS.

The NOMATCHACTION parameter specifies a comma-separated list of actions to take when a frame matches both the IPORT and EPORT values (if they are specified in the match) on an associated entry but there is no match for the frame contents.If SETPRIORITY is specified, the packet's 802.1p priority is set to the value specified by the PRIORITY parameter. If SENDCOS is specified, the packet is sent to the priority queue specified by the PRIORITY parameter. If SETTOS is specified, the packet's TOS (Type of Service) field is set to the value specified by the NEWTOS parameter. If DENY is specified, the packet is discarded. If SENDEPORT is specified and the frame is not a broadcast, multicast, DLF frame or marked for dropping or to be sent to the CPU, the new output port is set to the value of the PORT parameter. If SENDMIRROR is specified, the packet is sent to the mirror port. If FORWARD is specified, the packet is forwarded using the default Class of Service (priority). If MOVEPRIOTOTOS is specified the IP TOS field in the frame is replaced with the 802.1 priority value. If MOVETOSTOPRIO is specified, the 802.1 priority field in the frame is replaced with the IP TOS value, this also determines the egress priority queue. If SETIPDSCP is specified and the frame is an IPv4 frame, the Diffserv Codepoint field in the frame is set to the value specified by the NEWIPDSCP parameter. Actions that modify both the IP TOS and IP DSCP values in the frame are mutually exclusive. If SENDNONUNICASTTOPORT is specified, matching frames that are broadcast or DLF, multicast, marked for dropping or to be sent to the CPU are sent to the egress port specified by the PORT parameter. The default is FORWARD.

The NOMATCHDSCP parameter indicates the value to set in an IPv4 packet Diffserv CodePoint field when the NOMATCHACTION parameter is set to SETIPDSCP. The range of values for this parameter is from 0 to 63.

The NOMATCHPORT parameter specifies the new output port number. This port overrides the egress port selected by the Forwarding Database.

The NOMATCHPRIORITY parameter specifies the packet priority. There are eight levels of priority, from 0 to 7. This parameter is used only when the NOMATCHACTION parameter is set to SETPRIORITY or SENDCOS.

The NOMATCHTOS parameter specifies the new type of service value, assigning a new value to the TOS precedence field in the IP Header. This parameter is used only when the NOMATCHACTION parameter is set to SETTOS.

The PORT parameter specifies the new output port number. This port overrides the egress port selected by the Forwarding Database.

The PRIORITY parameter specifies the packet priority. There are eight levels of priority from 0 to 7. This parameter is used only when the ACTION parameter is set to SETPRIORITY or SENDCOS.

Examples To change the hardware packet filter that acts on traffic matched by classifier 1 so that it to denies this traffic, use the command:

```
SET SWITCH HWFILTER CLASSIFIER=1 ACTION=DENY
```

To set the transmit priority on all packets matching Classifier 100 to 3, and set the transmit priority on packets that partially match this classifier to 0, use the command:

```
SET SWITCH HWFILTER CLASSIFIER=100 ACTION=SENDCOS
   NOMATCHACTION=SENDCOS PRIORITY=3 NOMATCHPRIORITY=0
```

Related Commands
ADD SWITCH HWFILTER CLASSIFIER
DELETE SWITCH HWFILTER CLASSIFIER
SHOW SWITCH HWFILTER

SET SWITCH L3AGEINGTIMER

Syntax

```
SET SWITCH L3AGEINGTIMER=[30..43200]
```

Description This command sets the threshold value, in seconds, of the ageing timer for dynamic entries in the Layer 3 forwarding database. After a cycle of this timer, entries not used during the cycle remain in the table but their hit bits are reset to zero. After the next cycle, entries with hit bit still set to zero are deleted. Therefore, entries in the table are deleted when they are unused during two consecutive cycles of the timer. The default is 900.

This command can be executed only when the hardware forwarding entry ageing timer is enabled by using the ENABLE SWITCH AGEINGTIMER command. This ageing timer is enabled by default.

Examples To set the threshold of the Layer 3 forwarding table ageing timer to 30 minutes, use the command:

```
SET SWITCH L3AGEINGTIMER=1800
```

Related Commands

DISABLE SWITCH AGEINGTIMER
ENABLE SWITCH AGEINGTIMER
SHOW SWITCH

SET SWITCH L3FILTER ENTRY

Syntax

```
SET SWITCH L3FILTER=filter-id ENTRY=entry-id
   [ACTION={SETPRIORITY|SENDCOS|SETTOS|DENY|SENDEPORT|
   SENDMIRROR|MOVEPTIOTOTOS|MOVETOSTOPRIO|SETIPDSCP|
   SENDNONUNICASTTOPORT|FORWARD}[,…]] [DIPADDR=ipadd]
   [EPORT=port-number] [IPORT=port-number]
   [NEWIPDSCP=0..63] [NEWTOS=0..7] [PORT=port-number]
   [PRIORITY=0…7] [PROTOCOL={TCP|UDP|ICMP|IGMP|protocol}]
   [SIPADDR=ipadd] [TCPACK={TRUE|FALSE}]
   [TCPDPORT=port-id] [TCPFIN={TRUE|FALSE}]
   [TCPSPORT=port-id] [TCPSYN={TRUE|FALSE}] [TOS=0..7]
   [TTL=0..255] [TYPE=protocol-type] [UDPSPORT=port-id]
   [UDPDPORT=port-id]
```

where:

- *filter-id* is a decimal number in the range 1 to the number of filters defined.
- *entry-id* is a decimal number in the range 1 to the number of entries defined.
- *ipadd* is an IP address in dotted decimal notation.
- *port-number* is the switch port number, in the range 1 to m, where m is the highest numbered Ethernet switch port, including uplink ports.
- *protocol* is an IP protocol number in the range 1 to 255.
- *port-id* is an IP port number.
- *protocol-type* is a valid protocol-type number. A protocol type number is 2 bytes for Ethernet type II and 802.3 (DSAP/SSAP) encapsulation, or 5 bytes for SNAP encapsulation, and is specified in hexadecimal.

Description

This command modifies the selector values for an existing filter entry. The L3FILTER and ENTRY parameters specify the number of the filter and the filter entry to be modified, respectively. Filter and filter entry numbers are displayed in the output of the SHOW SWITCH L3FILTER command on page 3-138. The SWITCH HWFILTER CLASSIFIER commands may not be used with the SWITCH L3FILTER commands.

A port can belong to multiple STPs when the port is a member of more than one VLAN.

The ACTION parameter specifies a comma-separated list of actions to take when a packet matches the filter criteria specified in this entry. If SETPRIORITY is specified, the packet's 802.1p priority is set to the value specified by the PRIORITY parameter. If SENDCOS is specified, the packet's priority CoS queue is set to the value specified by the PRIORITY parameter. If SETTOS is specified, the packet's TOS field is set to the value specified by the NEWTOS parameter. If DENY is specified, the packet is discarded. If SENDEPORT is specified, and the new frame is not a broadcast, multicast, DLF frame or marked for dropping or to be sent to the CPU, the new output port is set to the value of the PORT parameter. If SENDMIRROR is specified, the

packet is sent to the mirror port. If MOVETOPRIOTOTOS is specified, the IP TOS field in the frame is replaced with the 802.1p priority value. If MOVETOSTOPRIO is specified, the 802.1p priority field in the frame is replaced with the IP TOS value - this also determines the egress priority queue. If SETIPDSCP is specified, and the frame is an IPv4 frame, the Diffserv Codepoint field in the frame is set to the value specified by the NEWIPDSCP parameter. Actions that modify both the IP TOS and IP DSCP values in the frame are mutually exclusive. If SENDNONUNICASTTOPORT is specified, matching frames that are broadcast or DLF, multicast, marked for dropping or to be sent to the CPU are sent to the egress port specified by the PORT parameter. If NODROP is specified, matching frames previously marked for dropping are not dropped. If FORWARD is specified, the packet is forwarded using the default Class of Service (priority). The default is FORWARD.

The ACTION parameter specifies a comma-separated list of actions to take when a packet matches the filter criteria specified in this entry. If SETPRIORITY is specified, the packet's 802.1p priority is set to the value specified by the PRIORITY parameter. If SENDCOS is specified, the packet's priority CoS queue is set to the value specified by the PRIORITY parameter. If SETTOS is specified, the packet's TOS field is set to the value specified by the NEWTOS parameter. If DENY is specified, the packet is discarded. If SENDEPORT is specified, the new output port is set to the value of the PORT parameter. If SENDMIRROR is specified, the packet is sent to the mirror port. If FORWARD is specified, the packet is forwarded using the default Class of Service (priority). The default is FORWARD.

The DIPADDR parameter specifies the destination IP addresses to match.

The EPORT parameter specifies the egress port number to be matched by this filter entry, if the EMPORT parameter in the filter match is set to TRUE. The default is no port, that is, the filter entry does not apply to any egress ports. If the EMPORT parameter in the filter match is set to FALSE, the EPORT parameter is ignored, and the filter entry applies to all egress ports.

The IPORT parameter specifies the ingress port number to be matched by this filter entry, if the IMPORT parameter in the filter match is set to TRUE. The default is no port, that is, the filter entry does not apply to any ingress ports. If the IMPORT parameter in the filter match is set to FALSE, the IPORT parameter is ignored, and the filter entry applies to all ingress ports.

The NEWIPDSCP parameter indicates the value to set in an IPv4 packet Diffserv CodePoint field when the ACTION parameter is set to SETIPDSCP. The range of values for this parameter is from 0 to 63.

The NEWTOS parameter specifies the new type of service value, assigning a new value to the TOS precedence field in the IP Header. This parameter is used only when the ACTION parameter is set to SETTOS.

The PORT parameter specifies the new output port number. This port overrides the egress port selected by the Forwarding Database.

The PRIORITY parameter specifies the new packet priority. There are eight levels of priority from 0 to 7. This parameter is used only when the ACTION parameter is set to SETPRIORITY or SENDCOS.

The PROTOCOL parameter specifies the IP protocol to match.

The SIPADDR parameter specifies the source IP address to match.

The TCPACK parameter specifies the ACK (acknowledgement) flag in the TCP header to match when the protocol is TCP. This parameter is required when TCPACK is specified in the ADD or SET SWITCH L3FILTER MATCH parameter, otherwise it is invalid.

The TCPDPORT parameter specifies the destination TCP port to match when the protocol is TCP.

The TCPFIN parameter specifies the FIN flag in the TCP header to match when the protocol is TCP. This parameter is required when TCPFIN is specified in the ADD or SET SWITCH L3FILTER MATCH parameter, otherwise it is invalid.

The TCPSPORT parameter specifies the source TCP port to match, if the protocol is TCP.

The TCPSYN parameter specifies the SYN flag in the TCP header to match, if the protocol is TCP. This parameter is required if TCPSYN is specified in the ADD or SET SWITCH L3FILTER MATCH parameter, otherwise it is invalid.

The TOS parameter specifies the type of service to match.

The TTL parameter specifies the *Time to Live* to match.

The TYPE parameter specifies a protocol-type number to match. The number is entered in hexadecimal, e.g. 0800 for an Ethernet type II IP packet. This parameter may not be used with any other packet field matching criteria, nor may it be used with the SETTOS action. With all other packet matching criteria there is an implicit match to an IP protocol Ethernet type II packet.

The UDPDPORT parameter specifies the UDP destination port to match, if the protocol is UDP.

The UDPSPORT parameter specifies the UDP source port to match, if the protocol is UDP.

Example To modify entry 2 of filter 1 to match UDP port 23, use the command:

```
SET SWITCH L3FILTER=1 ENTRY=2 PROT=udp TCPDPORT=23
```

Related Commands

ADD SWITCH L3FILTER ENTRY
DELETE SWITCH L3FILTER ENTRY
SHOW SWITCH L3FILTER

SET SWITCH L3FILTER MATCH

Syntax

```
SET SWITCH L3FILTER=filter-id MATCH={DIPADDR|IPDSCP|
   PROTOCOL|SIPADDR|TCPACK|TCPFIN|TCPDPORT|TCPSPORT|
   TCPSYN|TOS|TTL|UDPDPORT|UDPSPORT}[,...] [DCLASS={A|B|C|
   HOST}] [EMPORT={YES|NO|ON|OFF|TRUE|FALSE}]
   [IMPORT={YES|NO|ON|OFF|TRUE|FALSE}]
   [NOMATCHACTION={SETPRIORITY|SENDCOS|SETTOS|DENY|
   SENDEPORT|SENDMIRROR|MOVEPRIOTOTOS|MOVETOSTOPRIO|
   SETIPDSCP|SENDNONUNICASTTOPORT|FORWARD}[,...]]
   [NOMATCHDSCP=0..63] [NOMATCHPORT=port-number]
   [NOMATCHPRIORITY=0..7] [NOMATCHTOS=0..7] [SCLASS={A|B|
   C|HOST}] [TYPE={802|ETHII|SNAP}]
```

where:

- *filter-id* is a decimal number in the range 1 to the number of filters defined.
- *port-number* is the switch port number, in the range 1 to m.

Description

This command modifies an existing filter that specifies matching filter criteria for the packet filtering mechanism. The L3FILTER parameter specifies the number of the filter to be modified. Filter numbers are displayed in the output of the SHOW SWITCH L3FILTER command on page 3-138. The SWITCH HWFILTER CLASSIFIER commands may not be used with the SWITCH L3FILTER commands.

A port can belong to multiple STPs when the port is a member of more than one VLAN.

The MATCH parameter specifies a comma-separated list of packet fields and/or types to match. There is no default.

The DCLASS parameter specifies the IP destination address mask to apply to the destination IP address field in packets when matching destination IP addresses. If A is specified, a Class A mask of 255.0.0.0 is used. If B is specified, a Class B mask of 255.255.0.0 is used. If C is specified, a Class C mask of 255.255.255.0 is used. If HOST is specified, a host mask of 255.255.255.255 is used.

The EMPORT parameter specifies whether the filter applies to all egress ports or to a particular egress port specified in a filter entry. If NO, OFF, or FALSE is specified, the filter is applied to all egress ports. If YES, ON, or TRUE is specified, the filter is applied to the egress port specified by the EPORT parameter in the ADD or SET SWITCH L3FILTER ENTRY command. The default is FALSE, meaning the filter applies to all egress ports.

The IMPORT parameter specifies whether the filter applies to all ingress ports or to a particular ingress port specified in a filter entry. If NO, OFF, or FALSE is specified, the filter is applied to all ingress ports. If YES, ON, or TRUE is specified, the filter is applied to the ingress port specified by the IPORT parameter in the ADD or SET SWITCH L3FILTER ENTRY command. The default is FALSE, meaning the filter applies to all ingress ports.

The NOMATCHACTION parameter specifies a comma-separated list of actions to take when a frame matches both the IPORT and EPORT values (if they are specified in the match) on an associated entry but there is no match for the frame contents. If SETPRIORITY is specified, the packet's 802.1p priority is

set to the value specified by the PRIORITY parameter. If SENDCOS is specified, the packet is sent to the priority queue specified by the PRIORITY parameter. If SETTOS is specified, the packet's TOS (Type of Service) field is set to the value specified by the NEWTOS parameter. If DENY is specified, the packet is discarded. If SENDEPORT is specified and the frame is not a broadcast, multicast, DLF frame or marked for dropping or to be sent to the CPU, the new output port is set to the value of the PORT parameter. If SENDMIRROR is specified, the packet is sent to the mirror port. If FORWARD is specified, the packet is forwarded using the default Class of Service (priority). If MOVEPRIOTOTOS is specified, the IP TOS field in the frame is replaced with the 802.1p priority value. This also determines the egress priority queue. if SETIPDSCP is specified and the frame is an IPv4 frame, the Diffserv Codepoint field in the frame is set to the value specified by the NEWIPDSCP parameter. Actions that modify both the IP TOS and the IP DSCP values in the frame are mutually exclusive. If SENDNONUNICASTTOPORT is specified, matching frames that are broadcast or DLF, multicast, marked for dropping or to be sent to the CPU are sent to the egress port specified by the PORT parameter. The default is FORWARD.

The NOMATCHDSCP parameter indicates the value to set in an IPv4 packet DiffServe CodePoint field if the NOMATCHACTION parameter is set to SETIPDSCP. The range of values for this parameter is from 0 to 63.

The NOMATCHPORT parameter specifies the new output port number. This port overrides the egress port selected by the forwarding database.

The NOMATCHPRIORITY parameter specifies the packet priority. There are eight levels of priority, from 0 to 7. This parameter is used only when the NOMATCHACTION parameter is set to SETPRIORITY or SENDCOS.

The NOMATCHTOS parameter specifies the new Type of Service value, assigning a new value to the TOS precedence field in the IP header. This parameter is used only when the NOMATCHACTION parameter is set to SETTOS.

The SCLASS parameter specifies the IP source address mask to apply to the source IP address field in packets when matching source IP addresses. If A is specified, a Class A mask of 255.0.0.0 is used. If B is specified, a Class B mask of 255.255.0.0 is used. If C is specified, a Class C mask of 255.255.255.0 is used. If HOST is specified, a host mask of 255.255.255.255 is used.

The TYPE parameter specifies the format of the protocol-type. This parameter may be used with the EMPORT and IMPORT parameters, but not with the other packet matching criteria. When other criteria are used, there is an implicit match to an IP protocol Ethernet type II packet. If 802 is specified, then the match is on the 2-byte DSAP/SSAP field of an 802.3 packet. If ETHII is specified, then the match is on the 2-byte type field of an Ethernet type II packet. If SNAP is specified, then the match is on the 5-byte variable part of the identifier field of a SNAP packet (SNAP identifiers have the format *aa-aa-03-xx-xx-xx-xx-xx*).

Example To modify filter 1 to match UDP port, use the command:

```
SET SWITCH L3FILTER=1 MATCH=udpdport,prot
```

Related Commands

ADD SWITCH L3FILTER ENTRY
ADD SWITCH L3FILTER MATCH
DELETE SWITCH L3FILTER
SHOW SWITCH L3FILTER

SET SWITCH MIRROR

Syntax `SET SWITCH MIRROR={NONE|`*`port`*`}`

where:

- *port* is a single switch port number. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port.

Description This command sets the mirror port for the switch, and removes it from the default VLAN. If another port was previously set as the mirror port, this command returns it to the default VLAN as an untagged port. The mirror port is the port to which mirrored traffic is sent. The source of mirror traffic is configured with the SET SWITCH PORT command.

The MIRROR parameter specifies the switch port where mirror traffic is to be sent. The specified port must belong only to the default VLAN as an untagged or tagged port. The port cannot be part of a trunk group. If the value NONE is specified, no mirror port is defined for the switch and mirroring is disabled. The mirror port cannot be added to any VLAN.

Port mirroring does not duplicate packets. If one mirrored packet is captured in different ports, only one copy of the packet is sent to the mirror port.

If a packet is Layer 3 switched and mirrored, then the packet is always transmitted from the mirror port with a VLAN tag.

Example To set the mirror port to port 12, use the command:

```
SET SWITCH MIRROR=12
```

Related Commands

DISABLE SWITCH MIRROR
ENABLE SWITCH MIRROR
SET SWITCH PORT
SHOW SWITCH
SHOW SWITCH PORT

SET SWITCH PORT

Syntax

```
SET SWITCH PORT={port-list|ALL} [ACCEPTABLE={ALL|VLAN}]
   [BCLIMIT={NONE|limit}] [DESCRIPTION=description]
   [DLFLIMIT={NONE|limit}] [EGRESSLIMIT={NONE|DEFAULT|0|
   1000..127000|8..1016}] [INFILTERING={OFF|ON}]
   [INGRESSLIMIT={NONE|DEFAULT|0|64..127000|8..1016}]
   [LEARN={NONE|0|1..256] [INTRUSIONACTION={DISABLE|
   DISCARD|TRAP}] [MCLIMIT={NONE|limit}] [MIRROR={BOTH|
   NONE|RX|TX}] [MODE={AUTONEGOTIATE|MASTER|SLAVE}]
   [MULTICASTMODE={A|B|C}] [SPEED={AUTONEGOTIATE|10MHALF|
   10MFULL|10MHAUTO|10MFAUTO|100MHALF|100MFULL|100MHAUTO|
   100MFAUTO|1000MHALF|1000MFULL|1000MHAUTO|1000MFAUTO}]
```

where:

- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.
- *limit* is a decimal number, from 0 to the maximum value of the limit variable based on the particular switch hardware. The maximum packet storm protection limit is 262143.
- *description* is a string 1 to 47 characters long. Valid characters are any printable characters.

Description

This command modifies the value of parameters for switch ports.

The PORT parameter specifies the ports for which parameters are modified. When the command succeeds on a subset of the specified ports but causes errors on the others, then the command as a whole fails and has no effect. Reference in the descriptions below to an individual port should be taken as a reference to all ports selected by the PORT parameter. If packet storm protection limits are set on the switch, the PORT parameter must specify complete processing blocks (see the note after the BCLIMIT parameter description).

While the user may specify SET SWITCH PORT commands using groups of ports, the create config command on page 1-62 of Chapter 1, Operation generates a separate SET SWITCH PORT command for each port.

The ACCEPTABLE parameter sets the Acceptable Frame Types parameter, in the Ingress Rules, which controls reception of VLAN-tagged and VLAN-untagged frames on the port. If ALL is specified, then the Acceptable Frame Types parameter is set to Admit All Frames. If VLAN is specified, the parameter is set to Admit Only VLAN-tagged Frames, and any frame received that carries a null VLAN Identifier (VID) is discarded by the ingress rules. Untagged frames and priority-tagged frames carry a null VID. Untagged frames admitted according to the ACCEPTABLE parameter have the VID of the VLAN for which the port is untagged associated with them. The ACCEPTABLE parameter can be set only when the port is untagged for one VLAN. In this case, the default is ALL, admitting all tagged and untagged frames. If the port is tagged for all the VLANs to which it belongs, the ACCEPTABLE parameter is automatically set to VLAN, and cannot be changed to admit untagged frames.

The BCLIMIT parameter specifies a limit on the rate of reception of broadcast packets for the port(s). The value of this parameter represents a per second rate of packet reception above which packets are discarded for broadcast packets. If the value NONE or 0 is specified, then packet rate limiting for broadcast packets is turned off. If another value is specified, the reception of broadcast packets is limited to this number. See the note below for important information about packet rate limiting. The default is NONE.

Limiting packet reception rates for different classes of packets depends on the particular switch hardware. In particular, groups of ports may have to have the same limits set, and the same limit may be set for the different types of packets, depending on the hardware. When packet rate limits are set on switches with this type of constraint, the most current parameter values supersede earlier ones. When a command for specific ports changes parameters for other ports, a message reports these changes.
Packet storm protection limits cannot be set for each individual port on the switch, but can be set for each processing block of ports. The processing blocks are sets of 8 ports (e.g. as many as are applicable of ports 1-8, 9-16 and 17-24) and each uplink port is a further processing block. Therefore, a 16-port switch has four processing blocks and a 24-port switch has five. The two uplink ports are numbered sequentially after the last port, and therefore are 17 and 18 for a 16-port and 25 and 26 for a 24-port switch. Only one limit can be set per processing block, and then applies to all three packet types. Thus each of the packet types are either limited to this value, or unlimited (NONE).

The DESCRIPTION parameter can be used to describe the port. It is displayed by the SHOW SWITCH PORT command on page 3-140, but does not affect the operation of the switch in any way. The default is no description.

The DLFLIMIT parameter specifies a limit on the rate of reception of destination lookup failure packets for the port. The value of this parameter represents a per second rate of packet reception above which packets will be discarded for destination lookup failure packets. If the value NONE or 0 is specified, then packet rate limiting is turned off for these packets. If another value is specified, the reception of these packets is limited to this number. See the note after the BCLIMIT parameter description for important information about packet rate limiting. The default is NONE. If packet storm protection limits are set on the switch, the PORT parameter must specify complete processing blocks.

A destination lookup failure packet is one for which the switch hardware does not have a record of the destination address of the packet, either Layer 2 or Layer 3 address. These packets are passed to the CPU for further processing, so limiting the rate of reception of these packets may be a desirable feature to improve system performance.

The EGRESSLIMIT parameter specifies the maximum bandwidth for traffic egressing the specified port(s), in kbps (10/100 Mbps ports) or Mbps (Gigabit ports). If NONE or 0 (zero) is specified, egress limiting is disabled for the specified port. For 10/100 Mbps ports the input value (1000..127000) in kbps is rounded up to the nearest 1000 (or 1 Mbps). For Gigabit ports the input value (8..1016) in Mbps is rounded up to the nearest 8 Mbps. The default is NONE.

The INFILTERING parameter enables or disables Ingress Filtering of frames admitted according to the ACCEPTABLE parameter, on the specified ports. Each port on the switch belongs to one or more VLANs. If INFILTERING is set to ON, Ingress Filtering is enabled; frames received on a specified port are admitted when the port belongs to the VLAN with which the frames are associated. Conversely, frames are discarded when the port does not belong to

the VLAN with which the frames are associated. Untagged frames admitted by the ACCEPTABLE parameter are admitted since they have the numerical VLAN Identifier (VID) of the VLAN for which the port is an untagged member. If OFF is specified, Ingress Filtering is disabled, and no frames are discarded by this part of the Ingress Rules. The default is OFF.

The INGRESSLIMIT parameter specifies the maximum bandwidth for traffic ingressing the specified port(s), in kbps (10/100 Mbps ports) or Mbps (Gigabit ports). If NONE or 0 (zero) is specified, ingress limiting is disabled for the specified port. For 10/100 Mbps ports the input value (64..127000) in kbps is rounded up to the nearest 64kbps if below 1000, otherwise it is rounded up to the nearest 1000 (or 1 Mbps). For Gigabit ports the input value (8..1016) in Mbps is rounded up to the nearest 8 Mbps. The default is NONE.

The INTRUSIONACTION parameter specifies the action taken when the port receives packets from addresses that are not part of the learned list of addresses as specified by the LEARN parameter. If DISCARD is specified, packets are discarded that come from MAC addresses not on the port's learn list. If TRAP is specified, these packets are discarded and an SNMP trap is generated. If DISABLE is specified, the packet is discarded the first time it is received, an SNMP trap is generated, and the port is disabled. To re-enable the port, disable the Port Security function on the port. The default is DISCARD.

The LEARN parameter specifies whether the security feature of limiting the number of MAC addresses learned on this port is enabled. If NONE or zero is specified, all MAC addresses are learned on this port and the Port Security function is disabled. When a port has been automatically disabled by the switch's port security, setting the Learn parameter to 0 (zero) re-enables it. If a number from 1 to 256 is specified, the switch stops learning MAC addresses on this port when the number of MAC addresses is reached, and the port is locked. If the LEARN parameter is set to a value lower than the number of MAC addresses currently learned, then the port is unlocked if previously locked, all learned MAC addresses are cleared from the forwarding database for the port, and learning restarts. Packets from other addresses after this time are handled as intrusion packets (see the INTRUSIONACTION parameter). The default is NONE.

Learned addresses on locked ports can be saved as part of the switch configuration and become part of the configuration after a power cycle by using the create config command on page 1-62 of Chapter 1, Operation. If the configuration is not saved when there is a locked list for a port, the learning process begins again after the router is restarted.

The MCLIMIT parameter specifies a limit on the rate of reception of multicast packets for the port. The value of this parameter represents a per second rate of packet reception above which packets are discarded for multicast packets. If the value NONE or 0 is specified, then packet rate limiting for multicast packets is turned off. If another value is specified, the reception of multicast packets is limited to this number. See the note after the BCLIMIT parameter description for important information about packet rate limiting. The default is NONE. If packet storm protection limits are set on the switch, the PORT parameter must specify complete processing blocks.

The MIRROR parameter specifies the role of these ports as a source of mirror traffic. If NONE is specified, no traffic received or sent on these ports is mirrored. If RX is specified, all traffic received on these ports is mirrored. If TX is specified, all traffic transmitted is mirrored. If BOTH is specified, all traffic

received and transmitted is mirrored. Traffic is mirrored only when a mirror port is defined and mirroring is enabled. The default is NONE.

Caution: Four or more ports set to mirror traffic to the mirror port may significantly reduce switch performance.

The MULTICASTMODE parameter indicates how the switch handles traffic addressed to a multicast group to which the specified port or list of ports belongs. If A is specified, all traffic is flooded on all ports on the VLAN, irrespective of whether the ports have joined the multicast group. The effect of this option is to disable IGMP snooping without disabling IGMP. (See *Chapter 11, IP Multicasting*). If B is specified, the traffic is sent to ports that have joined the multicast group unless no ports have joined, in which case the traffic is flooded on all ports on the VLAN. If C is specified, the traffic is sent to ports that have joined the multicast group; if no ports have joined, the traffic is discarded. This option allows the manager more control over who receives traffic. The default is B.

The MODE parameter applies to gigabit copper interfaces only. It forces the interface to operate in master or slave mode by setting it to MASTER or SLAVE. This is not typically required and should be used when the link partner does not support autonegotiation of master/slave mode. The default is AUTONEGOTIATE.

The SPEED parameter specifies the configured line speed and duplex mode of the port(s) (Table 3-24 on page 3-112.) If AUTONEGOTIATE is specified, the port autonegotiate the highest mutually possible line speed and duplex mode with the link partner. If 10MFAUTO, 10MHAUTO, 100MFAUTO, 100MHAUTO, 1000MFAUTO, or 1000MHAUTO is specified, the port autonegotiates with the link partner and accepts operation at the specified speed and duplex mode. If 10MHALF, 10MFULL, 100MHALF, 100MFULL, 1000MHALF, or 1000MFULL is specified, then autonegotiation is disabled and the interface must operate at the specified speed and duplex mode regardless of whether the link partner is capable of working at that speed. When a port is included in a trunk group, it must operate at the speed specified for the trunk group and in full duplex mode. This speed is selected by autonegotiation with the link partner. If the port is removed from the trunk group, the previously configured speed and duplex mode are restored. The default is AUTONEGOTIATE. Gigabit fibre ports can operate at 1000Mbit/s full duplex, and gigabit copper ports on some units can only operate at 1000MBit/s half or full duplex.

Table 3-24: SWITCH PORT SPEED values.

Value	Meaning
10MHALF	10 Mbps, half duplex, fixed
10MFULL	10 Mbps, full duplex, fixed
10MHAUTO	10 Mbps, half duplex, autonegotiate
10MFAUTO	10 Mbps, full duplex, autonegotiate
100MHALF	100 Mbps, half duplex, fixed
100MFULL	10 Mbps, full duplex, fixed
100MHAUTO	100 Mbps, half duplex, autonegotiate
100MFAUTO	10 Mbps, full duplex, autonegotiate

Table 3-24: SWITCH PORT SPEED values. (Continued)

Value	Meaning
1000MHALF	1000 Mbps, half duplex, fixed
1000MFULL	1000 Mbps, full duplex, fixed
1000MHAUTO	1000 Mbps, half duplex, autonegotiate
1000MFAUTO	1000 Mbps, full duplex, autonegotiate

If you override a port's autonegotiation by setting it to a fixed speed/duplex setting, automatic MDI/MDI-X detection is also overridden. The port defaults to MDI-X.

Examples

To set the speed of port 5 to 10Mbps, half duplex, use the command:

```
SET SWITCH PORT=5 SPEED=10MHALF
```

To limit the rate of destination lookup failure packets to 1000 packets per second for the processing block of ports 17-24, use the command:

```
SET SWITCH PORT=17-24 DLFLIMIT=1000
```

To accept only VLAN-tagged frames on port 2, use the command:

```
SET SWITCH PORT=2 ACCEPTABLE=VLAN
```

To set the maximum bandwidth for port 1to 512Kbps, use the command:

```
SET SWITCH PORT=1 MAXBANDWIDTH=512
```

Related Commands

DISABLE SWITCH PORT
ENABLE SWITCH PORT
SHOW SWITCH PORT

SET SWITCH QOS

Syntax

`SET SWITCH QOS=`*`P0,P1,P2,P3,P4,P5,P6,P7`*

where:

- *P0-P7* are each numbers in the range 0-n where n+1 is the number of Quality of Service egress queues supported.

Description

This command maps user priority levels to Quality of Service egress queues.

This command also updates the Quality of Service module Hardware Priority settings (see the SET QOS HWPRIORITY command on page 7-29 and the SHOW QOS HWPRIORITY command on page 7-35 in *Chapter 7, Quality of Service (QoS)*.

The QOS parameter specifies a comma-separated list of eight values, all of which must be present. The first value, P0, represents the QOS queue for priority level 0. The last value, P7, represents the QOS queue for priority level 7. Similarly, values P1 to P6 represent the QOS queue for the corresponding priority level.

The switch has four QOS egress queues. Its default QOS values are 1,0,0,1,2,2,3,3 as shown in Table 3-25 on page 3-114.

Packets that originate on the switch or are routed by the switch's software have been assigned a Quality of Service priority of 7. To ensure that these packets are transmitted promptly, you should not assign priority 7 to a low-numbered egress queue.

Table 3-25: Default priority level to queue mapping for four QOS egress queues .

Priority level	Queue
0	1
1	0
2	0
3	1
4	2
5	2
6	3
7	3

Example To set the mapping shown in Table 3-26 on page 3-114, use the command:

```
SET SWITCH QOS=0,0,0,1,1,2,2,3
```

Table 3-26: Example priority level to QOS egress queue mapping .

Priority level	Queue
0	0
1	0
2	0
3	1
4	1
5	2
6	2
7	3

Related Commands SHOW SWITCH QOS

SET SWITCH TRUNK

Syntax

```
SET SWITCH TRUNK=trunk [SELECT={MACSRC|MACDEST|MACBOTH|
   IPSRC|IPDEST|IPBOTH}] [SPEED={10M|100M|1000M}]
```

where *trunk* is a character string 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-).

Description

This command sets parameters for the specified trunk group on the switch.

The TRUNK parameter specifies the name of the trunk group. The name is not case sensitive. The name uniquely identifies the trunk group. The specified trunk group must already exist.

The SELECT parameter specifies the port selection criterion for the trunk group. Each packet to be sent on the trunk group is checked by using the selection criterion, and a port in the trunk group is chosen to send the packet. If MACSRC is specified, the source MAC address is used. If MACDEST is specified, the destination MAC address is used. If MACBOTH is specified, both source and destination MAC addresses are used. If IPSRC is specified, the source IP address is used. If IPDEST is specified, the destination IP address is used. If IPBOTH is specified, both the source and destination IP addresses are used. The user of the switch should choose the value of this parameter to try to spread the load as evenly as possible on the trunk group. The default for this parameter is MACBOTH.

The SPEED parameter specifies the speed of the ports in the trunk group. For gigabit fibre ports, only the 1000M value is allowed. For gigabit copper ports, 10M, 100M, and 1000M values are allowed except that the uplink bays of some units are not 10/100M capable. For 10/100 switch ports, 10M and 100M values are allowed. The default is 100M. When a port is added to a trunk group, its current speed and duplex mode settings are ignored and the port uses the speed of the trunk group and full duplex mode. The ports that are members of the trunk group are constrained to autonegotiate to the trunk speed only.

Example

To set the speed of a trunk group called Trunk1 to 100 Mbps, use the command:

```
SET SWITCH TRUNK=Trunk1 SPEED=100M
```

Related Commands

ADD SWITCH TRUNK
CREATE SWITCH TRUNK
DELETE SWITCH TRUNK
DESTROY SWITCH TRUNK
SHOW SWITCH TRUNK

SET VLAN PORT

Syntax

```
SET VLAN={vlan-name|1..4094} PORT={port-list|ALL}
   FRAME={UNTAGGED|TAGGED}
```

where:

- *vlan-name* is a unique name for the VLAN 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *vlan-name* cannot be a number or ALL.
- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.

Description

This command changes the status of ports in a VLAN from tagged to untagged or vice-versa.

The VLAN parameter specifies the name of the VLAN or the numerical VLAN Identifier of the VLAN. The name is not case sensitive, although the case is preserved for display purposes. The VLAN specified must exist.

The PORT parameter specifies the port or ports to be changed. The ports must belong to the VLAN specified. When the command succeeds on a subset of the specified ports but causes errors on the others, then the command as a whole fails and has no effect. If ALL is specified, then all ports in the VLAN change.

The FRAME parameter specifies whether packets transmitted from a port for the specified VLAN include a VLAN tag header. If FRAME is set to UNTAGGED, the port becomes an untagged port for the specified VLAN, and the ACCEPTABLE switch parameter for the port is set to ALL. The user can then change the ACCEPTABLE parameter for the port. FRAME may only be set to UNTAGGED when the port was previously a tagged port in the same VLAN, and is not an UNTAGGED port of another VLAN. If FRAME is set to TAGGED, then the port becomes a tagged port for the VLAN and the ACCEPTABLE switch parameter for the port is set to VLAN. The user cannot change the ACCEPTABLE parameter for the tagged port. FRAME can be set to TAGGED only when the ports were previously untagged ports in the same VLAN.

Example

To change the status of port 1 of the default VLAN from untagged to tagged, use the command:

```
SET VLAN=DEFAULT PORT=1 FRAME=TAGGED
```

Related Commands

ADD VLAN PORT
DELETE VLAN PORT
SHOW VLAN

SHOW STP

Syntax `SHOW STP[={`*`stp-name`*`|ALL}] [SUMMARY]`

where *stp-name* is a character string 1 to 15 characters long.Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *stp-name* cannot be ALL.

Description This command displays information about the specified Spanning Tree Protocol instance (STP), or all STPs (Figure 3-10 on page 3-118, Table 3-27 on page 3-119).

If the SUMMARY parameter is specified, then a summary table of all configured STPs is displayed (Figure 3-11 on page 3-120, Figure 3-28 on page 3-120).

Figure 3-10: Example output from the SHOW STP command.

```
 STP Information
-------------------------------------------------------------
 Name ................. grey
 Mode ................. Rapid
 RSTP Type ............ Normal
 VLAN members ......... vlan4 (4)
 Status ............... ON
 Number of Ports ...... 2
   Number Enabled ..... 2
   Number Disabled ..... 0
 Bridge Identifier ..... 32768 : 00-00-cd-05-19-28
 Bridge Priority ....... 32768
 Root Bridge ........... 32768 : 00-00-cd-05-19-28
 Designated Bridge ..... 32768 : 00-00-cd-05-19-28
 Root Port ............ (n/a)
 Root Path Cost ........ 0
 Max Age ............... 20
 Hello Time ............ 2
 Forward Delay ......... 15
 Switch Max Age ........ 20
 Switch Hello Time ..... 2
 Switch Forward Delay .. 15
 Transmission Limit .... 3

 Name ................. default
 Mode ................. Standard
 RSTP Type ............ (n/a)
 VLAN members ......... default (1)
                        vlan5 (5)
                        vlan6 (6)
                        vlan7 (7)
                        vlan8 (8)
                        vlan9 (9)
                        vlan10 (10)
                        vlan11 (11)
                        vlan12 (12)
                        vlan13 (13)
                        vlan14 (14)
 Status ............... OFF
 Number of Ports ...... 22
   Number Enabled ..... 0
   Number Disabled ..... 22
 Bridge Identifier ..... 32768 : 00-00-cd-05-19-28
 Bridge Priority ....... 32768
 Designated Root ....... 32768 : 00-00-cd-05-19-28
 Root Port ............ (n/a)
 Root Path Cost ........ 0
 Max Age ............... 20
 Hello Time ............ 2
 Forward Delay ......... 15
 Switch Max Age ........ 20
 Switch Hello Time ..... 2
 Switch Forward Delay .. 15
 Hold Time ............. 1
-------------------------------------------------------------
```

Table 3-27: Parameters in the output of the SHOW STP command.

Parameter	Meaning
STP Name	The name of the Spanning Tree Protocol entity.
Mode	Whether STP is running in standard, or rapid mode.
RSTP Type	Whether RSTP is operating normally, or as STP compatible. In STP compatible mode, the rapid transitions to forwarding do not occur.
VLAN members	A list of the VLANs that are members of the STP. VLAN Identifiers are shown in brackets.
Status	The status of the STP; either ON or OFF.
Number of Ports	The number of ports belonging to the STP.
Number Enabled	The number of ports that have been enabled using the ENABLE STP command and are being considered by the Spanning Tree Algorithm.
Number Disabled	The number of ports that have been disabled using the DISABLE STP command and are not being considered by the Spanning Tree Algorithm.
Bridge Identifier	The unique Bridge Identifier of the switch. This parameter consists of two parts, one is derived from the unique Switch Address, and the other is the priority of the switch.
Bridge Priority	The settable priority component that permits the relative priority of bridges to be managed. The range of values is between 0 and 65535. A lower number indicates a higher priority.
Designated Root	The unique Bridge Identifier of the bridge assumed to be the root, (Standard Mode only).
Root Bridge	The unique Bridge Identifier of the bridge assumed to be the Root, (Rapid Mode only).
Designated Bridge	The unique Bridge Identifier of the bridge assumed to be the designated bridge. Displayed when STP is set to RAPID mode, (Rapid Mode only).
Root Port	The port number of the root port for the switch. If the switch is the Root Bridge this parameter is not valid, and (n/a) is shown.
Root Path Cost	The cost of the path to the Root from this switch. If the switch is the Root Bridge this parameter is not valid and is not shown.
Max Age	The maximum age of received Configuration Message information before it is discarded.
Hello Time	The time interval between successive transmissions of the Configuration Message information by a switch that is the Root or is trying to become the Root.
Forward Delay	In STP Standard mode, the time ports spend in the Listening state before moving to the Learning state and the Learning state before moving to the Forwarding state. In Rapid mode, the maximum time taken to transition from Discarding to Learning and Learning to Forwarding. In both modes, the value is also used for the ageing timer for the dynamic entries in the Forwarding Database.

Table 3-27: Parameters in the output of the SHOW STP command. (Continued)

Parameter	Meaning
Switch Max Age	The value of the Max Age parameter when this switch is the Root or is attempting to become the Root. This parameter is set by the MAXAGE parameter in the SET STP command.
Switch Hello Time	The value of the Hello Time parameter when this switch is the Root or is attempting to become the Root. This parameter is set by the HELLOTIME parameter in the SET STP command.
Switch Forward Delay	The value of the Forward Delay parameter when this switch is the Root or is attempting to become the Root. This parameter is set by the FORWARDDELAY parameter in the SET STP command.
Hold Time	The minimum time in seconds between the transmission of configuration BPDUs through a given LAN Port. The value of this fixed parameter is 1, as specified in IEEE 802.1d. This parameter applies only to STP running in standard mode.
Transmission Limit	In Rapid mode, this indicates the number of BPDUs that may be transmitted in the interval specified by Hello Time. The value of this fixed parameter is 3, as specified in IEEE 802.1t.

Figure 3-11: Example output from the SHOW STP SUMMARY command

```
STP Name  Mode  Ports Enabled  Ports Disabled   Bridge Role
-------------------------------------------------------------
Rstp1      Rapid       0           2            Root Bridge
Default    Standard    0          21            Root Bridge
-------------------------------------------------------------
```

Table 3-28: Parameters displayed in the output of the SHOW STP SUMMARY command.

Parameter	Meaning
STP name	The name of the Spanning Tree Protocol entry.
Mode	Whether STP is running in standard or rapid mode.
Ports Enabled	The number of ports that are being considered by the Spanning Tree Algorithm.
Ports Disabled	The number of ports that have been disabled and are not active in the Spanning Tree Algorithm.
Bridge Role	The role of the bridge in the STP, either None, Designated, or Root.

Example To show the current settings of the company STP, use the command:

```
SHOW STP=company
```

Related Commands

CREATE STP
DESTROY STP
DISABLE STP
ENABLE STP
SHOW STP COUNTER
SHOW STP PORT
SET STP

SHOW STP COUNTER

Syntax

```
SHOW STP[={stp-name|ALL}] COUNTER
```

where *stp-name* is a character string 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *stp-name* cannot be ALL.

Description

This command displays Spanning Tree Protocol counters for the specified STP or all STPs (Figure 3-12 on page 3-121, Table 3-28 on page 3-120). If no STP is specified, then counters for all STPs are displayed. If the port link status is DOWN, then no STP BPDUs are transmitted on the port.

Figure 3-12: Example output from the SHOW STP COUNTER command.

```
STP Counters
------------------------------------------------------------
STP Name: default
  Receive:                      Transmit:
  Total STP Packets      0     Total STP Packets          1677
  Configuration BPDU     0      Configuration BPDU          0
  TCN BPDU               0      TCN BPDU                    0
  RST BPDU               0      RSTP BPDU                1677
  Invalid BPDU           0

 Discarded:
  Port Disabled          0
  Invalid Protocol       0
  Invalid Type           0
  Invalid Message Age    0
  Config BPDU length     0
  TCN BPDU length        0
  RST BPDU length        0
------------------------------------------------------------
```

Table 3-29: Parameters in the output of the SHOW STP COUNTER command .

Parameter	Meaning
STP Name	The name of the STP.
Receive	STP packets received.
Total STP Packets	The total number of STP packets received. Valid STP packets comprise Configuration BPDUs and Topology Change Notification (TCN) BPDUs.

Table 3-29: Parameters in the output of the SHOW STP COUNTER command (Continued).

Parameter	Meaning
Configuration BPDU	The number of valid Configuration BPDUs received.
TCN BPDU	The number of valid Topology Change Notification BPDUs received.
RST BPDU	The number of valid Rapid Spanning Tree BPDUs received (RAPID mode only).
Invalid BPDU	The number of invalid STP packets received.
Transmit	STP packets transmitted.
Total STP packets	The total number of STP packets transmitted.
Configuration BPDU	The number of Configuration BPDUs transmitted.
TCN BPDU	The number of Topology Change Notification BPDUs transmitted.
RST BPDU	The number of valid Rapid Spanning Tree BPDUs transmitted (RAPID mode only).
Discarded	STP packets discarded.
Port Disabled	The number of BPDUs discarded because the port that the BPDU was received on was disabled.
Invalid Protocol	The number of STP packets that had an invalid Protocol Identifier field or invalid Protocol Version Identifier field.
Invalid Type	The number of STP packets that had an invalid Type field.
Invalid Message Age	The number of STP packets that had an invalid message age.
Config BPDU length	The number of Configuration BPDUs that had an incorrect length.
TCN BPDU length	The number of Topology Change Notification BPDUs that had an incorrect length.
RST BPDU length	The number of Rapid Spanning Tree BPDUs that had an incorrect length (RAPID mode only).

Example To show the counters for all STPs, use the command:

```
SHOW STP COUNTER
```

Related Commands
RESET STP
SHOW STP
SHOW STP PORT

SHOW STP DEBUG

Syntax `SHOW STP[={`*`stp-name`*`|ALL}] DEBUG`

where:

- *stp-name* is a character string 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-).

Description This command shows the debugging modes enabled on each port (Figure 3-13 on page 3-123, Table 3-30 on page 3-123).

An STP name can be specified. If no parameter is entered, then the default is ALL.

Figure 3-13: Example output from the SHOW STP DEBUG command.

```
STP Name
         Port   Enabled Debug Modes        Output         Timeout
-------------------------------------------------------------------------
default
         Port1  MSG, PKT, STATE           Console (16)   NONE
         Port2  STATE                     Console (16)   12345
         Port3  None
-------------------------------------------------------------------------
Admin
         Port1  MSG, PKT, STATE           TTY (12)       100
-------------------------------------------------------------------------
```

Table 3-30: Parameters displayed in the output of the SHOW STP DEBUG command .

Parameter	Meaning
Port	The port number on the switch.
Enabled Debug Modes	The debugging option for the port; either "MSG", "PKT", "STATE", or "NONE".
Output	The output device for the port.
Timeout	The time in seconds that the port stays in debug mode. If a timeout value is not set, "None" is shown.
STP name	Name of the STP instance.

Example To display the debug status for all ports in the switch, use the command:

```
SHOW STP DEBUG
```

To show STP on just the ADMIN network, use the command:

```
SHOW STP=ADMIN DEBUG
```

Related Commands
DISABLE STP DEBUG
ENABLE STP DEBUG
SHOW STP COUNTER

SHOW STP PORT

Syntax `SHOW STP[={stp-name|ALL}] PORT={port-list|ALL}`

where:

- *stp-name* is a character string 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen(-).
- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.

Description This command displays Spanning Tree Protocol port information for the specified ports, or all ports for the specified STP, or all STPs, (Figure 3-14 on page 3-125, Table 3-31 on page 3-126)

The STP parameter specifies an STP name. If no parameter is entered, the default is ALL.

Figure 3-14: Example output from the SHOW STP PORT command

```
STP Port Information
------------------------------------------------------------
 STP ..................... grey
   STP Status ............ ON
   Port .................. 3
     RSTP Port Role ...... Disabled
     State ............... Discarding
     Point To Point ...... No (Auto)
     Port Priority ....... 128
     Port Identifier ..... 8003
     Pathcost ............ 200000
     Designated Root ..... 32768 : 00-00-cd-05-19-28
     Designated Cost ..... 0
     Designated Bridge ... 32768 : 00-00-cd-05-19-28
     Designated Port ..... 8003
     EdgePort ............ No
     VLAN membership ..... 1

   Port .................. 4
     RSTP Port Role ...... Disabled
     State ............... Discarding
     Point To Point ...... No (Auto)
     Port Priority ....... 128
     Port Identifier ..... 8004
     Pathcost ............ 200000
     Designated Root ..... 32768 : 00-00-cd-05-19-28
     Designated Cost ..... 0
     Designated Bridge ... 32768 : 00-00-cd-05-19-28
     Designated Port ..... 8004
     EdgePort ............ No

STP ..................... default
   STP Status ............ OFF
   Port .................. 1
     State ............... Disabled
     Port Priority ....... 128
     Port Identifier ..... 8001
     Pathcost ............ 19
     Designated Root ..... 32768 : 00-00-cd-05-19-28
     Designated Cost ..... 0
     Designated Bridge ... 32768 : 00-00-cd-05-19-28
     Designated Port ..... 8001
```

Table 3-31: Parameters displayed in the output of the SHOW STP PORT command .

Parameter	Meaning
STP	The name of the STP that the port is a member of.
STP Status	Whether this STP is enabled or disabled; either ON or OFF.
Port	The number of the port.
RSTP Port Role	The role of the port; either Disabled, Alternate, Backup, Designated, or Root. (Rapid Mode only).
State	The state of the port; either "Disabled", "Blocking", "Listening", "Learning" or "Forwarding" for Standard mode, and either; "Disabled", "Discarding", "Learning", or "Forwarding" for Rapid mode.
Point To Point	Whether the port has a point to point connection with another bridge; either NO or YES. (Rapid Mode only).
Port Priority	The priority of the port. Used as part of the Port Identifier field. In Standard mode it forms the upper 8 bits of the Port Identifier field. In Rapid mode it forms the upper 4 bits of the Port Identifier field.
Port Identifier	The unique identifier of the port. This parameter determines the root port or designated port of the switch.
Pathcost	The path cost of the port.
Designated Root	The unique Bridge Identifier of the Root Bridge, as recorded in the configuration BPDU.
Designated Cost	The Designated Cost for the port.
Designated Bridge	Either the unique Bridge Identifier of the switch, or the unique Bridge Identifier of the switch believed to be the Designated Bridge for the LAN to which the port is attached.
Designated Port	The Port Identifier of the port on the Designated Bridge through which the Designated Bridge transmits Configuration BPDU information stored by this port.
Edge Port	An edge port is a port that attaches to a LAN that is known to have no other bridges attached; either YES, or NO.
VLAN membership	The number of VLANs the port is a member of within this STP instance.

Example To show STP information for port 2 on the STP named 'grey', use the command:

```
SHOW STP=grey PORT=2
```

Related Commands

DISABLE STP PORT
ENABLE STP PORT
SET STP PORT
SHOW STP

SHOW SWITCH

Syntax `SHOW SWITCH`

Description This command displays configuration information for the switch functions (Figure 3-15 on page 3-127, Table 3-32 on page 3-127).

Figure 3-15: Example output from the SHOW SWITCH command.

```
Switch Configuration
------------------------------------------------------------
Switch Address ............. 00-00-cd-04-e0-75
Learning ................... ON
Ageing Timer ............... ON
Number of Fixed Ports ...... 24
Number of Uplink Ports ..... 0
Mirroring .................. DISABLED
Mirror port ................ None
Ports mirroring on Rx ...... None
Ports mirroring on Tx ...... None
Ports mirroring on Both .... None
Number of WAN Interfaces ... 0
Name of Interface(s) ....... -
Ageingtime ................. 300
L3 Ageingtime .............. 900
UpTime ..................... 00:04:30
------------------------------------------------------------
```

Table 3-32: Parameters displayed in the output of the SHOW SWITCH command .

Parameter	Meaning
Switch Address	The MAC address of the switch, from which the Bridge Identifier used in the Spanning Tree Algorithm is derived.
Learning	Whether the switch's dynamic learning and updating of the Forwarding Database is enabled; either "ON" or "OFF".
Ageing Timer	Whether the ageing timer is enabled; either "ON" or "OFF".
Number of Fixed Ports	The number of fixed Ethernet switch ports.
Number of Uplink Ports	The number of Ethernet uplink ports.
Mirroring	The state of traffic mirroring, either "Enabled" or "Disabled".
Mirror port	The switch port where mirror traffic is sent.
Ports mirroring on Rx	The ports that are set to send all the traffic they receive to the mirror port.
Ports mirroring on Tx	The ports that are set to send all the traffic they transmit to the mirror port.
Ports mirroring on Both	The ports that are set to send all the traffic they both receive and transmit to the mirror port.
Number of WAN Interfaces	The total number of installed WAN interfaces.
Name of Interface(s)	The name of the installed WAN interface(s).

Table 3-32: Parameters displayed in the output of the SHOW SWITCH command (Continued).

Parameter	Meaning
Ageingtime	The value in seconds of the ageing timer, after which a dynamic entry is removed from the Forwarding Database.
L3 Ageingtime	The value in seconds of the Layer 3 ageing timer, after which a dynamic entry is removed from the Layer 3 Forwarding Database.
Uptime	The time in hours:minutes:seconds since the SWITCH was last powered up, rebooted, or restarted. This is the same as the value of the MIB object sysUpTime.
Uptime	The time in hours:minutes:seconds since the SWITCH was last powered up, rebooted, or restarted. This is the same as the value of the MIB object sysUpTime.

Example To display the configuration of the switch module, use the command:

```
SHOW SWITCH
```

Related Commands RESET SWITCH

SHOW SWITCH COUNTER

Syntax `SHOW SWITCH COUNTER`

Description This command displays information about the forwarding counters associated with the switch (Figure 3-16 on page 3-128, Table 3-33 on page 3-129).

To display reception and transmission packet counters for the switch, see the SHOW SWITCH PORT COUNTER command on page 3-143.

Figure 3-16: Example output from the SHOW SWITCH COUNTER command.

```
Switch Counters
----------------------------------------------------------------
Packet DMA counters

  Receive:                          Transmit:
Packets                  407     Packets                  708
Discards                   0     Discards                   0
TooFewBuffers              0     Aborts                     0
DescriptorsExhausteds      0     DescriptorAreaFilleds      0
QueueLength                0     QueueLength                0

  PCI bus counters:
ParityErrors               0     ErrorChannel               0
FatalErrors                0

  General counters:
Resets                     0
----------------------------------------------------------------
```

Table 3-33: Parameters in the output of the SHOW SWITCH COUNTER command .

Parameters	Meaning
Packet DMA counters	
Receive	Counters for packets received.
Packets	The number of packets received by the CPU from the switch chip.
Discards	The number of packets received from the switch chip that were discarded because either the receive queue was greater than 4096, or because the free buffers in the switch were below BufferLevel3, or because there were no data bytes in the packet.
TooFewBuffers	The number of packets received from the switch chip that were discarded because the free buffers in the switch were below BufferLevel3.
DescriptorsExhausteds	The number of times the switch chip reported that it could not transfer a packet by DMA to a switch buffer because there were no more receive buffer descriptors.
QueueLength	The number of packets received from the switch chip waiting to be processed by the CPU.
Transmit	Counters for packets transmitted.
Packets	The number of packets transferred from the CPU to the switch chip.
Discards	The number of packets waiting for transmission that were discarded when the DMA process was reset due to an error.
Aborts	The number of times transmission of a packet was aborted due to it taking an excessive length of time for the transmission to complete, perhaps due to a port being in a blocked state or due to a busy PCI bus.
DescriptorAreaFilleds	The number of times the transmit descriptor area filled due to a high rate of transfer of packets from the CPU to the switch chip or high PCI bus utilisation causing the DMA to proceed slowly.
QueueLength	The number of packets currently queued for transmission, or that have been transmitted and are waiting to be purged from the transmit queue.
PCI bus counters	
ParityErrors	The number of times the switch chip reported a parity error for a transaction on the PCI bus.
FatalErrors	The number of times the switch chip reported a fatal error for a transaction on the PCI bus.
ErrorChannel	The DMA channel for making the transaction for which the error occurred.
General counters	
Resets	The number of times the receive and transmit DMA channels have been reset due to the occurrence of an error.

Example To display the switching counters, use the command:

```
SHOW SWITCH COUNTER
```

Related Commands RESET SWITCH
SHOW SWITCH
SHOW SWITCH PORT COUNTER

SHOW SWITCH DEBUG

Syntax `SHOW SWITCH DEBUG`

Description This command displays debugging information for the switch (Figure 3-17 on page 3-130, Table 3-34 on page 3-130).

Figure 3-17: Example output from the SHOW SWITCH DEBUG command.

```
Enabled Switch Debug Modes    Output     Timeout
-----------------------------------------------------------
ARL, DMA                      16         12345
-----------------------------------------------------------
```

Table 3-34: Parameters in the output of the SHOW SWITCH DEBUG command.

Parameter	Meaning
Enabled Switch Debug Modes	The debugging option for the switch; either "ARL", "CMIC", "DMA", "QOS", "S5600", "PHY", or "None".
Output	The output device for the switch. This is shown when a debug mode is enabled.
Timeout	The time in seconds that debugging options for the switch are enabled. This is shown when a debug mode is enabled.

Example To display debugging information for the switch, use the command:

```
SHOW SWITCH DEBUG
```

Related Commands DISABLE SWITCH DEBUG
ENABLE SWITCH DEBUG

SHOW SWITCH FDB

Syntax

```
SHOW SWITCH FDB[={SW|HW}] [ADDRESS=macadd]
   [DISCARD={SOURCE|DESTINATION}] [HIT={YES|NO}] [L3={YES|
   NO}] [PORT={port-list|ALL}] [STATUS={STATIC|DYNAMIC}]
   [VLAN={vlan-name|1..4094}]
```

where:

- *macadd* is an Ethernet six-octet MAC address, expressed as six pairs of hexadecimal digits delimited by hyphens.
- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.
- *vlan-name* is a unique name for the VLAN 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *vlan-name* cannot be a number or ALL.

Description

This command displays the contents of the Forwarding Database (Figure 3-18 on page 3-132, Table 3-35 on page 3-133).

The FDB parameter specifies the version of the Forwarding Database that is displayed. The Forwarding Database is stored in hardware and a copy is held in software. If SW is specified, the software copy of the Forwarding Database is displayed; if HW is specified, the hardware version is displayed. Under normal circumstances, the two versions are identical. The default is SW.

The ADDRESS parameter specifies the MAC address of the device for which the contents of the Forwarding Database are to be displayed.

The DISCARD parameter specifies whether to display entries in the Forwarding Database where frames are discarded on the basis of the received frame's source or destination address.

The HIT parameter specifies whether to display filter entries in the Forwarding Database where a frame matching the entry either was or was not received during the latest Ageing Timer period.

The L3 parameter specifies whether to display filter entries in the Forwarding Database that were or were not created as part of a Layer 3 interface configuration.

The PORT parameter specifies that only those entries in the Forwarding Database that were learned from the specified port are to be displayed.

The STATUS parameter specifies whether to display only static filter entries or only dynamically-learned filter entries.

The VLAN parameter specifies the VLAN identifier of the VLAN for which the contents of the Forwarding Database are to be displayed.

Figure 3-18: Example output from the SHOW SWITCH FDB command.

```
Switch Forwarding Database (software)
-------------------------------------------------------------------------
 VLAN  MAC Address          Port  Status    Discard  L3  Hit  QOS  QSD
-------------------------------------------------------------------------
 1     00-00-cd-00-45-c7    CPU   static    -        y   y    0:0  dest
 42    00-00-c0-1d-2c-f8    1     dynamic   -        n   y    0:0  dest
 42    00-00-c0-71-e0-e4    1     dynamic   -        n   y    0:0  dest
 42    00-00-cd-00-a4-d6    1     dynamic   -        n   y    0:0  dest
 42    00-00-cd-00-ab-dc    1     dynamic   -        n   y    0:0  dest
 42    00-60-b0-ac-18-51    1     dynamic   -        n   y    0:0  dest
 42    00-90-27-23-a4-e9    1     dynamic   -        n   y    0:0  dest
 42    00-90-27-32-ad-61    1     dynamic   -        n   y    0:0  dest
 42    00-90-27-76-8a-55    1     dynamic   -        n   y    0:0  dest
 42    00-90-27-76-9a-99    1     dynamic   -        n   y    0:0  dest
 42    00-90-27-87-a5-22    1     dynamic   -        n   y    0:0  dest
 42    00-90-27-bd-c8-93    1     dynamic   -        n   y    0:0  dest
 42    00-90-27-bd-c9-7f    1     dynamic   -        n   y    0:0  dest
 42    00-90-27-d0-ae-c2    1     dynamic   -        n   y    0:0  dest
 42    00-90-27-d0-c7-12    1     dynamic   -        n   y    0:0  dest
 42    08-00-09-be-06-cd    1     dynamic   -        n   y    0:0  dest
-------------------------------------------------------------------------
```

Table 3-35: Parameters in the output of the SHOW SWITCH FDB command .

Parameter	Meaning
VLAN	The VLAN Identifier of the VLAN.
MAC Address	The MAC address as learned from the source address field of a frame, or entered as part of a static filter entry.
Port	The port from which the MAC address was learned.
Status	Whether the entry was a static filter entry or dynamically learned; either "dynamic" or "static".
Discard	If frames are to be discarded, discard on the basis of the source address or the destination address of the received frame; either "source" or "destination.".
L3	Whether the entry was created as part of a Layer 3 interface configuration; either "y" (yes) or "n" (no).
Hit	Whether a frame matching this filter entry was received during the latest Ageing Timer period; either "y" (yes) or "n" (no). If the Ageing Timer is enabled, entries with 'n' are purged from the Forwarding Database.
QOS	Quality of Service of the frame. The first number is the QoS based on the source address. The second number is the QoS based on the destination address.
QSD	Whether the source address QoS or the destination address QoS has priority in determining the QoS of frames received that do not contain priority information; either "source" or "dest".

Example To display the contents of the Forwarding Database, use the command:

```
SHOW SWITCH FDB
```

Related Commands ENABLE SWITCH LEARNING
SHOW SWITCH
SHOW SWITCH FILTER

SHOW SWITCH FILTER

Syntax

```
SHOW SWITCH FILTER [PORT={port-list|ALL}]
   [ACTION={FORWARD|DISCARD}] [DESTADDRESS=macadd]
   [ENTRY=entry-list] [VLAN={vlan-name|1..4094}]
```

where:

- *macadd* is an Ethernet six-octet MAC address, expressed as six pairs of hexadecimal digits delimited by hyphens.
- *entry-list* is an entry number, a range of entry numbers (specified as n-m), or a comma separated list of entry numbers and/or ranges. Entry numbers start at 0 and end at m, where m is the highest filter entry currently defined in the Permanent Forwarding Database. Each port has its own Permanent Forwarding Database.
- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.
- *vlan-name* is a unique name for the VLAN 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *vlan-name* cannot be a number or ALL.

Description

This command displays information about some or all of the static switch filter entries (Figure 3-19 on page 3-134, Table 3-36 on page 3-135). The output can be limited to display only entries matching the optional parameters as described below.

The ACTION parameter specifies whether frames matching the filter entry are forwarded or discarded.

The ENTRY parameter must specify an existing filter entry or entries in the Permanent Forwarding Database.

The DESTADDRESS parameter specifies the destination MAC address in the filter entry.

The PORT parameter specifies the outbound ports over which frames matching this filter entry are discarded or forwarded.

The VLAN parameter specifies the numerical VLAN Identifier with which the filter entry is associated.

Figure 3-19: Example output from the SHOW SWITCH FILTER command.

```
 Switch Filters
-------------------------------------------------------------------------
 Entry             VLAN         Destination Address   Port  Action  Source
-------------------------------------------------------------------------
    0           default (1)       aa-ab-cd-00-00-01     1  Forward  static
    1           default (1)       aa-ab-cd-00-00-02     1  Forward  static

    0           marketing (2)     aa-ab-cd-00-00-01     2  Discard  static
    1           marketing (2)     aa-ab-cd-00-00-02     2  Discard  learn
-------------------------------------------------------------------------
```

Table 3-36: Parameters in the output of the SHOW SWITCH FILTER command .

Parameter	Meaning
Entry	The number identifying the filter entry.
Destination Address	The destination MAC address for the entry.
VLAN	The VLAN name and identifier for the entry.
Port	The outbound port to match for the filter entry to be applied.
Action	The action specified by the filter entry; either "Forward" or "Discard".
Source	This parameter is either "static" (indicating the filter is a static filter) or "learned" (indicating the filter is present either because it has been added with the LEARN parameter of the SET SWITCH PORT command, or has been dynamically learned during normal intrusion detection operation).

Examples To display information about the entire Permanent Forwarding Database, use the command:

```
SHOW SWITCH FILTER PORT=ALL
```

To display information about the Permanent Forwarding Database for port 3, use the command:

```
SHOW SWITCH FILTER PORT=3
```

To display information about the Permanent Forwarding Database for the *marketing* VLAN, use the command:

```
SHOW SWITCH FILTER PORT=ALL VLAN=MARKETING
```

To display the port to which the MAC address 00-00-00-12-34-56 belongs, use the command:

```
SHOW SWITCH FILTER PORT=ALL DESTADDRESS=00-00-00-12-34-56
```

Related Commands ADD SWITCH FILTER
DELETE SWITCH FILTER

SHOW SWITCH HWFILTER

Syntax `SHOW SWITCH HWFILTER [CLASSIFIER=`*`classifier-list`*`]`

where:

- *classifier-list* is either an integer in the range 1 to 9999; a range of integers (specified as 0-4) or a comma separated list of classifier numbers and/or ranges (0, 3, 4-9).

Description This command displays hardware-based filtering entries created using the ADD SWITCH HWFILTER CLASSIFIER command on page 3-53 (Figure 3-20 on page 3-136, Figure 3-21 on page 3-136, Table 3-37 on page 3-137). All of the specified classifiers must exist and must already be incorporated into a filter entry. If CLASSIFIER is not specified, summary information is displayed for filters currently defined.

Figure 3-20: Example output from the SHOW SWITCH HWFILTER command.

```
Switch Hardware Filter Summary Information
------------------------------------------------------------
Status ............... ENABLED
Number of Filters .... 12

Filter ............... 1
Classifier ........... 3

Filter ............... 2
Classifier ........... 100

Filter ............... 3
Classifier ........... 101
------------------------------------------------------------
```

Figure 3-21: Example output from the SHOW SWITCH HWFILTER CLASSIFIER=3 command.

```
------------------------------------------------------------
Filter ............... 1
Classifier ........... 3
Action ............... sp
New IP DSCP .......... -
New TOS .............. -
Port ................. -
Priority ............. 5
No Match Action ...... st, sp
No Match DSCP ........ -
No Match TOS ......... 2
No Match Port ........ -
No Match Priority .... 1
------------------------------------------------------------
```

Table 3-37: Parameters displayed in the output of the SHOW SWITCH HWFILTER CLASSIFIER command .

Parameter	Meaning
Status	The current status for hardware filtering on the switch. Either "ENABLED" or "DISABLED"
Number of Filter	The current total of filters created using the ADD SWITCH HWFILTER command.
Filter	The filter number.
Classifier	The number of the classifier this filter entry is based on.
Action	The action to take when a packet matches this entry; one or more of "sp" (SETPRIORITY), "sc" (FORWARD or SENDCOS), "st" (SETTOS), "dn" (DENY), "se" (SENDEPORT), "sm" (SENDMIRROR), "mpt" (MOVEPRIOTOTOS) "mtp" (MOVETOSTOPRIO), "sds" (SETIPDSCP), "sn" (SENDNONUNICASTTOPORT), "nd" (NODROP).
New IP DSCP	The new IP DSCP value to assign to packets matching the entry.
New TOS	The new TOS value to assign to packets matching the entry.
Port	The new output port to use for packets matching the entry.
Priority	The new priority value to assign to packets matching the entry.
No Match Action	The action to take when a packet matches the specified ingress/ egress ports for this entry; one or more of "sp" (SETPRIORITY), "sc" (FORWARD or SENDCOS), "st" (SETTOS), "dn" (DENY), "se" (SENDEPORT), "sm" (SENDMIRROR), "mpt" (MOVEPRIOTOTOS) "mtp" (MOVETOSTOPRIO), "sds" (SETIPDSCP), "sn" (SENDNONUNICASTTOPORT).
No Match DSCP	The new IP DSCP value to assign to packets on a partial match.
No Match TOS	The new TOS value to assign to packets on a partial match.
No Match Port	The new output port to use for packets on a partial match.
No Match Priority	The new priority value to assign to packets on a partial match.

Example To display all filters, use the command:

```
SHOW SWITCH L3FILTER
```

To display entry 3 from filter 1, use the command:

```
SHOW SWITCH L3FILTER CLASSIFIER=1
```

Related Commands ADD SWITCH HWFILTER CLASSIFIER
DELETE SWITCH HWFILTER CLASSIFIER
SET SWITCH HWFILTER CLASSIFIER
SHOW CLASSIFIER in *Chapter 6, Generic Packet Classifier*

SHOW SWITCH L3FILTER

Syntax `SHOW SWITCH L3FILTER[=`*`filter-id`* `[ENTRY=`*`entry-id`*`]]`

where:

- *filter-id* is a decimal number in the range 1 to the number of filters defined.
- *entry-id* is a decimal number in the range 1 to the number of entries defined.

Description This command displays hardware-based Layer 3 filtering match criteria and their filter entries (Figure 3-22 on page 3-138, Table 3-38 on page 3-138).

Figure 3-22: Example output from the SHOW SWITCH L3FILTER command.

```
Filter ............... 1
Matched fields ....... tos, ttl, sipaddr, dipaddr, protocol
Source address mask .. 255.255.255.0
Dest. address mask ... 255.255.255.0
Ingress port mask .... true
Egress port mask ..... true
No match action ...... none

Ent.  S-Address           D-Address      Prot   TTL  TOS  NewTOS  Type
      S-Mask              D-Mask         Iport  Eport     Port    Syn/Ack/Fin
      S-Port              D-Port         Action
-------------------------------------------------------------------------------
1     192.168.1.0         192.168.2.0    ICMP   30   2    1       0
      255.255.255.0       255.255.255.0  2      3                 0/0/0
      -                   -              dn
-------------------------------------------------------------------------------
2     192.168.2.0         192.168.1.0    ICMP   30   2    1       0
      255.255.255.0       255.255.255.0  2      3                 0/0/0
      -                   -              sc
-------------------------------------------------------------------------------
```

Table 3-38: Parameters displayed in the output of the SHOW SWITCH L3FILTER command .

Parameter	Meaning
Filter	The filter number.
Match fields	A list of the fields matched by this filter; one or more of "tos", "ttl", "protocol", "sipaddr", "dipaddr", "tcpsport", "tcpdport", "tcpsyn", "tcpack", "tcpfin", "udpsport", or "udpdport".
Source address mask	The mask to apply to source IP address fields to determine a match.
Destination address mask	The mask to apply to destination IP address fields to determine a match.
Ingress port mask	Whether the filter applies to ingress ports. Either "TRUE" or "FALSE".
Egress port mask	Whether the filter applies to egress ports. Either "TRUE" or "FALSE".

Table 3-38: Parameters displayed in the output of the SHOW SWITCH L3FILTER command (Continued).

Parameter	Meaning
No Match Action	The action to take when a packet matches the specified ingress/egress ports for this entry; one or more of "sp" (SETPRIORITY), "sc" (FORWARD or SENDCOS), "st" (SETTOS), "dn" (DENY), "se" (SENDEPORT), "sm" (SENDMIRROR), "mpt" (MOVEPRIOTOTOS), "mtp" (MOVETOSTOPRIO), "sds" (SETIPDSCP), "sn" (SENDNONUNICASTTOPORT).
Ent.	The filter entry number.
S-Address, S-Mask, S-Port	The source IP address, source mask and source port to match.
D-Address, D-Mask, D-Port	The destination IP address, destination mask and destination port to match.
Prot	The protocol to match.
Iport	The ingress port number to match.
Action	The action to take when a packet matches this entry; either "sp" (SETPRIORITY), "sc" (FORWARD or SENDCOS), "st" (SETTOS), "dn" (DENY), "se" (SENDEPORT), or "sm" (SENDMIRROR).
TTL	The TTL value to match.
Eport	The egress port number to match.
TOS	The TOS value to match.
NewTOS	The new TOS value to assign to packets matching the entry.
Type	The value of the protocol-type to match. If a 5 byte hexadecimal number is shown then the packet type is SNAP, if 2 bytes are shown then the packet type is either Ethernet type II or 802.3 and (E-II) or (SNAP) is appended respectively.
Port	The new output port to use for packets matching the entry.
Priority	The new priority value to assign to packets matching the entry.

Example To display all filters, use the command:

```
SHOW SWITCH L3FILTER
```

To display entry 3 from filter 1, use the command:

```
SHOW SWITCH L3FILTER=1 ENTRY=3
```

Related Commands

ADD SWITCH L3FILTER MATCH
ADD SWITCH L3FILTER ENTRY
DELETE SWITCH L3FILTER
DELETE SWITCH L3FILTER ENTRY
DISABLE SWITCH L3FILTER
ENABLE SWITCH L3FILTER
SET SWITCH L3FILTER MATCH
SET SWITCH L3FILTER ENTRY

SHOW SWITCH PORT

Syntax `SHOW SWITCH PORT[={`*port-list*`|ALL}]`

where:

- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.

Description This command displays general information about the specified switch ports or all switch ports (Figure 3-23 on page 3-140, Table 3-39 on page 3-141).

Figure 3-23: Example output from the SHOW SWITCH PORT command.

```
Switch Port Information
------------------------------------------------------------------------------
 Port ........................ 1
   Description ................. To intranet hub, port 4
   Status ...................... ENABLED
   Link State .................. Up
   UpTime ...................... 00:10:49
   Port Media Type ............. ISO8802-3 CSMACD
   Configured speed/duplex ..... Autonegotiate
   Actual speed/duplex ......... 1000 Mbps, full duplex
   Configured master/slave mode .. Autonegotiate
   Actual master/slave mode ..... Master
   Acceptable Frame Types ...... Admit All Frames
   Broadcast rate limit ........ 1000/s
   Multicast rate limit ........ -
   DLF rate limit .............. -
   Learn limit ................. -
   Intrusion action ............ Discard
   Current learned, lock state ... 15, not locked
   Mirroring ................... Tx, to port 22
   Is this port mirror port ..... No
   Enabled flow control ........ Pause
   Send tagged pkts for VLAN(s) .. marketing (87)
                                   sales (321)
   Port-based VLAN ............. accounting (42)
   Ingress Filtering ........... OFF
   Trunk Group ................. -
   STP ......................... company
   Multicast filtering mode ..... (B) Forward all unregister groups

   GBIC vendor name ............ AGILENT
   GBIC part number ............ HFCT-5611
   GBIC vendor SN .............. 0111131243329572
   GBIC data code .............. 01111300
------------------------------------------------------------------------------
```

Table 3-39: Parameters in the output of the SHOW SWITCH PORT command .

Parameter	Meaning
Port	The number of the switch port.
Description	A description of the port.
Status	The state of the port; either "ENABLED" or "DISABLED".
Link state	The link state of the port; either "Up" or "Down".
Uptime	The count in hours:minutes:seconds of the elapsed time since the port was last reset or initialised.
Port Media Type	The MAC entity type as defined in the MIB object ifType.
Configured speed/duplex	The port speed mode configured for this port. Either "Autonegotiate" or a combination of a speed (one of "10 Mbps", "100 Mbps" or "1000 Mbps") and a duplex mode (one of "half duplex" or "full duplex"), and optionally "(by autonegotiation)".
Actual speed/duplex	The port speed and duplex mode that this port is actually running at. A combination of a speed (either "10 Mbps", "100 Mbps" or "1000 Mbps") and a duplex mode (either "half duplex" or "full duplex").
Configured master/slave mode	The master/slave mode configured for this port; either "Autonegotiate', "Master", "Slave", or "Not applicable".
Actual master/slave mode	The master/slave mode actually selected; either "-", "Master", "Slave", or "Not applicable".
Acceptable Frame Types	The value of the Acceptable Frame Types parameter, either: "Admit All Frames" or "Admit Only VLAN-tagged Frames".
Broadcast rate limit	The limit of the rate of reception of broadcast frames for this port, in frames per second.
Multicast cast rate limit	The limit of the rate of reception of multicast frames for this port, in frames per second.
DLF rate limit	The limit of the rate of reception of DLF (destination lookup failure) frames for this port, in frames per second.
Learn limit	The number of MAC addresses that may be learned for this port. Once the limit is reached, the port is locked against any new MAC addresses. Either "None" or a number from 1 to 256.
Intrusion action	The action taken on this port when a frame is received from an unknown MAC address when the port is locked. Either "Discard", "Trap", or "Disable".
Current learned, lock state	The number of MAC addresses currently learned on this port and the state of locking for this port. The current learned parameter is incremented when a Learn Limit is set for the port. The lock state is either "not locked", "locked by limit", or "locked by command".
Mirroring	The traffic mirroring for traffic in and out of this port. Either "None", "Rx" (for traffic received by this port), "Tx" (for traffic sent on this port), or "Both". The port where mirrored frames are sent is also displayed.
Is this port mirror port	Whether this port is a mirror port. Either "No" or "Yes".

Table 3-39: Parameters in the output of the SHOW SWITCH PORT command (Continued).

Parameter	Meaning
Enabled flow control	Flow control parameters set for the port; "Pause" or "-". If flow control is implemented on the switch, then Pause flow control is applied to the port.
Send tagged pkts for VLAN(s)	The name and VLAN Identifier (VID) of the tagged VLAN(s), if any, to which the port belongs.
Port-based VLAN	The name and VLAN Identifier (VID) of the port-based VLAN to which the port belongs.
Ingress Filtering	The state of Ingress Filtering; either "ON" or "OFF".
Trunk Group	Name of trunk group to which the port belongs, if any.
STP	The name of the STP to which the port belongs.
Multicast filtering mode	Either "(A) forward all groups", "(B) forward all unregistered groups", or "(C) filter all unregistered groups".
GBIC vendor name	The name of the GBIC vendor. This is shown when a valid GBIC is installed in the port.
GBIC part number	The vendor part number or product name.This is shown when a valid GBIC is installed in the port.
GBIC vendor SN	The vendor serial number. This is shown when a valid GBIC is installed in the port.
GBIC data code	The data code of this GBIC. This is shown when a valid GBIC is installed in the port.

Example To display the configuration for switch port 1, use the command:

```
SHOW SWITCH PORT=1
```

Related Commands SET SWITCH PORT

SHOW SWITCH PORT COUNTER

Syntax `SHOW SWITCH PORT[={`*`port-list`*`|ALL}] COUNTER`

where:

- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.

Description This command displays counters for the specified switch ports or all switch ports (Figure 3-24 on page 3-144, Table 3-40 on page 3-145).

Figure 3-24: Example output from the SHOW SWITCH PORT COUNTER command

```
Port 1. Fast Ethernet MAC counters:
 Combined receive/transmit packets by size (octets) counters:
  64                              65 512 - 1023                        0
  65 - 127                         5 1024 - MaxPktSz                   0
  128 - 255                        0 1519 - 1522                       0
  256 - 511                        0

General Counters:
 Receive                           Transmit
  Octets                         246 Octets                         4340
  Pkts                             3 Pkts                             67
  FCSErrors                        0 FCSErrors                         0
  MulticastPkts                    0 MulticastPkts                    65
  BroadcastPkts                    3 BroadcastPkts                     2
  PauseMACCtlFrms                  0 PauseMACCtrlFrm                   0
  OversizePkts                     0 OversizePkts                      0
  Fragments                        0 Fragments                         0
  Jabbers                          0 Jabbers                           0
  MACControlFrms                   0
  UnsupportOpcode                  0
  AlignmentErrors                  0
  OutOfRngeLenFld                  0
  SymErDurCarrier                  0
  CarrierSenseErr                  0
  UndersizePkts                    0
                                     PauseCtrlFrms                     0
                                     FrameWDeferrdTx                   0
                                     FrmWExcesDefer                    0
                                     SingleCollsnFrm                   0
                                     MultCollsnFrm                     0
                                     LateCollsns                       0
                                     ExcessivCollsns                   0
                                     CollisionFrames                   0

Layer 3 Counters:
  ifInUcastPkts                    0 ifOutUcastPkts                    0
  ifInDiscards                     0 ifOutErrors                       0
  ipInHdrErrors                    0

Miscellaneous Counters:
  DropEvents                       0
  ifOutDiscards                    0
  taggedPktTx                      0
  totalPktTxAbort                  0

HW Multicasting Counters:
  TTL expired                      0
  Bridged Frames                   0
  Routed Frames                    0
  Receive Drops                    0
  Transmit Drops                   0
```

Table 3-40: Parameters in output from SHOW SWITCH PORT COUNTER command .

Parameter	Description
Ethernet MAC counters	
Combined receive/transmit packets by size (octets) counters	The number of packets in each size range received and transmitted.
64	Number of 64 octet packets received and transmitted.
65 - 127	Number of 65 - 127 octet packets received and transmitted.
128 - 255	Number of 128 - 255 octet packets received and transmitted.
256 - 511	Number of 256 - 511 octet packets received and transmitted.
512 - 1023	Number of 512 - 1023 octet packets received and transmitted.
1024 - MaxPktSz	Number of packets received and transmitted with size 1024 octets to the maximum packet length.
1519 - 1522	Number of 1519 - 1522 octet frames received and transmitted.
General Counters	
Receive	Counters for traffic received.
Octets	The number of octets.
Pkts	The number of packets.
FCSErrors	The number of frames containing a Frame Check Sequence error.
MulticastPkts	The number of multicast packets.
BroadcastPkts	The number of broadcast packets.
PauseMACCtlFrms	The number of valid PAUSE MAC Control frames.
OversizePkts	The number of oversize packets.
Fragments	The number of fragments.
Jabbers	The number of jabber frames.
MACControlFrms	The number of MAC Control frames (Pause and Unsupported).
UnsupportOpcode	The number of MAC Control frames with unsupported opcode (i.e. not Pause).
AlignmentErrors	The number of frames with alignment errors.
OutOfRngeLenFld	The number of packets with length out of range.
SymErDurCarrier	The number of frames with invalid data symbols.
CarrierSenseErr	The number of false carrier conditions between frames.
UndersizePkts	The number of undersized packets.
Transmit	Counters for traffic transmitted
Octets	The number of octets.
Pkts	The number of packets.
FCSErrors	The number of frames containing a Frame Check Sequence error.

Table 3-40: Parameters in output from SHOW SWITCH PORT COUNTER command (Continued).

Parameter	Description
MulticastPkts	The number of multicast packets.
BroadcastPkts	The number of broadcast packets.
PauseMACCtlFrms	The number of valid PAUSE MAC Control frames.
OversizePkts	The number of oversize packets.
Fragments	The number of fragments.
Jabbers	The number of jabber frames.
PauseCtrlFrms	The number of Pause control frames.
FrameWDeferrdTx	The number of frames deferred once before successful transmission.
FrmWExcesDefer	The number of frame aborted after too many deferrals.
SingleCollsnFrm	The number of frames that experienced exactly one collision.
MultCollsnFrm	The number of frames that experienced 2 to 15 collisions (including late collisions).
LateCollsns	The number of frames that experienced late collisions.
ExcessivCollsns	The number of frames aborted before transmission after 16 collisions.
CollisionFrames	The total number of collisions.
Layer 3 Counters	Counters for Layer 3 switching. (These counters do not include packets sent to CPU for processing.)
ifInUcastPkts	The number of L3 switched unicast packets.
ifInDiscards	The number of packets for Layer 3 interfaces that are discarded.
ipInHdrErrors	The number of packets discarded due to IP header errors.
ifOutUcastPkts	The number of L3 switched unicast packets.
ifOutErrors	The number of L3 switched packets discarded at egress due to transmission errors.
Miscellaneous Counters	
DropEvents	The number of packets discarded at ingress port.
ifOutDiscards	The number of packets for transmission discarded due to ageing.
taggedPktTx	The number of VLAN tagged packets transmitted.
totalPktTxAbort	The number of Layer 2 and 3 packets aborted during transmission.
HW Multicasting Counters	
TTL expired	The number of packets dropped by the router because their IP multicasting Time to Live (TTL) counter was too low.
Bridged Frames	The number of IP multicasting packets received on this port and bridged (L2 switched) out another port.
Routed Frames	The number of IP multicasting packets received on this port and routed (L3 switched) out another port.

Table 3-40: Parameters in output from SHOW SWITCH PORT COUNTER command (Continued).

Parameter	Description
Receive Drops	The number of IP multicasting packets dropped by this port on ingress.
Transmit Drops	The number of IP multicasting packets dropped by this port on egress.

Example To display counters for switch port 1, use the command:

```
SHOW SWITCH PORT=1 COUNTER
```

Related Commands SET SWITCH PORT
SHOW SWITCH COUNTER
SHOW SWITCH PORT

SHOW SWITCH PORT INTRUSION

Syntax `SHOW SWITCH PORT={`*`port-list`*`|ALL} INTRUSION`

where:

- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered Ethernet switch port, including uplink ports.

Description This command shows a list of MAC addresses for devices that are active on a port, but which are not valid devices allowed or learned for the port. The list contains entries when the INTRUSIONACTION parameter (SET SWITCH PORT command) is of the type TRAP (Figure 3-25 on page 3-147).

Figure 3-25: Example output from the SHOW SWITCH PORT INTRUSION command.

```
Switch Port Information
------------------------------------------------------------------------
Port 2 -    13 intrusion(s) detected
    00-00-c0-1d-2c-f8  00-90-27-87-a5-22  00-00-cd-01-00-4a
    00-d0-b7-4d-93-c0  08-00-5a-a1-02-3f  00-d0-b7-d5-5f-a9
    00-b0-d0-20-d1-01  00-90-99-0a-00-49  00-10-83-05-72-83
    00-00-cd-00-45-9e  00-00-c0-ad-a3-d0  00-a0-24-8e-65-3c
    00-90-27-32-ad-61
------------------------------------------------------------------------
```

Example To display a list of MAC addresses for devices active on port 2, but which are not valid devices, use the command:

```
SHOW SWITCH PORT=2 INTRUSION
```

Related Commands SET SWITCH PORT

SHOW SWITCH QOS

Syntax `SHOW SWITCH QOS`

Description This command displays the current mapping of user priority level to QOS egress queue for the switch (Figure 3-26 on page 3-148, Table 3-41 on page 3-148).

Packets that originate on the switch or are routed by the switch's software have been assigned a Quality of Service priority of 7. To ensure that these packets are transmitted promptly, you should not assign priority 7 to a low-numbered egress queue.

Figure 3-26: Example output from the SHOW SWITCH QOS command

```
Priority Level       QOS egress queue
-------------------------------------
0 ................... 1
1 ................... 0
2 ................... 0
3 ................... 1
4 ................... 2
5 ................... 2
6 ................... 3
7 ................... 3
```

Table 3-41: Parameters displayed in the output of the SHOW SWITCH QOS command.

Parameter	Meaning
Priority level	The priority level of the received frame.
QOS egress queue	The Quality Of Service egress queue that frames with this priority level join.

Example To display the current configuration of the priority level to QOS egress queue mappings, use the command:

```
SHOW SWITCH QOS
```

Related Commands SET SWITCH QOS
SET QOS HWPRIORITY in *Chapter 7, Quality of Service (QoS)*
SHOW QOS HWPRIORITY in *Chapter 7, Quality of Service (QoS)*

SHOW SWITCH TRUNK

Syntax `SHOW SWITCH TRUNK[=`*`trunk`*`]`

where *trunk* is a character string 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-).

Description This command displays information about the specified trunk group, or all trunk groups on the switch (Figure 3-27 on page 3-149, Table 3-42 on page 3-149).

The TRUNK parameter specifies the name of the trunk group. The name is not case sensitive. The name uniquely identifies the trunk group. The trunk group specified must already exist.

Figure 3-27: Example output from the SHOW SWITCH TRUNK command

```
Switch trunk groups
-------------------------------------------------------------
Trunk group name ............... Uplink
  Speed ........................... 1000Mbps
  Selection criterion ............ Destination MAC address
  Ports ........................... 25,26
-------------------------------------------------------------
```

Table 3-42: Parameters in the output of the SHOW SWITCH TRUNK command.

Parameter	Meaning
Trunk group name	The name of the trunk group.
Speed	The configured speed of the trunk group ports, either "10Mbps", "100Mbps" or "1000Mbps", or "-" (speed has not been set yet).
Selection criterion	The selection criterion used to choose the trunk port on which a packet is to be sent.
Ports	A list of the ports in the trunk group, by port number.

Example To display information about all trunk groups, use the command:

```
SHOW SWITCH TRUNK
```

To display the settings for the *Uplink* trunk group, use the command:

```
SHOW SWITCH TRUNK=Uplink
```

Related Commands ADD SWITCH TRUNK
CREATE SWITCH TRUNK
DELETE SWITCH TRUNK
DESTROY SWITCH TRUNK
SET SWITCH TRUNK

SHOW VLAN

Syntax `SHOW VLAN[={`*`vlan-name`*`|1..4094|ALL}]`

where:

- *vlan-name* is a unique name for the VLAN 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-). The *vlan-name* cannot be a number or ALL.

Description This command displays information about the specified VLAN. If no VLAN or ALL is specified, then all VLANs are displayed (Figure 3-28 on page 3-150, Table 3-43 on page 3-151).

Figure 3-28: Example output from the SHOW VLAN command.

```
VLAN Information
------------------------------------------------------------------------------
 Name ............... default
 Identifier ......... 1
 Status ............. static
 Protected .......... No
 Untagged ports ..... 1,3-23
 Tagged ports ....... None
 Spanning Tree ...... default
 Trunk ports ........ None
 Mirror port ........ None
 Attachments:
 Module           Protocol         Format    Discrim     MAC address
 ---------------------------------------------------------------------
 GARP             Spanning tree    802.2     42              -
 IP               IP               Ethernet  0800            -
 IP               ARP              Ethernet  0806            -
 ---------------------------------------------------------------------

 Name ............... v2
 Identifier ......... 2
 Status ............. dynamic
 Protected .......... No
 Untagged ports ..... 2,24
 Tagged ports ....... None
 Spanning Tree ...... default
 Trunk ports ........ None
 Mirror port ........ None
 Attachments:
 Module           Protocol         Format    Discrim     MAC address
 ---------------------------------------------------------------------
 GARP             Spanning tree    802.2     42              -
 ---------------------------------------------------------------------
```

Table 3-43: Parameters displayed in the output of the SHOW VLAN command .

Parameter	Meaning
Name	The name of the VLAN.
Identifier	The numerical VLAN identifier of the VLAN.
Status	The status of the VLAN, either dynamic or static.
Protected	Whether the VLAN is a protected VLAN.
Untagged Ports	A list of untagged ports that belong to the VLAN.
Tagged Ports	A list of tagged ports that belong to the VLAN.
Spanning Tree	The name of the Spanning Tree Protocol to which the VLAN belongs.
Trunk ports	The list of switch ports that belong to trunk groups. This field is displayed when a port in the VLAN also belongs to a trunk group.
Mirror port	The mirror port for the switch, or "None". Displayed for the default VLAN only.
Attachments	This section contains information about attachments to the VLAN made by other modules in the switch.
Module	The name of the software module attached to the VLAN.
Protocol	The name of the protocol, which is determined from the format and identification number.
Format	The encapsulation format specified by the module.
Discrim	The discriminator specified by the module to identify which packets of the given format should be received.
MAC Address	The Media Access Control source address for which the module wants to receive packets. This is commonly known as the Ethernet address.

Examples To display information on the *marketing* VLAN, use the command:

```
SHOW VLAN=marketing
```

Related Commands CREATE VLAN
DESTROY VLAN

SHOW VLAN DEBUG

Syntax `SHOW VLAN DEBUG`

Description This command displays debug information for all VLANs (Figure 3-29 on page 3-152, Table 3-44 on page 3-152).

Figure 3-29: Example output from the SHOW VLAN DEBUG command

```
Vlan      Enabled Debug Modes       Output     Timeout
-------------------------------------------------------------
Vlan1     PKT                       16         NONE
-------------------------------------------------------------

Vlan      Enabled Debug Modes       Output     Timeout
-------------------------------------------------------------
Vlan4060  None
-------------------------------------------------------------
```

Table 3-44: Parameters in the output of the SHOW VLAN DEBUG command.

Parameter	Meaning
VLAN	A string comprising the constant "Vlan" and the VLAN Identifier of the VLAN.
Enabled Debug Modes	The debugging option for the VLAN; either "PKT" or "None".
Output	The output device for the VLAN. This is shown when a debug mode is enabled.
Timeout	The length of time in seconds that debugging options for the VLAN are enabled. This is shown when a debug mode is enabled. If a timeout value is not set, "None" is shown.

Examples To display debugging information for all VLANs, use the command:

```
SHOW VLAN DEBUG
```

Related Commands DISABLE VLAN DEBUG
ENABLE VLAN DEBUG

SHOW VLANRELAY

Syntax `SHOW VLANRELAY[=name]`

where *name* is a unique name for the VLAN relay entity 1 to 15 characters long. Valid characters are uppercase and lowercase letters, digits (0-9), the underscore character ("_"), and the hyphen (-).

Description This command displays information about one or all of the currently-configured VLAN relay entities (Figure 1, Table 1).

The VLANRELAY parameter specifies the name of the VLAN relay entity for which to show information. If the name is not given, information about all VLAN relay entities is displayed.

Figure 3-30: Example output from the SHOW VLANRELAY command.

```
VLAN relay entities
---------------------------------------------------------------
Name ............................... SNARelay
 Enabled ........................... Yes
 Debugging ......................... No
 Protocol .......................... 00
 Protocol .......................... 04
 VLAN .............................. 2 (Accounts)
 VLAN .............................. 5 (Admin)
 VLAN .............................. 16 (Sales)
Packet counters:
  VLAN 2 to VLAN 5 ................. 2345
           VLAN 16 ................. 148
  VLAN 5 to VLAN 2 ................. 2567
           VLAN 16 ................. 754
  VLAN 16 to VLAN 2 ................ 174
            VLAN 5 ................. 802
---------------------------------------------------------------
```

Table 3-45: Parameters displayed in the output of the SHOW VLANRELAY command .

Parameter	Meaning
Name	The name of the VLAN relay entity.
Enabled	Whether the VLAN relay entity is enabled or not.
Debugging	Whether packet debugging for the VLAN relay entity is enabled or not.
Protocol	The protocol number of each protocol that is relayed by the VLAN relay entity.
VLAN	The numerical VLAN Identifier and name of each VLAN that has been added to the VLAN relay entity.
Packet counters	The number of packets that have been relayed between VLANs by this VLAN relay entity.

Example To show the configuration and counters for the VLAN relay entity SNARelay, use the command:

```
SHOW VLANRELAY=SNARelay
```

Related Commands ADD VLANRELAY
CREATE VLANRELAY
DELETE VLANRELAY
DESTROY VLANRELAY