

Port Traffic Controls

Contents

Overview	13-3
Rate-Limiting	13-4
All Traffic Rate-Limiting	13-4
Configuring Rate-Limiting	13-5
Displaying the Current Rate-Limit Configuration	13-6
Operating Notes for Rate-Limiting	13-8
ICMP Rate-Limiting	13-10
Terminology	13-11
Guidelines for Configuring ICMP Rate-Limiting	13-12
Configuring ICMP Rate-Limiting	13-13
Using Both ICMP Rate-Limiting and All-Traffic Rate-Limiting on the Same Interface	13-14
Displaying the Current ICMP Rate-Limit Configuration	13-14
Operating Notes for ICMP Rate-Limiting	13-15
ICMP Rate-Limiting Trap and Event Log Messages	13-17
Guaranteed Minimum Bandwidth (GMB)	13-20
Introduction	13-20
Terminology	13-20
GMB Operation	13-20
Impacts of QoS Queue Configuration on GMB Operation	13-22
Configuring Guaranteed Minimum Bandwidth for Outbound Traffic	13-23
Displaying the Current Guaranteed Minimum Bandwidth Configuration	13-26
GMB Operating Notes	13-27
Jumbo Frames	13-28
Terminology	13-28
Operating Rules	13-29

Configuring Jumbo Frame Operation	13-30
Overview	13-30
Viewing the Current Jumbo Configuration	13-31
Enabling or Disabling Jumbo Traffic on a VLAN	13-33
Configuring a Maximum Frame Size	13-33
Configuring IP MTU	13-34
SNMP Implementation	13-34
Displaying the Maximum Frame Size	13-35
Operating Notes for Maximum Frame Size	13-35
Operating Notes for Jumbo Traffic-Handling	13-35
Troubleshooting	13-38

Overview

Feature	Default	Menu	CLI	Web
Rate-Limiting	None	n/a	13-4	n/a
Guaranteed Minimum Bandwidth	Per Queue (1-8 order): 2%-3%-30%-10%-10%- 10%-15%-20%	n/a	13-19	n/a
Jumbo Packets	Disabled	n/a	13-27	n/a

This chapter includes:

- **Rate-Limiting:** Enables a port to limit the amount of bandwidth a user or device may utilize for traffic on the switch.

Note

In earlier releases, all traffic rate-limiting applied to inbound traffic only, and was specified as a percentage of total bandwidth. Beginning with software release K.12.*xxx* or later, it is also possible to configure outbound rate-limiting for all traffic on a port, and specify bandwidth usage in terms of kilobits per second (kbps).

- **Guaranteed Minimum Bandwidth (GMB):** Provides a method for ensuring that each of a port's outbound queues has a specified minimum consideration for sending traffic out on the link to another device.
- **Jumbo Frames:** Enables ports operating at 1 Gbps or 10 Gbps speeds to accept inbound frames of up to 9220 bytes when configured for jumbo traffic.

Rate-Limiting

Feature	Default	Menu	CLI	Web
rate-limit all	none	n/a	page 13-5	n/a
show rate-limit all	n/a	n/a	page 13-6	n/a
rate-limit icmp	none	n/a	page 13-13	n/a
show rate-limit icmp	n/a	n/a	page 13-14	n/a

All Traffic Rate-Limiting

Rate-limiting for all traffic operates on a per-port basis to allow only the specified bandwidth to be used for inbound or outbound traffic. When traffic exceeds the configured limit, it is dropped. This effectively sets a usage level on a given port, and is a tool for enforcing maximum service level commitments granted to network users. This feature operates on a per-port level and is not configurable on port trunks. Note that rate-limiting is designed to be applied at the network edge to limit traffic from non-critical users or to enforce service agreements such as those offered by Internet Service Providers (ISPs) to provide only the bandwidth for which a customer has paid.

Note

Rate-limiting also can be applied by a RADIUS server during an authentication client session. For further details, refer to the chapter titled “*RADIUS Authentication and Accounting*” in the *Access Security Guide* for your switch.

Caution

Rate-limiting is intended for use on edge ports in a network. It is not recommended for use on links to other switches, routers, or servers within a network, or for use in the network core. Doing so can interfere with applications the network requires to function properly.

Note

The switches covered in this guide also support ICMP rate-limiting to mitigate the effects of certain ICMP-based attacks. For more information, refer to “ICMP Rate-Limiting” on page 13-10.

Configuring Rate-Limiting

Note

The mode using bits per second (bps) in releases before K.12.XX has been replaced by the kilobits per second (kbps) mode. Switches that have configurations with bps values will be automatically converted when you update your software to the new version. However, an older config file with bps values must be updated manually to kbps values or it will not load successfully onto a switch running later versions of the software (K.12.XX or greater).

The **rate-limit all** command controls the rate of traffic sent or received on a port by setting a limit on the bandwidth available. It includes options for:

- Rate-limiting on either inbound or outbound traffic.
- Specifying the traffic rate as either a percentage of bandwidth, or in terms of bits per second.

Syntax: [no] int <port-list> rate-limit all < in | out > <percent <0-100> | kbps < 0-10000000>>

Configures a traffic rate limit (on non-trunked ports) on the link. The "no" form of the command disables rate-limiting on the specified ports.

*(Default: **Disabled.**)*

Options include:

- **in** or **out** — *Specifies a traffic rate limit on inbound traffic passing through that port, or on outbound traffic.*
- **percent** or **kbps** — *Specifies the rate limit as a percentage of total available bandwidth, or in kilobits per second.*

Notes:

- *The **rate-limit icmp** command specifies a rate limit on inbound ICMP traffic only (see "ICMP Rate-Limiting" on page 13-9).*
- *Rate-limiting does not apply to trunked ports (including meshed ports).*

—Continued—

- *Kbps rate-limiting is done in segments of 1% of the lowest corresponding media speed. For example, if the media speed is 100 Kbps, the value would be 1 Mbps. A 1-100 Kbps rate-limit is implemented as a limit of 100 Kbps; a limit of 100-199 Kbps is also implemented as a limit of 100 Kbps, a limit of 200-299 Kbps is implemented as a limit of 200 Kbps, and so on.*
- *Percentage limits are based on link speed. For example, if a 100 Mbps port negotiates a link at 100 Mbps and the inbound rate-limit is configured at 50%, then the traffic flow through that port is limited to no more than 50 Mbps. Similarly, if the same port negotiates a 10 Mbps link, then it allows no more than 5 Mbps of inbound traffic.*

*Configuring a rate limit of 0 (zero) on a port blocks all traffic on that port. However, if this is the desired behavior on the port, ProCurve recommends using the **< port-list > disable** command instead of configuring a rate limit of 0.*

You can configure a rate limit from either the global configuration level or from the port context level. For example, either of the following commands configures an inbound rate limit of 60% on ports A3 - A5:

```
ProCurve (config)# int a3-a5 rate-limit all in percent 60
ProCurve (eth-A3-A5)# rate-limit all in percent 60
```

Displaying the Current Rate-Limit Configuration

The **show rate-limit all** command displays the per-port rate-limit configuration in the running-config file.

Syntax: show rate-limit all [*port-list*]

*Without [**port-list**], this command lists the rate-limit configuration for all ports on the switch. With [**port-list**], this command lists the rate-limit configuration for the specified port(s). This command operates the same way in any CLI context.*

For example, if you wanted to view the rate-limiting configuration on the first six ports in the module in slot “A”:

```
ProCurve(eth-A5)# show rate-limit all a1-a6
```

All-Traffic Rate Limit Maximum %

Port	Inbound Limit Mode	Radius Override	Outbound Limit Mode
A1	Disabled	No-override	200 kbps
A2	Disabled	No-override	200 kbps
A3	Disabled	No-override	200 kbps
A4	Disabled	No-override	200 kbps
A5	20 %	No-override	Disabled
A6	Disabled	No-override	Disabled

Ports A1-A4 are configured with an outbound rate limit of 200 Kbps; Port A5 is configured with an inbound rate limit of 20%. (Port A6 is not configured for rate-limiting.)

Figure 13-1. Example of Listing the Rate-Limit Configuration

The **show running** command displays the currently applied setting for any interfaces in the switch configured for all traffic rate-limiting and ICMP rate limiting. The **show config** command displays this information for the configuration currently stored in the startup-config file. (Note that configuration changes performed with the CLI, but not followed by a **write mem** command do not appear in the startup-config file.)

```
ProCurve(config)# show config

Startup configuration:

; J8697A Configuration Editor; Created on release #K.12.XX

hostname "ProCurve Switch 8212z1"
module 1 type J8705A
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24
  ip address dhcp-bootp
  exit
interface A1
  rate-limit all out kbps 100
  exit
interface A2
  rate-limit all out kbps 100
  exit
interface A3
  rate-limit all out kbps 100
  exit
interface A4
  rate-limit all out kbps 100
  exit
=====
interface A5
  rate-limit icmp 60
  exit
interface A6
  rate-limit icmp 60
  exit
```

Ports A1-A4 are configured with an outbound rate limit of 100 kbps.

Ports A3-A5 are configured with a rate limit for icmp traffic of 60 %.

Figure 13-2. Example of Rate-Limit Settings Listed in the “show config” Output

Operating Notes for Rate-Limiting

- **Rate-limiting operates on a per-port basis**, regardless of traffic priority. Rate-limiting is available on all types of ports (other than trunked ports) on the switches covered in this guide, and at all port speeds configurable for these devices.
- **Rate-limiting is not allowed on trunked ports:** Rate-limiting is not supported on ports configured in a trunk group (including mesh ports). Configuring a port for rate-limiting and then adding it to a trunk suspends

rate-limiting on the port while it is in the trunk. Attempting to configure rate-limiting on a port that already belongs to a trunk generates the following message:

```
<port-list>: Operation is not allowed for a trunked port.
```

- **Rate-limiting for inbound and outbound traffic are separate features:** The rate limits for each direction of traffic flow on the same port are configured separately—even the specified limits can be different.
- **Rate-limiting is visible as an outbound forwarding rate:** Because inbound rate-limiting is performed on packets during packet-processing, it is not shown via the inbound drop counters. Instead, this limit is verifiable as the ratio of outbound traffic from an inbound rate-limited port versus the inbound rate. For outbound rate-limiting, the rate is visible as the percentage of available outbound bandwidth (assuming that the amount of requested traffic to be forwarded is larger than the rate-limit).
- **Operation with other features:** Configuring rate-limiting on a port where other features affect port queue behavior (such as flow control) can result in the port not achieving its configured rate-limiting maximum. For example, in a situation where flow control is configured on a rate-limited port, there can be enough “back pressure” to hold high-priority inbound traffic from the upstream device or application to a rate that is lower than the configured rate limit. In this case, the inbound traffic flow does not reach the configured rate and lower priority traffic is not forwarded into the switch fabric from the rate-limited port. (This behavior is termed “head-of-line blocking” and is a well-known problem with flow-control.) In another type of situation, an outbound port can become oversubscribed by traffic received from multiple rate-limited ports. In this case, the actual rate for traffic on the rate-limited ports may be lower than configured because the total traffic load requested to the outbound port exceeds the port’s bandwidth, and thus some requested traffic may be held off on inbound.
- **Traffic filters on rate-limited ports:** Configuring a traffic filter on a port does not prevent the switch from including filtered traffic in the bandwidth-use measurement for rate-limiting when it is configured on the same port. For example, ACLs, source-port filters, protocol filters, and multicast filters are all included in bandwidth usage calculations.
- **Monitoring (Mirroring) rate-limited interfaces:** If monitoring is configured, packets dropped by rate-limiting on a monitored interface will still be forwarded to the designated monitor port. (Monitoring shows what traffic is inbound on an interface, and is not affected by “drop” or “forward” decisions.)

- **Optimum rate-limiting operation:** Optimum rate-limiting occurs with 64-byte packet sizes. Traffic with larger packet sizes can result in performance somewhat below the configured bandwidth. This is to ensure the strictest possible rate-limiting of all sizes of packets.

Note on Testing Rate-Limiting

Rate-limiting is applied to the available bandwidth on a port, and not to any specific applications running through the port. If the total bandwidth requested by all applications is less than the configured maximum rate, then no rate-limit can be applied. This situation occurs with a number of popular throughput-testing applications, as well as most regular network applications. Consider the following example that uses the minimum packet size:

The total available bandwidth on a 100 Mbps port “X” (allowing for Inter-packet Gap—IPG), with no rate-limiting restrictions, is:

$$(((100,000,000 \text{ bits}) / 8) / 84) \times 64 = 9,523,809 \text{ bytes per second}$$

where:

- The divisor (84) includes the 12-byte IPG, 8-byte preamble, and 64-bytes of data required to transfer a 64-byte packet on a 100 Mbps link.
- Calculated “bytes-per-second” includes packet headers and data. This value is the maximum “bytes-per-second” that 100 Mbps can support for minimum-sized packets.

Suppose port “X” is configured with a rate limit of 50% (4,761,904 bytes). If a throughput-testing application is the only application using the port, and transmits 1 Mbyte of data through the port, it uses only 10.5% of the port’s available bandwidth, and the rate-limit of 50% has no effect. This is because the maximum rate permitted (50%) exceeds the test application’s bandwidth usage (126,642-164,062 bytes, depending upon packet size, which is only 1.3-1.7% of the available total). Before rate-limiting can occur, the test application’s bandwidth usage must exceed 50% of the port’s total available bandwidth. That is, to test the rate-limit setting, the following must be true:

$$\text{bandwidth usage} > (0.50 \times 9,523,809)$$

ICMP Rate-Limiting

In IP networks, ICMP (Internet Control Message Protocol) messages are generated in response to either inquiries or requests from routing and diagnostic functions. These messages are directed to the applications originating the inquiries. In unusual situations, if the messages are generated rapidly with the intent of overloading network circuits, they can threaten network availability. This problem is visible in denial-of-service (DoS) attacks or other malicious behaviors where a worm or virus overloads the network with ICMP

messages to an extent where no other traffic can get through. (ICMP messages themselves can also be misused as virus carriers). Such malicious misuses of ICMP can include a high number of ping packets that mimic a valid source IP address and an invalid destination IP address (spoofed pings), and a high number of response messages (such as Destination Unreachable error messages) generated by the network.

ICMP rate-limiting provides a method for limiting the amount of bandwidth that may be utilized for inbound ICMP traffic on a switch port or trunk. This feature allows users to restrict ICMP traffic to percentage levels that permit necessary ICMP functions, but throttle additional traffic that may be due to worms or viruses (reducing their spread and effect). In addition, ICMP rate-limiting preserves inbound port bandwidth for non-ICMP traffic.

Caution

The ICMP protocol is necessary for routing, diagnostic, and error responses in an IP network. ICMP rate-limiting is primarily used for throttling worm or virus-like behavior, and should normally be configured to allow one to five per cent of available inbound bandwidth (at 10 Mbps or 100 Mbps speeds) or 100 - 10,000 kbps (1Gbps or 10 Gbps speeds) to be used for ICMP traffic. ***This feature should not be used to remove all ICMP traffic from a network.***

Note

ICMP rate-limiting does not throttle non-ICMP traffic. In cases where you want to throttle both ICMP traffic and all other inbound traffic on a given interface, you can separately configure both ICMP rate-limiting and all-traffic rate-limiting.

Beginning with software release K.12.*xx* or later, the all-traffic rate-limiting command (**rate-limit all**) and the ICMP rate-limiting command (**rate-limit icmp**) operate differently:

- All traffic rate-limiting applies to both inbound and outbound traffic, and can be specified either in terms of a percentage of total bandwidth or in terms of bits per second;
 - ICMP rate-limiting applies only to inbound traffic, and can only be specified as a percentage of total bandwidth.
-

Terminology

All-Traffic Rate-Limiting: Applies a rate-limit to all traffic (including ICMP traffic) on an interface. For details, see “Rate-Limiting” on page 13-4.

ICMP Rate-Limiting: Applies a rate-limit to all *inbound* ICMP traffic received on an interface, but does not limit other types of inbound traffic.

Spoofed Ping: An ICMP echo request packet intentionally generated with a valid source IP address and an invalid destination IP address. Spoofed pings are often created with the intent to oversubscribe network resources with traffic having invalid destinations.

Guidelines for Configuring ICMP Rate-Limiting

Apply ICMP rate-limiting on all connected interfaces on the switch to effectively throttle excessive ICMP messaging from any source. Figure 13-3 shows an example of how to configure this for a small to mid-sized campus though similar rate-limit thresholds are applicable to other network environments. On edge interfaces, where ICMP traffic should be minimal, a threshold of 1% of available bandwidth should be sufficient for most applications. On core interfaces, such as switch-to-switch and switch-to-router, a maximum threshold of 5% should be sufficient for normal ICMP traffic. (“Normal” ICMP traffic levels should be the maximums that occur when the network is rebooting.)

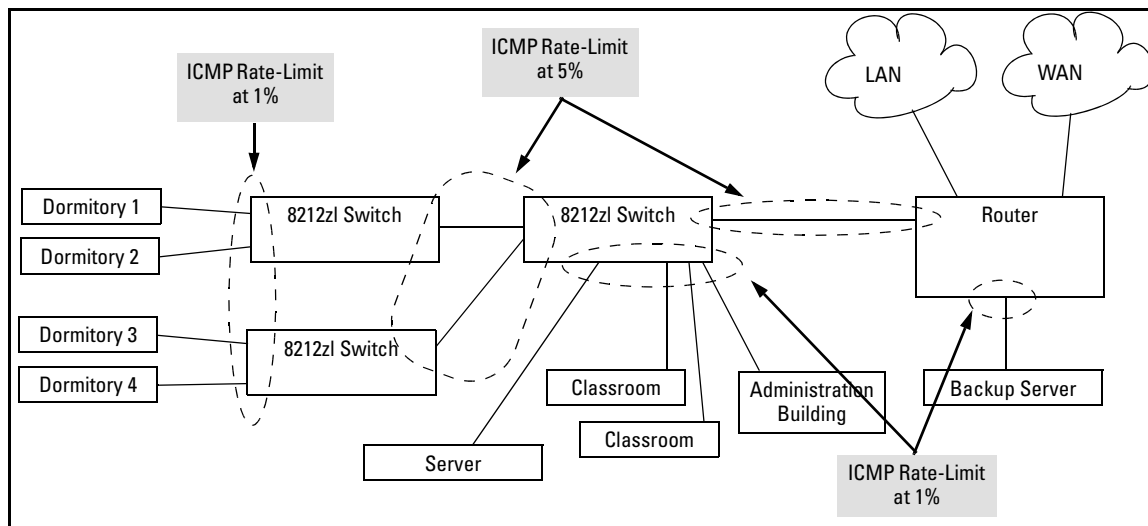


Figure 13-3. Example of ICMP Rate-Limiting

Configuring ICMP Rate-Limiting

The **rate-limit icmp** command controls inbound usage of a port by setting a limit on the bandwidth available for inbound ICMP traffic.

Syntax: [no] int < port-list > rate-limit icmp <percent < 0-100 > | kbps <0-10000000>>

*Configures inbound ICMP traffic rate limiting. You can configure a rate limit from either the global configuration level (as shown above) or from the interface context level. The **no** form of the command disables ICMP rate-limiting on the specified interface(s). (Default: **Disabled**.)*

percent <1-100>: Values in this range allow ICMP traffic as a percentage of the bandwidth available on the interface.

kbps <0-10000000>: Specifies the rate at which to forward traffic in kilobits per second.

0: This value causes an interface to drop all incoming ICMP traffic, and is not recommended. Refer to the Caution on page 13-11.

Note: ICMP rate-limiting is not supported on meshed ports. (Rate-limiting can reduce the efficiency of paths through a mesh domain).

For example, either of the following commands configures an inbound rate limit of 1% on ports A3 - A5, which are used as network edge ports:

```
ProCurve(config)# int a3-a5 rate-limit icmp 1
ProCurve (eth-A3-A5)# rate-limit icmp 1
```

Note

When using kbps-mode ICMP rate-limiting, the rate-limiting only operates on the IP payload part of the ICMP packet (as required by metering RFC 2698). This means that effective metering is at a rate greater than the configured rate, with the disparity increasing as the packet size decreases (the packet to payload ratio is higher).

Also, in kbps mode, metering accuracy is limited at low values, for example, less than 45 Kbps. This is to allow metering to function well at higher media speeds such as 10 Gbps.

Using Both ICMP Rate-Limiting and All-Traffic Rate-Limiting on the Same Interface

ICMP and all-traffic rate-limiting can be configured on the same interface. All-traffic rate-limiting applies to all inbound or outbound traffic (including ICMP traffic), while ICMP rate-limiting applies only to inbound ICMP traffic.

Note that if the all-traffic load on an interface meets or exceeds the currently configured all-traffic inbound rate-limit while the ICMP traffic rate-limit on the same interface has not been reached, then all excess traffic will be dropped, including any inbound ICMP traffic above the all-traffic limit (regardless of whether the ICMP rate-limit has been reached). Suppose, for example:

- The all-traffic inbound rate-limit on port “X” is configured at 55% of the port’s bandwidth.
- The ICMP traffic rate-limit on port “X” is configured at 2% of the port’s bandwidth.

If at a given moment:

- Inbound ICMP traffic on port “X” is using 1% of the port’s bandwidth, and
- Inbound traffic of all types on port “X” demands 61% of the ports’s bandwidth,

then all inbound traffic above 55% of the port’s bandwidth, including any additional ICMP traffic, will be dropped as long as all inbound traffic combined on the port demands 55% or more of the port’s bandwidth.

Displaying the Current ICMP Rate-Limit Configuration

The **show rate-limit icmp** command displays the per-interface ICMP rate-limit configuration in the running-config file.

Syntax: show rate-limit icmp [*port-list*]

Without [port-list], this command lists the ICMP rate-limit configuration for all ports on the switch. With [port-list], this command lists the rate-limit configuration for the specified interface(s). This command operates the same way in any CLI context.

For example, if you wanted to view the rate-limiting configuration on the first six ports in the module in slot “B”:

```
ProCurve(config)# show rate-limit icmp b1-b6

Inbound ICMP Rate Limit Maximum Percentage

Port | Mode      Rate
-----+-----
B1   | Disabled  Disabled
B2   | kbps      100
B3   | %         5
B4   | %         1
B5   | %         1
B6   | Disabled  Disabled
```

Figure 13-4. Example of Listing the Rate-Limit Configuration

The **show running** command displays the currently applied setting for any interfaces in the switch configured for all traffic rate-limiting and ICMP rate-limiting. The **show config** command displays this information for the configuration currently stored in the startup-config file. (Note that configuration changes performed with the CLI, but not followed by a **write mem** command do not appear in the startup-config file.)

Operating Notes for ICMP Rate-Limiting

ICMP rate-limiting operates on an interface (per-port) basis to allow, on average, the highest expected amount of legitimate, inbound ICMP traffic.

- **Interface support:** ICMP rate-limiting is available on all types of ports (other than trunk ports or mesh ports), and at all port speeds configurable for the switch.
- **Rate-limiting is not permitted on mesh ports:** Either type of rate-limiting (all traffic or ICMP) can reduce the efficiency of paths through a mesh domain.
- **Rate-limiting is not supported on port trunks:** Neither all-traffic nor ICMP rate-limiting are supported on ports configured in a trunk group.
- **ICMP percentage-based rate-limits are calculated as a percentage of the negotiated link speed:** For example, if a 100 Mbps port negotiates a link to another switch at 100 Mbps and is ICMP rate-limit configured at 5%, then the inbound ICMP traffic flow through that port is limited to 5 Mbps. Similarly, if the same port negotiates a 10 Mbps link, then it allows

0.5 Mbps of inbound traffic. If an interface experiences an inbound flow of ICMP traffic in excess of its configured limit, the switch generates a log message and an SNMP trap (if an SNMP trap receiver is configured).

- **ICMP rate-limiting is port-based:** ICMP rate-limiting reflects the available percentage of an interface's entire inbound bandwidth. The rate of inbound flow for traffic of a given priority and the rate of flow from an ICMP rate-limited interface to a particular queue of an outbound interface are not measures of the actual ICMP rate limit enforced on an interface.
- **Below-maximum rates:** ICMP rate-limiting operates on a per-interface basis, regardless of traffic priority. Configuring ICMP rate-limiting on an interface where other features affect inbound port queue behavior (such as flow control) can result in the interface not achieving its configured ICMP rate-limiting maximum. For example, in some situations with flow control configured on an ICMP rate-limited interface, there can be enough "back pressure" to hold high-priority inbound traffic from the upstream device or application to a rate that does not allow bandwidth for lower-priority ICMP traffic. In this case, the inbound traffic flow may not permit the forwarding of ICMP traffic into the switch fabric from the rate-limited interface. (This behavior is termed "head-of-line blocking" and is a well-known problem with flow-control.) In cases where both types of rate-limiting (**rate-limit all** and **rate-limit icmp**) are configured on the same interface, this situation is more likely to occur. In another type of situation, an outbound interface can become oversubscribed by traffic received from multiple ICMP rate-limited interfaces. In this case, the actual rate for traffic on the rate-limited interfaces may be lower than configured because the total traffic load requested to the outbound interface exceeds the interface's bandwidth, and thus some requested traffic may be held off on inbound.
- **Monitoring (Mirroring) ICMP rate-limited interfaces:** If monitoring is configured, packets dropped by ICMP rate-limiting on a monitored interface will still be forwarded to the designated monitor port. (Monitoring shows what traffic is inbound on an interface, and is not affected by "drop" or "forward" decisions.)
- **Optimum rate-limiting operation:** Optimum rate-limiting occurs with 64-byte packet sizes. Traffic with larger packet sizes can result in performance somewhat below the configured inbound bandwidth. This is to ensure the strictest possible rate-limiting of all sizes of packets.
- **Outbound Traffic Flow:** Configuring ICMP rate-limiting on an interface does not control the rate of outbound traffic flow on the interface.

**Note on Testing
ICMP Rate-Limiting**

ICMP rate-limiting is applied to the available bandwidth on an interface. If the total bandwidth requested by all ICMP traffic is less than the available, configured maximum rate, then no ICMP rate-limit can be applied. That is, an

interface must be receiving more inbound ICMP traffic than the configured bandwidth limit allows. If the interface is configured with both **rate-limit all** and **rate-limit icmp**, then the ICMP limit can be met or exceeded only if the rate limit for all types of inbound traffic has not already been met or exceeded. Also, to test the ICMP limit it is necessary to generate ICMP traffic that exceeds the configured ICMP rate limit. Using the recommended settings—1% for edge interfaces and 5% maximum for core interfaces—it is easy to generate sufficient traffic. However, if you are testing with higher maximums, it is necessary to ensure that the ICMP traffic volume exceeds the configured maximum.

Note also that testing ICMP rate-limiting where inbound ICMP traffic on a given interface has destinations on multiple outbound interfaces, the test results must be based on the received outbound ICMP traffic.

ICMP rate-limiting is not reflected in counters monitoring inbound traffic because inbound packets are counted before the ICMP rate-limiting drop action occurs.

ICMP Rate-Limiting Trap and Event Log Messages

If the switch detects a volume of inbound ICMP traffic on a port that exceeds the ICMP rate-limit configured for that port, it generates one SNMP trap and one informational Event Log message to notify the system operator of the condition. (The trap and Event Log message are sent within two minutes of when the event occurred on the port.)

For example:

```
I 06/30/05 11:15:42 RateLim: ICMP traffic exceeded  
configured limit on port A1
```

These trap and Event Log messages provide an advisory that inbound ICMP traffic on a given interface has exceeded the configured maximum. The additional ICMP traffic is dropped, but the excess condition may indicate an infected host (or other traffic threat or network problem) on that interface. The system operator should investigate the attached devices or network conditions further.

The switch does not send more traps or Event Log messages for excess ICMP traffic on the affected port until the system operator resets the port's ICMP trap function. The reset can be done through SNMP from a network management station or through the CLI with the following **setmib** command.

Syntax: `setmib hplcmpRatelimitPortAlarmflag.< internal-port-#> -i 1`

On a port configured with ICMP rate-limiting, this command resets the ICMP trap function, which allows the switch to generate a new SNMP trap and an Event Log message if ICMP traffic in excess of the configured limit is detected on the port.

For example, an operator noticing an ICMP rate-limiting trap or Event Log message originating with port A1 on a switch would use the following **setmib** command to reset the port to send a new message if the condition occurs again.

```
ProCurve(config)# setmib hpicmpratelimitportalarm-  
flag.1 -i 1
```

Determining the Switch Port Number Used in ICMP Port Reset

Commands: To enable excess ICMP traffic notification traps and Event Log messages, use the **setmib** command described on page 13-17. The port number included in the command corresponds to the internal number the switch maintains for the designated port, and not the port’s external (slot/number) identity.

To match the port’s external slot/number to the internal port number, use the **walkmib ifDescr** command, as shown in the following figure:

```
ProCurve# walkmib ifDescr  
┌ ifDescr.1 = A1 ───┐  
│ ifDescr.2 = A2 │  
│ ifDescr.3 = A3 │  
│ . │  
│ . │  
└ ifDescr.23 = A23 ─┘  
┌ ifDescr.24 = A24 ─┐  
│ ifDescr.27 = B1 │  
│ ifDescr.28 = B2 │  
│ ifDescr.29 = B3 │  
│ . │  
│ . │  
└ ifDescr.48 = B22 ─┘  
┌ ifDescr.49 = B23 ─┐  
│ ifDescr.50 = B24 │  
│ . │  
│ . │  
└ ───┘
```

The diagram shows the output of the `walkmib ifDescr` command. The output is enclosed in a large dashed-line box. Two grey callout boxes with arrows point to specific lines in the output. The first callout box, labeled "Beginning and Ending of Port Number Listing for Slot A", points to the line `ifDescr.23 = A23`. The second callout box, labeled "Beginning and Ending of Port Number Listing for Slot B", points to the line `ifDescr.48 = B22`. The output shows a sequence of port numbers: 1-3 (A1-A3), 23-24 (A23-A24), 27-29 (B1-B3), and 48-50 (B22-B24). Ellipses between 3 and 23, and between 29 and 48, indicate intermediate port numbers.

Figure 13-5. Matching Internal Port Numbers to External Slot/Port Numbers

Guaranteed Minimum Bandwidth (GMB)

Feature	Default	Menu	CLI	Web
bandwidth-min output	Per-Queue: 2%-3%-30%-10% 10%-10%-15%-20%	n/a	page 13-22	n/a
show bandwidth output [<i>port-list</i>]	n/a	n/a	page 13-25	n/a

Introduction

Guaranteed Minimum Bandwidth (GMB) provides a method for ensuring that each of a given port's outbound traffic priority queues has a specified minimum consideration for sending traffic out on the link to another device. This can prevent a condition where applications generating lower-priority traffic in the network are frequently or continually "starved" by high volumes of higher-priority traffic. You can configure GMB per-port.

Terminology

Oversubscribed Queue: The condition where there is insufficient bandwidth allocated to a particular outbound priority queue for a given port. If additional, unused bandwidth is not available, the port delays or drops the excess traffic.

GMB Operation

The switch services per-port outbound traffic in a descending order of priority; that is, from the highest priority to the lowest priority. By default, each port offers eight prioritized, outbound traffic queues. Tagged VLAN traffic is prioritized according to the 802.1p priority the traffic carries. Untagged VLAN traffic is assigned a priority of "0" (normal).

Table 13-1. Per-Port Outbound Priority Queues

802.1p Priority Settings in Tagged VLAN Packets*	Outbound Priority Queue for a Given Port
1 (low)	1
2 (low)	2
0 (normal)	3
3 (normal)	4
4 (medium)	5
5 (medium)	6
6 (high)	7
7 (high)	8

*The switch processes outbound traffic from an untagged port at the "0" (normal) priority level.

You can use GMB to reserve a specific percentage of each port's available outbound bandwidth for each of the eight priority queues. This means that regardless of the amount of high priority outbound traffic on a port, you can ensure that there will always be bandwidth reserved for lower-priority traffic.

Since the switch services outbound traffic according to priority (highest to lowest), the highest-priority outbound traffic on a given port automatically receives the first priority in servicing. Thus, in most applications, it is necessary only to specify the minimum bandwidth you want to allocate to the lower priority queues. In this case, the high-priority traffic automatically receives all unassigned bandwidth without starving the lower-priority queues.

Conversely, configuring a bandwidth minimum on only the high-priority outbound queue of a port (and not providing a bandwidth minimum for the lower-priority queues) is not recommended because it may "starve" the lower-priority queues. (See the **Note** on page 13-21.)

Note

For a given port, when the demand on one or more outbound queues exceeds the minimum bandwidth configured for those queues, the switch apportions unallocated bandwidth to these queues on a priority basis. As a result, specifying a minimum bandwidth for a high-priority queue but not specifying a minimum for lower-priority queues can starve the lower-priority queues during periods of high demand on the high priority queue. For example, if a port configured to allocate a minimum bandwidth of 80% for outbound high-priority traffic experiences a demand above this minimum, then this burst starves lower-priority queues that *do not have a minimum configured*. Normally, this will not altogether halt lower priority traffic on the network, but will likely cause delays in the delivery of the lower-priority traffic.

The sum of the GMB settings for all outbound queues on a given port cannot exceed 100%.

Impacts of QoS Queue Configuration on GMB Operation

The section on “*Configuring Guaranteed Minimum Bandwidth for Outbound Traffic*” assumes the ports on the switch offer eight prioritized, outbound traffic queues. This may not always be the case, however, since the switch supports a QoS queue configuration feature that allows you to reduce the number of outbound queues from eight (the default) to four queues or two.

Changing the number of queues affects the GMB commands (**interface bandwidth-min** and **show bandwidth output**) such that they operate only on the number of queues currently configured. If the queues are reconfigured, the guaranteed minimum bandwidth per queue is automatically re-allocated according to the following percentages:

Table 13-2. Default GMB Percentage Allocations per QoS Queue Configuration

802.1p Priority	8 Queues (default)	4 Queues	2 Queues
1 (lowest)	2%	10%	90%
2	3%		
0 (normal)	30%	70%	
3	10%		
4	10%	10%	10%
5	10%		
6	15%	10%	
7 (highest)	20%		

Note

For more information on queue configuration and the associated default minimum bandwidth settings, refer to the chapter titled “*Quality of Service (QoS): Managing Bandwidth More Effectively*” in the *Advanced Traffic Management Guide* for your switch.

Configuring Guaranteed Minimum Bandwidth for Outbound Traffic

For any port or group of ports you can configure either the default minimum bandwidth settings for each outbound priority queue or a customized bandwidth allocation. For most applications, ProCurve recommends configuring GMB with the same values on all ports on the switch so that the outbound traffic profile is consistent for all outbound traffic. However, there may be instances where it may be advantageous to configure special profiles on connections to servers or to the network infrastructure (such as links to routers, other switches, or to the network core).

Syntax: [no] int < port-list > bandwidth-min output

Configures the default minimum bandwidth allocation for the outbound priority queue for each port in < port-list >. The default values per priority queue are:

- Queue 1 (low priority): 2%
- Queue 2 (low priority): 3%
- Queue 3 (normal priority): 30%
- Queue 4 (normal priority): 10%
- Queue 5 (medium priority): 10%
- Queue 6 (medium priority): 10%
- Queue 7 (high priority): 15%
- Queue 8 (high priority): 20%

*The **no** form of the command disables GMB for all ports in < port-list >. In this state, which is the equivalent of setting all outbound queues on a port to **0** (zero), a high level of higher-priority traffic can starve lower-priority queues, which can slow or halt lower-priority traffic in the network. You can configure bandwidth minimums from either the global configuration level (as shown above) or from the port context level. For information on outbound port queues, refer to table 13-1, “Per-Port Outbound Priority Queues” on page 13-20.*

Syntax: [no] int < port-list > bandwidth-min output

[< queue1% > < queue2% > < queue3% > < queue4% > < queue5% >
< queue6% > < queue7% > < queue8% >]

For ports in < port-list >, specifies the minimum outbound bandwidth as a percent of the total bandwidth for each outbound queue. The queues receive service in descending order of priority. You must specify a bandwidth percent value for all eight queues, and the sum of the bandwidth percentages must not exceed 100%. (0 is a value for a queue percentage setting.) Configuring a total of less than 100% across the eight queues results in unallocated bandwidth that remains harmlessly unused unless a given queue becomes oversubscribed. In this case, the unallocated bandwidth is apportioned to oversubscribed queues in descending order of priority. For example, if you configure a minimum of 10% for queues 1 - 7, and 0% for queue 8, then the unallocated bandwidth will be available to all eight queues in the following prioritized order:

1. Queue 8 (high priority)
2. Queue 7 (high priority)
3. Queue 6 (medium priority)
4. Queue 5 (medium priority)
5. Queue 4 (normal priority)
6. Queue 3 (normal priority)
7. Queue 2 (low priority)
8. Queue 1 (low priority)

A setting of 0 (zero %) on a queue means that no bandwidth minimum is specifically reserved for that queue for each of the ports in < port-list >. Also, there is no benefit to setting the high-priority queue (queue 8) to 0 (zero) unless you want the medium queue (queue 4) to be able to support traffic bursts above its guaranteed minimum.

(continued)

***Notes:** Configuring 0% for a queue can result in that queue being starved if any higher queue becomes over-subscribed and is then given all unused bandwidth.*

The switch applies the bandwidth calculation to the link speed the port is currently using. For example, if a 10/100 Mbs port negotiates to 10 Mbps on the link, then it bases its GMB calculations on 10 Mbps; not 100 Mbps.

*Use **show bandwidth output < port-list >** to display the current GMB configuration. (The **show config** and **show running** commands do not include GMB configuration data.)*

For example, suppose you wanted to configure the following outbound minimum bandwidth availability for ports A1 and A2:

Priority of Outbound Port Queue	Minimum Bandwidth %	Effect on Outbound Bandwidth Allocation
8	20	Queue 8 has the first priority use of all outbound bandwidth not specifically allocated to queues 1 - 7. If, for example, bandwidth allocated to queue 5 is not being used and queues 7 and 8 become oversubscribed, queue 8 has first-priority use of the unused bandwidth allocated to queue 5.
7	15	Queue 7 has a guaranteed minimum bandwidth of 15% available for outbound traffic. If queue 7 becomes oversubscribed and queue 8 is not already using all of the unallocated bandwidth, then queue 7 can use the unallocated bandwidth. Also, any unused bandwidth allocated to queues 6 to queue 1 is available to queue 7 if queue 8 has not already claimed it.
6	10	Queue 6 has a guaranteed minimum bandwidth of 10% and, if oversubscribed, is subordinate to queues 8 and 7 in priority for any unused outbound bandwidth available on the port.
5	10	Queue 5 has a guaranteed minimum bandwidth of 10% and, if oversubscribed, is subordinate to queues 8, 7, and 6 for any unused outbound bandwidth available on the port.
4	10%	Queue 4 has a guaranteed minimum bandwidth of 10% and, if oversubscribed, is subordinate to queues, 8, 7, 6, and 5 for any unused outbound bandwidth available on the port.
3	30%	Queue 3 has a guaranteed minimum bandwidth of 30% and, if oversubscribed, is subordinate to queues, 8, 7, 6, 5, and 4 for any unused outbound bandwidth available on the port.
2	3%	Queue 2 has a guaranteed minimum bandwidth of 3% and, if oversubscribed, is subordinate to queues, 8, 7, 6, 5, 4, and 3 for any unused outbound bandwidth available on the port.
1	2%	Queue 1 has a guaranteed minimum bandwidth of 2% and, if oversubscribed, is subordinate to all the other queues for any unused outbound bandwidth available on the port.

Either of the following commands configures ports A1 through A5 with bandwidth settings:

```
ProCurve(config)#int a1-a5 bandwidth-min output 2 3 30 10  
10 10 15 20
```

```
ProCurve(eth-A1-A5)#bandwidth-min output 2 3 30 10 10 10  
15 20
```

Displaying the Current Guaranteed Minimum Bandwidth Configuration

This command displays the per-port GMB configuration in the running-config file.

Syntax: show bandwidth output [*port-list*]

*Without [**port-list**], this command lists the GMB configuration for all ports on the switch. With [**port-list**], this command lists the GMB configuration for the specified ports. This command operates the same way in any CLI context. If the command lists **Disabled** for a port, there are no bandwidth minimums configured for any queue on the port. (Refer to the description of the **no** form of the bandwidth-min output command on page 13-22.)*

For example, to display the GMB configuration resulting from either of the above commands:

```
ProCurve(config)# show bandwidth output a1-a5
```

Outbound	Guaranteed Minimum Bandwidth %							
Port	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8
A1	2	3	30	10	10	10	15	20
A2	2	3	30	10	10	10	15	20
A3	2	3	30	10	10	10	15	20
A4	2	3	30	10	10	10	15	20
A5	2	3	30	10	10	10	15	20

User-Configured Minimum Bandwidth Settings

Figure 13-6. Example of Listing the Guaranteed Minimum Bandwidth Configuration

This is how the preceding listing of the GMB configuration would appear in the startup-config file.

```
ProCurve(config)# show config status
Running configuration is same as the startup configuration.
ProCurve(config)# show config

Startup configuration:

; J8697A Configuration Editor; Created on release #K.11.00

hostname "ProCurve"
module 1 type J8697A
snmp-server community "public" Unrestricted
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24
  ip address dhcp-bootp
  exit
interface A1
  bandwidth-min output 2 3 30 10 10 10 15 20
  exit
interface A2
  bandwidth-min output 2 3 30 10 10 10 15 20
  exit
interface A3
  bandwidth-min output 2 3 30 10 10 10 15 20
  exit
interface A4
  bandwidth-min output 2 3 30 10 10 10 15 20
  exit
interface A5
  bandwidth-min output 2 3 30 10 10 10 15 20
  exit
```

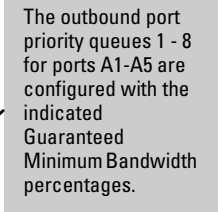


Figure 13-7. Example of GMB Settings Listed in the “show config” Output

GMB Operating Notes

Impact of QoS Queue Configuration on GMB commands. Changing the number of queues affects the GMB commands (**interface bandwidth-min** and **show bandwidth output**) to operate only on the number of queues currently configured. In addition, when the **qos queue-config** command is executed, any previously configured **bandwidth-min output** settings are removed from the startup configuration. Refer to Table 13-2 on page 13-21 for the default GMB percentage allocations per number of queues.

Jumbo Frames

Feature	Default	Menu	CLI	Web
display VLAN jumbo status	n/a	—	13-30	—
configure jumbo VLANs	Disabled	—	13-32	—

The *Maximum Transmission Unit* (MTU) is the maximum size IP frame the switch can receive for Layer 2 frames inbound on a port. The switch drops any inbound frames larger than the MTU allowed on the port. On ports operating at 10 Mbps or 100 Mbps, the MTU is fixed at 1522 bytes. However, ports operating at 1 Gbps or 10 Gbps speeds accept forward frames of up to 9220 bytes (including four bytes for a VLAN tag) when configured for jumbo traffic. You can enable inbound jumbo frames on a per-VLAN basis. That is, on a VLAN configured for jumbo traffic, all ports belonging to that VLAN and *operating* at 1 Gbps or 10 Gbps allow inbound jumbo frames of up to 9220 bytes. (Regardless of the mode configured on a given jumbo-enabled port, if the port is operating at only 10 Mbps or 100 Mbps, only frames that do not exceed 1522 bytes are allowed inbound on that port.)

Terminology

Jumbo Frame: An IP frame exceeding 1522 bytes in size. The maximum Jumbo frame size is 9220 bytes. (This size includes 4 bytes for the VLAN tag.)

Jumbo VLAN: A VLAN configured to allow inbound jumbo traffic. All ports belonging to a jumbo and operating at 1 Gbps or higher can receive jumbo frames from external devices. If the switch is in a meshed domain, then all meshed ports (operating at 1 Gbps or higher) on the switch will accept jumbo traffic from other devices in the mesh.

MTU (*Maximum Transmission Unit*): This is the maximum-size IP frame the switch can receive for Layer 2 frames inbound on a port. The switch allows jumbo frames of up to 9220 bytes.

Standard MTU: An IP frame of 1522 bytes in size. (This size includes 4 bytes for the VLAN tag.)

Operating Rules

- **Required Port Speed:** This feature allows inbound and outbound jumbo frames on ports operating at speeds of 1 gigabit or higher. At lower port speeds, only standard (1522-byte or smaller) frames are allowed, regardless of the jumbo configuration.
- **Switch Meshing:** If you enable jumbo traffic on a VLAN, then all meshed ports on the switch will be enabled to support jumbo traffic. (On a given meshed switch, every meshed port operating at 1 Gbps or higher becomes a member of every VLAN configured on the switch.)
- **GVRP Operation:** A VLAN enabled for jumbo traffic cannot be used to create a dynamic VLAN. A port belonging to a statically configured, jumbo-enabled VLAN cannot join a dynamic VLAN.
- **Port Adds and Moves:** If you add a port to a VLAN that is already configured for jumbo traffic, the switch enables that port to receive jumbo traffic. If you remove a port from a jumbo-enabled VLAN, the switch disables jumbo traffic capability on the port only if the port is not currently a member of another jumbo-enabled VLAN. This same operation applies to port trunks.
- **Jumbo Traffic Sources:** A port belonging to a jumbo-enabled VLAN can receive inbound jumbo frames through any VLAN to which it belongs, including non-jumbo VLANs. For example, if VLAN 10 (without jumbos enabled) and VLAN 20 (with jumbos enabled) are both configured on a switch, and port 1 belongs to both VLANs, then port 1 can receive jumbo traffic from devices on either VLAN. For a method to allow only some ports in a VLAN to receive jumbo traffic, refer to “Configuring a Maximum Frame Size” on page 13-32.

Configuring Jumbo Frame Operation

Command	Page
show vlans	13-30
show vlans ports < port-list >	13-31
show vlans < vid >	13-32
jumbo	13-32
jumbo max-frame-size	13-32

Overview

1. Determine the VLAN membership of the ports or trunks through which you want the switch to accept inbound jumbo traffic. For operation with GVRP enabled, refer to the GVRP topic under “Operating Rules”, above.
2. Ensure that the ports through which you want the switch to receive jumbo frames are operating at least at gigabit speed. (Check the **Mode** field in the output for the **show interfaces brief < port-list >** command.)
3. Use the **jumbo** command to enable jumbo frames on one or more VLANs statically configured in the switch. (All ports belonging to a jumbo-enabled VLAN can receive jumbo frames.)
4. Execute **write memory** to save your configuration changes to the startup-config file.

Viewing the Current Jumbo Configuration

Syntax: show vlans

*Lists the static VLANs configured on the switch and includes a **Jumbo** column to indicate which VLANs are configured to support inbound jumbo traffic. All ports belonging to a jumbo-enabled VLAN can receive jumbo traffic. (For more information refer to “Configuring a Maximum Frame Size” on page 13-32.) See Figure 13-8, below.*

```
ProCurve(config)# show vlans
```

Status and Counters - VLAN Information

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

802.1Q	VLAN ID	Name	Status	Voice	Jumbo
1		DEFAULT_VLAN	Port-based	No	Yes
5		VLAN5	Port-based	No	No
22		VLAN22	Port-based	No	No

Indicates which static VLANs are configured to enable jumbo frames.

Figure 13-8. Example Listing of Static VLANs To Show Jumbo Status Per VLAN

Syntax: show vlans ports < port-list >

*Lists the static VLANs to which the specified port(s) belong, including the **Jumbo** column to indicate which VLANs are configured to support jumbo traffic. Entering only one port in < port-list > results in a list of all VLANs to which that port belongs. Entering multiple ports in < port-list > results in a superset list that includes the VLAN memberships of all ports in the list, even though the individual ports in the list may belong to different subsets of the complete VLAN listing. For example, if port 1 belongs to VLAN 1, port 2 belongs to VLAN 10, and port 3 belongs to VLAN 15, then executing this command with a < port-list > of **1-3** results in a listing of all three VLANs, even though none of the ports belong to all three VLANs. (Refer to Figure 13-9.)*

```
ProCurve# show vlans ports 1-3
```

Status and Counters - VLAN Information - for ports 1-3

802.1Q VLAN ID	Name	Status	Voice	Jumbo
1	DEFAULT_VLAN	Port-based	No	Yes
10	VLAN10	Port-based	No	No
15	VLAN15	Port-based	No	No

Indicates which static VLANs are configured to enable jumbo frames.

Figure 13-9. Example of Listing the VLAN Memberships for a Range of Ports

Syntax: show vlans < vid >

This command shows port membership and jumbo configuration for the specified < vid >.

```
ProCurve(config)# show vlan 100
```

Status and Counters - VLAN Information - Ports - VLAN 100

802.1Q VLAN ID : 100
Name : VLAN100
Status : Port-based
Voice : No
Jumbo : No

Port	Information Mode	Unknown	VLAN	Status
1	Tagged	Learn		Up
2	Tagged	Learn		Up
3	Tagged	Learn		Up
4	Tagged	Learn		Down
5	Tagged	Learn		Up

Lists the ports belonging to VLAN 100 and whether the VLAN is enabled for jumbo frame traffic.

Figure 13-10. Example of Listing the Port Membership and Jumbo Status for a VLAN

Enabling or Disabling Jumbo Traffic on a VLAN

Syntax: vlan < vid > jumbo
[no] vlan < vid > jumbo

*Configures the specified VLAN to allow jumbo frames on all ports on the switch that belong to that VLAN. If the VLAN is not already configured on the switch, **vlan < vid > jumbo** also creates the VLAN. Note that a port belonging to one jumbo VLAN can receive jumbo frames through any other VLAN statically configured on the switch, regardless of whether the other VLAN is enabled for jumbo frames. The **[no]** form of the command disables inbound jumbo traffic on all ports in the specified VLAN that do not also belong to another VLAN that is enabled for jumbo traffic. In a VLAN context, the command forms are **jumbo** and **no jumbo**. (Default: Jumbos disabled on the specified VLAN.)*

Configuring a Maximum Frame Size

You can globally set a maximum frame size for Jumbo frames that will support values from 1518 bytes to 9216 bytes for untagged frames.

Syntax: jumbo max-frame-size <size>

Sets the maximum frame size for Jumbo frames. The range is from 1518 bytes to 9216 bytes.

Note: The **jumbo max-frame-size** is set on a **GLOBAL** level.

Default: 9216 bytes

Configuring IP MTU

Note

The following feature is available on the switches covered in this guide. Jumbos support is required. On switches that do not support this command, the IP MTU value is derived from the maximum frame size and is not configurable.

You can set the IP MTU globally by entering this command. The value of **max-frame-size** must be greater than or equal to 18 bytes more than the value selected for **ip-mtu**. For example, if **ip-mtu** is set to 8964, the **max-frame-size** is configured as 8982.

Syntax: jumbo ip-mtu <size>

*Globally sets the IP MTU size. Values range between 1500 and 9198 bytes. This value must be 18 bytes less than the value of **max-frame-size**.*

Default: 9198 bytes

SNMP Implementation

Jumbo Maximum Frame Size.

The maximum frame size for Jumbos is supported with the following proprietary MIB object:

hpSwitchMaxFrameSize OBJECT-TYPE

This is the value of the global **max-frame-size** supported by the switch. The default value is set to 9216 bytes.

Jumbo IP MTU.

The IP MTU for Jumbos is supported with the following proprietary MIB object:

hpSwitchIpMTU OBJECT-TYPE

This is the value of the global Jumbos IP MTU (or L3 MTU) supported by the switch. The default value is set to 9198 bytes (a value that is 18 bytes less than the largest possible maximum frame size of 9216 bytes). This object can only be used in switches which support **max-frame-size** and **ip-mtu** configuration.

Displaying the Maximum Frame Size

Use the **show jumbos** command to display the globally configured untagged maximum frame size for the switch.

```
ProCurve(config)# show jumbos

Jumbos Global Values

Configured : MaxFrameSize : 9216      Ip-MTU : 9198
In Use     : MaxFrameSize : 9216      Ip-MTU : 9198
```

Figure 14. Displaying the Maximum Frame Size and IP MTU Values

Operating Notes for Maximum Frame Size

- When you set a maximum frame size for Jumbo frames, it must be on a global level. You cannot use the **jumbo max-frame-size** command on a per-port or per-VLAN basis.
- The original way to configure Jumbo frames remains the same, which is per-VLAN, but you cannot set a maximum frame size per-VLAN.
- Jumbo support must be enabled for a VLAN from the CLI or through SNMP.
- Setting the maximum frame size does not require a reboot.
- When you upgrade to a version of software that supports setting the maximum frame size from a version that did not, the **max-frame-size** value is set automatically to 9216 bytes.
- Configuring a Jumbo maximum frame size on a VLAN allows frames up to **max-frame-size** even though other VLANs of which the port is a member are not enabled for Jumbo support.

Operating Notes for Jumbo Traffic-Handling

- ProCurve does not recommend configuring a voice VLAN to accept jumbo frames. Voice VLAN frames are typically small, and allowing a voice VLAN to accept jumbo frame traffic can degrade the voice transmission performance.
- You can configure the default, primary, and/or (if configured) the management VLAN to accept jumbo frames on all ports belonging to the VLAN.

- When the switch applies the default MTU (1522-bytes) to a VLAN, all ports in the VLAN can receive incoming frames of up to 1522 bytes in length. When the switch applies the jumbo MTU (9220 bytes) to a VLAN, all ports in that VLAN can receive incoming frames of up to 9220 bytes in length. A port receiving frames exceeding the applicable MTU drops such frames, causing the switch to generate an Event Log message and increment the “Giant Rx” counter (displayed by **show interfaces < port-list >**).
- The switch allows flow control and jumbo frame capability to co-exist on a port.
- The default MTU is 1522 bytes (including 4 bytes for the VLAN tag). The jumbo MTU is 9220 bytes (including 4 bytes for the VLAN tag).
- When a port is not a member of any jumbo-enabled VLAN, it drops all jumbo traffic. If the port is receiving “excessive” inbound jumbo traffic, the port generates an Event Log message to notify you of this condition. This same condition generates a Fault-Finder message in the Alert log of the switch’s web browser interface, and also increments the switch’s “Giant Rx” counter.
- If you do not want all ports in a given VLAN to accept jumbo frames, you can consider creating one or more jumbo VLANs with a membership comprised of only the ports you want to receive jumbo traffic. Because a port belonging to one jumbo-enabled VLAN can receive jumbo frames through any VLAN to which it belongs, this method enables you to include both jumbo-enabled and non-jumbo ports within the same VLAN. For example, suppose you wanted to allow inbound jumbo frames only on ports 6, 7, 12, and 13. However, these ports are spread across VLAN 100 and VLAN 200, and also share these VLANs with other ports you want excluded from jumbo traffic. A solution is to create a third VLAN with the sole purpose of enabling jumbo traffic on the desired ports, while leaving the other ports on the switch disabled for jumbo traffic. That is:

	VLAN 100	VLAN 200	VLAN 300
Ports	6-10	11-15	6, 7, 12, and 13
Jumbo-Enabled?	No	No	Yes

If there are security concerns with grouping the ports as shown for VLAN 300, you can either use source-port filtering to block unwanted traffic paths or create separate jumbo VLANs, one for ports 6 and 7, and another for ports 12 and 13.

- **Outbound Jumbo Traffic.** Any port operating at 1 Gbps or higher can transmit outbound jumbo frames through any VLAN, regardless of the jumbo configuration. The VLAN is not required to be jumbo-enabled, and the port is not required to belong to any other, jumbo enabled VLANs. This

can occur in situations where a non-jumbo VLAN includes some ports that do not belong to another, jumbo-enabled VLAN and some ports that do belong to another, jumbo-enabled VLAN. In this case, ports capable of receiving jumbo frames can forward them to the ports in the VLAN that do not have jumbo capability.

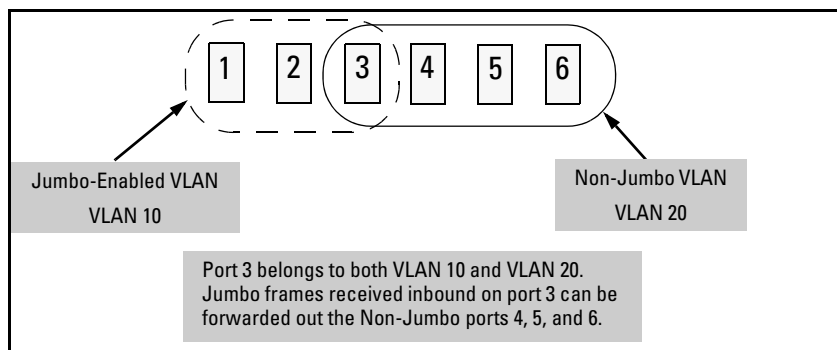


Figure 13-11. Forwarding Jumbo Frames Through Non-Jumbo Ports

Jumbo frames can also be forwarded out non-jumbo ports when the jumbo frames received inbound on a jumbo-enabled VLAN are routed to another, non-jumbo VLAN for outbound transmission on ports that have no memberships in other, jumbo-capable VLANs. Where either of the above scenarios is a possibility, the downstream device must be configured to accept the jumbo traffic. Otherwise, this traffic will be dropped by the downstream device.

- **Jumbo Traffic in a Switch Mesh Domain.** Note that if a switch belongs to a meshed domain, but does not have any VLANs configured to support jumbo traffic, then the meshed ports on that switch will drop any jumbo frames they receive from other devices. In this regard, if a mesh domain includes any ProCurve 1600M/2400M/2424M/4000M/8000M switches along with the switches covered in this guide configured to support jumbo traffic, only the switches covered in this guide will receive jumbo frames. The other switch models in the mesh will drop such frames. For more information on switch meshing, refer to the chapter titled “Switch Meshing” in the *Advanced Traffic Management Guide* for your switch.

Troubleshooting

A VLAN is configured to allow jumbo frames, but one or more ports drops all inbound jumbo frames. The port may not be operating at 1 gigabit or higher. Regardless of a port's configuration, if it is actually operating at a speed lower than 1 gigabit, it drops inbound jumbo frames. For example, if a port is configured for **Auto** mode (**speed-duplex auto**), but has negotiated a 100 Mbps speed with the device at the other end of the link, then the port cannot receive inbound jumbo frames. To determine the actual operating speed of one or more ports, view the **Mode** field in the output for the following command:

```
show interfaces brief < port-list >
```

A non-jumbo port is generating “Excessive undersize/giant frames” messages in the Event Log. The switches can transmit outbound jumbo traffic on any port, regardless of whether the port belongs to a jumbo VLAN. In this case, another port in the same VLAN on the switch may be jumbo-enabled through membership in a different, jumbo-enabled VLAN, and may be forwarding jumbo frames received on the jumbo VLAN to non-jumbo ports. Refer to “Outbound Jumbo Traffic” on page 13-35.

Port Traffic Controls
Jumbo Frames