HP 5820X & 5800 Switch Series
ACL and QoS

Configuration Guide

Abstract

This document describes the software features for the HP 5820X & 5800 Series products and guides you through the software configuration procedures. These configuration guides also provide configuration examples to help you apply software features to different network scenarios.

This documentation is intended for network planners, field technical support and servicing engineers, and network administrators working with the HP 5820X & 5800 Series products.

Part number: 5998-1618
Software version: Release 1211
Document version: 6W101-20121123

Legal and notice information

The information contained herein is subject to change without notice.

HEWLETT-PACKARD COMPANY MAKES NO WARRANTY OF ANY KIND WITH REGARD TO THIS MATERIAL, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Contents

Configuring ACL

Unless otherwise stated, ACLs refer to both IPv4 and IPv6 ACLs throughout this document.

The Layer 3 Ethernet interface in this document refers to the Ethernet port that can perform IP routing and inter-VLAN routing. You can set an Ethernet port as a Layer 3 Ethernet interface by using **port link-mode route** command (see *Layer 2—LAN Switching Configuration Guide*).

Overview

An ACL is a set of rules (or permit or deny statements) for identifying traffic based on criteria, such as source IP address, destination IP address, and port number.

ACLs are primarily used for packet filtering. A packet filter drops packets that match a deny rule and permits packets that match a permit rule. ACLs are also used by many modules, QoS and IP routing, for example, for traffic classification and identification.

Applications on the switch

An ACL is implemented in hardware or software, depending on the module that uses it. If the module, the packet filter or QoS module for example, is implemented in hardware, the ACL is applied to hardware to process traffic. If the module, the routing, or user interface access control module (Telnet, SNMP, or web) for example, is implemented in software, the ACL is applied to software to process traffic.

The user interface access control module denies packets that do not match any ACL. Some modules, QoS for example, ignore the permit or deny action in ACL rules and do not base their drop or forwarding decisions on the action set in ACL rules. See the specified module for information about ACL application.

ACL categories

Category	ACL number	IP version	Match criteria
Basic ACLs	2000 to 2999	IPv4	Source IPv4 address
		IPv6	Source IPv6 address
Advanced ACLs	3000 to 3999	IPv4	Source IPv4 address, destination IPv4 address, protocols over IPv4, and other Layer 3 and Layer 4 header fields
		IPv6	Source IPv6 address, destination IPv6 address, protocols over IPv6, and other Layer 3 and Layer 4 header fields
Ethernet frame header ACLs	4000 to 4999	IPv4 and IPv6	Layer 2 header fields, such as source and destination MAC addresses, 802.1p priority, and link layer protocol type

Numbering and naming ACLs

Each ACL category has a unique range of ACL numbers. When creating an ACL, you must assign it a number. In addition, assign the ACL a name for ease of identification. After creating an ACL with a name, you cannot rename it or delete its name.

For an Ethernet frame header ACL, the ACL number and name must be globally unique. For an IPv4 basic or advanced ACLs, its ACL number and name must be unique among all IPv4 ACLs, and for an IPv6 basic or advanced ACL, among all IPv6 ACLs. You can assign an IPv4 ACL and an IPv6 ACL the same number and name.

Match order

The rules in an ACL are sorted in a specific order. When a packet matches a rule, the device stops the matching process and performs the action defined in the rule. If an ACL contains overlapping or conflicting rules, the matching result and action to take depends on the rule order.

The following ACL match orders are available:

- **config**—Sorts ACL rules in ascending order of rule ID. A rule with a lower ID is matched before a rule with a higher ID. If you use this approach, carefully check the rules and their order.
- **auto**—Sorts ACL rules in depth-first order. Depth-first ordering ensures that any subset of a rule is always matched before the rule. Table 1 lists the sequence of tie breakers that depth-first ordering uses to sort rules for each type of ACL.

A wildcard mask, also called an "inverse mask," is a 32-bit binary and represented in dotted decimal notation. In contrast to a network mask, the 0 bits in a wildcard mask represent "do care" bits, and the 1 bits represent "don't care" bits. If the "do care" bits in an IP address are identical to the "do care" bits in an IP address criterion, the IP address matches the criterion. All "don't care" bits are ignored. The 0s and 1s in a wildcard mask can be noncontiguous. For example, 0.255.0.255 is a valid wildcard mask.

Table 1 Sort ACL rules in depth-first order

ACL category	Sequence of tie breakers
IPv4 basic ACL	**1.** VPN instance **2.** More 0s in the source IP address wildcard (more 0s means a narrower IP address range) **3.** Smaller rule ID
IPv4 advanced ACL	**1.** VPN instance **2.** Specific protocol type rather than IP (IP represents any protocol over IP) **3.** More 0s in the source IP address wildcard mask **4.** More 0s in the destination IP address wildcard **5.** Narrower TCP/UDP service port number range **6.** Smaller rule ID
IPv6 basic ACL	**1.** VPN instance **2.** Longer prefix for the source IP address (a longer prefix means a narrower IP address range) **3.** Smaller rule ID

ACL category	Sequence of tie breakers
IPv6 advanced ACL	1. VPN instance 2. Specific protocol type rather than IP (IP represents any protocol over IPv6) 3. Longer prefix for the source IPv6 address 4. Longer prefix for the destination IPv6 address 5. Narrower TCP/UDP service port number range 6. Smaller rule ID
Ethernet frame header ACL	1. More 1s in the source MAC address mask (more 1s means a smaller MAC address) 2. More 1s in the destination MAC address mask 3. Smaller rule ID

Rule comments and rule range remarks

Add a comment about an ACL rule to make it easy to understand. The rule comment appears below the rule statement.

In addition, add a rule range remark to indicate the start or end of a range of rules created for the same purpose. A rule range remark always appears above the specified ACL rule. If the specified rule has not been created yet, the position of the remark in the ACL is as follows:

- If the match order is **config**, the remark is inserted into the ACL in descending order of rule ID.
- If the match order is **auto**, the remark is placed at the end of the ACL. After you create the rule, the remark appears above the rule.

For more information about how to use rule range remarks, see the **rule remark** command in *ACL and QoS Command Reference* for your switch.

Rule numbering

ACL rules can be manually numbered or automatically numbered. This section describes how automatic ACL rule numbering works.

Rule numbering step

If you do not assign an ID to the rule you are creating, the system automatically assigns it a rule ID. The rule numbering step sets the increment by which the system automatically numbers rules. For example, the default ACL rule numbering step is 5. If you do not assign IDs to rules you are creating, they are automatically numbered 0, 5, 10, 15, and so on. The wider the numbering step, the more rules you can insert between two rules.

By introducing a gap between rules rather than contiguously numbering rules, you have the flexibility of inserting rules in an ACL. This feature is important for a configuration order ACL, where ACL rules are matched in ascending order of rule ID.

Automatic rule numbering and renumbering

The ID automatically assigned to an ACL rule takes the nearest higher multiple of the numbering step to the current highest rule ID, starting with 0.

For example, if the numbering step is 5 (the default), and there are five ACL rules numbered 0, 5, 9, 10, and 12, the newly defined rule is numbered 15. If the ACL does not contain any rules, the first rule is numbered 0.

Whenever the step changes, the rules are renumbered, starting from 0. For example, if there are five rules numbered 5, 10, 13, 15, and 20, changing the step from 5 to 2 causes the rules to be renumbered 0, 2, 4, 6, and 8.

Implementing time-based ACL rules

You can implement ACL rules based on the time of day by applying a time range to them. A time-based ACL rule only takes effect in the time periods specified by the time range.

The following basic types of time range are available:

- **Periodic time range**—Recurs periodically on a day or days of the week.
- **Absolute time range**—Only represents a period of time and does not recur.

You can specify a time range in ACL rules before or after you create it. However, the rules using the time range take effect only after you define the time range.

Fragments filtering with IPv4 ACLs

To prevent attackers from fabricating fragments, an ACL packet filter on the switch by default matches all fragments. To improve efficiency, configure the **fragment** keyword to only apply an IPv4 ACL rule to non-first fragments.

Configuration task list

Complete the following tasks to configure an IPv4 ACL:

Task	Remarks
Configuring a time range	Optional
Configuring an IPv4 basic ACL	Required Configure at least one task
Configuring an IPv6 basic ACL	
Configuring an IPv4 advanced ACL	
Configuring an IPv6 advanced ACL	
Configuring an Ethernet frame header ACL	
Copying an ACL	Optional
Packet filtering with ACLs	Optional

Configuring a time range

You can create up to 256 time ranges, each having a maximum of 32 periodic statements and 12 absolute statements. If a time range has multiple statements, its active period is calculated as follows:

1. Combining all periodic statements.
2. Combining all absolute statements.
3. Taking the intersection of the two statement sets as the active period of the time range.

To configure a time range:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Configure a time range.	**time-range** *time-range-name* { *start-time* **to** *end-time days* [**from** *time1 date1*] [**to** *time2 date2*] | **from** *time1 date1* [**to** *time2 date2*] | **to** *time2 date2* }	Required. By default, no time range exists. Repeat this command with the same time range name to create multiple statements for a time range.

Configuring a basic ACL

Configuring an IPv4 basic ACL

IPv4 basic ACLs match packets based only on the source IP address.

To configure an IPv4 basic ACL:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Create an IPv4 basic ACL and enter its view.	**acl number** *acl-number* [**name** *acl-name*] [**match-order** { **auto** | **config** }]	Required. By default, no ACL exists. IPv4 basic ACLs are numbered ranging from 2000 to 2999. Use the **acl name** *acl-name* command to enter the view of a named IPv4 ACL.
3. Configure a description for the IPv4 basic ACL.	**description** *text*	Optional. By default, an IPv4 basic ACL has no ACL description.
4. Set the rule numbering step.	**step** *step-value*	Optional. 5 by default.
5. Create or edit a rule.	**rule** [*rule-id*] { **deny** | **permit** } [**counting** | **fragment** | **logging** | **source** { *sour-addr sour-wildcard* | **any** } | **time-range** *time-range-name* | **vpn-instance** *vpn-instance-name*] *	Required. By default, an IPv4 basic ACL does not contain any rules. To create or edit multiple rules, repeat this step. If the ACL is for QoS traffic classification, do not specify the **vpn-instance** keyword. This keyword can cause ACL application failure. The **logging** and **counting** keywords (even if specified) do not take effect for QoS policies.

To do…	Use the command…	Remarks
6. Add or edit a rule comment.	**rule** *rule-id* **comment** *text*	Optional. By default, an IPv4 ACL rule has no rule comment.
7. Add or edit a rule range remark.	**rule** [*rule-id*] **remark** *text*	Optional. By default, no rule range remarks are configured.
8. Enable counting ACL rule matches performed in hardware.	**hardware-count enable**	Optional. Disabled by default. When the ACL is referenced by a QoS policy, this command does not take effect.

Configuring an IPv6 basic ACL

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Create an IPv6 basic ACL view and enter its view.	**acl ipv6 number** *acl6-number* [**name** *acl6-name*] [**match-order** { **auto** | **config** }]	Required. By default, no ACL exists. IPv6 basic ACLs are numbered ranging from 2000 to 2999. Use the **acl ipv6 name** *acl6-name* command to enter the view of a named IPv6 ACL.
3. Configure a description for the IPv6 basic ACL.	**description** *text*	Optional. By default, an IPv6 basic ACL has no ACL description.
4. Set the rule numbering step.	**step** *step-value*	Optional. 5 by default.
5. Create or edit a rule.	**rule** [*rule-id*] { **deny** | **permit** } [**counting** | **fragment** | **logging** | **source** { *ipv6-address prefix-length* | *ipv6-address/prefix-length* | **any** } | **time-range** *time-range-name*] *	Required. By default, an IPv6 basic ACL does not contain any rules. To create or edit multiple rules, repeat this step. If the ACL is for QoS traffic classification, do not specify the **fragment** keyword. This keyword can cause ACL application failure. The **logging** and **counting** keywords (even if specified) do not take effect for QoS.
6. Configure or edit a rule comment.	**rule** *rule-id* **comment** *text*	Optional. By default, an IPv6 basic ACL rule has no rule comment.

To do…	Use the command…	Remarks
7. Add or edit a rule range remark.	**rule** [*rule-id*] **remark** *text*	Optional. By default, no rule range remarks are configured.
8. Enable counting ACL rule matches performed in hardware.	**hardware-count enable**	Optional. Disabled by default. When the ACL is referenced by a QoS policy, this command does not take effect.

Configuring an advanced ACL

Configuring an IPv4 advanced ACL

IPv4 advanced ACLs match packets based on source IP addresses, destination IP addresses, protocols over IP, and other protocol header information, such as TCP/UDP source and destination port numbers, TCP flags, ICMP message types, and ICMP message codes.

IPv4 advanced ACLs also allow you to filter packets based on these priority criteria: ToS, IP precedence, and DSCP priority.

Compared to IPv4 basic ACLs, IPv4 advanced ACLs allow for more flexible and accurate filtering.

To configure an IPv4 advanced ACL:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Create an IPv4 advanced ACL and enter its view.	**acl number** *acl-number* [**name** *acl-name*] [**match-order** { **auto** | **config** }]	Required. By default, no ACL exists. IPv4 advanced ACLs are numbered ranging from 3000 to 3999. Use the **acl name** *acl-name* command to enter the view of a named IPv4 ACL.
3. Configure a description for the IPv4 advanced ACL.	**description** *text*	Optional. By default, an IPv4 advanced ACL has no ACL description.
4. Set the rule numbering step.	**step** *step-value*	Optional. 5 by default.

To do…	Use the command…	Remarks
5. Create or edit a rule.	**rule** [*rule-id*] { **deny** | **permit** } *protocol* [{ { **ack** *ack-value* | **fin** *fin-value* | **psh** *psh-value* | **rst** *rst-value* | **syn** *syn-value* | **urg** *urg-value* } * | **established** } | **counting** | **destination** { *dest-addr dest-wildcard* | **any** } | **destination-port** *operator port1* [*port2*] | **dscp** *dscp* | **fragment** | **icmp-type** { *icmp-type icmp-code* | *icmp-message* } | **logging** | **precedence** *precedence* | **reflective** | **source** { *sour-addr sour-wildcard* | **any** } | **source-port** *operator port1* [*port2*] | **time-range** *time-range-name* | **tos** *tos* | **vpn-instance** *vpn-instance-name*] *	Required. By default, an IPv4 advanced ACL does not contain any rules. To create or edit multiple rules, repeat this step. The **reflective** keyword is not supported. If the ACL is for packet filtering, the *operator* argument cannot be **neq**. If the ACL is for QoS traffic classification, do not specify the **vpn-instance** keyword or specify **neq** for the *operator* argument. The keywords can cause ACL application failure. The **logging** and **counting** keywords (even if specified) do not take effect for QoS.
6. Configure or edit a rule comment.	**rule** *rule-id* **comment** *text*	Optional. By default, an IPv4 advanced ACL rule has no rule comment.
7. Add or edit a rule range remark.	**rule** [*rule-id*] **remark** *text*	Optional. By default, no rule range remarks are configured.
8. Enable counting ACL rule matches performed in hardware.	**hardware-count enable**	Optional. Disabled by default. When the ACL is referenced by a QoS policy, this command does not take effect.

Configuring an IPv6 advanced ACL

IPv6 advanced ACLs match packets based on the source IPv6 addresses, destination IPv6 addresses, protocols carried over IPv6, and other protocol header fields such as the TCP/UDP source port number, TCP/UDP destination port number, ICMPv6 message type, and ICMPv6 message code.

Compared to IPv6 basic ACLs, IPv6 advanced ACLs allow for more flexible and accurate filtering.

To configure an IPv6 advanced ACL:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Create an IPv6 advanced ACL and enter its view.	**acl ipv6 number** *acl6-number* [**name** *acl6-name*] [**match-order** { **auto** | **config** }]	Required. By default, no ACL exists. IPv6 advanced ACLs are numbered ranging from 3000 to 3999. Use the **acl ipv6 name** *acl6-name* command to enter the view of a named IPv6 ACL.
3. Configure a description for the IPv6 advanced ACL.	**description** *text*	Optional. By default, an IPv6 advanced ACL has no ACL description.
4. Set the rule numbering step.	**step** *step-value*	Optional. 5 by default.
5. Create or edit a rule.	**rule** [*rule-id*] { **deny** | **permit** } *protocol* [{ { **ack** *ack-value* | **fin** *fin-value* | **psh** *psh-value* | **rst** *rst-value* | **syn** *syn-value* | **urg** *urg-value* } * | **established** } | **counting** | **destination** { *dest dest-prefix* | *dest/dest-prefix* | **any** } | **destination-port** *operator port1* [*port2*] | **dscp** *dscp* | **flow-label** *flow-label-value* | **fragment** | **icmp6-type** { *icmp6-type icmp6-code* | *icmp6-message* } | **logging** | **source** { *source source-prefix* | *source/source-prefix* **|** **any** } | **source-port** *operator port1* [*port2*] | **time-range** *time-range-name*] *	Required. By default IPv6 advanced ACL does not contain any rules. To create or edit multiple rules, repeat this step. If the ACL is for packet filtering, the *operator* argument cannot be **neq**. If the ACL is for QoS traffic classification, do not specify the **fragment** keyword or specify **neq** for the *operator* argument. The keywords can cause ACL application failure. The **logging** and **counting** keywords (even if specified) do not take effect for QoS.
6. Configure or edit a rule comment.	**rule** *rule-id* **comment** *text*	Optional. By default, an IPv6 advanced ACL rule has no rule comment.
7. Add or edit a rule range remark.	**rule** [*rule-id*] **remark** *text*	Optional. By default, no rule range remarks are configured.
8. Enable counting ACL rule matches performed in hardware.	**hardware-count enable**	Optional. Disabled by default. When the ACL is referenced by a QoS policy, this command does not take effect.

Configuring an Ethernet frame header ACL

Ethernet frame header ACLs, also called "Layer 2 ACLs," match packets based on Layer 2 protocol header fields, such as source MAC address, destination MAC address, 802.1p priority (VLAN priority), and link layer protocol type.

To configure an Ethernet frame header ACL:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Create an Ethernet frame header ACL and enter its view.	**acl number** *acl-number* [**name** *acl-name*] [**match-order** { **auto** | **config** }]	Required. By default, no ACL exists. Ethernet frame header ACLs are numbered ranging from 4000 to 4999. Use the **acl name** *acl-name* command to enter the view of a named Ethernet frame header ACL.
3. Configure a description for the Ethernet frame header ACL.	**description** *text*	Optional. By default, an Ethernet frame header ACL has no ACL description.
4. Set the rule numbering step.	**step** *step-value*	Optional. 5 by default.
5. Create or edit a rule.	**rule** [*rule-id*] { **deny** | **permit** } [**cos** *vlan-pri* | **counting** | **dest-mac** *dest-addr dest-mask* | { **lsap** *lsap-type lsap-type-mask* | **type** *protocol-type protocol-type-mask* } | **source-mac** *sour-addr source-mask* | **time-range** *time-range-name*] *	Required. By default, an Ethernet frame header ACL does not contain any rules. To create or edit multiple rules, repeat this step.
6. Configure or edit a rule comment.	**rule** *rule-id* **comment** *text*	Optional. By default, an Ethernet frame header ACL rule has no rule comment.
7. Add or edit a rule range remark.	**rule** [*rule-id*] **remark** *text*	Optional. By default, no rule range remarks are configured.
8. Enable counting ACL rule matches performed in hardware.	**hardware-count enable**	Optional. Disabled by default. When the ACL is referenced by a QoS policy, this command does not take effect.

Copying an ACL

Create an ACL by copying an existing ACL (source ACL). The new ACL (destination ACL) has the same properties and content as the source ACL, but not the same ACL number and name.

To copy an IPv4 or IPv6 ACL successfully, make sure that:

- The destination ACL number is from the same category as the source ACL number.
- The source IPv4 or IPv6 ACL already exists, but the destination IPv4 or IPv6 ACL does not.

Copying an IPv4 ACL

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Copy an existing IPv4 ACL to create a new IPv4 ACL.	**acl copy** { *source-acl-number* | **name** *source-acl-name* } **to** { *dest-acl-number* | **name** *dest-acl-name* }	Required

Copying an IPv6 ACL

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Copy an existing IPv6 ACL to generate a new one of the same category.	**acl ipv6 copy** { *source-acl6-number* | **name** *source-acl6-name* } **to** { *dest-acl6-number* | **name** *dest-acl6-name* }	Required

Packet filtering with ACLs

Use an ACL to filter incoming or outgoing IPv4 or IPv6 packets. With a basic or advanced ACL, you can log filtering events by specifying the **logging** keyword in the ACL rules and enabling the counting function.

You can set the packet filter to periodically send packet filtering logs to the information center as informational messages. The interval for generating and outputting packet filtering logs is configurable. The log information includes the number of matching packets and the ACL rules used in an interval. For more information about the information center, see *Network Management and Monitoring Configuration Guide*.

NOTE:

- An ACL on a VLAN interface filters packets forwarded at Layer 3 and multicast packets forwarded at Layer 2 in the VLAN.
- The packet filter does not support ACLs that have a **vpn-instance** criterion.

Applying an IPv4 ACL for packet filtering

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Enter Ethernet interface view or VLAN interface view.	**interface** *interface-type interface-number*	—
3. Apply an IPv4 ACL to the interface to filter IPv4 packets.	**packet-filter** { *acl-number* | **name** *acl-name* } { **inbound** | **outbound** }	Required. By default, no IPv4 ACL is applied to the interface.
4. Exit to system view.	**quit**	—
5. Set the interval for generating and outputting IPv4 packet filtering logs.	**acl logging frequence** *frequence*	Required. By default, the interval is 0. No IPv4 packet filtering logs are generated.

Applying an IPv6 ACL for packet filtering

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Enter Ethernet interface view or VLAN interface view.	**interface** *interface-type interface-number*	—
3. Apply an IPv6 ACL to the interface to filter IPv6 packets.	**packet-filter ipv6** { *acl6-number* | **name** *acl6-name* } { **inbound** | **outbound** }	Required. By default, no IPv6 ACL is applied to the interface.
4. Exit to system view.	**quit**	—
5. Set the interval for generating and outputting IPv6 packet filtering logs.	**acl ipv6 logging frequence** *frequence*	Required. By default, the interval is 0. No IPv6 packet filtering logs are generated.

Displaying and maintaining ACLs

Task	Command	Remarks
Display configuration and match statistics for one or all IPv4 ACLs.	**display acl** { *acl-number* | **all** | **name** *acl-name* } [**slot** *slot-number*] [| { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display configuration and match statistics for one or all IPv6 ACLs.	**display acl ipv6** { *acl6-number* | **all** | **name** *acl6-name* } [**slot** *slot-number*] [| { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the usage of ACL resources.	**display acl resource** [**slot** *slot-number*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view

Task	Command	Remarks
Display the application status of packet filtering ACLs on interfaces.	**display packet-filter** { { **all** | **interface** *interface-type interface-number* } [**inbound** | **outbound**] | **interface vlan-interface** *vlan-interface-number* [**inbound** | **outbound**] [**slot** *slot-number*] } [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the configuration and status of one or all time ranges.	**display time-range** { *time-range-name* | **all** } [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Clear statistics on one or all IPv4 ACLs.	**reset acl counter** { *acl-number* | **all** | **name** *acl-name* }	Available in user view
Clear statistics on one or all IPv6 basic and advanced ACLs.	**reset acl ipv6 counter** { *acl6-number* | **all** | **name** *acl6-name* }	Available in user view

Configuration examples

IPv4 packet filtering configuration example

Network requirements

As shown in Figure 1, apply an ACL to the inbound direction of interface GigabitEthernet 1/0/1 on Device A so that every day from 08:00 to 18:00 the interface only allows packets sourced from Host A to pass. Configure Device A to output packet filtering logs to the console at 10-minute intervals.

Figure 1 Network diagram for applying an IPv4 ACL to an interface for packet filtering

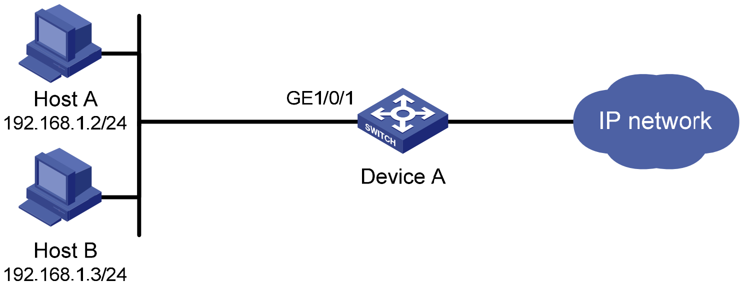

Configuration procedure

# Create a time range from 08:00 to 18:00 every day.

```
<DeviceA> system-view
[DeviceA] time-range study 8:00 to 18:00 daily
```

# Create IPv4 ACL 2009, and configure two rules in the ACL. One rule permits packets sourced from Host A at 192.168.1.2 and the other rule denies packets sourced from any other host during the time range **study**. Enable logging for both rules.

```
[DeviceA] acl number 2009
[DeviceA-acl-basic-2009] rule permit source 192.168.1.2 0 time-range study logging
[DeviceA-acl-basic-2009] rule deny source any time-range study logging
[DeviceA-acl-basic-2009] quit
```

# Apply IPv4 ACL 2009 to filter incoming packets on GigabitEthernet 1/0/1.

```
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] packet-filter 2009 inbound
[DeviceA-GigabitEthernet1/0/1] quit
```

# Enable the device to generate and output IPv4 packet filtering logs at 10-minute intervals.

```
[DeviceA] acl logging frequence 10
```

# Configure the device to output informational log messages to the console.

```
[DeviceA] info-center source default channel 0 log level informational
```

IPv6 packet filtering configuration example

Network requirements

As shown in Figure 2, apply an IPv6 ACL to the incoming traffic of GigabitEthernet 1/0/1 on Device A so that every day from 08:00 to 18:00 the interface only allows packets from Host A to pass through. Configure Device A to output packet filtering logs to the console at 10-minute intervals.

Figure 2 Network diagram for applying an IPv6 ACL to an interface for packet filtering

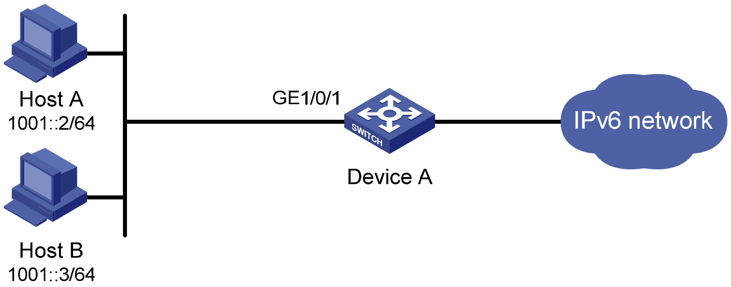

Configuration procedure

# Create a time range from 08:00 to 18:00 every day.

```
<DeviceA> system-view
[DeviceA] time-range study 8:0 to 18:0 daily
```

# Create IPv6 ACL 2009 and configure two rules for the ACL. One rule permits packets sourced from Host A with the IPv6 address 1001::2 and the other rule denies packets sourced from any other host during the time range **study**. Enable logging for both rules.

```
[DeviceA] acl ipv6 number 2009
[DeviceA-acl6-basic-2009] rule permit source 1001::2 128 time-range study logging
[DeviceA-acl6-basic-2009] rule deny source any time-range study logging
[DeviceA-acl6-basic-2009] quit
```

# Apply IPv6 ACL 2009 to filter incoming packets on GigabitEthernet 1/0/1.

```
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] packet-filter ipv6 2009 inbound
[DeviceA-GigabitEthernet1/0/1] quit
```

# Configure the device to collect and output IPv6 packet filtering logs at 10-minute intervals.

```
 [DeviceA] acl ipv6 logging frequence 10
```

# Configure the device to output informational log messages to the console.

```
[DeviceA] info-center source default channel 0 log level informational
```

QoS overview

In data communications, QoS is a network's ability to provide differentiated service guarantees for diversified traffic in terms of bandwidth, delay, jitter, and drop rate.

Network resources are scarce. The contention for resources requires that QoS prioritize important traffic flows over trivial ones. For example, in the case of fixed bandwidth, if a traffic flow gets more bandwidth, the other traffic flows gets less bandwidth and may be affected. When making a QoS scheme, you must consider the characteristics of various applications to balance the interests of diversified users and to utilize network resources.

The following sections describe some typical QoS service models and widely used mature QoS techniques.

Service models

The Best-effort service model, InServ model, and DiffServ model are the most common QoS service models.

Best-effort service model

The best-effort model is a single-service model and also the simplest service model. In this model, the network does its best to deliver packets, but does not guarantee delivery or control delay.

The best-effort service model uses the FIFO queuing mechanism and is the basic model for the Internet.

IntServ model

The IntServ model is a multiple-service model that can accommodate diverse QoS requirements. This service model provides the most granularly differentiated QoS as it identifies and guarantees definite QoS for each data flow.

In the IntServ model, an application must request service from the network before it sends data. IntServ signals the service request with the RSVP. All nodes receiving the request reserve resources as requested and maintain state information for the application flow.

The IntServ model demands high storage and processing capabilities because it requires all nodes along the transmission path to maintain resource state information for each flow. This model is suitable for small-sized or edge networks, but not large-sized networks, for example, the core layer of the Internet, where billions of flows are present.

For more information about RSVP, see *MPLS Configuration Guide*.

DiffServ model

The DiffServ model is a multiple-service model that can satisfy diverse QoS requirements. It is easy to implement and extend. DiffServ does not signal the network to reserve resources before sending data, as IntServ does.

All QoS techniques in this document are based on the DiffServ model.

Techniques

The QoS techniques include traffic classification, traffic policing, traffic shaping, line rate, congestion management, and congestion avoidance. They address problems that arise at different positions of a network.

Figure 3 Placement of QoS techniques in a network

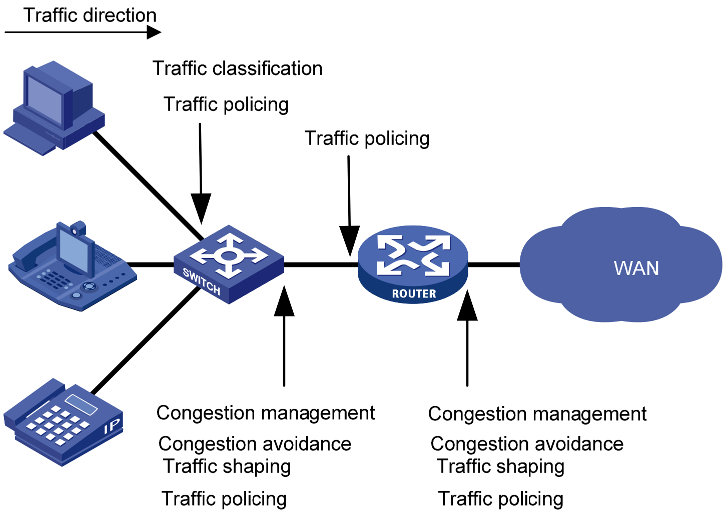

As shown in Figure 3, traffic classification, traffic shaping, traffic policing, congestion management, and congestion avoidance mainly implement the following functions:

- Traffic classification uses certain match criteria to assign packets with the same characteristics to a class. You can provide differentiated services based on classes.
- Traffic policing polices flows entering or leaving a device, and imposes penalties on traffic flows that exceed the preset threshold to prevent aggressive use of network resources. You can apply traffic policing to both incoming and outgoing traffic of a port.
- To eliminate packet drops, traffic shaping proactively adapts the output rate of traffic to the network resources available on the downstream device. Traffic shaping usually applies to the outgoing traffic of a port.
- Congestion management provides a resource scheduling policy to determine the packet forwarding sequence when congestion occurs. Congestion management usually applies to the outgoing traffic of a port.
- Congestion avoidance monitors the network resource usage, and is usually applied to the outgoing traffic of a port. When congestion worsens, congestion avoidance reduces the queue length by dropping packets.

Figure 4 provides an overview of how the QoS module processes traffic.

Figure 4 QoS processing flow

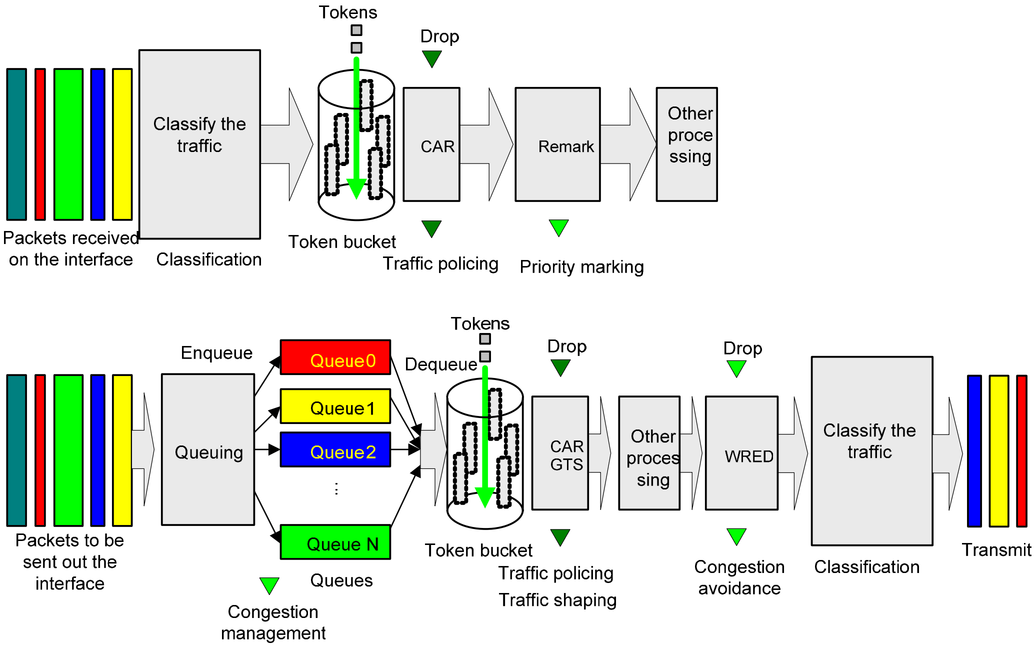

Configuration approaches

You can configure QoS in the following approaches:

- MQC approach
- Non-MQC approach

Some features support both approaches, but some support only one.

MQC approach

In modular QoS configuration (MQC) approach, you configure QoS service parameters by using QoS policies (see "Configuring a QoS policy").

Non-MQC approach

In the non-MQC approach, you configure QoS service parameters without using a QoS policy. For example, use the line rate feature to set a rate limit on an interface without using a QoS policy.

Configuring a QoS policy

A QoS policy is a set of class-behavior associations and defines the shaping, policing, or other QoS actions to take on different classes of traffic.

A class is a set of match criteria for identifying traffic and it uses the AND or OR operator:

- **AND**—A packet must match all criteria to match the class.
- **OR**—A packet matches the class if it matches any of the criteria in the class.

A traffic behavior defines a set of QoS actions to take on packets, such as priority marking and redirect.

By associating a traffic behavior with a class in a QoS policy, you apply the specific set of QoS actions to the class of traffic.

Figure 5 shows the procedure of configuring and using a QoS policy.

Figure 5 QoS policy configuration procedure

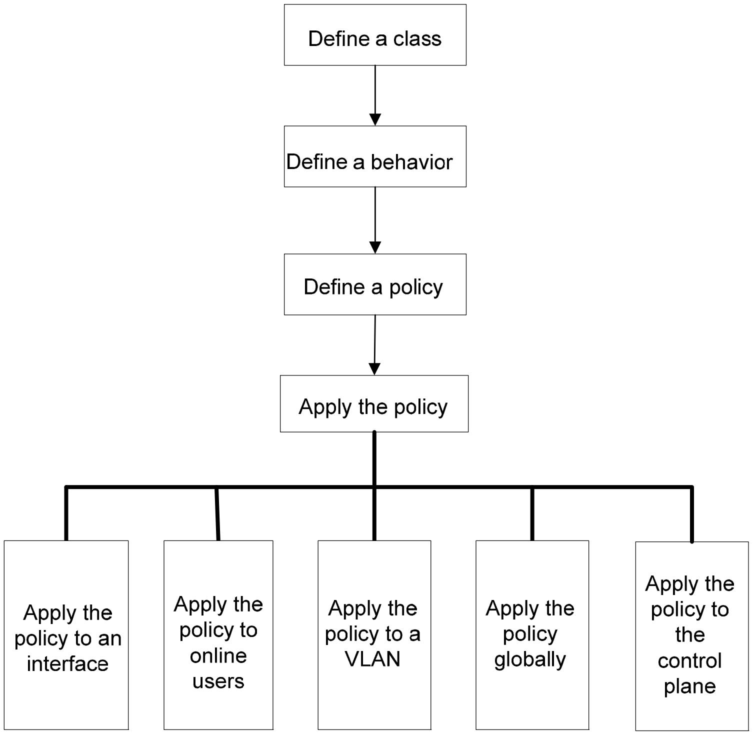
Define a class
Define a behavior
Define a policy
Apply the policy
Apply the policy to an interface
Apply the policy to online users
Apply the policy to a VLAN
Apply the policy globally
Apply the policy to the control plane

Defining a class

To define a class, specify its name and then configure the match criteria in class view.

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Create a class and enter class view.	**traffic classifier** *tcl-name* [**operator** { **and** | **or** }]	Required. By default, the operator of a class is AND. The operator of a class can be AND or OR. • **AND**—A packet is only assigned to a class when the packet matches all criteria in the class. • **OR**—A packet is assigned to a class if it matches any of the criteria in the class
3. Configure match criteria.	**if-match** *match-criteria*	Required. For more information, see *ACL and QoS Command Reference*.

match-criteria—Match criterion. Table 2 shows the value range for the *match-criteria* argument.

Table 2 Value range for the *match-criteria* argument

Option	Description
acl [**ipv6**] { *acl-number* | **name** *acl-name* }	Matches an ACL. The *acl-number* argument ranges from 2000 to 4999 for an IPv4 ACL, and 2000 to 3999 for an IPv6 ACL. The *acl-name* argument is a case-insensitive string of 1 to 32 characters, which must start with an English letter from a to z or A to Z, and to avoid confusion, cannot be **all**.
any	Matches all packets.
customer-dot1p *8021p-list*	Matches the 802.1p priority of the customer network. The *8021p-list* argument is a list of up to eight 802.1p priority values. An 802.1p priority ranging from 0 to 7.
customer-vlan-id { *vlan-id-list* | *vlan-id1* **to** *vlan-id2* }	Matches the VLAN IDs of customer networks. The *vlan-id-list* argument is a list of up to eight VLAN IDs. The *vlan-id1* **to** *vlan-id2* specifies a VLAN ID range, where the *vlan-id1* must be smaller than the *vlan-id2*. A VLAN ID ranging from 1 to 4094.
destination-mac *mac-address*	Matches a destination MAC address.
dscp *dscp-list*	Matches DSCP values. The *dscp-list* argument is a list of up to eight DSCP values. A DSCP value can be a number from 0 to 63 or any keyword in Table 13.
ip-precedence *ip-precedence-list*	Matches IP precedence. The *ip-precedence-list* argument is a list of up to eight IP precedence values. An IP precedence ranges from 0 to 7.

Option	Description
protocol *protocol-name*	Matches a protocol. The *protocol-name* argument can be IP or IPv6.
qos-local-id *local-id-value*	Matches a local QoS ID, which ranges from 1 to 4095. The local QoS IDs supported on the 5800 Switch Series and the 5820X Switch Series are from 1 to 3999.
service-dot1p *8021p-list*	Matches the 802.1p priority of the service provider network. The *8021p-list* argument is a list of up to eight 802.1p priority values. An 802.1p priority ranges from 0 to 7.
service-vlan-id { *vlan-id-list* | *vlan-id1* **to** *vlan-id2* }	Matches the VLAN IDs of ISP networks. The *vlan-id-list* is a list of up to eight VLAN IDs. The *vlan-id1* **to** *vlan-id2* specifies a VLAN ID range, where the *vlan-id1* must be smaller than the *vlan-id2*. A VLAN ID ranges from 1 to 4094.
source-mac *mac-address*	Matches a source MAC address.
system-index *index-value-list*	Matches a pre-defined match criterion (system-index) for packets sent to the control plane. The *index-value-list* argument specifies a list of up to eight system indexes. The system index ranges from 1 to 128.

In a class using the **AND** operator, when using the **if-match** command to define match criteria, if multiple match criteria are used consider the following:

- If the **acl** or **acl ipv6** keyword specified are defined in a class, the relationship between these match criteria is **or** when the policy is applied.
- If the **customer-vlan-id** or **service-vlan-id** keyword specified are defined in a class, the relationship between these match criteria is **or**.

To successfully execute the traffic behavior associated with a traffic class that uses the AND operator, define only one **if-match** clause for any of the following match criteria and enter only one value for any of the following *list* arguments, for example, the *8021p-list* argument:

- **customer-dot1p** *8021p-list*
- **destination-mac** *mac-address*
- **dscp** dscp-list
- **ip-precedence** ip-precedence-list
- **service-dot1p** *8021p-list*
- **source-mac** *mac-address*
- **system-index** index-value-list

To create multiple if-match clauses or specify multiple values for a list argument for any of the match criteria previously listed, use the **OR** operator.

A QoS policy referencing an **if match customer-dot1p** clause cannot be applied to the outgoing traffic.

Defining a traffic behavior

A traffic behavior is a set of QoS actions (such as traffic filtering, shaping, policing, and priority marking) to take on a class of traffic. To define a traffic behavior, first create it and then configure QoS actions, such as priority marking and traffic redirecting, in traffic behavior view.

To define a traffic behavior:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Create a traffic behavior and enter traffic behavior view.	**traffic behavior** *behavior-name*	Required.
3. Configure actions in the traffic behavior.	See the following chapters, depending on the purpose of the traffic behavior: traffic policing, traffic filtering, traffic redirecting, priority marking, traffic accounting, and so on.	

Associating a class with behavior in a policy

You associate a behavior with a class in a QoS policy to perform the actions defined in the behavior for the class of packets.

To associate a class with a behavior in a policy:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Create a policy and enter policy view.	**qos policy** *policy-name*	Required.
3. Associate a class with a behavior in the policy.	**classifier** *tcl-name* **behavior** *behavior-name* [**mode** { **dcbx** | **dot1q-tag-manipulation** }]	Required. Repeat this step to create more class-behavior associations.

If an ACL is referenced by a QoS policy for defining traffic match criteria, packets matching the ACL are organized as a class and the behavior defined in the QoS policy applies to the class regardless of whether the match mode of the if-match clause is deny or permit.

In a QoS policy with multiple class-to-traffic-behavior associations, if the action of creating an outer VLAN tag, setting customer network VLAN ID, or setting service provider network VLAN ID is configured in a traffic behavior, do not configure any other action in this traffic behavior; otherwise, the QoS policy may not function as expected after it is applied. For more information about the action of setting customer network VLAN ID or service provider network VLAN ID, see *Layer 2—LAN Switching Configuration Guide*.

The **mode dcbx** keyword specifies that the class-behavior association is for the DCBX purposes. For more information about DCBX, see *Layer 2—LAN Switching Configuration Guide*.

The **do1q-tag-manipulation** keyword applies to only many-to-one VLAN mapping configuration. For more information about many-to-one VLAN mapping, see *Layer 2—LAN Switching Configuration Guide*.

Applying the QoS policy

Apply a QoS policy to the following occasions:

- **An interface**—The policy takes effect on the traffic sent or received on the interface.
- **A user profile**—The policy takes effect on the traffic sent or received by the online users of the user profile.
- **A VLAN**—The policy takes effect on the traffic sent or received on all ports in the VLAN.
- **Globally**—The policy takes effect on the traffic sent or received on all ports.
- **Control plane**—The policy takes effect on the traffic received on the control plane.

NOTE:

A port QoS policy takes priority over a global QoS policy.

- You can modify classes, behaviors, and class-behavior associations in a QoS policy applied to an interface, VLAN, or inactive user profile, or globally. If a class references an ACL for traffic classification, you can delete or modify the ACL (such as add rules to, delete rules from, and modify rules of the ACL).

Applying the QoS policy to an interface

A policy can be applied to multiple interfaces, but only one policy can be applied in one direction (inbound or outbound) of an interface.

Both bridge mode (Layer 2) and route mode (Layer 3) Ethernet ports support QoS policies. The term "interface" in this section collectively refers to these types of ports. Use **port link-mode** command to set an Ethernet port to operate in bridge or route mode (see *Layer 2—LAN Switching Configuration Guide*).

To apply the QoS policy to an interface:

To do…		Use the command…	Remarks
1. Enter system view.		**system-view**	—
2. Enter interface view or port group view:	Enter interface view	**interface** *interface-type interface-number*	Use either command. Settings in interface view take effect on the current interface. Settings in port group view take effect on all ports in the port group.
	Enter port group view	**port-group manual** *port-group-name*	
3. Apply the policy to the interface or port group.		**qos apply policy** *policy-name* { **inbound** | **outbound** }	Required.

NOTE:

The QoS policy applied to the outgoing traffic on an interface does not regulate local packets, which are critical protocol packets sent by the local system for maintaining the normal operation of the device. To avoid dropping local packets, QoS does not process them. The most common used local packets include: link maintenance packets, IS-IS packets, OSPF packets, RIP packets, BGP packets, LDP packets, RSVP packets, and SSH packets.

Applying the QoS policy to online users

You can apply a QoS policy to multiple online users. In one direction of each online user, only one policy can be applied. To modify a QoS policy already applied in a certain direction, remove the QoS policy application first.

To apply the QoS policy to online users:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Enter user profile view.	**user-profile** *profile-name*	Required. The configuration made in user profile view takes effect when the user-profile is activated and the users of the user profile are online. For more information about user profiles, see *Security Configuration Guide*.
3. Apply the QoS policy.	**qos apply policy** *policy-name* { **inbound** | **outbound** }	Required. Use the **inbound** keyword to apply the QoS policy to the incoming traffic of the switch (traffic sent by the online users). Use the **outbound** keyword to apply the QoS policy to the outgoing traffic (traffic received by the online users).
4. Return to system view.	**quit**	—
5. Activate the user profile.	**user-profile** *profile-name* **enable**	Required. Inactive by default.

NOTE:

If a user profile is active, the QoS policy applied to it cannot be configured or removed, except ACLs referenced in the QoS policy. However, when the users of the user profile are online, the referenced ACLs cannot be modified either.

The QoS policy applied to a user profile only supports the **remark**, **car**, and **filter** actions.

- Do not apply a null policy to a user profile. The user profile using a null policy cannot be activated.

Applying the QoS policy to a VLAN

You can apply a QoS policy to a VLAN to regulate traffic of the VLAN.

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Apply the QoS policy to VLANs.	**qos vlan-policy** *policy-name* **vlan** *vlan-id-list* { **inbound** | **outbound** }	Required

NOTE:

QoS policies cannot be applied to dynamic VLANs, such as VLANs created by GVRP.

Do not apply a QoS policy to a VLAN and the ports in the VLAN at the same time.

Applying the QoS policy globally

You can apply a QoS policy globally to the inbound or outbound direction of all ports.

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Apply the QoS policy globally.	**qos apply policy** *policy-name* **global** { **inbound** | **outbound** }	Required

Applying the QoS policy to the control plane

A device provides the data plane and the control plane:

- The data plane has, units responsible for receiving, transmitting, and switching (forwarding) packets, such as various dedicated forwarding chips. They deliver super processing speeds and throughput.
- The control plane has processing units running most routing and switching protocols and is responsible for protocol packet resolution and calculation, such as CPUs. Compared with data plane units, the control plane units allow for great packet processing flexibility but have lower throughput.
- When the data plane receives packets that it cannot recognize or process, it transmits them to the control plane. If the transmission rate exceeds the processing capability of the control plane, which very likely occurs at times of DoS attacks, the control plane is busy handling undesired packets and fails to handle legitimate packets correctly or timely. As a result, protocol performance is affected.
- To address this problem, apply a QoS policy to the control plane to take QoS actions, such as traffic filtering or rate limiting, on inbound traffic. This action ensures that the control plane can receive, transmit, and process packets properly.

To the control plane:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Enter control plane view.	**control-plane slot** *slot-number*	Required
3. Apply the QoS policy to the control plane.	**qos apply policy** *policy-name* **inbound**	Required

△ CAUTION:

By default, some devices are configured with pre-defined control plane policies, which take effect on the control planes by default. A pre-defined control plane QoS policy uses the system-index to identify the type of packets sent to the control plane. You can reference system-indexes in **if-match** commands in class view for traffic classification and then re-configure traffic behaviors for these classes as required. Use **display qos policy control-plane pre-defined** command to display them.

In a QoS policy for control planes, if a system index classifier is configured, the associated traffic behavior can contain only the CAR or accounting action. In addition, if the CAR action is configured, only its CIR setting can be applied.

In the QoS policy for a control plane, if a system index classifier is not configured, the associated traffic behaviors also take effect on the data traffic of the card where the control plane resides.

Displaying and maintaining QoS policies

Task	Command	Remarks
Display traffic class configuration.	**display traffic classifier user-defined** [*tcl-name*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display traffic behavior configuration.	**display traffic behavior user-defined** [*behavior-name*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display user-defined QoS policy configuration.	**display qos policy user-defined** [*policy-name* [**classifier** *tcl-name*]] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display QoS policy configuration on the specified or all interfaces.	**display qos policy interface** [*interface-type interface-number*] [**inbound** | **outbound**] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display VLAN QoS policy configuration.	**display qos vlan-policy** { **name** *policy-name* | **vlan** *vlan-id* } [**slot** *slot-number*] [**inbound** | **outbound**] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display information about QoS policies applied globally.	**display qos policy global** [**slot** *slot-number*] [**inbound** | **outbound**] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display information about QoS policies applied to a control plane.	**display qos policy control-plane** [**slot** *slot-number*] [**inbound**] [| { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display information about pre-defined QoS policies applied to a control plane.	**display qos policy control-plane pre-defined** [**slot** *slot-number*] [| { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Clear VLAN QoS policy statistics.	**reset qos vlan-policy** [**vlan** *vlan-id*] [**inbound** | **outbound**]	Available in user view
Clear the statistics for a QoS policy applied globally.	**reset qos policy global** [**inbound** | **outbound**]	Available in user view
Clear the statistics for the QoS policy applied to a control plane.	**reset qos policy control-plane** [**slot** *slot-number*] [**inbound**]	Available in user view

Configuring priority mapping

Overview

When a packet enters a device, depending on your configuration, the device assigns a set of QoS priority parameters to the packet based on either a certain priority field carried in the packet or the port priority of the incoming port. This process is called "priority mapping." During this process, the device can modify the priority of the packet depending on the status of the device. The set of QoS priority parameters determines the scheduling priority and forwarding priority of the packet.

Priority mapping is implemented with priority mapping tables and involves priorities such as 802.1p priority, DSCP, EXP, IP precedence, local precedence, and drop precedence.

Types of priorities

Priorities fall into the following types: priorities carried in packets and priorities locally assigned only for scheduling.

The packet carried priorities include 802.1p priority, DSCP precedence, IP precedence, EXP, and so on. These priorities have global significance and affect the forwarding priority of packets across the network. For more information about these priorities, see "Appendix B Introduction to packet precedence."

The locally assigned priorities only have local significance. They are only assigned by the device for scheduling. These priorities include the local precedence and drop precedence, as follows:

- **Local precedence is used for queuing**—A local precedence value corresponds to an output queue. A packet with a higher local precedence is assigned to a higher priority output queue to be preferentially scheduled.
- **Drop precedence is used for making packet drop decisions**—Packets with the highest drop precedence are dropped preferentially.

Priority mapping tables

Priority mapping is implemented with priority mapping tables. By looking up a priority mapping table, the device determines priority value to assign to a packet for subsequent packet processing. The switch provides the following priority mapping tables:

- **dot1p-dp**—802.1p-to-drop priority mapping table.
- **dot1p-exp**—802.1p-to-EXP priority mapping table.
- **dot1p-lp**—802.1p-to-local priority mapping table.
- **dscp-dot1p**—DSCP-to-802.1p priority mapping table, which only applies to IP packets.
- **dscp-dp**—DSCP-to-drop priority mapping table, which only applies to IP packets.
- **dscp-dscp**—DSCP-to-DSCP priority mapping table, which only applies to IP packets.
- **exp-dot1p**—EXP-to-802.1p priority mapping table.
- **exp-dp**—EXP-to-drop priority mapping table.

The default priority mapping tables (see "Appendix A Default priority mapping tables") are available for priority mapping. In most cases, they are adequate for priority mapping. If a default priority mapping table cannot meet your requirements, you can modify the priority mapping table as required.

Priority trust mode on a port

The priority trust mode on a port determines which priority is used for priority mapping table lookup. Port priority was introduced so that you can use it for priority mapping in addition to the priority fields carried in packets. The HP 5800 Switch Series and 5820X Switch Series provide the following priority trust modes:

- **dot1p**—Uses the 802.1p priority carried in packets for priority mapping.

Table 3 Priority mapping results of trusting the 802.1p priority (when the default dot1p-lp priority mapping table is used)

802.1p priority carried in packets	Local precedence	Queue ID
0	2	2
1	0	0
2	1	1
3	3	3
4	4	4
5	5	5
6	6	6
7	7	7

NOTE:

When the 802.1p priority carried in packets is trusted, the port priority is used for priority mapping for packets which do not carry VLAN tags (namely, do not carry 802.1p priorities.) The priority mapping results are the same as not trusting packet priority, as shown in Table 5.

- **dscp**—Uses the DSCP carried in packets for priority mapping.

Table 4 Priority mapping results of trusting the DSCP (when the default dscp-dot1p and dot1p-lp priority mapping tables are used)

DSCP value carried in packets	Local precedence	Queue ID
0 to 7	0	0
8 to 15	1	1
16 to 23	2	2
24 to 31	3	3
32 to 39	4	4
40 to 47	5	5
48 to 55	6	6
56 to 63	7	7

- **undo qos trust**—Uses the port priority as the 802.1p priority for priority mapping. The port priority is user configurable.

Table 5 Priority mapping results of not trusting packet priority (when the default dot1p-lp priority mapping table is used)

Port priority	Local precedence	Queue ID
0 (default)	2	2
1	0	0
2	1	1
3	3	3
4	4	4
5	5	5
6	6	6
7	7	7

The priority mapping procedure varies with the priority modes. For more information, see "Priority mapping procedure."

Priority mapping procedure

On receiving an Ethernet packet on a port, the switch marks the scheduling priorities (local precedence and drop precedence) for the Ethernet packet. This procedure is done according to the priority trust mode of the receiving port and the 802.1q tagging status of the packet, as shown in Figure 6.

NOTE:

The priority mapping procedures shown in Figure 6 and Figure 7 apply in the absence of priority marking. If priority marking is configured, the switch performs priority marking before priority mapping. The switch then uses the remarked packet-carried priority for priority mapping or directly uses the remarked scheduling priority for traffic scheduling depending on your configuration. Neither priority trust mode configuration on the port nor port priority configuration takes effect.

Figure 6 Priority mapping procedure for an Ethernet packet

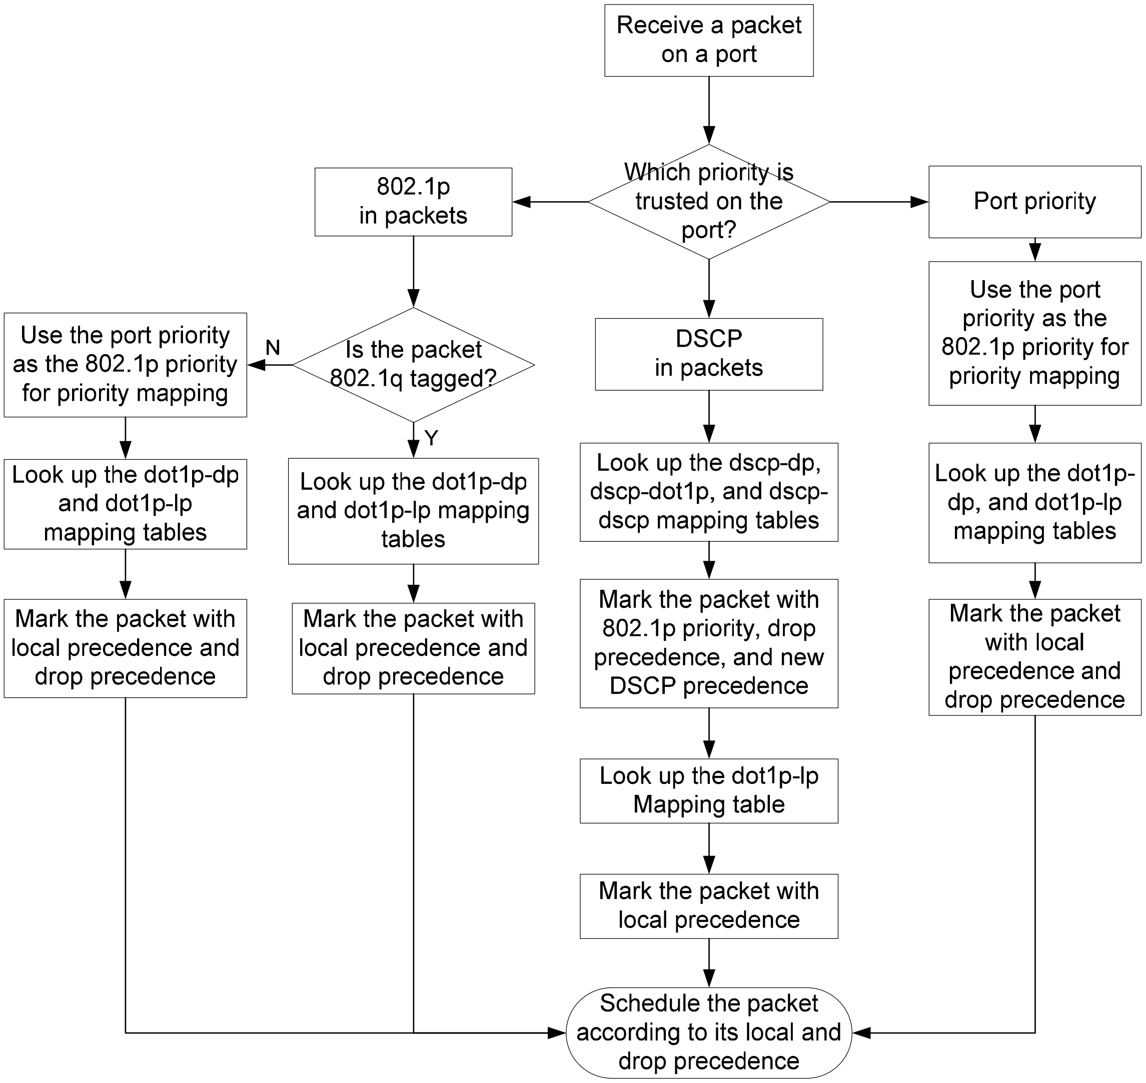
Receive a packet on a port
Which priority is trusted on the port?
802.1p in packets
DSCP in packets
Port priority
Is the packet 802.1q tagged?
N
Y
Use the port priority as the 802.1p priority for priority mapping
Look up the dot1p-dp and dot1p-lp mapping tables
Mark the packet with local precedence and drop precedence
Look up the dot1p-dp and dot1p-lp mapping tables
Mark the packet with local precedence and drop precedence
Look up the dscp-dp, dscp-dot1p, and dscp-dscp mapping tables
Mark the packet with 802.1p priority, drop precedence, and new DSCP precedence
Look up the dot1p-lp Mapping table
Mark the packet with local precedence
Use the port priority as the 802.1p priority for priority mapping
Look up the dot1p-dp, and dot1p-lp mapping tables
Mark the packet with local precedence and drop precedence
Schedule the packet according to its local and drop precedence

For an MPLS packet, the priority mapping procedure as shown in Figure 7 is adopted:

Figure 7 Priority mapping procedure for an MPLS packet

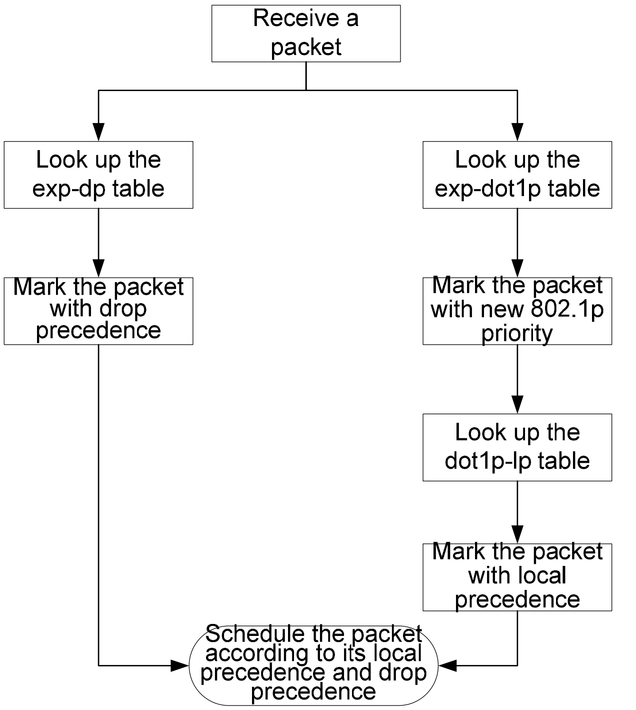
Receive a packet
Look up the exp-dp table
Look up the exp-dot1p table
Mark the packet with drop precedence
Mark the packet with new 802.1p priority
Look up the dot1p-lp table
Mark the packet with local precedence
Schedule the packet according to its local precedence and drop precedence

Configuration guidelines

Modify priority mappings by modifying priority mapping tables, priority trust mode on a port, and port priority.

HP recommends planning QoS throughout the network before making QoS configuration.

Both bridge mode (Layer 2) and route mode (Layer 3) Ethernet ports support the priority mapping function. The term "interface" in this chapter collectively refers to these types of ports. Use the **port link-mode** command to set an Ethernet port to operate in bridge or route mode (see *Layer 2—LAN Switching Configuration Guide*)

Configuring a priority mapping table

To configure a priority mapping table:

To do...	Use the command...	Remarks
1. Enter system view.	**system-view**	—
2. Enter priority mapping table view.	**qos map-table** { **dot1p-dp** | **dot1p-exp** | **dot1p-lp** | **dscp-dot1p** | **dscp-dp** | **dscp-dscp** | **exp-dot1p** | **exp-dp** }	Required.
3. Configure the priority mapping table.	**import** *import-value-list* **export** *export-value*	Required. Newly configured mappings overwrite the old ones.

Configuring a port to trust packet priority for priority mapping

When configuring the trusted packet priority type on an interface or port group, use the following priority trust modes:

- **dot1p**—Uses the 802.1p priority of received packets for mapping.
- **dscp**—Uses the DSCP precedence of received IP packets for mapping.
- **untrust**—Uses port priority as the 802.1p priority for priority mapping.

To configure the trusted packet priority type on an interface or port group:

To do…		Use the command…	Remarks
1. Enter system view.		**system-view**	—
2. Enter interface view or port group view:	Enter interface view.	**interface** *interface-type interface-number*	Use either command. Settings in interface view take effect on the current interface. Settings in port group view take effect on all ports in the port group.
	Enter port group view.	**port-group manual** *port-group-name*	
3. Configure the priority trust mode:	Trust the 802.1p or DSCP priority in packets.	**qos trust** { **dot1p** | **dscp** }	Use either command. By default, the device trusts the port priority.
	Trust the port priority.	**undo qos trust**	

Changing the port priority of an interface

To change the port priority of an interface:

To do…		Use the command…	Remarks
1. Enter system view		**system-view**	—
2. Enter interface view or port group view:	Enter interface view.	**interface** *interface-type interface-number*	Use either command. Settings in interface view take effect on the current interface. Settings in port group view take effect on all ports in the port group.
	Enter port group view.	**port-group manual** *port-group-name*	
3. Set the port priority of the interface.		**qos priority** *priority-value*	Required. The default port priority is 0.

Displaying priority mappings

To display priority mappings, complete the following tasks:

Task	Command	Remarks
Display priority mapping table configuration	**display qos map-table** [**dot1p-dp** | **dot1p-exp** | **dot1p-lp** | **dscp-dot1p** | **dscp-dp** | **dscp-dscp** | **exp-dot1p** | **exp-dp**] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display the trusted packet priority type on a port	**display qos trust interface** [*interface-type interface-number*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view

Priority mapping configuration examples

Priority trust mode and port priority configuration example

Network requirements

As shown in Figure 8, Device A is connected to GigabitEthernet 1/0/1 of Device C. Device B is connected to GigabitEthernet 1/0/2 of Device C, and the packets from Device A and Device B to Device C are not VLAN tagged.

Configure Device C to preferentially process packets from Device A to Server when GigabitEthernet 1/0/3 of Device C is congested.

Figure 8 Network diagram for priority trust mode configuration

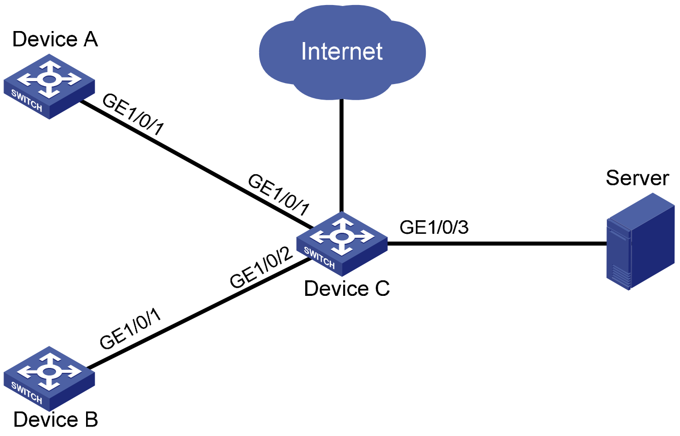

Configuration procedure

# Assign port priority to GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2. Make sure that the priority of GigabitEthernet 1/0/1 is higher than that of GigabitEthernet 1/0/2 and that no trusted packet priority type is configured on GigabitEthernet 1/0/1 or GigabitEthernet 1/0/2.

```
<DeviceC> system-view
[DeviceC] interface gigabitethernet 1/0/1
[DeviceC-GigabitEthernet1/0/1] qos priority 3
[DeviceC-GigabitEthernet1/0/1] quit
[DeviceC] interface gigabitethernet 1/0/2
[DeviceC-GigabitEthernet1/0/2] qos priority 1
[DeviceC-GigabitEthernet1/0/2] quit
```

Priority mapping table and priority marking configuration example

For more information about priority marking, see "Configuring priority marking."

Network requirements

As shown in Figure 9, the company's enterprise network interconnects all departments through Device. The network is described as follows:

- The marketing department connects to GigabitEthernet 1/0/1 of Device, which sets the 802.1p priority of traffic from the marketing department to 3.
- The R&D department connects to GigabitEthernet 1/0/2 of Device, which sets the 802.1p priority of traffic from the R&D department to 4.
- The management department connects to GigabitEthernet 1/0/3 of Device, which sets the 802.1p priority of traffic from the management department to 5.

Configure port priority, 802.1p-to-local priority mapping table, and priority marking to implement the plan as described in Table 6.

Table 6 Configuration plan

Traffic destination	Traffic priority order	Queuing plan		
		Traffic source	**Output queue**	**Queue priority**
Public servers	R&D department > management department > marketing department	R&D department	6	High
		Management department	4	Medium
		Marketing department	2	Low
Internet through HTTP	Management department > marketing department > R&D department	R&D department	2	Low
		Management department	6	High
		Marketing department	4	Medium

Figure 9 Network diagram for priority mapping table and priority marking configuration

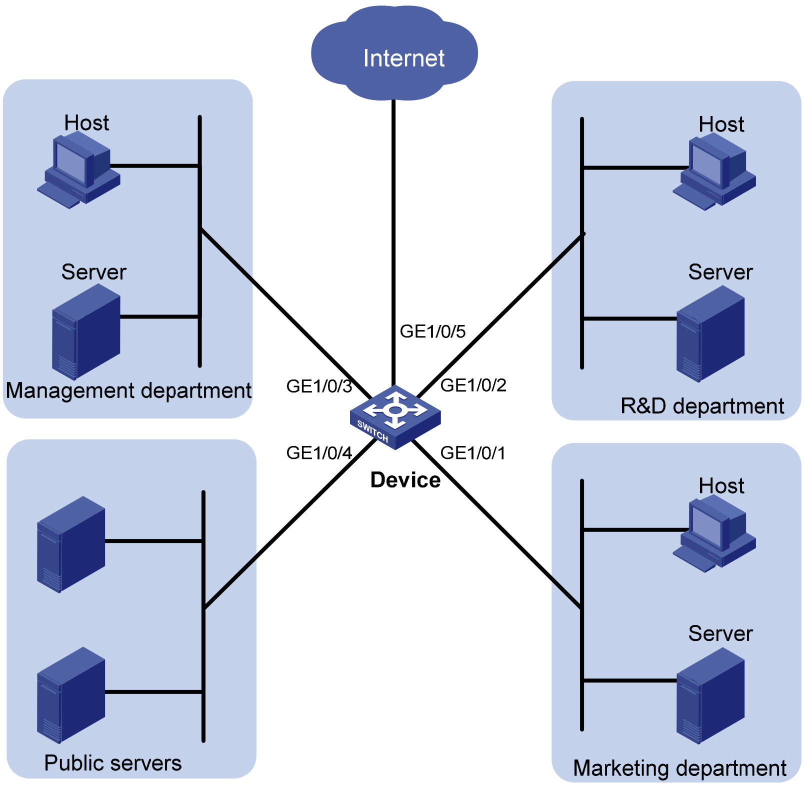

Configuration procedure

1. Configure the trusting port priority.

# Set the port priority of GigabitEthernet 1/0/1 to 3.

```
<Device> system-view
[Device] interface gigabitethernet 1/0/1
[Device-GigabitEthernet1/0/1] qos priority 3
[Device-GigabitEthernet1/0/1] quit
```

# Set the port priority of GigabitEthernet 1/0/2 to 4.

```
[Device] interface gigabitethernet 1/0/2
[Device-GigabitEthernet1/0/2] qos priority 4
[Device-GigabitEthernet1/0/2] quit
```

# Set the port priority of GigabitEthernet 1/0/3 to 5.

```
[Device] interface gigabitethernet 1/0/3
[Device-GigabitEthernet1/3] qos priority 5
[Device-GigabitEthernet1/3] quit
```

2. Configure the priority mapping table.

# Configure the 802.1p-to-local priority mapping table to map 802.1p priority values 3, 4, and 5 to local precedence values 2, 6, and 4.

```
[Device] qos map-table dot1p-lp
[Device-maptbl-dot1p-lp] import 3 export 2
[Device-maptbl-dot1p-lp] import 4 export 6
[Device-maptbl-dot1p-lp] import 5 export 4
[Device-maptbl-dot1p-lp] quit
```

3. Configure priority marking.

# Mark the HTTP traffic of the management department, marketing department, and R&D department to the Internet with 802.1p priorities 4, 5, and 3 respectively. Use the priority mapping table you have configured to map the 802.1p priorities to local precedence values 6, 4, and 2, respectively for differentiated traffic treatment.

# Create ACL 3000 to match HTTP traffic.

```
[Device] acl number 3000
[Device-acl-adv-3000] rule permit tcp destination-port eq 80
[Device-acl-adv-3000] quit
```

# Create class **http** and reference ACL 3000 in the class.

```
[Device] traffic classifier http
[Device-classifier-http] if-match acl 3000
[Device-classifier-http] quit
```

# Configure a priority marking policy for the management department and apply the policy to the incoming traffic of GigabitEthernet 1/0/3.

```
[Device] traffic behavior admin
[Device-behavior-admin] remark dot1p 4
[Device-behavior-admin] quit
[Device] qos policy admin
[Device-qospolicy-admin] classifier http behavior admin
[Device-qospolicy-admin] quit
[Device] interface gigabitethernet 1/0/3
[Device-GigabitEthernet1/0/3] qos apply policy admin inbound
```

# Configure a priority marking policy for the marketing department and apply the policy to the incoming traffic of GigabitEthernet 1/0/1.

```
[Device] traffic behavior market
[Device-behavior-market] remark dot1p 5
[Device-behavior-market] quit
[Device] qos policy market
[Device-qospolicy-market] classifier http behavior market
[Device-qospolicy-market] quit
[Device] interface gigabitethernet 1/0/1
[Device-GigabitEthernet1/0/1] qos apply policy market inbound
```

# Configure a priority marking policy for the R&D department and apply the policy to the incoming traffic of GigabitEthernet 1/0/2.

```
[Device] traffic behavior rd
[Device-behavior-rd] remark dot1p 3
[Device-behavior-rd] quit
[Device] qos policy rd
[Device-qospolicy-rd] classifier http behavior rd
[Device-qospolicy-rd] quit
[Device] interface gigabitethernet 1/0/2
[Device-GigabitEthernet1/0/2] qos apply policy rd inbound
```

Configuring traffic policing, traffic shaping, and line rate

Both bridge mode (Layer 2) and route mode (Layer 3) Ethernet ports support traffic policing, traffic shaping, and line rate. The term "interface" in this chapter collectively refers to these types of ports. Use the **port link-mode** command to set an Ethernet port to operate in bridge or route mode (see *Layer 2—LAN Switching Configuration Guide*).

Overview

Traffic policing, traffic shaping, and rate limit are QoS technologies that help assign network resources, such as assigning bandwidth. They increase network performance and user satisfaction. For example, you can configure a flow to only use the resources committed to it in a certain time range. This avoids network congestion caused by bursty traffic.

Traffic policing, GTS, and line rate limit the traffic rate and resource usage according to traffic specifications. Once a particular flow exceeds its specifications, such as assigned bandwidth, the flow is shaped or policed to make sure that it is under the specifications. You can use token buckets for evaluating traffic specifications.

Traffic evaluation and token buckets

A token bucket is analogous to a container that holds a certain number of tokens. Each token represents a certain forwarding capacity. The system puts tokens into the bucket at a constant rate. When the token bucket is full, the extra tokens cause the token bucket to overflow.

Token bucket features

Evaluating traffic with the token bucket

A token bucket mechanism evaluates traffic by looking at the number of tokens in the bucket. If the number of tokens in the bucket is enough for forwarding packets, the traffic conforms to the specification, and is called "conforming traffic." If the traffic does not conform to the specification it is called "excess traffic."

A token bucket has the following configurable parameters:

- Mean rate at which tokens are put into the bucket, which is the permitted average rate of traffic. It is usually set to the CIR.
- Burst size or the capacity of the token bucket. It is the maximum traffic size permitted in each burst. It is usually set to the CBS. The set burst size must be greater than the maximum packet size.

Each arriving packet is evaluated. In each evaluation, if the number of tokens in the bucket is enough, the traffic conforms to the specification and the tokens for forwarding the packet are taken away; if the number of tokens in the bucket is not enough, the traffic is excessive.

Complicated evaluation

Set two token buckets, bucket C and bucket E, to evaluate traffic in a more complicated environment and achieve more policing flexibility. Traffic policing uses the following parameters:

- **CIR**—Rate at which tokens are put into bucket C. It sets the average packet transmission or forwarding rate allowed by bucket C.
- **CBS**—Size of bucket C, which specifies the transient burst of traffic that bucket C can forward.
- **EBS**—Size of bucket E, which specifies the transient burst of traffic that bucket E can forward.
- **PIR**—The rate at which tokens are put into bucket E, which specifies the average packet transmission or forwarding rate allowed by bucket E.

CBS is implemented with bucket C, and EBS with bucket E. In each evaluation, packets are measured against the following bucket scenarios:

- If bucket C has enough tokens, packets are colored green.
- If bucket C does not have enough tokens but bucket E has enough tokens, packets are colored yellow.
- If neither bucket C nor bucket E has sufficient tokens, packets are colored red.

Traffic policing

Traffic policing polices inbound and outbound traffic. The outbound traffic is taken for example.

A typical application of traffic policing is to supervise the specification of certain traffic entering a network and limit it within a reasonable range, or to "discipline" the extra traffic to prevent aggressive use of network resources by a certain application. For example, you can limit bandwidth for HTTP packets to less than 50% of the total. If the traffic of a certain session exceeds the limit, traffic policing can drop the packets or reset the IP precedence of the packets. Figure 10 shows an example of policing outbound traffic on an interface.

Figure 10 Schematic diagram for traffic policing

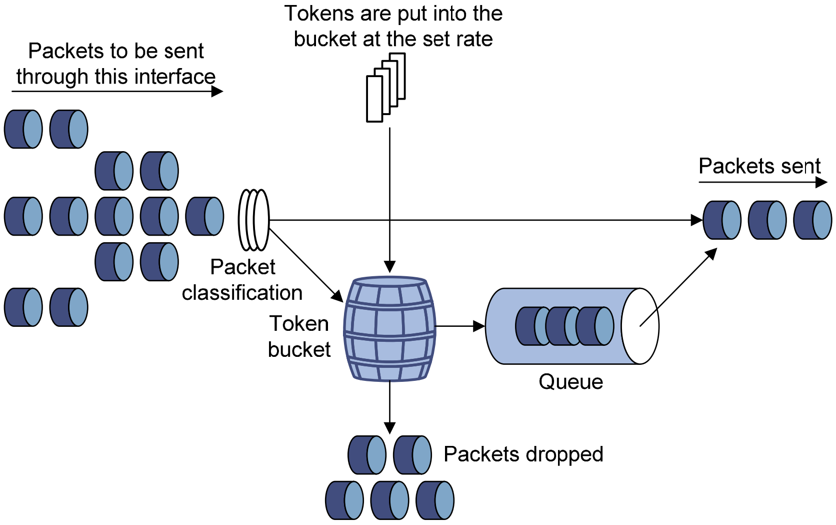
Tokens are put into the bucket at the set rate
Packets to be sent through this interface
Packets sent
Packet classification
Token bucket
Queue
Packets dropped

Traffic policing is widely used in policing traffic entering the networks of ISPs. It can classify the policed traffic and depending on the evaluation result, take pre-defined policing actions on each packet:

- Forwarding the packet if the evaluation result is "conforming."
- Dropping the packet if the evaluation result is "excess."
- Forwarding the packet with its precedence re-marked if the evaluation result is "conforming."

Traffic shaping

Traffic shaping shapes the outbound traffic.

Traffic shaping limits the outbound traffic rate by buffering exceeding traffic. Use traffic shaping to adapt the traffic output rate on a device to the input traffic rate of its connected device to avoid packet loss.

The difference between traffic policing and GTS is that packets to be dropped with traffic policing are retained in a buffer or queue with GTS, as shown in Figure 11. When enough tokens are in the token bucket, the buffered packets are sent at an even rate. Traffic shaping can result in additional delay, while traffic policing does not.

Figure 11 Schematic diagram for GTS

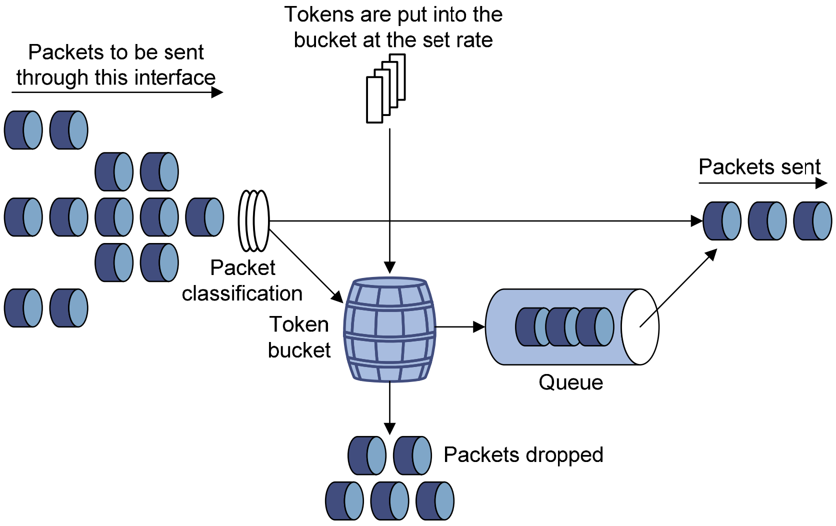
Tokens are put into the bucket at the set rate
Packets to be sent through this interface
Packets sent
Packet classification
Token bucket
Queue
Packets dropped

For example, in Figure 12, Device B performs traffic policing on packets from Device A and drops packets exceeding the limit.

Figure 12 GTS application

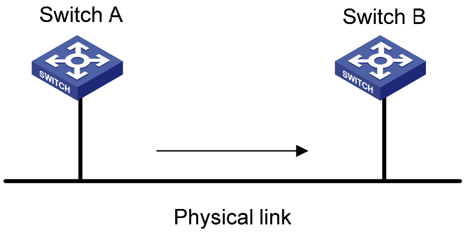

To avoid packet loss, you can perform traffic shaping on the outgoing interface of Switch A so packets exceeding the limit are cached in Switch A. Once resources are released, traffic shaping retrieves the cached packets and sends them out.

Line rate

Line rate rate-limits the inbound and outbound traffic. The outbound traffic is taken for example.

The line rate of a physical interface specifies the maximum rate for forwarding packets (including critical packets).

Line rate also uses token buckets for traffic control. With line rate configured on an interface, all packets to be sent through the interface are handled by the token bucket at the line rate. If enough tokens are in the token bucket, packets can be forwarded. Otherwise, packets are put into QoS queues for congestion management. In this way, the traffic passing the physical interface is controlled.

Figure 13 Line rate implementation

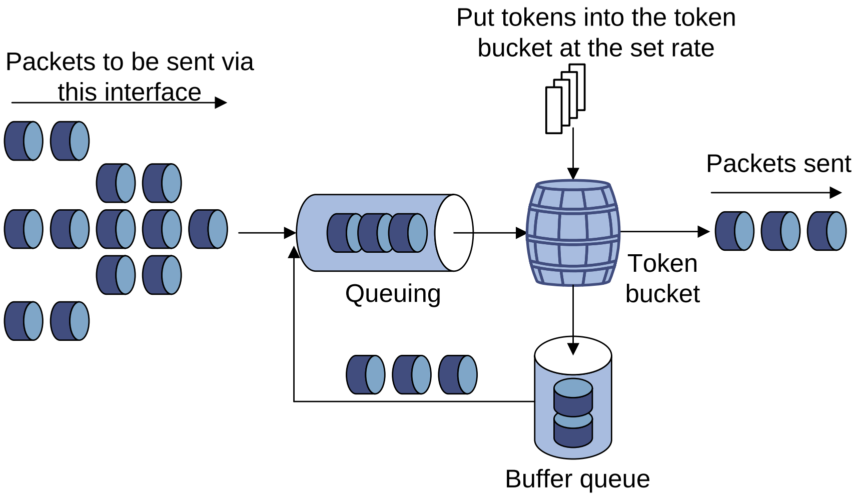
Put tokens into the token bucket at the set rate
Packets to be sent via this interface
Packets sent
Queuing
Token bucket
Buffer queue

The token bucket mechanism limits traffic rate when accommodating bursts. It allows bursty traffic to be transmitted if enough tokens are available. If tokens are scarce, packets cannot be transmitted until efficient tokens are generated in the token bucket. It restricts the traffic rate to the rate for generating tokens.

Line rate can only limit traffic rate on a physical interface, and traffic policing can limit the rate of a flow on an interface. It is easier to use line rate to limit the rate of all packets on interfaces.

Configuring traffic policing

To configure traffic policing:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Create a class and enter class view.	**traffic classifier** *tcl-name* [**operator** { **and** | **or** }]	—
3. Configure the match criteria.	**if-match** *match-criteria*	—
4. Exit class view.	**quit**	—
5. Create a behavior and enter behavior view.	**traffic behavior** *behavior-name*	—

To do…		Use the command…	Remarks
6. Configure a traffic policing action.		**car cir** *committed-information-rate* [**cbs** *committed-burst-size* [**ebs** *excess-burst-size*]] [**pir** *peak-information-rate*] [**green** *action*] [**red** *action*] [**yellow** *action*] [**hierarchy-car** *hierarchy-car-name* [**mode** { **and** | **or** }]]	Required. For more information about hierarchical CAR, see "Configuring global CAR."
7. Exit behavior view.		**quit**	—
8. Create a policy and enter policy view.		**qos policy** *policy-name*	—
9. Associate the class with the traffic behavior in the QoS policy.		**classifier** *tcl-name* **behavior** *behavior-name*	—
10. Exit policy view.		**quit**	—
11. Apply the QoS policy:	To an interface.	Applying the QoS policy to an interface	—
	To online users.	Applying the QoS policy to online users	—
	To a VLAN.	Applying the QoS policy to a VLAN	—
	Globally.	Applying the QoS policy globally	—
	To the control plane.	Applying the QoS policy to the control plane	—

Configuring GTS

The HP 5800 Switch Series and 5820X Switch Series, supports queue-based GTS, which shapes traffic of a specific queue.

To configure queue-based GTS:

To do…		Use the command…	Remarks
1. Enter system view.		**system-view**	—
2. Enter interface view or port group view:	Enter interface view.	**interface** *interface-type interface-number*	Use either command. Settings in interface view take effect on the current interface. Settings in port group view take effect on all ports in the port group.
	Enter port group view.	**port-group manual** *port-group-name*	
3. Configure GTS for a queue.		**qos gts queue** *queue-number* **cir** *committed-information-rate* [**cbs** *committed-burst-size*]	Required.

Configuring the line rate

The line rate of a physical interface specifies the maximum rate of outgoing packets.

To configure the line rate:

To do…		Use the command…	Remarks
1. Enter system view.		**system-view**	—
2. Enter interface view or port group view:	Enter interface view.	**interface** *interface-type interface-number*	Use either command. Settings in interface view take effect on the current interface. Settings in port group view take effect on all ports in the port group.
	Enter port group view.	**port-group manual** *port-group-name*	
3. Configure the inbound or outbound line rate for the interface or port group.		**qos lr** { **inbound** | **outbound** } **cir** *committed-information-rate* [**cbs** *committed-burst-size*]	Required.

Displaying and maintaining traffic policing, GTS, and line rate

On the HP 5800 Switch Series and HP 5820X Switch Series, you can configure traffic policing in MQC approach. For more information about the displaying and maintaining commands, see "Configuration approaches."

To display GTS and line rate complete the following tasks:

Task	Command	Remarks
Display interface GTS configuration information.	**display qos gts interface** [*interface-type interface-number*] [| { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display interface line rate configuration information.	**display qos lr interface** [*interface-type interface-number*] [| { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view

Configuring congestion management

Both bridge mode (Layer 2) and route mode (Layer 3) Ethernet ports support congestion management. The term "interface" in this chapter collectively refers to these types of ports. Use the **port link-mode** command to set an Ethernet port to operate in bridge or route mode (see *Layer 2—LAN Switching Configuration Guide*).

Overview

Network congestion degrades service quality on a traditional network. Congestion is a situation where the forwarding rate decreases due to insufficient resources, resulting in extra delay.

Congestion is more likely to occur in complex packet switching circumstances. Figure 14 shows two common cases:

Figure 14 Traffic congestion causes

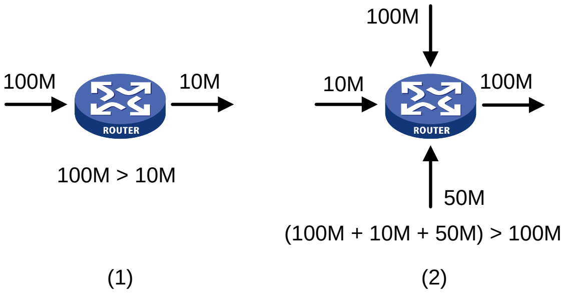

Congestion can bring the following negative results:

- Increased delay and jitter during packet transmission
- Decreased network throughput and resource use efficiency
- Network resource (memory in particular) exhaustion and system breakdown

Congestion is unavoidable in switched networks and multiuser application environments. To improve the service performance of your network, you must take proper measures to address the congestion issues.

The key to congestion management is defining a dispatching policy for resources to decide the order of forwarding packets when congestion occurs.

Techniques

Congestion management uses queuing and scheduling algorithms to classify and sort traffic leaving a port. Each queuing algorithm addresses a particular network traffic problem, and has a different impact on bandwidth resource assignment, delay, and jitter.

Queue scheduling processes packets by their priorities, preferentially forwarding high-priority packets. The following section describes SP, WFQ, WRR, SP+WRR, and SP+WFQ queuing.

SP queuing

SP queuing is designed for mission-critical applications that require preferential service to reduce the response delay when congestion occurs.

Figure 15 Schematic diagram for SP queuing

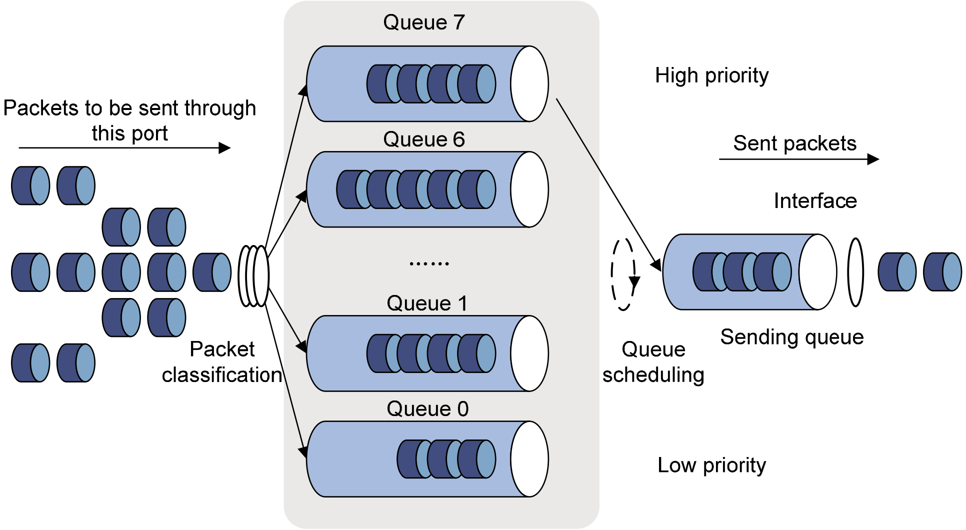

In Figure 15, SP queuing classifies eight queues on a port into eight classes, numbered 7 to 0 in descending priority order.

SP queuing schedules the eight queues in the descending order of priority. SP queuing sends packets in the queue with the highest priority first. When the queue with the highest priority is empty, it sends packets in the queue with the second highest priority, and so on. You can assign mission-critical packets to the high priority queue to make sure that they are always served first, and assign common service packets to the low priority queues, to be transmitted when the high priority queues are empty.

SP guarantees that high priority traffic is always preferentially treated over low priority traffic, without considering the fairness of transmission. Use of SP could, in the worst case, result in lower-priority traffic never being sent.

WRR queuing

WRR queuing schedules all queues in turn to make sure every queue is served for a certain time, as shown in Figure 16.

Figure 16 Schematic diagram for WRR queuing

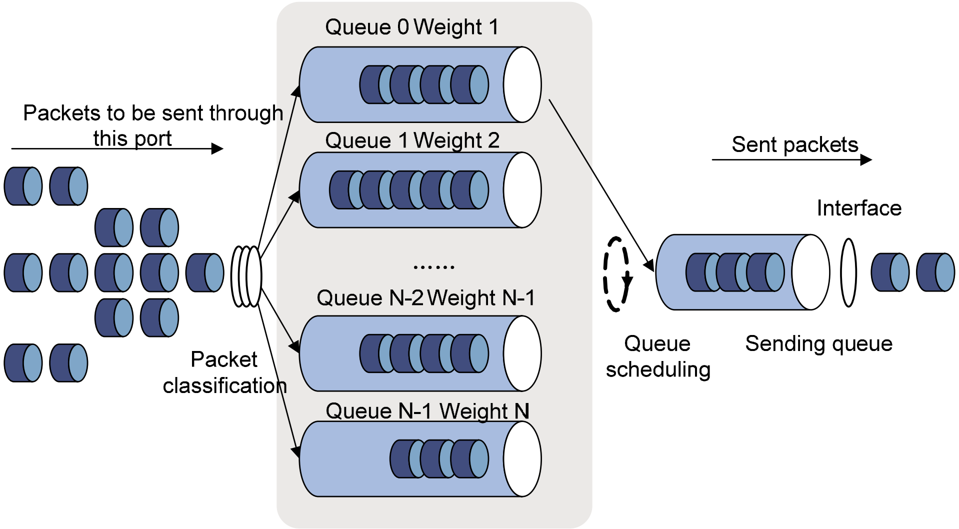

In contrast to SP, WRR queuing schedules queues in a round-robin way to guarantee each queue certain service time in each scheduling cycle. In WRR queuing, each queue has a scheduling weight, which determines the percentage of resources assigned to the queue. The HP 5800 Switch Series and the HP 5820X Switch Series support packet-based WRR, which allocates bandwidth to queues in terms of packets, and byte-count WRR, which allocates bandwidth to queues in terms of bytes.

Take byte-count WRR for example. On a 1000 Mbps port, configure the scheduling weights of queues 0 through 7 as 5, 5, 3, 3, 1, 1, 1, and 1. Then, the lowest-priority queue can get a minimum of 1/(5+5+3+3+1+1+1+1)×1000Mbps = 50 Mbps of bandwidth.

Another advantage of WRR queuing is that when the queues are scheduled in turn, the service time for each queue is not fixed. If a queue is empty, the next queue is scheduled immediately. This improves bandwidth resource use efficiency

WFQ queuing

WFQ queuing is shown in Figure 17.

Figure 17 Schematic diagram for WFQ queuing

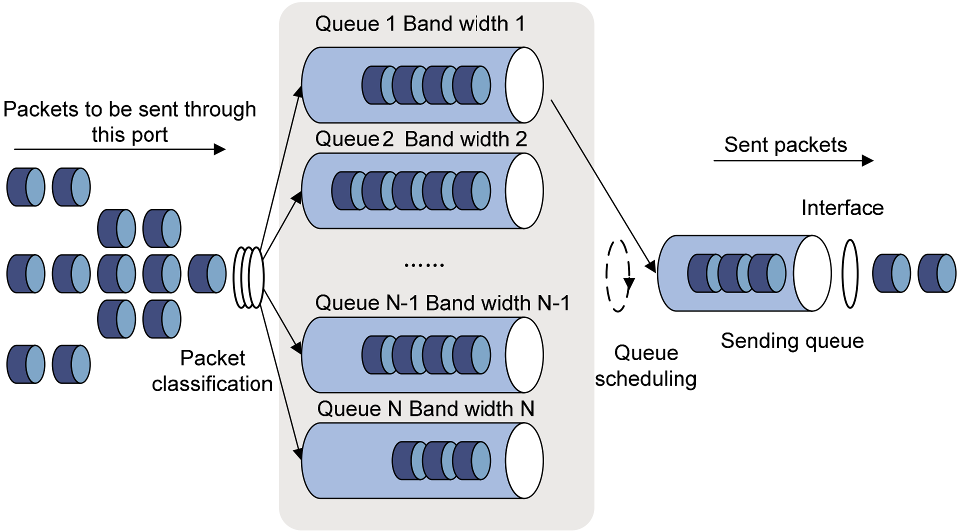

WFQ is similar to WRR. The difference between WFQ and WRR is that WFQ supports byte-count and packet-based scheduling weights. You can use WFQ as an alternative to WRR.

Compared with WRR, WFQ can work with the minimum guaranteed bandwidth as follows:

- By setting the minimum guaranteed bandwidth, you can make sure that each WFQ queue is assured of a certain bandwidth.
- The assignable bandwidth is allocated based on the priority of each queue (assignable bandwidth = total bandwidth – the sum of minimum guaranteed bandwidth of each queue).

For example, assume the total bandwidth of a port is 10 Mbps, and the port has five flows, with the precedence being 0, 1, 2, 3, and 4 and the minimum guaranteed bandwidth being 128 kbps, 128 kbps, 128 kbps, 64 kbps, and 64 kbps, respectively.

- The assignable bandwidth = 10 Mbps – (128 kbps + 128 kbps + 128 kbps + 64 kbps + and 64 kbps) = 9.5 Mbps
- The total assignable bandwidth quota is the sum of all (precedence value + 1)s, 1 + 2 + 3 + 4 + 5 = 15.
- The bandwidth percentage assigned to each flow is (precedence value of the flow + 1)/total assignable bandwidth quota. The bandwidth percentages for the flows are 1/15, 2/15, 3/15, 4/15, and 5/15, respectively.
- The bandwidth assigned to a queue = the minimum guaranteed bandwidth + the bandwidth allocated to the queue from the assignable bandwidth.

WFQ is effectively applied in some special occasions, because WFQ can balance delay and jitter among flows when congestion occurs. For example, WFQ is used for the assured forwarding (AF) services of the RSVP. In GTS, WFQ schedules buffered packets.

SP+WRR queuing

You can assign some queues on a port to the SP scheduling group and the others to the WRR scheduling group (group 1) to implement SP + WRR queue scheduling. The switch schedules packets in the SP scheduling group preferentially, and when the SP scheduling group is empty, schedules the packets in the WRR scheduling group. Queues in the SP scheduling group are scheduled with the SP queue scheduling algorithm. Queues in the WRR scheduling group are scheduled with WRR.

Configuring SP queuing

Configuration procedure

To do…		Use the command…	Remarks
1. Enter system view.		**system-view**	—
2. Enter interface view or port group view:	Enter interface view.	**interface** *interface-type interface-number*	Use either command. Settings in interface view take effect on the current interface. Settings in port group view take effect on all ports in the port group.
	Enter port group view.	**port-group manual** *port-group-name*	
3. Configure SP queuing.		**qos sp**	Required. The default queuing algorithm on an interface is WRR queuing.
4. Display SP queuing configuration.		**display qos sp interface** [*interface-type interface-number*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Optional. Available in any view.

Configuration example

Network requirements

Configure GigabitEthernet 1/0/1 to use SP queuing.

Configuration procedure

# Enter system view.

```
<Sysname> system-view
```

# Configure GigabitEthernet1/0/1 to use SP queuing.

```
[Sysname]interface gigabitethernet 1/0/1
[Sysname-GigabitEthernet1/0/1] qos sp
```

Configuring WRR queuing

Configuration procedure

To do...		Use the command...	Remarks
1. Enter system view.		**system-view**	—
2. Enter interface view or port group view:	Enter interface view.	**interface** *interface-type interface-number*	Use either command. Settings in interface view take effect on the current interface. Settings in port group view take effect on all ports in the port group.
	Enter port group view.	**port-group manual** *port-group-name*	
3. Enable byte-count or packet-based WRR queuing.		**qos wrr** [**byte-count** | **weight**]	Optional. By default, byte-count WRR queuing is enabled.
4. Configure the scheduling weight for a queue:	For a byte-count WRR queue.	**qos wrr** *queue-id* **group 1 byte-count** *schedule-value*	Select an approach according to the WRR queuing type. By default, byte-count WRR is used, and the weights of queues 0 through 7 are 1, 2, 3, 4, 5, 9, 13, and 15.
	For a packet-based WRR queue.	**qos wrr** *queue-id* **group 1 weight** *schedule-value*	
5. Display WRR queuing configuration information.		**display qos wrr interface** [*interface-type interface-number*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Optional. Available in any view.

Configuration example

Network requirements

- Enable byte-count WRR on port GigabitEthernet 1/0/1.
- Assign queues 0 through 7 to the WRR group, with their weights being 1, 2, 4, 6, 8, 10, 12, and 14, respectively.

Configuration procedure

# Enter system view.

```
<Sysname> system-view
```

# Configure WRR queuing on GigabitEthernet 1/0/1.

```
[Sysname] interface GigabitEthernet 1/0/1
[Sysname-GigabitEthernet1/0/1] qos wrr
[Sysname-GigabitEthernet1/0/1] qos wrr 0 group 1 byte-count 1
[Sysname-GigabitEthernet1/0/1] qos wrr 1 group 1 byte-count 2
[Sysname-GigabitEthernet1/0/1] qos wrr 2 group 1 byte-count 4
[Sysname-GigabitEthernet1/0/1] qos wrr 3 group 1 byte-count 6
[Sysname-GigabitEthernet1/0/1] qos wrr 4 group 1 byte-count 8
```

```
[Sysname-GigabitEthernet1/0/1] qos wrr 5 group 1 byte-count 10
[Sysname-GigabitEthernet1/0/1] qos wrr 6 group 1 byte-count 12
[Sysname-GigabitEthernet1/0/1] qos wrr 7 group 1 byte-count 14
```

Configuring WFQ queuing

Configuration procedure

To do…		Use the command…	Remarks
1. Enter system view.		**system-view**	—
2. Enter interface view or port group view:	Enter interface view.	**interface** *interface-type interface-number*	Use either command. Settings in interface view take effect on the current interface. Settings in port group view take effect on all ports in the port group.
	Enter port group view.	**port-group manual** *port-group-name*	
3. Enable WFQ queuing.		**qos wfq**	Required. The default queuing algorithm on an interface is WRR queuing.
4. Configure the minimum guaranteed bandwidth for a WFQ queue.		**qos bandwidth queue** *queue-id* **min** *bandwidth-value*	Optional. 64 kbps by default.
5. Specify the queue scheduling weight for a WFQ queue.		**qos wfq** *queue-id* **weight** *schedule-value*	Optional. 1 by default.
6. Display WFQ queuing configuration.		**display qos wfq interface** [*interface-type interface-number*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Optional. Available in any view.

Configuration example

Network requirements

- Enable WFQ on GigabitEthernet 1/0/1, and set the weights of queues 0 through 7 to 1, 2, 4, 6, 8, 10, 12, and 14, respectively.
- Set the minimum guaranteed bandwidth of queue 0 to 128 kbps.

Configuration procedure

# Enter system view.

```
<Sysname> system-view
```

# Configure WFQ queues on GigabitEthernet 1/0/1.

```
[Sysname] interface GigabitEthernet 1/0/1
[Sysname-GigabitEthernet1/0/1] qos wfq
[Sysname-GigabitEthernet1/0/1] qos wfq 0 weight 1
[Sysname-GigabitEthernet1/0/1] qos wfq 1 weight 2
[Sysname-GigabitEthernet1/0/1] qos wfq 2 weight 4
```

```
[Sysname-GigabitEthernet1/0/1] qos wfq 3 weight 6
[Sysname-GigabitEthernet1/0/1] qos wfq 4 weight 8
[Sysname-GigabitEthernet1/0/1] qos wfq 5 weight 10
[Sysname-GigabitEthernet1/0/1] qos wfq 6 weight 12
[Sysname-GigabitEthernet1/0/1] qos wfq 7 weight 14
```

# Set the minimum guaranteed bandwidth of queue 0 to 128 kbps.

```
[Sysname-GigabitEthernet1/0/1] qos bandwidth queue 0 min 128
```

Configuring SP+WRR queues

Configuration procedure

To do…		Use the command…	Remarks
1. Enter system view.		**system-view**	—
2. Enter interface view or port group view:	Enter interface view.	**interface** *interface-type interface-number*	Use either command. Settings in interface view take effect on the current interface. Settings in port group view take effect on all ports in the port group.
	Enter port group view.	**port-group manual** *port-group-name*	
3. Enable byte-count or packet-based WRR queuing.		**qos wrr** [**byte-count** | **weight**]	Optional. By default, byte-count WRR queuing is enabled.
4. Configure SP queue scheduling.		**qos wrr** *queue-id* **group sp**	Required. By default, all ports use the WRR queue scheduling algorithm.
5. Configure the scheduling weight for a queue:	For a byte-count WRR queue.	**qos wrr** *queue-id* **group 1 byte-count** *schedule-value*	Select an approach according to the WRR queuing type. By default, byte-count WRR is used, and the weights of queues 0 through 7 are 1, 2, 3, 4, 5, 9, 13, and 15.
	For a packet-based WRR queue.	**qos wrr** *queue-id* **group 1 weight** *schedule-value*	

Configuration example

Network requirements

- Configure SP+WRR queue scheduling algorithm on GigabitEthernet 1/0/1.
- Configure queue 0, queue 1, queue 2, and queue 3 on GigabitEthernet 1/0/1 to be in SP queue scheduling group.
- Configure queue 4, queue 5, queue 6, and queue 7 on GigabitEthernet 1/0/1 to use byte-count WRR queuing, with the weight 2, 4, 6, and 8, respectively.

Configuration procedure

# Enter system view.

```
<Sysname> system-view
```

# Enable the SP+WRR queue scheduling algorithm on GigabitEthernet1/0/1.

```
[Sysname] interface GigabitEthernet 1/0/1
[Sysname-GigabitEthernet1/0/1] qos wrr
[Sysname-GigabitEthernet1/0/1] qos wrr 0 group sp
[Sysname-GigabitEthernet1/0/1] qos wrr 1 group sp
[Sysname-GigabitEthernet1/0/1] qos wrr 2 group sp
[Sysname-GigabitEthernet1/0/1] qos wrr 3 group sp
[Sysname-GigabitEthernet1/0/1] qos wrr 4 group 1 byte-count 2
[Sysname-GigabitEthernet1/0/1] qos wrr 5 group 1 byte-count 4
[Sysname-GigabitEthernet1/0/1] qos wrr 6 group 1 byte-count 6
[Sysname-GigabitEthernet1/0/1] qos wrr 7 group 1 byte-count 8
```

Configuring congestion avoidance

Both bridge mode (Layer 2) and route mode (Layer 3) Ethernet ports support congestion avoidance. The term "interface" in this chapter collectively refers to these types of ports. Use the **port link-mode** command to set an Ethernet port to operate in bridge or route mode (see *Layer 2—LAN Switching Configuration Guide*).

Overview

Avoiding congestion before it occurs is a proactive approach to improving network performance. As a flow control mechanism, congestion avoidance actively monitors network resources (such as queues and memory buffers), and drops packets when congestion is expected to occur or deteriorate.

Compared with end-to-end flow control, this flow control mechanism controls the load of more flows in a device. When dropping packets from a source end, it cooperates with the flow control mechanism (such as TCP flow control) at the source end to regulate the network traffic size. The combination of the local packet drop policy and the source-end flow control mechanism helps maximize throughput and network use efficiency and minimize packet loss and delay.

Tail drop

Congestion management techniques drop all packets that are arriving at a full queue. This tail drop mechanism results in global TCP synchronization. If packets from multiple TCP connections are dropped, these TCP connections go into the state of congestion avoidance and slow start to reduce traffic, but traffic peak occurs later. Consequently, the network traffic jitters all the time.

RED and WRED

You can use RED or WRED to avoid global TCP synchronization.

Both RED and WRED avoid global TCP synchronization by randomly dropping packets. When the sending rates of some TCP sessions slow down after their packets are dropped, other TCP sessions remain at high sending rates. Link bandwidth is efficiently used because TCP sessions at high sending rates always exist.

The RED or WRED algorithm sets an upper threshold and lower threshold for each queue, and processes the packets in a queue as follows:

- When the queue size is shorter than the lower threshold, no packets are dropped.
- When the queue size reaches the upper threshold, all subsequent packets are dropped.
- When the queue size is between the lower threshold and the upper threshold, the received packets are dropped at random. The drop probability in a queue increases along with the queue size under the maximum drop probability.

WRED determines differentiated drop policies for packets with different IP precedence values. Packets with a lower IP precedence are more likely to be dropped.

If the current queue length is compared with the upper threshold and lower threshold to determine the drop policy, bursty traffic is not fairly treated. To solve this problem, WRED compares the average queue length with the upper threshold and lower threshold to determine the drop probability.

The average queue length reflects the queue length change trend but is not sensitive to bursty queue length changes, and bursty traffic can be fairly treated. The average queue length is calculated using the formula: average queue length = previous average queue length × $(1-2^{-n})$ + current queue length × 2^{-n}, where n can be configured with the **queue weighting-constant** command.

WRED parameters

Before configuring WRED, determine the following parameters:

- **Upper threshold and lower threshold**—When the average queue length is smaller than the lower threshold, no packet is dropped. When the average queue length is between the lower threshold and the upper threshold, the packets are dropped at random according to certain drop probability. When the average queue length exceeds the upper threshold, subsequent packets are dropped.
- **Drop precedence**—Parameter used in packet drop. Value 0 represents green packets, 1 represents yellow packets, and 2 represents red packets. Red packets are preferentially dropped.
- **Exponent used for average queue length calculation**—The bigger the exponent is, the less sensitive the average queue length is to real-time queue length changes.
- **Drop probability**—Probability of dropping packets, represented in percentage. A greater value indicates a higher probability of dropping packets.

Configuring WRED

To implement WRED on an HP 5800 or HP 5820X switch, configure WRED tables globally in system view and apply them to interfaces. After a WRED table is applied to an interface, you can modify the values of the WRED table, but cannot remove the WRED table.

In the WRED tables, drop parameters are configured on a per queue basis because WRED regulates packets on a per queue basis.

Configuration procedure

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Create a WRED table.	**qos wred queue table** *table-name*	—
3. Set the WRED exponent for average queue length calculation.	**queue** *queue-id* **weighting-constant** *exponent*	Optional. 9 by default.
4. Configure the drop parameters for each queue in the WRED table.	**queue** *queue-id* [**drop-level** *drop-level*] **low-limit** *low-limit* **high-limit** *high-limit* [**discard-probability** *discard-prob*]	Optional. By default, the *low-limit* argument is 100, the *high-limit* is 1000, and the *discard-prob* argument is 10.
5. Enter interface view or Enter interface view.	**interface** *interface-type interface-number*	Use either command. Settings in interface view take

To do…	Use the command…	Remarks
port group view: Enter port group view.	**port-group manual** *port-group-name*	effect on the current interface. Settings in port group view take effect on all ports in the port group.
6. Apply the WRED table.	**qos wred apply** *table-name*	Required.

Configuration example

Network requirements

Configure a WRED table, set the following parameters for yellow packets (with drop precedence 1) in queue 1: upper threshold 100, lower threshold 30, and drop probability 50%, and then apply the WRED to port GigabitEthernet 1/0/1.

Configuration procedure

# Enter system view.

```
<Sysname> system-view
```

# Create a WRED table, and configure parameters for the WRED table as required.

```
[Sysname] qos wred queue table queue-table1
[Sysname-wred-table-queue-table1] queue 1 drop-level 1 low-limit 30 high-limit 100 discard-probability 50
[Sysname-wred-table-queue-table1] quit
```

# Enter the view of GigabitEthernet 1/0/1.

```
[Sysname] interface gigabitethernet 1/0/1
```

# Apply the WRED table to GigabitEthernet 1/0/1.

```
[Sysname-GigabitEthernet1/0/1] qos wred apply queue-table1
```

Displaying WRED

To display WRED complete the following tasks:

Task	Command	Remarks
Display WRED configuration information on an interface or all interfaces.	**display qos wred interface** [*interface-type interface-number*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display configuration information about a WRED table or all WRED tables.	**display qos wred table** [*table-name*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view

Configuring traffic filtering

Overview

You can filter in or filter out a class of traffic by associating the class with a traffic filtering action. For example, you can filter packets sourced from a specific IP address according to network status.

Configuration procedure

To do...		Use the command...	Remarks
1. Enter system view.		**system-view**	—
2. Create a class and enter class view.		**traffic classifier** *tcl-name* [**operator** { **and** | **or** }]	—
3. Configure match criteria.		**if-match** *match-criteria*	—
4. Return to system view.		**quit**	—
5. Create a behavior and enter behavior view.		**traffic behavior** *behavior-name*	—
6. Configure the traffic filtering action.		**filter** { **deny** | **permit** }	Required: • **deny**—Drops packets. • **permit**—Permits packets to pass through.
7. Return to system view.		**quit**	—
8. Create a policy and enter policy view.		**qos policy** *policy-name*	—
9. Associate the class with the traffic behavior in the QoS policy.		**classifier** *tcl-name* **behavior** *behavior-name*	—
10. Return to system view.		**quit**	—
11. Apply the QoS policy:	To an interface.	Applying the QoS policy to an interface	—
	To online users.	Applying the QoS policy to online users	—
	To a VLAN.	Applying the QoS policy to a VLAN	—
	Globally.	Applying the QoS policy globally	—
	To the control plane.	Applying the QoS policy to the control plane	—
12. Display the traffic filtering configuration		**display traffic behavior user-defined** [*behavior-name*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Optional. Available in any view.

NOTE:

With **filter deny** configured for a traffic behavior, the other actions (except class-based accounting) in the traffic behavior do not take effect.

Configuration example

Network requirements

As shown in Figure 18, Host is connected to GigabitEthernet 1/0/1 of Device.

Configure traffic filtering to filter the packets with source port 21 and received on GigabitEthernet 1/0/1.

Figure 18 Network diagram for traffic filtering configuration

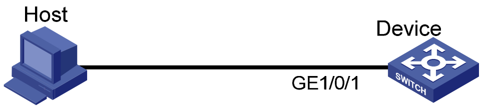

Configuration procedure

# Create advanced ACL 3000, and configure a rule to match packets whose source port number is 21.

```
<DeviceA> system-view
[DeviceA] acl number 3000
[DeviceA-acl-basic-3000] rule 0 permit tcp source-port eq 21
[DeviceA-acl-basic-3000] quit
```

# Create a class named **classifier_1** and reference us ACL 3000 as the match criterion in the class.

```
[DeviceA] traffic classifier classifier_1
[DeviceA-classifier-classifier_1] if-match acl 3000
[DeviceA-classifier-classifier_1] quit
```

# Create a behavior named **behavior_1** and configure the traffic filtering action for the behavior to drop packets.

```
[DeviceA] traffic behavior behavior_1
[DeviceA-behavior-behavior_1] filter deny
[DeviceA-behavior-behavior_1] quit
```

# Create a policy named **policy** and associate class **classifier_1** with behavior **behavior_1** in the policy.

```
[DeviceA] qos policy policy
[DeviceA-qospolicy-policy] classifier classifier_1 behavior behavior_1
[DeviceA-qospolicy-policy] quit
```

# Apply the policy named **policy** to the incoming traffic of GigabitEthernet 1/0/1.

```
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] qos apply policy policy inbound
```

Configuring priority marking

Overview

Priority marking sets the priority fields or flag bits of packets to modify the priority of traffic. For example, you can use priority marking to set IP precedence or DSCP for a class of IP traffic to change its transmission priority in the network.

Priority marking can be used together with priority mapping. For more information, see "Configuring priority mapping."

Color-based priority marking

The switch colors a packet to indicate its transmission priority after evaluating the status of processing resources and the priority of the packet. You can re-mark the priority of a packet depending on its color.

Coloring a packet

The switch can color a packet by using one of the following approaches:

- Use the token bucket mechanism (bucket C and bucket E) of traffic policing:
  - If bucket C has enough tokens, the packet is colored green.
  - If bucket C does not have enough tokens but bucket E has enough tokens, the packet is colored yellow.
  - If neither bucket C nor bucket E has enough tokens, the packet is colored red.
- If traffic policing is not configured, look up the 802.1p priority of a packet in the 802.1p-to-drop priority mapping table, allocates drop precedence to the packet and colors the packet according to the drop precedence:
  - 0 represents green packets.
  - 1 represents yellow packets.
  - 2 represents red packets.

NOTE:

- The traffic policing functions that support coloring packets include common CAR and aggregation CAR. For more information, see "Configuring traffic policing, traffic shaping, and line rate" and "Configuring global CAR."
- For more information about priority mapping tables, see "Configuring priority mapping."

Marking packets based on their colors

Color-based priority marking supports re-marking DSCP, 802.1p priority, and local precedence.

You can configure color-based marking in the following ways:

- Mark packets based on a color set during traffic policing, configure a priority marking action for the color in the traffic policing command **car** or **qos car aggregative**. For more information, see "Configuring traffic policing, traffic shaping, and line rate" and "Configuring global CAR."
- Mark packets based on their drop precedence, configure a priority marking action for a color by using the **remark** command as described in the subsequent section.

IMPORTANT:

Do not use the **remark** command together with the **car** or **car name** command in a traffic behavior to perform color-based marking.

Configuration procedure

To do...	Use the command...	Remarks
1. Enter system view.	**system-view**	—
2. Create a class and enter class view.	**traffic classifier** *tcl-name* [**operator** { **and** | **or** }]	—
3. Configure match criteria.	**if-match** *match-criteria*	—
4. Return to system view.	**quit**	—
5. Create a behavior and enter behavior view.	**traffic behavior** *behavior-name*	—
6. Set the DSCP value for packets.	**remark** [**green** | **red** | **yellow**] **dscp** *dscp-value*	Optional. To guarantee that the QoS policy can be applied normally, make sure that the priority marking actions do not conflict with those in the common CAR action or aggregation CAR.
7. Set the 802.1p priority for packets.	**remark** [**green** | **red** | **yellow**] **dot1p** *802.1p*	Optional. To guarantee that the QoS policy can be applied normally, make sure that the priority marking actions do not conflict with those in the common CAR action or aggregation CAR.
8. Configure the inner-to-outer tag priority copying function.	**remark dot1p customer-dot1p-trust**	

To do…	Use the command…	Remarks
9. Set the drop precedence for packets.	**remark drop-precedence** *drop-precedence-value*	Optional. Applicable to only the outbound direction.
10. Set the IP precedence for packets.	**remark ip-precedence** *ip-precedence-value*	Optional.
11. Set the local precedence for packets.	**remark** [**green** | **red** | **yellow**] **local-precedence** *local-precedence*	Optional. To guarantee that the QoS policy can be applied normally, make sure that the priority marking actions do not conflict with those in the common CAR action or aggregation CAR. If packets are colored by using aggregation CAR, do not mark local precedence values for these packets based on colors.
12. Set the QoS-local ID for packets.	**remark qos-local-id** *local-id-value*	Optional. The QoS-local-ID is used for identifying services and only has local significance. By marking different classes of traffic with the same QoS local ID, you can reclassify them to apply a uniform set of QoS actions on them.
13. Return to system view.	**quit**	—
14. Create a policy and enter policy view.	**qos policy** *policy-name*	—
15. Associate the class with the traffic behavior in the QoS policy.	**classifier** *tcl-name* **behavior** *behavior-name*	—
16. Return to system view.	**quit**	—
17. Apply the QoS policy: To an interface.	Applying the QoS policy to an interface	—
To online users.	Applying the QoS policy to online users	—
To a VLAN	Applying the QoS policy to a VLAN	—
Globally.	Applying the QoS policy globally	—
To the control plane.	Applying the QoS policy to the control plane	—
18. Display the priority marking configuration.	**display traffic behavior user-defined** [*behavior-name*] [| { **begin** | **exclude** | **include** } *regular-expression*]	Optional. Available in any view.

Table 7 Support for priority marking actions in the inbound and outbound directions

Action	Inbound	Outbound
Marking 802.1p priority	Supported	Supported
Marking DSCP precedence	Supported	Supported
Marking drop precedence	Supported	Not supported
Marking IP precedence	Not supported	Supported
Marking local precedence	Supported	Not supported
Marking qos-local-id	Supported	Not supported

Priority marking configuration examples

Basic priority marking configuration example

Network requirements

As shown in Figure 19, the enterprise network of a company interconnects hosts with servers through Device. The network is described as follows:

- Host A and Host B are connected to GigabitEthernet 1/0/1 of Device.
- The data server, mail server, and file server are connected to GigabitEthernet 1/0/2 of Device.

To configure priority marking on a Device to satisfy the following requirements:

Traffic source	Destination	Processing priority
Host A, B	Data server	High
Host A, B	Mail server	Medium
Host A, B	File server	Low

Figure 19 Network diagram for priority marking configuration

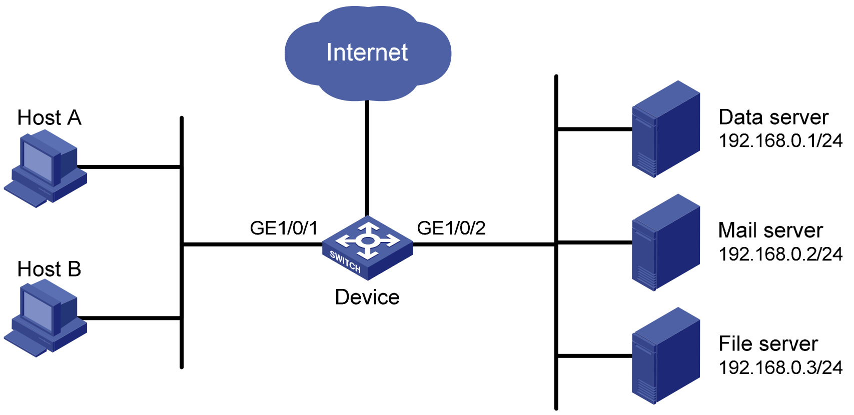

Configuration procedure

# Create advanced ACL 3000 and configure a rule to match packets with destination IP address 192.168.0.1.

```
<Device> system-view
[Device] acl number 3000
[Device-acl-adv-3000] rule permit ip destination 192.168.0.1 0
[Device-acl-adv-3000] quit
```

# Create advanced ACL 3001 and configure a rule to match packets with destination IP address 192.168.0.2.

```
[Device] acl number 3001
[Device-acl-adv-3001] rule permit ip destination 192.168.0.2 0
[Device-acl-adv-3001] quit
```

# Create advanced ACL 3002 and configure a rule to match packets with destination IP address 192.168.0.3.

```
[Device] acl number 3002
[Device-acl-adv-3002] rule permit ip destination 192.168.0.3 0
[Device-acl-adv-3002] quit
```

# Create a class named **classifier_dbserver** and reference ACL 3000 in the class.

```
[Device] traffic classifier classifier_dbserver
[Device-classifier-classifier_dbserver] if-match acl 3000
[Device-classifier-classifier_dbserver] quit
```

# Create a class named **classifier_mserver** and reference ACL 3001 in the class.

```
[Device] traffic classifier classifier_mserver
[Device-classifier-classifier_mserver] if-match acl 3001
[Device-classifier-classifier_mserver] quit
```

# Create a class named **classifier_fserver** and reference ACL 3002 in the class.

```
[Device] traffic classifier classifier_fserver
[Device-classifier-classifier_fserver] if-match acl 3002
[Device-classifier-classifier_fserver] quit
```

# Create a behavior named **behavior_dbserver** and configure the action of setting the local precedence value to 4 for the behavior.

```
[Device] traffic behavior behavior_dbserver
[Device-behavior-behavior_dbserver] remark local-precedence 4
[Device-behavior-behavior_dbserver] quit
```

# Create a behavior named **behavior_mserver** and configure the action of setting the local precedence value to 3 for the behavior.

```
[Device] traffic behavior behavior_mserver
[Device-behavior-behavior_mserver] remark local-precedence 3
[Device-behavior-behavior_mserver] quit
```

# Create a behavior named **behavior_fserver** and configure the action of setting the local precedence value to 2 for the behavior.

```
[Device] traffic behavior behavior_fserver
[Device-behavior-behavior_fserver] remark local-precedence 2
[Device-behavior-behavior_fserver] quit
```

# Create a policy named **policy_server** and associate classes with behaviors in the policy.

```
[Device] qos policy policy_server
[Device-qospolicy-policy_server]      classifier      classifier_dbserver      behavior
behavior_dbserver
[Device-qospolicy-policy_server] classifier classifier_mserver behavior behavior_mserver
[Device-qospolicy-policy_server] classifier classifier_fserver behavior behavior_fserver
[Device-qospolicy-policy_server] quit
```

# Apply the policy named **policy_server** to the incoming traffic of GigabitEthernet 1/0/1.

```
[Device] interface gigabitethernet 1/0/1
[Device-GigabitEthernet1/0/1] qos apply policy policy_server inbound
[Device-GigabitEthernet1/0/1] quit
```

QoS-local-ID marking configuration example

Network requirements

QoS-local-ID marking enables reclassifying packets of multiple classes to perform a uniform set of actions on them as a reclassified class.

Use QoS-local-ID marking to limit the total rate of packets with source MAC address 0001-0001-0001 and packets with source IP address 1.1.1.1 to 128 kbps.

Without QoS-local-ID marking, you can only assign fixed bandwidth to the two classes by associating each of them with a rate-limit traffic behavior. With QoS-local-ID marking, traffic limit applies to the two classes as a whole, allowing the switch to dynamically assign the bandwidth to the two classes depending on their traffic size.

To configure QoS-local-ID marking to limit the total rate of the two classes, you must mark packets of the two classes with the same QoS-local-ID, create a class to match the QoS local ID, and associate this class with the traffic policing action.

Configuration procedure

# Create ACL 2000 to match packets with source IP address 1.1.1.1.

```
<Sysname> system-view
[Sysname] acl number 2000
[Sysname-acl-basic-2000] rule permit source 1.1.1.1 0
[Sysname-acl-basic-2000] quit
```

# Create a class **class_a** to match both packets with source MAC address 0001-0001-0001 and packets with source IP 1.1.1.1.

```
<Sysname> system-view
[Sysname] traffic classifier class_a operator or
[Sysname-classifier-class_a] if-match source-mac 1-1-1
[Sysname-classifier-class_a] if-match acl 2000
[Sysname-classifier-class_a] quit
```

# Create a behavior **behavior_a** and configure the action of marking packets with QoS-local-ID 100 for the behavior.

```
[Sysname] traffic behavior behavior_a
[Sysname-behavior-behavior_a] remark qos-local-id 100
[Sysname-behavior-behavior_a] quit
```

# Create a class **class_b** to match packets with QoS-local-ID 100.

```
[Sysname] traffic classifier class_b
[Sysname-classifier-class_b] if-match qos-local-id 100
[Sysname-classifier-class_b] quit
```

# Create a behavior **behavior_b** and configure the action of limiting traffic rate to 128 kbps for the behavior.

```
[Sysname] traffic behavior behavior_b
[Sysname-behavior-behavior_b] car cir 128
[Sysname-behavior-behavior_b] quit
```

# Create a QoS policy **car_policy**. In the QoS policy, associate class **class_a** with behavior **behavior_a** and associate class **class_b** with behavior **behavior_b**.

```
[Sysname] qos policy car_policy
[Sysname-qospolicy-car_policy] classifier class_a behavior behavior_a
[Sysname-qospolicy-car_policy] classifier class_b behavior behavior_b
```

Apply the QoS policy **car_policy** to the interface, satisfying the network requirements.

Configuring color-based priority marking with traffic policing example

Network requirements

As shown in Figure 20, Switch serves as an edge device of the MPLS domain, and connects to an IPv4 network through port GigabitEthernet 1/0/1. Switch encapsulates traffic entering the MPLS domain with MPLS labels.

Rate-limit and color the traffic with destination IP address 192.168.0.1/24 received on port GigabitEthernet 1/0/1, by using the following parameters: CIR = 1024 kbps, CBS = 8000 bytes, EBS = 8000 bytes, and PIR = 2048 kbps. Transmit all packets evaluated by traffic policing. However, encapsulate green packets with MPLS labels with EXP value 6, encapsulate yellow packets with MPLS labels with EXP value 3, and encapsulate red packets with MPLS labels with EXP value 1.

Figure 20 Network diagram for priority marking based on colors obtained through traffic policing

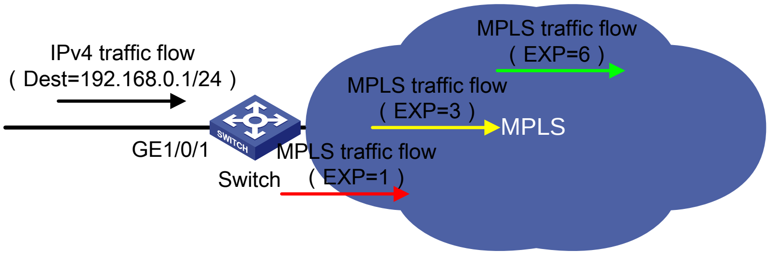

Configuration procedure

1. Configure basic MPLS functions.

For more information about the basic MPLS function configuration, see *MPLS Configuration Guide*.

2. Configure a traffic policing policy.

# Configure advanced ACL 3000 to match the traffic with destination IP address 192.168.0.1/24.

```
<Sysname> system-view
[Sysname] acl number 3000
[Sysname-acl-adv-3000] rule permit ip destination 192.168.1.0 0.0.0.255
[Sysname-acl-adv-3000] quit
```

# Create class **class1** and use ACL 3000 as the match criterion of the class.

```
[Sysname] traffic classifier class1
[Sysname-classifier-class1] if-match acl 3000
[Sysname-classifier-class1] quit
```

# Create behavior **behavior1** and configure a common CAR action according to the network requirements. An 802.1p priority value is mapped to the same EXP value by default. You can mark packets in different colors with different 802.1p priority values, which are mapped to EXP values according to the 802.1p-to-EXP priority mapping table, and mark packets in different colors with different EXP values in the MPLS labels.

```
[Sysname] traffic behavior behavior1
[Sysname-behavior-behavior1] car cir 1024 cbs 8000 ebs 8000 pir 2048 green remark-dot1p-
pass 6 red remark-dot1p-pass 1 yellow remark-dot1p-pass 3
[Sysname-behavior-behavior1] quit
```

# Create QoS policy **policy1** and associate class **class1** with behavior **behavior1** in the QoS policy.

```
[Sysname] qos policy policy1
[Sysname-qospolicy-policy1] classifier class1 behavior behavior1
[Sysname-qospolicy-policy1] quit
```

# Apply QoS policy **policy1** to the incoming traffic of port GigabitEthernet 1/0/1.

```
[Sysname] interface gigabitethernet 1/0/1
[Sysname-GigabitEthernet1/0/1] qos apply policy policy1 inbound
```

Configuring traffic redirecting

Overview

Traffic redirecting is the action of redirecting the packets matching the specific match criteria to a certain location for processing.

Supported redirection actions

The following redirect actions are supported:

- **Redirecting traffic to the CPU**—Redirects packets that require processing by the CPU to the CPU.
- **Redirecting traffic to an interface**—Redirects packets that require processing by an interface to the interface. Note that this action only applies to Layer 2 packets, and the target interface must be a Layer 2 interface.
- **Redirecting traffic to the next hop**—Redirects packets that require processing by an interface to the interface. This action only applies to Layer 3 packets.
- Configuration procedure The actions of redirecting traffic to the CPU, redirecting traffic to an interface, and redirecting traffic to the next hop are mutually exclusive in the same traffic behavior.
- The default of the **fail-action** keyword, if supported, is **forward**.
- Use the **display traffic behavior user-defined** command to view the traffic redirecting configuration.

To configure traffic redirection:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Create a class and enter class view.	**traffic classifier** *tcl-name* [**operator** { **and** | **or** }]	—
3. Configure match criteria.	**if-match** *match-criteria*	—
4. Return to system view.	**quit**	—
5. Create a behavior and enter behavior view.	**traffic behavior** *behavior-name*	Required
6. Configure a traffic redirecting action.	**redirect** { **cpu** | **interface** *interface-type interface-number* | **next-hop** { *ipv4-add1* [*ipv4-add2*] | *ipv6-add1* [*interface-type interface-number*] [*ipv6-add2* [*interface-type interface-number*]] } [**fail-action** { **discard** | **forward** }] }	Optional
7. Return to system view	**quit**	—
8. Create a policy and enter policy view.	**qos policy** *policy-name*	—
9. Associate the class with the traffic behavior in the QoS policy.	**classifier** *tcl-name* **behavior** *behavior-name*	—

To do…		Use the command…	Remarks
10. Return to system view.		**quit**	—
11. Apply the QoS policy:	To an interface.	Applying the QoS policy to an interface	—
	To a VLAN.	Applying the QoS policy to a VLAN	—
	Globally.	Applying the QoS policy globally	—
	To the control plane.	Applying the QoS policy to the control plane	—

Configuring global CAR

Overview

Global CAR polices traffic flows globally. It adds flexibility to common CAR where traffic policing is performed only on a per-class or per-interface basis. In this approach, CAR actions are created in system view and each can be referenced to police multiple traffic flows as a whole. Global CAR provides the following CAR actions: aggregation CAR and hierarchical CAR.

Aggregation CAR

An aggregation CAR action is created globally and can be directly applied to interfaces or referenced in the traffic behaviors associated with different traffic classes to police multiple traffic flows as a whole. The total rate of the traffic flows must conform to the traffic policing specifications set in the aggregation CAR action.

Hierarchical CAR

A hierarchical CAR action is created globally. It must be used in conjunction with a common CAR or aggregation CAR action. With a hierarchical CAR action, you can limit the total traffic of multiple classes in addition to configuring a separate common CAR action (or aggregation CAR) for each class.

A hierarchical CAR action can be referenced in the common or aggregation CAR action for a class in either AND mode or OR mode:

- **AND**—The rate of the traffic class is strictly limited under the common or aggregation CAR. This mode applies to flows that must be strictly rate limited.
- **OR**—The traffic class can use idle bandwidth of other classes associated with the hierarchical CAR. This mode applies to high priority, bursty traffic like video.

By using the two modes appropriately, you can improve bandwidth efficiency.

For example, suppose two flows exist: a low priority data flow and a high priority, bursty video flow. Their total traffic rate cannot exceed 4096 and the video flow must be assured of at least 2048 kbps bandwidth. You can configure common CAR actions to set the traffic rate to 2048 kbps for the two flows, configure a hierarchical CAR action to limit their total traffic rate to 4096 kbps, and reference the action in AND mode in the common CAR action for the data flow, and in OR mode in the common CAR action for the video flow. The video flow is assured of 2048 kbps bandwidth and can use idle bandwidth of the data flow.

In a bandwidth oversubscription scenario, where the bandwidth of an uplink port is lower than the total traffic rate of its downlink ports for example, you can use hierarchical CAR to limit the total rate of the traffic from the downlink ports while allowing each downlink port to forward traffic at the maximum rate when the other ports are idle. For example, use common CAR actions to limit the rate of Internet access flow 1 and that of flow 2 to 128 kbps respectively, and use a hierarchical CAR action to limit their total traffic rate to 192 kbps. Reference the hierarchical CAR action for both flow 1 and flow 2 in AND mode. When flow 1 is not present, flow 2 is transmitted at the maximum rate, 128 kbps. When both flows are present, the traffic rate of flow 2 may drop below this rate, because the total rate of the two flows cannot exceed 192 kbps.

Configuring aggregation CAR

Configuration procedure

To do...	Use the command...	Remarks
1. Enter system view.	**system-view**	—
2. Configure an aggregation CAR action.	**qos car** *car-name* **aggregative cir** *committed-information-rate* [**cbs** *committed-burst-size* [**ebs** *excess-burst-size*]] [**pir** *peek-information-rate*] [**red** *action*]	Required
3. Enter behavior view.	**traffic behavior** *behavior-name*	Required
4. Reference the aggregation CAR in the traffic behavior.	**car name** *car-name*	Required
5. Display the traffic behavior configuration information.	**display traffic behavior user-defined** [*behavior-name*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Optional Available in any view
6. Display the configuration and statistics for the specified aggregation CAR.	**display qos car name** [*car-name*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	

Configuration example

# Create an aggregation CAR action named **aggcar-1**, and configure the following parameters for it: CIR 256, CBS 2000, and dropping red packets. Reference aggregation CAR **aggcar-1** in behavior **be1**.

```
<Sysname> system-view
[Sysname] qos car aggcar-1 aggregative cir 256 cbs 2000 red discard
[Sysname] traffic behavior be1
[Sysname-behavior-be1] car name aggcar-1
```

Configuring hierarchical CAR

To do...		Use the command...	Remarks
1. Enter system view		**system-view**	—
2. Configure a hierarchical CAR action		**qos car** *car-name* **hierarchy cir** *committed-information-rate* [**cbs** *committed-burst-size*]	Required.
3. Enter behavior view		**traffic behavior** *behavior-name*	Required.
4. Reference the hierarchic	Aggregation CAR	**car name** *car-name* **hierarchy-car** *hierarchy-car-name* [**mode** { **and** | **or** }]	Required. Use either approach.

To do…		Use the command…	Remarks
al CAR in the traffic behavior to cooperate with an aggregation or common CAR:	Common CAR.	**car cir** *committed-information-rate* [**cbs** *committed-burst-size* [**ebs** *excess-burst-size*]] [**pir** *peak-information-rate*] [**green** *action*] [**yellow** *action*] [**red** *action*] **hierarchy-car** *hierarchy-car-name* [**mode** { **and** | **or** }]	For more information about common CAR, see "Configuring traffic policing, traffic shaping, and line rate."
5. Display the traffic behavior configuration.		**display traffic behavior user-defined** [*behavior-name*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Optional. Available in any view.
6. Display the configuration and statistics for the specified global CAR.		**display qos car name** [*car-name*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	

Displaying and maintaining global CAR configuration

To display global CAR complete the following tasks:

Task	Command	Remarks
Display statistics for global CAR actions.	**display qos car name** [*car-name*] [**|** { **begin** | **exclude** | **include** } *regular-expression*]	Required Available in any view
Clear statistics for global CAR actions..	**reset qos car name** [*car-name*]	Required Available in user view

Global CAR configuration examples

Aggregation CAR configuration example

Network requirements

As shown in Figure 21, configure an aggregation CAR to rate-limit the traffic of VLAN 10 and VLAN 100 received on GigabitEthernet 1/0/1 using these parameters: CIR is 256 kbps, CBS is 2000 bytes, and the action for red packets is **discard**.

Figure 21 Network diagram for aggregation CAR configuration

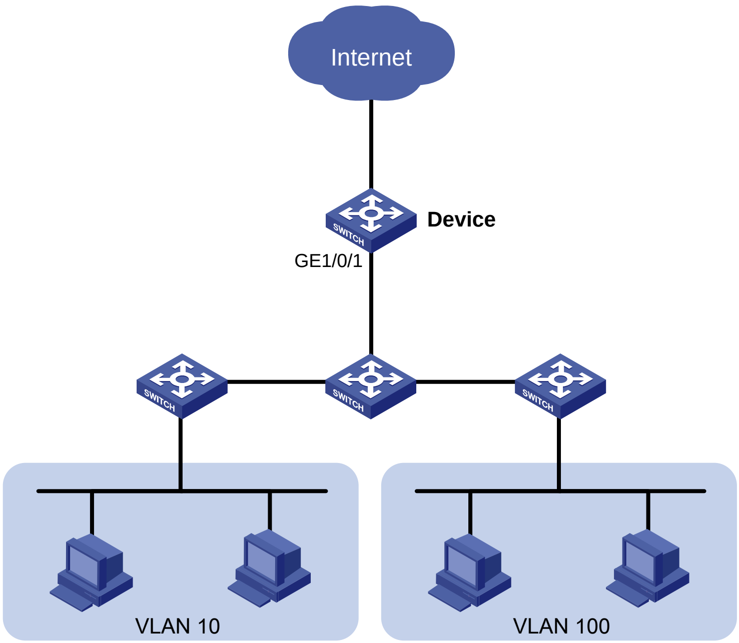

Configuration procedure

# Configure an aggregation CAR according to the rate limit requirements.

```
<Sysname> system-view
[Sysname] qos car aggcar-1 aggregative cir 256 cbs 2000 red discard
```

# Create class 1 to match traffic of VLAN 10; create behavior 1 and reference the aggregation CAR in the behavior.

```
[Sysname] traffic classifier 1
[Sysname-classifier-1] if-match customer-vlan-id 10
[Sysname-classifier-1] quit
[Sysname] traffic behavior 1
[Sysname-behavior-1] car name aggcar-1
[Sysname-behavior-1] quit
```

# Create class 2 to match traffic of VLAN 100; create behavior 2 and reference the aggregation CAR in the behavior.

```
[Sysname] traffic classifier 2
[Sysname-classifier-2] if-match customer-vlan-id 100
[Sysname-classifier-2] quit
[Sysname] traffic behavior 2
[Sysname-behavior-2] car name aggcar-1
[Sysname-behavior-2] quit
```

# Create QoS policy **car**, associate class 1 with behavior 1 and associate class 2 with behavior 2.

```
[Sysname] qos policy car
[Sysname-qospolicy-car] classifier 1 behavior 1
[Sysname-qospolicy-car] classifier 2 behavior 2
[Sysname-qospolicy-car] quit
```

# Apply the QoS policy to the incoming traffic of GigabitEthernet 1/0/1.

```
[Sysname] interface GigabitEthernet 1/0/1
[Sysname-GigabitEthernet1/0/1]qos apply policy car inbound
```

AND-mode hierarchical CAR configuration example

Network requirements

As shown in Figure 22, configure rate limiting for HTTP traffic received on GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 to limit the rate of HTTP traffic received on each port to 192 kbps. At the same time, configure hierarchical CAR to limit the rate of HTTP traffic received on the two ports to 256 kbps and drop the exceeding packets.

Figure 22 Network diagram for AND-mode hierarchical CAR configuration

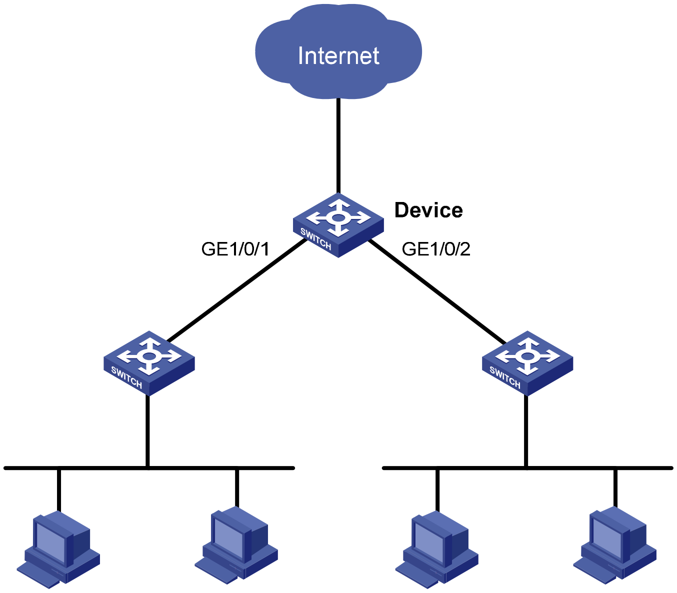

Configuration procedure

# Configure a hierarchical CAR action according to the rate limit requirements.

```
<Device> system-view
[Device] qos car http hierarchy cir 256 red discard
```

# Configure ACL 3000 to match HTTP packets.

```
[Device] acl number 3000
[Device-acl-basic-3000] rule permit tcp destination-port eq 80
[Device-acl-basic-3000] quit
```

# Create class 1 and use ACL 3000 as the match criterion in the class; create behavior 1, configure the common CAR action in the behavior, and reference the hierarchical CAR, with the collaborating mode being AND.

```
[Device] traffic classifier 1
[Device-classifier-1] if-match acl 3000
[Device-classifier-1] quit
[Device] traffic behavior 1
[Device-behavior-1] car cir 192 hierarchy-car http mode and
[Device-behavior-1] quit
```

# Create a QoS policy named **http** and associate class 1 with traffic behavior 1 in the QoS policy.

```
[Device] qos policy http
[Device-qospolicy-http] classifier 1 behavior 1
[Device-qospolicy-http] quit
```

# Apply the policy named **http** to the incoming traffic of GigabitEthernet 1/0/1 and GigabitEthernet 1/0/2 respectively.

```
[Device] interface gigabitethernet 1/0/1
[Device-GigabitEthernet1/0/1] qos apply policy http inbound
[Device-GigabitEthernet1/0/1] quit
[Device] interface gigabitethernet 1/0/2
[Device-GigabitEthernet1/0/2] qos apply policy http inbound
```

OR-mode hierarchical CAR configuration example

Network requirements

As shown in Figure 23, configure rate limiting for video traffic received from 192.168.0.2 and 192.168.0.3 on GigabitEthernet 1/0/1. Set the CIR to 256 kbps for both video streams according to their regular average rates. To guarantee that occasional large bursts can pass through, configure hierarchical CAR to limit the video traffic rate to 640 kbps, and drop the exceeding traffic.

Figure 23 Network diagram for or mode configuration

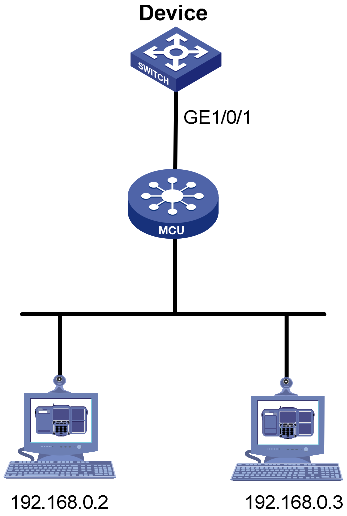

Configuration procedure

# Configure a hierarchical CAR action according to the rate limit requirements.

```
<Device> system-view
[Device] qos car video hierarchy cir 640 red discard
```

# Create class 1 and use ACL 2000 as the match criterion to match packets sourced from 192.168.0.2 in the class; create behavior 1, configure a common CAR action, and reference hierarchical CAR named **video**, with the collaborating mode being OR.

```
[Device] acl number 2000
[Device-acl-basic-2000] rule permit source 192.168.0.2 0.0.0.0
[Device-acl-basic-2000] quit
[Device] traffic classifier 1
```

```
[Device-classifier-1] if-match acl 2000
[Device-classifier-1] quit
[Device] traffic behavior 1
[Device-behavior-1] car cir 256 hierarchy-car video mode or
[Device-behavior-1] quit
```

# Create class 2 and use ACL 2001 as the match criterion to match packets sourced from 192.168.0.3 in the class; create behavior 2, configure a common CAR action, and reference hierarchical CAR named **video**, with the collaborating mode being OR.

```
[Device] acl number 2001
[Device-acl-basic-2001] rule permit source 192.168.0.3 0.0.0.0
[Device-acl-basic-2001] quit
[Device] traffic classifier 2
[Device-classifier-2] if-match acl 2001
[Device-classifier-2] quit
[Device] traffic behavior 2
[Device-behavior-2] car cir 256 hierarchy-car video mode or
[Device-behavior-2] quit
```

# Create a QoS policy named **video** and associate class 1 with traffic behavior 1 and class 2 with behavior 2 in the QoS policy.

```
[Device] qos policy video
[Device-qospolicy-video] classifier 1 behavior 1
[Device-qospolicy-video] classifier 2 behavior 2
[Device-qospolicy-video] quit
```

# Apply the policy named **video** to the incoming traffic of GigabitEthernet 1/0/1.

```
[Device] interface gigabitethernet 1/0/1
[Device-GigabitEthernet1/0/1] qos apply policy video inbound
```

Configuring class-based accounting

Overview

Class-based accounting collects statistics (in packets or bytes) on a per-traffic class basis. For example, you can define the action to collect statistics for traffic sourced from a certain IP address. By analyzing the statistics, you can determine whether anomalies have occurred and what action to take.

Configuration procedure

To do…		Use the command…	Remarks
1. Enter system view.		**system-view**	—
2. Create a class and enter class view.		**traffic classifier** *tcl-name* [**operator** { **and** | **or** }]	—
3. Configure the match criteria.		**if-match** *match-criteria*	—
4. Exit class view.		**quit**	—
5. Create a behavior and enter behavior view.		**traffic behavior** *behavior-name*	Required
6. Configure the accounting action.		**accounting** { **byte** | **packet** }	Optional. • **byte**—Counts traffic in bytes. • **packet**—Counts traffic in packets.
7. Exit behavior view.		**quit**	—
8. Create a policy and enter policy view.		**qos policy** *policy-name*	—
9. Associate the class with the traffic behavior in the QoS policy.		**classifier** *tcl-name* **behavior** *behavior-name*	—
10. Exit policy view.		**quit**	—
11. Apply the QoS policy:	To an interface.	Applying the QoS policy to an interface	—
	To a VLAN.	Applying the QoS policy to a VLAN	—
	Globally.	Applying the QoS policy globally	—
	To the control plane.	Applying the QoS policy to the control plane	—

Displaying and maintaining class-based traffic accounting

You can verify the configuration with the **display qos policy global**, **display qos policy interface**, or **display qos vlan-policy** command depending on the occasion where the QoS policy is applied.

Configuration example

Network requirements

As shown in Figure 24, Host is connected to GigabitEthernet 1/0/1 of Device.

Configure class-based accounting to collect statistics for traffic sourced from 1.1.1.1/24 and received on GigabitEthernet 1/0/1.

Figure 24 Network diagram for traffic accounting configuration

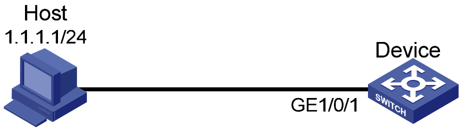

Configuration procedure

# Create basic ACL 2000 and configure a rule to match packets with source IP address 1.1.1.1.

```
<DeviceA> system-view
[DeviceA] acl number 2000
[DeviceA-acl-basic-2000] rule permit source 1.1.1.1 0
[DeviceA-acl-basic-2000] quit
```

# Create a class named **classifier_1** and reference ACL 2000 in the class.

```
[DeviceA] traffic classifier classifier_1
[DeviceA-classifier-classifier_1] if-match acl 2000
[DeviceA-classifier-classifier_1] quit
```

# Create behavior **behavior_1** and configure a byte-based accounting action in the behavior.

```
[DeviceA] traffic behavior behavior_1
[DeviceA-behavior-behavior_1] accounting byte
[DeviceA-behavior-behavior_1] quit
```

# Create a policy named **policy** and associate class **classifier_1** with behavior **behavior_1** in the policy.

```
[DeviceA] qos policy policy
[DeviceA-qospolicy-policy] classifier classifier_1 behavior behavior_1
[DeviceA-qospolicy-policy] quit
```

# Apply the policy named **policy** to the incoming traffic of GigabitEthernet 1/0/1.

```
[DeviceA] interface gigabitethernet 1/0/1
[DeviceA-GigabitEthernet1/0/1] qos apply policy policy inbound
[DeviceA-GigabitEthernet1/0/1] quit
```

Display traffic statistics to verify the configuration.

```
[DeviceA] display qos policy interface gigabitethernet 1/0/1

 Interface: GigabitEthernet1/0/1

 Direction: Inbound

 Policy: policy
  Classifier: classifier_1
   Operator: AND
   Rule(s) : If-match acl 2000
   Behavior: behavior_1
   Accounting Enable:
    28529 (Packets)
```

Configuring the data buffer

Overview

An HP 5800 or HP 5820X switch provides a data buffer for buffering packets to be sent out ports to avoid packet loss when bursty traffic causes congestion.

Buffer resources

The switch controls how a port uses the data buffer by allocating the cell resource and packet resource (called "buffer resources").

- The cell resource is the physical storage space in cells for the data buffer. The cell resource allocated to a port indicates the maximum buffer space that the port can occupy in the buffer.
- The packet resource is the logical buffering space in packets. A packet is counted as one regardless of its size. A packet in the packet resource uses a certain amount of cell resources in the cell resource, depending on its size. The packet resource allocated to a port indicates the maximum number of packets that the port can store in the buffer.

Set independently, the packet resource and the cell resource work simultaneously to regulate data buffering. A packet can be buffered only when both resources are adequate.

NOTE:

On the HP 5820X Switch Series, the data buffer is only used through allocating cell resources.

Data buffer allocation

To handle bursty traffic flexibly, the HP 5800 Switch Series and the HP 5820X Switch Series divide the cell resource and packet resource into a shared resource and a dedicated resource. A port whose dedicated resource is used out can temporarily use the shared resource for transmitting packets. You can manually set the shared resource size by a percentage of the cell resource and packet resource. The remaining buffer space is then used as the dedicated resource area, as shown in Figure 25.

NOTE:

Both the cell resources and the packet resources are allocated as described in this section, but you can configure different allocation schemes for them.

Figure 25 Buffer resource allocation on the HP 5800 Switch Series and the HP 5820X Switch Series

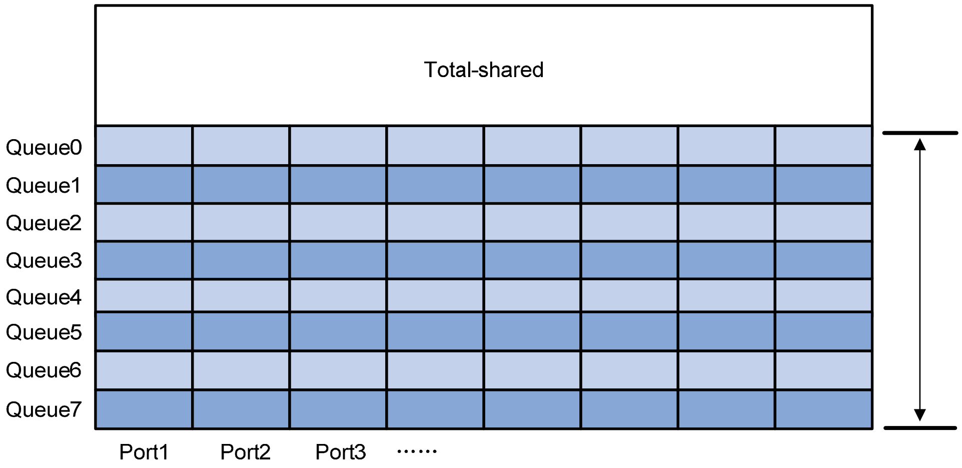

The dedicated buffer is allocated following these rules:

- **On a per-port basis**—Illustrated by the vertical lines in Figure 25, the switch automatically divides the dedicated resource evenly among all ports.
- **On a per-queue basis**—Illustrated by the horizontal lines in Figure 25, the dedicated resource of each port is proportionately allocated among the queues on it and all ports use the same allocation scheme. The percentage of the resource allocated to a queue is called the "minimum guaranteed resource percentage of the queue."

Using the shared resource

When a port needs to send jumbo packets or a large quantity of packets in a short period, its dedicated cell resource or packet resource may become inadequate.

To address the occasional requirement of ports for buffering space, a portion of space is taken out of the cell resource and the packet resource as the shared resource to provide a temporary buffering space that can be used by any queue on any port. When the dedicated cell or packet resource of a certain port or queue is insufficient to accommodate a traffic burst, it can temporarily use the shared resource. After the burst traffic is transmitted, the occupied shared resource is reclaimed for other ports or queues to use.

How queues use the shared resource

When a certain queue of a port is congested because its dedicated cell resource or packet resource gets full, it can use a certain portion of the shared resource. The maximum shared resource size available for a queue is defined as a percentage of the shared resource, which is user configurable and applied globally to the queue with the same number on each port. For example, you can configure queue 0 to use up to 6% of the shared resource space in the cell resource when its dedicated cell resource gets full.

How ports use the shared resource

When all queues of a port are congested because the dedicated cell resource or packet resource space of the port gets full, it can use a certain portion of the shared resource. The maximum shared resource size available for a port is defined as a percentage of the shared resource, which is user configurable and applied globally to all ports.

NOTE:

The total percentage of the shared resource used by the eight queues on a port cannot exceed the maximum shared resource size per port, despite their respective maximum shared resource size settings.

Configuration approaches

You can configure the data buffer in one of the following approaches:

- Using the burst function to configure the data buffer setup
- Manually configuring the data buffer setup

NOTE:

The two approaches are mutually exclusive. If the data buffer setup has been configured in one approach, you must remove the present configuration first before you use the other approach.

Using the burst function to configure the data buffer setup

The burst function allows the switch to automatically determine the shared resource size, the minimum guaranteed resource size for each queue, the maximum shared resource size for each queue, and the maximum shared resource size per port.

The burst function helps optimize packet buffering to ameliorate forwarding performance in the following scenarios:

- Broadcast or multicast traffic is dense and bursts of traffic are usually large.
- High-speed traffic is forwarded over low-speed links or traffic received from multiple ports is forwarded through a port operating at the same speed.

To configure the data buffer setup using the burst function:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Enable the burst function.	**burst-mode enable**	Required Disabled by default

Manually configuring the data buffer setup

IMPORTANT:

Data buffer configuration is complicated and has significant impacts on the forwarding performance of a device. HP does not recommend modifying the data buffer parameters unless you are sure that your device will benefit from the change. If a larger buffer is needed, HP recommends that you enable the burst function to allocate the buffer automatically.

Configuration task list

Perform the following tasks to manually configure the data buffer setup:

Task	Remarks
Configuring the shared resource size	All these tasks are optional. If any of them is performed, you must apply the configuration to have it take effect.
Configuring the minimum guaranteed resource size for a queue	
Configuring the maximum shared resource size for a queue	
Configuring the maximum shared resource size per port	
Applying the data buffer settings	Required

Configuring the shared resource size

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Configure the shared resource area of the cell resource in percentage.	**buffer egress** [**slot** *slot-number*] **cell total-shared ratio** *ratio*	Configure at least one command. By default, on the HP 5800 Switch Series, 73% of the cell resource is the shared resource and 74% of the packet resource is the shared resource; on the HP 5820X Switch Series, 66% of the cell resource is the shared resource. The HP 5820X Switch Series does not support the packet resource.
3. Configure the shared resource area of the packet resource in percentage.	**buffer egress** [**slot** *slot-number*] **packet total-shared ratio** *ratio*	

Configuring the minimum guaranteed resource size for a queue

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Configure the minimum guaranteed cell resource size for a queue as a percentage of the dedicated cell resource per port.	**buffer egress** [**slot** *slot-number*] **cell queue** *queue-id* **guaranteed ratio** *ratio*	Configure at least one command. By default, the minimum guaranteed resource size for a queue is 12% of the dedicated resource of the port in both the cell resource and the packet resource. The HP 5820X Switch Series does not support the packet resource.
3. Configure the minimum guaranteed packet resource size for a queue as a percentage of the dedicated packet resource per port.	**buffer egress** [**slot** *slot-number*] **packet queue** *queue-id* **guaranteed ratio** *ratio*	

NOTE:

Modifying the minimum guaranteed resource size for a queue can affect other queues, because the dedicated resource of a port is shared by eight queues. The system will automatically allocate the remaining dedicated resource space among all queues that are not manually assigned a minimum guaranteed resource space. For example, if you set the minimum guaranteed resource size to 30% for a queue, the remaining seven queues will each share 10% of the dedicated resource of the port.

The minimum guaranteed resource settings of a queue apply to the queue with the same number on each port.

Configuring the maximum shared resource size for a queue

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Configure the maximum shared cell resource size for a queue as a percentage of the shared cell resource.	**buffer egress** [**slot** *slot-number*] **cell queue** *queue-id* **shared ratio** *ratio*	Configure at least one command. By default, a queue can use up to 33% of both the shared cell resource and the shared packet resource. The HP 5820X Switch Series does not support the packet resource.
3. Configure the maximum shared packet resource size for a queue as a percentage of the shared packet resource.	**buffer egress** [**slot** *slot-number*] **packet queue** *queue-id* **shared ratio** *ratio*	

NOTE:

The maximum shared resource settings for a queue apply to the queue with the same number on each port.

Configuring the maximum shared resource size per port

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Configure the maximum shared cell resource size per port as a percentage of the shared cell resource.	**buffer egress** [**slot** *slot-number*] **cell shared ratio** *ratio*	Configure at least one command. By default, a port can use up to 33% of both the shared cell resource and the share packet resource. The HP 5820X Switch Series does not support the packet resource.
3. Configure the maximum shared packet resource size per port as a percentage of the shared packet resource.	**buffer egress** [**slot** *slot-number*] **packet shared ratio** *ratio*	

NOTE:

The maximum shared resource per port settings apply to all ports.

Applying the data buffer settings

To have the manual data buffer settings you made take effect, apply them globally.

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Apply the data buffer settings.	**buffer apply**	Required

Configuring HQoS

HQoS is available on the HP 5800 Switch Series, but not the HP 5820X Switch Series.

Prerequisites

Before reading this chapter, see *ACL and QoS Configuration Guide* to become familiar with the following:

- SP and WRR queue scheduling algorithms in "Configuring congestion management."
- Packet priorities and priority mapping tables in "Configuring priority mapping."
- Traffic shaping in "Configuring traffic policing, traffic shaping, and line rate."

Overview

HQoS uniformly manages traffic and hierarchically schedules traffic by user, network service, and application. It controls internal resources based on policies at different levels.

HQoS provides greater granular traffic control and quality assurance services than common QoS implementations. Common QoS implementations assign packets to output queues only according to packet priorities. They do not differentiate packets by user.

Functionality of HQoS on the HP 5800 Switch Series

The HQoS function on the HP 5800 Switch Series supports hierarchically scheduling packets by VLAN and packet priority. It preferentially schedules packets from certain VLANs and then performs priority-based queue scheduling for the scheduled traffic together with other traffic. Figure 26 shows two-layer scheduling of HQoS prioritizes traffic transmission for high-priority users.

Figure 26 How HQoS works

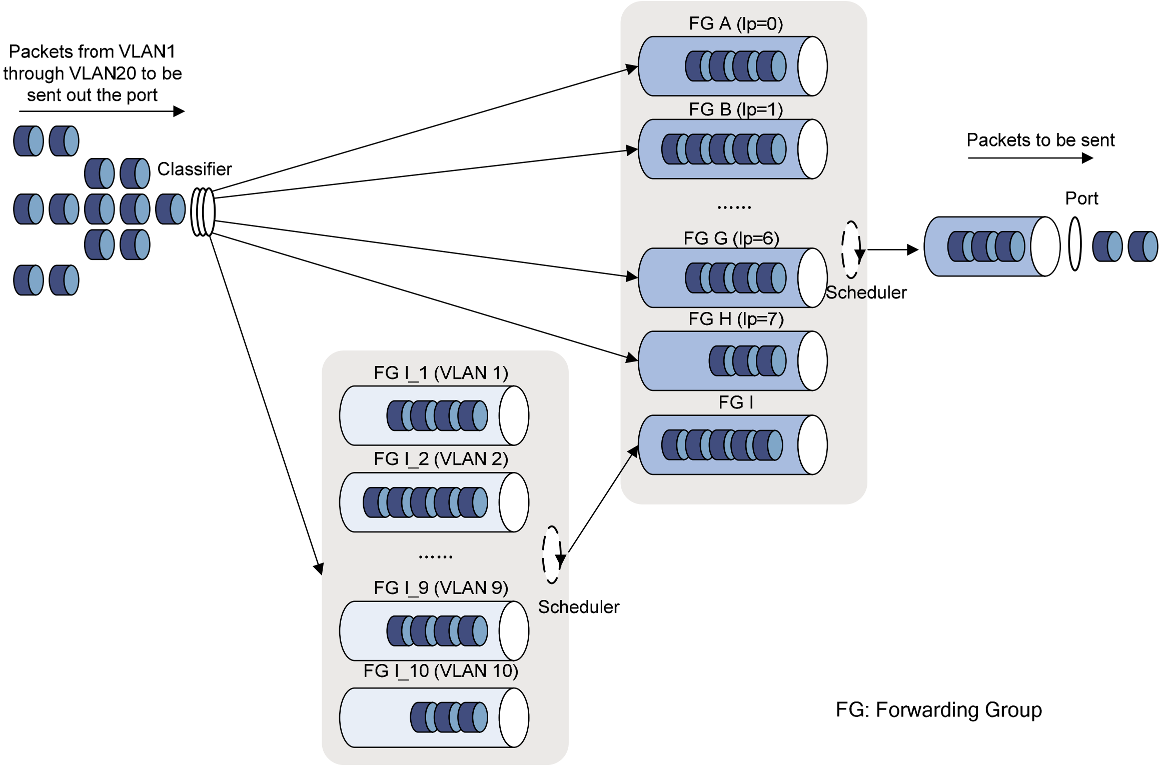

HQoS uses forwarding groups in place of queues in common QoS implementations. It uses the SP or WRR algorithm to schedule forwarding groups.

In this example, the port forwards packets for VLAN 1 through VLAN 20, among which VLAN 1 through VLAN 10 have higher priority. HQoS works on the port as follows:

1. Maps packets from VLAN 1 through VLAN 10 to forwarding groups FG I_1 through FG I_10, dequeues packets from the forwarding groups according to the scheduling algorithm, and then delivers the dequeued packets to forwarding group FG I. This scheduling process is called "layer-2 scheduling."
2. Maps packets from VLANs 11 through VLAN 20 to forwarding groups FG A through FG H according to their local precedence values.
3. Schedules forwarding groups FG A through FG H together with forwarding group FG I. This scheduling process is called "layer-1 scheduling." To preferentially forward packets from VLANs 1 through 10, assign forwarding group FG I a higher scheduling priority than other forwarding groups. Then, the layer-1 scheduler prioritizes the packets from VLAN 1 through 10, and schedules packets from the other VLANs according to their priorities.

In this example, forwarding group FG I is a parent forwarding group. It contains packets delivered from other forwarding groups, FG I_1 through FG I_10 (in this example), which are called "child forwarding groups."

NOTE:

- The HP 5800 Switch Series supports only one parent forwarding group on a port. A parent group can nest up to 16 child forwarding groups.
- You can schedule HQoS forwarding groups together with QoS queues. In this example, you can also assign packets from VLANs 11 through 20 to QoS queues through their 802.1p priority or local precedence. For more information, see "Instantiating a forwarding group."

Figure 27 shows a simplified tree-shaped diagram for this HQoS scheduling example.

Figure 27 How HQoS works

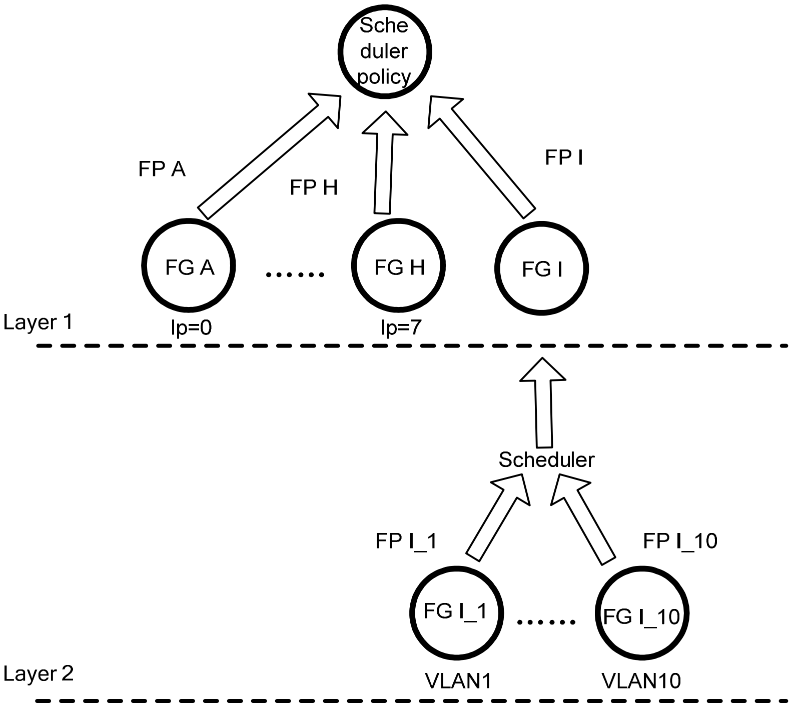

As shown in Figure 27, HQoS first performs layer-2 scheduling and then performs layer-1 scheduling for the scheduled packets together with other packets.

HQoS enables you to achieve hierarchical traffic management by configuring match criteria for each forwarding group at each layer to classify traffic by user or traffic type.

You can configure a set of forwarding and scheduling actions (including the scheduling algorithm, scheduling weight, and traffic shaping parameters) for each forwarding group at layer 2 of the scheduling hierarchy. For their parent forwarding group at the layer 1 of the scheduling hierarchy, you can configure different traffic control parameters.

HQoS concepts

The section describes key terms and concepts of the HQoS configuration procedure.

Forwarding group

A forwarding group represents a class of traffic with certain traffic characteristics (VLAN ID or local precedence value). All traffic assigned to a forwarding group shares a common set of scheduling and forwarding parameters.

Forwarding profile

A forwarding profile comprises a set of traffic control actions, including the queue scheduling algorithm, scheduling weight, and traffic shaping parameters. You can associate a forwarding profile with a forwarding group to set the scheduling priority and bandwidth resources for the forwarding group.

Scheduler policy

A scheduler policy is a set of forwarding group and forwarding profile associations. To implement HQoS on a port, apply a scheduler policy to the port.

Scheduling layer

The scheduling layer is where a forwarding group is scheduled by HQoS. The HQoS scheduling hierarchy comprises layer 1 and layer 2, as shown in Figure 27.

Instantiation

Instantiation maps packets to different forwarding groups according to match criteria, for example, maps packets from VLAN 1 through VLAN 10 to forwarding groups FG I_1 through FG I_10, as shown in Figure 27. You can classify traffic and map them to different forwarding groups by configuring instantiation rules (match criteria). Then, HQoS schedules packets in a forwarding group according to the forwarding profile associated with the forwarding group.

Configuration task list

HQoS configuration organizes a tree-shaped scheduler policy model. First create forwarding groups, then nest forwarding groups and specify a forwarding profile for each forwarding group, and finally, instantiate each forwarding group to assign traffic to forwarding groups.

Perform the following tasks to configure HQoS:

Task	Remarks
Configuring a forwarding profile	Required.
Configuring a forwarding group	Required. Nest forwarding groups and specify a forwarding profile for each child forwarding group.
Configuring a scheduler policy	Required. Create a scheduler policy, nest layer-1 forwarding groups in the scheduler policy, and specify a forwarding profile for each layer-1 forwarding group.
Instantiating a forwarding group	Required. Instantiate the forwarding groups and specify forwarding profiles for forwarding groups. For example, configure layer-1 forwarding group A1 to match traffic with local precedence 0 and use forwarding profile A1 to process the matching traffic, configure layer-2 forwarding group B1 to match traffic from VLAN 1, and use forwarding profile B1 to process the matching traffic.
Applying a scheduler policy to a port	Required.
Copying a forwarding group	Optional.

Task	Remarks
Copying a scheduler policy	Optional.

Configuring a forwarding profile

A forwarding file comprises the following contents:

- A queue scheduling method for its associated forwarding groups. Available scheduling algorithms include SP and WRR.
- Optionally, traffic shaping parameters and minimum guaranteed bandwidth. These parameters limit the traffic forwarding rate within a certain range for each forwarding group associated with the forwarding profile.

A forwarding profile configured with SP is called a "SP forwarding profile," and a forwarding profile configured with WRR is called a "WRR forwarding profile."

To modify the parameters of an existing user-defined forwarding profile:

To do…		Use the command…	Remarks
1. Enter system view.		**system-view**	—
2. Create a forwarding profile or enter the view of an existing forwarding profile.		**qos forwarding-profile** *fp-name* [**id** *fp-id*]	Required.
3. Configure a queue scheduling method:	Specify SP queue scheduling.	**sp**	Use either approach. By default, a forwarding profile is not configured with any queue scheduling method. IMPORTANT: For more information about SP and WRR, see "Configuring congestion management."
	Configure WRR queue scheduling.	**wrr** [**weight** *weight-value*]	
4. Configure GTS.		**gts cir** *cir-value* [**cbs** *cbs-value*] [**ebs** *ebs-value*] [**pir** *pir-value*]	Optional. By default, a forwarding profile has no GTS parameter for rate limiting traffic.
5. Configure the minimum guaranteed bandwidth.		**bandwidth** *bandwidth-value*	Optional. By default, a forwarding profile has no minimum guaranteed bandwidth.

NOTE:

- The WRR queue scheduling algorithm configured in a forwarding profile schedules packets in the unit of bytes. The WRR scheduling weight has the same meaning as that for a generic WRR queue. The HQoS forwarding groups can be scheduled together with generic WRR queues with byte-based weights. WRR directly uses the weights of the HQoS forwarding groups for scheduling packets. For more information about how forwarding groups are scheduled together with generic WRR queues, see "Applying a scheduler policy to a port." For more information about the scheduling units of WRR queues, see "Configuring congestion management."
- You cannot modify a forwarding profile that has been associated with one or more forwarding groups, if the modification is invalid for any of its associated forwarding group.
- When you modify a forwarding profile already applied to a port, the modification may fail because of insufficient hardware resources.

Configuring a forwarding group

A forwarding group is a basic scheduling entity in a scheduler policy and is also an instantiation object.

Forwarding group configuration tasks include the following steps:

- Creating a forwarding group
- Nesting a forwarding group

Creating a forwarding group

To configure a forwarding group:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Create a forwarding group.	**qos forwarding-group** *fg-name* [**id** *fg-id*]	Required You can also use this command to enter the view of an existing forwarding group

Nesting a forwarding group

To achieve two-level traffic scheduling for some traffic, nest several forwarding groups in a parent forwarding group, for example, forwarding groups FG I_1 through FG I_10 in forwarding group FG I in Figure 26.

In a parent forwarding group, associate each child forwarding group with an SP or WRR forwarding profile.

If some child forwarding groups use SP scheduling and others use WRR scheduling, HQoS schedules traffic for the child forwarding groups that use SP scheduling first, and then schedules traffic for groups associated with WRR forwarding profiles in a round robin fashion.

When performing SP scheduling, HQoS schedules traffic for the child forwarding groups in the reverse configuration order: the last nested child forwarding group is scheduled first.

NOTE:

The HP 5800 Switch Series supports only one parent group on a port. A parent group can nest up to 16 child forwarding groups.

To nest a forwarding group:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Enter forwarding group view.	**qos forwarding-group** *fg-name* [**id** *fg-id*]	—
3. Nest a child forwarding group in the parent forwarding group and specify a forwarding profile for the child forwarding group.	**forwarding-group** *sub-fg-name* **profile** *fp-name*	Required. Repeat this step to nest multiple child forwarding groups in the parent forwarding group. You can also use this command to change forwarding profile for a child forwarding group that has been nested.

NOTE:

- If the forwarding profile assigned for it contains parameters that conflict with the forwarding group, the child forwarding group will fail to nest.
- A forwarding group cannot nest itself.

Configuring a scheduler policy group

A scheduler policy is a set of forwarding group and forwarding profile associations. By organizing these associations and specify a scheduling layer for each association, you set up a tree-shaped HQoS scheduler policy.

Configuration guidelines

You can nest a forwarding group in multiple scheduler policies, and associate it with different forwarding profiles.

In a scheduler policy, you associate each forwarding group with an SP or WRR forwarding profile.

If some forwarding groups use SP scheduling and others use WRR scheduling, HQoS first schedules traffic for the forwarding groups that use SP scheduling and then schedules traffic for groups associated with WRR forwarding profiles in a round robin fashion.

When performing SP scheduling, HQoS schedules traffic for the forwarding groups in the ascending order of their local precedence values specified in the instantiation rules (match criteria). For more information about instantiation, see "Instantiating a forwarding group." If the scheduler policy nests a parent forwarding group that uses SP scheduling, HQoS always schedules this forwarding group prior to any other forwarding group.

Configuration procedure

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Create a scheduler policy or enter the view of an existing scheduler policy.	**qos scheduler-policy** *sp-name* [**id** *sp-id*]	Required.
3. Nest a forwarding group in the scheduler policy and specify a forwarding profile for the forwarding group.	**forwarding-group** *fg-name* **profile** *fp-name*	Required. Repeat this step to nest multiple forwarding groups in the scheduler policy. You can also use this command to change forwarding profile for a forwarding group that has been nested.

NOTE:

- You will fail to nest a forwarding group, if the forwarding profile assigned for it contains parameters that conflict with the forwarding group.
- A scheduler policy can nest up to nine forwarding groups. Each forwarding group nested in a scheduler policy must be unique.

Instantiating a forwarding group

When instantiating a forwarding group, you can perform the following operations:

- Specify its scheduling layer in a scheduler policy (see Figure 27 for reference).
- Classify traffic by user (VLAN) or service type (packet priority).

Instantiation modes

Instantiate forwarding groups in one of these two instantiation modes:

Instantiation mode	Applicable to	Configuration guidelines
Match mode	• Child forwarding groups • Forwarding groups that have no child forwarding groups	You must specify instantiation rules. Available match criteria include local precedence and service provider VLAN IDs.
Group mode	Parent forwarding groups	You do not need to specify match criteria for a parent forwarding group. However, you must configure match criteria for its child forwarding groups in match mode.

Configuration guidelines

Use the following table when instantiating forwarding groups for a scheduler policy:

Configuration items (right) **Forwarding group type (below)**	**Scheduling layer**	**Instantiation mode**	**Match criteria**	**Remarks**
Forwarding groups that have no children (Childless forwarding groups)	Layer 1	Match mode	Local precedence	The local precedence values must be unique within the scheduler policy.
The parent forwarding group	Layer 1	Group mode	None	Instantiate the parent forwarding group before its child groups. To cancel instantiation for a parent forwarding group, first cancel instantiation for all child forwarding groups in it.
Child forwarding groups	Layer 2	Match mode	Service provider VLAN IDs	The VLAN IDs must be unique within the scheduler policy.

NOTE:

The childless forwarding groups of HQoS share resources with QoS queues on a port. After you apply a scheduler policy to a port, each childless forwarding group in the policy replaces the QoS queue with the same local precedence on the port. QoS queues and HQoS forwarding groups are scheduled together at layer 1. The QoS queues must adopt WRR scheduling.

Configuration procedure

To do…			**Use the command…**	**Remarks**
1. Enter system view.			**system-view**	—
2. Enter scheduler policy view.			**qos scheduler-policy** *sp-name* [**id** *sp-id*]	—
3. Enter scheduling layer view.			**layer** { **1** | **2** }	—
4. Instantiate a forwarding group:	In match mode.	Layer 1	**forwarding-group** *fg-name* **match local-precedence** *local-precedence*	Required. Use either approach.
		Layer 2	**forwarding-group** *fg-name* **match service-vlan-id** { *vlan-id-list* | *vlan-id1* **to** *vlan-id2* } }	
	In group mode.		**forwarding-group** *fg-name* **group**	

NOTE:

If you perform the **forwarding-group match** command multiple times for a forwarding group, only the first configuration applies. To change the match criteria for a forwarding group that has been instantiated, first remove the instantiation from the scheduler policy.

Applying a scheduler policy to a port

To use a scheduler policy to control traffic on a port, you must apply the scheduler policy to the port. On the HP 5800 Switch Series, only some ports support HQoS.

HQoS-capable ports on the HP 5800 Switch Series

Switch model	HQoS-capable ports
• 5800-24G Switch (JC100A) • 5800-24G TAA Switch (JG255A) • 5800-24G-PoE+ Switch (JC099A) • 5800-24G-PoE+TAA Switch (JG254A) • 5800-48G Switch with 1 Interface Slot (JC105A) • 5800-48G TAA Switch with 1 Interface Slot (JG258A) • 5800-48G-PoE+ Switch with 1 Interface Slot (JC104A) • 5800-48G-PoE+ TAA Switch with 1 Interface Slot (JG257A) • 5800-24G-SFP Switch with 1 Interface Slot (JC103A) • 5800-24G-SFP TAA Switch with 1 Interface Slot (JG256A)	• The four fixed 10GE ports on the front panel • 10GE ports on interface card LSW1SP4P0/LSW1SP2P0 • GE ports numbered 3, 4, 7, 8, 11, 12, 15, and 16 on interface card LSW1GP16P0/LSW1GT16P
• 5800-48G-PoE+ Switch with 2 Interface Slots (JC101A) • 5800-48G-PoE+ TAA Switch with 2 Interface Slots (JG242A)	• Port GigabitEthernet 1/0/49 • 10GE ports on interface card LSW1SP4P0/LSW1SP2P0 • GE ports numbered 3, 4, 7, 8, 11, 12, 15, and 16 on interface card LSW1GP16P0/LSW1GT16P
5800AF-48G Switch (JG225A)	The six fixed 10GE ports on the front panel

Configuration guidelines

You cannot apply a scheduler policy to a port that has been configured with any of these QoS features: traffic shaping, congestion avoidance, and queue scheduling algorithm other than the default WRR algorithm. Also, you cannot configure any of these QoS features on a port if a scheduler policy has been applied to it.

You can apply a scheduler policy only to the outbound direction of a port.

Configuration prerequisites

Instantiate each forwarding group nested in the scheduler policy.

Configuration procedure

To apply a scheduler policy to a port:

To do…		Use the command…	Remarks
1. Enter system view.		**system-view**	—
2. Enter port view or port group view:	Enter port view.	**interface** *interface-type interface-number*	Use either approach. To apply a scheduler policy to one port, enter port view. To apply it to a group of ports, enter port group view.
	Enter port group view.	**port-group manual** *group-name*	
3. Apply a scheduler policy.		**qos apply scheduler-policy** *sp-name* **outbound**	Required. You can apply only one scheduler policy to a port.

NOTE:

In a scheduler policy that has been applied to a port, you can modify or replace the forwarding profile associated with a forwarding group, but you cannot add or remove forwarding groups. To add or remove forwarding groups, or change the instantiation rules, remove the scheduler policy from the port first.

Copying a forwarding group

You can copy a forwarding group (the source forwarding group) to create multiple destination forwarding groups, which are automatically numbered.

The destination forwarding groups have the same configurations as the source forwarding group except for their group name and number.

The copying process stops when the maximum number of forwarding groups (90), on the switch is exceeded.

To copy a forwarding group:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Copy a forwarding group.	**qos copy forwarding-group** *fg-source* **to** *fg-dest*&<1-8>	Required

Copying a scheduler policy

You can copy a scheduler policy (called "source scheduler policy") to create one destination scheduler policy, which is automatically numbered.

The destination scheduler policy has the same configuration as the source scheduler policy (including the nested forwarding groups and their instantiation rules) except the group name and number.

To copy a scheduler policy:

To do…	Use the command…	Remarks
1. Enter system view.	**system-view**	—
2. Copy a scheduler policy.	**qos copy scheduler-policy** *sp-source* **to** *sp-dest*	Required

Displaying and maintaining HQoS

To display different HQoS information complete the following tasks:

Task	Command	Remarks
Display forwarding group information.	**display qos forwarding-group** [*fg-name*] [| { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display forwarding profile information.	**display qos forwarding-profile** [*fp-name*] [| { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display scheduler policy information.	**display qos scheduler-policy name** [*sp-name*] [| { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display scheduler policy information on a port.	**display qos scheduler-policy interface** [*interface-type interface-number* [**outbound**]] [| { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view
Display scheduler policy debugging information on a port.	**display qos scheduler-policy diagnosis interface** [*interface-type interface-number* [**outbound**]] [| { **begin** | **exclude** | **include** } *regular-expression*]	Available in any view

Configuration example

Network requirements

The metropolitan area network, shown in Figure 28, assigns traffic of different user groups to different VLANs, as shown in Table 8. The network does not differentiate personal users. It applies the same service and traffic management scheme for all personal users. The network differentiates business users A and B, and applies different service and traffic management scheme for them.

Table 8 VLAN assignment and traffic management scheme

Item	VLAN assignment and traffic management scheme
Traffic policy for personal user group A	Assign service VLANs 100 through 299 for personal user group A. • **VoIP**—High priority, 802.1p priority 6 or 7. Set the minimum guaranteed bandwidth to 20 Mbps for the traffic with 802.1p 6, and 30 Mbps for traffic with 802.1p 7. • **VoD**—Middle priority, 802.1p priority 4 or 5. • **VPN**—Normal priority, 802.1p priority 2 or 3. • **Web service**—Low priority, 802.1p priority 0 or 1.
Traffic policy for business users	• VLANs 501 through 503 for business user A. • VLANs 504 through 506 for business user B. Rate-limit all business user traffic to 100 Mbps, and assign 50 Mbps of minimum guaranteed bandwidth for all business user traffic.

Configure Switch **Core** as follows:

- Always schedule business traffic before personal user traffic.
- Schedule traffic for business users A and B using a ratio of 2:1.
- For personal user traffic, always schedule VoIP traffic first, and then schedule VoD, VPN, and Web traffic in turn using a ratio of 3:2:1, three VoD packets, two VPN packets, and one HTTP packet in each poll.

Figure 28 A service VLAN-mode metropolitan area network

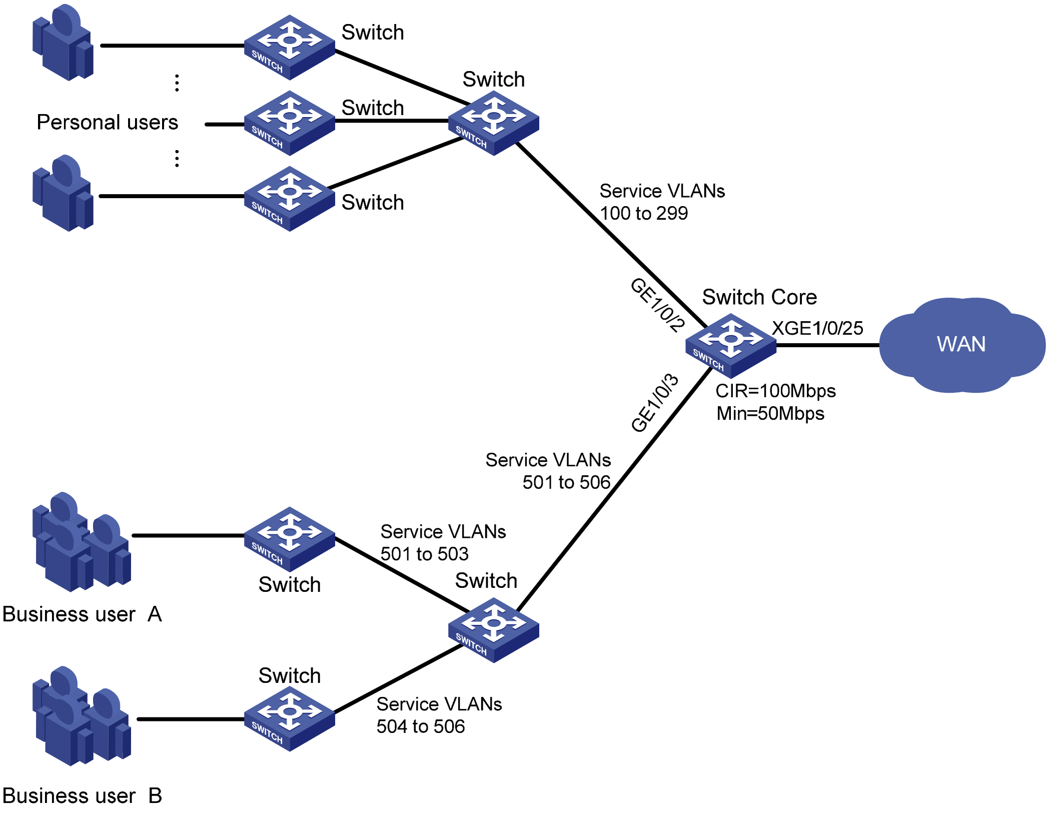

Configuration considerations

According to the network requirements, business user traffic is preferentially sent regardless of traffic types, and personal user traffic is scheduled by priority. Configure HQoS as follows:

- Assign all business user traffic to a parent forwarding group. The group nests six child forwarding groups, which match the six VLANs for business user A and business user B. Use WRR to schedule traffic in the required ratio for these child forwarding groups. Finally, associate the parent forwarding group with an SP forwarding profile to prioritize the business traffic, and configure GTS and minimum guaranteed bandwidth to rate-limit the traffic.
- For personal user traffic, map different 802.1p priority values to different local precedence values, and allocate them to the eight forwarding groups. Prioritize VoIP traffic by associating the forwarding group for VoIP traffic with an SP forwarding profile configured with minimum guaranteed bandwidth. Associate other forwarding groups with WRR forwarding profiles configured with different scheduling weights to implement proportional scheduling.

Figure 29 shows the HQoS structure for this example.

Figure 29 HQoS structure

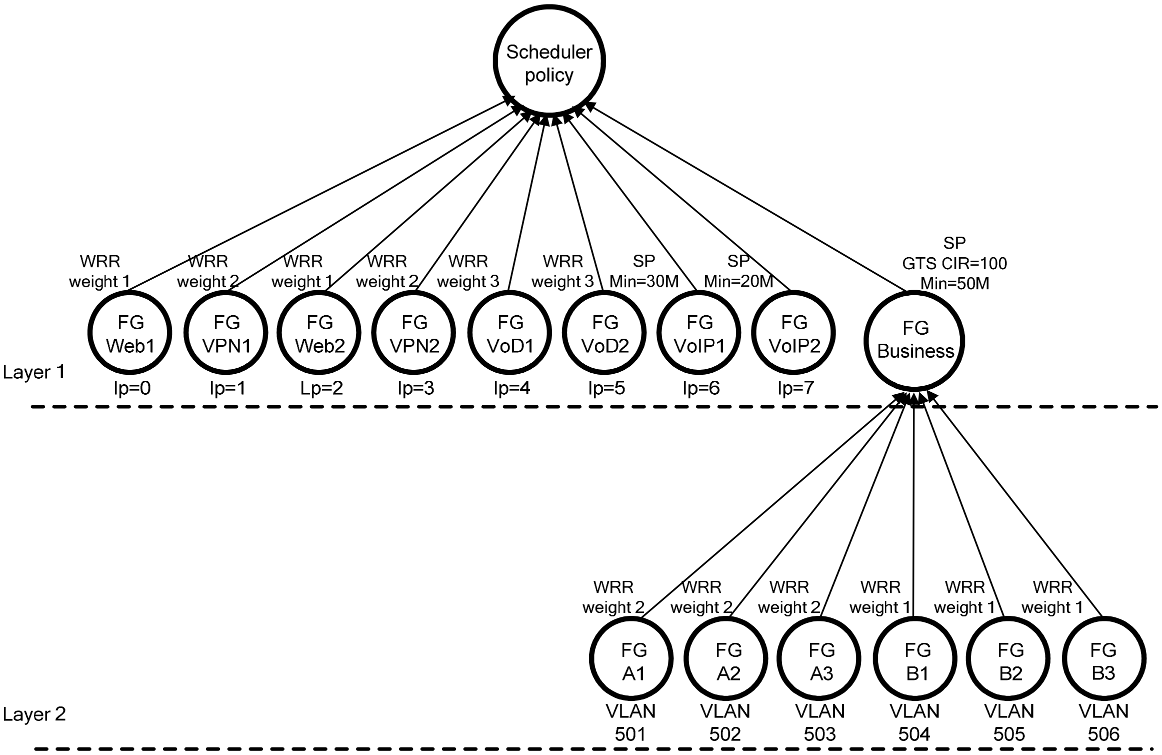

NOTE:

According to the dot1p-lp priority mapping table, 802.1p priority values 0 and 1 (Web traffic) are mapped to local precedence values 0 and 2, 802.1p priority values 2 and 3 (VPN traffic) are mapped to local precedence values 1 and 3, and other 802.1p priority values are mapped to their identical local precedence values.

Configuration procedure

1. Create child forwarding groups for business users.

```
<Sysname> system-view
[Sysname] qos forwarding-group A1
```

```
[Sysname-hqos-fg-A1] quit
[Sysname] qos forwarding-group A2
[Sysname-hqos-fg-A2] quit
[Sysname] qos forwarding-group A3
[Sysname-hqos-fg-A3] quit
[Sysname] qos forwarding-group B1
[Sysname-hqos-fg-B1] quit
[Sysname] qos forwarding-group B2
[Sysname-hqos-fg-B2] quit
[Sysname] qos forwarding-group B3
[Sysname-hqos-fg-B3] quit
```

2. Configure forwarding profiles for the child forwarding groups for the business users, according to Table 8.

TIP:

You only need to configure one forwarding profile for each business users, because you can associate a forwarding profile with multiple forwarding groups.

```
[Sysname] qos forwarding-profile A
[Sysname-hqos-fp-A] wrr weight 2
[Sysname-hqos-fp-A] quit
[Sysname] qos forwarding-profile B
[Sysname-hqos-fp-B] wrr weight 1
[Sysname-hqos-fp-B] quit
```

3. Create a parent forwarding group for the business users, nest the child forwarding groups in the parent forwarding group, and then configure forwarding profiles for the child forwarding groups.

```
[Sysname] qos forwarding-group business
[Sysname-hqos-fg-business] forwarding-group A1 profile A
[Sysname-hqos-fg-business] forwarding-group A2 profile A
[Sysname-hqos-fg-business] forwarding-group A3 profile A
[Sysname-hqos-fg-business] forwarding-group B1 profile B
[Sysname-hqos-fg-business] forwarding-group B2 profile B
[Sysname-hqos-fg-business] forwarding-group B3 profile B
[Sysname-hqos-fg-business] quit
```

4. Create a forwarding profile for the parent forwarding group.

```
[Sysname] qos forwarding-profile business
[Sysname-hqos-fp-business] sp
[Sysname-hqos-fp-business] gts cir 102400
[Sysname-hqos-fp-business] bandwidth 51200
[Sysname-hqos-fp-business] quit
```

5. Create forwarding groups for personal users.

```
[Sysname] qos forwarding-group VoIP1
[Sysname-hqos-fg-VoIP1] quit
[Sysname] qos forwarding-group VoIP2
[Sysname-hqos-fg-VoIP2] quit
[Sysname] qos forwarding-group VoD1
[Sysname-hqos-fg-VoD1] quit
[Sysname] qos forwarding-group VoD2
```

```
[Sysname-hqos-fg-VoD2] quit
[Sysname] qos forwarding-group VPN1
[Sysname-hqos-fg-VPN1] quit
[Sysname] qos forwarding-group VPN2
[Sysname-hqos-fg-VPN2] quit
[Sysname] qos forwarding-group WEB1
[Sysname-hqos-fg-WEB1] quit
[Sysname] qos forwarding-group WEB2
[Sysname-hqos-fg-WEB2] quit
```

6. Configure forwarding profiles for the forwarding groups for the personal users, according to Table 8.

TIP:

You must configure two forwarding profiles that use different minimum guaranteed bandwidth settings for VoIP traffic. However, you only need to configure one forwarding profile for each type of other traffic.

```
[Sysname] qos forwarding-profile VoIP1
[Sysname-hqos-fp-VoIP1] sp
[Sysname-hqos-fp-VoIP1] bandwidth 30720
[Sysname-hqos-fp-VoIP1] quit
[Sysname] qos forwarding-profile VoIP2
[Sysname-hqos-fp-VoIP2] sp
[Sysname-hqos-fp-VoIP2] bandwidth 20480
[Sysname-hqos-fp-VoIP2] quit
[Sysname] qos forwarding-profile VoD
[Sysname-hqos-fp-VoD] wrr weight 3
[Sysname-hqos-fp-VoD] quit
[Sysname] qos forwarding-profile VPN
[Sysname-hqos-fp-VPN] wrr weight 2
[Sysname-hqos-fp-VPN] quit
[Sysname] qos forwarding-profile WEB
[Sysname-hqos-fp-WEB] wrr weight 1
[Sysname-hqos-fp-WEB] quit
```

7. Create a scheduler policy, nest the forwarding groups for the personal users and the parent forwarding group for the business users in the scheduler policy, and specify a forwarding profile for each forwarding group.

```
[Sysname] qos scheduler-policy hqos
[Sysname-hqos-sp-hqos] forwarding-group VoIP1 profile VoIP1
[Sysname-hqos-sp-hqos] forwarding-group VoIP2 profile VoIP2
[Sysname-hqos-sp-hqos] forwarding-group VoD1 profile VoD
[Sysname-hqos-sp-hqos] forwarding-group VoD2 profile VoD
[Sysname-hqos-sp-hqos] forwarding-group VPN1 profile VPN
[Sysname-hqos-sp-hqos] forwarding-group VPN2 profile VPN
[Sysname-hqos-sp-hqos] forwarding-group WEB1 profile WEB
[Sysname-hqos-sp-hqos] forwarding-group WEB2 profile WEB
[Sysname-hqos-sp-hqos] forwarding-group business profile business
```

8. At layer 1 of the scheduler policy, configure instantiation rules for the forwarding groups for the personal users and the parent forwarding group for the business users.

```
[Sysname-hqos-sp-hqos] layer 1
[Sysname-hqos-sp-hqos-layer1] forwarding-group VoIP2 match local-precedence 7
[Sysname-hqos-sp-hqos-layer1] forwarding-group VoIP1 match local-precedence 6
[Sysname-hqos-sp-hqos-layer1] forwarding-group VoD2 match local-precedence 5
[Sysname-hqos-sp-hqos-layer1] forwarding-group VoD1 match local-precedence 4
[Sysname-hqos-sp-hqos-layer1] forwarding-group VPN2 match local-precedence 3
[Sysname-hqos-sp-hqos-layer1] forwarding-group VPN1 match local-precedence 1
[Sysname-hqos-sp-hqos-layer1] forwarding-group WEB2 match local-precedence 2
[Sysname-hqos-sp-hqos-layer1] forwarding-group WEB1 match local-precedence 0
[Sysname-hqos-sp-hqos-layer1] forwarding-group business group
[Sysname-hqos-sp-hqos-layer1] quit
[Sysname-hqos-sp-hqos] quit
```

9. At layer 2 of the scheduler policy, configure instantiation rules for the six child forwarding groups for the business users.

```
[Sysname-hqos-sp-hqos] layer 2
[Sysname-hqos-sp-hqos-layer2] forwarding-group A1 match service-vlan-id 501
[Sysname-hqos-sp-hqos-layer2] forwarding-group A2 match service-vlan-id 502
[Sysname-hqos-sp-hqos-layer2] forwarding-group A3 match service-vlan-id 503
[Sysname-hqos-sp-hqos-layer2] forwarding-group B1 match service-vlan-id 504
[Sysname-hqos-sp-hqos-layer2] forwarding-group B2 match service-vlan-id 505
[Sysname-hqos-sp-hqos-layer2] forwarding-group B3 match service-vlan-id 506
[Sysname-hqos-sp-hqos-layer2] quit
```

10. Apply scheduler policy **hqos** to the outgoing traffic on port GigabitEthernet 1/0/25.

```
[Sysname] interface Ten-GigabitEthernet 1/0/25
[Sysname-Ten-GigabitEthernet1/0/25] qos apply scheduler-policy hqos outbound
```

Appendix A Default priority mapping tables

For the default **dot1p-exp**, **dscp-dscp**, and **exp-dot1p** mapping tables, an input value yields a target value equal to it.

Table 9 Default dot1p-lp and dot1p-dp priority mapping tables

Input priority value	dot1p-lp mapping	dot1p-dp mapping
802.1p priority (dot1p)	**Local precedence (lp)**	**Drop precedence (dp)**
0	2	0
1	0	0
2	1	0
3	3	0
4	4	0
5	5	0
6	6	0
7	7	0

Table 10 Default dscp-dp and dscp-dot1p priority mapping tables

Input priority value	dscp-dp mapping	dscp-dot1p mapping
DSCP	**Drop precedence (dp)**	**802.1p priority (dot1p)**
0 to 7	0	0
8 to 15	0	1
16 to 23	0	2
24 to 31	0	3
32 to 39	0	4
40 to 47	0	5
48 to 55	0	6
56 to 63	0	7

Table 11 Default exp-dp priority mapping table

Input priority value	exp-dp mapping
EXP value	Drop precedence (dp)
0	0
1	0
2	0

Input priority value	exp-dp mapping
3	0
4	0
5	0
6	0
7	0

Appendix B Introduction to packet precedence

IP precedence and DSCP values

Figure 30 ToS and DS fields

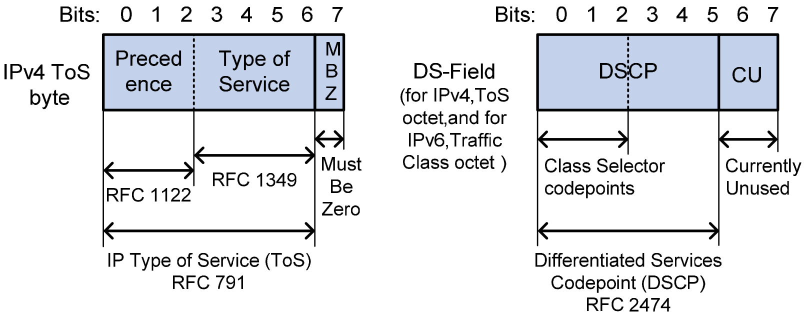
Bits: 0 1 2 3 4 5 6 7
IPv4 ToS byte
Preced ence
Type of Service
M B Z
RFC 1122
RFC 1349
Must Be Zero
IP Type of Service (ToS) RFC 791
Bits: 0 1 2 3 4 5 6 7
DS-Field (for IPv4,ToS octet,and for IPv6,Traffic Class octet)
DSCP
CU
Class Selector codepoints
Currently Unused
Differentiated Services Codepoint (DSCP) RFC 2474

As shown in Figure 30, the ToS field in the IP header contains eight bits. The first three bits (0 to 2) represent IP precedence from 0 to 7. According to RFC 2474, the ToS field is redefined as the DS field, where a DSCP value is represented by the first six bits (0 to 5), and range from 0 to 63. The remaining two bits (6 and 7) are reserved.

Table 12 Description on IP precedence

IP precedence (decimal)	IP precedence (binary)	Description
0	000	Routine
1	001	priority
2	010	immediate
3	011	flash
4	100	flash-override
5	101	critical
6	110	internet
7	111	network

Table 13 Description on DSCP values

DSCP value (decimal)	DSCP value (binary)	Description
46	101110	ef
10	001010	af11
12	001100	af12
14	001110	af13
18	010010	af21

DSCP value (decimal)	DSCP value (binary)	Description
20	010100	af22
22	010110	af23
26	011010	af31
28	011100	af32
30	011110	af33
34	100010	af41
36	100100	af42
38	100110	af43
8	001000	cs1
16	010000	cs2
24	011000	cs3
32	100000	cs4
40	101000	cs5
48	110000	cs6
56	111000	cs7
0	000000	be (default)

802.1p priority

802.1p priority lies in the Layer 2 header and applies to occasions where Layer 3 header analysis is not needed and QoS must be assured at Layer 2.

Figure 31 An Ethernet frame with an 802.1Q tag header

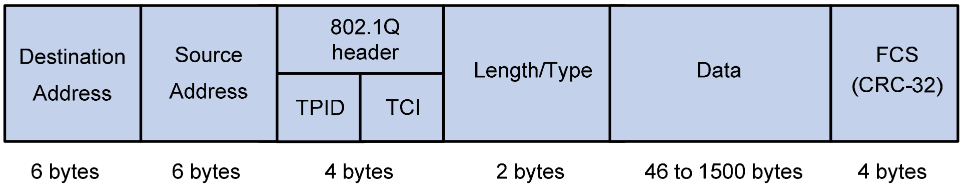

As shown in Figure 31, the four-byte 802.1Q tag header consists of the TPID (two bytes in length), whose value is 0x8100, and the TCI (two bytes in length). Figure 32 shows the format of the 802.1Q tag header. The Priority field in the 802.1Q tag header is called the "802.1p priority," because its use is defined in IEEE 802.1p. Table 14 shows the values for 802.1p priority.

Figure 32 802.1Q tag header

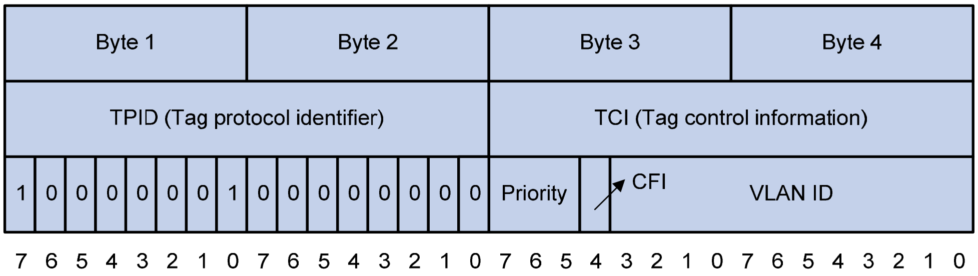

Table 14 Description on 802.1p priority

802.1p priority (decimal)	802.1p priority (binary)	Description
0	000	best-effort
1	001	background
2	010	spare
3	011	excellent-effort
4	100	controlled-load
5	101	video
6	110	voice
7	111	network-management

EXP values

The EXP field is in MPLS labels for MPLS QoS purposes.

Figure 33 MPLS label structure

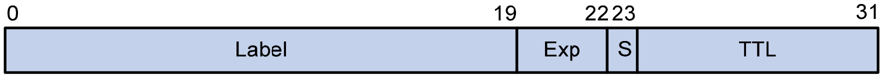

As shown in Figure 33, the EXP field is 3 bits long and ranges from 0 to 7.

Appendix C Support and other resources

Contacting HP

For worldwide technical support information, see the HP support website:

http://www.hp.com/support

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

http://www.hp.com/go/wwalerts

After registering, you will receive email notification of product enhancements, new driver versions, firmware updates, and other product resources.

Related information

Documents

To find related documents, browse to the Manuals page of the HP Business Support Center website:

http://www.hp.com/support/manuals

- For related documentation, navigate to the Networking section, and select a networking category.
- For a complete list of acronyms and their definitions, see *HP A-Series Acronyms.*

Websites

- HP.com http://www.hp.com
- HP Networking http://www.hp.com/go/networking
- HP manuals http://www.hp.com/support/manuals
- HP download drivers and software http://www.hp.com/support/downloads
- HP software depot http://www.software.hp.com

Conventions

This section describes the conventions used in this documentation set.

Command conventions

Convention	Description
Boldface	**Bold** text represents commands and keywords that you enter literally as shown.
Italic	*Italic* text represents arguments that you replace with actual values.
[ ]	Square brackets enclose syntax choices (keywords or arguments) that are optional.
{ x | y | ... }	Braces enclose a set of required syntax choices separated by vertical bars, from which you select one.
[x | y | ...]	Square brackets enclose a set of optional syntax choices separated by vertical bars, from which you select one or none.
{ x | y | ... } *	Asterisk-marked braces enclose a set of required syntax choices separated by vertical bars, from which you select at least one.
[x | y | ...] *	Asterisk-marked square brackets enclose optional syntax choices separated by vertical bars, from which you select one choice, multiple choices, or none.
&<1-n>	The argument or keyword and argument combination before the ampersand (&) sign can be entered 1 to n times.
#	A line that starts with a pound (#) sign is comments.

GUI conventions

Convention	Description
Boldface	Window names, button names, field names, and menu items are in bold text. For example, the **New User** window appears; click **OK**.
>	Multi-level menus are separated by angle brackets. For example, **File** > **Create** > **Folder**.

Symbols

Convention	Description
⚠ WARNING	An alert that calls attention to important information that if not understood or followed can result in personal injury.
△ CAUTION	An alert that calls attention to important information that if not understood or followed can result in data loss, data corruption, or damage to hardware or software.
❗ IMPORTANT	An alert that calls attention to essential information.
NOTE	An alert that contains additional or supplementary information.
💡 TIP	An alert that provides helpful information.

Network topology icons

	Represents a generic network device, such as a router, switch, or firewall.
	Represents a routing-capable device, such as a router or Layer 3 switch.
	Represents a generic switch, such as a Layer 2 or Layer 3 switch, or a router that supports Layer 2 forwarding and other Layer 2 features.

Port numbering in examples

The port numbers in this document are for illustration only and might be unavailable on your device.

Index