



White Paper

Quality of Service on Cisco Catalyst 4500 Series Supervisor Engines

INTRODUCTION

Quality of Service (QoS) provides the capability to differentiate among different classes of traffic and to prioritize the traffic in times of network congestion, according to its relative importance. The primary goals of QoS are to provide guaranteed minimum bandwidth for identified traffic, control jitter and latency (required by some real-time and interactive traffic) and to improve loss characteristics. This paper assumes the reader is familiar with basic QoS concepts such as classification, congestion avoidance, and bandwidth allocation. The objective of this paper is to provide an overview and configuration examples of the rich QoS capabilities available on the Cisco Catalyst 4500 Series supervisor engines.

These QoS capabilities include:

- Traffic Classification and Marking
- Ingress and Egress Policing
 - Individual
 - Aggregate
 - Microflow Policing
 - Hierarchical Policing
- Active Queue Management for congestion avoidance
 - Dynamic Buffer Limiting
- Per port Egress Queue Scheduling based on:
 - Bandwidth Sharing
 - Traffic Shaping
 - Strict Priority
- Auto QoS (introduced in Cisco IOS Software Release 12.1.19EW)

Figure 1. Cisco Catalyst 4500 Series Supervisor Engine QoS Processing

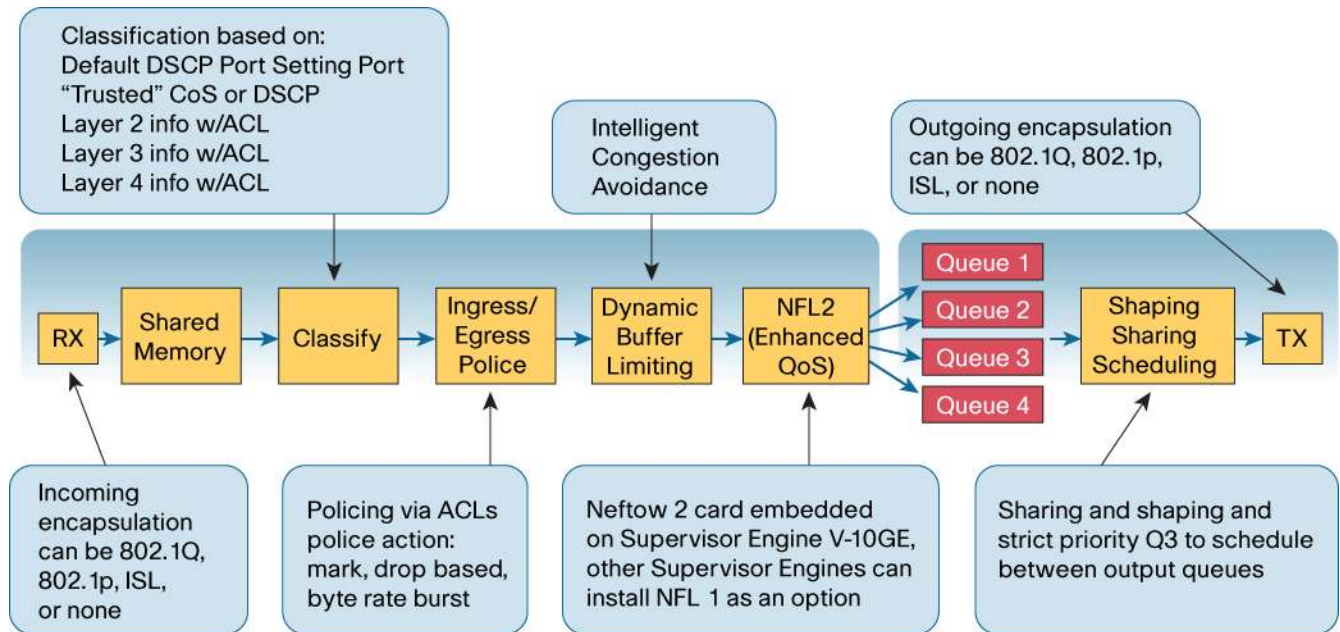


Figure 1 shows the different QoS stages implemented on the supervisor engines. The Cisco Catalyst 4500 Series uses a centralized switching architecture. As soon as a packet is received, it is placed into a shared memory buffer until the forwarding decision is made. The forwarding decision includes any QoS actions taken on the packet. When the forwarding decision has been made, the packet is rewritten including the appropriate Layer 2 or Layer 3 marking. The packet is then placed into the appropriate transmit queue based on Differentiated Services Code Point (DSCP) value and is scheduled according to the scheduling criteria configured on the interface queue: strict priority, sharing, or shaping. The actual QoS packet processing flow is depicted in Figure 1. However, for the purposes of clarity, this paper will discuss the beginning and the ending processes first: Traffic Classification and Egress Queue Scheduling followed by Policing and Dynamic Buffer Limiting.

TRAFFIC CLASSIFICATION AND MARKING

Ingress traffic is classified at Layers 3 and 4 for IP traffic as well as at Layer 2 for non-IP traffic or traffic coming in on trunk ports. Regardless of whether the incoming packet is Layer 2 or Layer 3, the Supervisor Engine QoS assigns an “internal DSCP” tag to the packet. This internal DSCP value is then used to place the packet into one of four queues for egress scheduling. This internal DSCP value can be derived in several ways. When QoS is globally disabled (default), the incoming packet’s DSCP value is trusted. However a Layer 2 packet’s class-of-service (CoS) value would be considered untrusted and a CoS value of zero would be used. This CoS value of 0 also maps to a DSCP value of zero. Example 3 shows CoS to DSCP mappings. The egress port queue selection is then based on the global DSCP-to-Transmit (Tx) queue mapping as shown in Example 1.

Example 1. Global DSCP to Egress Port Queue Mapping

```
r3_4507R_S4#sh qos maps dscp tx-queue
DSCP-TxQueue Mapping Table (dscp = d1d2)
d1 : d2  0  1  2  3  4  5  6  7  8  9
-----
0 :    01 01 01 01 01 01 01 01 01 01
1 :    01 01 01 01 01 01 02 02 02 02
2 :    02 02 02 02 02 02 02 02 02 02
3 :    02 02 03 03 03 03 03 03 03 03
4 :    03 03 03 03 03 03 03 03 04 04
5 :    04 04 04 04 04 04 04 04 04 04
6 :    04 04 04 04
```

For DSCP of 46 Tx Queue is 3

For example, if the internal DSCP tag is 46 (D1= 4, D2=6), Tx queue 3 would be used on egress. With QoS globally disabled, the four Tx queues are serviced Round-Robin. The DSCP to Tx queue mappings are fully configurable and again are enabled on a global basis.

With QoS globally enabled, the internal DSCP value used to determine the egress port queue can be derived in several ways:

1. **CoS/DSCP port trust state:** With QoS globally enabled, the default port trust state for CoS and DSCP is untrusted. This can be changed to trust either DSCP or CoS using the interface command. Trusting the incoming DSCP value is shown in Example 2.

Example 2. Changing the Trust State of a Port

```
r3_4507R_S4(config-if)#qos trust dscp
```

When the incoming CoS value is trusted, for non-IP traffic or traffic arriving on a trunk port, the internal DSCP value used for egress scheduling depends on the CoS to DSCP table mapping, which is configurable. The default mapping is can be shown in Example 3.

Example 3. CoS-to-DSCP Mapping

```
r3_4507R_S4#sh qos maps CoS dscp
CoS-DSCP Mapping Table
CoS:  0  1  2  3  4  5  6  7
-----
DSCP: 0  8 16 24 32 40 48 56
```

Trust can also be extended to Cisco IP phones attached to the port as shown in Example 4 (this will be done automatically using Auto QoS, which is discussed later):

Example 4. Trusting a Cisco IP Phone

```
r3_4507R_S4(config-if)#qos trust device cisco-phone
```

- 2. CoS/DSCP value on the ingress port:** With QoS globally enabled, the second way to derive the internal DSCP value used to determine the egress port queue is by using the DSCP or CoS value assigned to the ingress port. The default setting for CoS and DSCP is zero. This may be changed as shown in Example 5.

Example 5. Changing the Port DSCP/CoS Setting

```
r3_4507R_S4(config-if)#qos CoS ?
<0-7> CoS value
r3_4507R_S4(config-if)#qos dscp ?
<0-63> DSCP value
```

- 3. Using access control lists (ACLs) and service policy class maps:** With QoS globally enabled, the third way to derive the internal DSCP value that is used to determine the egress port queue is by using a service policy. Service policies are configured using Cisco's Modular QoS CLI (MQC) where a class map is used to identify the traffic of interest, a policy is defined and then assigned to an interface (Example 6). Service policy configuration details can be found in the documentation referenced at the end of this paper.

Example 6. Using a Service Policy to Mark a Packet

```
access-list 100 permit udp any any
mac access-list extended classify permit host 1234.5678.9abc any protocol-family ipx
class-map match-all class_setprec
  match ip access group 100
class-map match-all mac
  match access-group name classify
policy-map pol_setprec
  class class_setprec
    set ip dscp 26
  class mac
    set ip dscp 24
class mac
  set ip dscp 24
interface vlan 2
service-policy input pol_setprec
```

In Example 6, the **policy-map pol_setprec** was defined as an input policy. All incoming UDP packets will be classified with a DSCP value of 26. Classification can also be applied using Layer 2 ACLs for Layer 2 or non-IP traffic. In this example, Internetwork Packet Exchange (IPX) packets from a host are being classified with a DSCP of 24. If the service policy was applied as an output policy, the resultant marking would be used by the next hop QoS enabled device. Policy maps may be applied to individual interfaces or on a VLAN wide basis.

VLAN BASED QoS

VLAN based QoS is typically used for trunk ports. Prior to Cisco IOS Software Release 12.2(25)EWA, VLAN based QoS was the only QoS feature used when a single trunk port was part of the system. With Release 12.2(25)EWA, per-Port per-VLAN QoS was introduced, allowing for multiple trunk ports to be defined in the system and have further VLAN QoS granularity configurations. This is discussed in the next section.

The default QoS behavior is per port. Applying a service policy to a VLAN requires changing this default behavior as shown in Example 7.

Example 7. Changing the default QoS port behavior to VLAN Based

```
Switch(config)# interface FastEthernet3/2
Switch(config-if)# qos vlan-based
```

Also note that a policy map will override a port trust status. Note that port-based and VLAN based QoS provides flexibility in defining QoS policies. With VLAN based QoS, if one service policy is attached to a port and another to an SVI, the SVI service-policy takes precedence. Note, however, if port-based QoS is specified on certain ports within a VLAN and VLAN based QoS is configured on other ports, the port-based service-policy will take precedence over the SVI service-policy. This provides QoS flexibility because the more general service-policy may be defined for the VLAN and exception service-policies would be defined on the port. Examples 8 and 9 contain a traditional VLAN based QoS configuration on an access port and trunk port.

Example 8. Traditional VLAN Based QoS

```
Switch(config)# interface FastEthernet3/2
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 2
Switch(config-if)# qos vlan-based
Switch(config-if)# service-policy input P1
!
Switch(config-)# int vlan 2
Switch(config-if)# service-policy input P2
```

In Example 8, FastEthernet 3/2 has been configured as a QoS VLAN based interface. By doing this, service policy input P2 configuration will take precedence over service policy P1. If a service policy is not defined under the interface vlan, the service policy defined at the port level will take effect.

Example 9. VLAN Based QoS on a Trunk Interface.

```
Switch# configure terminal
Switch(config)# interface gigabitethernet 3/3
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# qos vlan-based
Switch(config-if)# exit
Switch(config)# interface vlan 2
Switch(config-if)# service-policy input P2
Switch(config-if)# interface vlan 3
Switch(config-if)# service-policy input P3
Switch(config-if)# interface vlan 4
Switch(config-if)# service-policy input P4
Switch(config-if)# exit
```

In Example 9, FastEthernet 3/3 has been configured as a trunk interface with interface VLANs 2–4. Each VLAN has a different service policy defined providing good granularity. However, this feature is limited to only one trunk interface per system. For example, if FastEthernet 3/4 is configured as a trunk interface as well, VLANs 2–4 under this interface would also be subjected to the same service policies defined for VLANs 2–4 in interface FastEthernet 3/3. Per Port per VLAN QoS is an extension to the traditional VLAN based QoS feature, which provides the network administrator the capability to have identical VLANs with different QoS configurations.

PER-PORT PER-VLAN QoS

As stated above, with per-Port per-VLAN QoS, different QoS configurations can be specified on different VLANs for a given set of trunk interfaces. For example, vlan 10 which exists in trunk interface 3/2 and 3/3 can have different service policies P1 and P2 assigned respectively. Per Port Per VLAN QoS is most often used by service providers wanting to provide guaranteed service level agreements (SLAs) with great QoS.

Example 10 shows how to configure multiple trunk interfaces (GigabitEthernet 6/1 and GigabitEthernet 6/2) with different service policies per QoS VLAN.

Example 10. Configuring Multiple Trunk Interfaces with Different Service Policies

```
Switch# configure terminal
Switch(config)# interface gigabitethernet 6/1
Switch(config-if)# vlan-range 20,400
! The above command specifies the range or specific vlans to which the service policy will be
attached
Switch(config-if-vlan-range)# service-policy input p1
! The above command assigns the service policy p1 to the vlans defined above.
Switch(config-if-vlan-range)# exit
Switch(config-if)# exit
Switch(config)# interface gigabitethernet 6/2
Switch(config-if)# vlan-range 20,400
Switch(config-if-vlan-range)# service-policy input p2
Switch(config-if-vlan-range)# end
```

In Example 10, VLANs 20 and 400 are voice and video VLANs for two different customers attached to two different trunk ports. Each customer receives a different set of QoS configurations using the same VLANs.

When the traffic has been classified and placed in the appropriate Tx queue, it is ready to be scheduled. The packet is scheduled out of the four Tx queues based on Tx queue characteristics such as Strict Priority, Traffic Shaping, or Bandwidth Sharing. This is discussed in the next section.

EGRESS TRAFFIC QUEUING AND SCHEDULING

The Catalyst 4500 Series supervisor engines use a shared memory architecture to provide queuing and scheduling features to the connected line cards. There are two types of gigabit linecards used: blocking and non-blocking (see Table 1). “Non-blocking” simply means none of the front end user ports are oversubscribed in the backplane connection to the Supervisor Engine switch fabric. “Blocking” means there is an oversubscription between the front end user connected ports and the backplane connection to the Supervisor Engine switch fabric. Depending on the linecard, the oversubscription ratio can be either 4: 1 or 8:1. These cards are typically used for high-density Gigabit Ethernet configurations, either directly to a user or in a server farm.

Table 1 shows the blocking and non-blocking Gigabit Ethernet line cards.

Table 1. Blocking and Non-Blocking Gigabit Ethernet Line Cards

Non-Blocking Gigabit Ethernet Line Card	Blocking Gigabit Ethernet Line Card	Oversubscription Ratio for Blocking Line Cards
Supervisor Gigabit Ethernet uplink Ports	All ports on the WS-X4424-GB-RJ45	4:1
WS-X4306-GB-all ports	All ports on the WS-X4448-GB-LX	8:1
Two 1000-BASE-X ports on the WS-X4232-GB-RJ	All ports on the WS-X4448-GB-RJ45	8:1
First two ports on WS-X4418-GB	Last 16 ports on the WS-X4418-GB	4:1
Two 1000-BASE-X ports on the WS-X4412-2GB-T (end of sale card)	1000-BASE-T ports on the WS-X4412-2GB-TX	4:1
WS-X4302-GB—all ports	WS-X4548-GB-RJ45V	8:1
WS-X4506-GB-T		

For Supervisor Engines IV, II-Plus, and II-Plus-TS, the transmit queue size for Fast Ethernet ports and blocking gigabit ports is fixed at 240 packets per queue, and 1,920 packets per queue for non-blocking gigabit ports. For the Supervisor Engine V, the transmit queue size for Fast Ethernet ports and blocking Gigabit ports has been increased to 292 packets per queue, and 2,336 packets for non-blocking Gigabit ports.

For Supervisor Engines II-Plus-10GE, and V-10GE, the transmit queue size for Fast Ethernet ports and blocking gigabit ports is fixed at 260 packets per queue, and 2,080 packets per queue for non-blocking gigabit ports and ten gigabit ports.

Note that the queue size is based on the number, not the size of packets. Input queuing on the Cisco Catalyst 4500 Series is unnecessary because the switch fabric is non-blocking (see Table 2).

Table 2. Cisco Catalyst 4500 Switch Fabric and Queue Allocations

	Supervisor Engine II-Plus	Supervisor Engine II-Plus-10GE	Supervisor Engine II-Plus-TS	Supervisor Engine IV	Supervisor Engine V	Supervisor Engine V-10GE
Switch fabric	32 Gigabit non-blocking	54 Gigabit non-blocking	32 Gigabit non-blocking	32 Gigabit non-blocking	48 Gigabit non-blocking	68 Gigabit non-blocking
Number of packets per transmit queue Fast Ethernet/gigabit blocking Ports	240	260	240	240	292	260
Number of packets per transmit queue non-blocking gigabit Ports	1,920	2,080	1,920	1,920	2,336	2,080

Note: The Cisco Catalyst 4500 Series supports IEEE 802.3z flow control to throttle back on the edge devices connected to the blocking Gigabit Ethernet ports in times of congestion when using the blocking line cards.

BANDWIDTH SHARING

All Cisco Catalyst 4500 Series supervisor engines support two ways of allocating bandwidth among the four Tx queues on a port: sharing and shaping. Bandwidth sharing provides a guaranteed **minimum** bandwidth to each of the four queues. Sharing is supported only on the non-blocking gigabit ports for Supervisor Engines IV, II-Plus, and II-Plus-TS. It is supported on all ports for the Supervisor V, II-Plus-10GE, and V-10GE. Traffic can either be below share, ie. less than its guaranteed minimum, or above share. This fact is important when discussing transmit scheduling.

Bandwidth sharing would typically be used to guarantee some minimum bandwidth to a specific application, such as VoIP which has stringent latency and loss requirements, or a high priority business application. When QoS is enabled globally, all four queues on all the non-blocking ports are assigned equal bandwidth for all supervisor engines. With Supervisor Engines V, II-Plus-10GE, and V-10GE, the default bandwidth allocated to each transmit queue on a blocking port depends on the oversubscription ratio of the line card. For a line card that has a 4:1 oversubscription ratio, the default bandwidth allocation will be 62.5 Mbps per queue. For line cards with 8:1 oversubscription, the default bandwidth allocation will be 31.25 Mbps per queue. The default settings may be changed on individual interfaces. In Example 11, the default 250 Mbps per queue setting is changed for the first three queues. The default bandwidth for blocking ports with Supervisor Engine V would be changed identically.

Example 11. Changing the Default Tx Queue Bandwidth Allocations

```
r3_4506-sup3#sh run int gig 5/6
interface GigabitEthernet5/6
  tx-queue 1
    bandwidth 300 mbps or bandwidth percent 30
  tx-queue 2
    bandwidth 400 mbps
  tx-queue 3
    bandwidth 50 mbps
```

Starting with Cisco IOS Software Release 12.1.19EW1, a bandwidth percentage, as shown in Example 11, can be used to specify the minimum guaranteed Tx queue bandwidth.

TRAFFIC SHAPING

Egress Traffic Shaping limits the transmitted traffic to no more than the configured shape rate. This can be looked at as the **maximum** bandwidth allowed on the port queue on which it is configured. It is used for smoothing traffic helping to reduce downstream device buffering. Traffic Shaping can be configured on all Gigabit Ethernet and Fast Ethernet ports. It is disabled by default. For Supervisor Engines IV, II-Plus, II-Plus-TS, II-Plus-10GE, and V-10E, traffic shaping is typically used on a blocking port's Strict Priority Queue to prevent this queue from potentially starving other queues. Queue scheduling will be explained in more detail in the following section. Traffic Shaping buffers traffic that exceeds the shape limit. In contrast traffic policing drops traffic that exceeds the rate limit. So although both methods are used to limit traffic flow, they use two different means of enforcement.

The Supervisor Engines II-Plus-10GE, V, and V-10GE have some limitations regarding traffic shaping. For Fast Ethernet port transmit queues, traffic can be accurately shaped up to 7 Mbps. For Gigabit Ethernet port transmit queues, traffic can be accurately shaped up to 50 Mbps. When shaping beyond these rates, it should be noted the shaped rate may not be achieved in certain circumstances. These circumstances would include: for example, when all the transmit queues are active; or when the packet sizes on the shaped queue include many small packets and the competing traffic on the other transmit queues has a large proportion of large packets.

STRICT PRIORITY QUEUING

For all supervisor engines, Strict Priority Queuing can be configured on **Transmit Queue # 3** of both Gigabit Ethernet and Fast Ethernet ports. For Supervisor Engines IV, II-Plus, II-Plus-TS, II-Plus-10GE, V, and V-10GE, traffic shaping is typically used together with Strict Priority Queuing on blocking ports to ensure the Strict Priority Queue does not starve the other queues. On non-blocking ports on Supervisor Engines IV, II-Plus, and II-Plus-TS, and for all ports on the Supervisor Engine V, II-Plus-10GE, and V-10GE traffic in the Strict Priority Queue will be scheduled ahead of traffic in other queues as long as the configured queue bandwidth share value is not exceeded. Example 12 shows Transmit Queue # 3 being defined as a Strict Priority Queue and the maximum bandwidth or shape rate is set at 50 Mbps.

Example 12. Traffic Shaping on the Strict Priority Queue

```
interface GigabitEthernet5/6 (truncated output)
  tx-queue 1
    bandwidth 300 mbps
  tx-queue 2
    bandwidth 300 mbps
  tx-queue 3
    bandwidth 100 mbps
    priority high
    shape 50 mbps
```

TRANSMIT QUEUE SCHEDULING

Before presenting the hierarchy of queue scheduling, it is important to provide queue definitions:

- **Strict Priority Queue:** For non-blocking ports on Supervisor Engines IV, II-Plus, and II-Plus-TS and for all ports on a Supervisor Engines V, II-Plus-10GE, and V-10GE the queue remains a Strict Priority Queue as long as it has not exceeded its configured minimum bandwidth share. Therefore, Strict Priority Queue does not mean the queue will always be serviced as long as there are packets in it. For blocking Gigabit Ethernet ports and Fast Ethernet ports on Supervisor Engines IV, II-Plus, and II-Plus-TS the Strict Priority Queue will always be serviced as long as there are packets in the queue.
- **High Priority Queue:** This is not configured as a Strict Priority Queue and has not exceeded its configured minimum bandwidth share. High priority does not mean a higher CoS/DSCP or ToS value. This definition applies to non-blocking ports on Supervisor Engines IV, II-Plus, II-Plus-TS and to all ports on Supervisor Engines V, II-Plus-10GE, and V-10GE.
- **Low Priority Queue:** This includes all ports on Supervisor Engines V, II-Plus-10GE, V-10GE and non blocking ports on Supervisor Engines IV, II-Plus, and II-Plus-TS whose queue has exceeded its guaranteed minimum bandwidth share. This applies to both a Strict Priority Queue and a High Priority Queue. On Supervisor's IV, II-Plus, and II-Plus-TS for blocking Gigabit Ethernet ports and Fast Ethernet ports, all queues except the Strict Priority Queue are Low Priority Queues.

Packets are placed into one of the four transmit queues based on the internal DSCP, as discussed previously. Transmit queues are then serviced as follows:

- Strict Priority Queue when the traffic rate is below the shape rate and the share rate. This applies to all ports on Supervisor Engines V, II-Plus-10GE, and V-10GE and the non-blocking gigabit ports on Supervisor Engines IV, II-Plus, and II-Plus-TS. For blocking gigabit ports and fast Ethernet ports on Supervisor Engines IV, II-Plus, and II-Plus-TS the Strict Priority Queue will always be serviced as long as there are packets in the queue.
- High Priority Queues with a traffic rate that is below the shape rate and the share rate. This applies to all ports on Supervisor Engines V, II-Plus-10GE, and V-10GE and to the non-blocking gigabit ports on Supervisor Engines IV, II-Plus, and II-Plus-TS.
- If there is more than one High Priority Queue, round robin between them.
- If there is more than one Low Priority Queue, round robin between them.
- When any queue's traffic rate exceeds the shape rate (if configured), the packet is buffered and not transmitted until the traffic rate is again below the shape rate.

The following command shown in Example 13 is useful in determining the QoS settings on a particular port on the Cisco Catalyst 4500 Series Switch with a Supervisor Engine V. This will show trust settings and any shaping and sharing configured on the port as well as the transmit queue size. Please note that for Supervisor Engines IV, II-Plus or II-Plus-TS the queue size would be 1,920 packets.

Example 13. Interface QoS Information

```
r3_4510R_S4#sh qos int gig 5/6
```

```
QoS is enabled globally
```

```
Port QoS is enabled
```

```
Port Trust State: 'untrusted'
```

```
Trust device: none
```

```
Default DSCP: 0 Default CoS: 0
```

```
Appliance trust: none
```

```
! Output continue below.
```

Tx-Queue	Bandwidth (bps)	ShapeRate (bps)	Priority	QueueSize (packets)
1	300000000	disabled	N/A	2336
2	300000000	disabled	N/A	2336
3	100000000	50000000	high	2336
4	300000000	disabled	N/A	2336

QOS POLICING AND MARKING

A policer is typically used to rate limit traffic. Traffic that exceeds the specified rate is considered out of profile and can either be dropped or have its DSCP marked down. The Cisco Catalyst 4500 Series supervisor engines support ingress, egress, and microflow policing (Supervisor Engine V-10E only) in hardware using a single rate policer. The policer also includes a burst parameter to allow for bursting of traffic above the specified rate. The policer supports rates from 32 kbps to 32 Gbps. The burst size is measured in bytes and can range from 1KB to 512MB. The burst size should take into account the type of traffic being policed. If the burst size is set too small, there may be some over policing. If the burst size is set too large, the entire transfer may fit into the burst and hence there may be over policing, particularly for smaller packet sizes. In general smaller burst sizes should be used for voice and video traffic and larger burst sizes should be used for TCP traffic. A good starting point for TCP traffic is:

- $Burst = 2 * (Packet\ Round\ Trip\ Time) * (Data\ rate)$

For example, assume a policing rate of 960 kbps and a TCP RTT of 0.05 seconds, the burst size would be:

- $Burst\ size = 2 * (.05) * (960\ kbps/8\ or\ 120\ kbps) = 12,000\ bytes$

The exact burst size to be used in practice would depend on the type of traffic being policed.

Policing is configured on the Cisco Catalyst 4500 Series supervisor engines using the previously discussed, Modular QoS CLI. Traffic to be policed is identified using ACLs or match statements in a class map. The supervisor engines support the following match statements in a class map:

- Match access-group (can be Layer 2/3/4)
- Match any
- Match IP DSCP
- Match IP Precedence
- Match flow ip (applicable only when configuring microflow policing to deliver user based rate limiting)

The class maps are then used to identify the type of traffic to police. These policers would then be applied via a service policy to either an individual port(s) or VLAN(s). A single policy map can support up to the Cisco IOS Software system wide limit of 255 class maps. Remember that class maps can be shared among policy maps. For example, in a Metro Ethernet environment, where a service provider may offer four different types of services to 240 different customers, only four class maps are required to accommodate the service policies for these customers.

There are three types of policers supported on the Cisco Catalyst 4500 IOS Supervisors: individual (per interface), aggregate (multiple interfaces), and micro flow policers.

- **Individual policer**—QoS applies the bandwidth limits specified in the policer separately to each matched traffic class for each port/VLAN to which the policy map is attached. In order to configure an individual policer within a policy map the command **police** under policy-map class configuration mode needs to be specified. Example 14 shows an individual policer configuration

Example 14. Individual Policer Configuration Applied to VLAN 2

```
class-map match-all limit_host
  match access-group 101
! The above class map, limit_host, is looking for traffic that matches ACL 101 below:
Extended IP access list 101
  permit ip any host 10.5.1.2
! The policy map limit_10, below, will be applied to rate limit this type of traffic
policy-map limit_10
  class limit_host
    police 10m 8k conform-action transmit exceed-action drop
! The policy map is then applied to Vlan 2 as an input policy to rate limit the specified traffic
inbound traffic on Vlan 2 to 10Mbps with a burst size of 8k:
interface Vlan2
  ip address 10.20.1.1 255.255.255.0
  service-policy input limit_10
```

- **Aggregate policer**—With aggregate policers, QoS applies the bandwidth limits specified in an aggregate policer cumulatively to all matched traffic flows. Before any policy map is configured on an interface, the bandwidth limits of the policer must be defined using the **qos aggregate-policer** global configuration command. In order to configure an aggregate policer, within a policy map, the command **police aggregate** under policy-map class configuration mode needs to be specified. Example 15 shows an aggregate policer configuration.

Example 15. Aggregate Policer Configuration Applied to VLAN 2

```
qos aggregate-policer pol_10meg 10 mbps 1 kbyte conform-action transmit exceed-action drop
! The above defines the aggregate policer which rate limits traffic to 10 Mbps
class-map match-all limit_host
  match access-group 101
! the above class map, limit_host, is looking for traffic that matches ACL 101 below:
Extended IP access list 101
  permit ip any host 10.5.1.2
! The policy map limit_10, below, will be applied to rate limit this type of traffic
policy-map limit_10
  class limit_host
    police aggregate pol_10meg
! The policy map is then applied to Vlan 2 as an input policy to rate limit the specified traffic
inbound traffic on Vlan 2:
interface Vlan2
  ip address 10.20.1.1 255.255.255.0
  service-policy input limit_10
```

- **Flow or microflow policer**—With flow-based policing, all the identified flows are policed to the specified rate individually. Because the flows are dynamic, distinguishing fields must be configured in class maps. Through the use of microflow policers, two important features can be configured: User Based Rate Limiting (UBRL), and Hierarchical Policing.
 - *User Based Rate Limiting*—This feature adopts the microflow policing capability to dynamically learn traffic flows and rate limit unique flow to an individual rate. UBRL is supported only on the Supervisor Engine V-10GE with the built-in NetFlow support. UBRL is typically used in environments where a per-user granular rate-limiting mechanism is required, such as, per-user outbound traffic rate versus a per-user inbound traffic rate. UBRL can police IP and non-IP traffic. Example 16 shows a typical user based rate limiting configuration.
 - *Hierarchical Policing*—This defines the ability to configure a set of policers (parent and child) to further define the rate limits of a particular flow. A parent policer can be considered as an aggregate policer, while a child policers can be defined as an individual policer within the aggregate policer. A flow-based policy map is defined as a child policy map by default. A parent policy map cannot be a flow-based policy map. Furthermore, both the child policy map and parent policy map must have **match-all** in their class-map configuration. Hierarchical Policing is supported only with the Supervisor Engine V-10GE in conjunction with UBRL. Example 17 shows a typical hierarchical policer configuration.

Example 16. UBRL Configuration Applied to FastEthernet 6/1 Matching the IP Source Address

```
Switch(config)# class-map c1
! Defining the class map c1
Switch(config-cmap)# match flow ip source-address
! The above command matches against any flow in FastEthernet 6/1.
Switch(config-cmap)# exit
!
Switch(config)# policy-map p1
! The above command configures the policy map p1
Switch(config-pmap)# class c1
```

```

! The above command assigns the class map c1 to the policy map p1
Switch(config-pmap-c)# police 1000000 9000
! The above command specifies the traffic rate to which the flows are being policed
Switch(config-pmap-c)# exit
!
Switch(config-pmap)# exit
Switch(config)# interface fa6/1
Switch(config-if)# service-policy input p1
! The above command assigns the policy p1 to FastEthernet 6/1 in the ingress direction

```

Example 17. Hierarchical Policer Configuration Applied to FastEthernet 6/1

```

Switch(config)# ip access-list 10 permit 101.237.0.0 0.0.255.255
Switch(config)# ip access-list 20 permit 0.0.10.0 255.255.0.255
Switch(config)# class-map match-all flow-class
! The above command defines the class map to be used with child policy
Switch(config-cmap)# match flow ip source-address
! The above command matches all flows to a source address
Switch(config-cmap)# match access-group 20
! The above command matches the class map used for the child policy to access list 20
Switch(config-cmap)# class-map match-all aggregate-class
! The above command defines the class map to be used with the parent policy
Switch(config-cmap)# match-access group 10
! The above command matches class map aggregate-class to access list 10
Switch(config-cmap)# exit
Switch(config)# policy-map flow-policy
! The above command defines the child policy flow-policy
Switch(config-pmap)# class flow-class
! The above command assigns the class map flow-class to the child policy flow-policy
Switch(config-pmap-c)# police 2m 1k conform-action transmit exceed-action drop
! The above command defines the rate limits of the flows identified by the class map flow-class
Switch(config-pmap-c)# end
Switch# config t
Switch(config)# policy-map aggregate-policy
! The above command defines the parent policy aggregate-policy
Switch(config-pmap)# class aggregate-class
! The above command assigns the class map aggregate-class to the parent policy aggregate-policy
Switch(config-pmap-c)# police 50m 4k conform-action transmit exceed-action drop
! The above command defines the rate limits of the aggregate flows defined via access list 10.
Switch(config-pmap-c)#end
Switch# config t
Switch(config)# policy-map aggregate-policy

```

```
Switch(config-pmap)# class aggregate-class
! The above command assigns the class map aggagate-class to the policy map
Switch(config-pmap-c)# service-policy flow-policy
! The above command assigns the child policy to the parent policy providing hierarchical policing
```

Note: The Cisco Catalyst 4500 Series supervisor engines also support egress policing. This can be implemented by specifying service-policy output to the interface VLAN 2 configuration shown in Example 8.

The following command, shown in Example 18 would be used to monitor policing operations.

Example 18. Monitoring Policing Operations

```
r3_4507R_S4#sh policy-map int vlan 2
Vlan2
  service-policy input: limit_10
    class-map: limit_host (match-all)
      1503268 packets
    match: access-group 101
    police: Per-interface ! or police for an aggregate policer
      Conform: 8758484 bytes Exceed: 76688002 bytes
```

When monitoring policers, there are several considerations to know:

- The per class packet counter, counts all packets matching the class among all interfaces where this class is applied in a service policy.
- Policers maintain byte counters not packet counters.
- There is no command to verify the offered or outgoing traffic rate per policer. A rough approximation can be obtained by dividing the conform bytes by the total bytes (conform + exceed). Using the policer results shown in Example 18, this yields a conform ratio of 8758484/85446486 or ~10.2 percent. The policer's rate limit was set at 10 Mbps on a 100 Mbps link, so this provides a good approximation.
- The counters are updated approximately every 2 seconds, keep this in mind when executing the command shown in Example 18, in rapid succession.
- Starting with Cisco IOS Software Release 12.1.19EW1 and later, the policer calculations can include the 14-byte Ethernet header field and 4-byte FCS field when policing packets. This is enabled by the global command: **QoS account layer2 encapsulation length 18**. Releases prior to this do not include these fields. When determining the policing rate and burst parameters, the Layer 2 encapsulation length needs to be deducted, otherwise "under-policing" will result, particularly for smaller packet sizes in the 64 byte to 256 byte range.

An aggregate policer takes a single entry in the hardware table. An individual policer takes an entry per policer. The Supervisor Engines IV and V support 1,024 (1,020 usable) ingress and 1,024 (1,020 usable) egress policers. The Supervisor Engines II-Plus, II-Plus-TS, and II-Plus-10GE support 512 (508 usable) ingress and 512 (508 usable) egress policers, and the Supervisor Engine V-10GE supports 8,192 (8,188 usable) ingress and 8,192 (8,188 usable) egress policers, and 511 flow policers. There are 4 input and 4 output policers being reserved by the system for internal use.

Example 19 shows a useful command to determine how many policer entries are being consumed in the hardware.

Example 19. Cisco Catalyst 4500 Supervisor Engine IV Policer Hardware Utilization Output

```
r3_4507R_S4#sh platform hardware qos policers utilization
  Software Usage Statistics

```

	Used (%)	Free (%)	Total
Input Policers	5 (0.4)	1019 (99.5)	1024
Output Policers	4 (0.3)	1020 (99.6)	1024

The actual number of policers that can be supported involves an interaction between the resources used to define a policer and their individual limits:

- Total number of maximum class-maps (current Cisco IOS Software system limit is 1,024, with maximum 255 class maps per policy map)
- Total number of maximum policy-maps (current Cisco IOS Software system limit is 8,192)
- Number of Layer 4 operators per ACL in the class-map match statements (limited by hardware resources to no more than six)
- Number of hardware ternary content addressable memory (TCAM) entries that could be consumed by the ACLs used in class-map match statements
 - 32,000 access control entries (ACEs) (16,000 input and 16,000 output) in for the Supervisor Engines IV, V, V-10GE
 - 16,000 ACEs (8,000 input and 8,000 output) for the Supervisor Engines II Plus, II-Plus-TS, II-Plus-10GE
- Total number of hardware policers:
 - 1,020 ingress and 1,020 egress for the Supervisor Engines IV, and V
 - 8,188 ingress and 8,188 egress for the Supervisor Engine V-10GE
 - 508 ingress and 508 egress for the Supervisor II Plus, II-Plus-TS, II-Plus-10GE

Another exceed action to take when the traffic is out of profile is to mark down the DSCP value in the packet. Example 20 defines a service policy that identifies ingress traffic marked as DSCP 24 and when the profile is exceeded, the traffic is marked down to DSCP 16.

Example 20. Marking Down an Out of Profile Packet

```
interface FastEthernet3/1
  switchport access vlan 2
  qos vlan-based
  qos trust dscp
! The command above will trust the inbound DSCP marking and the QoS policy will be VLAN based
qos map dscp policed 24 to dscp 16
! The command above will adjust the DSCP to policed DSCP table
class-map match-all class_setprec
match ip dscp 24
! The command above will match any incoming traffic with DSCP of 24
qos aggregate-policer po2_20meg 20 mbps 1 kbyte conform-action transmit exceed-action policed-
dscp-transmit
! The command above defines the aggregate policer po2_20meg to mark down the non-conforming
traffic
Policy Map pol_setdscp
```

```

class match_dscp
  police aggregate po2_20meg
! The command above defines the service policy po2_20meg
interface Vlan2
  ip address 10.20.1.1 255.255.255.0
  service-policy input pol_setdscp

! The command above applies the service policy to the interface, in this case vlan 2.

```

If ordering of packets is important, a best practice is to mark down out of profile packets to a DSCP value that would place it in the same output queue as an in-profile packet.

Packets can also be marked based on the criteria specified in a class map. Refer to the earlier Example 6 for this scenario.

In the Cisco Catalyst 4500 Series supervisor engines, ingress policies are applied before egress policies. For example, if both ingress and egress policing are applied to a packet, the most severe decision will be taken. As an example, if the ingress policer action is to drop the packet and the egress policer action is to mark the packet down, the packet will be dropped. If an ingress service policy re-marks a packet, the egress policy will see the packet as if no change had been made, because the actual marking of the packet occurs after the egress policy. Therefore, egress marking overrides ingress marking and an egress policy cannot match on the new QoS levels that were changed by the ingress marking.

Table 3 summarizes the QoS actions taken on a packet when both ingress and egress policies are applied.

Table 3. QoS Actions

Egress Policy	Ingress Policy			
	Transmit	Drop	Markdown	Mark
Transmit	Transmit	Drop	Markdown	Mark
Drop	Drop	Drop	Drop	Drop
Markdown	Markdown	Drop	Markdown	Markdown
Mark	Mark	Drop	Mark	Mark

DYNAMIC BUFFER LIMITING

Dynamic Buffer Limiting (DBL) is a congestion avoidance technique, used to drop packets before the congestion occurs. Weighted Random Early Discard (WRED) randomly discards packets at specified queue thresholds. Lower priority ToS/DSCP packets are dropped before higher priority packets, hence the weighting in WRED. This action reduces the average queue size and thus allows the switch or router to detect congestion before the queue overflows. WRED is packet based not flow based like DBL. Because DBL is flow based and not random, it does not impact “well behaved” flows that are not causing the queues to congest. DBL uses logical flow tables per port/per queue on each interface. A DBL flow is comprised of a source/destination IP address, Layer 4 TCP/UDP ports, and a VLAN. During congestion, the logic in the hardware for dropping a packet is based on the flow. DBL is particularly effective with non-adaptive flows (NAFs). A NAF is any flow that does not reduce the traffic rate in response to packet drops. Usually NAFs use UDP connectionless protocol. Some examples include UDP music or video flows, Internet streaming multimedia, and multicast traffic.

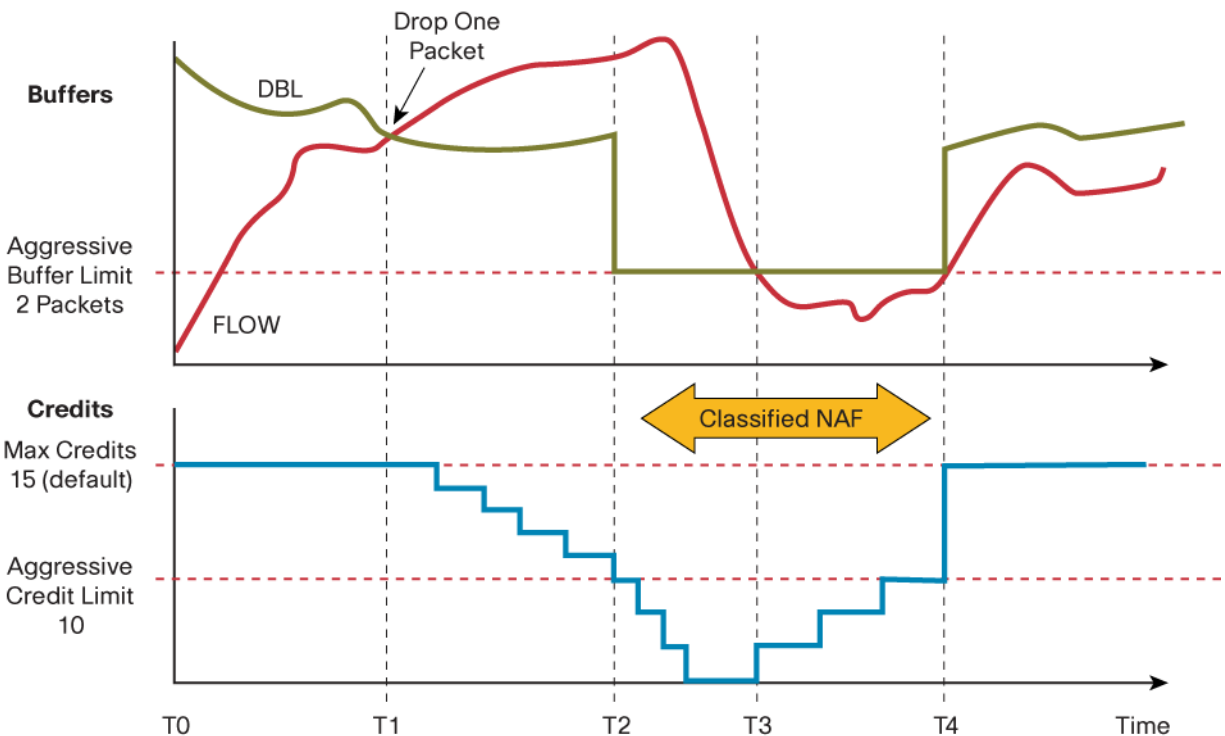
DBL is similar to the Flow Based WRED (FRED) that is used on Cisco IOS routers, except it is implemented in hardware at full line rate on all supervisors engines. The hardware implementation is important when deployed in Gigabit Ethernet networks versus typical WANs where maximum bandwidth is usually T-1 speed (1.5 Mbps), DS-3 (45 Mbps) or perhaps OC-4 (155 Mbps).

DBL is supported on all the ports of a Cisco Catalyst 4500 Series Switch. The DBL transmit queue logic is on all 4 transmit queues on every Fast Ethernet, Gigabit Ethernet, or 10 Gigabit Ethernet port. In addition, the protocols on the port are transparent to DBL, which means that you can have routed, switched, access, or trunk ports with Cisco EtherChannel technology or any other protocol configured on them. A Cisco Catalyst 4500 Series Switch with DBL can be used with a Cisco Catalyst 6500 Series Switch that supports WRED. The DiffServ Internet architecture is structured around Per Hop Behavior (PHB). For example, with a Cisco Catalyst 4500 Series Switch in the wiring closet with uplinks to a Cisco Catalyst 6500 Series Switch in the distribution/core network, the non-adaptive flows (NAF's) have already been controlled using DBL prior to reaching the Cisco Catalyst 6500 Switch. The Cisco Catalyst 6500 Series could then use WRED on those flows that will respond to packet drops.

Referring to the QoS flow shown at the beginning of this paper in Figure 1, DBL acts on a packet flow before the packet is enqueued, avoiding tail drops. The DBL function also occurs after the policing function. DBL and policing are not alternatives to each other. Policing is used to control selected traffic flows by rate limiting them. It is still possible to have Tx queue congestion particularly when bursting occurs. DBL is a QoS congestion avoidance technique specifically designed to prevent Tx queue congestion.

Figure 2 shows how DBL works.

Figure 2. DBL Operation



For every active flow, the switch maintains two parameters—the number of credits and the number of buffers used. When a flow takes more buffers than the dynamically computed limit, DBL drops one packet (starting at T1 in the plot). A single packet is dropped to signal congestion so that the flow has a chance to retreat without losing more packets. If, however, the flow does not respond to the single packet drop and continues to send packets at the same high rate, it loses its credits one by one. The flow is considered an aggressive flow (NAF) when its credits are below the aggressive credit limit. At this point the buffer usage is also reduced to the aggressive buffer limit. On the other hand, at time T3 when the flow uses fewer buffers than the dynamically computed limit, the number of credits is increased one at a time for aggressive flows. Well-behaved flows go back to full credits immediately. Note that when DBL is used with TCP flows, the TCP phenomenon known as global synchronization is avoided.

because there is a 15 percent drop probability. That is, if many TCP flows arrive at the same time, a packet will be dropped from approximately 1 of every 6 flows, thereby avoiding a congestion collapse and a ramp up of all of the flows at the same time.

Example 21 shows DBL configuration parameters. These parameters were carefully selected and altering them is not recommended.

Example 21. Dynamic Buffer Limiting Configuration Parameters

```
r3_4507R_S4#sh qos dbl
QoS is enabled globally
DBL is enabled globally
DBL flow includes vlan
DBL flow includes layer4-ports
DBL uses ecn to indicate congestion
DBL exceed-action probability: 15%
DBL max credits: 15
DBL aggressive credit limit: 10
DBL aggressive buffer limit: 2 packets
```

DBL is configured globally as shown in Example 22.

Example 22. DBL Configuration Using Explicit Congestion Notification (ECN) Option

```
r3_4507R_S4(config)#qos dbl exceed-action ecn
```

In Example 22, DBL was enabled with the Explicit Congestion Notification (ECN) option. This option is recommended for end devices supporting ECN, such as Linux or Solaris hosts.

DBL is then configured via the Modular QoS CLI used to configure service policies, and then applied to the interface where congestion is expected. Example 23 is a sample configuration.

Example 23. DBL Service Policy

```
r3_4507R_S4# (config)# policy-map LAB-POLICY
r3_4507R_S4# (config-pmap)# class UDP-LargePkts
r3_4507R_S4# (config-pmap-c)# db1
! DBL will be applied to the flows identified in the class map
r3_4507R_S4# (config-pmap)# class FTP
r3_4507R_S4# (config-pmap-c)# db1
r3_4507R_S4(config-if)#service-policy output LAB-POLICY
! Apply the DBL service policy to the interface you are expecting the congestion on
```

In this example, the classes UDP-LargePkts and FTP both are using a DSCP value of 0, so they will be placed in Tx queue 1.

The following command shown in Example 24 is useful to monitor DBL. The output shown is before DBL is applied.

Example 24. Interface Output Without DBL

```
r3_4507R_S4# show interface gi1/1 counters detail (truncated output)
Port      Tx-Bytes-Queue-1  Tx-Bytes-Queue-2  Tx-Bytes-Queue-3  Tx-Bytes-Queue-4
Gi1/1     315142608        28919476         0                 430984
Port      Tx-Drops-Queue-1  Tx-Drops-Queue-2  Tx-Drops-Queue-3  Tx-Drops-Queue-4
Gi1/1     14489            0                 0                 0
Port      Dbl-Drops-Queue-1 Dbl-Drops-Queue-2 Dbl-Drops-Queue-3 Dbl-Drops-Queue-4
Gi1/1     0                0                 0                 0
```

Notice the taildrops in Tx queue 1. There are also no DBL drops indicated. Example 25 shows the output after the service policy in Example 23 has been implemented.

Example 25. Interface Output with DBL Implemented

```
r3_4507R_S4# show interface gi1/1 counters detail (truncated output)
Port      Tx-Bytes-Queue-1  Tx-Bytes-Queue-2  Tx-Bytes-Queue-3  Tx-Bytes-Queue-4
Gi1/1     10250756         1656656          0                 24204
Port      Tx-Drops-Queue-1  Tx-Drops-Queue-2  Tx-Drops-Queue-3  Tx-Drops-Queue-4
Gi1/1     0                0                 0                 0
Port      Dbl-Drops-Queue-1 Dbl-Drops-Queue-2 Dbl-Drops-Queue-3 Dbl-Drops-Queue-4
Gi1/1     8201             0                 0                 0
```

AUTO QoS

The Cisco Catalyst 4500 Series Auto QoS feature is used to simplify QoS for VoIP deployments. It is available in Cisco IOS Software Release 12.1.19 EW and later on the Cisco Catalyst 4500 Series supervisors engines. With Auto QoS enabled, ingress traffic is automatically classified and then placed into the appropriate egress queue. Auto QoS should be enabled on ports directly connected to Cisco IP phones as well as uplink ports that will be transmitting VoIP traffic. After being configured, Auto QoS performs the following functions:

- Detects the presence or absence of a Cisco IP Phone
- Configures ingress traffic classification
- Configures egress queues
- Enables a service policy that matches all traffic and enables DBL on the interface for congestion avoidance
- Automatically shapes the VoIP traffic

Table 4 summarizes how Auto QoS classifies ingress traffic types and assigns the appropriate egress queues.

Table 4. Auto QoS Classification

	VoIP Traffic from Cisco IP Phones	VoIP Control Traffic from Cisco IP Phones	Routing Protocol Traffic	Spanning Tree Protocol BPDUs Traffic	All Other Traffic
Ingress DSCP	46	26	—	—	—
Ingress CoS	5	3	6	7	—
DiffServ	EF	AF31	—	—	—

	VoIP Traffic from Cisco IP Phones	VoIP Control Traffic from Cisco IP Phones	Routing Protocol Traffic	Spanning Tree Protocol BPDU Traffic	All Other Traffic
Assigned DSCP	46	26	48	56	0
Assigned CoS	5	3	6	7	0
Tx Queue	3	4	4	4	1

Example 26 shows Auto QoS configured on a port connected to a Cisco IP Phone, with the Auto QoS keyword, “cisco-phone” used. The minimum bandwidth percent shown for tx-queue 3 is supported on Supervisor Engines V, V-10GE, and II-Plus-10GE on any 10/100/1000 Ethernet port and non-blocking Gigabit Ethernet ports. With Supervisor Engines IV, II-Plus, and II-Plus-TS, the minimum bandwidth percent command is supported on non-blocking Gigabit Ethernet ports only.

Example 26. Auto QoS Configuration for a Port Connected to a Cisco IP Phone

```
interface FastEthernet3/4
  switchport access vlan 2
  switchport voice vlan 20
  qos trust device cisco-phone
  qos trust CoS
  service-policy output autoqos-voip-policy
  auto qos voip cisco-phone
  tx-queue 3
    bandwidth percent 33 (Supervisor V, V-10GE, and II-Plus-10GE only)
    priority high
    shape percent 33
```

A similar configuration is shown for an uplink in Example 27. Here the Auto QoS keyword trust was used.

Example 27. Auto QoS Configuration for an Uplink Port

```
interface GigabitEthernet5/1
  switchport trunk encapsulation dot1q
  qos trust CoS
  service-policy output autoqos-voip-policy
  auto qos voip trust
  tx-queue 3
  bandwidth percent 33 (Supported on all Supervisor Engines on non-blocking gigabit ports)
    priority high
    shape percent 33
```

Auto QoS can be enabled on static, dynamic-access, voice VLAN access, and trunk ports. It is not supported on Cisco EtherChannel ports.

CPU QoS

The Cisco Catalyst 4500 Series supervisor engines have an extra QoS feature for traffic that is sent to the CPU. This would typically be control plane traffic such as bridge protocol data units (BPDUs) or routing protocol updates. It can also be used for packets that are not switched in hardware such as Address Resolution Protocol (ARP) packets. In times of high CPU usage, priority is given to the control packets such as routing updates and BPDUs to provide stability and to prevent routing flaps or Spanning Tree events. This capability is not user configurable.

APPENDIX A

Table 5 correlates DSCP PHB labels to their binary and decimal representations and their IP precedences.

Table 5. DSCP and IP Precedence Decimal and Binary Equivalents

DSCP PHB	DSCP Binary	DSCP Decimal	IP Precedence Decimal	IP Precedence Binary
–	000010	2	0	000
–	000100	4	0	000
–	000110	6	0	000
BE	000000	0	0	000
AF11	001010	10	1	001
AF12	001100	12	1	001
AF13	001110	14	1	001
AF21	010010	18	2	010
AF22	010100	20	2	010
AF23	010110	22	2	010
AF31	011010	26	3	011
AF32	011100	28	3	011
AF33	011110	30	3	011
AF41	100010	34	4	100
AF42	100100	36	4	100
AF43	100110	38	4	100
EF	101110	46	5	101

FOR MORE INFORMATION

For more information about Cisco IOS Software Release 12.2(25)SG on the Cisco Catalyst 4500 Series Switch, please visit:
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/12_2_25s/index.htm

**Corporate Headquarters**

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel
Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan
Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2006 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StrataView Plus, TeleRouter, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0502R) 205377.BE_ETMG_CC_1.06

