

Classification Overview

Classification entails using a traffic descriptor to categorize a packet within a specific group to define that packet and make it accessible for QoS handling on the network. Using packet classification, you can partition network traffic into multiple priority levels or classes of service. When traffic descriptors are used to classify traffic, the source agrees to adhere to the contracted terms and the network promises a quality of service. Traffic policers, such as committed access rate's (CAR's) rate-limiting feature, and traffic shapers, such as Frame Relay Traffic Shaping (FRTS) and Generic Traffic Shaping (GTS), use a packet's traffic descriptor—that is, its classification—to ensure adherence to the contract.

Packet classification is pivotal to policy techniques that select packets traversing a network element or a particular interface for different types of QoS service. For example, you can use classification to mark certain packets for IP Precedence and you can identify others as belonging to a Resource Reservation Protocol (RSVP) flow.

Methods of classification were once limited to use of the contents of the packet header. Today's methods of marking a packet with its classification allow you to set information in the Layer 2, 3, or 4 headers, or even by setting information within the packet's payload. Criteria for classification of a group might be as broad as "traffic destined for subnetwork X" or as narrow as a single flow.


This chapter explains IP Precedence, then it gives a brief description of the kinds of traffic classification provided by the Cisco IOS QoS features. It discusses the following features:

- Policy-Based Routing
- QoS Policy Propagation via Border Gateway Protocol
- Committed Access Rate

About IP Precedence

Use of IP Precedence allows you to specify the class of service (CoS) for a packet. You use the three precedence bits in the IPv4 header's type of service (ToS) field for this purpose. Figure 2 shows the ToS field.

Figure 2 IPv4 Packet Type of Service Field



Using the ToS bits, you can define up to six classes of service. Other features configured throughout the network can then use these bits to determine how to treat the packet in regard to the type of service to grant it. These other QoS features can assign appropriate traffic-handling policies including congestion management strategy and bandwidth allocation. For example, although IP Precedence is not a queueing method, queueing methods such as weighted fair queueing (WFQ) and Weighted Random Early Detection (WRED) can use the IP Precedence setting of the packet to prioritize traffic.

By setting precedence levels on incoming traffic and using them in combination with the Cisco IOS QoS queueing features, you can create differentiated service. You can use features such as policy-based routing (PBR) and CAR to set precedence based on extended access list classification. These features afford considerable flexibility for precedence assignment. For example, you can assign precedence based on application or user, or by destination and source subnetwork.

So that each subsequent network element can provide service based on the determined policy, IP Precedence is usually deployed as close to the edge of the network or the administrative domain as possible. You can think of IP Precedence as an edge function that allows core, or backbone, QoS features, such as WRED, to forward traffic based on CoS. IP Precedence can also be set in the host or network client, but this setting can be overridden by policy within the network.

The following QoS features can use the IP Precedence field to determine how traffic is treated:

- Distributed Weighted Random Early Detection (Distributed-WRED)
- Weighted Fair Queueing (WFQ)
- Committed Access Rate (CAR)

How the IP Precedence Bits Are Used to Classify Packets

You use the three IP Precedence bits in the ToS field of the IP header to specify CoS assignment for each packet. You can partition traffic into up to six classes—the remaining two are reserved for internal network use—and then use policy maps and extended ACLs to define network policies in terms of congestion handling and bandwidth allocation for each class.

For historical reasons, each precedence corresponds to a name. These names, which continue to evolve, are defined in the RFC 791 document. Table 3 lists the numbers and their corresponding names, from least to most important.

Table 3 IP Precedence Values

Number	Name
0	routine
1	priority
2	immediate
3	flash
4	flash-override
5	critical
6	internet
7	network

However, the IP Precedence feature allows you considerable flexibility for precedence assignment. That is, you can define your own classification mechanism. For example, you might want to assign precedence based on application or access router.

Note IP Precedence bit settings 6 and 7 are reserved for network control information, such as routing updates.

Setting or Changing the IP Precedence Value

By default, the Cisco IOS software leaves the IP Precedence value untouched, preserving the precedence value set in the header, allowing all internal network devices to provide service based on the IP Precedence setting. This policy follows the standard approach stipulating that network traffic should be sorted into various types of service at the basic perimeter of the network and that those types of service should be implemented in the core of the network. Routers in the core of the network can then use the precedence bits, for example, to determine the order of transmission, the likelihood of packet drop, and so on.

However, because traffic coming into your network can have precedence set by outside devices, Cisco recommends you reset the precedence for all traffic entering your network. By controlling IP Precedence settings, you prohibit users who have already set the IP Precedence from acquiring better service for their traffic simply by setting a high precedence for all of their packets.

You can use any of the following features to set the IP precedence in packets:

- Policy-Based Routing
- QoS Policy Propagation via Border Gateway Protocol (PB-BGP)
- Committed Access Rate (CAR)

As mentioned previously, after a packet has been classified, you can use other QoS features such as CAR and WRED to specify and enforce business policies to fit your business model.

Policy-Based Routing

PBR gives you a flexible means of routing packets by allowing you to configure a defined policy for traffic flows, lessening reliance on routes derived from routing protocols. To this end, PBR gives you more control over routing by extending and complementing the existing mechanisms provided by routing protocols. PBR allows you to set the IP Precedence. It also allows you to specify a path for certain traffic, such as priority traffic over a high-cost link.

You can set up PBR as a way to route packets based on configured policies. For example, you can implement routing policies to allow or deny paths based on the identity of a particular end system, an application protocol, or the size of packets.

PBR allows you to do the following:

- Classify traffic based on extended ACL criteria. ACLs, then, establish the match criteria.
- Set IP Precedence bits, giving the network the ability to enable differentiated classes of service.
- Route packets to specific traffic-engineered paths; you might need to route them to allow a specific QoS service through the network.

Policies can be based on IP address, port numbers, protocols, or size of packets. For a simple policy, you can use any one of these descriptors; for a complicated policy, you can use all of them.

For example, classification of traffic through PBR allows you to identify traffic for different classes of service at the edge of the network and then implement QoS defined for each CoS in the core of the network using priority, custom, or weighted fair queueing techniques. This process obviates the need to classify traffic explicitly at each WAN interface in the core-backbone network.

How It Works

All packets received on an interface with PBR enabled are passed through enhanced packet filters known as route maps. The route maps used by PBR dictate the policy, determining to where the packets are forwarded.

Route maps are composed of statements. The route map statements can be marked as permit or deny, and they are interpreted in the following way:

- If the packets do not match any route map statements, then all the set clauses are applied.
- If a statement is marked as deny, the packets meeting the match criteria are sent back through the normal forwarding channels and destination-based routing is performed.
- If the statement is marked as permit and the packets do not match any route map statements, the packets are sent back through the normal forwarding channels and destination-based routing is performed.

You specify PBR on the interface that receives the packet, not on the interface from which the packet is sent.

When Should You Use Policy-Based Routing?

You might enable PBR if you want certain packets to be routed some way other than the obvious shortest path. Some possible applications for PBR are to provide equal access, protocol-sensitive routing, source-sensitive routing, routing based on interactive versus batch traffic, or routing based on dedicated links.

Some applications or traffic can benefit from QoS-specific routing; for example, you could transfer stock records to a corporate office on a higher-bandwidth, higher-cost link for a short time while transmitting routine application data such as e-mail over a lower-bandwidth, lower-cost link.

QoS Policy Propagation via Border Gateway Protocol

BGP is an interdomain routing protocol that exchanges reachability information with other BGP systems. It is defined by RFC 1163.

Policy Propagation via BGP allows you to classify packets based on the following:

- Access lists
- BGP community lists. A community is a group of destinations that share some common attribute. You use community lists to create groups of communities to use in a match clause of a route map. As with access lists, a series of community lists can be created.
- BGP autonomous system paths. An autonomous system path is a collection of networks under a common administration sharing a common routing strategy. BGP carries the autonomous system path in its routing updates. You can filter routing updates by specifying an access list on both incoming and outbound updates based on the BGP autonomous system path.
- IP Precedence. See the section “About IP Precedence” earlier in this chapter.
- Source and destination address lookup. You can specify whether the IP Precedence level is obtained from the source (input) address or destination (output) address entry in the route table.

After a packet has been classified using BGP, you can use other QoS features such as CAR and WRED to specify and enforce business policies to fit your business model.

BGP Policy Propagation leverages BGP to distribute QoS policy to remote routers in your network. It allows ingress routers to prioritize incoming traffic.

Policy Propagation via BGP is supported on these platforms:

- Cisco 7000 series routers with the RSP7000 and RSP7000CI
- Cisco 7200 series
- Cisco 7500 series

For the Policy Propagation via BGP feature to work, you must enable BGP and Cisco Express Forwarding (CEF)/Distributed CEF (DCEF) on the router.

Subinterfaces on an ATM interface that has the **bgp-policy** command enabled must use CEF mode because Distributed CEF is not supported. (Note that DCEF uses the VIP rather than the RSP to perform forwarding functions.)

Committed Access Rate

CAR is a multifaceted feature that implements both classification services and policing through rate limiting. This section describes its classification capability. For information on its rate limiting features, see the chapter “Policing and Shaping Overview.”

You can use CAR’s classification services to set the IP Precedence for packets entering the network. This capability of CAR allows you to partition your network into multiple priority levels or classes of service. Networking devices within your network can then use the adjusted IP Precedence to determine how to treat the traffic. For example, VIP-Distributed WRED uses the IP Precedence to determine the probability of whether a packet will be dropped.

As discussed in the section “About IP Precedence,” you can use the three precedence bits in the ToS field of the IP header to define up to six classes of service.

You can classify packets using policies based on physical port, source or destination IP or MAC address, application port, IP protocol type, or other criteria specifiable by access lists or extended access lists. You can even classify packets by categories external to the network, for example, by customer. After a packet has been classified, a network can either accept or override and reclassify the packet according to a specified policy. CAR includes commands you can use to classify and reclassify packets.

CAR is supported on these routers:

- Cisco 2600 series
- Cisco 3600 series
- Cisco 4500 series
- Cisco 4700 series
- Cisco 7200 series

VIP-Distributed CAR is a version of CAR that runs on the Versatile Interface Processor (VIP). It is supported on the following routers with a VIP2-40 or greater interface processor:

- Cisco 7000 series with RSP7000
- Cisco 7500 series