

Quality of Service Policy Propagation via Border Gateway Protocol

Feature Summary

The Quality of Service (QoS) policy propagation via Border Gateway Protocol (BGP) feature allows you to classify packets based on access lists, BGP community lists, and BGP autonomous system (AS) paths. The supported classification policies include Internet Protocol (IP) precedence setting and the ability to tag the packet with a QoS class identifier internal to the router. After a packet has been classified, you can use other QoS features such as Committed Access Rate (CAR) and Weighted Random Early Detection (WRED) to specify and enforce business policies to fit your business model.

The QoS policy propagation via BGP feature was introduced in Cisco IOS Release 11.1(17)CC. With Release 11.1(20)CC, the QoS policy propagation via BGP feature has the following enhancements:

- **QoS group ID**—You can set an internal QoS group ID that can be used later to perform rate-limiting or weighted fair queuing based on the QoS group ID. In the previous release you could only set up to eight IP precedence level to classify packets. By setting the QoS group ID in addition to the IP precedence, you can now have more than eight classes on which to perform rate-limiting or weighted fair queuing.
- **Source and destination address lookup**—You can specify whether the IP precedence level or QoS group ID used is obtained from the source (input) address or destination (output) address entry in the route table. In the previous release you could only use the destination address. You can now specify the input or output address.

Benefits

BGP policy propagation provides the following benefits:

- Allows you to classify packets using access lists, community lists, and AS paths.
- Leverages BGP to distribute QoS policy to remote routers in your network.
- Allows ingress routers to prioritize incoming and outgoing traffic.
- Allows you to classify packets based on IP precedence or QoS group ID.

List of Terms

Autonomous system (AS) path—A collection of networks under a common administration sharing a common routing strategy. BGP carries the AS 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 AS path.

Border Gateway Protocol (BGP)—Interdomain routing protocol that replaces EGP. BGP exchanges reachability information with other BGP systems. It is defined by RFC 1163.

Cisco Express Forwarding (CEF)—CEF is an advanced Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks with large and dynamic traffic patterns, such as the Internet, on networks characterized by intensive Web-based applications, or interactive sessions. Although you can use CEF in any part of a network, it is designed for high-performance, highly resilient Layer 3 IP backbone switching.

Committed Access Rate (CAR)—CAR limits the input or output transmission rate on an interface or subinterface based on a flexible set of criteria. In addition, CAR classifies packets by setting the IP precedence. CAR can be used to rate-limit traffic based on packet characteristics such as access list, incoming interface, or IP precedence. CAR provides configurable actions, such as transmit, drop, or set precedence, when traffic conforms to or exceeds the rate limit.

Community list—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. Just like an access list, a series of community lists can be created.

Internet Protocol (IP) precedence—Bits within the ToS (type of service) field of the IP header that can be used to classify packets.

QoS group ID—User-specified number that is assigned to a packet when that packet matches user-specified criteria. The packet can then be classified based on that number.

Weighted Random Early Detection (WRED)—Drops packets selectively based on IP precedence. Packets with a higher IP precedence are less likely to be dropped than packets with a lower precedence. Thus, higher priority traffic is delivered with a higher probability than lower priority traffic. WRED is useful on any output interface where you expect to have congestion. However, WRED is usually used in the core routers of a network, rather than the edge. Edge routers assign IP precedences to packets as they enter the network. WRED uses these precedences to determine how it treats different types of traffic.

Document Conventions

Command descriptions use these conventions:

- **Boldface** indicates commands and keywords that are entered literally as shown.
- *Italics* indicate arguments for which you supply values; in contexts that do not allow italics, arguments are enclosed in angle brackets (>).
- Square brackets ([]) indicate optional elements.
- Braces ({ }) group required choices, and vertical bars (|) separate alternative elements.
- Braces and vertical bars within square brackets ({ | }) indicate a required choice within an optional element.

Platforms

This feature is supported on these platforms:

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

Supported MIBs and RFCs

None

Restrictions

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

Prerequisites

For the QoS policy propagation via BGP feature to work, you must enable BGP and CEF/dCEF on the router.

Configuration Tasks

This section describes the tasks required to configure QoS policy propagation via BGP and how to verify the information is correct. You can propagate QoS policy using access lists, BGP community lists, and BGP AS paths. You can use any combination of these methods. The tasks are discussed in the following sections:

- Configure Policy Propagation Based on Community Lists
- Configure Policy Propagation Based on the AS Path Attribute
- Configure Policy Propagation Based on an Access List
- Verify the Configuration

Configuring QoS policy propagation via BGP consists of the following steps:

- Step 1** Configure BGP and CEF or DCEF.
- Step 2** Define the policy.
- Step 3** Apply the policy through BGP.
- Step 4** Configure the access list, BGP community list, or BGP AS path.
- Step 5** Enable the policy on an interface.
- Step 6** Enable CAR, DWRED, or DWFQ to use the policy.

This document discusses steps 2 through 5. To configure BGP, refer to the *Network Protocols Configuration Guide, Part 1*. To configure CEF/dCEF, CAR, and WRED, refer to the appropriate feature module. Cisco IOS Release 11.1 documents and Release 11.1 CC feature documents can be found on the Documentation CD-ROM and on Cisco Connection Online (CCO).

For examples of propagating QoS policy using access lists, BGP community lists, and BGP AS paths, see the “Configuration Examples” section later in this chapter.

Configure Policy Propagation Based on Community Lists

This section describes how to configure QoS policy propagation via BGP using community lists. The tasks listed in this section are required unless noted as optional. This section assumes you have already configured CEF/dCEF and BGP on your router.



Caution If you specify both **source** and **destination** on the interface, the software looks up the source address in the routing table and classifies the packet based on the source address first; then the software looks up the destination address in the routing table and reclassifies it based on the destination address.

To configure the router to propagate the IP precedence and/or the QoS group ID based on the community lists, perform the following steps beginning in global configuration mode:

Task	Command
Step 1 Define a route map to control redistribution and enter route-map configuration mode.	route-map <i>route-map-name</i> [permit deny [<i>sequence-number</i>]]
Step 2 Match a BGP community list.	match community-list <i>community-list-number</i> [exact]
Step 3 Set the IP precedence field when the community list matches.	set ip precedence [<i>value</i> <i>name</i>]
and/or	
Step 4 Set the QoS group ID when the community list matches.	set ip qos-group <i>group-id</i>
Step 5 Enter router configuration mode.	router bgp <i>autonomous-system</i>
Step 6 Modify the metric and tag values when the IP routing table is updated with BGP learned routes.	table-map <i>route-map-name</i>
Step 7 Create a community list for BGP and control access to it.	ip community-list <i>community-list-number</i> { permit deny } <i>community-number</i>
Step 8 Specify the interfaces (or subinterface) and enter interface configuration mode.	interface <i>type number</i>
Step 9 Classify packets using the IP precedence based on the packet’s source address and/or destination address.	bgp-policy source ip-prec-map bgp-policy destination ip-prec-map
and/or	
Step 10 Classify packets using the QoS group ID based on the packet’s source address and/or destination address.	bgp-policy source ip-qos-map bgp-policy destination ip-qos-map
Step 11 Optionally configure new community format so that the community number is displayed in the short form.	ip bgp-community new-format

Task	Command
Step 12 Exit configuration mode.	end

Configure Policy Propagation Based on the AS Path Attribute

This section describes how to configure QoS policy propagation via BGP based on the AS path. The tasks listed in this section are required unless noted as optional. This section assumes you have already configured CEF/dCEF and BGP on your router.



Caution If you specify both **source** and **destination** on the interface, the software looks up the source address in the routing table and classifies the packet based on the source address first; then the software looks up the destination address in the routing table and reclassifies it based on the destination address.

To configure the router to propagate the IP precedence and QoS group ID based on the AS-path attribute, perform the following steps beginning in global configuration mode:

Task	Command
Step 1 Define a route map to control redistribution and enter route-map configuration mode.	route-map <i>route-map-name</i> [permit deny [<i>sequence-number</i>]]
Step 2 Match a BGP autonomous system path access list.	match as-path <i>path-list-number</i>
Step 3 Set the IP precedence field when the AS path matches.	set ip precedence [<i>value</i> <i>name</i>]
and/or	
Step 4 Set the QoS group ID when the AS path matches.	set ip qos-group <i>group-id</i>
Step 5 Enter router configuration mode.	router bgp <i>autonomous-system</i>
Step 6 Modify the metric and tag values when the IP routing table is updated with BGP learned routes.	table-map <i>route-map-name</i>
Step 7 Define an AS path access list.	ip as-path access-list <i>access-list-number</i> { permit deny } <i>as-regular-expression</i>
Step 8 Specify the interfaces (or subinterface) and enter interface configuration mode.	interface <i>type number</i>
Step 9 Classify packets using the IP precedence based on the packet's source address and/or destination address.	bgp-policy source ip-prec-map bgp-policy destination ip-prec-map
and/or	
Step 10 Classify packets using the QoS group ID based on the packet's source address and/or destination address.	bgp-policy source ip-qos-map bgp-policy destination ip-qos-map
Step 11 Exit configuration mode.	end

Configure Policy Propagation Based on an Access List

This section describes how to configure QoS precedence propagation via BGP based on an access list. The tasks listed in this section are required unless noted as optional. This section assumes you have already configured CEF/dCEF and BGP on your router.



Caution If you specify both **source** and **destination** on the interface, the software looks up the source address in the routing table and classifies the packet based on the source address first; then the software looks up the destination address in the routing table and reclassifies it based on the destination address.

To configure the router to propagate the IP precedence and QoS group ID based on an access list, perform the following steps beginning in global configuration mode:

Task	Command
Step 1 Define a route map to control redistribution and enter route-map configuration mode.	route-map <i>route-map-name</i> [permit deny [<i>sequence-number</i>]]
Step 2 Match an access list.	match ip address <i>access-list-number</i>
Step 3 Set the IP precedence field when the access list matches.	set ip precedence [<i>value</i> <i>name</i>]
and/or	
Step 4 Set the QoS group ID when the access list matches.	set ip qos-group <i>group-id</i>
Step 5 Enter router configuration mode.	router bgp <i>autonomous-system</i>
Step 6 Modify the metric and tag values when the IP routing table is updated with BGP learned routes.	table-map <i>route-map-name</i>
Step 7 Define an access list.	access-list <i>access-list-number</i> { permit deny } <i>source</i>
Step 8 Specify the interfaces (or subinterface) and enter interface configuration mode.	interface <i>type number</i>
Step 9 Classify packets using the IP precedence based on the packet's source address and/or destination address.	bgp-policy source ip-prec-map bgp-policy destination ip-prec-map
and/or	
Step 10 Classify packets using the QoS group ID based on the packet's source address and/or destination address.	bgp-policy source ip-qos-map bgp-policy destination ip-qos-map
Step 11 Exit configuration mode.	end

Verify the Configuration

This section describes how to verify that QoS policy propagation via BGP is configured correctly. The tasks listed in this section are optional.

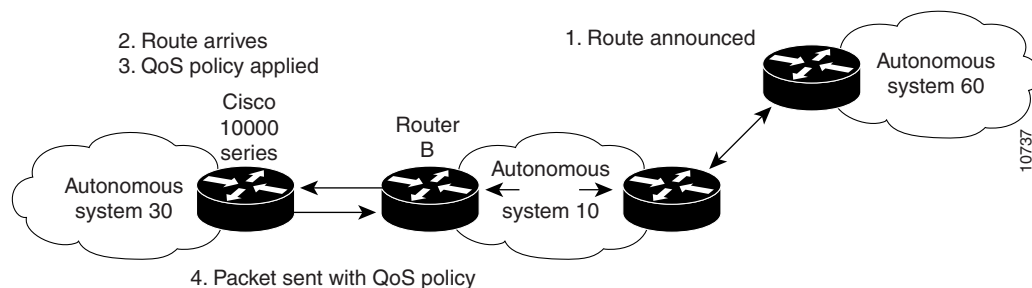
To verify the configuration, perform any of the following steps in EXEC mode:

Task	Command
To verify the correct community is set on the prefixes.	show ip bgp
To verify that the correct prefixes are selected.	show ip bgp community-list <i>community-list-number</i>
To verify that CEF has the correct precedence value for the prefix.	show ip cef <i>prefix</i>
To display information about the interface	show ip interface
To verify that the correct precedence values are set on the prefixes.	show ip route <i>prefix</i>

Configuration Examples

The following example shows how to create route maps to match access lists, BGP community lists, and BGP AS paths and apply IP precedence to routes learned from neighbors.

In this example, Router A learns routes from AS 10 and AS 60. QoS policy is applied to all packets that match the defined route maps. Any packets from Router A to AS 10 or AS 60 are sent to the appropriate QoS policy.



Router A's Configuration

```

router bgp 30
  table-map precedence-map
  neighbor 20.20.20.1 remote-as 10
  neighbor 20.20.20.1 send-community
  neighbor 20.20.20.1 route-map precedence-map out
!
ip bgp-community new-format
!
! Match community 1 and set the IP precedence to priority and set the QoS group to 1
route-map precedence-map permit 10
  match community 1
  set ip precedence priority
  set ip qos-group 1
!

```

```
! Match community 2 and set the IP precedence to immediate
route-map precedence-map permit 20
  match community 2
  set ip precedence immediate
!
! Match community 3 and set the IP precedence to flash
route-map precedence-map permit 30
  match community 3
  set ip precedence flash
!
! Match community 4 and set the IP precedence to flash-override
route-map precedence-map permit 40
  match community 4
  set ip precedence flash-override
!
! Match community 5 and set the IP precedence to critical
route-map precedence-map permit 50
  match community 5
  set ip precedence critical
!
! Match community 6 and set the IP precedence to internet
route-map precedence-map permit 60
  match community 6
  set ip precedence internet
!
! Match community 7 and set the IP precedence to network
route-map precedence-map permit 70
  match community 7
  set ip precedence network
!
! Match ip address access list 69 or match AS path 1, set the IP precedence to
! critical, and set the Qos group to 9
route-map precedence-map permit 75
  match ip address 69
  match as-path 1
  set ip precedence critical
  set ip qos-group 9
!
! For everything else, set the IP precedence to routine
route-map precedence-map permit 80
  set ip precedence routine
!
! Define the community lists
ip community-list 1 permit 60:1
ip community-list 2 permit 60:2
ip community-list 3 permit 60:3
ip community-list 4 permit 60:4
ip community-list 5 permit 60:5
ip community-list 6 permit 60:6
ip community-list 7 permit 60:7
!
! Define the AS path
ip as-path access-list 1 permit ^10_60
! Define the access list
access-list 69 permit 69.0.0.0
```

Router B's Configuration

```
router bgp 10
  neighbor 30.30.30.1 remote-as 30
  neighbor 30.30.30.1 send-community
  neighbor 30.30.30.1 route-map send_community out
!
ip bgp-community new-format
```

```

! Match prefix 10 and set community to 60:1
route-map send_community permit 10
  match ip address 10
  set community 60:1
!
! Match prefix 20 and set community to 60:2
route-map send_community permit 20
  match ip address 20
  set community 60:2
!
! Match prefix 30 and set community to 60:3
route-map send_community permit 30
  match ip address 30
  set community 60:3
!
! Match prefix 40 and set community to 60:4
route-map send_community permit 40
  match ip address 40
  set community 60:4
!
! Match prefix 50 and set community to 60:5
route-map send_community permit 50
  match ip address 50
  set community 60:5
!
! Match prefix 60 and set community to 60:6
route-map send_community permit 60
  match ip address 60
  set community 60:6
!
! Match prefix 70 and set community to 60:7
route-map send_community permit 70
  match ip address 70
  set community 60:7
!
! For all others, set community to 60:8
route-map send_community permit 80
  set community 60:8
!
! Define the access lists
access-list 10 permit 61.0.0.0
access-list 20 permit 62.0.0.0
access-list 30 permit 63.0.0.0
access-list 40 permit 64.0.0.0
access-list 50 permit 65.0.0.0
access-list 60 permit 66.0.0.0
access-list 70 permit 67.0.0.0

```

The following example shows how to configure several interfaces to classify packets based on the IP precedence and QoS group ID.

```

interface Hssi5/0/0.1 point-to-point
  ip address 200.28.38.2 255.255.255.0
  bgp-policy source ip-prec-map
  no ip mroute-cache
  no cdp enable
  frame-relay interface-dlci 20 IETF

interface Hssi6/0/0.1 point-to-point
  ip address 200.28.28.2 255.255.255.0
  bgp-policy source qos-group
  no ip mroute-cache
  no cdp enable
  frame-relay interface-dlci 20 IETF

```

Command Reference

This section documents new commands. All other commands used with this feature are documented in the Cisco IOS Release 11.1 command reference publications.

- **bgp-policy**
- **set ip qos-group**
- **show ip cef**
- **show ip interface**

bgp-policy

To enable QoS policy propagation via BGP on the interface, use the **bgp-policy** interface configuration command. To disable QoS policy propagation via BGP, use the **no** form of the command.

```
bgp-policy {source | destination} {ip-prec-map | ip-qos-map}
```

Syntax Description

source	The IP precedence bit or QoS group ID from the source address entry in the route table.
destination	The IP precedence bit or QoS group ID from the destination address entry in the route table.
ip-prec-map	The QoS policy based on the IP precedence.
ip-qos-map	The QoS policy based on the QoS group ID.

Default

Disabled

Command Mode

Interface configuration

Usage Guidelines

This command first appeared in Cisco IOS Release 11.1 CC. This command was modified in Cisco IOS Release 11.1(20)CC to include the **input**, **output**, and **ip-qos-map** keywords. This command was modified in Cisco IOS Release 11.1(21)CC to change the **input** keyword to **source** and the **output** keyword to **destination**.

For the QoS policy propagation via BGP feature to work, you must enable BGP and CEF/dCEF. In addition, the proper route-map configuration must be in place to specify the IP precedence or QoS group ID (for example, **set ip precedence** route-map configuration command).



Caution If you specify both **source** and **destination** on the interface, the software looks up the source address in the routing table and classifies the packet based on the source address first; then the software looks up the destination address in the routing table and reclassifies it based on the destination address.

To display QoS policy information for the interface, use the **show ip interface** command.

Examples

The following example enables QoS policy propagation via BGP on an interface based on the source address and the IP precedence setting. For a complete configuration example, refer to the “Configuration Examples” section earlier in this document.

```
router# configure terminal
router(config)# interface ethernet 4/0/0
router(config-if)# bgp-policy source ip-prec-map
router(config-if)# end
router#
```

set ip qos-group

To set a group ID that can be used later to classify packets, use the **set ip qos-group** route-map configuration command. To remove the group ID, use the **no** form of this command.

```
set ip qos-group group-id  
no set ip qos-group group-id
```

Syntax Description

group-id Group ID number. Range is 0 to 99

Default

No group ID is specified.

Command Mode

Route map configuration

Usage Guidelines

This command first appeared in Cisco IOS Release 11.1 CC.

This feature allows you to set a group ID in the routing table that can be used later to classify packets into QoS groups based on prefix, AS, and community string. These packets can then be rate limited or weighted fair queued based on the QoS group ID.

To display QoS group information, use the **show ip cef** command.

Examples

The following example sets the QoS group to 1 for all packets that match community 1. These packets are then rate limited based on the QoS group ID. For a complete configuration example, refer to the “Configuration Examples” section earlier in this document.

```
router# configure terminal  
router(config)# route-map precedence-map permit 10  
router(config)# match community 1  
router(config)# set ip qos-group 1  
router(config)# interface hssi0/0/0  
router(config-if)# bgp-policy source qos-group  
router(config-if)# end
```

show ip cef

To display entries in the FIB table based on the IP address, use the **show ip cef** EXEC command.

```
show ip cef network [mask [longer-prefix] [detail]
```

Syntax Description

<i>network</i>	Displays the FIB entry for the specific destination network.
<i>mask</i>	(Optional) Displays the FIB entry for the specified destination network and mask.
longer-prefix	(Optional) Displays the FIB entries for all more specific destinations.
detail	(Optional) Displays detailed FIB information.

Command Mode

EXEC

Usage Guidelines

This command was updated in Cisco IOS Release 11.1 CC to add information on the QoS group ID.

Sample Display

The following is sample output from the **show ip cef** command for the network address 51.0.0.0:

```
Router# show ip cef 51.0.0.0  
51.0.0.0/8, version 161, cached adjacency 200.31.51.2  
0 packets, 0 bytes, precedence priority (1), qos-group 1  
via 50.50.50.1, 0 dependencies, recursive  
next hop 200.31.51.2, FastEthernet5/1/0 via 50.0.0.0/8  
valid cached adjacency
```

show ip interface

To display the usability status of interfaces configured for IP, use the **show ip interface EXEC** command.

```
show ip interface [type number]
```

Syntax Description

<i>type</i>	(Optional) Interface type.
<i>number</i>	(Optional) Interface number.

Command Mode

EXEC

Usage Guidelines

This command first appeared in Cisco IOS Release 10.0. This command was modified in Cisco IOS Release 11.1 CC to add information on QoS policy propagation.

The Cisco IOS software automatically enters a directly connected route in the routing table if the interface is usable. A usable interface is one through which the software can send and receive packets. If the software determines that an interface is not usable, it removes the directly connected routing entry from the routing table. Removing the entry allows the software to use dynamic routing protocols to determine backup routes to the network (if any).

If the interface can provide two-way communication, the line protocol is marked “up.” If the interface hardware is usable, the interface is marked “up.”

If you specify an optional interface type, you will see only information on that specific interface.

If you specify no optional arguments, you will see information on all the interfaces.

When an asynchronous interface is encapsulated with PPP or SLIP, IP fast switching is enabled. A **show ip interface** command on an asynchronous interface encapsulated with PPP or SLIP displays a message indicating that IP fast switching is enabled.

Sample Display

The following is sample output from the **show ip interface** command:

```
Router# show ip interface hssi 5/0/0.1  
Hssi5/0/0.1 is up, line protocol is up  
Internet address is 200.28.38.2/24  
Broadcast address is 255.255.255.255  
Address determined by non-volatile memory  
MTU is 4470 bytes  
Helper address is not set  
Directed broadcast forwarding is enabled  
Outgoing access list is not set  
Inbound access list is not set  
Proxy ARP is enabled  
Security level is default  
Split horizon is enabled  
ICMP redirects are always sent  
ICMP unreachable are always sent  
ICMP mask replies are never sent
```

```

IP fast switching is enabled
IP fast switching on the same interface is enabled
IP Optimum switching is disabled
IP Flow switching is enabled
IP CEF switching is enabled
IP Distributed switching is enabled
IP LES Flow switching turbo vector
IP Flow CEF switching turbo vector
IP multicast fast switching is disabled
IP multicast distributed fast switching is disabled
Router Discovery is disabled
IP output packet accounting is disabled
IP access violation accounting is disabled
TCP/IP header compression is disabled
Probe proxy name replies are disabled
Gateway Discovery is disabled
Policy routing is disabled
Web Cache Redirect is disabled
BGP Policy Mapping is enabled (source ip-prec-map)

```

Table 1 describes the fields in the display.

Table 1 Show IP Interface Field Descriptions

Field	Description
Hssi5/0/0.1 is up	If the interface hardware is usable, the interface is marked "up." For an interface to be usable, both the interface hardware and line protocol must be up.
line protocol is up	If the interface can provide two-way communication, the line protocol is marked "up." For an interface to be usable, both the interface hardware and line protocol must be up.
Internet address and subnet mask	IP Internet address and subnet mask of the interface.
Broadcast address	Shows the broadcast address.
Address determined by...	Indicates how the IP address of the interface was determined.
MTU	Shows the MTU value set on the interface.
Helper address	Shows a helper address, if one has been set.
Secondary address	Shows a secondary address, if one has been set.
Directed broadcast forwarding	Indicates whether directed broadcast forwarding is enabled.
Multicast groups joined	Indicates the multicast groups this interface is a member of.
Outgoing access list	Indicates whether the interface has an outgoing access list set.
Inbound access list	Indicates whether the interface has an incoming access list set.
Proxy ARP	Indicates whether Proxy ARP is enabled for the interface.
Security level	Specifies the IPSO security level set for this interface.
Split horizon	Indicates split horizon is enabled.
ICMP redirects	Specifies whether redirects will be sent on this interface.
ICMP unreachable	Specifies whether unreachable messages will be sent on this interface.
ICMP mask replies	Specifies whether mask replies will be sent on this interface.
IP fast switching	Specifies whether fast switching has been enabled for this interface. It is generally enabled on serial interfaces, such as this one.

Table 1 Show IP Interface Field Descriptions (Continued)

Field	Description
IP fast switching on the same interface	Specifies whether fast switching has been enabled on the same interface.
IP Optimum switching	Specifies whether IP Optimum switching is enabled.
IP Flow switching	Specifies whether IP Flow switching is enabled.
IP CEF switching	Specifies whether IP CEF switching is enabled.
IP LES Flow switching	Specifies whether the IP LES Flow switching is enabled.
IP Flow CEF switching	Specifies whether the IP Flow CEF switching is enabled.
IP multicast fast switching	Specifies whether IP multicast fast switching is enabled.
IP multicast distributed fast switching	Specifies whether IP multicast distributed fast switching is enabled.
Router Discovery	Specifies whether the discovery process has been enabled for this interface. It is generally disabled on serial interfaces.
IP output packet accounting	Specifies whether IP accounting is enabled for this interface and what the threshold (maximum number of entries) is.
TCP/IP header compression	Indicates whether compression is enabled or disabled.
Probe proxy name	Indicates whether HP Probe proxy name replies are generated.
Gateway Discovery	Specifies whether gateway discover is enabled.
Policy routing	Specifies whether policy routing is enabled.
Web Cache Redirect	Specifies whether web cache redirect is enabled.
BGP Policy Mapping	Specified whether BGP policy mapping is enabled and the current settings on the interface.