

BGP Policy Accounting

Feature History

Release	Modification
12.0(9)S	This feature was introduced.

This document describes the BGP Policy Accounting feature in Cisco IOS Release 12.0(9)S. It includes the following sections:

- [Feature Overview, page 1](#)
- [Supported Platforms, page 3](#)
- [Supported Standards, MIBs, and RFCs, page 3](#)
- [Prerequisites, page 4](#)
- [Configuration Tasks, page 4](#)
- [Monitoring and Maintaining BGP Policy Accounting, page 7](#)
- [Configuration Examples, page 7](#)
- [Command Reference, page 8](#)
- [Glossary, page 21](#)


Feature Overview

Border Gateway Protocol (BGP) policy accounting measures and classifies IP traffic that is sent to, or received from, different peers. Policy accounting is enabled on an input interface, and counters based on parameters such as community list, autonomous system number, or autonomous system path are assigned to identify the IP traffic.

Using the BGP **table-map** command, prefixes added to the routing table are classified by BGP attribute, autonomous system number, or autonomous system path. Packet and byte counters are incremented per input interface. A Cisco IOS policy-based classifier maps the traffic into one of eight possible buckets, representing different traffic classes.

Using BGP policy accounting, you can account for traffic according to the route it traverses. Service providers (SPs) can identify and account for all traffic by customer and bill accordingly. In [Figure 1](#), BGP policy accounting can be implemented in Router A to measure packet and byte volumes in autonomous system buckets. Customers are billed appropriately for traffic that is routed from a domestic, international, or satellite source.

Figure 1 Sample Topology for BGP Policy Accounting



BGP policy accounting using autonomous system numbers can be used to improve the design of network circuit peering and transit agreements between Internet service providers (ISPs).

Benefits

Account for IP Traffic Differentially

BGP policy accounting classifies IP traffic by autonomous system number, autonomous system path, or community list string, and increments packet and byte counters. Service providers can account for traffic and apply billing, according to the route specific traffic traverses.

Efficient Network Circuit Peering and Transit Agreement Design

Implementing BGP policy accounting on an edge router can highlight potential design improvements for peering and transit agreements.

Related Features and Technologies

Additional BGP command and configuration information is documented in the “Configuring BGP” chapter of the *Cisco IOS Release 12.0 Network Protocols Configuration Guide, Part 1* and *Cisco IOS Release 12.0 Network Protocols Command Reference, Part 1*.

Additional Cisco Express Forwarding (CEF) and distributed CEF (dCEF) command and configuration information is documented in the “Cisco Express Forwarding Overview” and in the “Configuring Cisco Express Forwarding” chapters of the Release 12.0 *Cisco IOS Switching Services Configuration Guide*, and in the *Cisco IOS Switching Services Command Reference*.

Related Documents

- Cisco IOS Release 12.0 *Network Protocols Configuration Guide, Part 1*:
- Cisco IOS Release 12.0 *Network Protocols Command Reference, Part 1*:

- Cisco IOS Release 12.0 *Cisco IOS Switching Services Configuration Guide*:
- Cisco IOS Release 12.0 *Cisco IOS Switching Services Command Reference*:

Supported Platforms

The BGP Policy Accounting feature is supported by the following platforms that support Cisco IOS Release 12.0(9)S:

- Cisco 7200 series
- Cisco 7500 series
- Cisco 12000 series

Determining Platform Support Through Cisco Feature Navigator

Cisco IOS software is packaged in feature sets that are supported by specific platforms. To get updated information regarding platform support for this feature, access Cisco Feature Navigator. Cisco Feature Navigator dynamically updates the list of supported platforms as new platform support is added for the feature.

Cisco Feature Navigator is a web-based tool that enables you to determine which Cisco IOS software images support a specific set of features and which features are supported in a specific Cisco IOS image. You can search by feature or release. Under the release section, you can compare releases side by side to display both the features unique to each software release and the features in common.

To access Cisco Feature Navigator, you must have an account on Cisco.com. If you have forgotten or lost your account information, send a blank e-mail to cco-locksmith@cisco.com. An automatic check will verify that your e-mail address is registered with Cisco.com. If the check is successful, account details with a new random password will be e-mailed to you. Qualified users can establish an account on Cisco.com by following the directions at <http://www.cisco.com/register>.

Cisco Feature Navigator is updated regularly when major Cisco IOS software releases and technology releases occur. For the most current information, go to the Cisco Feature Navigator home page at the following URL:

<http://www.cisco.com/go/fn>

Availability of Cisco IOS Software Images

Platform support for particular Cisco IOS software releases is dependent on the availability of the software images for those platforms. Software images for some platforms may be deferred, delayed, or changed without prior notice. For updated information about platform support and availability of software images for each Cisco IOS software release, refer to the online release notes or, if supported, Cisco Feature Navigator.

Supported Standards, MIBs, and RFCs

Standards

No new or modified standards are supported by this feature.

MIBs

- CISCO-BGP-POLICY-ACCOUNTING-MIB

To obtain lists of supported MIBs by platform and Cisco IOS release, and to download MIB modules, go to the Cisco MIB website on Cisco.com at the following URL:

<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

RFCs

No new or modified RFCs are supported by this feature.

Prerequisites

Before using the BGP Policy Accounting feature you must enable BGP and CEF or dCEF on the router.

Configuration Tasks

See the following sections for configuration tasks for the BGP Policy Accounting feature. Each task in the list is identified as either required or optional.

- [Specifying the Match Criteria for BGP Policy Accounting, page 4](#) (required)
- [Classifying the IP Traffic and Enabling BGP Policy Accounting, page 5](#) (required)
- [Verifying BGP Policy Accounting, page 6](#) (optional)

Specifying the Match Criteria for BGP Policy Accounting

The first task in configuring BGP policy accounting is to specify the criteria that must be matched. Community lists, autonomous system paths, or autonomous system numbers are examples of BGP attributes that can be specified and subsequently matched using a route map.

To specify the BGP attribute to use for BGP policy accounting and create the match criteria in a route map, use the following commands in global configuration mode:

	Command	Purpose
Step 1	Router(config)# ip community-list <i>community-list-number</i> { permit deny } <i>community-number</i>	Creates a community list for BGP and controls access to it. This step must be repeated for each community to be specified.
Step 2	Router(config)# route-map <i>map-name</i> [permit deny] [<i>sequence-number</i>]	Enters route-map configuration mode and defines the conditions for policy routing. The <i>map-name</i> argument identifies a route map. The optional permit and deny keywords work with the match and set criteria to control how the packets are accounted for. The optional <i>sequence-number</i> argument indicates the position a new route map is to have in the list of route maps already configured with the same name.

	Command	Purpose
Step 3	Router(config-route-map)# match community-list <i>community-list-number</i> [exact]	Matches a BGP community.
Step 4	Router(config-route-map)# set traffic-index <i>bucket-number</i>	Indicates where to output packets that pass a match clause of a route map for BGP policy accounting.

Classifying the IP Traffic and Enabling BGP Policy Accounting

After a route map has been defined to specify match criteria, you must configure a way to classify the IP traffic before enabling BGP policy accounting.

Using the **table-map** command, BGP classifies each prefix it adds to the routing table based on the match criteria. When the **bgp-policy accounting** command is configured on an interface, BGP policy accounting is enabled.

To classify the IP traffic and enable BGP policy accounting, use the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	Router(config)# router bgp <i>as-number</i>	Configures a BGP routing process and enters router configuration mode for the specified routing process.
Step 2	Router(config-router)# table-map <i>route-map-name</i>	Classifies BGP prefixes entered in the routing table.
Step 3	Router(config-router)# network <i>network-number</i> [mask <i>network-mask</i>]	Specifies a network to be advertised by the BGP routing process.
Step 4	Router(config-router)# neighbor <i>ip-address</i> remote-as <i>as-number</i>	Specifies a BGP peer by adding an entry to the BGP routing table.
Step 5	Router(config-router)# exit	Exits to global configuration mode.
Step 6	Router(config)# interface <i>interface-type</i> <i>interface-number</i>	Specifies the interface type and number and enters interface configuration mode.
Step 7	Router(config-if)# no ip directed-broadcast	Configures the interface to drop directed broadcasts destined for the subnet to which that interface is attached, rather than being broadcast. This is a security issue.
Step 8	Router(config-if)# ip address <i>ip-address</i> <i>mask</i>	Configures the interface with an IP address.
Step 9	Router(config-if)# bgp-policy accounting	Enables BGP policy accounting for the interface. Note If the prefixes classified by the route-map in Step 2 were installed in the Forwarding Information Base (FIB) table prior to adding or modifying the BGP policy accounting configuration, the FIB must be refreshed using the clear ip route command and specifying the classified prefixes for the command arguments.

Verifying BGP Policy Accounting

To verify that BGP policy accounting is operating, perform the following steps:

- Step 1** Enter the **show ip cef EXEC** command with the **detail** keyword to learn which accounting bucket is assigned to a specified prefix.

In this example, the output is displayed for the prefix 192.168.5.0. It shows that the accounting bucket number 4 (traffic_index 4) is assigned to this prefix.

```
Router# show ip cef 192.168.5.0 detail

192.168.5.0/24, version 21, cached adjacency to POS7/2
0 packets, 0 bytes, traffic_index 4
  via 10.14.1.1, 0 dependencies, recursive
  next hop 10.14.1.1, POS7/2 via 10.14.1.0/30
  valid cached adjacency
```

- Step 2** Enter the **show ip bgp EXEC** command for the same prefix used in Step 1—192.168.5.0— to learn which community is assigned to this prefix.

In this example, the output is displayed for the prefix 192.168.5.0. It shows that the community of 100:197 is assigned to this prefix.

```
Router# show ip bgp 192.168.5.0

BGP routing table entry for 192.168.5.0/24, version 2
Paths: (1 available, best #1)
  Not advertised to any peer
  100
    10.14.1.1 from 10.14.1.1 (32.32.32.32)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: 100:197
```

- Step 3** Enter the **show cef interface policy-statistics EXEC** command to display the per-interface traffic statistics.

In this example, the output shows the number of packets and bytes that have been assigned to each accounting bucket:

```
LC-Slot7# show cef interface policy-statistics

POS7/0 is up (if_number 8)
Bucket   Packets           Bytes
-----
1         0                 0
2         0                 0
3         50                5000
4        100               10000
5        100               10000
6         10                1000
7         0                 0
8         0                 0
```

Monitoring and Maintaining BGP Policy Accounting

To monitor and maintain the BGP Policy Accounting feature, use the following commands in EXEC mode, as needed:

Command	Purpose
Router# show cef interface [<i>type number</i>] policy-statistics	Displays detailed CEF policy statistical information for all interfaces.
Router# show ip bgp [<i>network</i>] [<i>network mask</i>] [longer-prefixes]	Displays entries in the BGP routing table.
Router# show ip cef [<i>network [mask]</i>] [detail]	Displays entries in the Forwarding Information Base (FIB) or FIB summary information.

Configuration Examples

This section provides the following configuration examples:

- [Specifying the Match Criteria for BGP Policy Accounting Example](#)
- [Classifying the IP Traffic and Enabling BGP Policy Accounting Example](#)

Specifying the Match Criteria for BGP Policy Accounting Example

In the following example, BGP communities are specified in community lists, and a route map named `set_bucket` is configured to match each of the community lists to a specific accounting bucket using the `set traffic-index` command:

```
ip community-list 30 permit 100:190
ip community-list 40 permit 100:198
ip community-list 50 permit 100:197
ip community-list 60 permit 100:296
!
route-map set_bucket permit 10
match community 30
set traffic-index 2
!
route-map set_bucket permit 20
match community 40
set traffic-index 3
!
route-map set_bucket permit 30
match community 50
set traffic-index 4
!
route-map set_bucket permit 40
match community 60
set traffic-index 5
```

Classifying the IP Traffic and Enabling BGP Policy Accounting Example

In the following example, BGP policy accounting is enabled on POS interface 7/0 and the **table-map** command is used to modify the bucket number when the IP routing table is updated with routes learned from BGP:

```
router bgp 65000
  table-map set_bucket
  network 10.15.1.0 mask 255.255.255.0
  neighbor 10.14.1.1 remote-as 65100
!
ip classless
ip bgp-community new-format
!
interface POS7/0
  ip address 10.15.1.2 255.255.255.0
  no ip directed-broadcast
  bgp-policy accounting
  no keepalive
  crc 32
  clock source internal
```

Command Reference

This section documents new and modified commands. All other commands used with this feature are documented in the Cisco IOS Release 12.0 command reference publications.

- [bgp-policy](#)
- [set traffic-index](#)
- [show cef interface policy-statistics](#)
- [show ip bgp](#)
- [show ip cef](#)

bgp-policy

To enable Border Gateway Protocol (BGP) policy accounting or policy propagation on an interface, use the **bgp-policy** command in interface configuration mode. To disable BGP policy propagation or policy accounting, use the **no** form of this command.

bgp-policy {**accounting** | **ip-prec-map**}

no bgp-policy {**accounting** | **ip-prec-map**}

Syntax Description

accounting	Accounting policy based on community lists, autonomous system numbers, or autonomous system paths.
ip-prec-map	Quality of service (QoS) policy based on the IP precedence.

Defaults

BGP policy accounting and policy propagation are disabled.

Command Modes

Interface configuration

Command History

Release	Modification
11.1 CC	This command was introduced.
12.0(9) S	This command was integrated into Cisco IOS Release 12.0(9)S and the accounting keyword was added.

Usage Guidelines

For BGP policy propagation to function, you must enable BGP and either Cisco Express Forwarding (CEF) or distributed CEF (dCEF).

To specify the QoS policy based on the IP precedence, the proper route-map configuration must be in place (for example, the **set ip precedence** route-map configuration command). To display QoS policy information for the interface, use the **show ip interface** command.



Note

If you specify both the source and destination addresses when configuring policy propagation based on an access control list (ACL), 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 the packet based on the destination address.

To specify the accounting policy, the proper route-map configuration must be in place matching specific BGP attributes using the **set traffic-index** command. In BGP router configuration mode use the **table-map** command to modify the accounting buckets when the IP routing table is updated with routes learned from BGP. To display accounting policy information, use the **show cef interface policy-statistics**, the **show ip bgp**, and the **show ip cef detail EXEC** commands.

Examples

The following example enables the BGP policy propagation feature on an interface based on the source address and the IP precedence setting:

```
configure terminal
interface ethernet 4/0/0
  bgp-policy ip-prec-map
end
```

The following example enables the BGP policy accounting feature on GE-WAN interface 9/1. The policy is classified by autonomous system paths.

```
router bgp 65000
no synchronization
  table-map buckets
!
ip as-path access-list 1 permit _10_
ip as-path access-list 2 permit _11_
!
route-map buckets permit 10
  match as-path 1
  set traffic-index 1
!
route-map buckets permit 20
  match as-path 2
  set traffic-index 2
!
route-map buckets permit 80
  set traffic-index 7
!
interface GE-WAN9/1
  ip address 10.0.2.2 255.255.255.0
  bgp-policy accounting
  no negotiation auto
```

Related Commands

Command	Description
set ip precedence	Sets the precedence values in the IP header.
set traffic-index	Defines where to output packets that pass a match clause of a route map for BGP policy accounting.
show cef interface policy-statistics	Displays detailed CEF policy statistical information for all interfaces.
show ip bgp	Displays entries in the BGP routing table.
show ip cef	Displays entries in the FIB or FIB summary information.
show ip interface	Displays the usability status of interfaces.
table-map	Classifies routes according to a route map.

set traffic-index

To indicate where to output packets that pass a match clause of a route map for Border Gateway Protocol (BGP) policy accounting, use the **set traffic-index** command in route-map configuration mode. To delete an entry, use the **no** form of this command.

set traffic-index *bucket-number*

no set traffic-index *bucket-number*

Syntax Description	<i>bucket-number</i> Number, in the range from 1 to 8, representing a bucket into which packet and byte statistics are collected for a specific traffic classification.
---------------------------	---

Defaults	Routing traffic is not classified.
-----------------	------------------------------------

Command Modes	Route-map configuration
----------------------	-------------------------

Command History	Release	Modification
	12.0(9)S	This command was introduced.

Usage Guidelines	Use the set traffic-index route-map configuration command, the route-map global configuration command, and the match route-map configuration command to define the conditions for BGP policy accounting. The match commands specify the <i>match criteria</i> —the conditions under which policy routing occurs. The set traffic-index command specifies the <i>set actions</i> —the particular routing actions to perform if the criteria enforced by the match commands are met.
-------------------------	--

Examples	In the following example, an index for BGP policy accounting is set according to autonomous system path criteria:
-----------------	---

```
route-map buckets permit 10
 match as-path 1
 set traffic-index 1
```

Related Commands	Command	Description
	bgp-policy	Enables BGP policy accounting or policy propagation on an interface.
	route-map	Defines the conditions for redistributing routes from one routing protocol to another, or enables policy routing.

show cef interface policy-statistics

To display detailed Cisco Express Forwarding (CEF) policy statistical information for all interfaces, use the **show cef interface policy-statistics** command in EXEC mode.

show cef interface [*type number*] **policy-statistics**

Syntax Description	<i>type number</i> (Optional) Interface type and number.
---------------------------	--

Command Modes	EXEC
----------------------	------

Command History	Release	Modification
	12.0(9)S	This command was introduced to support the Cisco 12000 series Internet routers.
	12.0(14)ST	This command was introduced to support the Cisco 12000 series Internet routers.
	12.2(11)S	This command was integrated into Cisco IOS Release 12.2(11)S.

Usage Guidelines	This command is available only on distributed switching platforms.
-------------------------	--

The *type* and *number* arguments display CEF status information for the specified interface type and number.

Examples	The following is sample output from the show cef interface policy-statistics command:
-----------------	--

```
Router# show cef interface policy-statistics
```

```
POS7/0 is up (if_number 8)
Index   Packets           Bytes
-----
1        0                   0
2        0                   0
3        50                  5000
4       100                 10000
5       100                 10000
6        10                  1000
7         0                   0
8         0                   0
```

The following is sample output from the **show cef interface policy-statistics** command showing policy statistics for Ethernet interface 1/0:

```
Router# show cef interface ethernet 1/0 policy-statistics
```

```
Ethernet1/0 is up (if_number 3)
  Corresponding hwidb fast_if_number 3
  Corresponding hwidb firstsw->if_number 3
Index   Packets           Bytes
-----
1        0                   0
```

2	0	0
3	0	0
4	0	0
5	0	0
6	0	0
7	0	0
8	0	0

Table 1 describes the significant fields shown in the display.

Table 1 *show cef interface policy-statistics Field Descriptions*

Field	Description
Index	Traffic index set with the route-map command.
Packts	Number of packets switched matching the index definition.
Bytes	Number of bytes switched matching the index definition.

Related Commands

Command	Description
show cef	Displays which packets the line cards dropped or displays which packets were not express-forwarded.
show cef linecard	Displays CEF-related interface information by line card.

show ip bgp

To display entries in the Border Gateway Protocol (BGP) routing table, use the **show ip bgp** command in EXEC mode.

show ip bgp [*network*] [*network-mask*] [**longer-prefixes**]

Syntax Description

<i>network</i>	(Optional) Network number, entered to display a particular network in the BGP routing table.
<i>network-mask</i>	(Optional) Network mask address.
longer-prefixes	(Optional) Displays the route and more specific routes.

Command Modes

EXEC

Command History

Release	Modification
10.0	This command was introduced.
12.0	The display of prefix advertisement statistics was added.
12.0(6)T	The display of a message indicating support for route refresh capability was added.
12.0(9)S	The display of a message indicating support for BGP policy accounting was added.

Examples

The following is sample output from the **show ip bgp** command in privileged EXEC mode:

```
Router# show ip bgp

BGP table version is 5, local router ID is 10.0.33.34
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0            0         32768 ?
* 2.0.0.0          10.0.33.35         10         0 35 ?
*>                  0.0.0.0            0         32768 ?
* 10.0.0.0         10.0.33.35         10         0 35 ?
*>                  0.0.0.0            0         32768 ?
*> 192.168.0.0/16  10.0.33.35         10         0 35 ?
```

[Table 2](#) describes the significant fields shown in the display.

Table 2 show ip bgp Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.

Table 2 *show ip bgp Field Descriptions (continued)*

Field	Description
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	IP address of a network entity.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.
Metric	If shown, the value of the interautonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

The following is sample output from the **show ip bgp** command in privileged EXEC mode when you specify the **longer-prefixes** keyword:

```
Router# show ip bgp 198.92.0.0 255.255.0.0 longer-prefixes
```

```
BGP table version is 1738, local router ID is 198.168.72.24
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 198.168.0.0	198.168.72.30	8896		32768	?
*	198.168.72.30			0	109 108 ?
*> 198.168.1.0	198.168.72.30	8796		32768	?
*	198.168.72.30			0	109 108 ?
*> 198.168.11.0	198.168.72.30	42482		32768	?
*	198.168.72.30			0	109 108 ?
*> 198.168.14.0	198.168.72.30	8796		32768	?
*	198.168.72.30			0	109 108 ?
*> 198.168.15.0	198.168.72.30	8696		32768	?
*	198.168.72.30			0	109 108 ?
*> 198.168.16.0	198.168.72.30	1400		32768	?
*	198.168.72.30			0	109 108 ?
*> 198.168.17.0	198.168.72.30	1400		32768	?
*	198.168.72.30			0	109 108 ?
*> 198.168.18.0	198.168.72.30	8876		32768	?
*	198.168.72.30			0	109 108 ?
*> 198.168.19.0	198.168.72.30	8876		32768	?
*	198.168.72.30			0	109 108 ??

The following is sample output from the **show ip bgp** command in privileged EXEC mode, showing information for prefix 10.3.0.0:

```
Router# show ip bgp 10.3.0.0
```

```
BGP routing table entry for 10.3.0.0/8, version 628
Paths: (1 available, best #1)
Advertised to peer-groups:
  ebgp
Advertised to non-peer-group peers:
  172.17.232.162
  109 65000 297 701 80
  172.17.233.56 from 172.17.233.56 (172.19.185.32)
  Origin incomplete, localpref 100, valid, external, best, ref 2
```

The following is sample output from the **show ip bgp** command for the prefix 192.168.5.0, showing that the community of 100:197 is assigned to this prefix:

```
Router# show ip bgp 192.168.5.0
```

```
BGP routing table entry for 192.168.5.0/24, version 2
Paths: (1 available, best #1)
Not advertised to any peer
100
  10.14.1.1 from 10.14.1.1 (172.18.32.32)
  Origin IGP, metric 0, localpref 100, valid, external, best
  Community: 100:197
```



Note

If a prefix has not been advertised to any peer, the display shows “Not advertised to any peer.”

Related Commands

Command	Description
clear ip bgp	Resets a BGP connection or session.
neighbor soft-reconfiguration	Configures the Cisco IOS software to start storing updates.

show ip cef

To display entries in the Forwarding Information Base (FIB) or to display a summary of the FIB, use the **show ip cef** command in EXEC mode:

```
show ip cef [vrf vrf-name] [unresolved [detail] | {detail | summary}]
```

Specific FIB Entries Based on IP Address Information

```
show ip cef [vrf vrf-name] [network [mask]] [longer-prefixes] [detail]
```

Specific FIB Entries Based on Interface Information

```
show ip cef [vrf vrf-name] [type number] [detail]
```

Specific FIB Entries Based on Nonrecursive Routes

```
show ip cef [vrf vrf-name] non-recursive [detail]
```

Syntax Description		
vrf	(Optional) A Virtual Private Network (VPN) routing and forwarding (VRF) instance.	
<i>vrf-name</i>	(Optional) Name assigned to the VRF.	
unresolved	(Optional) Displays unresolved FIB entries.	
detail	(Optional) Displays detailed FIB entry information.	
summary	(Optional) Displays a summary of the FIB.	
<i>network</i>	(Optional) Network number for which to display a FIB entry.	
<i>mask</i>	(Optional) Network mask to be used with the specified <i>network</i> .	
longer-prefixes	(Optional) Displays FIB entries for more specific destinations.	
<i>type number</i>	(Optional) Interface type and number for which to display FIB entries.	
non-recursive	Displays only nonrecursive routes.	

Command Modes EXEC

Command History	Release	Modification
	11.2 GS	This command was introduced to support the Cisco 12012 Internet router.
	11.1 CC	Multiple platform support was added.
	12.0(5)T	The vrf keyword was added.
	12.0(9)S	The display of a message indicating support for Border Gateway Protocol (BGP) policy accounting was added.

Usage Guidelines The **show ip cef** command without any keywords or arguments shows a brief display of all FIB entries. The **show ip cef detail** command shows detailed FIB entry information for all FIB entries.

Examples

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

```
Router# show ip cef unresolved

IP Distributed CEF with switching (Table Version 136632)
45776 routes, 13 unresolved routes (0 old, 13 new)
45776 leaves, 2868 nodes, 8441480 bytes, 136632 inserts, 90856 invalidations
1 load sharing elements, 208 bytes, 1 references
1 CEF resets, 1 revisions of existing leaves
refcounts: 527292 leaf, 465617 node

10.214.0.0/16, version 136622
0 packets, 0 bytes
  via 172.17.233.56, 0 dependencies, recursive
  unresolved
10.215.0.0/16, version 136623
0 packets, 0 bytes
  via 172.17.233.56, 0 dependencies, recursive
  unresolved
10.218.0.0/16, version 136624
0 packets, 0 bytes
```

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

```
Router# show ip cef summary

IP Distributed CEF with switching (Table Version 135165)
45788 routes, 0 reresolve, 4 unresolved routes (0 old, 4 new)
45788 leaves, 2868 nodes, 8442864 bytes, 135165 inserts, 89377 invalidations
0 load sharing elements, 0 bytes, 0 references
1 CEF resets, 0 revisions of existing leaves
refcounts: 527870 leaf, 466167 node
```

The following is sample output from the **show ip cef detail** command for Ethernet interface 0. It shows all the prefixes resolving through adjacency pointing to next hop Ethernet interface 0/0 and next hop interface IP address 172.19.233.33.

```
Router# show ip cef e0/0 172.19.233.33 detail

IP Distributed CEF with switching (Table Version 136808)
45800 routes, 8 unresolved routes (0 old, 8 new) 45800 leaves, 2868 nodes, 8444360 bytes,
136808 inserts, 91008 invalidations 1 load sharing elements, 208 bytes, 1 references 1 CEF
resets, 1 revisions of existing leaves refcounts: 527343 leaf, 465638 node

172.19.233.33/32, version 7417, cached adjacency 172.19.233.33 0 packets, 0 bytes,
Adjacency-prefix
via 172.19.233.33, Ethernet0/0, 0 dependencies
next hop 172.19.233.33, Ethernet0/0
valid cached adjacency
```

The following is sample output from the **show ip cef detail** command for the prefix 192.168.5.0, showing that the BGP policy accounting bucket number 4 (traffic_index 4) is assigned to this prefix:

```
Router# show ip cef 192.168.5.0 detail

192.168.5.0/24, version 21, cached adjacency to POS7/2
0 packets, 0 bytes, traffic_index 4
via 10.14.1.1, 0 dependencies, recursive
next hop 10.14.1.1, POS7/2 via 10.14.1.0/30
valid cached adjacency
```

The following example shows the forwarding table associated with the VRF named vrf1:

```
Router# show ip cef vrf vrf1

Prefix          Next Hop          Interface
0.0.0.0/32      receive
10.11.0.0/16    10.50.0.1        Ethernet1/3
10.12.0.0/16    10.52.0.2        POS6/0
10.50.0.0/16    attached         Ethernet1/3
10.50.0.0/32    receive
10.50.0.1/32    10.50.0.1        Ethernet1/3
10.50.0.2/32    receive
10.255.255.255/32 receive
10.51.0.0/16    10.52.0.2        POS6/0
224.0.0.0/24    receive
255.255.255.255/32 receive
```

Table 3 describes the significant fields shown in the display.

Table 3 show ip cef vrf Field Descriptions

Field	Description
Prefix	Specifies the network prefix.
Next Hop	Specifies the BGP next hop address.
Interface	Specifies the VRF interface.

Related Commands

Command	Description
show cef	Displays which packets the line cards dropped or displays which packets were not express-forwarded.
show cef interface	Displays CEF-related interface information.

Glossary

AS—autonomous system. An IP term to describe a routing domain that has its own independent routing policy, and is administered by a single authority.

BGP—Border Gateway Protocol. Interdomain routing protocol that exchanges reachability information with other BGP systems.

CEF—Cisco Express Forwarding.

dCEF—distributed Cisco Express Forwarding.

