



BGP Policy Accounting

Feature History

Release	Modification
12.0(9)S	This feature was introduced.
12.0(17)ST	This feature was integrated into Cisco IOS Release 12.0(17)ST.
12.2(13)T	This feature was integrated into Cisco IOS Release 12.2(13)T.

This document describes the BGP Policy Accounting feature in Cisco IOS Release 12.2(13)T. 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 8](#)

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.




Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

© 2004–2008 Cisco Systems, Inc. All rights reserved.

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

To locate BGP configuration information, use the [BGP Feature Roadmap](#) module chapter of the *Cisco IOS IP Routing Protocols Configuration Guide* and to locate BGP command information, use the [Cisco IOS IP Routing Protocols Command Reference](#).

Additional Cisco Express Forwarding (CEF) and distributed CEF (dCEF) command and configuration information is documented in the “[Cisco Express Forwarding Overview](#)” module of the *Cisco IOS Switching Services Configuration Guide* and in the [Cisco IOS Switching Services Command Reference](#).

Related Documents

- [Cisco IOS IP Routing Protocols Command Reference](#)
- [Cisco IOS Switching Services Command Reference](#)

Supported Platforms

The BGP Policy Accounting feature is supported by the following platforms that support Cisco IOS Release 12.2(13)T:

- Cisco 1400 series
- Cisco 1600 series
- Cisco 1700 series
- Cisco 2600 series
- Cisco 3600 series
- Cisco 7100 series
- Cisco 7200 series
- Cisco 7500 series
- Cisco AS5300
- Cisco AS5350
- Cisco AS5400
- Cisco AS5800
- Cisco AS5850
- Cisco ICS7750
- Cisco IGX 8400 URM
- Cisco MC3810
- Cisco MGX 8850
- Cisco uBR7200 series

Finding Support Information for Platforms and Cisco IOS and Catalyst OS Software Images

Use Cisco Feature Navigator to find information about platform support and Cisco IOS and Catalyst OS software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.

Supported Standards, MIBs, and RFCs

Standards

No new or modified standards are supported by this feature.

MIBs

- CISCO-BGP-POLICY-ACCOUNTING-MIB

**Note**

The CISCO-BGP-POLICY-ACCOUNTING-MIB is only available in the Cisco IOS Release 12.0(9)S, 12.0(17)ST, and later releases. This MIB is not available on any mainline and T-train release.

To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL:

<http://tools.cisco.com/ITDIT/MIBS/servlet/index>

If Cisco MIB Locator does not support the MIB information that you need, you can also obtain a list of supported MIBs and download MIBs from the Cisco MIBs page at the following URL:

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

To access Cisco MIB Locator, 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 found at this URL:

<http://www.cisco.com/register>

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.
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.

	Command	Purpose
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.

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

The following commands are introduced or modified in the feature or features documented in this module. For information about these commands, see the *Cisco IOS IP Routing Protocols Command Reference* at http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html. For information about all Cisco IOS commands, go to the Command Lookup Tool at <http://tools.cisco.com/Support/CLILookup> or to the *Cisco IOS Master Commands List*.

- **bgp-policy**
- **set traffic-index**
- **show cef interface policy-statistics**
- **show ip bgp**
- **show ip cef**

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.

CCDE, CCSI, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Stackpower, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX,

PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo 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. (0903R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

© 2004–2008 Cisco Systems, Inc. All rights reserved.

