



BGP Policy Accounting Output Interface Accounting

Border Gateway Protocol (BGP) policy accounting (PA) measures and classifies IP traffic that is sent to, or received from, different peers. Policy accounting was previously available on an input interface only. The BGP Policy Accounting Output Interface Accounting feature introduces several extensions to enable BGP PA on an output interface and to include accounting based on a source address for both input and output traffic on an interface. Counters based on parameters such as community list, autonomous system number, or autonomous system path are assigned to identify the IP traffic.

Feature History for BGP PA Output Interface Accounting

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.
12.0(22)S	Output interface accounting was added, and the bucket size was increased.
12.3(4)T	This feature was integrated into Cisco IOS Release 12.3(4)T.
12.2(22)S	This feature was integrated into Cisco IOS Release 12.2(22)S.

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.

Contents

- [Prerequisites for BGP PA Output Interface Accounting, page 2](#)
- [Restrictions for BGP PA Output Interface Accounting, page 2](#)
- [Information About BGP PA Output Interface Accounting, page 2](#)
- [How to Configure BGP PA Output Interface Accounting, page 3](#)



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

© 2005–2008 Cisco Systems, Inc. All rights reserved.

- [Configuration Examples for BGP PA Output Interface Accounting, page 10](#)
- [Where to Go Next, page 11](#)
- [Additional References, page 11](#)
- [Command Reference, page 12](#)
- [Glossary, page 13](#)

Prerequisites for BGP PA Output Interface Accounting

Before using the BGP Policy Accounting Output Interface Accounting feature, you must enable BGP and Cisco Express Forwarding (CEF) or distributed CEF (dCEF) on the router.

Restrictions for BGP PA Output Interface Accounting

The CISCO-BGP-POLICY-ACCOUNTING-MIB is only available in the Cisco IOS Release 12.0(9)S, 12.0(17)ST, 12.2(22)S, and later releases. This MIB is not available on any mainline and T-train release.

Information About BGP PA Output Interface Accounting

To configure BGP PA output interface accounting, you should understand the following concepts:

- [BGP PA Output Interface Accounting, page 2](#)
- [Benefits of BGP PA Output Interface Accounting, page 3](#)


BGP PA Output Interface Accounting

Policy accounting using BGP measures and classifies IP traffic that is sent to, or received from, different peers. Originally, BGP PA was available on an input interface only. BGP PA output interface accounting introduces several extensions to enable BGP PA on an output interface and to include accounting based on a source address for both input and output traffic on an interface. 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 or output interface. A Cisco IOS policy-based classifier maps the traffic into one of eight possible buckets that represent different traffic classes.

Using BGP PA, you can account for traffic according to its origin or the route it traverses. Service providers (SPs) can identify and account for all traffic by customer and can bill accordingly. In [Figure 1](#), BGP PA 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 of BGP PA Output Interface Accounting

Accounting 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. Policy accounting can also be based on the source address. Service providers can account for traffic and apply billing according to the origin of the traffic or the route that 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.

How to Configure BGP PA Output Interface Accounting

This section contains the following tasks:

- [Specifying the Match Criteria for BGP PA, page 4](#) (required)
- [Classifying the IP Traffic and Enabling BGP PA, page 5](#) (required)
- [Verifying BGP Policy Accounting, page 7](#) (optional)

Specifying the Match Criteria for BGP PA

The first task in configuring BGP PA 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. Perform this task to specify the BGP attribute to use for BGP PA and to create the match criteria in a route map.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip community-list** *{standard-list-number | expanded-list-number [regular-expression] | {standard | expanded} community-list-name}* **{permit | deny}** *{community-number | regular-expression}*
4. **route-map** *map-name* **[permit | deny]** *[sequence-number]*
5. **match community-list** *community-list-number* **[exact]**
6. **set traffic-index** *bucket-number*
7. **exit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	ip community-list <i>{standard-list-number expanded-list-number [regular-expression] {standard expanded} community-list-name}</i> {permit deny} <i>{community-number regular-expression}</i> Example: Router(config)# ip community-list 30 permit 100:190	Creates a community list for BGP and controls access to it. <ul style="list-style-type: none"> • Repeat this step for each community to be specified.

	Command or Action	Purpose
Step 4	<pre>route-map map-name [permit deny] [sequence-number]</pre> <p>Example: Router(config)# route-map set_bucket permit 10</p>	<p>Enters route-map configuration mode and defines the conditions for policy routing.</p> <ul style="list-style-type: none"> 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 that a new route map is to have in the list of route maps already configured with the same name.
Step 5	<pre>match community-list community-list-number [exact]</pre> <p>Example: Router(config-route-map)# match community-list 30</p>	<p>Matches a BGP community.</p>
Step 6	<pre>set traffic-index bucket-number</pre> <p>Example: Router(config-route-map)# set traffic-index 2</p>	<p>Indicates where to output packets that pass a match clause of a route map for BGP policy accounting.</p>
Step 7	<pre>exit</pre> <p>Example: Router(config-route-map)# exit</p>	<p>Exits route-map configuration mode and returns to global configuration mode.</p>

Classifying the IP Traffic and Enabling BGP PA

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 that it adds to the routing table according to the match criteria. When the **bgp-policy accounting** command is configured on an interface, BGP policy accounting is enabled.

Perform this task to classify the IP traffic and enable BGP policy accounting.

SUMMARY STEPS

- enable**
- configure terminal**
- router bgp** *as-number*
- table-map** *route-map-name*
- network** *network-number* [**mask** *network-mask*]
- neighbor** *ip-address* **remote-as** *as-number*
- exit**
- interface** *type number*

9. `ip address ip-address mask`
10. `bgp-policy accounting [input | output] [source]`
11. `exit`

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>enable</code> Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> Enter your password if prompted.
Step 2	<code>configure terminal</code> Example: Router# configure terminal	Enters global configuration mode.
Step 3	<code>router bgp as-number</code> Example: Router(config)# router bgp 65000	Configures a BGP routing process and enters router configuration mode for the specified routing process. <ul style="list-style-type: none"> The <i>as-number</i> argument identifies a BGP autonomous system number.
Step 4	<code>table-map route-map-name</code> Example: Router(config-router)# table-map set_bucket	Classifies BGP prefixes entered in the routing table.
Step 5	<code>network network-number [mask network-mask]</code> Example: Router(config-router)# network 10.15.1.0 mask 255.255.255.0	Specifies a network to be advertised by the BGP routing process.
Step 6	<code>neighbor ip-address remote-as as-number</code> Example: Router(config-router)# neighbor 10.14.1.1 remote-as 65100	Specifies a BGP peer by adding an entry to the BGP routing table.
Step 7	<code>exit</code> Example: Router(config-router)# exit	Exits router configuration mode and returns to global configuration mode.
Step 8	<code>interface type number</code> Example: Router(config)# interface POS 7/0	Specifies the interface type and number and enters interface configuration mode. <ul style="list-style-type: none"> The <i>type</i> argument identifies the type of interface. The <i>number</i> argument identifies the slot and port numbers of the interface. The space between the interface type and number is optional.

	Command or Action	Purpose
Step 9	<p>ip address <i>ip-address mask</i></p> <p>Example: Router(config-if)# ip-address 10.15.1.2 255.255.255.0</p>	Configures the interface with an IP address.
Step 10	<p>bgp-policy accounting [input output] [source]</p> <p>Example: Router(config-if)# bgp-policy accounting input source</p>	<p>Enables BGP policy accounting for the interface.</p> <ul style="list-style-type: none"> Use the optional input or output keyword to account for traffic either entering or leaving the router. By default, BGP policy accounting is based on traffic entering the router. Use the optional source keyword to account for traffic based on source address.
Step 11	<p>exit</p> <p>Example: Router(config-if)# exit</p>	Exits interface configuration mode and returns to global configuration mode.

Verifying BGP Policy Accounting

Perform this task to verify that BGP policy accounting is operating.

SUMMARY STEPS

- show ip cef** [*network [mask]*] [**detail**]
- show ip bgp** [*network*] [*network-mask*] [**longer-prefixes**]
- show cef interface** [*type number*] **policy-statistics** [**input** | **output**]
- show cef interface** [*type number*] [**statistics**] [**detail**]

DETAILED STEPS

Step 1 **show ip cef** [*network [mask]*] [**detail**]

Enter the **show ip cef** 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 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 **show ip bgp** [*network*] [*network-mask*] [**longer-prefixes**]

Enter the **show ip bgp** 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 show cef interface [type number] policy-statistics [input | output]

Enter the **show cef interface policy-statistics** 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:

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

```
FastEthernet1/0/0 is up (if_number 6)
  Corresponding hwidb fast_if_number 6
  Corresponding hwidb firstsw->if_number 6
  BGP based Policy accounting on input is enabled
```

Index	Packets	Bytes
1	9999	999900
2	0	0
3	0	0
4	0	0
5	0	0
6	0	0
7	0	0
8	0	0
9	0	0
10	0	0
11	0	0
12	0	0
13	0	0
14	0	0
15	0	0
16	0	0
17	0	0
18	0	0
19	0	0
20	0	0
21	0	0
22	0	0
23	0	0
24	0	0
25	0	0
26	0	0
27	0	0
28	0	0
29	0	0
30	0	0
31	0	0
32	0	0
33	0	0
34	1234	123400
35	0	0
36	0	0
37	0	0

38	0	0
39	0	0
40	0	0
41	0	0
42	0	0
43	0	0
44	0	0
45	1000	100000
46	0	0
47	0	0
48	0	0
49	0	0
50	0	0
51	0	0
52	0	0
53	0	0
54	5123	1198782
55	0	0
56	0	0
57	0	0
58	0	0
59	0	0
60	0	0
61	0	0
62	0	0
63	0	0
64	0	0

Step 4 **show cef interface** [*type number*] [*statistics*] [*detail*]

Enter the **show cef interface EXEC** command to display the state of BGP policy accounting on a specified interface.

In this example, the output shows that BGP policy accounting has been configured to be based on input traffic at Fast Ethernet interface 1/0/0:

```
Router# show cef interface Fast Ethernet 1/0/0

FastEthernet1/0/0 is up (if_number 6)
  Corresponding hwidb fast_if_number 6
  Corresponding hwidb firstsw->if_number 6
  Internet address is 10.1.1.1/24
  ICMP redirects are always sent
  Per packet load-sharing is disabled
  IP unicast RPF check is disabled
  Inbound access list is not set
  Outbound access list is not set
  IP policy routing is disabled
  BGP based policy accounting on input is enabled
  BGP based policy accounting on output is disabled
  Hardware idb is FastEthernet1/0/0 (6)
  Software idb is FastEthernet1/0/0 (6)
  Fast switching type 1, interface type 18
  IP Distributed CEF switching enabled
  IP Feature Fast switching turbo vector
  IP Feature CEF switching turbo vector
  Input fast flags 0x100, Output fast flags 0x0, Flags 0x0
  ifindex 7(7)
  Slot 1 Slot unit 0 VC -1
  Transmit limit accumulator 0xE8001A82 (0xE8001A82)
  IP MTU 1500
```

Configuration Examples for BGP PA Output Interface Accounting

This section contains the following configuration examples:

- [Specifying the Match Criteria for BGP Policy Accounting: Example, page 10](#)
- [Classifying the IP Traffic and Enabling BGP Policy Accounting: Example, page 10](#)

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-list 30
  set traffic-index 2
!
route-map set_bucket permit 20
 match community-list 40
  set traffic-index 3
!
route-map set_bucket permit 30
 match community-list 50
  set traffic-index 4
!
route-map set_bucket permit 40
 match community-list 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. The policy accounting criteria is based on the source address of the input traffic, 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
 bgp-policy accounting input source
 no keepalive
 crc 32
 clock source internal
```

Where to Go Next

Additional BGP, CEF, and dCEF command and configuration information is available from the appropriate Cisco IOS command reference or configuration guide documents. For more details, see the “[Related Documents](#)” section.

Additional References

The following sections provide references related to BGP policy accounting.

Related Documents

Related Topic	Document Title
BGP commands: complete command syntax, command mode, defaults, usage guidelines, and examples	Cisco IOS IP Routing Protocols Command Reference
Locate BGP configuration features	BGP Feature Roadmap
Switching commands: complete command syntax, command mode, defaults, usage guidelines, and examples	Cisco IOS IP Switching Command Reference
CEF and dCEF configuration information	“Cisco Express Forwarding Overview” module

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIBs	MIBs Link
CISCO-BGP-POLICY-ACCOUNTING-MIB Note This MIB is available only in Cisco IOS Release 12.0(9)S, 12.0(17)ST, 12.2(22)S, 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://www.cisco.com/go/mibs

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature.	—

Technical Assistance

Description	Link
<p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p>	http://www.cisco.com/techsupport

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**
- **show cef interface policy-statistics**

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.

© 2005–2008 Cisco Systems, Inc. All rights reserved.

