BGP Policy Accounting and BGP Policy Accounting Output Interface Accounting Features

Document ID: 13760

Introduction

The Border Gateway Protocol (BGP) policy accounting (PA) feature allows you to account for IP traffic differentially by assigning counters based on community list, AS number, and/or AS_PATH on a per input interface basis.

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.

Prerequisites

Requirements

Before you use the BGP PA feature, enable Cisco Express Forwarding (CEF) or distributed Cisco Express Forwarding (dCEF) on the router.

Components Used

The BGP PA feature is first supported by the following platforms that support Cisco IOS Release 12.0(9)S.

- Cisco 7200, 7500, and 12000 Series Routers

The number of platforms that support this feature in Cisco IOS Release 12.2(13)T is much larger. The platforms include:

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

BGP PA Output Interface Accounting was added in 12.0(22)S and first introduced as new feature in 12.3(4)T. There are many Cisco platfoms that support this feature.

Note: To get updated information regarding platform support for this feature, access Cisco Feature Navigator

II (registered customers only) .

The information presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If you are working in a live network, ensure that you understand the potential impact of any command before using it.

Background Theory

This feature was first available on the Cisco IOS® Software release 12.0(9)S ED. For the policy accounting feature to work, you must enable BGP and CEF/dCEF on the router.

Using BGP policy accounting, you can account for traffic (and apply billing) according to the route it traverses. For example, you can account for traffic that's routed domestic, international, terrestrial, or satellite. In this way, you can identify and account for all traffic on a per−customer basis.

This feature takes advantage of the BGP **table−map** command, which classifies the prefixes it puts into the routing table according to community list, AS number, AS_PATH, and so on. Based on these match criteria, the BGP accounting policy sets a bucket number (currently 1 to 64) of an accounting table associated with each interface. Each bucket represents a traffic classification, which allows IP traffic to be accounted for differentially by community list, AS number, or AS_PATH per input interface.

For more information, refer to BGP Policy Accounting.

Note: BGP PA measures and classifies IP traffic that is sent to, or received from, different peers. PA 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. The output interface accounting was added in Cisco IOS release 12.0(22)S.

Conventions

For more information on document conventions, refer to Cisco Technical Tips Conventions.

Configuring BGP Policy Accounting

1. Specify communities in community lists (or define AS_PATH lists) that classify traffic for accounting.

```
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
  ip community−list 70 permit 100:201
  !
```

2. Define a **route−map** to match community lists and set appropriate bucket numbers.

```
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
!
route-map set_bucket permit 50
match community 70
set traffic-index 6
```

3. Use the **table-map** command under BGP to modify the bucket number when the IP routing table is updated with routes learned from BGP.

```
router bgp 110
  table-map set_bucket
  network 15.1.1.0 mask 255.255.255.0
  neighbor 14.1.1.1 remote-as 100
  !
  ip classless
  ip bgp-community new-format
```

4. Enable the policy accounting feature on the input interface connected to the customer.

```
interface POS7/0
  ip address 15.1.1.2 255.255.255.0
  no ip directed-broadcast
  bgp-policy accounting
  no keepalive
  crc 32
  clock source internal
```

Configuring BGP Policy Accounting Output Interface Accounting

The configuration of BGP PA Output Interface Accounting is very similar to BGP PA. The first three step described in the previous section are exactly the same. The only change is in the **bgp-policy accounting** command that is used to enable the PA feature on the interface. In the example below BGP PA is enabled on POS interface 7/0. The PA criteria is based on the source address of the output traffic

```
interface POS7/0
 ip address 10.15.1.2 255.255.255.0
 bgp-policy accounting output source
 no keepalive
 crc 32
 clock source internal
```

show Commands that Monitor Policy Accounting

To inspect which prefix is assigned to which bucket and which community (or communities), use the **show ip cef** and **show ip bgp** commands.

```
Router# show ip cef 196.240.5.0 detail
196.240.5.0/24, version 21, cached adjacency to POS7/2
0 packets, 0 bytes, traffic_index 4
  via 14.1.1.1, 0 dependencies, recursive
   next hop 14.1.1.1, POS7/2 via 14.1.1.0/30
   valid cached adjacency
```

```
Router# show ip bgp 196.240.5.0
BGP routing table entry for 196.240.5.0/24, version 2
Paths: (1 available, best #1)
  Not advertised to any peer
  100
    14.1.1.1 from 14.1.1.1 (32.32.32.32)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: 100:197
```

To look at per−interface traffic statistics, use the **show cef interface policy−statistics** command.

```
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
```

Related Information

- **BGP Policy Accounting**
- **BGP Policy Accounting Output Interface Accounting**
- **XC: Part 1: Cisco IOS Switching Paths**
- **IPC: Part 2: IP Routing Protocols**
- **BGP Support Page**
- **Technical Support − Cisco Systems**

Updated: Dec 29, 2005 Document ID: 13760