show ip as-path-access-list

To display the contents of all current autonomous system (AS) path access lists, use the **show ip as-path-access-list** command in user EXEC or privileged EXEC mode.

show ip as-path-access-list [*number*]

Syntax Description

number	(Optional) Specifies the AS path access list number. The range is from 1 to 500.

Command Default

If the *number* argument is not specified, command output is displayed for all AS path access lists.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
11.3	This command was introduced.
12.2(33)SRC	This command was integrated into Cisco IOS Release 12.2(33)SRC.
12.2(33)SXI	This command was integrated into Cisco IOS Release 12.2(33)SXI.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB.
Cisco IOS XE Release 2.1	This command was integrated into Cisco IOS XE Release 2.1.

Examples

The following is sample output from the **show ip as-path-access-list** command:

```
Router# show ip as-path-access-list

AS path access list 34
    deny RTR$
AS path access list 100
    permit 100$
```

Table 27 describes the fields shown in the display.

Table 27 show ip as-path-access-list Field Descriptions

Field	Description
AS path access list	Indicates the AS path access list number.
deny	Indicates the number of packets that are rejected since the regular expression failed to match the representation of the AS path of the route as an ASCII string.
permit	Indicates the number of packets that are forwarded since the regular expression matched the representation of the AS path of the route as an ASCII string.

Related Commands

Command	Description
ip as-path access-list	Configures an autonomous system path filter using a regular expression.

show ip bgp

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

show ip bgp [*ip-address* [*mask* [**longer-prefixes** [**injected**] | **shorter-prefixes** [*length*] | **bestpath** | **multipaths** | **subnets**] | **bestpath** | **multipaths**] | **all** | **oer-paths** | **prefix-list** *name* | **pending-prefixes** | **route-map** *name*]

Syntax Description

ip-address	(Optional) IP address entered to filter the output to display only a particular host or network in the BGP routing table.
mask	(Optional) Mask to filter or match hosts that are part of the specified network.
longer-prefixes	(Optional) Displays the specified route and all more specific routes.
injected	(Optional) Displays more specific prefixes injected into the BGP routing table.
shorter-prefix	(Optional) Displays the specified route and all less specific routes.
length	(Optional) The prefix length. The value for this argument is a number from 0 to 32.
bestpath	(Optional) Displays the bestpath for this prefix
multipaths	(Optional) Displays multipaths for this prefix.
subnets	(Optional) Displays the subnet routes for the specified prefix.
all	(Optional) Displays all address family information in the BGP routing table.
oer-paths	(Optional) Displays Optimized Edge Routing (OER) controlled prefixes in the BGP routing table.
prefix-list *name*	(Optional) Filters the output based on the specified prefix list.
pending-prefixes	(Optional) Displays prefixes that are pending deletion from the BGP routing table.
route-map *name*	(Optional) Filters the output based on the specified route map.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
10.0	This command was introduced.
12.0	This command was modified. The display of prefix advertisement statistics was added.
12.0(6)T	This command was modified. The display of a message indicating support for route refresh capability was added.
12.0(14)ST	This command was modified. The **prefix-list**, **route-map**, and **shorter-prefixes** keywords were added.
12.2(2)T	This command was modified. The output was modified to display multipaths and a best path to the specified network.

Release	Modification
12.0(21)ST	The output was modified to show the number of Multiprotocol Label Switching (MPLS) labels that arrive at and depart from the prefix.
12.0(22)S	This command was modified. A new status code indicating stale routes was added to support BGP graceful restart.
12.2(14)S	This command was modified. A message indicating support for BGP policy accounting was added and this command was integrated into Cisco IOS Release 12.2(14)S.
12.2(14)SX	This command was integrated into Cisco IOS Release 12.2(14)SX.
12.2(15)T	This command was modified. A new status code indicating stale routes was added to support BGP graceful restart.
12.3(2)T	This command was modified. The **all** keyword was added.
12.2(17b)SXA	This command was integrated into Cisco IOS Release 12.2(17b)SXA.
12.3(8)T	This command was modified. The **oer-paths** keyword was added.
12.4(15)T	This command was modified. The **pending-prefixes**, **bestpath**, **multipaths**, and **subnets** keywords were added
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(31)SB2	This command was integrated into Cisco IOS Release 12.2(31)SB2.
12.0(32)S12	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
12.0(32)SY8	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.4(24)T	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
Cisco IOS XE Release 2.3	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
12.2(33)SXI1	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.0(33)S3	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format is now asplain.
Cisco IOS XE Release 2.4	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format is now asplain.
12.2(33)SRE	This command was modified. The command output was modified to show the backup path and the best external path information. Support for the best external route and backup path was added. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)XNE	This command was integrated into Cisco IOS Release 12.2(33)XNE.
15.0(1)S	This command was integrated into Cisco IOS Release 15.0(1)S.

Usage Guidelines

The **show ip bgp** command is used to display the contents of the BGP routing table. The output can be filtered to display entries for a specific prefix, prefix length, and prefixes injected through a prefix list, route map, or conditional advertisement.

In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, the Cisco implementation of 4-byte autonomous system numbers uses asplain—65538 for example—as the default regular expression match and output display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain format and the asdot format as described in RFC 5396. To change the default regular expression match and output display of 4-byte autonomous system numbers to asdot format, use the **bgp asnotation dot** command followed by the **clear ip bgp *** command to perform a hard reset of all current BGP sessions.

In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, the Cisco implementation of 4-byte autonomous system numbers uses asdot—1.2 for example—as the only configuration format, regular expression match, and output display, with no asplain support.

oer-paths Keyword

In Cisco IOS Release 12.3(8)T, and later releases, BGP prefixes that are monitored and controlled by OER are displayed by entering the **show ip bgp** command with the **oer-paths** keyword.

Examples

- **show ip bgp: Example, page 454**
- **show ip bgp (4-Byte Autonomous System Numbers): Example, page 456**
- **show ip bgp ip-address: Example, page 456**
- **show ip bgp all: Example, page 457**
- **show ip bgp longer-prefixes: Example, page 459**
- **show ip bgp shorter-prefixes: Example, page 459**
- **show ip bgp prefix-list: Example, page 459**
- **show ip bgp route-map: Example, page 460**

show ip bgp: Example

The following sample output shows the BGP routing table:

```
Router# show ip bgp

BGP table version is 22, local router ID is 10.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.1.1/32      0.0.0.0                  0         32768 i
*>i10.2.2.2/32      172.16.1.2               0    100      0 i
*bi10.9.9.9/32      192.168.3.2              0    100      0 10 10 i
*>                  192.168.1.2                            0 10 10 i
* i172.16.1.0/24    172.16.1.2               0    100      0 i
*>                  0.0.0.0                  0         32768 i
*> 192.168.1.0      0.0.0.0                  0         32768 i
*>i192.168.3.0      172.16.1.2               0    100      0 i
*bi192.168.9.0      192.168.3.2              0    100      0 10 10 i
*>                  192.168.1.2                            0 10 10 i
*bi192.168.13.0     192.168.3.2              0    100      0 10 10 i
*>                  192.168.1.2                            0 10 10 i
```

Table 28 describes the significant fields shown in the display.

Table 28 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.
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 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. • r—The table entry is a RIB-failure. • S—The table entry is stale. • m—The table entry has multipath to use for that network. • b—The table entry has backup path to use for that network. • x—The table entry has best external route to use for the network.
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 an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. • e—Entry originated from an 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.
(stale)	Indicates that the following path for the specified autonomous system is marked as “stale” during a graceful restart process.

show ip bgp (4-Byte Autonomous System Numbers): Example

The following sample output shows the BGP routing table with 4-byte autonomous system numbers, 65536 and 65550, shown under the Path field. This example requires Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, or a later release.

```
RouterB# show ip bgp

BGP table version is 4, local router ID is 172.17.1.99
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.1.0/24      192.168.1.2              0             0 65536  i
*> 10.2.2.0/24      192.168.3.2              0             0 65550  i
*> 172.17.1.0/24    0.0.0.0                  0         32768 i
```

show ip bgp ip-address: Example

The following sample output displays information about the 192.168.1.0 entry in the BGP routing table:

```
Router# show ip bgp 192.168.1.0

BGP routing table entry for 192.168.1.0/24, version 22
Paths: (2 available, best #2, table default)
  Additional-path
  Advertised to update-groups:
     3
  10 10
    192.168.3.2 from 172.16.1.2 (10.2.2.2)
      Origin IGP, metric 0, localpref 100, valid, internal, backup/repair
  10 10
    192.168.1.2 from 192.168.1.2 (10.3.3.3)
      Origin IGP, localpref 100, valid, external, best , recursive-via-connected
```

The following sample output displays information about the 10.3.3.3 255.255.255.255 entry in the BGP routing table:

```
Router# show ip bgp 10.3.3.3 255.255.255.255

BGP routing table entry for 10.3.3.3/32, version 35
Paths: (3 available, best #2, table default)
Multipath: eBGP
Flag: 0x860
  Advertised to update-groups:
     1
  200
    10.71.8.165 from 10.71.8.165 (192.168.0.102)
      Origin incomplete, localpref 100, valid, external, backup/repair
      Only allowed to recurse through connected route
  200
    10.71.11.165 from 10.71.11.165 (192.168.0.102)
      Origin incomplete, localpref 100, weight 100, valid, external, best
      Only allowed to recurse through connected route
  200
    10.71.10.165 from 10.71.10.165 (192.168.0.104)
      Origin incomplete, localpref 100, valid, external,
      Only allowed to recurse through connected route
```

Table 29 describes the significant fields shown in the display.

Table 29 show ip bgp Field Descriptions

Field	Description
BGP routing table entry for	IP address or network number of the routing table entry.
version	Internal version number of the table. This number is incremented whenever the table changes.
Paths	The number of available paths, and the number of installed best paths. This line displays "Default-IP-Routing-Table" when the best path is installed in the IP routing table.
Multipath	This field is displayed when multipath loadsharing is enabled. This field will indicate if the multipaths are iBGP or eBGP.
Advertised to update-groups	The number of each update group for which advertisements are processed.
Origin	Origin of the entry. The origin can be IGP, EGP, or incomplete. This line displays the configured metric (0 if no metric is configured), the local preference value (100 is default), and the status and type of route (internal, external, multipath, best).
Extended Community	This field is displayed if the route carries an extended community attribute. The attribute code is displayed on this line. Information about the extended community is displayed on a subsequent line.

show ip bgp all: Example

The following is sample output from the **show ip bgp** command entered with the **all** keyword. Information about all configured address families is displayed.

```
Router# show ip bgp all

For address family: IPv4 Unicast   *****
BGP table version is 27, local router ID is 10.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.1.0/24      0.0.0.0                  0         32768 ?
*> 10.13.13.0/24    0.0.0.0                  0         32768 ?
*> 10.15.15.0/24    0.0.0.0                  0         32768 ?
*>i10.18.18.0/24    172.16.14.105         1388  91351      0 100 e
*>i10.100.0.0/16    172.16.14.107          262    272      0 1 2 3 i
*>i10.100.0.0/16    172.16.14.105         1388  91351      0 100 e
*>i10.101.0.0/16    172.16.14.105         1388  91351      0 100 e
*>i10.103.0.0/16    172.16.14.101         1388    173    173 100 e
*>i10.104.0.0/16    172.16.14.101         1388    173    173 100 e
*>i10.100.0.0/16    172.16.14.106         2219  20889      0 53285 33299 51178 47751 e
*>i10.101.0.0/16    172.16.14.106         2219  20889      0 53285 33299 51178 47751 e
*  10.100.0.0/16    172.16.14.109         2309             0 200 300 e
*>                  172.16.14.108         1388             0 100 e
*  10.101.0.0/16    172.16.14.109         2309             0 200 300 e
*>                  172.16.14.108         1388             0 100 e
*> 10.102.0.0/16    172.16.14.108         1388             0 100 e
*> 172.16.14.0/24   0.0.0.0                  0         32768 ?
*> 192.168.5.0      0.0.0.0                  0         32768 ?
*> 10.80.0.0/16     172.16.14.108         1388             0 50 e
*> 10.80.0.0/16     172.16.14.108         1388             0 50 e
```

```
For address family: VPNv4 Unicast   *****
BGP table version is 21, local router ID is 10.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf vpn1)
*> 10.1.1.0/24      192.168.4.3           1622             0 100 53285 33299 51178
{27016,57039,16690} e
*> 10.1.2.0/24      192.168.4.3           1622             0 100 53285 33299 51178
{27016,57039,16690} e
*> 10.1.3.0/24      192.168.4.3           1622             0 100 53285 33299 51178
{27016,57039,16690} e
*> 10.1.4.0/24      192.168.4.3           1622             0 100 53285 33299 51178
{27016,57039,16690} e
*> 10.1.5.0/24      192.168.4.3           1622             0 100 53285 33299 51178
{27016,57039,16690} e
*>i172.17.1.0/24    10.3.3.3                 10     30      0 53285 33299 51178 47751 ?
*>i172.17.2.0/24    10.3.3.3                 10     30      0 53285 33299 51178 47751 ?
*>i172.17.3.0/24    10.3.3.3                 10     30      0 53285 33299 51178 47751 ?
*>i172.17.4.0/24    10.3.3.3                 10     30      0 53285 33299 51178 47751 ?
*>i172.17.5.0/24    10.3.3.3                 10     30      0 53285 33299 51178 47751 ?

For address family: IPv4 Multicast   *****
BGP table version is 11, local router ID is 10.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.40.40.0/26    172.16.14.110         2219             0 21 22 {51178,47751,27016} e
*                   10.1.1.1              1622             0 15 20 1 {2} e
*> 10.40.40.64/26   172.16.14.110         2219             0 21 22 {51178,47751,27016} e
*                   10.1.1.1              1622             0 15 20 1 {2} e
*> 10.40.40.128/26  172.16.14.110         2219             0 21 22 {51178,47751,27016} e
*                   10.1.1.1              2563             0 15 20 1 {2} e
*> 10.40.40.192/26  10.1.1.1              2563             0 15 20 1 {2} e
*> 10.40.41.0/26    10.1.1.1              1209             0 15 20 1 {2} e
*>i10.102.0.0/16    10.1.1.1               300    500       0 5 4 {101,102} e
*>i10.103.0.0/16    10.1.1.1               300    500       0 5 4 {101,102} e

For address family: NSAP Unicast *****
BGP table version is 1, local router ID is 10.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
* i45.0000.0002.0001.000c.00
                    49.0001.0000.0000.0a00
                                                100      0 ?
* i46.0001.0000.0000.0000.0a00
                    49.0001.0000.0000.0a00
                                                100      0 ?
* i47.0001.0000.0000.000b.00
                    49.0001.0000.0000.0a00
                                                100      0 ?
* i47.0001.0000.0000.000e.00
                    49.0001.0000.0000.0a00
```

show ip bgp longer-prefixes: Example

The following is sample output from the **show ip bgp** command entered with the **longer-prefixes** keyword:

```
Router# show ip bgp 10.92.0.0 255.255.0.0 longer-prefixes

BGP table version is 1738, local router ID is 192.168.72.24
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 10.92.0.0         10.92.72.30         8896         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.1.0         10.92.72.30         8796         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.11.0        10.92.72.30        42482         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.14.0        10.92.72.30         8796         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.15.0        10.92.72.30         8696         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.16.0        10.92.72.30         1400         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.17.0        10.92.72.30         1400         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.18.0        10.92.72.30         8876         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.19.0        10.92.72.30         8876         32768 ?
*                    10.92.72.30                          0 109 108 ?
```

show ip bgp shorter-prefixes: Example

The following is sample output from the **show ip bgp** command entered with the **shorter-prefixes** keyword. An 8-bit prefix length is specified.

```
Router# show ip bgp 172.16.0.0/16 shorter-prefixes 8

*> 172.16.0.0         10.0.0.2                               0 ?
*                     10.0.0.2                 0             0 200 ?
```

show ip bgp prefix-list: Example

The following is sample output from the **show ip bgp** command entered with the **prefix-list** keyword:

```
Router# show ip bgp prefix-list ROUTE

BGP table version is 39, local router ID is 10.0.0.1
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
*> 192.168.1.0      10.0.0.2                               0 ?
*                   10.0.0.2                 0             0 200 ?
```

show ip bgp route-map: Example

The following is sample output from the **show ip bgp** command entered with the **route-map** keyword:

```
Router# show ip bgp route-map LEARNED_PATH

BGP table version is 40, local router ID is 10.0.0.1
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
*> 192.168.1.0      10.0.0.2                             0 ?
*                   10.0.0.2                  0          0 200 ?
```

Related Commands

Command	Description
bgp asnotation dot	Changes the default display and the regular expression match format of BGP 4-byte autonomous system numbers from asplain (decimal values) to dot notation.
ip prefix-list	Creates a prefix list or adds a prefix-list entry.
route-map	Defines the conditions for redistributing routes from one routing protocol into another routing protocol.
router bgp	Configures the BGP routing process.

show ip bgp all dampening

To display BGP dampening information, use the **show ip bgp all dampening** command in user EXEC or privileged EXEC mode.

show ip bgp all dampening {**dampened-paths** | **flap-statistics** [**filter-list** *filter-list* | **quote-regexp** *regexp* | **regexp** *regexp*] | **parameters**}

Syntax Description

dampened-paths	Display routes suppressed due to dampening.
flap-statistics	Displays flap statistics of routes.
filter-list *filter-list*	(Optional) Used with the **flap-statistics** keyword, displays routes that conform to the specified filter list in the range 1-500.
quote-regexp *regexp*	(Optional) Used with the **flap-statistics** keyword, displays routes matching the AS path "regular expression".
regexp *regexp*	(Optional) Used with the **flap-statistics** keyword, displays routes matching the AS path regular expression.
parameters	Display details of configured dampening parameters.

Command Modes

User EXEC (>)
Privileged EXEC

Command History

Release	Modification
15.0(1)M	This command was introduced.

Usage Guidelines

Use this command to display BGP dampening information.

Examples

The following example show how to display the BGP dampening parameters.

```
Router# show ip bgp all dampening parameters
For address family: IPv4 Unicast

% dampening not enabled for base

For address family: VPNv4 Unicast

% dampening not enabled for base

For vrf: Cust_A
 dampening 15 750 2000 60 (DEFAULT)
  Half-life time      : 15 mins       Decay Time       : 2320 secs
  Max suppress penalty: 12000         Max suppress time: 60 mins
  Suppress penalty    :  2000         Reuse penalty    : 750


For vrf: Cust_B

 dampening 15 750 2000 60 (DEFAULT)
```

```
  Half-life time      : 15 mins       Decay Time       : 2320 secs
  Max suppress penalty: 12000         Max suppress time: 60 mins
  Suppress penalty    :  2000         Reuse penalty    : 750


For address family: IPv4 Multicast

% dampening not enabled for base
Router#
```

Table 30 describes the significant fields shown in the display.

Table 30 show ip bgp all dampening Field Descriptions

Field	Description
Half-life time	Time after which a penalty is decreased, in minutes. Once the interface has been assigned a penalty, the penalty is decreased by half after the half-life period. The process of reducing the penalty happens every 5 seconds. The range of the half-life is 1 to 45 minutes. The default is 1 minute.
Decay Time	Penalty value below which an unstable interface is unsuppressed, in seconds. The process of unsuppressing routers occurs at 10-second increments. The range of the reuse value is 1 to 20000 seconds. The default value is 750 seconds.
Max suppress penalty	Limit at which an interface is suppressed when its penalty exceeds that limit, in seconds. The default value is 2000 seconds.
Max suppress time	Maximum time that an interface can be suppressed, in minutes. This value effectively acts as a ceiling that the penalty value cannot exceed. The default value is four times the half-life period.

The following is sample output for the **show ip bgp all dampening dampened-paths** command. The output includes dampened paths for individual VRFs.

```
Router# show ip bgp all dampening dampened-paths

For address family: IPv4 Unicast

% dampening not enabled for base

For address family: VPNv4 Unicast

% dampening not enabled for base

For vrf: Cust_A

BGP table version is 42, local router ID is 144.124.23.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          From             Reuse    Path
Route Distinguisher: 1:100 (Cust_A)
*d 10.10.10.10/32   172.16.1.2       00:04:49 65001 ?
*d 20.20.20.20/32   172.16.1.2       00:04:59 65001 ?

For address family: IPv4 Multicast

% dampening not enabled for base
```

Related Commands

Command	Description
bgp dampening	Enables BGP route dampening or changes various BGP route dampening factors.
show dampening interface	Displays a summary of the dampening parameters and status.

show ip bgp cidr-only

To display routes with classless interdomain routing (CIDR), use the **show ip bgp cidr-only** command in EXEC mode.

show ip bgp cidr-only

Syntax Description

This command has no arguments or keywords.

Command Modes

EXEC

Command History

Release	Modification
10.0	This command was introduced.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Examples

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

```
Router# show ip bgp cidr-only

BGP table version is 220, local router ID is 172.16.73.131
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 192.168.0.0/8    172.16.72.24                         0 1878 ?
*> 172.16.0.0/16    172.16.72.30                         0 108 ?
```

Table 31 describes the significant fields shown in the display.

Table 31 show ip bgp cidr-only Field Descriptions

Field	Description
BGP table version is 220	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
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. *—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.

Table 31 show ip bgp cidr-only Field Descriptions (continued)

Field	Description
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 an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an 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	Internet address of the network the entry describes.
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 access server has some non-BGP route 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. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a **network** router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP.

show ip bgp community

To display routes that belong to specified BGP communities, use the **show ip bgp community** command in EXEC mode.

show ip bgp community *community-number* [**exact**]

Syntax Description

community-number	Valid value is a community number in the range from 1 to 4294967200, or AA:NN (autonomous system-community number/2-byte number), **internet**, **no-export**, **local-as**, or **no-advertise**.
exact	(Optional) Displays only routes that have the same specified communities.

Command Modes

EXEC

Command History

Release	Modification
10.3	This command was introduced.
12.0	The **local-as** community was added.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Examples

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

```
Router# show ip bgp community 111:12345 local-as

 BGP table version is 10, local router ID is 224.0.0.10
 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
 *> 172.16.2.2/32    10.43.222.2           0             0 222 ?
 *> 10.0.0.0         10.43.222.2           0             0 222 ?
 *> 10.43.0.0        10.43.222.2           0             0 222 ?
 *> 10.43.44.44/32   10.43.222.2           0             0 222 ?
 *  10.43.222.0/24   10.43.222.2           0             0 222 i
 *> 172.17.240.0/21  10.43.222.2           0             0 222 ?
 *> 192.168.212.0    10.43.222.2           0             0 222 i
 *> 172.31.1.0       10.43.222.2           0             0 222 ?
```

Table 32 describes the significant fields shown in the display.

Table 32 show ip bgp community 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.
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. *—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 an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an 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, this is the value of the interautonomous system metric. This field is frequently not used.
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.

show ip bgp community-list

To display routes that are permitted by the Border Gateway Protocol (BGP) community list, use the **show ip bgp community-list** command in user or privileged EXEC mode.

show ip bgp community-list {*community-list-number* | *community-list-name* [**exact-match**]}

Syntax Description

community-list-number	A standard or expanded community list number in the range from 1 to 500.
community-list-name	Community list name. The community list name can be standard or expanded.
exact-match	(Optional) Displays only routes that have an exact match.

Command Modes

User EXEC
Privileged EXEC

Command History

Release	Modification
10.3	This command was introduced.
12.0(10)S	Named community list support was added.
12.0(16)ST	Named community lists support was integrated into Cisco IOS Release 12.0(16)ST.
12.1(9)E	Named community lists support was integrated into Cisco IOS Release 12.1(9)E.
12.2(8)T	Named community lists support was integrated into Cisco IOS Release 12.2(8)T.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB to support the Cisco 10000 Series Routers.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

This command requires you to specify an argument when used. The **exact-match** keyword is optional.

Examples

The following is sample output of the **show ip bgp community-list** command in privileged EXEC mode:

```
Router# show ip bgp community-list 20

BGP table version is 716977, local router ID is 192.168.32.1
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
   Network          Next Hop          Metric LocPrf Weight Path
* i10.3.0.0         10.0.22.1             0    100      0 1800 1239 ?
*>i                 10.0.16.1             0    100      0 1800 1239 ?
* i10.6.0.0         10.0.22.1             0    100      0 1800 690 568 ?
*>i                 10.0.16.1             0    100      0 1800 690 568 ?
* i10.7.0.0         10.0.22.1             0    100      0 1800 701 35 ?
*>i                 10.0.16.1             0    100      0 1800 701 35 ?
*                   10.92.72.24                    0 1878 704 701 35 ?
* i10.8.0.0         10.0.22.1             0    100      0 1800 690 560 ?
*>i                 10.0.16.1             0    100      0 1800 690 560 ?
*                   10.92.72.24                    0 1878 704 701 560 ?
* i10.13.0.0        10.0.22.1             0    100      0 1800 690 200 ?
*>i                 10.0.16.1             0    100      0 1800 690 200 ?
*                   10.92.72.24                    0 1878 704 701 200 ?
* i10.15.0.0        10.0.22.1             0    100      0 1800 174 ?
*>i                 10.0.16.1             0    100      0 1800 174 ?
* i10.16.0.0        10.0.22.1             0    100      0 1800 701 i
*>i                 10.0.16.1             0    100      0 1800 701 i
*                   10.92.72.24                    0 1878 704 701 i
```

Table 33 describes the significant fields shown in the display.

Table 33 show ip bgp community-list 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.
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. *—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 an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an 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, this is the value of the interautonomous system metric. This field is frequently not used.
LocPrf	Local preference value as set with the **set local-preference** route-map configuration command. The default value is 100.

Table 33 show ip bgp community-list Field Descriptions (continued)

Field	Description
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.

show ip bgp dampened-paths

To display BGP dampened routes, use the **show ip bgp dampened-paths** command in EXEC mode.

show ip bgp dampened-paths

Syntax Description

This command has no arguments or keywords.

Command Modes

EXEC

Command History

Release	Modification
11.0	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

On the Cisco 10000 series router, use the **show ip bgp dampening dampened-paths** command to display BGP dampened routes.

Examples

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

```
Router# show ip bgp dampened-paths

BGP table version is 10, local router ID is 172.29.232.182
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          From             Reuse   Path
*d 10.0.0.0         172.16.232.177   00:18:4 100 ?
*d 10.2.0.0         172.16.232.177   00:28:5 100 ?
```

Table 34 describes the significant fields shown in the display.

Table 34 show ip bgp dampened-paths Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router	IP address of the router where route dampening is enabled.
*d	Route to the network indicated is dampened.
From	IP address of the peer that advertised this path.

Table 34 show ip bgp dampened-paths Field Descriptions (continued)

Field	Description
Reuse	Time (in hours:minutes:seconds) after which the path will be made available.
Path	Autonomous system path of the route that is being dampened.

Related Commands

Command	Description
bgp dampening	Enables BGP route dampening or changes various BGP route dampening factors.
clear ip bgp dampening	Clears BGP route dampening information and unsuppresses the suppressed routes.

show ip bgp dampening dampened-paths

To display Border Gateway Protocol (BGP) dampened routes on the Cisco 10000 series router, use the **show ip bgp dampening dampened-paths** command in EXEC mode.

show ip bgp dampening dampened-paths [*community-list-number* | *community-list-name* [**exact-match**]]

Syntax Description

community-list-number	(Optional) Community list number. The range is from 1 to 500.
community-list-name	(Optional) Community list name.
exact-match	(Optional) Displays only routes that have an exact match.

Command Modes

EXEC

Command History

Release	Modification
12.2S	This command was introduced.

Usage Guidelines

For router platforms other than the Cisco 10000 series router, use the **show ip bgp dampened-paths** command to display BGP dampened routes.

Examples

The following example show how to display BGP dampened routes information:

```
Router# show ip bgp dampening dampened-paths

BGP table version is 10, local router ID is 172.29.232.182
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          From             Reuse    Path
*d 10.0.0.0         172.16.232.177   00:18:4 100 ?
*d 10.2.0.0         172.16.232.177   00:28:5 100 ?
```

Table 35 describes the significant fields shown in the display.

Table 35 show ip bgp dampening dampened-paths 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 where route dampening is enabled.
*d	Route to the network indicated is dampened.
From	IP address of the peer that advertised this path.

Table 35 show ip bgp dampening dampened-paths Field Descriptions (continued)

Field	Description
Reuse	Time (in hours:minutes:seconds) after which the path will be made available.
Path	Autonomous system (AS) path of the route that is being dampened.

Related Commands

Command	Description
bgp dampening	Enables BGP route dampening or changes various BGP route dampening factors.
clear ip bgp dampening	Clears BGP route dampening information and unsuppresses the suppressed routes.
show dampening interface	Displays a summary of the dampening parameters and status.

show ip bgp dampening flap-statistics

To display Border Gateway Protocol (BGP) flap statistics for all paths on the Cisco 10000 series router, use the **show ip bgp dampening flap-statistics** command in privileged EXEC mode.

show ip bgp dampening flap-statistics [*ip-address* [**mask**] | **cidr-only** | **filter-list** *access-list-number* | i**njected-paths** | **labels** | **prefix-list** *prefix-list* | **quote-regexp** *regexp* | **regexp** *regexp* | **route-map** *route-map-name* | **template** {**peer-policy** *template-name* | **peer-session** *template-name*}]

Syntax Description

ip-address	Specifies the IP address for the flap statistics you want to display.
mask	Specifies the mask to filter or match hosts that are part of the specified network.
cidr-only	Displays flap statistics for routes with classless interdomain routing (CIDR).
filter-list *access-list-number*	Displays flap statistics for routes that conform to the specified autonomous system (AS) path access list number.
injected-paths	Displays flap statistics for all injected paths.
labels	Displays flap statistics for IPv4 Network Layer Reachability Information (NLRI) labels.
prefix-list *prefix-list*	Filters output based on the specified prefix list.
quote-regexp *regexp*	Filters output based on the specified quoted expression.
regexp *regexp*	Filters output based on the specified regular expression.
route-map *route-map-name*	Filters output based on the specified route map.
template	Displays peer-policy or peer-session template information.
peer-policy *template-name*	Used with the **template** keyword, displays peer-policy template information for the specified template name.
peer-session *template-name*	Used with the **template** keyword, displays peer-session template information for the specified template name.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.2S	This command was introduced.

Usage Guidelines

For router platforms other than the Cisco 10000 series router, use the **show ip bgp flap-statistics** command to display BGP flap statistics.

Examples

The following example show how to display the BGP flap statistics for routes with nonnatural network masks (CIDR):

```
Router# show ip bgp dampening flap-statistics cidr-only

BGP table version is 56, local router ID is 100.10.7.11
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i205.0.5.0/30     100.10.5.11              0    100      0 i
*>i205.0.5.4/30     205.0.5.1                0    100      0 105 ?
*>i205.10.5.9/32    205.0.5.1                2    100      0 105 ?
*>i205.10.5.13/32   205.0.5.1                2    100      0 105 ?
*>i206.0.6.0/30     100.10.5.11              0    100      0 i
*>i206.0.6.4/30     206.0.6.1                0    100      0 106 ?
*>i206.10.6.9/32    206.0.6.1                2    100      0 106 ?
*>i206.10.6.13/32   206.0.6.1                2    100      0 106 ?
*> 207.0.7.0/30     0.0.0.0                  0         32768 i
*> 207.0.7.4/30     207.0.7.1                0             0 107 ?
*> 207.10.7.9/32    207.0.7.1                2             0 107 ?
*> 207.10.7.13/32   207.0.7.1                2             0 107 ?
*> 208.0.8.0/30     0.0.0.0                  0         32768 i
*> 208.0.8.4/30     208.0.8.1                0             0 108 ?
*> 208.10.8.9/32    208.0.8.1                2             0 108 ?
*> 208.10.8.13/32   208.0.8.1                2             0 108 ?
```

Table 35 describes the significant fields shown in the display.

Table 36 show ip bgp dampening flap-statistics cidr-only 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 where route dampening is enabled.
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. *—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.
Network	Internet address of the network the entry describes.
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 access server has some non-BGP route 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.

Table 36 show ip bgp dampening flap-statistics cidr-only Field Descriptions (continued)

Field	Description
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. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP.

Related Commands

Command	Description
bgp dampening	Enables BGP route dampening or changes various BGP route dampening factors.
clear ip bgp flap-statistics	Clears BGP flap statistics.
show dampening interface	Displays a summary of the dampening parameters and status.

show ip bgp dampening parameters

To display detailed Border Gateway Protocol (BGP) dampening information on the Cisco 10000 series router, use the **show ip bgp dampening parameters** command in privileged EXEC mode.

show ip bgp dampening parameters

Syntax Description

This command has no arguments or keywords.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.2S	This command was introduced.

Examples

The following example shows how to display detailed BGP dampening information:

```
Router# show ip bgp dampening parameters

dampening 15 750 2000 60 (DEFAULT)
  Half-life time      : 15 mins       Decay Time       : 2320 secs
  Max suppress penalty: 12000         Max suppress time: 60 mins
```

Table 37 describes the significant fields shown in the display.

Table 37 show ip bgp dampening parameters Field Descriptions

Field	Description
Half-life time	Time after which a penalty is decreased, in minutes. Once the interface has been assigned a penalty, the penalty is decreased by half after the half-life period. The process of reducing the penalty happens every 5 seconds. The range of the half-life period is 1 to 45 minutes. The default is 1 minute.
Decay Time	Penalty value below which an unstable interface is unsuppressed, in seconds. The process of unsupressing routers occurs at 10 second increments. The range of the reuse value is 1 to 20000 seconds. The default value is 750 seconds.
Max suppress penalty	Limit at which an interface is suppressed when its penalty exceeds that limit, in seconds. The default value is 2000 seconds.
Max suppress time	Maximum time that an interface can be suppressed, in minutes. This value effectively acts as a ceiling that the penalty value cannot exceed. The default value is four times the half-life period.

Related Commands

Command	Description
bgp dampening	Enables BGP route dampening or changes various BGP route dampening factors.
clear ip bgp dampening	Clears BGP dampening information.
show dampening interface	Displays a summary of the dampening parameters and status.

show ip bgp extcommunity-list

To display routes that match the extended community list in the Border Gateway Protocol (BGP) routing table, use the **show ip bgp extcommunity-list** command in user EXEC or privileged EXEC mode.

show ip bgp extcommunity-list [*list-name*]

Syntax Description

list-name	(Optional) Specifies an extended community list name.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.3(11)T	This command was introduced.
12.2(27)SBC	This command was integrated into the Cisco IOS Release 12.2(27)SBC.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.
Cisco IOS XE Release 2.3	This command was integrated into Cisco IOS XE Release 2.3.

Usage Guidelines

You need to configure the extended community lists by using the **ip extcommunity-list** command for the **show ip bgp extcommunity-list** command to display the output.

Examples

The following is sample output from the **show ip bgp extcommunity-list** command:

```
Router# show ip bgp extcommunity-list 1

Standard extended community-list list1
     9 permit RT:1:100 RT:2:100
     19 deny RT:5:100 RT:6:200
     29 permit RT:4:100
     39 permit RT:5:900
     49 permit RT:4:100 RT:6:200
```

show ip bgp filter-list

To display routes that conform to a specified filter list, use the **show ip bgp filter-list** command in EXEC mode.

show ip bgp filter-list *access-list-number*

Syntax Description

access-list-number	Number of an autonomous system path access list. It can be a number from 1 to 199, or on the Cisco 10000 series router this is a number from 1 to 500.

Command Modes

EXEC

Command History

Release	Modification
10.0	This command was introduced.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Examples

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

```
Router# show ip bgp filter-list 2

BGP table version is 1738, local router ID is 172.16.72.24
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*  172.16.0.0       172.16.72.30                         0 109 108 ?
*  172.16.1.0       172.16.72.30                         0 109 108 ?
*  172.16.11.0      172.16.72.30                         0 109 108 ?
*  172.16.14.0      172.16.72.30                         0 109 108 ?
*  172.16.15.0      172.16.72.30                         0 109 108 ?
*  172.16.16.0      172.16.72.30                         0 109 108 ?
*  172.16.17.0      172.16.72.30                         0 109 108 ?
*  172.16.18.0      172.16.72.30                         0 109 108 ?
*  172.16.19.0      172.16.72.30                         0 109 108 ?
*  172.16.24.0      172.16.72.30                         0 109 108 ?
*  172.16.29.0      172.16.72.30                         0 109 108 ?
*  172.16.30.0      172.16.72.30                         0 109 108 ?
*  172.16.33.0      172.16.72.30                         0 109 108 ?
*  172.16.35.0      172.16.72.30                         0 109 108 ?
*  172.16.36.0      172.16.72.30                         0 109 108 ?
*  172.16.37.0      172.16.72.30                         0 109 108 ?
*  172.16.38.0      172.16.72.30                         0 109 108 ?
*  172.16.39.0      172.16.72.30                         0 109 108 ?
```

Table 38 describes the significant fields shown in the display.

Table 38 show ip bgp filter-list 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.
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. *—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 an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an 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	Internet address of the network the entry describes.
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 route to this network.
Metric	If shown, this is the value of the interautonomous system metric. This field is frequently not used.
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. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a **network** router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP.

show ip bgp flap-statistics

To display BGP flap statistics, use the **show ip bgp flap-statistics** command in EXEC mode.

show ip bgp flap-statistics [**regexp** *regexp* | **filter-list** *access-list* | *ip-address mask* [**longer-prefix**]]

Syntax Description

regexp *regexp*	(Optional) Clears flap statistics for all the paths that match the regular expression.
filter-list *access-list*	(Optional) Clears flap statistics for all the paths that pass the access list.
ip-address	(Optional) Clears flap statistics for a single entry at this IP address.
mask	(Optional) Network mask applied to the value.
longer-prefix	(Optional) Displays flap statistics for more specific entries.

Command Modes

EXEC

Command History

Release	Modification
11.0	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

If no arguments or keywords are specified, the router displays flap statistics for all routes.

Examples

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

```
Router# show ip bgp flap-statistics

BGP table version is 10, local router ID is 172.29.232.182
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          From            Flaps Duration Reuse    Path
*d 10.0.0.0         172.29.232.177  4     00:13:31 00:18:10 100
*d 10.2.0.0         172.29.232.177  4     00:02:45 00:28:20 100
```

Table 39 describes the significant fields shown in the display.

Table 39 show ip bgp flap-statistics 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 where route dampening is enabled.
Network	Route to the network indicated is dampened.
From	IP address of the peer that advertised this path.
Flaps	Number of times the route has flapped.
Duration	Time (in hours:minutes:seconds) since the router noticed the first flap.
Reuse	Time (in hours:minutes:seconds) after which the path will be made available.
Path	Autonomous system path of the route that is being dampened.

Related Commands

Command	Description
bgp dampening	Enables BGP route dampening or changes various BGP route dampening factors.
clear ip bgp flap-statistics	Clears BGP flap statistics.

show ip bgp inconsistent-as

To display routes with inconsistent originating autonomous systems, use the **show ip bgp inconsistent-as** command in EXEC mode.

show ip bgp inconsistent-as

Syntax Description

This command has no arguments or keywords.

Command Modes

EXEC

Command History

Release	Modification
11.0	This command was introduced.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Examples

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

```
Router# show ip bgp inconsistent-as

BGP table version is 87, local router ID is 172.19.82.53
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*  10.1.0.0         172.29.232.55          0             0 300 88 90 99 ?
*>                  172.29.232.52       2222             0 400 ?
*  172.29.0.0       172.29.232.55          0             0 300 90 99 88 200 ?
*>                  172.29.232.52       2222             0 400 ?
*  10.200.199.0     172.29.232.55          0             0 300 88 90 99 ?
*>                  172.29.232.52       2222             0 400 ?
```

show ip bgp injected-paths

To display all the injected paths in the Border Gateway Protocol (BGP) routing table, use the **show ip bgp injected-paths** command in user or privileged EXEC mode.

show ip bgp injected-paths

Syntax Description

This command has no arguments or keywords.

Command Modes

User EXEC
Privileged EXEC

Command History

Release	Modification
12.0(14)ST	This command was introduced.
12.2(4)T	This command was integrated into Cisco IOS Release 12.2(4)T.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB.

Examples

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

```
Router# show ip bgp injected-paths

BGP table version is 11, local router ID is 10.0.0.1
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
*> 172.16.0.0       10.0.0.2                             0 ?
*> 172.17.0.0/16    10.0.0.2                             0 ?
```

Table 40 describes the significant fields shown in the display.

Table 40 show ip bgp injected-paths 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 40 show ip bgp injected-paths 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 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 an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an 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	The Multi Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.)
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.

show ip bgp ipv4

To display entries in the IP version 4 (IPv4) Border Gateway Protocol (BGP) routing table, use the **show ip bgp ipv4** command in privileged EXEC mode.

show ip bgp ipv4 {**mdt** {**all** | **rd** | **vrf**} | **multicast** | **tunnel** | **unicast**}

Syntax Description

mdt	Displays entries for multicast discovery tree sessions.
all	Displays all multicast discovery tree information.
rd	Displays information about the VPN route distinguisher in the MDT session.
vrf	Displays information about the VRF in the MDT session.
multicast	Displays entries for multicast sessions.
tunnel	Displays entries for tunnel sessions.
unicast	Displays entries for unicast sessions.

Command Modes

Privileged EXEC (#)

Command History

Release	Modification
12.0(7)T	This command was introduced.
12.0(29)S	The **mdt** keyword was added.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(31)SB2	This command was integrated into Cisco IOS Release 12.2(31)SB2.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.
12.4(20)T	The **mdt** keyword was added.

Examples

The following is sample output from the **show ip bgp ipv4 unicast** command:

```
Router# show ip bgp ipv4 unicast

 BGP table version is 4, local router ID is 10.0.40.1
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
*> 10.10.10.0/24   172.16.10.1              0             0  300 i
*> 10.10.20.0/24   172.16.10.1              0             0  300 i
*  10.20.10.0/24   172.16.10.1              0             0  300 i
```

The following is sample output from the **show ip bgp ipv4 multicast** command:

```
Router# show ip bgp ipv4 multicast

 BGP table version is 4, local router ID is 10.0.40.1
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
*> 10.10.10.0/24    172.16.10.1              0             0  300 i
*> 10.10.20.0/24    172.16.10.1              0             0  300 i
*  10.20.10.0/24    172.16.10.1              0             0  300 i
```

Table 41 describes the significant fields shown in the display.

Table 41 show ip bgp ipv4 unicast 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.
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 damped. • h—The table entry 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 displayed at the end of each line in the table. It can be one of the following values: • i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. • e—Entry originated from an 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.

Related Commands

Command	Description
clear ip bgp ipv4 mdt	Resets multicast discovery tree IPv4 BGP address-family sessions.
show ip bgp	Displays entries in the BGP routing table.

show ip bgp ipv4 multicast

To display IP Version 4 multicast database-related information, use the **show ip bgp ipv4 multicast** command in EXEC mode.

show ip bgp ipv4 multicast [*command*]

Syntax Description

command	(Optional) Any multiprotocol BGP command supported by the **show ip bgp ipv4 multicast** *command*.

Command Modes

EXEC

Command History

Release	Modification
12.0(7)T	This command was introduced.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

Use this command in conjunction with the **show ip rpf** command to determine if IP multicast routing is using multiprotocol BGP routes.

To determine which multiprotocol BGP commands are supported by the **show ip bgp ipv4 multicast** command, enter the following command while in EXEC mode:

```
Router# show ip bgp ipv4 multicast ?
```

The **show ip bgp ipv4 multicast** command replaces the **show ip mbgp** command.

Examples

The following is sample output from the **show ip bgp ipv4 multicast** command:

```
Router# show ip bgp ipv4 multicast

MBGP table version is 6, local router ID is 192.168.200.66
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
*> 10.0.20.16/28     0.0.0.0                0      0 32768 i
*> 10.0.35.16/28     0.0.0.0                0      0 32768 i
*> 10.0.36.0/28      0.0.0.0                0      0 32768 i
*> 10.0.48.16/28     0.0.0.0                0      0 32768 i
*> 10.2.0.0/16       0.0.0.0                0      0 32768 i
*> 10.2.1.0/24       0.0.0.0                0      0 32768 i
*> 10.2.2.0/24       0.0.0.0                0      0 32768 i
*> 10.2.3.0/24       0.0.0.0                0      0 32768 i
*> 10.2.7.0/24       0.0.0.0                0      0 32768 i
*> 10.2.8.0/24       0.0.0.0                0      0 32768 i
```

```
*> 10.2.10.0/24      0.0.0.0                0      0 32768 i
*> 10.2.11.0/24      0.0.0.0                0      0 32768 i
*> 10.2.12.0/24      0.0.0.0                0      0 32768 i
*> 10.2.13.0/24      0.0.0.0                0      0 32768 i
```

Table 42 describes the significant fields shown in the display.

Table 42 show ip bgp ipv4 multicast Field Descriptions

Field	Description
MBGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
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 historical. *—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 an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration or address family configuration command. e—Entry originated from an 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.

Related Commands

Command	Description
show ip rpf	Displays how IP multicast routing does RPF.

show ip bgp ipv4 multicast summary

To display a summary of IP Version 4 multicast database-related information, use the **show ip bgp ipv4 multicast summary** command in EXEC mode.

show ip bgp ipv4 multicast summary

Syntax Description

This command has no arguments or keywords.

Command Modes

EXEC

Command History

Release	Modification
12.0(7)T	This command was introduced.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

The **show ip bgp ipv4 multicast summary** command replaces the **show ip mbgp summary** command.

Examples

The following is sample output from the **show ip bgp ipv4 multicast summary** command:

```
Router# show ip bgp ipv4 multicast summary

BGP router identifier 10.0.33.34, local AS number 34
BGP table version is 5, main routing table version 1
4 network entries and 6 paths using 604 bytes of memory
5 BGP path attribute entries using 260 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
2 BGP community entries using 48 bytes of memory
2 BGP route-map cache entries using 32 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP activity 8/28 prefixes, 12/0 paths, scan interval 15 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.33.35      4    35     624     624        5    0    0 10:13:46        3
```

Table 43 describes the significant fields shown in the display.

Table 43 show ip bgp ipv4 multicast summary Field Descriptions

Field	Description
Neighbor	IP address of configured neighbor in the multicast routing table.
V	Version of multiprotocol BGP used.
AS	Autonomous system to which the neighbor belongs.

Table 43 show ip bgp ipv4 multicast summary Field Descriptions (continued)

Field	Description
MsgRcvd	Number of messages received from the neighbor.
MsgSent	Number of messages sent to the neighbor.
TblVer	Number of the table version, which is incremented each time the table changes.
InQ	Number of messages received in the input queue.
OutQ	Number of messages ready to go in the output queue.
Up/Down	Days and hours that the neighbor has been up or down (no information in the State column means the connection is up).
State/PfxRcd	State of the neighbor/number of routes received. If no state is indicated, the state is up.

Related Commands

Command	Description
show ip rpf	Displays how IP multicast routing does RPF.

show ip bgp l2vpn

To display Layer 2 Virtual Private Network (L2VPN) address family information from the Border Gateway Protocol (BGP) table, use the **show ip bgp l2vpn** command in user EXEC or privileged EXEC mode.

With BGP show Command Argument

show ip bgp l2vp vpls {**all** | **rd** *route-distinguisher*} [*bgp-keyword*]

With IP Prefix and Mask Length Syntax

show ip bgp l2vp vpls {**all** | **rd** *route-distinguisher*} [*ip-prefix*/*length* [**bestpath**] [**longer-prefixes** [**injected**]] [**multipaths**] [**shorter-prefixes** [*mask-length*]] [**subnets**]]

With Network Address Syntax

show ip bgp l2vp vpls {**all** | **rd** *route-distinguisher*} [*network-address* [*mask* | **bestpath** | **multipaths**] [**bestpath**] [**longer-prefixes** [**injected**]] [**multipaths**] [**shorter-prefixes** [*mask-length*]] [**subnets**]]

Syntax Description

vpls	Displays L2VPN address family database information for the Virtual Private LAN Service (VPLS) subsequent address family identifier (SAFI).
all	Displays the complete L2VPN database.
rd *route-distinguisher*	Displays prefixes that match the specified route distinguisher.
bgp-keyword	(Optional) Argument representing a **show ip bgp** command keyword that can be added to this command. See Table 44.
ip-prefix/*length*	(Optional) The IP prefix address (in dotted decimal format) and the length of the mask (0 to 32). The slash mark must be included.
bestpath	(Optional) Displays the best path for the specified prefix.
longer-prefixes	(Optional) Displays the route and more specific routes.
injected	(Optional) Displays more specific routes that were injected because of the specified prefix.
multipaths	(Optional) Displays the multipaths for the specified prefix.
shorter-prefixes	(Optional) Displays the less specific routes.
mask-length	(Optional) The length of the mask as a number in the range from 0 to 32. Prefixes longer than the specified mask length are displayed.
subnets	(Optional) Displays the subnet routes for the specified prefix.
network-address	(Optional) The IP address of a network in the BGP routing table.
mask	(Optional) The mask of the network address, in dotted decimal format.

Command Default

If no arguments or keywords are specified, this command displays the complete L2VPN database.

Command Modes

User EXEC
Privileged EXEC

Command History

Release	Modification
12.2(33)SRB	This command was introduced.
Cisco IOS XE 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Usage Guidelines

Table 44 displays optional **show ip bgp** command keywords that can be configured with the **show ip bgp l2vpn** command. Replace the *bgp-keyword* argument with the appropriate keyword from the table. For more details about each command in its **show ip bgp** *bgp-keyword* form, see the *Cisco IOS IP Routing Protocols Command Reference*, Release 12.2SR.

Table 44 Optional show ip bgp Command Keywords and Descriptions

Keyword	Description
community	Displays routes that match a specified community
community-list	Displays routes that match a specified community list.
dampening	Displays paths suppressed because of dampening (BGP route from peer is up and down).
extcommunity-list	Displays routes that match a specified extcommunity list.
filter-list	Displays routes that conform to the filter list.
inconsistent-as	Displays only routes that have inconsistent autonomous systems of origin.
neighbors	Displays details about TCP and BGP neighbor connections.
oer-paths	Displays all OER-managed path information.
paths [*regexp*]	Displays autonomous system path information. If the optional *regexp* argument is entered, the autonomous system paths that are displayed match the autonomous system path regular expression.
peer-group	Displays information about peer groups.
pending-prefixes	Displays prefixes that are pending deletion.
prefix-list	Displays routes that match a specified prefix list.
quote-regexp	Displays routes that match the quoted autonomous system path regular expression.
regexp	Displays routes that match the autonomous system path regular expression.
replication	Displays the replication status update groups.
route-map	Displays routes that match the specified route map.
rt-filter-list	Displays the specified inbound route target filter list.
summary	Displays a summary of BGP neighbor status.
update-group	Displays information on update groups.

Examples

The following example shows output for the **show ip bgp l2vpn** command when the **vpls** and **all** keywords are used to display the complete L2VPN database:

```
Router# show ip bgp l2vpn vpls all

BGP table version is 5, local router ID is 192.168.3.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 45000:100
*> 45000:100:172.17.1.1/96
                    0.0.0.0                            32768 ?
*>i45000:100:172.18.2.2/96
                    172.16.1.2               0    100      0 ?
Route Distinguisher: 45000:200
*> 45000:200:172.17.1.1/96
                    0.0.0.0                            32768 ?
*>i45000:200:172.18.2.2/96
                    172.16.1.2               0    100      0 ?
```

Table 45 describes the significant fields shown in the display.

Table 45 show ip bgp l2vpn vpls all 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.
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 a historical entry. • *—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. • r—The table entry failed to install in the routing information base (RIB) table. • S—The table entry is Stale (old). This entry is useful in BGP graceful restart situations.

Table 45 show ip bgp l2vpn vpls all Field Descriptions (continued)

Field	Description
Origin codes	Origin of the entry. The origin code is displayed at the end of each line in the table. It can be one of the following values: • i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. • e—Entry originated from an 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.
Route Distinguisher	Route distinguisher that identifies a set of routing and forwarding tables used in virtual private networks.

The following example shows output for the **show ip bgp l2vpn** command when the **vpls** and **rd** keywords are used to display the L2VPN information that matches the route distinguisher 45000:100. Note that the information displayed is a subset of the information displayed using the **all** keyword.

```
Router# show ip bgp l2vpn vpls rd 45000:100

BGP table version is 5, local router ID is 192.168.3.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 45000:100
*> 45000:100:172.17.1.1/96
                    0.0.0.0                              32768 ?
*>i45000:100:172.18.2.2/96
                    172.16.1.2               0    100      0 ?
```

Related Commands

Command	Description
address-family l2vpn	Enters address family configuration mode to configure a routing session using L2VPN endpoint provisioning information.

show ip bgp neighbors

To display information about Border Gateway Protocol (BGP) and TCP connections to neighbors, use the **show ip bgp neighbors** command in user or privileged EXEC mode.

show ip bgp [**ipv4** {**multicast** | **unicast**} | **vpnv4 all** | **vpnv6 unicast all**] **neighbors** [**slow** | *ip-address* | *ipv6-address* [**advertised-routes** | **dampened-routes** | **flap-statistics** | **paths** [*reg-exp*] | **policy** [**detail**] | **received prefix-filter** | **received-routes** | **routes**]]

Syntax Description

ipv4 {**multicast** | **unicast**}	(Optional) Displays peers in the IPv4 address family.
vpnv4 all	(Optional) Displays peers in the VPNv4 address family.
vpnv6 unicast all	(Optional) Displays peers in the VPNv6 address family.
slow	(Optional) Displays information about dynamically configured slow peers.
ip-address	(Optional) Displays information about the IPv4 neighbor. If this argument is omitted, information about all neighbors is displayed.
ipv6-address	(Optional) Displays information about the IPv6 neighbor.
advertised-routes	(Optional) Displays all routes that have been advertised to neighbors.
dampened-routes	(Optional) Displays the dampened routes received from the specified neighbor.
flap-statistics	(Optional) Displays the flap statistics of the routes learned from the specified neighbor (for external BGP peers only).
paths *reg-exp*	(Optional) Displays autonomous system paths learned from the specified neighbor. An optional regular expression can be used to filter the output.
policy	(Optional) Displays the policies applied to this neighbor per address family.
detail	(Optional) Displays detailed policy information such as route maps, prefix lists, community lists, access control lists (ACLs), and autonomous system path filter lists.
received prefix-filter	(Optional) Displays the prefix-list (outbound route filter [ORF]) sent from the specified neighbor.
received-routes	(Optional) Displays all received routes (both accepted and rejected) from the specified neighbor.
routes	(Optional) Displays all routes that are received and accepted. The output displayed when this keyword is entered is a subset of the output displayed by the **received-routes** keyword.

Command Default

The output of this command displays information for all neighbors.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

OS Release	Modification
12.0(18)S	The output was modified to display the no-prepend configuration option, and this command was integrated into Cisco IOS Release 12.0(18)S.
12.0(21)ST	The output was modified to display Multiprotocol Label Switching (MPLS) label information.
12.0(22)S	Support for the BGP graceful restart capability was integrated into the output. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was also added.
12.0(25)S	The **policy** and **detail** keywords were added.
12.0(27)S	The command output was modified to support the BGP TTL Security Check feature and to display explicit-null label information.
12.0(31)S	Support for the Bidirectional Forwarding Detection (BFD) feature was integrated into the output.
12.0(32)S12	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
12.0(32)SY8	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.0(33)S3	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format became asplain.
S Release	**Modification**
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(17b)SXA	This command was integrated into Cisco IOS Release 12.2(17b)SXA.
12.2(18)SXE	Support for the Bidirectional Forwarding Detection (BFD) feature was integrated into the output.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA, and the output was modified to support BGP TCP path MTU discovery.
12.2(33)SRB	Support for the **policy** and **detail** keywords was integrated into Cisco IOS Release 12.2(33)SRB.
12.2(33)SXH	The output was modified to support BGP dynamic neighbors.
12.2(33)SRC	The output was modified to support BGP graceful restart per peer.
12.2(33)SB	The output was modified to support the BFD and the BGP graceful restart per peer features, and support for the **policy** and **detail** keywords was integrated into Cisco IOS Release 12.2(33)SB.
12.2(33)SXI1	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)SRE	This command was modified. The command output was modified to support the BGP best external and BGP additional path features. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)XNE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
15.0(1)S	This command was modified. The **slow** keyword was added.
15.1(1)S	This command was modified to display the Layer 2 VPN address family if graceful restart (GR) or nonstop forwarding (NSF) is enabled.

Mainline and T Release	Modification
10.0	This command was introduced.
11.2	The **received-routes** keyword was added.
12.2(4)T	The **received** and **prefix-filter** keywords were added, and this command was integrated into Cisco IOS Release 12.2(4)T.
12.2(15)T	Support for the BGP graceful restart capability was integrated into the output.
12.3(7)T	The command output was modified to support the BGP TTL Security Check feature and to display explicit-null label information.
12.4(4)T	Support for the Bidirectional Forwarding Detection (BFD) feature was integrated into the output.
12.4(11)T	Support for the **policy** and **detail** keywords was integrated into Cisco IOS Release 12.4(11)T.
12.4(20)T	The output was modified to support BGP TCP path MTU discovery.
12.4(24)T	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
Cisco IOS XE	**Modification**
Cisco IOS XE Release 2.1	This command was integrated into Cisco IOS XE Release 2.1.
Cisco IOS XE Release 2.4	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format became asplain.
Cisco IOS XE Release 3.1S	This command was modified. The **slow** keyword was added.

Usage Guidelines

Use the **show ip bgp neighbors** command to display BGP and TCP connection information for neighbor sessions. For BGP, this includes detailed neighbor attribute, capability, path, and prefix information. For TCP, this includes statistics related to BGP neighbor session establishment and maintenance.

Prefix activity is displayed based on the number of prefixes that are advertised and withdrawn. Policy denials display the number of routes that were advertised but then ignored based on the function or attribute that is displayed in the output.

In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, the Cisco implementation of 4-byte autonomous system numbers uses asplain—65538 for example—as the default regular expression match and output display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain format and the asdot format as described in RFC 5396. To change the default regular expression match and output display of 4-byte autonomous system numbers to asdot format, use the **bgp asnotation dot** command followed by the **clear ip bgp *** command to perform a hard reset of all current BGP sessions.

In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, the Cisco implementation of 4-byte autonomous system numbers uses asdot—1.2 for example—as the only configuration format, regular expression match, and output display, with no asplain support.

Cisco IOS Releases 12.0(25)S, 12.4(11)T, 12.2(33)SRB, 12.2(33)SB, and Later Releases

When BGP neighbors use multiple levels of peer templates, it can be difficult to determine which policies are applied to the neighbor.

In Cisco IOS Releases 12.0(25)S, 12.4(11)T, 12.2(33)SRB, 12.2(33)SB, and later releases, the **policy** and **detail** keywords were added to display the inherited policies and the policies configured directly on the specified neighbor. Inherited policies are policies that the neighbor inherits from a peer-group or a peer-policy template.

Examples

Example output is different for the various keywords available for the **show ip bgp neighbors** command. Examples using the various keywords appear in the following sections:

- show ip bgp neighbors: Example, page 501
- show ip bgp neighbors (4-Byte Autonomous System Numbers): Example, page 507
- show ip bgp neighbors advertised-routes: Example, page 507
- show ip bgp neighbors paths: Example, page 509
- show ip bgp neighbors received prefix-filter: Example, page 509
- show ip bgp neighbors policy: Example, page 509
- Cisco IOS Release 12.0(31)S, 12.4(4)T, 12.2(18)SXE, and 12.2(33)SB: Example, page 510
- Cisco IOS Release 12.2(33)SRA and 12.4(20)T: Example, page 510
- Cisco IOS Release 12.2(33)SXH: Example, page 510
- Cisco IOS Releases 12.2(33)SRC and 12.2(33)SB: Example, page 511
- Cisco IOS Release 15.1(1)S: Example, page 511

show ip bgp neighbors: Example

The following example shows output for the BGP neighbor at 10.108.50.2. This neighbor is an internal BGP (iBGP) peer. This neighbor supports the route refresh and graceful restart capabilities.

```
Router# show ip bgp neighbors 10.108.50.2

BGP neighbor is 10.108.50.2,  remote AS 1, internal link
  BGP version 4, remote router ID 192.168.252.252
  BGP state = Established, up for 00:24:25
  Last read 00:00:24, last write 00:00:24, hold time is 180, keepalive interval is
   60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    MPLS Label capability: advertised and received
    Graceful Restart Capability: advertised
    Address family IPv4 Unicast: advertised and received
  Message statistics:
    InQ depth is 0
    OutQ depth is 0
                         Sent       Rcvd
    Opens:                  3          3
    Notifications:          0          0
    Updates:                0          0
    Keepalives:           113        112
    Route Refresh:          0          0
    Total:                116        115
  Default minimum time between advertisement runs is 5 seconds

 For address family: IPv4 Unicast
  BGP additional-paths computation is enabled
  BGP advertise-best-external is enabled
  BGP table version 1, neighbor version 1/0
 Output queue size : 0
```

```
  Index 1, Offset 0, Mask 0x2
  1 update-group member
                                Sent       Rcvd
  Prefix activity:              ----       ----
    Prefixes Current:              0          0
    Prefixes Total:                0          0
    Implicit Withdraw:             0          0
    Explicit Withdraw:             0          0
    Used as bestpath:            n/a          0
    Used as multipath:           n/a          0

                                  Outbound    Inbound
  Local Policy Denied Prefixes:   --------    -------
    Total:                               0          0
  Number of NLRIs in the update sent: max 0, min 0

  Connections established 3; dropped 2
  Last reset 00:24:26, due to Peer closed the session
External BGP neighbor may be up to 2 hops away.
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled
Local host: 10.108.50.1, Local port: 179
Foreign host: 10.108.50.2, Foreign port: 42698

Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)

Event Timers (current time is 0x68B944):
Timer          Starts    Wakeups            Next
Retrans            27          0             0x0
TimeWait            0          0             0x0
AckHold            27         18             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0

iss: 3915509457  snduna: 3915510016  sndnxt: 3915510016     sndwnd:  15826
irs:  233567076  rcvnxt:  233567616  rcvwnd:      15845  delrcvwnd:    539

SRTT: 292 ms, RTTO: 359 ms, RTV: 67 ms, KRTT: 0 ms
minRTT: 12 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: passive open, nagle, gen tcbs
IP Precedence value : 6

Datagrams (max data segment is 1460 bytes):
Rcvd: 38 (out of order: 0), with data: 27, total data bytes: 539
Sent: 45 (retransmit: 0, fastretransmit: 0, partialack: 0, Second Congestion: 08
```

Table 46 describes the significant fields shown in the display. Fields that are preceded by the asterisk character (*) are displayed only when the counter has a nonzero value.

Table 46 show ip bgp neighbors Field Descriptions

Field	Description
BGP neighbor	IP address of the BGP neighbor and its autonomous system number.
remote AS	Autonomous system number of the neighbor.

Table 46 show ip bgp neighbors Field Descriptions (continued)

Field	Description
local AS 300 no-prepend (not shown in display)	Verifies that the local autonomous system number is not prepended to received external routes. This output supports the hiding of the local autonomous systems when migrating autonomous systems.
internal link	"internal link" is displayed for iBGP neighbors. "external link" is displayed for external BGP (eBGP) neighbors.
BGP version	BGP version being used to communicate with the remote router.
remote router ID	IP address of the neighbor.
BGP state	Finite state machine (FSM) stage of session negotiation.
up for	Time, in hhmmss, that the underlying TCP connection has been in existence.
Last read	Time, in hhmmss, since BGP last received a message from this neighbor.
last write	Time, in hhmmss, since BGP last sent a message to this neighbor.
hold time	Time, in seconds, that BGP will maintain the session with this neighbor without receiving a messages.
keepalive interval	Time interval, in seconds, at which keepalive messages are transmitted to this neighbor.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor. "advertised and received" is displayed when a capability is successfully exchanged between two routers.
Route Refresh	Status of the route refresh capability.
MPLS Label Capability	Indicates that MPLS labels are both sent and received by the eBGP peer.
Graceful Restart Capability	Status of the graceful restart capability.
Address family IPv4 Unicast	IP Version 4 unicast-specific properties of this neighbor.
Message statistics	Statistics organized by message type.
InQ depth is	Number of messages in the input queue.
OutQ depth is	Number of messages in the output queue.
Sent	Total number of transmitted messages.
Received	Total number of received messages.
Opens	Number of open messages sent and received.
notifications	Number of notification (error) messages sent and received.
Updates	Number of update messages sent and received.
Keepalives	Number of keepalive messages sent and received.
Route Refresh	Number of route refresh request messages sent and received.
Total	Total number of messages sent and received.
Default minimum time between...	Time, in seconds, between advertisement transmissions.

Table 46 show ip bgp neighbors Field Descriptions (continued)

Field	Description
For address family:	Address family to which the following fields refer.
BGP table version	Internal version number of the table. This is the primary routing table with which the neighbor has been updated. The number increments when the table changes.
neighbor version	Number used by the software to track prefixes that have been sent and those that need to be sent.
...update-group	Number of update-group member for this address family.
Prefix activity	Prefix statistics for this address family.
Prefixes current	Number of prefixes accepted for this address family.
Prefixes total	Total number of received prefixes.
Implicit Withdraw	Number of times that a prefix has been withdrawn and readvertised.
Explicit Withdraw	Number of times that prefix has been withdrawn because it is no longer feasible.
Used as bestpath	Number of received prefixes installed as bestpaths.
Used as multipath	Number of received prefixes installed as multipaths.
* Saved (soft-reconfig)	Number of soft resets performed with a neighbor that supports soft reconfiguration. This field is displayed only if the counter has a nonzero value.
* History paths	This field is displayed only if the counter has a nonzero value.
* Invalid paths	Number of invalid paths. This field is displayed only if the counter has a nonzero value.
Local Policy Denied Prefixes	Prefixes denied due to local policy configuration. Counters are updated for inbound and outbound policy denials. The fields under this heading are displayed only if the counter has a nonzero value.
* route-map	Displays inbound and outbound route-map policy denials.
* filter-list	Displays inbound and outbound filter-list policy denials.
* prefix-list	Displays inbound and outbound prefix-list policy denials.
* Ext Community	Displays only outbound extended community policy denials.
* AS_PATH too long	Displays outbound AS-path length policy denials.
* AS_PATH loop	Displays outbound AS-path loop policy denials.
* AS_PATH confed info	Displays outbound confederation policy denials.
* AS_PATH contains AS 0	Displays outbound denials of autonomous system (AS) 0.
* NEXT_HOP Martian	Displays outbound martian denials.
* NEXT_HOP non-local	Displays outbound non-local next-hop denials.
* NEXT_HOP is us	Displays outbound next-hop-self denials.
* CLUSTER_LIST loop	Displays outbound cluster-list loop denials.
* ORIGINATOR loop	Displays outbound denials of local originated routes.

Table 46 show ip bgp neighbors Field Descriptions (continued)

Field	Description
* unsuppress-map	Displays inbound denials due to an unsuppress-map.
* advertise-map	Displays inbound denials due to an advertise-map.
* VPN Imported prefix	Displays inbound denials of VPN prefixes.
* Well-known Community	Displays inbound denials of well-known communities.
* SOO loop	Displays inbound denials due to site-of-origin.
* Bestpath from this peer	Displays inbound denials because the bestpath came from the local router.
* Suppressed due to dampening	Displays inbound denials because the neighbor or link is in a dampening state.
* Bestpath from iBGP peer	Deploys inbound denials because the bestpath came from an iBGP neighbor.
* Incorrect RIB for CE	Deploys inbound denials due to RIB errors for a CE router.
* BGP distribute-list	Displays inbound denials due to a distribute list.
Number of NLRIs...	Number of network layer reachability attributes in updates.
Connections established	Number of times a TCP and BGP connection has been successfully established.
dropped	Number of times that a valid session has failed or been taken down.
Last reset	Time since this peering session was last reset. The reason for the reset is displayed on this line.
External BGP neighbor may be... (not shown in the display)	Indicates that the BGP TTL security check is enabled. The maximum number of hops that can separate the local and remote peer is displayed on this line.
Connection state	Connection status of the BGP peer.
Connection is ECN Disabled	Explicit congestion notification status (enabled or disabled).
Local host: 10.108.50.1, Local port: 179	IP address of the local BGP speaker. BGP port number 179.
Foreign host: 10.108.50.2, Foreign port: 42698	Neighbor address and BGP destination port number.
Enqueued packets for retransmit:	Packets queued for retransmission by TCP.
Event Timers	TCP event timers. Counters are provided for starts and wakeups (expired timers).
Retrans	Number of times a packet has been retransmitted.
TimeWait	Time waiting for the retransmission timers to expire.
AckHold	Acknowledgment hold timer.
SendWnd	Transmission (send) window.
KeepAlive	Number of keepalive packets.
GiveUp	Number times a packet is dropped due to no acknowledgment.
PmtuAger	Path MTU discovery timer.

Table 46 show ip bgp neighbors Field Descriptions (continued)

Field	Description
DeadWait	Expiration timer for dead segments.
iss:	Initial packet transmission sequence number.
snduna:	Last transmission sequence number that has not been acknowledged.
sndnxt:	Next packet sequence number to be transmitted.
sndwnd:	TCP window size of the remote neighbor.
irs:	Initial packet receive sequence number.
rcvnxt:	Last receive sequence number that has been locally acknowledged.
rcvwnd:	TCP window size of the local host.
delrcvwnd:	Delayed receive window—data the local host has read from the connection, but has not yet subtracted from the receive window the host has advertised to the remote host. The value in this field gradually increases until it is larger than a full-sized packet, at which point it is applied to the rcvwnd field.
SRTT:	A calculated smoothed round-trip timeout.
RTTO:	Round-trip timeout.
RTV:	Variance of the round-trip time.
KRTT:	New round-trip timeout (using the Karn algorithm). This field separately tracks the round-trip time of packets that have been re-sent.
minRTT:	Smallest recorded round-trip timeout (hard-wire value used for calculation).
maxRTT:	Largest recorded round-trip timeout.
ACK hold:	Length of time the local host will delay an acknowledgment to carry (piggyback) additional data.
IP Precedence value:	IP precedence of the BGP packets.
Datagrams	Number of update packets received from a neighbor.
Rcvd:	Number of received packets.
with data	Number of update packets sent with data.
total data bytes	Total amount of data received, in bytes.
Sent	Number of update packets sent.
Second Congestion	Number of update packets with data sent.
Datagrams: Rcvd	Number of update packets received from a neighbor.
out of order:	Number of packets received out of sequence.
with data	Number of update packets received with data.
Last reset	Elapsed time since this peering session was last reset.
unread input bytes	Number of bytes of packets still to be processed.
retransmit	Number of packets retransmitted.

Table 46 show ip bgp neighbors Field Descriptions (continued)

Field	Description
fastretransmit	Number of duplicate acknowledgments retransmitted for an out of order segment before the retransmission timer expires.
partialack	Number of retransmissions for partial acknowledgements (transmissions before or without subsequent acknowledgments).
Second Congestion	Number of second retransmissions sent due to congestion.

show ip bgp neighbors (4-Byte Autonomous System Numbers): Example

The following partial example shows output for several external BGP neighbors in autonomous systems with 4-byte autonomous system numbers, 65536 and 65550. This example requires Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, or a later release.

```
Router# show ip bgp neighbors

BGP neighbor is 192.168.1.2,  remote AS 65536, external link
  BGP version 4, remote router ID 0.0.0.0
  BGP state = Idle
  Last read 02:03:38, last write 02:03:38, hold time is 120, keepalive interval is 70
seconds
  Configured hold time is 120, keepalive interval is 70 seconds
  Minimum holdtime from neighbor is 0 seconds
.
.
.
BGP neighbor is 192.168.3.2,  remote AS 65550, external link
 Description: finance
  BGP version 4, remote router ID 0.0.0.0
  BGP state = Idle
  Last read 02:03:48, last write 02:03:48, hold time is 120, keepalive interval is 70
seconds
  Configured hold time is 120, keepalive interval is 70 seconds
  Minimum holdtime from neighbor is 0 seconds
```

show ip bgp neighbors advertised-routes: Example

The following example displays routes advertised for only the 172.16.232.178 neighbor:

```
Router# show ip bgp neighbors 172.16.232.178 advertised-routes

BGP table version is 27, local router ID is 172.16.232.181
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
*>i10.0.0.0      172.16.232.179         0    100      0 ?
*> 10.20.2.0     10.0.0.0               0         32768 i
```

Table 47 describes the significant fields shown in the display.

Table 47 show ip bgp neighbors advertised-routes Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This is the primary routing table with which the neighbor has been updated. The number increments when the table changes.
local router ID	IP address of the local BGP speaker.
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 and will not be advertised to BGP neighbors. • h—The table entry does not contain the best path based on historical information. • *—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 route that is redistributed into BGP from an IGP.
Network	IP address of a network entity.
Next Hop	IP address of the next system used to forward a packet to the destination network. An entry of 0.0.0.0 indicates that there are non-BGP routes in the path to the destination network.
Metric	If shown, this is the value of the interautonomous system metric. This field is not used frequently.
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.

show ip bgp neighbors paths: Example

The following is example output from the **show ip bgp neighbors** command entered with the **paths** keyword:

```
Router# show ip bgp neighbors 172.29.232.178 paths ^10

Address    Refcount Metric Path
0x60E577B0        2     40 10 ?
```

Table 48 describes the significant fields shown in the display.

Table 48 show ip bgp neighbors paths Field Descriptions

Field	Description
Address	Internal address where the path is stored.
Refcount	Number of routes using that path.
Metric	Multi Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.)
Path	Autonomous system path for that route, followed by the origin code for that route.

show ip bgp neighbors received prefix-filter: Example

The following example shows that a prefix-list that filters all routes in the 10.0.0.0 network has been received from the 192.168.20.72 neighbor:

```
Router# show ip bgp neighbors 192.168.20.72 received prefix-filter

Address family:IPv4 Unicast
ip prefix-list 192.168.20.72:1 entries
   seq 5 deny 10.0.0.0/8 le 32
```

Table 49 describes the significant fields shown in the display.

Table 49 show ip bgp neighbors received prefix-filter Field Descriptions

Field	Description
Address family	Address family mode in which the prefix filter is received.
ip prefix-list	Prefix list sent from the specified neighbor.

show ip bgp neighbors policy: Example

The following sample output shows the policies applied to the neighbor at 192.168.1.2. The output displays both inherited policies and policies configured on the neighbor device. Inherited polices are policies that the neighbor inherits from a peer-group or a peer-policy template.

```
Router# show ip bgp neighbors 192.168.1.2 policy

Neighbor: 192.168.1.2, Address-Family: IPv4 Unicast
Locally configured policies:
 route-map ROUTE in
Inherited polices:
 prefix-list NO-MARKETING in
 route-map ROUTE in
 weight 300
 maximum-prefix 10000
```

Cisco IOS Release 12.0(31)S, 12.4(4)T, 12.2(18)SXE, and 12.2(33)SB: Example

The following is sample output from the **show ip bgp neighbors** command that verifies that Bidirectional Forwarding Detection (BFD) is being used to detect fast fallover for the BGP neighbor that is a BFD peer:

```
Router# show ip bgp neighbors

BGP neighbor is 172.16.10.2,  remote AS 45000, external link
.
.
.
 Using BFD to detect fast fallover
```

Cisco IOS Release 12.2(33)SRA and 12.4(20)T: Example

The following is sample output from the **show ip bgp neighbors** command that verifies that BGP TCP path maximum transmission unit (MTU) discovery is enabled for the BGP neighbor at 172.16.1.2:

```
Router# show ip bgp neighbors 172.16.1.2

BGP neighbor is 172.16.1.2,  remote AS 45000, internal link
  BGP version 4, remote router ID 172.16.1.99
.
.
.
 For address family: IPv4 Unicast
  BGP table version 5, neighbor version 5/0
.
.
.
  Address tracking is enabled, the RIB does have a route to 172.16.1.2
  Address tracking requires at least a /24 route to the peer
  Connections established 3; dropped 2
  Last reset 00:00:35, due to Router ID changed
  Transport(tcp) path-mtu-discovery is enabled
.
.
.
SRTT: 146 ms, RTTO: 1283 ms, RTV: 1137 ms, KRTT: 0 ms
minRTT: 8 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: higher precedence, retransmission timeout, nagle, path mtu capable
```

Cisco IOS Release 12.2(33)SXH: Example

The following is sample output from the **show ip bgp neighbors** command that verifies that the neighbor 192.168.3.2 is a member of the peer group, group192, and belongs to the subnet range group 192.168.0.0/16, which shows that this BGP neighbor was dynamically created:

```
Router# show ip bgp neighbors 192.168.3.2

BGP neighbor is *192.168.3.2,  remote AS 50000, external link
 Member of peer-group group192 for session parameters
 Belongs to the subnet range group: 192.168.0.0/16
  BGP version 4, remote router ID 192.168.3.2
  BGP state = Established, up for 00:06:35
  Last read 00:00:33, last write 00:00:25, hold time is 180, keepalive intervals
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Address family IPv4 Unicast: advertised and received
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

                         Sent       Rcvd
```

```
    Opens:                  1          1
    Notifications:          0          0
    Updates:                0          0
    Keepalives:             7          7
    Route Refresh:          0          0
    Total:                  8          8
  Default minimum time between advertisement runs is 30 seconds

 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1/0
  Output queue size : 0
  Index 1, Offset 0, Mask 0x2
  1 update-group member
  group192 peer-group member
.
.
.
```

Cisco IOS Releases 12.2(33)SRC and 12.2(33)SB: Example

The following is partial output from the **show ip bgp neighbors** command that verifies the status of the BGP graceful restart capability for the external BGP peer at 192.168.3.2. Graceful restart is shown as disabled for this BGP peer.

```
Router# show ip bgp neighbors 192.168.3.2

BGP neighbor is 192.168.3.2,  remote AS 50000, external link
 Inherits from template S2 for session parameters
  BGP version 4, remote router ID 192.168.3.2
  BGP state = Established, up for 00:01:41
  Last read 00:00:45, last write 00:00:45, hold time is 180, keepalive intervals
  Neighbor sessions:
    1 active, is multisession capable
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Address family IPv4 Unicast: advertised and received
.
.
.
Address tracking is enabled, the RIB does have a route to 192.168.3.2
  Connections established 1; dropped 0
  Last reset never
  Transport(tcp) path-mtu-discovery is enabled
  Graceful-Restart is disabled
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
```

Cisco IOS Release 15.1(1)S: Example

The following is partial output from the **show ip bgp neighbors** command. For this release, the display includes the Layer 2 VFN address family information if graceful restart or NSF is enabled.

For more information about the other fields shown in the display, see Table 46 on page 502.

```
Router# show ip bgp neighbors

Load for five secs: 2%/0%; one minute: 0%; five minutes: 0%
Time source is hardware calendar, *21:49:17.034 GMT Wed Sep 22 2010

BGP neighbor is 10.1.1.3,  remote AS 2, internal link
  BGP version 4, remote router ID 10.1.1.3
  BGP state = Established, up for 00:14:32
  Last read 00:00:30, last write 00:00:43, hold time is 180, keepalive interval is 60
seconds
```

```
 Neighbor sessions:
   1 active, is not multisession capable (disabled)
 Neighbor capabilities:
   Route refresh: advertised and received(new)
   Four-octets ASN Capability: advertised and received
   Address family IPv4 Unicast: advertised and received
   Address family L2VPN Vpls: advertised and received
   Graceful Restart Capability: advertised and received
     Remote Restart timer is 120 seconds
     Address families advertised by peer:
       IPv4 Unicast (was not preserved), L2VPN Vpls (was not preserved)
   Multisession Capability:
 Message statistics:
   InQ depth is 0
   OutQ depth is 0

                        Sent       Rcvd
   Opens:                  1          1
   Notifications:          0          0
   Updates:                4         16
   Keepalives:            16         16
   Route Refresh:          0          0
   Total:                 21         33
 Default minimum time between advertisement runs is 0 seconds

For address family: IPv4 Unicast
 Session: 10.1.1.3
 BGP table version 34, neighbor version 34/0
 Output queue size : 0
 Index 1, Advertise bit 0
 1 update-group member
 Slow-peer detection is disabled
 Slow-peer split-update-group dynamic is disabled
                               Sent       Rcvd
 Prefix activity:              ----       ----
   Prefixes Current:              2         11 (Consumes 572 bytes)
   Prefixes Total:                4         19
   Implicit Withdraw:             2          6
   Explicit Withdraw:             0          2
   Used as bestpath:             n/a         7
   Used as multipath:            n/a         0

                                  Outbound    Inbound
 Local Policy Denied Prefixes:    --------    -------
   NEXT_HOP is us:                     n/a          1
   Bestpath from this peer:             20        n/a
   Bestpath from iBGP peer:              8        n/a
   Invalid Path:                        10        n/a
   Total:                               38          1
 Number of NLRIs in the update sent: max 2, min 0
 Last detected as dynamic slow peer: never
 Dynamic slow peer recovered: never

For address family: L2VPN Vpls
 Session: 10.1.1.3
 BGP table version 8, neighbor version 8/0
 Output queue size : 0
 Index 1, Advertise bit 0
 1 update-group member
 Slow-peer detection is disabled
 Slow-peer split-update-group dynamic is disabled
                               Sent       Rcvd
 Prefix activity:              ----       ----
   Prefixes Current:              1          1 (Consumes 68 bytes)
```

```
    Prefixes Total:                 2          1
    Implicit Withdraw:              1          0
    Explicit Withdraw:              0          0
    Used as bestpath:             n/a          1
    Used as multipath:            n/a          0

                                  Outbound    Inbound
  Local Policy Denied Prefixes:   --------    -------
    Bestpath from this peer:             4        n/a
    Bestpath from iBGP peer:             1        n/a
    Invalid Path:                        2        n/a
    Total:                               7          0
  Number of NLRIs in the update sent: max 1, min 0
  Last detected as dynamic slow peer: never
  Dynamic slow peer recovered: never

  Address tracking is enabled, the RIB does have a route to 10.1.1.3
  Connections established 1; dropped 0
  Last reset never
  Transport(tcp) path-mtu-discovery is enabled
  Graceful-Restart is enabled, restart-time 120 seconds, stalepath-time 360 seconds
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled
Mininum incoming TTL 0, Outgoing TTL 255
Local host: 10.1.1.1, Local port: 179
Foreign host: 10.1.1.3, Foreign port: 48485
Connection tableid (VRF): 0

Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)

Event Timers (current time is 0xE750C):
Timer          Starts    Wakeups            Next
Retrans            18          0             0x0
TimeWait            0          0             0x0
AckHold            22         20             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0
Linger              0          0             0x0
iss: 3196633674  snduna: 3196634254  sndnxt: 3196634254     sndwnd:  15805
irs: 1633793063  rcvnxt: 1633794411  rcvwnd:      15037  delrcvwnd:   1347

SRTT: 273 ms, RTTO: 490 ms, RTV: 217 ms, KRTT: 0 ms
minRTT: 2 ms, maxRTT: 300 ms, ACK hold: 200 ms
Status Flags: passive open, gen tcbs
Option Flags: nagle, path mtu capable

Datagrams (max data segment is 1436 bytes):
Rcvd: 42 (out of order: 0), with data: 24, total data bytes: 1347
Sent: 40 (retransmit: 0 fastretransmit: 0),with data: 19, total data bytes: 579
```

Related Commands

Command	Description
bgp asnotation dot	Changes the default display and the regular expression match format of BGP 4-byte autonomous system numbers from asplain (decimal values) to dot notation.
neighbor send-label	Enables a BGP router to send MPLS labels with BGP routes to a neighboring BGP router.
neighbor send-label explicit-null	Enables a BGP router to send MPLS labels with explicit-null information for a CSC-CE router and BGP routes to a neighboring CSC-PE router.
router bgp	Configures the BGP routing process.

show ip bgp paths

To display all the BGP paths in the database, use the **show ip bgp paths** command in EXEC mode.

show ip bgp paths

Cisco 10000 Series Router

show ip bgp paths *regexp*

Syntax Description

regexp	Regular expression to match the BGP autonomous system paths.

Command Modes

EXEC

Command History

Release	Modification
10.0	This command was introduced.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.
12.0(33)S3	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format is now asplain.
Cisco IOS XE Release 2.4	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format is now asplain.
12.2(33)SRE	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)XNE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.

Examples

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

```
Router# show ip bgp paths

Address    Hash Refcount Metric Path
0x60E5742C    0        1      0 i
0x60E3D7AC    2        1      0 ?
0x60E5C6C0   11        3      0 10 ?
0x60E577B0   35        2     40 10 ?
```

Table 50 describes the significant fields shown in the display.

Table 50 show ip bgp paths Field Descriptions

Field	Description
Address	Internal address where the path is stored.
Hash	Hash bucket where path is stored.
Refcount	Number of routes using that path.
Metric	The Multi Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.)
Path	The autonomous system path for that route, followed by the origin code for that route.

show ip bgp peer-group

To display information about BGP peer groups, use the **show ip bgp peer-group** command in user EXEC or privileged EXEC mode.

show ip bgp peer-group [*peer-group-name*] [**summary**]

Syntax Description

peer-group-name	(Optional) Displays information about a specific peer group.
summary	(Optional) Displays a summary of the status of all the members of a peer group.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
11.0	This command was introduced.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH, and the output was modified to support BGP dynamic neighbors.
15.0(1)S	This command was integrated into Cisco IOS Release 15.0(1)S, with the modified output to support BGP dynamic neighbors.
Cisco IOS XE Release 3.1S	This command was integrated into Cisco IOS XE Release 3.1S, with the modified output to support BGP dynamic neighbors.

Examples

The following is sample output from the **show ip bgp peer-group** command for a peer group named internal in privileged EXEC mode:

```
Router# show ip bgp peer-group internal

BGP peer-group is internal, remote AS 100
  BGP version 4
  Minimum time between advertisement runs is 5 seconds

 For address family:IPv4 Unicast
  BGP neighbor is internal, peer-group internal, members:
           10.1.1.1         10.1.1.2
  Index 3, Offset 0, Mask 0x8
  Incoming update AS path filter list is 53
  Outgoing update AS path filter list is 54
  Route map for incoming advertisements is MAP193
  Route map for outgoing advertisements is MAP194
  Update messages formatted 0, replicated 0
```

The following output from the **show ip bgp peer-group** command shows information about a configured listen range group, group192. In Cisco IOS Release 12.2(33)SXH, 15.0(1)S, and XE Release 3.1S and later releases, the BGP dynamic neighbor feature introduced the ability to support the dynamic creation of BGP neighbor peers using a subnet range associated with a peer group (listen range group).

```
Router# show ip bgp peer-group group192

BGP peer-group is group192,  remote AS 40000
  BGP peergroup group192 listen range group members:
  192.168.0.0/16
  BGP version 4
  Default minimum time between advertisement runs is 30 seconds

 For address family: IPv4 Unicast
  BGP neighbor is group192, peer-group external, members:
  *192.168.3.2
  Index 0, Offset 0, Mask 0x0
  Update messages formatted 0, replicated 0
  Number of NLRIs in the update sent: max 0, min 0
```

show ip bgp quote-regexp

To display routes matching the autonomous system path regular expression, use the **show ip bgp quote-regexp** command in privileged EXEC mode.

show ip bgp quote-regexp *regexp*

Syntax Description

regexp	The regular expression to match the Border Gateway Protocol (BGP) autonomous system paths. • In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, 4-byte autonomous system numbers are supported in the range from 65536 to 4294967295 in asplain notation and in the range from 1.0 to 65535.65535 in asdot notation. • In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, 4-byte autonomous system numbers are supported in the range from 1.0 to 65535.65535 in asdot notation only. For more details about autonomous system number formats, see the **router bgp** command. **Note** The regular expression has to be an exact match.

Command Modes

Privileged EXEC (#)

Command History

Release	Modification
11.1	This command was introduced.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(14)SX	This command was integrated into Cisco IOS Release 12.2(14)SX.
12.0(32)S12	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
12.0(32)SY8	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.4(24)T	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
Cisco IOS XE Release 2.3	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
12.2(33)SXI1	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.0(33)S3	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format is now asplain.

Release	Modification
Cisco IOS XE Release 2.4	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format is now asplain.
12.2(33)SRE	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)XNE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.

Usage Guidelines

In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, the Cisco implementation of 4-byte autonomous system numbers uses asplain—65538 for example—as the default regular expression match and output display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain format and the asdot format as described in RFC 5396. To change the default regular expression match and output display of 4-byte autonomous system numbers to asdot format, use the **bgp asnotation dot** command followed by the **clear ip bgp *** command to perform a hard reset of all current BGP sessions.

In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, the Cisco implementation of 4-byte autonomous system numbers uses asdot—1.2 for example—as the only configuration format, regular expression match, and output display, with no asplain support.

Examples

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

```
Router# show ip bgp quote-regexp "^10_" | begin 10.40

*> 10.40.0.0/20     10.10.10.10                            0 10 2548 1239 10643 i
*> 10.40.16.0/20    10.10.10.10                            0 10 2548 6172 i
*> 10.40.32.0/19    10.10.10.10                            0 10 2548 6172 i
*> 10.41.0.0/19     10.10.10.10                            0 10 2548 3356 3703 ?
*> 10.42.0.0/17     10.10.10.10                            0 10 2548 6172 i
```

Note Although the columns in the above display are not labeled, see Table 51 for detailed information.

Table 51 describes the significant fields shown in the display from left to right.

Table 51 show ip bgp quote-regexp Field Descriptions

Field	Description
Status codes	Status of the table entry; for example, * in the above display. 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 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. r—The table entry failed to install in the routing table. S—The table entry is a stale route.
Network	IP address of a network entity; for example, 24.40.0.0/20 in the above display.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network; for example, 10.10.10.10. in the above display. 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.; for example, 0 in the above display.
LocPrf	Local preference value as set with the **set local-preference** route-map configuration command; for example, 10 in the above display. The default value is 100.
Weight	Weight of the route as set via autonomous system filters; for example, 2548 in the above display.
Path	Autonomous system paths to the destination network; for example, 1239 in the above display. There can be one entry in this field for each autonomous system in the path.
Origin codes	Origin of the entry; for example, ? in the above display. 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 an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.

The following output from the **show ip bgp quote-regexp** command shows routes that match the quoted regular expression for the 4-byte autonomous system number 65550. The 4-byte autonomous system number is displayed in the default asplain format. This example requires Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, or a later release.

```
Router# show ip bgp quote-regexp "^65550$"

BGP table version is 4, local router ID is 172.17.1.99
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.2.2.0/24      192.168.3.2              0             0 65550 i
```

Related Commands

Command	Description
bgp asnotation dot	Changes the default display and the regular expression match format of BGP 4-byte autonomous system numbers from asplain (decimal values) to dot notation.
router bgp	Configures the BGP routing process.
show ip bgp regexp	Displays routes matching the autonomous system path regular expression.

show ip bgp regexp

To display routes matching the autonomous system path regular expression, use the **show ip bgp regexp** command in EXEC mode.

show ip bgp regexp *regexp*

Syntax Description

regexp	Regular expression to match the BGP autonomous system paths. • In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, 4-byte autonomous system numbers are supported in the range from 65536 to 4294967295 in asplain notation and in the range from 1.0 to 65535.65535 in asdot notation. • In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, 4-byte autonomous system numbers are supported in the range from 1.0 to 65535.65535 in asdot notation only. For more details about autonomous system number formats, see the **router bgp** command.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
10.0	This command was introduced.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(14)SX	This command was integrated into Cisco IOS Release 12.2(14)SX.
12.0(32)S12	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
12.0(32)SY8	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.4(24)T	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
Cisco IOS XE Release 2.3	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
12.2(33)SXI1	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.0(33)S3	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format is now asplain.
Cisco IOS XE Release 2.4	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format is now asplain.

Release	Modification
12.2(33)SRE	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)XNE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.

Usage Guidelines

In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, the Cisco implementation of 4-byte autonomous system numbers uses asplain—65538 for example—as the default regular expression match and output display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain format and the asdot format as described in RFC 5396. To change the default regular expression match and output display of 4-byte autonomous system numbers to asdot format, use the **bgp asnotation dot** command followed by the **clear ip bgp *** command to perform a hard reset of all current BGP sessions.

In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, the Cisco implementation of 4-byte autonomous system numbers uses asdot—1.2 for example—as the only configuration format, regular expression match, and output display, with no asplain support.

To ensure a smooth transition we recommend that all BGP speakers within an autonomous system that is identified using a 4-byte autonomous system number, are upgraded to support 4-byte autonomous system numbers.

Examples

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

```
Router# show ip bgp regexp 108$

BGP table version is 1738, local router ID is 172.16.72.24
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*  172.16.0.0       172.16.72.30                         0 109 108 ?
*  172.16.1.0       172.16.72.30                         0 109 108 ?
*  172.16.11.0      172.16.72.30                         0 109 108 ?
*  172.16.14.0      172.16.72.30                         0 109 108 ?
*  172.16.15.0      172.16.72.30                         0 109 108 ?
*  172.16.16.0      172.16.72.30                         0 109 108 ?
*  172.16.17.0      172.16.72.30                         0 109 108 ?
*  172.16.18.0      172.16.72.30                         0 109 108 ?
*  172.16.19.0      172.16.72.30                         0 109 108 ?
*  172.16.24.0      172.16.72.30                         0 109 108 ?
*  172.16.29.0      172.16.72.30                         0 109 108 ?
*  172.16.30.0      172.16.72.30                         0 109 108 ?
*  172.16.33.0      172.16.72.30                         0 109 108 ?
*  172.16.35.0      172.16.72.30                         0 109 108 ?
*  172.16.36.0      172.16.72.30                         0 109 108 ?
*  172.16.37.0      172.16.72.30                         0 109 108 ?
*  172.16.38.0      172.16.72.30                         0 109 108 ?
*  172.16.39.0      172.16.72.30                         0 109 108 ?
```

The following example requires Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, or a later release. After the **bgp asnotation dot** command is configured, the regular expression match format for 4-byte autonomous system paths is changed to asdot notation format. Although a 4-byte autonomous system number can be configured in a

regular expression using either asplain or asdot format, only 4-byte autonomous system numbers configured using the current default format are matched. In the first example, the **show ip bgp regexp** command is configured with a 4-byte autonomous system number in asplain format. The match fails because the default format is currently asdot format and there is no output. In the second example using asdot format, the match passes and the information about the 4-byte autonomous system path is shown using the asdot notation.

Note The asdot notation uses a period which is a special character in Cisco regular expressions. to remove the special meaning, use a backslash before the period.

```
Router# show ip bgp regexp ^65536$


Router# show ip bgp regexp ^1\.0$

BGP table version is 2, local router ID is 172.17.1.99
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.1.0/24      192.168.1.2              0             0 1.0 i
```

The following is sample output from the **show ip bgp regexp** command after the **bgp asnotation dot** command has been entered to display 4-byte autonomous system numbers in dot notation in Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, or later release. The dot notation is the only format for 4-byte autonomous system numbers in Cisco IOS Release 12.0(32)S12, 12.4(24)T, or Cisco IOS XE Release 2.3.

Note The asdot notation uses a period which is a special character in Cisco regular expressions. to remove the special meaning, use a backslash before the period.

```
Router# show ip bgp regexp ^1\.14$

BGP table version is 4, local router ID is 172.17.1.99
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.1.0/24      192.168.1.2              0             0 1.14   i
```

Related Commands

Command	Description
bgp asnotation dot	Changes the default display and the regular expression match format of BGP 4-byte autonomous system numbers from asplain (decimal values) to dot notation.
router bgp	Configures the BGP routing process.
show ip bgp quote-regexp	Displays routes matching the autonomous system path regular expression.

show ip bgp replication

To display update replication statistics for Border Gateway Protocol (BGP) update groups, use the **show ip bgp replication** command in EXEC mode.

show ip bgp replication [*index-group* | *ip-address*]

Syntax Description

index-group	(Optional) Displays update replication statistics for the update group with the corresponding index number. The range of update-group index numbers is from 1 to 4294967295.
ip-address	(Optional) Displays update replication statistics for this neighbor.

Command Modes

EXEC

Command History

Release	Modification
12.0(24)S	This command was introduced.
12.2(18)S	This command was integrated into Cisco IOS Release 12.2(18)S.
12.3(4)T	This command was integrated into Cisco IOS Release 12.3(4)T.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

The output of this command displays BGP update-group replication statistics.

When a change to outbound policy occurs, the router automatically recalculates update-group memberships and applies the changes by triggering an outbound soft reset after a 3-minute timer expires. This behavior is designed to provide the network operator with time to change the configuration if a mistake is made. You can manually enable an outbound soft reset before the timer expires by entering the **clear ip bgp** *ip-address* **soft out** command.

Examples

The following sample output from the **show ip bgp replication** command shows update-group replication information for all neighbors:

```
Router# show ip bgp replication

BGP Total Messages Formatted/Enqueued : 0/0

     Index     Type  Members          Leader   MsgFmt  MsgRepl  Csize  Qsize
         1 internal        1       10.4.9.21        0        0      0      0
         2 internal        2        10.4.9.5        0        0      0      0
```

The following sample output from the **show ip bgp replication** command shows update-group statistics for the 10.4.9.5 neighbor:

```
Router# show ip bgp replication 10.4.9.5

     Index     Type  Members          Leader   MsgFmt  MsgRepl  Csize  Qsize
         2 internal        2        10.4.9.5        0        0      0      0
```

Table 52 describes the significant fields shown in the display.

Table 52 show ip bgp replication Field Descriptions

Field	Description
Index	Index number of the update group.
Type	Type of peer (internal or external).
Members	Number of members in the dynamic update peer group.
Leader	First member of the dynamic update peer group.

Related Commands

Command	Description
clear ip bgp	Resets a BGP connection or session.
clear ip bgp update-group	Clears BGP update-group member sessions.
debug ip bgp groups	Displays information related to the processing of BGP update groups.
show ip bgp peer-group	Displays information about BGP update groups.

show ip bgp rib-failure

To display Border Gateway Protocol (BGP) routes that failed to install in the Routing Information Base (RIB) table, use the **show ip bgp rib-failure** command in privileged EXEC mode.

show ip bgp rib-failure

Syntax Description

This command has no keywords or arguments.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.3	This command was introduced.
12.0(26)S	This command was integrated into Cisco IOS Release 12.0(26)S.
12.2(25)S	This command was integrated into Cisco IOS Release 12.2(25)S.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.

Examples

The following is sample output from the **show ip bgp rib-failure** command:

```
Router# show ip bgp rib-failure

Network           Next Hop                    RIB-failure   RIB-NH Matches
10.1.15.0/24      10.1.35.5          Higher admin distance             n/a
10.1.16.0/24      10.1.15.1          Higher admin distance             n/a
```

Table 53 describes the significant fields shown in the display.

Table 53 show ip bgp rib-failure Field Descriptions

Field	Description
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.

Table 53 show ip bgp rib-failure Field Descriptions (continued)

Field	Description
RIB-failure	Cause of RIB failure. Higher admin distance means that a route with a better (lower) administrative distance such as a static route already exists in the IP routing table.
RIB-NH Matches	Route status that applies only when Higher admin distance appears in the RIB-failure column and **bgp suppress-inactive** is configured for the address family being used. There are three choices: • Yes—Means that the route in the RIB has the same next hop as the BGP route or next hop recurses down to the same adjacency as the BGP nexthop. • No—Means that the next hop in the RIB recurses down differently from the next hop of the BGP route. • n/a—Means that **bgp suppress-inactive** is not configured for the address family being used.

Related Commands

Command	Description
bgp suppress-inactive	Configures a router to suppress the advertisement of BGP routes that are not installed in the RIB and FIB tables.
clear ip bgp	Resets a BGP connection or session.
neighbor soft-reconfiguration	Configures the Cisco IOS software to start storing updates.

show ip bgp rtfilter

To display information about BGP route target (RT) filtering, use the **show ip bgp rtfilter** command in user EXEC or privileged EXEC mode.

show ip bgp rtfilter unicast {**all** | **default** | **rt** {*ASN*:*nn* | *ip-address*:*nn*}}

Syntax Description

unicast	Display unicast information.
all	Display RT information for all VPNs.
default	Display the default RT filter.
rt	Display a specific RT filter prefix.
ASN:*nn*	Autonomous system number, followed by a colon and number.
ip-address:*nn*	IP address, followed by a colon and a number.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
15.1(1)S	This command was introduced.
Cisco IOS XE Release 3.2S	This command was integrated into Cisco IOS XE Release 3.2S.

Usage Guidelines

Use this command if you have configured the BGP: RT Constrained Route Distribution feature and you want to display RT filter information.

Note If you enter the **all** keyword, there are many more optional keywords available that are not shown here.

Examples

The following is sample output from the **show ip bgp rtfilter unicast all** command:

```
Router# show ip bgp rtfilter unicast all

BGP table version is 14, local router ID is 192.168.7.7
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external, f
RT-Filter
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network Next HopMetricLocPrf Weight Path
*>i0:0:0:0192.168.2.201000 i
*>i1:2:1:100192.168.6.601000 i
* i1:2:3:3192.168.2.201000 i
*>  0.0.0.0    32768 i
*>i1:2:150:1192.168.6.601000 i
* i1:2:200:200192.168.2.201000 i
*>  0.0.0.0    32768 i
Router#
```

Table 54 describes the fields shown in the display.

Table 54 show ip bgp rtfilter Field Descriptions

Field	Description
Network	RT filter prefix.
Next Hop	Next hop in the RT filter prefix.
Metric	BGP metric associated with the RT filter prefix.
LocPref	BGP local preference.
Weight	BGP weight.
Path	Path information associated with the RT prefix.

The following is sample output from the **show ip bgp rtfilter all summary** command:

```
Router# show ip bgp rtfilter all summary

BGP router identifier 192.168.7.7, local AS number 1
BGP table version is 14, main routing table version 14
5 network entries using 820 bytes of memory
7 path entries using 336 bytes of memory
2/2 BGP path/bestpath attribute entries using 256 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
2 BGP extended community entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1484 total bytes of memory
BGP activity 7/0 prefixes, 14/5 paths, scan interval 60 secs

NeighborVASMsgRcvdMsgSentTblVerInQOutQUp/Down  State/PfxRcd
192.168.2.2411312140   0   00:03:21        5
Router#
```

Related Commands

Command	Description
address-family rtfilter unicast	Enters address family configuration mode and enables Automated Route Target Filtering with a BGP peer.
neighbor default-originate	Allows a BGP speaker (the local router) to send the default route 0.0.0.0 to a neighbor for use as a default route.
show ip bgp rtfilter all summary	Displays summary information about RT filtering.

show ip bgp summary

To display the status of all Border Gateway Protocol (BGP) connections, use the **show ip bgp summary** command in user EXEC or privileged EXEC mode.

show ip bgp [**ipv4** {**multicast** | **unicast**} | **vpnv4 all** | **vpnv6 unicast all** | **topology**{*|*routing-topology-instance-name*}] [**update-group**] **summary** [**slow**]

Syntax Description

ipv4 {**multicast**	**unicast**}	(Optional) Displays peers in the IPv4 address family.
vpnv4 all	(Optional) Displays peers in the VPNv4 address family.
vpnv6 unicast all	(Optional) Displays peers in the VPNv6 address family.
topology	(Optional) Displays routing topology information.
*	(Optional) Displays allrouting topology instances.
routing-topology-instance-name	(Optional) Displays routing topology information for that instance.
update-group	(Optional) Includes information about the update group of the peers.
slow	(Optional) Displays only information about dynamically configured slow peers.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
10.0	This command was introduced.
12.0	Support for the **neighbor maximum-prefix** command was added to the output.
12.2	• The number of networks and paths displayed in the output was split out to two separate lines. • A field was added to display multipath entries in the routing table.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.4(11)T	A line was added to the output to display the advertised bitfield cache entries and associated memory usage.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH, and the output was modified to support BGP dynamic neighbors.
12.0(32)S12	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
12.0(32)SY8	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.4(24)T	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
Cisco IOS XE Release 2.3	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.

Release	Modification
12.2(33)SXI1	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.0(33)S3	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format is now asplain.
Cisco IOS XE Release 2.4	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format is now asplain.
12.2(33)SRE	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)XNE	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
15.0(1)S	This command was modified. The **slow** keyword was added.
Cisco IOS XE Release 3.1S	This command was modified. The **slow** keyword was added.

Usage Guidelines

The **show ip bgp summary** command is used to display BGP path, prefix, and attribute information for all connections to BGP neighbors.

A prefix is an IP address and network mask. It can represent an entire network, a subset of a network, or a single host route. A path is a route to a given destination. By default, BGP will install only a single path for each destination. If multipath routes are configured, BGP will install a path entry for each multipath route, and only one multipath route will be marked as the bestpath.

BGP attribute and cache entries are displayed individually and in combinations that affect the bestpath selection process. The fields for this output are displayed when the related BGP feature is configured or attribute is received. Memory usage is displayed in bytes.

In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, the Cisco implementation of 4-byte autonomous system numbers uses asplain—65538 for example—as the default regular expression match and output display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain format and the asdot format as described in RFC 5396. To change the default regular expression match and output display of 4-byte autonomous system numbers to asdot format, use the **bgp asnotation dot** command followed by the **clear ip bgp *** command to perform a hard reset of all current BGP sessions.

In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, the Cisco implementation of 4-byte autonomous system numbers uses asdot—1.2 for example—as the only configuration format, regular expression match, and output display, with no asplain support.

Examples

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

```
Router# show ip bgp summary

BGP router identifier 172.16.1.1, local AS number 100
BGP table version is 199, main routing table version 199
37 network entries using 2850 bytes of memory
59 path entries using 5713 bytes of memory
18 BGP path attribute entries using 936 bytes of memory
```

```
2 multipath network entries and 4 multipath paths
10 BGP AS-PATH entries using 240 bytes of memory
7 BGP community entries using 168 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
90 BGP advertise-bit cache entries using 1784 bytes of memory
36 received paths for inbound soft reconfiguration
BGP using 34249 total bytes of memory
Dampening enabled. 4 history paths, 0 dampened paths
BGP activity 37/2849 prefixes, 60/1 paths, scan interval 15 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down State/PfxRcd
10.100.1.1      4   200      26      22      199    0    0 00:14:23 23
10.200.1.1      4   300      21      51      199    0    0 00:13:40 0
```

Table 55 describes the significant fields shown in the display. Fields that are preceded by the asterisk character (*) are not shown in the above output.

Table 55 show ip bgp summary Field Descriptions

Field	Description
BGP router identifier	In order of precedence and availability, the router identifier specified by the **bgp router-id** command, a loopback address, or the highest IP address.
BGP table version	Internal version number of BGP database.
main routing table version	Last version of BGP database that was injected into the main routing table.
...network entries	Number of unique prefix entries in the BGP database.
...using ... bytes of memory	Amount of memory, in bytes, that is consumed for the path, prefix, or attribute entry displayed on the same line.
...path entries using	Number of path entries in the BGP database. Only a single path entry will be installed for a given destination. If multipath routes are configured, a path entry will be installed for each multipath route.
...multipath network entries using	Number of multipath entries installed for a given destination.
* ...BGP path/bestpath attribute entries using	Number of unique BGP attribute combinations for which a path is selected as the bestpath.
* ...BGP rrinfo entries using	Number of unique ORIGINATOR and CLUSTER_LIST attribute combinations.
...BGP AS-PATH entries using	Number of unique AS_PATH entries.
...BGP community entries using	Number of unique BGP community attribute combinations.
*...BGP extended community entries using	Number of unique extended community attribute combinations.
BGP route-map cache entries using	Number of BGP route-map match and set clause combinations. A value of 0 indicates that the route cache is empty.
...BGP filter-list cache entries using	Number of filter-list entries that match an AS-path access list permit or deny statements. A value of 0 indicates that the filter-list cache is empty.

Table 55 show ip bgp summary Field Descriptions (continued)

Field	Description
BGP advertise-bit cache entries using	(Cisco IOS Release 12.4(11)T and later releases only) Number of advertised bitfield entries and the associated memory usage. A bitfield entry represents a piece of information (one bit) that is generated when a prefix is advertised to a peer. The advertised bit cache is built dynamically when required.
...received paths for inbound soft reconfiguration	Number paths received and stored for inbound soft reconfiguration.
BGP using...	Total amount of memory, in bytes, used by the BGP process.
Dampening enabled...	Indicates that BGP dampening is enabled. The number of paths that carry an accumulated penalty and the number of dampened paths are displayed on this line.
BGP activity...	Displays the number of times that memory has been allocated or released for a path or prefix.
Neighbor	IP address of the neighbor.
V	BGP version number spoken to the neighbor.
AS	Autonomous system number.
MsgRcvd	Number of messages received from the neighbor.
MsgSent	Number of messages sent to the neighbor.
TblVer	Last version of the BGP database that was sent to the neighbor.
InQ	Number of messages queued to be processed from the neighbor.
OutQ	Number of messages queued to be sent to the neighbor.
Up/Down	The length of time that the BGP session has been in the Established state, or the current status if not in the Established state.
State/PfxRcd	Current state of the BGP session, and the number of prefixes that have been received from a neighbor or peer group. When the maximum number (as set by the **neighbor maximum-prefix** command) is reached, the string "PfxRcd" appears in the entry, the neighbor is shut down, and the connection is set to Idle. An (Admin) entry with Idle status indicates that the connection has been shut down using the **neighbor shutdown** command.

The following output from the **show ip bgp summary** command shows that the BGP neighbor 192.168.3.2 was dynamically created and is a member of the listen range group, group192. The output also shows that the IP prefix range of 192.168.0.0/16 is defined for the listen range group named group192. In Cisco IOS Release 12.2(33)SXH and later releases, the BGP dynamic neighbor feature introduced the ability to support the dynamic creation of BGP neighbor peers using a subnet range associated with a peer group (listen range group).

```
Router# show ip bgp summary

BGP router identifier 192.168.3.1, local AS number 45000
BGP table version is 1, main routing table version 1
```

```
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
*192.168.3.2    4 50000       2       2        0    0    0 00:00:37        0
* Dynamically created based on a listen range command
Dynamically created neighbors: 1/(200 max), Subnet ranges: 1

BGP peergroup group192 listen range group members:
  192.168.0.0/16
```

The following output from the **show ip bgp summary** command shows two BGP neighbors, 192.168.1.2 and 192.168.3.2, in different 4-byte autonomous system numbers, 65536 and 65550. The local autonomous system 65538 is also a 4-byte autonomous system number and the numbers are displayed in the default asplain format. This example requires Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, or a later release.

```
Router# show ip bgp summary

BGP router identifier 172.17.1.99, local AS number 65538
BGP table version is 1, main routing table version 1

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  Statd
192.168.1.2     4       65536       7       7        1    0    0 00:03:04      0
192.168.3.2     4       65550       4       4        1    0    0 00:00:15      0
```

The following output from the **show ip bgp summary** command shows the same two BGP neighbors, but the 4-byte autonomous system numbers are displayed in asdot notation format. To change the display format the **bgp asnotation dot** command must be configured in router configuration mode. This example requires Cisco IOS Release 12.0(32)SY8, 12.0(32)S12, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, 12.4(24)T, or Cisco IOS XE Release 2.3 or later releases.

```
Router# show ip bgp summary

BGP router identifier 172.17.1.99, local AS number 1.2
BGP table version is 1, main routing table version 1

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  Statd
192.168.1.2     4         1.0       9       9        1    0    0 00:04:13      0
192.168.3.2     4        1.14       6       6        1    0    0 00:01:24      0
```

The following example displays sample output of the **show ip bgp summary slow** command:

```
Router> show ip bgp summary slow
BGP router identifier 2.2.2.2, local AS number 100
BGP table version is 37, main routing table version 37
36 network entries using 4608 bytes of memory
36 path entries using 1872 bytes of memory
1/1 BGP path/bestpath attribute entries using 124 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
1 BGP extended community entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 6700 total bytes of memory
BGP activity 46/0 prefixes, 48/0 paths, scan interval 60 secs

Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
6.6.6.6 4 100 11 10 1 0 0 00:44:20 0
```

Related Commands

Command	Description
bgp asnotation dot	Changes the default display and the regular expression match format of BGP 4-byte autonomous system numbers from asplain (decimal values) to dot notation.
bgp router-id	Configures a fixed router ID for the local BGP routing process.
neighbor maximum-prefix	Controls how many prefixes can be received from a BGP neighbor.
neighbor shutdown	Disables a BGP neighbor or peer group.
neighbor slow-peer split-update-group dynamic	Causes a dynamically detected slow peer to be moved to a slow update group.
router bgp	Configures the BGP routing process.

show ip bgp template peer-policy

To display locally configured peer policy templates, use the **show ip bgp template peer-policy** command in user EXEC or privileged EXEC mode.

show ip bgp template peer-policy [*policy-template-name* [**detail**]]

Syntax Description

policy-template-name	(Optional) Name of a locally configured peer policy template.
detail	(Optional) Displays detailed policy information such as route maps, prefix lists, community lists, access control lists (ACLs), and AS-path filter lists.

Command Default

If a peer policy template is not specified using the *policy-template-name* argument, all peer policy templates will be displayed.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.0(24)S	This command was introduced.
12.0(25)S	The **detail** keyword was added.
12.2(18)S	This command was integrated into Cisco IOS Release 12.2(18)S.
12.3(4)T	This command was integrated into Cisco IOS Release 12.3(4)T.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.4(11)T	Support for the **detail** keyword was integrated into Cisco IOS Release 12.4(11)T.
12.2(33)SRB	This command and support for the **detail** keyword were integrated into Cisco IOS Release 12.2(33)SRB.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.
12.2(33)SB	Support for the **detail** keyword was integrated into Cisco IOS Release 12.2(33)SB.

Usage Guidelines

This command is used to display locally configured peer policy templates. The output can be filtered to display a single peer policy template using the *policy-template-name* argument. This command also supports all standard output modifiers.

When BGP neighbors use multiple levels of peer templates it can be difficult to determine which policies are associated with a specific template. In Cisco IOS Release 12.0(25)S, 12.4(11)T, 12.2(33)SRB, 12.2(33)SB, and later releases, the **detail** keyword was added to display the detailed configuration of local and inherited policies associated with a specific template. Inherited policies are policies that the template inherits from other peer-policy templates.

Examples

The **show ip bgp template peer-policy** command is used to verify the configuration of local peer policy templates. The following sample output shows the peer policy templates named GLOBAL and NETWORK1. The output also shows that the GLOBAL template was inherited by the NETWORK1 template.

```
Router# show ip bgp template peer-policy

Template:GLOBAL, index:1.
Local policies:0x80840, Inherited polices:0x0
 *Inherited by Template NETWORK1, index:2
Locally configured policies:
  prefix-list NO-MARKETING in
  weight 300
  maximum-prefix 10000
Inherited policies:

Template:NETWORK1, index:2.
Local policies:0x1, Inherited polices:0x80840
This template inherits:
  GLOBAL, index:1, seq_no:10, flags:0x1
Locally configured policies:
  route-map ROUTE in
Inherited policies:
  prefix-list NO-MARKETING in
  weight 300
  maximum-prefix 10000
```

Table 56 describes the significant fields shown in the display.

Table 56 show ip bgp template peer-policy Field Descriptions

Field	Description
Template	Name of the peer template.
index	The sequence number in which the displayed template is processed.
Local policies	Displays the hexadecimal value of locally configured policies.
Inherited polices	Displays the hexadecimal value of inherited policies. The 0x0 value is displayed when no templates are inherited.
Locally configured policies	Displays a list of commands that are locally configured in a peer policy template.
Inherited policies	Displays a list of commands that are inherited from a peer template.

The following sample output of the **show ip bgp template peer-policy** command with the **detail** keyword displays details of the template named NETWORK1, which includes the inherited template named GLOBAL. The output in this example displays the configuration commands of the locally configured route map and prefix list and the inherited prefix list.

```
Router# show ip bgp template peer-policy NETWORK1 detail

Template:NETWORK1, index:2.
Local policies:0x1, Inherited polices:0x80840
This template inherits:
  GLOBAL, index:1, seq_no:10, flags:0x1
Locally configured policies:
  route-map ROUTE in
```

```
Inherited policies:
  prefix-list NO-MARKETING in
  weight 300
  maximum-prefix 10000

Template:NETWORK1 <detail>
Locally configured policies:
  route-map ROUTE in
route-map ROUTE, permit, sequence 10
  Match clauses:
    ip address prefix-lists: DEFAULT
ip prefix-list DEFAULT: 1 entries
   seq 5 permit 10.1.1.0/24

  Set clauses:
  Policy routing matches: 0 packets, 0 bytes

Inherited policies:
  prefix-list NO-MARKETING in
ip prefix-list NO-MARKETING: 1 entries
   seq 5 deny 10.2.2.0/24
```

Related Commands

Command	Description
inherit peer-policy	Configures a peer policy template to inherit the configuration from another peer policy template.
template peer-policy	Creates a peer policy template and enters policy-template configuration mode.

show ip bgp template peer-session

To display peer policy template configurations, use the **show ip bgp template peer-session** command in user EXEC and privileged EXEC mode.

show ip bgp template peer-session [*session-template-name*]

Syntax Description

session-template-name	(Optional) Name of a locally configured peer session template.

Defaults

If a peer session template is not specified with the *session-template-name* argument, all peer session templates will be displayed.

Command Modes

User EXEC
Privileged EXEC

Command History

Release	Modification
12.0(24)S	This command was introduced.
12.3(4)T	This command was integrated into Cisco IOS Release 12.3(4)T.
12.2(18)S	This command was integrated into Cisco IOS Release 12.2(18)S.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(31)SB	This command was integrated into Cisco IOS Release 12.2(31)SB.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

This command is used to display locally configured peer session templates. The output can be filtered to display a single peer session template with the *peer-session-name* argument. This command also supports all standard output modifiers.

Examples

The **show ip bgp template peer-session** command is used to verify the configuration of local peer session templates. The following example shows the peer session templates named INTERNAL-BGP and CORE1. The output also shows that INTERNAL-BGP is inherited by CORE1.

```
Router# show ip bgp template peer-session

Template:INTERNAL-BGP, index:1
Local policies:0x21, Inherited polices:0x0
 *Inherited by Template CORE1, index= 2
Locally configured session commands:
 remote-as 202
 timers 30 300
Inherited session commands:
```

```
Template:CORE1, index:2
Local policies:0x180, Inherited polices:0x21
This template inherits:
  INTERNAL-BGP index:1 flags:0x0
Locally configured session commands:
 update-source loopback 1
 description CORE-123
Inherited session commands:
 remote-as 202
 timers 30 300
```

Table 57 describes the significant fields shown in the display.

Table 57 show ip bgp template peer-session Field Descriptions

Field	Description
Template:	Name of the peer template.
index:	The sequence number in which the displayed template is processed.
Local policies:	Displays the hexadecimal value of locally configured policies.
Inherited polices:	Displays the hexadecimal value of inherited policies. The 0x0 value is displayed when no templates are inherited.
Locally configured session commands:	Displays a list of commands that are locally configured in a peer template.
Inherited session commands:	Displays a list of commands that are inherited from a peer session template.

Related Commands

Command	Description
inherit peer-session	Configures a peer session template to inherit the configuration from another peer session template.
template peer-session	Creates a peer session template and enters session-template configuration mode.

show ip bgp unicast route-server

To display on a BGP route server which paths are chosen for a route server context, in particular if the normal bestpath was overridden or suppressed, use the **show ip bgp unicast route-server** command in privileged EXEC mode.

show ip bgp {**ipv4** | **ipv6**} **unicast route-server** {**all** | **context** *context-name*} [**summary**]

Syntax Description

ipv4	Displays only IPv4 prefixes.
ipv6	Displays only IPv6 prefixes.
all	Displays information for all route server contexts.
context *context-name*	Displays information for the specified route server context only.
summary	(Optional) Displays the neighbor state for route server clients.

Command Modes

Privileged EXEC (#)

Command History

Release	Modification
Cisco IOS XE 3.3S	This command was introduced.

Usage Guidelines

Use this command on a BGP route server to see the next hop to network prefixes and additional information about the path.

Examples

The following output displays all the routes chosen by the policy for the context named example-context:

```
Route-Server# show ip bgp ipv4 unicast route-server context example-context

Networks for route server context example-context:
   Network          Next Hop            Metric LocPrf Weight Path
*  1.1.1.1/32       10.10.10.22            123             0 22 ?
*  1.1.2.0/24       10.10.10.22            123             0 22 ?
*  1.3.0.0/16       10.10.10.22            123             0 22 ?
*  8.8.0.0/16       10.10.10.22            123             0 22 ?
   100.100.100.21/32 (suppressed)
*> 100.100.100.22/32 10.10.10.22           123             0 22 ?
*  100.100.100.23/32 10.10.10.23           123             0 23 ?
*> 100.100.100.24/32 10.10.10.24           123             0 24 ?
*> 100.100.100.25/32 10.10.10.25           123             0 25 ?
*> 100.100.100.26/32 10.10.10.26           123             0 26 ?
```

Three types of routes can be in a context, as shown in the preceding output. They are:

- Those where the policy for the context chooses the same path as the regular BGP best path algorithm (for example, 100.100.100.25/32, denoted by ">").
- Those where the policy for the context excluded the regular best path, but found a suitable alternative path to advertise to the client (for example, 1.1.1.1/32, not denoted with ">", but still valid "*").

- Those where the policy for the context excluded all available paths and therefore those routes will not be sent to the client; for example, 100.100.100.21/32, denoted by "(suppressed)".

In the following example, specifying **all** instead of a specific context reveals that different contexts may have differing routes due to the configured policy:

```
Route-Server# show ip bgp ipv4 unicast route-server all

Networks for route server context all-base:
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.1.1/32       10.10.10.21             23             0 21 ?
*> 1.1.2.0/24       10.10.10.21             23             0 21 ?
*> 1.3.0.0/16       10.10.10.21             23             0 21 ?
*> 8.8.0.0/16       10.10.10.21             23             0 21 ?
*> 100.100.100.21/32 10.10.10.21            23             0 21 ?
*> 100.100.100.22/32 10.10.10.22           123             0 22 ?
*> 100.100.100.23/32 10.10.10.21            23             0 21 ?
*  100.100.100.24/32 10.10.10.24           123             0 24 ?
*> 100.100.100.25/32 10.10.10.25           123             0 25 ?
*> 100.100.100.26/32 10.10.10.26           123             0 26 ?

Networks for route server context all-policy-deny:
   Network          Next Hop            Metric LocPrf Weight Path
   1.1.1.1/32       (suppressed)
   1.1.2.0/24       (suppressed)
   1.3.0.0/16       (suppressed)
   8.8.0.0/16       (suppressed)
   100.100.100.21/32 (suppressed)
   100.100.100.22/32 (suppressed)
   100.100.100.23/32 (suppressed)
   100.100.100.24/32 (suppressed)
   100.100.100.25/32 (suppressed)
   100.100.100.26/32 (suppressed)

Networks for route server context all-policy:
   Network          Next Hop            Metric LocPrf Weight Path
*  1.1.1.1/32       10.10.10.27             878             0 27 ?
*  1.1.2.0/24       10.10.10.27             878             0 27 ?
*  1.3.0.0/16       10.10.10.27             878             0 27 ?
*  8.8.0.0/16       10.10.10.27             878             0 27 ?
*  100.100.100.21/32 10.10.10.27            878             0 27 ?
*  100.100.100.22/32 10.10.10.27            878             0 27 ?
*  100.100.100.23/32 10.10.10.27            878             0 27 ?
*  100.100.100.24/32 10.10.10.27            878             0 27 ?
*  100.100.100.25/32 10.10.10.27            878             0 27 ?
*  100.100.100.26/32 10.10.10.27            878             0 27 ?

Networks for route server context example-context:
   Network          Next Hop            Metric LocPrf Weight Path
*  1.1.1.1/32       10.10.10.23             123             0 23 ?
*  1.1.2.0/24       10.10.10.23             123             0 23 ?
*  1.3.0.0/16       10.10.10.23             123             0 23 ?
*  8.8.0.0/16       10.10.10.23             123             0 23 ?
   100.100.100.21/32 (suppressed)
*> 100.100.100.22/32 10.10.10.22           123             0 22 ?
*  100.100.100.23/32 10.10.10.23           123             0 23 ?
*  100.100.100.24/32 10.10.10.24           123             0 24 ?
*> 100.100.100.25/32 10.10.10.25           123             0 25 ?
*> 100.100.100.26/32 10.10.10.26           123             0 26 ?
```

In the following example, the **summary** keyword displays output similar to the **show ip bgp summary** command in that it shows the neighbor state for route server clients in the specified context (or all contexts):

```
Route-Server# show ip bgp ipv4 unicast route-server context example-context summary

Route server clients assigned to context example-context:
Neighbor     V   AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.10.10.18  4   18     283     291       13    0    0 04:13:21        0
```

In the following example, the **all** keyword and the **summary** keyword display summary output for all contexts:

```
Route-Server# show ip bgp ipv4 unicast route-server all summary

Route server clients without assigned contexts:
Neighbor     V   AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.10.10.12  4   12      12      17       12    0    0 00:08:29        0

Route server clients assigned to context all-base:
Neighbor     V   AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.10.10.14  4   14      12      17       12    0    0 00:08:25        0

Route server clients assigned to context all-policy-deny:
Neighbor     V   AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.10.10.16  4   16      12      13       12    0    0 00:08:24        0

Route server clients assigned to context all-policy:
Neighbor     V   AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.10.10.13  4   13      11      14       12    0    0 00:08:22        0

Route server clients assigned to context example-context:
Neighbor     V   AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.10.10.18  4   18      12      17       12    0    0 00:08:30        0
```

Related Commands

Command	Description
neighbor route-server-client	Specifies on a BGP route server that a neighbor is a route server client.

show ip bgp update-group

To display information about the Border Gateway Protocol (BGP) update groups, use the **show ip bgp update-group** command in user EXEC or privileged EXEC mode.

show ip bgp update-group [*index-group* | *ip-address* | *ipv6-address*] [**summary**]

Syntax Description

index-group	(Optional) Update group type with its corresponding index number. The range of update-group index numbers is from 1 to 4294967295.
ip-address	(Optional) IP address of a single neighbor who is a member of an update group.
ipv6-address	(Optional) IPv6 address of a single neighbor who is member of an update group.
summary	(Optional) Displays a summary of update-group member information. The output can be filtered to show information for a single index group or peer with the *index-group*, *ip-address*, or *ipv6-address* argument.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.0(24)S	This command was introduced.
12.2(18)S	This command was integrated into Cisco IOS Release 12.2(18)S. The *ipv6-address* argument was added.
12.3(4)T	This command was integrated into Cisco IOS Release 12.3(4)T.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

Use this command to display information about BGP update groups. When a change to BGP outbound policy occurs, the router automatically recalculates update group memberships and applies the changes by triggering an outbound soft reset after a 1-minute timer expires. This behavior is designed to provide the network operator with time to change the configuration if a mistake is made. You can manually enable an outbound soft reset before the timer expires by entering the **clear ip bgp** *ip-address* **soft out** command.

Note In Cisco IOS Release 12.0(25)S, 12.3(2)T, and prior releases, the update group recalculation delay timer is set to 3 minutes.

Examples

The following sample output from the **show ip bgp update-group** command shows update group information for all neighbors:

```
Router# show ip bgp update-group

BGP version 4 update-group 1, internal, Address Family: IPv4 Unicast
  BGP Update version : 0, messages 0/0
  Route map for outgoing advertisements is COST1
  Update messages formatted 0, replicated 0
  Number of NLRIs in the update sent: max 0, min 0
  Minimum time between advertisement runs is 5 seconds
  Has 1 member:
  10.4.9.21

BGP version 4 update-group 2, internal, Address Family: IPv4 Unicast
  BGP Update version : 0, messages 0/0
  Update messages formatted 0, replicated 0
  Number of NLRIs in the update sent: max 0, min 0
  Minimum time between advertisement runs is 5 seconds
  Has 2 members:
  10.4.9.5 10.4.9.8
```

Table 58 describes the significant fields shown in the display.

Table 58 show ip bgp update-group Field Descriptions

Field	Description
BGP version	BGP version.
update-group	Update-group number and type (internal or external).
Update messages formatted..., replicated...	Number of update messages that have been formatted and replicated.
Number of NLRIs...	NLRI information sent in update.
Minimum time between...	Minimum time, in seconds, between update advertisements.
Has...	Number of member listed by IP address in the update group.

The following sample output from the **show ip bgp update-group** command shows a summary of update-group information for the 10.4.9.8 neighbor:

```
Router# show ip bgp update-group 10.4.9.8 summary

Summary for Update-group 2 :
------------------------------
BGP router identifier 10.4.9.4, local AS number 101
BGP table version is 1, main routing table version 1

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.4.9.5        4   101      35      35        1    0    0 00:26:22        0
10.4.9.8        4   101      39      39        1    0    0 00:26:21        0
```

Table 59 describes the significant fields shown in the display.

Table 59 show ip bgp update-group summary Field Descriptions

Field	Description
Summary for Update-group...	Update-group number.
BGP router identifier...	IP address and AS number for specified peer.

Table 59 show ip bgp update-group summary Field Descriptions (continued)

Field	Description
update messages formatted..., replicated...	Number of update messages that have been formatted and replicated.
BGP table version...	Displays incremental changes in the BGP routing table.
Neighbor...	Specific peer information and statistics, including IP address and AS number.

Related Commands

Command	Description
clear ip bgp	Resets a BGP connection or session.
clear ip bgp update-group	Clears BGP update-group member sessions.
debug ip bgp groups	Displays information related to the processing of BGP update groups.
show ip bgp replication	Displays BGP update-group replication statistics.

show ip bgp vpnv4 all sso summary

To display information about Border Gateway Protocol (BGP) peers that support BGP nonstop routing (NSR) with stateful switchover (SSO), use the **show ip bgp vpn4 sso summary** command in privileged EXEC mode.

show ip bgp vpnv4 all sso summary

Syntax Description

This command has no arguments or keywords.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.2(28)SB	This command was introduced.
15.0(1)S	This command was integrated into Cisco IOS Release 15.0(1)S.
Cisco IOS XE 3.1S	This command was integrated into Cisco IOS XE Release 3.1S.

Usage Guidelines

The **show ip bgp vpnv4 all sso summary** command is used to display the number of BGP neighbors that are in SSO mode.

Examples

The following is sample output from the **show ip bgp vpnv4 all sso summary** command:

```
Router# show ip bgp vpnv4 all sso summary

   Stateful switchover support enabled for 40 neighbors
```

Table 60 describes the significant fields shown in the display.

Table 60 show ip bgp vpnv4 all sso summary Field Descriptions

Field	Description
Stateful Switchover support enabled for	Indicates the number of BGP neighbors that are in SSO mode.

Related Commands

Command	Description
neighbor ha-mode sso	Configures a BGP neighbor to support SSO.

show ip bgp vpnv4

To display Virtual Private Network Version 4 (VPNv4) address information from the Border Gateway Protocol (BGP) table, use the **show ip bgp vpnv4** command in user EXEC or privileged EXEC mode.

show ip bgp vpnv4 {**all** | **rd** *as_number*:*nn* | *ip_address*:*nn* | **vrf** *vrf-name*} [*ip-prefix/length* | [*network-address* [*mask*] [[**longer-prefixes**] | [**multipaths**] | [**shorter-prefixes** [*mask-length*]] | [**subnets**]] | [**cidr-only**] | [**community-list** *community-list-number* | *community-list-name*] | [**dampening** {**dampened-paths** | **flap-statistics** | **parameters**}] | [**filter-list** *regular_expression_access_list_number*] | [**inconsistency nexthop-label**] | [**inconsistent-as**] | [**labels**] | [**neighbors** [*ip-address* | *ipv6-address*] | [**paths** [*line*]] | [**peer-group** [*peer-group-name* [**summary**]]] | [**quote-regexp** *regexp*] | [**regexp**] | [**rib-failure**] | [**summary**]]

Syntax Description

all	Displays the complete VPNv4 database.
rd *as_number*:*nn* | *ip_address*:*nn*	Displays Network Layer Reachability Information (NLRI) prefixes that match the specified route distinguisher.
vrf *vrf-name*	Displays NLRI prefixes associated with the named VPN routing and forwarding (VRF) instance.
ip-prefix/length	(Optional) IP prefix address (in dotted decimal format) and the length of the mask (0 to 32). The slash mark must be included.
network-address	(Optional) IP address of a network in the BGP routing table.
mask	(Optional) Mask of the network address, in dotted decimal format.
longer-prefixes	(Optional) Displays the entry, if any, that exactly matches the specified prefix parameter and all entries that match the prefix in a "longest-match" sense. That is, prefixes for which the specified prefix is an initial substring.
multipaths	(Optional) Displays the multipaths for this prefix.
shorter-prefixes	(Optional) Displays less specific routes.
mask-length	(Optional) Displays prefixes longer than this mask length.
subnets	(Optional) Displays route and more specific routes.
cidr-only	(Optional) Displays only routes that have nonclassful net masks.
community-list *community-list-number* | *community-list-name*	(Optional) Displays routes that pass the specified community list.
dampening	(Optional) Displays paths suppressed because of dampening (BGP route from peer is up and down).
dampened-paths	(Optional) Displays paths suppressed due to dampening.
flap-statistics	(Optional) Displays flap statistics of routes.
parameters	(Optional) Displays details of configured dampening parameters.
filter-list *regular_expresssion_access_list_number*	(Optional) Displays routes that conform to the filter list.
inconsistency nexthop-label	(Optional) Displays routes that have a nexthop-label inconsistency found when the **bgp consistency-checker** command is configured.

inconsistent-as	(Optional) Displays only routes that have inconsistent autonomous systems of origin.
labels	(Optional) Displays incoming and outgoing BGP labels for each NLRI prefix.
neighbors	(Optional) Displays details about TCP and BGP neighbor connections.
ip-address | *ipv6-address*	(Optional) Displays information about specific neighbor.
paths	(Optional) Displays path information.
line	(Optional) A regular expression to match the BGP autonomous system paths.
peer-group	(Optional) Displays information about peer groups.
peer-group-name	(Optional) Displays information about specific peer group.
summary	(Optional) Displays summary of peer-group member status.
quote-regexp *regexp*	(Optional) Displays routes that match the autonomous system path regular expression.
regexp *line*	(Optional) Displays routes that match the autonomous system path regular expression. The *line* argument is a regular expression to match BGP AS paths.
rib-failure	(Optional) Displays BGP routes that failed to install in the VRF table.
summary	(Optional) Displays summary of BGP neighbor status.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.2(2)T	The output of the **show ip bgp vpnv4 all** *ip-prefix* command was enhanced to display attributes including multipaths and a best path to the specified network.
12.0(21)ST	The **tags** keyword was replaced by the **labels** keyword to conform to the MPLS guidelines. This command was integrated into Cisco IOS Release 12.0(21)ST.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.2(13)T	This command was integrated into Cisco IOS Release 12.2(13)T.
12.0(27)S	The output of the **show ip bgp vpnv4 all labels** command was enhanced to display explicit-null label information.
12.3	The **rib-failure** keyword was added for VRFs.
12.2(22)S	The output of the **show ip bgp vpnv4 vrf** *vrf-name* **labels** command was modified so that directly connected VRF networks no longer display as aggregate; no label appears instead.
12.2(25)S	This command was updated to display MPLS VPN nonstop forwarding information.

Release	Modification
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB and implemented on the Cisco 10000 series router. The display output was modified to indicate whether BGP Nonstop Routing (NSR) with stateful switchover (SSO) is enabled and the reason the last BGP lost SSO capability.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA, and the output was modified to support per-VRF assignment of the BGP router ID.
12.2(31)SB2	The output was modified to support per-VRF assignment of the BGP router ID.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH, and the output was modified to support per-VRF assignment of the BGP router ID. **Note** In Cisco IOS Release 12.2(33)SXH, the command output does not display on the standby route processor in NSF/SSO mode.
12.4(20)T	The output was modified to support per-VRF assignment of the BGP router ID.
15.0(1)M	This command was modified. The output was modified to support BGP Event-Based VPN Import.
12.2(33)SRE	This command was modified. The command output was modified to support the BGP Event-Based VPN Import, BGP best external and BGP additional path features.
12.2(33)XNE	This command was integrated into Cisco IOS Release 12.2(33)XNE.
Cisco IOS XE Release 2.5	This command was integrated into Cisco IOS XE Release 2.5.
15.0(1)S	This command was integrated into Cisco IOS Release 15.0(1)S.
15.1(2)S	This command was modified. The **inconsistency nexthop-label** keyword was added.
Cisco IOS XE 3.3S	This command was modified. The **inconsistency nexthop-label** keyword was added.

Usage Guidelines

Use this command to display VPNv4 information from the BGP database. The **show ip bgp vpnv4 all** command displays all available VPNv4 information. The **show ip bgp vpnv4 all summary** command displays BGP neighbor status. The **show ip bgp vpnv4 all labels** command displays explicit-null label information.

Examples

The following example shows all available VPNv4 information in a BGP routing table:

```
Router# show ip bgp vpnv4 all

BGP table version is 18, local router ID is 10.14.14.14
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
Route Distinguisher: 1:101 (default for vrf vpn1)
*>i10.6.6.6/32      10.0.0.21                11    100      0 ?
*> 10.7.7.7/32      10.150.0.2               11          32768 ?
*>i10.69.0.0/30     10.0.0.21                 0    100      0 ?
*> 10.150.0.0/24    0.0.0.0                   0          32768 ?
```

Table 61 describes the significant fields shown in the display.

Table 61 show ip bgp vpnv4 all Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Displays the address of the BGP next hop.
Metric	Displays the BGP metric.
LocPrf	Displays the local preference.
Weight	Displays the BGP weight.
Path	Displays the BGP path per route.

The following example shows how to display a table of labels for NLRI prefixes that have a route distinguisher value of 100:1.

```
Router# show ip bgp vpnv4 rd 100:1 labels

Network            Next Hop       In label/Out label
Route Distinguisher: 100:1 (vrf1)
   10.0.0.0         10.20.0.60      34/nolabel
   10.0.0.0         10.20.0.60      35/nolabel
   10.0.0.0         10.20.0.60      26/nolabel
                    10.20.0.60      26/nolabel
   10.0.0.0         10.15.0.15      nolabel/26
```

Table 62 describes the significant fields shown in the display.

Table 62 show ip bgp vpnv4 rd labels Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Specifies the BGP next hop address.
In label	Displays the label (if any) assigned by this router.
Out label	Displays the label assigned by the BGP next-hop router.

The following example shows VPNv4 routing entries for the VRF named vpn1:

```
Router# show ip bgp vpnv4 vrf vpn1

BGP table version is 18, local router ID is 10.14.14.14
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 100:1 (default for vrf test1)
*> 10.1.1.1/32      192.168.1.1              0             0 100 i
*bi                 10.4.4.4                 0    100      0 100 i
*> 10.2.2.2/32      192.168.1.1                           0 100 i
*bi                 10.4.4.4                 0    100      0 100 i
*> 172.16.1.0/24    192.168.1.1              0             0 100 i
* i                 10.4.4.4                 0    100      0 100 i
r> 192.168.1.0      192.168.1.1              0             0 100 i
rbi                 10.4.4.4                 0    100      0 100 i
*> 192.168.3.0      192.168.1.1                           0 100 i
*bi                 10.4.4.4                 0    100      0 100 i
```

Table 63 describes the significant fields shown in the display.

Table 63 show ip bgp vpnv4 vrf Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Displays the address of the BGP next hop.
Metric	Displays the BGP metric.
LocPrf	Displays the local preference.
Weight	Displays the BGP weight.
Path	Displays the BGP path per route.

The following example shows attributes for network 192.168.9.0 that include multipaths, best path, and a recursive-via-host flag:

```
Router# show ip bgp vpnv4 vrf vpn1 192.168.9.0 255.255.255.0

BGP routing table entry for 100:1:192.168.9.0/24, version 44
Paths: (2 available, best #2, table test1)
  Additional-path
  Advertised to update-groups:
     2
  100, imported path from 400:1:192.168.9.0/24
    10.8.8.8 (metric 20) from 10.5.5.5 (10.5.5.5)
      Origin IGP, metric 0, localpref 100, valid, internal, backup/repair
      Extended Community: RT:100:1 RT:200:1 RT:300:1 RT:400:1
      Originator: 10.8.8.8, Cluster list: 10.5.5.5 , recursive-via-host
      mpls labels in/out nolabel/17
  100, imported path from 300:1:192.168.9.0/24
    10.7.7.7 (metric 20) from 10.5.5.5 (10.5.5.5)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:100:1 RT:200:1 RT:300:1 RT:400:1
      Originator: 10.7.7.7, Cluster list: 10.5.5.5 , recursive-via-host
      mpls labels in/out nolabel/17
```

Table 64 describes the significant fields shown in the display.

Table 64 show ip bgp vpnv4 all network-address Field Descriptions

Field	Description
BGP routing table entry for ... version	Internal version number of the table. This number is incremented whenever the table changes.
Paths	Number of autonomous system paths to the specified network. If multiple paths exist, one of the multipaths is designated the best path.
Multipath	Indicates the maximum paths configured (iBGP or eBGP).
Advertised to non peer-group peers	IP address of the BGP peers to which the specified route is advertised.
10.22.7.8 (metric 11) from 10.11.3.4 (10.0.0.8)	Indicates the next hop address and the address of the gateway that sent the update.

Table 64 show ip bgp vpnv4 all network-address Field Descriptions (continued)

Field	Description
Origin	Indicates the origin of the entry. It can be one of the following values: • IGP—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. • incomplete—Entry originated from other than an IGP or Exterior Gateway Protocol (EGP) and was advertised with the **redistribute** router configuration command. • EGP—Entry originated from an EGP.
metric	If shown, the value of the interautonomous system metric.
localpref	Local preference value as set with the **set local-preference route-map** configuration command. The default value is 100.
valid	Indicates that the route is usable and has a valid set of attributes.
internal/external	The field is *internal* if the path is learned via iBGP. The field is *external* if the path is learned via eBGP.
multipath	One of multiple paths to the specified network.
best	If multiple paths exist, one of the multipaths is designated the best path and this path is advertised to neighbors.
Extended Community	Route Target value associated with the specified route.
Originator	The router ID of the router from which the route originated when route reflector is used.
Cluster list	The router ID of all the route reflectors that the specified route has passed through.

The following example shows routes that BGP could not install in the VRF table:

```
Router# show ip bgp vpnv4 vrf xyz rib-failure

Network            Next Hop                      RIB-failure   RIB-NH Matches
Route Distinguisher: 2:2 (default for vrf bar)
10.1.1.2/32        10.100.100.100      Higher admin distance              No
10.111.111.112/32  10.9.9.9            Higher admin distance             Yes
```

Table 65 describes the significant fields shown in the display.

Table 65 show ip bgp vpnv4 vrf rib-failure Field Descriptions

Field	Description
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.

Table 65 show ip bgp vpnv4 vrf rib-failure Field Descriptions (continued)

Field	Description
RIB-failure	Cause of the Routing Information Base (RIB) failure. Higher admin distance means that a route with a better (lower) administrative distance, such as a static route, already exists in the IP routing table.
RIB-NH Matches	Route status that applies only when Higher admin distance appears in the RIB-failure column and the **bgp suppress-inactive** command is configured for the address family being used. There are three choices: • Yes—Means that the route in the RIB has the same next hop as the BGP route or that the next hop recurses down to the same adjacency as the BGP next hop. • No—Means that the next hop in the RIB recurses down differently from the next hop of the BGP route. • n/a—Means that the **bgp suppress-inactive** command is not configured for the address family being used.

The following example shows the information displayed on the active and standby Route Processors when they are configured for NSF/SSO: MPLS VPN.

Note In Cisco IOS Release 12.2(33)SXH, the Cisco IOS Software Modularity: MPLS Layer 3 VPNs feature incurred various infrastructure changes. The result of those changes affects the output of this command on the standby Route Processor (RP). In Cisco IOS Release 12.2(33)SXH, the standby RP does not display any output from the **show ip bgp vpnv4** command.

Active Route Processor

```
Router# show ip bgp vpnv4 all labels

Network          Next Hop   In label/Out label
Route Distinguisher: 100:1 (vpn1)
10.12.12.12/32   0.0.0.0    16/aggregate(vpn1)
10.0.0.0/8       0.0.0.0    17/aggregate(vpn1)
Route Distinguisher: 609:1 (vpn0)
10.13.13.13/32   0.0.0.0    18/aggregate(vpn0)

Router# show ip bgp vpnv4 vrf vpn1 labels

Network          Next Hop   In label/Out label
Route Distinguisher: 100:1 (vpn1)
10.12.12.12/32   0.0.0.0    16/aggregate(vpn1)
10.0.0.0/8       0.0.0.0    17/aggregate(vpn1)
```

Standby Route Processor

```
Router# show ip bgp vpnv4 all labels

Network       Masklen   In label
Route Distinguisher: 100:1
10.12.12.12   /32       16
10.0.0.0      /8        17
Route Distinguisher: 609:1
10.13.13.13   /32       18
```

```
Router# show ip bgp vpnv4 vrf vpn1 labels

Network        Masklen   In label
Route Distinguisher: 100:1
10.12.12.12    /32       16
10.0.0.0       /8        17
```

Table 66 describes the significant fields shown in the display.

Table 66 show ip bgp vpn4 labels Field Descriptions

Field	Description
Network	The network address from the BGP table.
Next Hop	The BGP next-hop address.
In label	The label (if any) assigned by this router.
Out label	The label assigned by the BGP next-hop router.
Masklen	The mask length of the network address.

The following example displays output, including the explicit-null label, from the **show ip bgp vpnv4 all labels** command on a CSC-PE router:

```
Router# show ip bgp vpnv4 all labels

   Network          Next Hop      In label/Out label
Route Distinguisher: 100:1 (v1)
   10.0.0.0/24       10.0.0.0        19/aggregate(v1)
   10.0.0.1/32       10.0.0.0        20/nolabel
   10.1.1.1/32       10.0.0.0        21/aggregate(v1)
   10.10.10.10/32    10.0.0.1        25/exp-null
   10.168.100.100/32
                     10.0.0.1        23/exp-null
   10.168.101.101/32
                     10.0.0.1        22/exp-null
```

Table 67 describes the significant fields shown in the display.

Table 67 show ip bgp vpnv4 all labels Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Displays the address of the BGP next hop.
In label	Displays the label (if any) assigned by this router.
Out label	Displays the label assigned by the BGP next-hop router.
Route Distinguisher	Displays an 8-byte value added to an IPv4 prefix to create a VPN IPv4 prefix.

The following example displays separate router IDs for each VRF in the output from an image in Cisco IOS Release 12.2(31)SB2, 12.2(33)SRA, 12.2(33)SXH, 12.4(20)T, Cisco IOS XE Release 2.1, and later releases with the Per-VRF Assignment of BGP Router ID feature configured. The router ID is shown next to the VRF name.

```
Router# show ip bgp vpnv4 all
```

```
BGP table version is 5, local router ID is 172.17.1.99
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf vrf_trans) VRF Router ID 10.99.1.2
*> 192.168.4.0      0.0.0.0                  0         32768 ?
Route Distinguisher: 42:1 (default for vrf vrf_user) VRF Router ID 10.99.1.1
*> 192.168.5.0      0.0.0.0                  0         32768 ?
```

Table 68 describes the significant fields shown in the display.

Table 68 show ip bgp vpnv4 all (VRF Router ID) Field Descriptions

Field	Description
Route Distinguisher	Displays an 8-byte value added to an IPv4 prefix to create a VPN IPv4 prefix.
vrf	Name of the VRF.
VRF Router ID	Router ID for the VRF.

In this example, the BGP Event-Based VPN Import feature is configured in Cisco IOS Release 15.0(1)M, 12.2(33)SRE, and later releases. When the **import path selection** command is configured, but the **strict** keyword is not included, then a safe import path selection policy is in effect. When a path is imported as the best available path (when the bestpath or multipaths are not eligible for import), the imported path includes the wording "imported safety path," as shown in the output.

```
Router# show ip bgp vpnv4 all 172.17.0.0

BGP routing table entry for 45000:1:172.17.0.0/16, version 10
Paths: (1 available, best #1, table vrf-A)
Flag: 0x820
   Not advertised to any peer
   2, imported safety path from 50000:2:172.17.0.0/16
     10.0.101.1 from 10.0.101.1 (10.0.101.1)
       Origin IGP, metric 200, localpref 100, valid, internal, best
       Extended Community: RT:45000:100
```

In this example the BGP Event-Based VPN Import feature is configured in Cisco IOS Release 15.0(1)M, 12.2(33)SRE, and later releases. When the **import path selection** command is configured with the **all** keyword, any path that matches an RD of the specified VRF will be imported, even though the path does not match the RTs imported by the specified VRF. In this situation, the imported path is marked as "not-in-vrf" as shown in the output. Note that on the net for vrf-A, this path is not the bestpath as any paths that are not in the VRFs appear less attractive than paths in the VRF.

```
Router# show ip bgp vpnv4 all 172.17.0.0

BBGP routing table entry for 45000:1:172.17.0.0/16, version 11
Paths: (2 available, best #2, table vrf-A)
Flag: 0x820
   Not advertised to any peer
   2
     10.0.101.2 from 10.0.101.2 (10.0.101.2)
       Origin IGP, metric 100, localpref 100, valid, internal, not-in-vrf
       Extended Community: RT:45000:200
       mpls labels in/out nolabel/16
   2
     10.0.101.1 from 10.0.101.1 (10.0.101.1)
       Origin IGP, metric 50, localpref 100, valid, internal, best
```

```
Extended Community: RT:45000:100
mpls labels in/out nolabel/16
```

Related Commands

Command	Description
import path limit	Specifies the maximum number of BGP paths, per VRF importing net, that can be imported from an exporting net.
import path selection	Specifies the BGP import path selection policy for a specific VRF instance.
show ip vrf	Displays the set of defined VRFs and associated interfaces.

show ip bgp vpnv4 all dampening

To display BGP dampening information for the Virtual Private Network Version 4 (VPNv4) address family, use the **show ip bgp vpnv4 all dampening** command in user EXEC or privileged EXEC mode.

show ip bgp vpnv4 all dampening {dampened-paths | flap-statistics [*network-address* **[***mask* **| bestpath | multipaths]] |** *ip-prefix/length* **| cidr-only | filter-list** *filter-list* **| oer-paths | prefix-list** *prefix-list* **| quote-regexp** *regexp* **| regexp** *regexp* **| route-map** *map-name* **| version** {*number* **| recent** }] **| parameters**}

Syntax Description

dampened-paths	Display routes suppressed due to dampening.
flap-statistics	Displays flap statistics of routes.
network-address	(Optional) Used with the **flap-statistics** keyword, network in the BGP routing table to display.
mask	(Optional) Used with the *network-address* argument, network mask that determines the networks displayed.
bestpath	(Optional) Used with the *network-address* argument, displays the bestpath for this prefix.
multipaths	(Optional) Used with the *network-address* argument, displays the multipaths for this prefix.
ip-prefix/length	(Optional) Used with the **flap-statistics** keyword, IP prefix/network length, such as 10.0.0.0/8.
cidr-only	(Optional) Used with the **flap-statistics** keyword, displays only routes with non-natural netmasks.
filter-list *filter-list*	(Optional) Used with the **flap-statistics** keyword, displays routes that conform to the specified filter list in the range 1-500.
oer-paths	(Optional) Used with the **flap-statistics** keyword, displays all OER controlled paths.
prefix-list *prefix-list*	(Optional) Used with the **flap-statistics** keyword, displays routes allowed by the prefix list.
quote-regexp *regexp*	(Optional) Used with the **flap-statistics** keyword, displays routes matching the AS path “regular expression”.
regexp *regexp*	(Optional) Used with the **flap-statistics** keyword, displays routes matching the AS path regular expression.
route-map *map-name*	(Optional) Used with the **flap-statistics** keyword, displays routes allowed by the route map.
version *number* **| recent**	(Optional) Used with the **flap-statistics** keyword, displays version of BGP table.
parameters	Display details of configured dampening parameters.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
15.0(1)M	This command was introduced.

Usage Guidelines

Use this command to display dampening information for the VPNv4 address family.

Examples

The following example shows dampening flap-statistics for the VPNv4 address family:

```
Router# show ip bgp vpnv4 all dampening flap-statistics

For_address_family: VPNv4 Unicast

% dampening not enabled for base

For vrf: Cust_A

BGP table version is 15, local router ID is 144.124.23.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external, f
RT-Filter
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          From            Flaps Duration Reuse    Path
*>  20.20.20.20/32   172.16.1.2      1     00:01:05          65001

For vrf: Cust_B

*d  11.11.11.11/32   192.168.1.2     3     00:04:22 00:04:49 65001
Router#
```

Related Commands

Command	Description
bgp dampening	Enables BGP route dampening or changes BGP route dampening parameters.

show ip bgp vpnv6 unicast all dampening

To display BGP dampening information for the Virtual Private Network Version 6 (VPNv6) address family, use the **show ip bgp vpnv6 unicast all dampening** command in user EXEC or privileged EXEC mode.

show ip bgp vpnv6 unicast all dampening {**dampened-paths** | **flap-statistics** [*network/length* | **filter-list** *filter-list* | **injected-paths** | **prefix-list** *prefix-list* | **quote-regexp** *regexp* | **regexp** *regexp* | **route-map** *map-name*] | **parameters**}

Syntax Description

dampened-paths	Display routes suppressed due to dampening.
flap-statistics	Displays flap statistics of routes.
network/length	(Optional) Used with the **flap-statistics** keyword, IPv6 prefix network/length in the format *X:X:X:X::X/<0-128>*.
filter-list *filter-list*	(Optional) Used with the **flap-statistics** keyword, displays routes that conform to the specified filter list in the range 1-500.
injected-paths	(Optional) Used with the **flap-statistics** keyword, displays all injected paths.
prefix-list *list*	(Optional) Used with the **flap-statistics** keyword, displays routes allowed by the prefix list.
quote-regexp *regexp*	(Optional) Used with the **flap-statistics** keyword, displays routes matching the AS path “regular expression”.
regexp *regexp*	(Optional) Used with the **flap-statistics** keyword, displays routes matching the AS path regular expression.
route-map *map-name*	(Optional) Used with the **flap-statistics** keyword, displays routes allowed by the route map.
parameters	Display details of configured dampening parameters.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
15.0(1)S	This command was introduced.

Usage Guidelines

Use this command to display BGP dampening information for the VPNv6 address family.

Examples

The following example shows dampening VPNv6 information:

```
Router# show ip bgp vpnv6 unicast all dampening flap-statistics

For_address_family: VPNv6 Unicast

% dampening not enabled for base
```

```
For vrf: RED


For vrf: BLUE

BGP table version is 36, local router ID is 10.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external, f
RT-Filter
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          From            Flaps Duration Reuse    Path
*d  11::/64          20::2           3     00:03:17 00:05:59 2
*d  22::/64          20::2           3     00:03:17 00:05:59 2
*d  33::/64          20::2           3     00:03:17 00:05:59 2
*d  44::/64          20::2           3     00:03:17 00:05:59 2
*d  55::/64          20::2           3     00:03:17 00:05:59 2
R1#
```

Related Commands

Command	Description
bgp dampening	Enables BGP route dampening or changes BGP route dampening parameters.

show ip community-list

To display configured community lists, use the **show ip community-list** command in user or privileged EXEC mode.

show ip community-list [*community-list-number* | *community-list-name*] [**exact-match**]

Syntax Description

community-list-number	(Optional) A standard or expanded community list number in the range from 1 to 500.
community-list-name	(Optional) Community list name. The community list name can be standard or expanded.
exact-match	(Optional) Displays only routes that have an exact match.

Command Modes

User EXEC
Privileged EXEC

Command History

Release	Modification
11.0	This command was introduced.
12.0(10)S	Named community list support was added.
12.0(16)ST	Named community lists support was integrated into Cisco IOS Release 12.0(16)ST.
12.1(9)E	Named community lists support was integrated into Cisco IOS Release 12.1(9)E.
12.2(8)T	Named community lists support was integrated into Cisco IOS Release 12.2(8)T.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

This command can be used without any arguments or keywords. If no arguments are specified, this command will display all community lists. However, the community list name or number can be specified when entering the **show ip community-list** command. This option can be useful for filtering the output of this command and verifying a single named or numbered community list.

Examples

The following sample output is similar to the output that will be displayed when the **show ip community-list** command is entered in privileged EXEC mode:

```
Router# show ip community-list

Community standard list 1
     permit 3
```

```
     deny 5
Community (expanded) access list 101
    deny 4
    permit 6
Named Community standard list COMMUNITY_LIST_NAME
    permit 1
    deny 7
Named Community expanded list COMMUNITY_LIST_NAME_TWO
    deny 2
    permit 8
```

Table 69 describes the significant fields shown in the display.

Table 69 show ip community-list Field Descriptions

Field	Description
Community standard list	If shown, this value will display a standard community list number (1 to 99). The standard community list number will immediately follow this value.
Community (expanded) access list	If shown, this value will display an expanded community list number (100 to 500). The expanded community list number will immediately follow this value.
Named community standard list	If shown, this value will display a standard community list name. The standard community list name will immediately follow this value.
Named community expanded list	If shown, this value will display an expanded community list name. The expanded community list name will immediately follow this value.

show ip extcommunity-list

To display routes that are permitted by an extended community list, use the **show ip extcommunity-list** command in user EXEC or privileged EXEC mode.

show ip extcommunity-list [*list-number* | *list-name*]

Syntax Description

list-number	(Optional) Specifies an extended community list number from 1 to 500. A standard extended community list number is from 1 to 99. An expanded extended list is from 100 to 500.
list-name	(Optional) Specifies an extended community list name. If a specific extended community list number is not specified, all locally configured extended community lists will be displayed by default.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.1	This command was introduced.
12.2(25)S	Support for named extended community lists was added. Minor formatting changes were made to the output.
12.3(11)T	Support for named extended community lists was added. Minor formatting changes were made to the output.
12.2(27)SBC	This command was integrated into the Cisco IOS Release 12.2(27)SBC.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.
12.0(32)S12	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
12.0(32)SY8	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.4(24)T	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
Cisco IOS XE Release 2.3	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation only was added.
12.2(33)SXI1	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.0(33)S3	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format is now asplain.
Cisco IOS XE Release 2.4	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format is now asplain.

Release	Modification
12.2(33)SRE	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)XNE	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.

Usage Guidelines

In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, the Cisco implementation of 4-byte autonomous system numbers uses asplain—65538 for example—as the default regular expression match and output display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain format and the asdot format as described in RFC 5396. To change the default regular expression match and output display of 4-byte autonomous system numbers to asdot format, use the **bgp asnotation dot** command followed by the **clear ip bgp *** command to perform a hard reset of all current BGP sessions.

In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, the Cisco implementation of 4-byte autonomous system numbers uses asdot—1.2 for example—as the only configuration format, regular expression match, and output display, with no asplain support.

If the route target—RT in the output—contains a 4-byte autonomous system number as part of the extended community list, it will be displayed in the appropriate format.

Examples

The following is sample output from the **show ip extcommunity-list** command:

```
Router# show ip extcommunity-list

Standard extended community-list 1
     10 permit RT:64512:10
     20 permit SoO:65400:20
     30 deny RT:65424:30 SoO:64524:40
Standard extended community-list 99
     10 permit RT:65504:40 SoO:65505:50
     20 deny RT:65406:60 SoO:65307:70
Expanded extended community-list LIST_NAME
     10 permit 0-9* A-Z* a-z*
```

Table 70 describes the significant fields shown in the display.

Table 70 show ip extcommunity-list Field Descriptions

Field	Description
... extended community-list....	The type of extended community-list (standard or expanded), and the name or number of the extended community list.
10	The sequence number of the extended community list entry. 10 is the lowest default sequence number. Extended community lists increment by 10 when default values are configured.
permit/deny	Indicates a permit or deny sequence entry.
RT/SoO	Indicates the route target or the site of origin used in a standard extended community list.
0-9* A-Z* a-z*	Regular expression used in an expanded extended community list.

The following output is from the **show ip extcommunity-list** command after a 4-byte autonomous system number has been configured as part of the route target. The 4-byte autonomous system number, 65537, is displayed in the default asplain format. This example requires Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, or a later release.

```
Router# show ip extcommunity-list 1

Extended community standard list 1
     permit RT:65537:100
```

The following output displays a 4-byte autonomous system number that has been configured as part of the route target. The 4-byte autonomous system number—1.1—is displayed in asdot notation. The dot notation is the only format for 4-byte autonomous system numbers in Cisco IOS Release 12.0(32)S12, 12.4(24)T, or Cisco IOS XE Release 2.3. This output can also be seen in Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, or later releases. after the **bgp asnotation dot** command has been entered to display 4-byte autonomous system numbers in dot notation.

```
Router# show ip extcommunity-list 1

Extended community standard list 1
     permit RT:1.1:100
```

Related Commands

Command	Description
bgp asnotation dot	Changes the default display and the regular expression match format of BGP 4-byte autonomous system numbers from asplain (decimal values) to dot notation.
router bgp	Configures the BGP routing process.
show route-map	Displays configured route maps.

show ip policy-list

To display information about a configured policy list and policy list entries, use the **show ip policy-list command in user** EXEC mode.

show ip policy-list [*policy-list-name*]

Syntax Description

policy-list-name	(Optional) Displays information about the specified policy list with this argument.

Command Modes

EXEC

Command History

Release	Modification
12.0(22)S	This command was introduced.
12.2(15)T	This command was integrated into 12.2(15)T.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.

Examples

The following is sample output from the **show ip policy-list** command. The output of this command will display the policy-list name and configured match clauses. The following sample output is similar to the output that will be displayed:

```
Router> show ip policy-list

policy-list POLICY-LIST-NAME-1 permit
  Match clauses:
    metric 20
policy-list POLICY-LIST-NAME-2 permit
  Match clauses:
    as-path (as-path filter): 1
```

Related Commands

Command	Description
show route-map	Displays configured route maps and information about referenced policy maps.

show ip prefix-list

To display information about a prefix list or prefix list entries, use the **show ip prefix-list** command in user or privileged EXEC mode.

show ip prefix-list [**detail | summary**][*prefix-list-name* [**seq** *sequence-number* | *network/length* [**longer| first-match**]]]

Syntax Description

detail | summary	(Optional) Displays detailed or summarized information about all prefix lists.
prefix-list-name	(Optional) Displays the entries in a specific prefix list.
seq *sequence-number*	(Optional) Displays only the prefix list entry with the specified sequence number in the specified prefix-list.
network/length	(Optional) Displays all entries in the specified prefix list that use this network address and netmask length (in bits).
longer	(Optional) Displays all entries of the specified prefix list that match or are more specific than the given *network/length.*
first-match	(Optional) Displays the first entry of the specified prefix list that matches the given *network/length*.

Command Modes

User EXEC
Privileged EXEC

Command History

Release	Modification
12.0	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command was integrated into Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Examples

The following example shows the output of the **show ip prefix-list** command with details about the prefix list named test:

```
Router# show ip prefix-list detail test

ip prefix-list test:
Description: test-list
 count: 1, range entries: 0, sequences: 10 - 10, refcount: 3
 seq 10 permit 10.0.0.0/8 (hit count: 0, refcount: 1)
```

Related Commands

Command	Description
clear ip prefix-list	Resets the hit count of the prefix list entries.
distribute-list in (BGP)	Filters networks received in updates.

distribute-list out (BGP)	Suppresses networks from being advertised in updates.
ip prefix-list	Creates an entry in a prefix list.
ip prefix-list description	Adds a text description of a prefix list.
match ip address	Distributes any routes that have a destination network number address that is permitted by a standard or extended access list, and performs policy routing on packets.
neighbor prefix-list	Distributes BGP neighbor information as specified in a prefix list.

show ip route

To display the current state of the routing table, use the **show ip route** command in user EXEC or privileged EXEC mode.

show ip route [*ip-address* [**repair-paths** | **next-hop-override** [**dhcp**] | *mask* [**longer-prefixes**]] | *protocol* [*process-id*] | **list** [*access-list-number* | *access-list-name*] | **static download** | **update-queue**]

Syntax Description

ip-address	(Optional) IP address about which routing information should be displayed.
repair-paths	(Optional) Displays the repair paths.
next-hop-override	(Optional) Displays the Next Hop Resolution Protocol (NHRP) overrides associated with a particular route, along with the corresponding default next hops.
dhcp	(Optional) Displays routes added by the Dynamic Host Configuration Protocol (DHCP) server.
mask	(Optional) The subnet mask.
longer-prefixes	(Optional) Specifies that only routes matching the *ip-address* and *mask* pair should be displayed.
protocol	(Optional) The name of a routing protocol, or the keyword **connected**, **mobile**, **static**, or **summary**. If you specify a routing protocol, use one of the following keywords: **bgp**, **eigrp**, **hello**, **isis**, **odr**, **ospf**, **nhr**, and **rip**.
process-id	(Optional) The number used to identify a process of the specified protocol.
list	(Optional) Filters output by an access list name or number.
access-list-number	(Optional) Specific access list number for which output from the routing table should be displayed.
access-list-name	(Optional) Specific access list name for which output from the routing table should be displayed.
static	(Optional) Displays static routes.
download	(Optional) Displays the route installed using the authentication, authorization, and accounting (AAA) route download function. This keyword is used only when AAA is configured.
update-queue	(Optional) Displays Routing Information Base (RIB) queue updates.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
9.2	This command was introduced.
10.0	The "D—EIGRP, EX—EIGRP, N1—OSPF NSSA external type 1 route" and "N2—OSPF NSSA external type 2 route" codes were added to the command output.
10.3	The *process-id* argument was added.
11.0	The **longer-prefixes** keyword was added.
11.1	The "U—per-user static route" code was added to the command output.
11.2	The "o—on-demand routing" code was added to the command output.
12.2(33)SRA	This command was modified. The **update-queue** keyword was added.
11.3	The output from the **show ip route** *ip-address* command was enhanced to display the origination of an IP route in Intermediate System-to-Intermediate System (IS-IS) networks.
12.0(1)T	The "M—mobile" code was added to the command output.
12.0(3)T	The "P—periodic downloaded static route" code was added to the command output.
12.0(4)T	The "ia—IS-IS" code was added to the command output.
12.2(2)T	The output from the **show ip route** *ip-address* command was enhanced to display information on the multipaths to the specified network.
12.2(13)T	The *egp* and *igrp* arguments were removed because the exterior gateway protocol (EGP) and the Interior Gateway Routing Protocol (IGRP) are no longer available in Cisco IOS software.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(14)SX	This command was integrated into Cisco IOS Release 12.2(14)SX.
12.3(2)T	The output was enhanced to display route tag information.
12.3(8)T	The output was enhanced to display static routes using DHCP.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
12.2(33)SRE	This command was modified. The **dhcp** and **repair-paths** keywords were added. Support for the Border Gateway Protocol (BGP) best external and BGP additional path features was added.
12.2(24)T	This command was modified. The "L" code was added.
12.2(33)XNE	This command was integrated into Cisco IOS Release 12.2(33)XNE.
Cisco IOS XE Release 2.5	This command was modified. The **next-hop-override** and **nhrp** keywords were added.
15.0(1)S	This command was integrated into Cisco IOS Release 15.0(1)S.

Usage Guidelines

The **show ip route static download** command provides a way to display all dynamic static routes with name and distance information, including active and inactive ones. You can display all active dynamic static routes with both the **show ip route** and **show ip route static** commands after these active routes are added in the main routing table.

Examples

Routing Table Examples

The following examples show the standard routing tables displayed by the **show ip route** command. Use the codes displayed at the beginning of each report and the information in Table 71 to understand the type of route.

The following is sample output from the **show ip route** command when entered without an address:

```
Router# show ip route

Codes: R - RIP derived, O - OSPF derived,
       C - connected, S - static, B - BGP derived,
       * - candidate default route, IA - OSPF inter area route,
       i - IS-IS derived, ia - IS-IS, U - per-user static route,
       o - on-demand routing, M - mobile, P - periodic downloaded static route,
       D - EIGRP, EX - EIGRP external, E1 - OSPF external type 1 route,
       E2 - OSPF external type 2 route, N1 - OSPF NSSA external type 1 route,
       N2 - OSPF NSSA external type 2 route

Gateway of last resort is 10.119.254.240 to network 10.140.0.0

O E2 10.110.0.0 [160/5] via 10.119.254.6, 0:01:00, Ethernet2
E    10.67.10.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
O E2 10.68.132.0 [160/5] via 10.119.254.6, 0:00:59, Ethernet2
O E2 10.130.0.0 [160/5] via 10.119.254.6, 0:00:59, Ethernet2
E    10.128.0.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.129.0.0 [200/129] via 10.119.254.240, 0:02:22, Ethernet2
E    10.65.129.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.10.0.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.75.139.0 [200/129] via 10.119.254.240, 0:02:23, Ethernet2
E    10.16.208.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.84.148.0 [200/129] via 10.119.254.240, 0:02:23, Ethernet2
E    10.31.223.0 [200/128] via 10.119.254.244, 0:02:22, Ethernet2
E    10.44.236.0 [200/129] via 10.119.254.240, 0:02:23, Ethernet2
E    10.141.0.0 [200/129] via 10.119.254.240, 0:02:22, Ethernet2
E    10.140.0.0 [200/129] via 10.119.254.240, 0:02:23, Ethernet2
```

The following is sample output that includes IS-IS Level 2 routes learned:

```
Router# show ip route

Codes: L- Local R - RIP derived, O - OSPF derived,
       C - connected, S - static, B - BGP derived,
       * - candidate default route, IA - OSPF inter area route,
       i - IS-IS derived, ia - IS-IS, U - per-user static route,
       o - on-demand routing, M - mobile, P - periodic downloaded static route,
       D - EIGRP, EX - EIGRP external, E1 - OSPF external type 1 route,
       E2 - OSPF external type 2 route, N1 - OSPF NSSA external type 1 route,
       N2 - OSPF NSSA external type 2 route

Gateway of last resort is 192.168.1.2 to network 0.0.0.0
S*    0.0.0.0/0 [1/0] via 192.168.1.2
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        10.10.10.0/24 is directly connected, Vlan1
L        10.10.10.1/32 is directly connected, Vlan1
192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.1.0/30 is directly connected, GigabitEthernet0
L        192.168.1.1/32 is directly connected, GigabitEthernet0
```

The following is sample output using the **longer-prefixes** keyword. When the **longer-prefixes** keyword is included, the address and mask pair becomes the prefix, and any address that matches that prefix is displayed. Therefore, multiple addresses are displayed.

In the following example, the logical AND operation is performed on the source address 10.0.0.0 and the mask 10.0.0.0, resulting in 10.0.0.0. Each destination in the routing table is also logically ANDed with the mask and compared to that result of 10.0.0.0. Any destinations that fall into that range are displayed in the output.

```
Router# show ip route 10.0.0.0 10.0.0.0 longer-prefixes

Codes: L - Local R - RIP derived, O - OSPF derived,
       C - connected, S - static, B - BGP derived,
       * - candidate default route, IA - OSPF inter area route,
       i - IS-IS derived, ia - IS-IS, U - per-user static route,
       o - on-demand routing, M - mobile, P - periodic downloaded static route,
       D - EIGRP, EX - EIGRP external, E1 - OSPF external type 1 route,
       E2 - OSPF external type 2 route, N1 - OSPF NSSA external type 1 route,
       N2 - OSPF NSSA external type 2 route

Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        10.4.9.0/24 is directly connected, GigabitEthernet0/1
L        10.4.9.134/32 is directly connected, GigabitEthernet0/1
      171.69.0.0/16 is variably subnetted, 2 subnets, 2 masks
S        171.69.0.0/16 [1/0] via 10.4.9.1
S        171.69.1.129/32 [1/0] via 10.4.9.1
```

The following examples display all downloaded static routes. A P designates which route was installed using AAA route download.

```
Router# show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR, P - periodic downloaded static route
       T - traffic engineered route

Gateway of last resort is 172.21.17.1 to network 0.0.0.0

        172.31.0.0/32 is subnetted, 1 subnets
P       172.31.229.41 is directly connected, Dialer1 20.0.0.0/24 is subnetted, 3 subnets
P       10.1.1.0 [200/0] via 172.31.229.41, Dialer1
P       10.1.3.0 [200/0] via 172.31.229.41, Dialer1
P       10.1.2.0 [200/0] via 172.31.229.41, Dialer1

Router# show ip route static

     172.27.4.0/8 is variably subnetted, 2 subnets, 2 masks
P       172.16.1.1/32 is directly connected, BRI0
P       172.27.4.0/8 [1/0] via 10.1.1.1, BRI0
S    172.31.0.0/16 [1/0] via 172.21.114.65, Ethernet0
S    10.0.0.0/8 is directly connected, BRI0
P    10.0.0.0/8 is directly connected, BRI0
     172.21.0.0/16 is variably subnetted, 5 subnets, 2 masks
S       172.21.114.201/32 is directly connected, BRI0
S       172.21.114.205/32 is directly connected, BRI0
S       172.21.114.174/32 is directly connected, BRI0
S       172.21.114.12/32 is directly connected, BRI0
P    10.0.0.0/8 is directly connected, BRI0
P    10.1.0.0/16 is directly connected, BRI0
P    10.2.2.0/24 is directly connected, BRI0
S*   0.0.0.0/0 [1/0] via 172.21.114.65, Ethernet0
```

```
S    172.29.0.0/16 [1/0] via 172.21.114.65, Ethernet0
```

The following example shows how to use the **show ip route static download** command to display all active and inactive routes installed using AAA route download:

```
Router# show ip route static download

Connectivity: A - Active, I - Inactive

A     10.10.0.0 255.0.0.0 BRI0
A     10.11.0.0 255.0.0.0 BRI0
A     10.12.0.0 255.0.0.0 BRI0
A     10.13.0.0 255.0.0.0 BRI0
I     10.20.0.0 255.0.0.0 172.21.1.1
I     10.22.0.0 255.0.0.0 Serial0
I     10.30.0.0 255.0.0.0 Serial0
I     10.31.0.0 255.0.0.0 Serial1
I     10.32.0.0 255.0.0.0 Serial1
A     10.34.0.0 255.0.0.0 192.168.1.1
A     10.36.1.1 255.255.255.255 BRI0 200 name remote1
I     10.38.1.9 255.255.255.0 192.168.69.1
```

The following example shows how to use the **show ip route nhrp** command to enable shortcut switching on the tunnel interface:

```
Router# show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP

Gateway of last resort is not set

       10.0.0.0/16 is variably subnetted, 3 subnets, 2 masks
C       10.1.1.0/24 is directly connected, Tunnel0
C       172.16.22.0 is directly connected, Ethernet1/0
H       172.16.99.0 [250/1] via 10.1.1.99, 00:11:43, Tunnel0
     10.11.0.0/24 is subnetted, 1 subnets
C       10.11.11.0 is directly connected, Ethernet0/0

Router# show ip route nhrp

H       172.16.99.0 [250/1] via 10.1.1.99, 00:11:43, Tunnel0
```

The following is sample output using the **next-hop-override** keyword. When the **next-hop-override** keyword is included, the NHRP Nexthop-overrides associated with a particular route, along with the corresponding default next hops, are displayed.

```
==============================================================
1) Initial configuration
==============================================================
Router# show ip route

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
```

```
       + - replicated route

Gateway of last resort is not set

      10.2.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        10.2.1.0/24 is directly connected, Loopback1
L        10.2.1.1/32 is directly connected, Loopback1
      10.0.0.0/24 is subnetted, 1 subnets
S        10.10.10.0 is directly connected, Tunnel0
      10.11.0.0/24 is subnetted, 1 subnets
S        10.11.11.0 is directly connected, Ethernet0/0

Router# show ip route next-hop-override

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
       + - replicated route

Gateway of last resort is not set

      10.2.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        10.2.1.0/24 is directly connected, Loopback1
L        10.2.1.1/32 is directly connected, Loopback1
      10.0.0.0/24 is subnetted, 1 subnets
S        10.10.10.0 is directly connected, Tunnel0
      10.11.0.0/24 is subnetted, 1 subnets
S        10.11.11.0 is directly connected, Ethernet0/0

Router# show ip cef

Prefix               Next Hop             Interface
.
.
.
10.2.1.255/32         receive              Loopback1
10.10.10.0/24        attached             Tunnel0  <<<<<<<<
10.11.11.0/24        attached             Ethernet0/0
127.0.0.0/8          drop
.
.
.
================================================================
2) Add a Nexthop-override
    address = 10.10.10.0
    mask = 255.255.255.0
    gateway = 10.1.1.1
    interface = Tunnel0
================================================================
Router# show ip route

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
       + - replicated route
```

```
Gateway of last resort is not set

      10.2.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        10.2.1.0/24 is directly connected, Loopback1
L        10.2.1.1/32 is directly connected, Loopback1
      10.0.0.0/24 is subnetted, 1 subnets
% S        10.10.10.0 is directly connected, Tunnel0
      10.11.0.0/24 is subnetted, 1 subnets
S        10.11.11.0 is directly connected, Ethernet0/0

Router# show ip route next-hop-override

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
       + - replicated route

Gateway of last resort is not set

      10.2.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        10.2.1.0/24 is directly connected, Loopback1
L        10.2.1.1/32 is directly connected, Loopback1
      10.0.0.0/24 is subnetted, 1 subnets
% S        10.10.10.0 is directly connected, Tunnel0
                    [NHO][1/0] via 10.1.1.1, Tunnel0
      10.11.0.0/24 is subnetted, 1 subnets
S        10.11.11.0 is directly connected, Ethernet0/0

Router# show ip cef

Prefix               Next Hop             Interface
.
.
.
10.2.1.255/32         receive              Loopback110.10.10.0/24

10.10.10.0/24       10.1.1.1             Tunnel0
10.11.11.0/24       attached            Ethernet0/0
10.12.0.0/16 drop
.
.
.

================================================================
3) Delete a Nexthop-override
    address = 10.10.10.0
    mask = 255.255.255.0
    gateway = 10.11.1.1
    interface = Tunnel0
================================================================
Router# show ip route

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
       + - replicated route
```

```
Gateway of last resort is not set

      10.2.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        10.2.1.0/24 is directly connected, Loopback1
L        10.2.1.1/32 is directly connected, Loopback1
      10.0.0.0/24 is subnetted, 1 subnets
S        10.10.10.0 is directly connected, Tunnel0
      10.11.0.0/24 is subnetted, 1 subnets
S        10.11.11.0 is directly connected, Ethernet0/0

Router# show ip route next-hop-override

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
       + - replicated route

Gateway of last resort is not set

      10.2.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        10.2.1.0/24 is directly connected, Loopback1
L        10.2.1.1/32 is directly connected, Loopback1
      10.0.0.0/24 is subnetted, 1 subnets
S        10.10.10.0 is directly connected, Tunnel0
      10.11.0.0/24 is subnetted, 1 subnets
S        10.11.11.0 is directly connected, Ethernet0/0

Router# show ip cef

Prefix               Next Hop             Interface
.
.
.
10.2.1.255/32         receive              Loopback110.10.10.0/24

10.10.10.0/24        attached             Tunnel0
10.11.11.0/24        attached             Ethernet0/0
10.120.0.0/16 drop
.
.
.
```

Table 71 show ip route Field Descriptions

Field	Description
Codes	Indicates the protocol that derived the route. It can be one of the following values: • B—BGP derived • C—connected • D—Enhanced Interior Gateway Routing Protocol (EIGRP) • EX—EIGRP external • H— NHRP • i—IS-IS derived • ia—IS-IS • L—local • M—mobile • O—Open Shortest Path First (OSPF) derived • P—periodic downloaded static route • R—Routing Information Protocol (RIP) derived • S—static • U—per-user static route • o—on-demand routing • +—replicated route
Codes	Type of route. It can be one of the following values: • *—Indicates the last path used when a packet was forwarded. It pertains only to the nonfast-switched packets. However, it does not indicate which path will be used next when forwarding a nonfast-switched packet, except when the paths are equal cost. • E1—OSPF external type 1 route • E2—OSPF external type 2 route • IA—OSPF inter area route • L1—IS-IS Level 1 route • L2—IS-IS Level 2 route • N1—OSPF not-so-stubby area (NSSA) external type 1 route • N2—OSPF NSSA external type 2 route
10.110.0.0	Indicates the address of the remote network.
[160/5]	The first number in the brackets is the administrative distance of the information source; the second number is the metric for the route.
via 10.119.254.6	Specifies the address of the next router to the remote network.
0:01:00	Specifies the last time the route was updated (in hours:minutes:seconds).
Ethernet2	Specifies the interface through which the specified network can be reached.

Specific Route Information

When you specify that you want information about a specific network displayed, more detailed statistics are shown. The following is sample output from the **show ip route** command when entered with the IP address 10.0.0.1:

```
Router# show ip route 10.0.0.1

Routing entry for 10.0.0.1/32
    Known via "isis", distance 115, metric 20, type level-1
    Redistributing via isis
    Last update from 10.191.255.251 on Fddi1/0, 00:00:13 ago
    Routing Descriptor Blocks:
    * 10.22.22.2, from 10.191.255.247, via Serial2/3
       Route metric is 20, traffic share count is 1
       10.191.255.251, from 10.191.255.247, via Fddi1/0
       Route metric is 20, traffic share count is 1
```

When an IS-IS router advertises its link-state information, it includes one of its own IP addresses to be used as the originator IP address. When other routers calculate IP routes, they can store the originator IP address with each route in the routing table.

The preceding example shows the output from the **show ip route** command for an IP route generated by IS-IS. Each path that is shown under the Routing Descriptor Blocks report displays two IP addresses. The first address (10.22.22.2) is the next hop address. The second is the originator IP address from the advertising IS-IS router. This address helps you determine where a particular IP route has originated in your network. In the example the route to 10.0.0.1/32 was originated by a router with IP address 10.191.255.247.

Table 72 describes the significant fields shown when using the **show ip route** command with an IP address.

Table 72 show ip route with IP Address Field Descriptions

Field	Description
Routing entry for 10.0.0.1/32	Network number and mask.
Known via...	Indicates how the route was derived.
Tag	Integer that is used to implement the route.
type	Indicates the IS-IS route type (Level 1 or Level 2).
Redistributing via...	Indicates the redistribution protocol.
Last update from 10.191.255.251	Indicates the IP address of a router that is the next hop to the remote network and the router interface on which the last update arrived.
Routing Descriptor Blocks:	Displays the next hop IP address followed by the information source.
Route metric	This value is the best metric for this routing descriptor block.
traffic share count	Number of uses for this routing descriptor block.

The following is sample output using the **longer-prefixes** keyword. When the **longer-prefixes** keyword is included, the address and mask pair becomes the prefix, and any address that matches that prefix is displayed. Therefore, multiple addresses are displayed.

In the following example, the logical AND operation is performed on the source address 10.0.0.0 and the mask 10.0.0.0, resulting in 10.0.0.0. Each destination in the routing table is also logically ANDed with the mask and compared to that result of 10.0.0.0. Any destinations that fall into that range are displayed in the output.

```
Router# show ip route 10.0.0.0 10.0.0.0 longer-prefixes

Codes: R - RIP derived, O - OSPF derived,
       C - connected, S - static, B - BGP derived,
       * - candidate default route, IA - OSPF inter area route,
       i - IS-IS derived, ia - IS-IS, U - per-user static route,
       o - on-demand routing, M - mobile, P - periodic downloaded static route,
       D - EIGRP, EX - EIGRP external, E1 - OSPF external type 1 route,
       E2 - OSPF external type 2 route, N1 - OSPF NSSA external type 1 route,
       N2 - OSPF NSSA external type 2 route

Gateway of last resort is not set

S    10.134.0.0 is directly connected, Ethernet0
S    10.10.0.0 is directly connected, Ethernet0
S    10.129.0.0 is directly connected, Ethernet0
S    10.128.0.0 is directly connected, Ethernet0
S    10.49.246.0 is directly connected, Ethernet0
S    10.160.97.0 is directly connected, Ethernet0
S    10.153.88.0 is directly connected, Ethernet0
S    10.76.141.0 is directly connected, Ethernet0
S    10.75.138.0 is directly connected, Ethernet0
S    10.44.237.0 is directly connected, Ethernet0
S    10.31.222.0 is directly connected, Ethernet0
S    10.16.209.0 is directly connected, Ethernet0
S    10.145.0.0 is directly connected, Ethernet0
S    10.141.0.0 is directly connected, Ethernet0
S    10.138.0.0 is directly connected, Ethernet0
S    10.128.0.0 is directly connected, Ethernet0
     10.19.0.0 255.255.255.0 is subnetted, 1 subnets
C       10.19.64.0 is directly connected, Ethernet0
     10.69.0.0 is variably subnetted, 2 subnets, 2 masks
C       10.69.232.32 255.255.255.240 is directly connected, Ethernet0
S       10.69.0.0 255.255.0.0 is directly connected, Ethernet0
```

The following output includes the tag 120 applied to the route 10.22.0.0/16. You must specify an IP prefix in order to see the tag value.

```
Router# show ip route 10.22.0.0

Routing entry for 10.22.0.0/16
  Known via "isis", distance 115, metric 12
  Tag 120, type level-1
  Redistributing via isis
  Last update from 172.19.170.12 on Ethernet2, 01:29:13 ago
  Routing Descriptor Blocks:
    * 172.19.170.12, from 10.3.3.3, via Ethernet2
        Route metric is 12, traffic share count is 1
        Route tag 120
```

Static Routes Using a DHCP Gateway Examples

The following example shows that IP route 10.8.8.0 is directly connected to the Internet and is the next-hop (option 3) default gateway. Routes 10.1.1.1 [1/0], 10.3.2.1 [24/0], and 172.2.2.2 [1/0] are static, and route 10.0.0.0/0 is a default route candidate.

```
Router# show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
```

```
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is 10.0.19.14 to network 0.0.0.0

10.0.0.0/24 is subnetted, 1 subnets
C 10.8.8.0 is directly connected, Ethernet1
  10.0.0.0/32 is subnetted, 1 subnets
S 10.1.1.1 [1/0] via 10.8.8.1
  10.0.0.0/32 is subnetted, 1 subnets
S 10.3.2.1 [24/0] via 10.8.8.1
  172.16.0.0/32 is subnetted, 1 subnets
S 172.2.2.2 [1/0] via 10.8.8.1
  10.0.0.0/28 is subnetted, 1 subnets
C 10.0.19.0 is directly connected, Ethernet0
  10.0.0.0/24 is subnetted, 1 subnets
C 10.15.15.0 is directly connected, Loopback0

S* 10.0.0.0/0 [1/0] via 10.0.19.14
```

The following sample output from the **show ip route repair-paths** command shows the repair paths marked with the tag [RPR]:

```
Router# show ip route repair-paths

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/32 is subnetted, 3 subnets
C        10.1.1.1 is directly connected, Loopback0
B        10.2.2.2 [200/0] via 172.16.1.2, 00:31:07
                  [RPR][200/0] via 192.168.1.2, 00:31:07
B        10.9.9.9 [20/0] via 192.168.1.2, 00:29:45
                  [RPR][20/0] via 192.168.3.2, 00:29:45
      172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C        172.16.1.0/24 is directly connected, Ethernet0/0
L        172.16.1.1/32 is directly connected, Ethernet0/0
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.1.0/24 is directly connected, Serial2/0
L        192.168.1.1/32 is directly connected, Serial2/0
B     192.168.3.0/24 [200/0] via 172.16.1.2, 00:31:07
                     [RPR][200/0] via 192.168.1.2, 00:31:07
B     192.168.9.0/24 [20/0] via 192.168.1.2, 00:29:45
                     [RPR][20/0] via 192.168.3.2, 00:29:45
B     192.168.13.0/24 [20/0] via 192.168.1.2, 00:29:45
                      [RPR][20/0] via 192.168.3.2, 00:29:45

Router# show ip route repair-paths 10.9.9.9

>Routing entry for 10.9.9.9/32
>  Known via "bgp 100", distance 20, metric 0
>  Tag 10, type external
```

```
>  Last update from 192.168.1.2 00:44:52 ago
>  Routing Descriptor Blocks:
>  * 192.168.1.2, from 192.168.1.2, 00:44:52 ago, recursive-via-conn
>      Route metric is 0, traffic share count is 1
>      AS Hops 2
>      Route tag 10
>      MPLS label: none
>    [RPR]192.168.3.2, from 172.16.1.2, 00:44:52 ago
>      Route metric is 0, traffic share count is 1
>      AS Hops 2
>      Route tag 10
>      MPLS label: none
```

Related Commands

Command	Description
show dialer	Displays general diagnostic information for interfaces configured for DDR.
show interfaces tunnel	Displays a list of tunnel interface information.
show ip route summary	Displays the current state of the routing table in summary format.

show ip route vrf

To display the IP routing table associated with a Virtual Private Network (VPN) routing and forwarding (VRF) instance, use the **show ip route vrf** command in user EXEC or privileged EXEC mode.

show ip route vrf *vrf-name* [**connected**] [*protocol* [*as-number*] [*tag*] [*output-modifiers*]] [**list** *number* [*output-modifiers*]] [**profile**] [**static** [*output-modifiers*]] [**summary** [*output-modifiers*]] [**supernets-only** [*output-modifiers*]] [*ip-address* [**repair-paths** [**dhcp** | *mask* [**longer-prefixes**]]]] [**supernets-only**]

Syntax Description

vrf-name	Name assigned to the VRF.
connected	(Optional) Displays all connected routes in a VRF.
protocol	(Optional) To specify a routing protocol, use one of the following keywords: **bgp**, **egp**, **eigrp**, **hello**, **igrp**, **isis**, **ospf**, or **rip**.
as-number	(Optional) Autonomous system number.
tag	(Optional) Cisco IOS routing area label.
output-modifiers	(Optional) For a list of associated keywords and arguments, use context-sensitive help.
ip-prefix	(Optional) Specifies a network to display.
list *number*	(Optional) Specifies the IP access list to display.
profile	(Optional) Displays the IP routing table profile.
static	(Optional) Displays static routes.
summary	(Optional) Displays a summary of routes.
ip-address	(Optional) Address about which routing information should be displayed.
repair-paths	(Optional) Displays the repair paths.
dhcp	(Optional) Displays routes added by the DHCP server.
longer-prefixes	(Optional) Specifies that only routes matching the *ip-address* and *mask* pair should be displayed.
supernets-only	(Optional) Displays supernet entries only.

Command Modes

User EXEC
Privileged EXEC

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.2(2)T	The *ip-prefix* argument was added. The output from the **show ip route vrf** *vrf-name ip-prefix* command was enhanced to display information on the multipaths to the specified network.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.0(22)S	Enhanced Interior Gateway Routing Protocol (EIGRP) VRF support was added.

Release	Modification
12.2(15)T	EIGRP VRF support was integrated into Cisco IOS Release 12.2(15)T.
12.2(18)S	EIGRP VRF support was integrated into Cisco IOS Release 12.2(18)S.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SXH	The output was enhanced to display remote label information and corresponding MPLS flags for prefixes that have remote labels stored in the Routing Information Base (RIB).
12.2(33)SRE	This command was modified. The **repair-paths**, **dhcp**, and **supernets-only** keywords were added. Support for the BGP best external and BGP additional path features was added.
12.2(33)XNE	This command was integrated into Cisco IOS Release 12.2(33)XNE.
Cisco IOS XE Release 2.5	This command was integrated into Cisco IOS XE Release 2.5.

Usage Guidelines

This command displays specified information from the IP routing table of a VRF.

Examples

This example shows the IP routing table associated with the VRF named vrf1:

```
Router# show ip route vrf vrf1

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR
       T - traffic engineered route

Gateway of last resort is not set

B    10.0.0.0/8 [200/0] via 10.13.13.13, 00:24:19
C    10.0.0.0/8 is directly connected, Ethernet1/3
B    10.0.0.0/8 [20/0] via 10.0.0.1, 02:10:22
B    10.0.0.0/8 [200/0] via 10.13.13.13, 00:24:20
```

This example shows BGP entries in the IP routing table associated with the VRF named vrf1:

```
Router# show ip route vrf vrf1 bgp

B  10.0.0.0/8 [200/0] via 10.13.13.13, 03:44:14
B  10.0.0.0/8 [20/0] via 10.0.0.1, 03:44:12
B  10.0.0.0/8 [200/0] via 10.13.13.13, 03:43:14
```

This example shows the IP routing table associated with a VRF named PATH and network 10.22.22.0:

```
Router# show ip route vrf PATH 10.22.22.0

Routing entry for 10.22.22.0/24
  Known via "bgp 1", distance 200, metric 0
  Tag 22, type internal
  Last update from 10.22.5.10 00:01:07 ago
  Routing Descriptor Blocks:
  * 10.22.7.8 (Default-IP-Routing-Table), from 10.11.3.4, 00:01:07 ago
```

```
    Route metric is 0, traffic share count is 1
    AS Hops 1
  10.22.1.9 (Default-IP-Routing-Table), from 10.11.1.2, 00:01:07 ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
  10.22.6.10 (Default-IP-Routing-Table), from 10.11.6.7, 00:01:07 ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
  10.22.4.10 (Default-IP-Routing-Table), from 10.11.4.5, 00:01:07 ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
  10.22.5.10 (Default-IP-Routing-Table), from 10.11.5.6, 00:01:07 ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
```

The following are sample outputs from the **show ip route vrf** command to include the recursive-via-host and recursive-via-connected flags.

```
Router# show ip route vrf v2 10.2.2.2

Routing Table: v2
Routing entry for 10.2.2.2/32
  Known via "bgp 10", distance 20, metric 0
  Tag 100, type external
  Last update from 192.168.1.1 00:15:54 ago
  Routing Descriptor Blocks:
  * 192.168.1.1, from 192.168.1.1, 00:15:54 ago, recursive-via-conn
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 100
      MPLS label: none

Router# show ip route vrf v2 10.2.2.2

Routing Table: v2
Routing entry for 10.2.2.2/32
  Known via "bgp 10", distance 200, metric 0
  Tag 100, type internal
  Last update from 10.3.3.3 00:18:11 ago
  Routing Descriptor Blocks:
  * 10.3.3.3 (default), from 10.5.5.5, 00:18:11 ago, recursive-via-host
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 100
      MPLS label: 16
      MPLS Flags: MPLS Required
```

Table 73 describes the significant fields shown when the **show ip route vrf** *vrf-name ip-prefix* command is used.

Table 73 show ip route vrf Field Descriptions

Field	Description
Routing entry for 10.22.22.0/24	Network number.
Known via ...	Indicates how the route was derived.
distance	Administrative distance of the information source.
metric	The metric to reach the destination network.
Tag	Integer that is used to implement the route.

Table 73 show ip route vrf Field Descriptions (continued)

Field	Description
type	Indicates that the route is an L1 type or L2 type route.
Last update from 10.22.5.10	Indicates the IP address of a router that is the next hop to the remote network and the router interface on which the last update arrived.
00:01:07 ago	Specifies the last time the route was updated (in hours:minutes:seconds).
Routing Descriptor Blocks:	Displays the next hop IP address followed by the information source.
10.22.6.10, from 10.11.6.7, 00:01:07 ago	Indicates the next hop address, the address of the gateway that sent the update, and the time that has elapsed since this update was received (in hours:minutes:seconds).
Route metric	This value is the best metric for this routing descriptor block.
traffic share count	Number of uses for this routing descriptor block.
AS Hops	Number of hops to the destination or to the router where the route first enters internal BGP (iBGP).

Example of Output Using the Cisco IOS Software Modularity for Layer 3 VPNs Feature

The following is sample output from the **show ip route vrf** command on routers using the Cisco IOS Software Modularity for Layer 3 VPNs feature. The output includes remote label information and corresponding MPLS flags for prefixes that have remote labels stored in the RIB, if BGP is the label distribution protocol:

```
Router# show ip route vrf v2 10.2.2.2

Routing entry for 10.2.2.2/32
  Known via "bgp 1", distance 200, metric 0, type internal
  Redistributing via ospf 2
  Advertised by ospf 2 subnets
  Last update from 10.0.0.4 00:22:59 ago
  Routing Descriptor Blocks:
  * 10.0.0.4 (Default-IP-Routing-Table), from 10.0.0.31, 00:22:59 ago
      Route metric is 0, traffic share count is 1
      AS Hops 0
      MPLS label: 1300
      MPLS Flags: MPLS Required
```

Table 74 describes the significant fields shown in the display.

Table 74 show ip route vrf Field Descriptions

Field	Description
MPLS label	Displays the BGP prefix from the BGP peer. The output shows one of the following values: • A label value (16 - 1048575) • A reserved label value, such as explicit-null or implicit-null • The word “none” if no label is received from the peer The MPLS label field does not display if any of the following conditions is true: • BGP is not the LDP. However, OSPF prefixes learned via sham link display an MPLS label. • MPLS is not supported. • The prefix was imported from another VRF, where the prefix was an IGP prefix and LDP provided the remote label for it.
MPLS Flags	The name of one of the following MPLS flags is displayed if any is set: • MPLS Required—Packets are forwarded to this prefix because the MPLS label stack is present. If MPLS is disabled in the outgoing interface, the packets are dropped. • No Global—MPLS packets for this prefix are forwarded from the VRF interface, not from the interface in global table. Using the VRF interface prevents loops in scenarios that use ieBGP multipath. • NSF—The prefix is from an NSF-aware neighbor. If the routing information temporarily disappears due to a disruption in the control plane, packets for this prefix are preserved.

The following sample output from the **show ip route vrf** command shows the repair paths marked with the tag [RPR], the best path, and the repair path in the routing table:

```
Router> show ip route vrf test1 repair-paths 192.168.3.0

Routing Table: test1
Routing entry for 192.168.3.0/24
  Known via "bgp 10", distance 20, metric 0
  Tag 100, type external
  Last update from 192.168.1.1 00:49:39 ago
  Routing Descriptor Blocks:
  * 192.168.1.1, from 192.168.1.1, 00:49:39 ago, recursive-via-conn
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 100
      MPLS label: none
    [RPR]10.4.4.4 (default), from 10.5.5.5, 00:49:39 ago, recursive-via-host
      Route metric is 0, traffic share count is 1
      AS Hops 1
      Route tag 100
      MPLS label: 29
      MPLS Flags: MPLS Required, No Global
```

Related Commands

Command	Description
show ip cache	Displays the Cisco Express Forwarding table associated with a VRF.
show ip vrf	Displays the set of defined VRFs and associated interfaces.

show tcp ha connections

To display connection-ID-to-TCP mapping data, use the **show tcp ha connections** command in privileged EXEC mode.

show tcp ha connections

Syntax Description

This command has no arguments or keywords.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.2(28)SB	This command was introduced.
15.0(1)S	This command was integrated into Cisco IOS Release 15.0(1)S.
Cisco IOS XE 3.1S	This command was integrated into Cisco IOS XE Release 3.1S.

Usage Guidelines

The **show tcp ha connections** command is used to display connection-ID-to-TCP mapping data.

Examples

The following is sample output from the **show tcp ha connections** command:

```
Router# show tcp ha connections

SSO enabled for 40 connections
TCB       Local Address           Foreign Address        (state)    Conn Id
71EACE60  10.0.56.1.179           10.0.56.3.58671         ESTAB      37
71EA9320  10.0.53.1.179           10.0.53.3.58659         ESTAB      34
71EA35F8  10.0.41.1.179           10.0.41.3.58650         ESTAB      22
71A21FE0  10.0.39.1.179           10.0.39.3.58641         ESTAB      20
71EAA6E0  10.0.54.1.179           10.0.54.3.58663         ESTAB      35
71EA2238  10.0.40.1.179           10.0.40.3.58646         ESTAB      21
71EABAA0  10.0.55.1.179           10.0.55.3.58667         ESTAB      36
71EAE710  10.0.28.1.179           10.0.28.3.58676         ESTAB      9
71EA2728  10.0.50.1.179           10.0.50.3.58647         ESTAB      31
720541D8  10.0.49.1.179           10.0.49.3.58642         ESTAB      30
71EAA1F0  10.0.44.1.179           10.0.44.3.58662         ESTAB      25
2180B3A8  10.0.33.1.179           10.0.33.3.58657         ESTAB      14
71EAB5B0  10.0.45.1.179           10.0.45.3.58666         ESTAB      26
21809FE8  10.0.32.1.179           10.0.32.3.58653         ESTAB      13
71EA8E30  10.0.43.1.179           10.0.43.3.58658         ESTAB      24
71EAD350  10.0.27.1.179           10.0.27.3.58672         ESTAB      8
2180A9C8  10.0.52.1.179           10.0.52.3.58655         ESTAB      33
2180A4D8  10.0.42.1.179           10.0.42.3.58654         ESTAB      23
71EABF90  10.0.26.1.179           10.0.26.3.58668         ESTAB      7
71EA3AE8  10.0.51.1.179           10.0.51.3.58651         ESTAB      32
720546C8  10.0.59.1.179           10.0.59.3.58643         ESTAB      40
```

Table 75 describes the significant fields shown in the display.

Table 75 show tcp ha connections Field Descriptions

Field	Description
SSO enabled for	Displays the number of TCP connections that support BGP Nonstop Routing (NSR) with SSO.
TCB	An internal identifier for the endpoint.
Local Address	The local IP address and port.
Foreign Address	The foreign IP address and port (at the opposite end of the connection).
(state)	TCP connection state. A connection progresses through a series of states during its lifetime. The states that follow are shown in the order in which a connection progresses through them. • LISTEN—Waiting for a connection request from any remote TCP and port. • SYNSENT—Waiting for a matching connection request after having sent a connection request. • SYNRCVD—Waiting for a confirming connection request acknowledgment after having both received and sent a connection request. • ESTAB—Indicates an open connection; data received can be delivered to the user. This is the normal state for the data transfer phase of the connection. • FINWAIT1—Waiting for a connection termination request from the remote TCP or an acknowledgment of the connection termination request previously sent.
Conn id	Identifying number of the TCP connection.

slow-peer detection

To use a policy template to specify a threshold time that dynamically determines a BGP slow peer, use the **slow-peer detection** command in policy template configuration mode. To restore the default value, use the **no** form of this command.

slow-peer detection [**threshold** *seconds*]

no slow-peer detection

Syntax Description

threshold *seconds*	(Optional) Specifies the threshold time in seconds that the timestamp of the oldest message in a peers queue can be lagging behind the current time before the BGP peer is determined to be a slow peer. The range is from 120 to 3600; the default is 300.

Command Default

300 seconds

Command Modes

Policy template configuration (config-router-ptmp)

Command History

Release	Modification
15.0(1)S	This command was introduced.
Cisco IOS XE 3.1S	This command was introduced.

Usage Guidelines

Update messages are timestamped when they are formatted. The timestamp of the oldest update message in a peers queue is compared to the current time to determine if the peer is lagging more than the configured number of seconds. When a peer is dynamically detected to be a slow peer, the system will send a syslog message. The peer will be marked as recovered and another syslog message will be generated only after the peer's update group converges.

Note The **neighbor slow-peer detection** command performs the same function as the **bgp slow-peer detection** command (at the address-family level), except that the **neighbor slow-peer detection** command overrides the address-family level command. When the **neighbor slow-peer detection** command is unconfigured, the system will function according to the address-family level configuration. The **slow-peer detection** command performs the same function through a peer policy template.

Examples

The following example specifies that if the timestamp on a peer's update message is more than 360 seconds before the current time, the peer that sent the update message is considered to be slow. The commands configured under the peer-policy template will be applied to the neighbor once it inherits the peer-policy.

```
Router(config)# router bgp 13
Router(config-router)# template peer-policy ipv4_ucast_pp1
Router(config-router-ptmp)# slow-peer detection threshold 360
```

```
Router(config-router-ptmp)# slow-peer split-update-group dynamic
```

Related Commands

Command	Description
bgp slow-peer detection	Specifies a threshold time that dynamically determines a slow peer.
bgp slow-peer split-update-group dynamic	Moves a dynamically detected slow peer to a slow update group.
clear ip bgp slow	Moves dynamically configured slow peers back to their original update groups.
neighbor slow-peer split-update-group dynamic	Moves a dynamically detected slow peer to a slow update group.
slow-peer split-update-group dynamic	Moves a dynamically detected slow peer to a slow update group.

slow-peer split-update-group dynamic

To use a policy template to move a dynamically detected slow peer to a slow update group, use the **slow-peer split-update-group dynamic** command in policy template configuration mode. To disable dynamically detected slow peers from being moved to a slow update group, use the **no** form of this command.

slow-peer split-update-group dynamic [**permanent**]

no slow-peer split-update-group dynamic

Syntax Description

permanent	(Optional) Specifies that after the slow peer becomes a regular peer (converges), it is not moved back to its original update group automatically. It remains in the slow update group until the network administrator uses one of the **clear slow** commands to move the peer to its original update group.

Command Default

No dynamically detected slow peer is moved to a slow peer update group.

Command Modes

Policy template (config-router-ptmp)

Command History

Release	Modification
15.0(1)S	This command was introduced.
Cisco IOS XE 3.1S	This command was introduced.

Usage Guidelines

When a peer is dynamically detected to be a slow peer, the slow peer is moved to a slow update group. If a *static* slow peer update group exists, the dynamic slow peer is moved to the static slow peer update group; otherwise, a new slow peer update group is created and the peer is moved to that group.

- We recommend you configure the **permanent** keyword. If the **permanent** keyword is configured, the peer is not automatically moved to its original update group. After you resolve the root cause of the slow peer, you can use the **clear bgp slow** command to move the peer back to its original update group.
- If the **permanent** keyword is not configured, the slow peer will be moved back to its regular original update group after it becomes a regular peer (converges).

Note The **neighbor slow-peer split-update-group dynamic** command performs the same function as the **bgp slow-peer split-update-group dynamic** command (at the address-family level), except that the **neighbor slow-peer split-update-group dynamic** command overrides the address-family level command. When the **neighbor slow-peer split-update-group dynamic** command is unconfigured, the system will function according to the address-family level configuration. The **slow-peer split-update-group dynamic** command performs the same function through a policy template.

If **slow-peer split-update-group dynamic** is configured, but no slow peer detection is configured, the detection will be done at the default threshold of 300 seconds. That is, detection is enabled automatically with its default threshold.

Examples

In the following example, the timestamp of the oldest message in a peers queue is compared to the current time to determine if the peer is lagging more than 360 seconds. If it is, the neighbor that sent the message is determined to be a slow peer, and is put in the slow peer update group. Because the **permanent** keyword is not configured, the slow peer will be moved back to its regular original update group after it becomes a regular peer (converges).

```
Router(config)# router bgp 13
Router(config-router)# template peer-policy ipv4_ucast_pp1
Router(config-router-ptmp)# slow-peer detection threshold 360
Router(config-router-ptmp)# slow-peer split-update-group dynamic
```

Related Commands

Command	Description
slow-peer detection	Specifies a threshold time that dynamically determines a slow peer.
show ip bgp template peer-policy	Displays locally configured peer policy templates.

slow-peer split-update-group static

To mark a BGP neighbor as a slow peer and move it to a slow update group, use the **slow-peer split-update-group static** command by using a peer policy template. To unmark the slow peer and return it to its original update group, use the **no** form of this command.

slow-peer split-update-group static

no slow-peer split-update-group static

Syntax Description

This command has no arguments or keywords.

Command Default

No peer is marked as slow and moved to a slow peer update group in a static manner using a peer policy template.

Command Modes

Peer policy template (config-router-ptmp)

Command History

Release	Modification
15.0(1)S	This command was introduced.
Cisco IOS XE 3.1S	This command was introduced.

Usage Guidelines

Configure a static slow peer when the peer is known to be slow (perhaps due to a slow link or low processing power).

The **neighbor slow-peer split-update-group static** command performs the same function in address-family mode.

Examples

In the following example, the neighbor is marked as a slow peer and is moved to a slow update group.

```
Router(config)# router bgp 13
Router(config-router)# template peer-policy ipv4_ucast_pp1
Router(config-router-ptmp)# slow-peer split-update-group static
```

Related Commands

Command	Description
neighbor slow-peer split-update-group static	Marks a BGP neighbor as a slow peer and moves it to a slow update group.

soo

To set the site-of-origin (SoO) value for a Border Gateway Protocol (BGP) peer policy template, use the **soo** command in policy-template configuration mode. To remove the SoO value, use the **no** form of this command.

soo *extended-community-value*

no soo

Syntax Description

extended-community-value	Specifies the VPN extended community value. The value takes one of the following formats: • A 16-bit autonomous system number, a colon, and a 32-bit number, for example: 45000:3 • A 32-bit IP address, a colon, and a 16-bit number, for example: 192.168.10.2:51 In Cisco IOS Release 12.4(24)T, 4-byte autonomous system numbers are supported in the range from 1.0 to 65535.65535 in asdot notation only. In Cisco IOS XE Release 2.4, and later releases, 4-byte autonomous system numbers are supported in the range from 65536 to 4294967295 in asplain notation and in the range from 1.0 to 65535.65535 in asdot notation. For more details about autonomous system number formats, see the **router bgp** command.

Command Default

No SoO value is set for a BGP peer policy template.

Command Modes

Policy-template configuration (config-router-ptmp)

Command History

Release	Modification
12.4(11)T	This command was introduced.
12.2(33)SRB	This command was integrated into Cisco IOS Release 12.2(33)SRB.
12.2(33)SB	This command was integrated into Cisco IOS Release 12.2(33)SB.
12.4(24)T	Support for 4-byte autonomous system numbers in asdot notation only was added.
Cisco IOS XE Release 2.4	This command was modified. Support for asplain notation was added and the default format for 4-byte autonomous system numbers is now asplain.

Release	Modification
12.2(33)SRE	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)XNE	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.

Usage Guidelines

Use this command to set the SoO value for a BGP peer policy template that a BGP neighbor can inherit. The SoO value is set for a peer policy template, and a BGP neighbor is identified under address family IPv4 VRF configuration mode to inherit the peer policy that contains the SoO value.

The SoO extended community is a BGP extended community attribute that is used to identify routes that have originated from a site so that the readvertisement of that prefix back to the source site can be prevented. The SoO extended community uniquely identifies the site from which a router has learned a route. BGP can use the SoO value associated with a route to prevent routing loops.

In releases prior to Cisco IOS Release 12.4(11)T, 12.2(33)SRB, and 12.2(33)SB, the SoO extended community attribute is configured using an inbound route map that sets the SoO value during the update process. The introduction of the **neighbor soo** and **soo** commands simplifies the SoO value configuration.

In Cisco IOS Release 12.4(24)T, the Cisco implementation of 4-byte autonomous system numbers uses asdot—1.2 for example—as the only configuration format, regular expression match, and output display, with no asplain support.

In Cisco IOS XE Release 2.4, and later releases, the Cisco implementation of 4-byte autonomous system numbers uses asplain—65538 for example—as the default regular expression match and output display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain format and the asdot format as described in RFC 5396. To change the default regular expression match and output display of 4-byte autonomous system numbers to asdot format, use the **bgp asnotation dot** command followed by the **clear ip bgp *** command to perform a hard reset of all current BGP sessions.

Note If a BGP peer inherits from several peer policy templates that specify different SoO values, the SoO value in the last template applied takes precedence and is applied to the peer. However, direct configuration of the SoO value on the BGP neighbor overrides any inherited template configurations of the SoO value.

Examples

The following example shows how to create a peer policy template and configure an SoO value as part of the peer policy. Under address family IPv4 VRF, a neighbor is identified and configured to inherit the peer policy that contains the SoO value.

```
router bgp 45000
 template peer-policy SOO_POLICY
  soo 45000:3
  exit-peer-policy
 address-family ipv4 vrf SOO_VRF
  neighbor 192.168.3.2 remote-as 50000
  neighbor 192.168.3.2 activate
  neighbor 192.168.3.2 inherit peer-policy SOO_POLICY
  end
```

The following example shows how to create a peer policy template and configure an SoO value using a 4-byte autonomous system number, 1.2 in asdot format, as part of the peer policy. Under address family IPv4 VRF, a neighbor is identified and configured to inherit the peer policy that contains the SoO value. This example requires Cisco IOS Release 12.4(24)T, Cisco IOS XE Release 2.4, or a later release.

```
router bgp 1.2
 template peer-policy SOO_POLICY
  soo 1.2:3
  exit-peer-policy
 address-family ipv4 vrf SOO_VRF
  neighbor 192.168.3.2 remote-as 1.14
  neighbor 192.168.3.2 activate
  neighbor 192.168.3.2 inherit peer-policy SOO_POLICY
  end
```

Related Commands

Command	Description
address-family ipv4 (BGP)	Enters address family configuration mode to configure a routing session using standard IP Version 4 address prefixes.
neighbor soo	Sets the SoO value for a BGP neighbor or peer group.
router bgp	Configures the BGP routing process.
template peer-policy	Creates a peer policy template and enters policy-template configuration mode.

synchronization

To enable the synchronization between BGP and your Interior Gateway Protocol (IGP) system, use the **synchronization** command in address family or router configuration mode. To enable the Cisco IOS software to advertise a network route without waiting for the IGP, use the **no** form of this command.

synchronization

no synchronization

Syntax Description

This command has no arguments or keywords.

Defaults

The behavior of this command is disabled by default.

Command Modes

Address family configuration
Router configuration

Command History

Release	Modification
10.0	This command was introduced.
12.0(7)T	Address family configuration mode was added.
12.2(8)T	Command default behavior changed to disabled.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

Usually, a BGP speaker does not advertise a route to an external neighbor unless that route is local or exists in the IGP. By default, synchronization between BGP and the IGP is turned off to allow the Cisco IOS software to advertise a network route without waiting for route validation from the IGP. This feature allows routers and access servers within an autonomous system to have the route before BGP makes it available to other autonomous systems.

Use the **synchronization** command if routers in the autonomous system do not speak BGP.

Examples

The following example shows how to enable synchronization in router configuration mode. The router validates the network route in its IGP before advertising the route externally.

```
router bgp 65120
 synchronization
```

The following example shows how to enable synchronization in address family configuration mode. The router validates the network route in its IGP before advertising the route externally.

```
router bgp 65120
address-family ipv4 unicast
 synchronization
```

Related Commands

Command	Description
address-family ipv4 (BGP)	Places the router in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard IP Version 4 address prefixes.
address-family vpnv4	Places the router in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard VPN Version 4 address prefixes.

table-map

To modify metric and tag values when the IP routing table is updated with BGP learned routes, use the **table-map** command in address family or router configuration mode. To disable this function, use the **no** form of the command.

table-map *map-name*

no table-map *map-name*

Syntax Description

map-name	Route map name from the **route-map** command.

Defaults

This command is disabled by default.

Command Modes

Address family configuration
Router configuration

Command History

Release	Modification
10.0	This command was introduced.
12.0(7)T	Address family configuration mode was added.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

This command adds the route map name defined by the **route-map** command to the IP routing table. This command is used to set the tag name and the route metric to implement redistribution.

You can use **match** clauses of route maps in the **table-map** command. IP access list, autonomous system paths, and next hop match clauses are supported.

Examples

In the following router configuration mode example, the Cisco IOS software is configured to automatically compute the tag value for the BGP learned routes and to update the IP routing table:

```
route-map tag
 match as path 10
 set automatic-tag
!
router bgp 100
 table-map tag
```

In the following address family configuration mode example, the Cisco IOS software is configured to automatically compute the tag value for the BGP learned routes and to update the IP routing table:

```
route-map tag
 match as path 10
 set automatic-tag
!
router bgp 100
address-family ipv4 unicast
 table-map tag
```

Related Commands

Command	Description
address-family ipv4 (BGP)	Places the router in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard IP Version 4 address prefixes.
address-family vpn4	Places the router in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard VPN Version 4 address prefixes.
match as-path	Matches a BGP autonomous system path access list.
match ip address	Distributes any routes that have a destination network number address that is permitted by a standard or extended access list, and performs policy routing on packets.
match ip next-hop	Redistributes any routes that have a next hop router address passed by one of the access lists specified.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.

template peer-policy

To create a peer policy template and enter policy-template configuration mode, use the **template peer-policy** command in router configuration mode. To remove a peer policy template, use the **no** form of this command.

template peer-policy *policy-template-name*

no template peer-policy *policy-template-name*

Syntax Description

policy-template-name	Name or tag for the peer policy template.

Defaults

Removing a peer policy template by using the **no** form of this command removes all policy configurations inside of the template.

Command Modes

Router configuration

Command History

Release	Modification
12.0(24)S	This command was introduced.
12.2(18)S	This command was integrated into Cisco IOS Release 12.2(18)S.
12.3(4)T	This command was integrated into Cisco IOS Release 12.3(4)T.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

Peer policy templates are used to group and apply the configuration of commands that are applied within specific address-families and NLRI configuration mode. Peer policy templates are created and configured in peer policy configuration mode. BGP policy commands that are configured for specific address-families or NLRI configuration modes are configured in a peer policy template. The following BGP policy commands are supported by peer policy templates:

- **advertisement-interval**
- **allowas-in**
- **as-override**
- **capability**
- **default-originate**
- **distribute-list**
- **dmzlink-bw**
- **exit-peer-policy**
- **filter-list**

- **inherit peer-policy**
- **maximum-prefix**
- **next-hop-self**
- **next-hop-unchanged**
- **prefix-list**
- **remove-private-as**
- **route-map**
- **route-reflector-client**
- **send-community**
- **send-label**
- **soft-reconfiguration**
- **unsuppress-map**
- **weight**

Peer policy templates are used to configure BGP policy commands that are configured for neighbors that belong to specific address-families and NLRI configuration modes. Like peer session templates, peer policy templates are configured once and then applied to many neighbors through the direct application of a peer policy template or through inheritance from peer policy templates. The configuration of peer policy templates simplifies the configuration of BGP policy commands that are applied to all neighbors within an autonomous system.

Peer policy templates support direct and indirect inheritance from up to eight peer policy templates. Inherited peer policy templates are configured with sequence numbers like route-maps. An inherited peer policy template, like a route-map, is evaluated starting with the inherit statement with the lowest sequence number and ending with the highest sequence number. However, there is a difference; a peer policy template will not fall through like a route-map. Every sequence is evaluated, and if a BGP policy command is reapplied with different value, it will overwrite any previous value from a lower sequence number.

Peer policy templates support only general policy commands. BGP policy configuration commands that are configured only for specific address families or NLRI configuration modes are configured with peer policy templates.

Note A BGP neighbor cannot be configured to work with both peer groups and peer templates. A BGP neighbor can be configured to belong only to a peer group or to inherit policies from only peer templates.

Examples

The following example creates a peer policy template named CUSTOMER-A. This peer policy template is configured to inherit the configuration from the peer policy templates named PRIMARY-IN and GLOBAL.

```
Router(config-router)# template peer-policy CUSTOMER-A
Router(config-router-ptmp)# route-map SET-COMMUNITY in
Router(config-router-ptmp)# filter-list 20 in
Router(config-router-ptmp)# inherit peer-policy PRIMARY-IN 20
Router(config-router-ptmp)# inherit peer-policy GLOBAL 10
Router(config-router-ptmp)# exit-peer-policy
Router(config-router)#
```

Related Commands

Command	Description
advertisement-interval	Sets the minimum interval between the sending of BGP routing updates.
allowas-in	Configures PE routers to allow readvertisement of all prefixes containing duplicate autonomous system numbers.
as-override	Configures a PE router to override the ASN of a site with the ASN of a provider.
capability orf prefix-list	Configures outbound route filtering and advertises the capability to send and receive ORF updates to the neighbor routers.
default-originate	Originates a default route to the local router.
distribute-list	Distributes BGP neighbor information as specified in an access list.
dmzlink-bw	Advertises the bandwidth of links that are used to exit an autonomous system.
exit peer-policy	Exits policy-template configuration mode and enters router configuration mode.
filter-list	Sets up a BGP filter.
inherit peer-policy	Configures a peer policy template to inherit the configuration from another peer policy template.
maximum-prefix	Controls how many prefixes can be received from a neighbor.
neighbor inherit peer-policy	Configures a router to send a peer policy template to a neighbor so that the neighbor can inherit the configuration.
neighbor send-label	Enables a BGP router to send MPLS labels with BGP routes to a neighboring BGP router.
next-hop-self	Disables next-hop processing of BGP updates on the router.
next-hop-unchanged	Propagates the next- hop unchanged for iBGP paths to this router.
prefix-list	Specifies a prefix list, a CLNS filter set, or a CLNS filter expression to be used to filter BGP advertisements.
remove-private-as	Removes the private autonomous system number from outbound routing updates.
route-map	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.
route-reflector-client	Configures the router as a BGP route reflector and configures the specified neighbor as its client.
send-community	Specifies that the BGP community attribute should be sent to the specified neighbor.
show ip bgp template peer-policy	Displays locally configured peer policy templates.
show ip bgp template peer-session	Displays locally configured peer session templates.
soft-reconfiguration	Configures the Cisco IOS software to start storing updates.
template peer-session	Creates a peer session template and enters session-template configuration mode.
unsuppress-map	Selectively unsuppresses surpressed routes.
weight	Assigns a weight to a neighbor connection.

template peer-session

To create a peer session template and enter session-template configuration mode, use the **template peer-session** command in router configuration mode. To remove a peer session template, use the **no** form of this command.

template peer-session *session-template-name*

no template peer-session *session-template-name*

Syntax Descriptions

session-template-name	Name or tag for the peer session template.

Defaults

Removing a peer session template by using the **no** form of this command removes all session command configurations inside of the template.

Command Modes

Address family configuration
Router configuration

Command History

Release	Modification
12.0(24)S	This command was introduced.
12.2(18)S	This command was integrated into Cisco IOS Release 12.2(18)S.
12.3(4)T	This command was integrated into Cisco IOS Release 12.3(4)T.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

Peer session templates are used to group and apply the configuration of general session commands to groups of neighbors that share common session configuration elements. General session commands that are common for neighbors that are configured in different address families can be configured within the same peer session template. Peer session templates are created and configured in peer session configuration mode. Only general session commands can be configured in a peer session template. The following general session commands are supported by peer session templates:

- **description**
- **disable-connected-check**
- **ebgp-multihop**
- **exit peer-session**
- **inherit peer-session**
- **local-as**
- **password**

- **remote-as**
- **shutdown**
- **timers**
- **translate-update**
- **update-source**
- **version**

General session commands can be configured once in a peer session template and then applied to many neighbors through the direct application of a peer session template or through indirect inheritance from a peer session template. The configuration of peer session templates simplify the configuration of general session commands that are commonly applied to all neighbors within an autonomous system.

Peer session templates support direct and indirect inheritance. A peer can be configured with only one peer session template at a time, and that peer session template can contain only one indirectly inherited peer session template. However, each inherited session template can also contain one indirectly inherited peer session template. So, only one directly applied peer session template and up to seven additional indirectly inherited peer session templates can be applied, allowing you to apply up to a maximum of eight peer session configurations to a neighbor: the configuration from the directly inherited peer session template and the configurations from up to seven indirectly inherited peer session templates. Inherited peer session templates are evaluated first, and the directly applied template will be evaluated and applied last. So, if a general session command is reapplied with a different value, the subsequent value will have priority and overwrite the previous value that was configured in the indirectly inherited template.

Peer session templates support only general session commands. BGP policy configuration commands that are configured only for specific address families or NLRI configuration modes are configured with peer policy templates.

Note A BGP neighbor cannot be configured to work with both peer groups and peer templates. A BGP neighbor can be configured only to belong to a peer group or to inherit policies from peer templates.

Examples

The following example creates a peer session template named CORE1. This example inherits the configuration of the peer session template named INTERNAL-BGP.

```
Router(config-router)# template peer-session CORE1
Router(config-router-stmp)# description CORE-123
Router(config-router-stmp)# update-source loopback 1
Router(config-router-stmp)# inherit peer-session INTERNAL-BGP
Router(config-router-stmp)# exit-peer-session
Router(config-router)#
```

Related Commands

Command	Description
description	Configures a description to be displayed by the local or a peer router.
disable-connected-check	Disables connection verification for eBGP peers no more than one hop away when the eBGP peer is configured with a loopback interface.
ebgp-multihop	Accepts or initiates BGP connections to external peers residing on networks that are not directly connected.
exit peer-session	Exits session-template configuration mode and enters router configuration mode.

Command	Description
inherit peer-session	Configures a peer session template to inherit the configuration from another peer session template.
local-as	Allows the customization of the autonomous system number for eBGP peer groupings.
neighbor inherit peer-session	Configures a router to send a peer session template to a neighbor so that the neighbor can inherit the configuration.
neighbor translate-update	Upgrades a router running BGP in the NLRI format to support multiprotocol BGP.
password	Enables MD5 authentication on a TCP connection between two BGP peers.
remote-as	Adds an entry to the BGP or multiprotocol BGP neighbor table.
show ip bgp template peer-policy	Displays locally configured peer policy templates.
show ip bgp template peer-session	Displays locally configured peer session templates.
shutdown	Disables a neighbor or peer group.
timers bgp	Adjusts BGP network timers.
update-source	Specifies that the Cisco IOS software allow internal BGP sessions to use any operational interface for TCP connections.
version	Configures the Cisco IOS software to accept only a particular BGP version.

timers bgp

To adjust BGP network timers, use the **timers bgp** command in router configuration mode. To reset the BGP timing defaults, use the **no** form of this command.

timers bgp *keepalive holdtime* [*min-holdtime*]

no timers bgp

Syntax Description

keepalive	Frequency (in seconds) with which the Cisco IOS software sends *keepalive* messages to its peer. The default is 60 seconds. The range is from 0 to 65535.
holdtime	Interval (in seconds) after not receiving a *keepalive* message that the software declares a peer dead. The default is 180 seconds. The range is from 0 to 65535.
min-holdtime	(Optional) Interval (in seconds) specifying the minimum acceptable hold-time from a BGP neighbor. The minimum acceptable hold-time must be less than, or equal to, the interval specified in the *holdtime* argument. The range is from 0 to 65535.

Defaults

keepalive: 60 seconds
holdtime: 180 seconds

Command Modes

Router configuration

Command History

Release	Modification
10.0	This command was introduced.
12.0(26)S	The *min-holdtime* argument was added.
12.3(7)T	The *min-holdtime* argument was added.
12.2(22)S	The *min-holdtime* argument was added.
12.2(27)SBC	The *min-holdtime* argument was added and this command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(33)SRA	The *min-holdtime* argument was added and this command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SXH	The *min-holdtime* argument was added and this command was integrated into Cisco IOS Release 12.2(33)SXH.

Usage Guidelines

When configuring the *holdtime* argument for a value of less than twenty seconds, the following warning is displayed:

```
% Warning: A hold time of less than 20 seconds increases the chances of peer flapping
```

If the minimum acceptable hold-time interval is greater than the specified hold-time, a notification is displayed:

```
% Minimum acceptable hold time should be less than or equal to the configured hold time
```

Note When the minimum acceptable hold-time is configured on a BGP router, a remote BGP peer session is established only if the remote peer is advertising a hold-time that is equal to, or greater than, the minimum acceptable hold-time interval. If the minimum acceptable hold-time interval is greater than the configured hold-time, the next time the remote session tries to establish, it will fail and the local router will send a notification stating "unacceptable hold time."

Examples

The following example changes the keepalive timer to 70 seconds, the hold-time timer to 130 seconds, and the minimum acceptable hold-time interval to 100 seconds:

```
router bgp 45000
 timers bgp 70 130 100
```

Related Commands

Command	Description
clear ip bgp peer-group	Removes all the members of a BGP peer group.
router bgp	Configures the BGP routing process.
show ip bgp	Displays entries in the BGP routing table.