



BGP Commands_ show ip through Z

- [show ip bgp, page 2](#)
- [show ip bgp ipv4, page 14](#)
- [show ip bgp neighbors, page 18](#)
- [show ip bgp paths, page 39](#)
- [show ip bgp summary, page 41](#)
- [show ip bgp template peer-policy, page 49](#)
- [show ip bgp template peer-session, page 52](#)
- [show ip community-list, page 55](#)
- [show ip extcommunity-list, page 57](#)
- [show ip route, page 61](#)
- [template peer-session, page 73](#)
- [timers bgp, page 77](#)

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

<i>ip-address</i>	(Optional) IP address entered to filter the output to display only a particular host or network in the BGP routing table.
<i>mask</i>	(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-prefixes	(Optional) Displays the specified route and all less-specific routes.
<i>length</i>	(Optional) The prefix length. The range is a number from 0 to 32.
bestpath	(Optional) Displays the best path 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 <i>name</i>	(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 <i>name</i>	(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.
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.
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.

Release	Modification
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)SX11	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 was changed 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 was changed 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.
15.2(1)S	This command was modified to display an RPKI validation code per network, if one applies.
Cisco IOS XE Release 3.5S	This command was modified to display an RPKI validation code per network, if one applies.
15.1(1)SG	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
Cisco IOS XE Release 3.3SG	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
15.2(4)S	This command was modified. Output about discarded or unknown path attributes was added for the BGP Attribute Filter feature. Output about additional path selection was added for the BGP Additional Paths feature. Output about paths imported from a VRF table to the global table was added for the BGP Support for IP Prefix Export from a VRF Table into the Global Table.

Release	Modification
Cisco IOS XE Release 3.7S	This command was modified. Output about discarded or unknown path attributes was added for the BGP Attribute Filter feature. Output about additional path selection was added for the BGP Additional Paths feature. Output about paths imported from a VRF table to the global table was added for the BGP Support for IP Prefix Export from a VRF Table into the Global Table.
15.1(1)SY	This command was integrated into Cisco IOS Release 15.1(1)SY.
15.2(1)E	This command was integrated into Cisco IOS Release 15.2(1)E.

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)SX11, 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

Examples

The following sample output shows the BGP routing table:

```
Router# show ip bgp

BGP table version is 6, local router ID is 10.0.96.2
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, a additional-path
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

      Network          Next Hop              Metric LocPrf Weight Path
-----
N*  10.0.0.1           10.0.0.3                  0             0 3 ?
N*>
Nr  10.0.0.0/8         10.0.0.3                  0             0 3 ?
Nr>
```

```

Nr> 10.0.0.0/24      10.0.0.3          0          0 3 ?
V*> 10.0.2.0/24      0.0.0.0           0          32768 i
Vr> 10.0.3.0/24      10.0.3.5          0          0 4 ?

```

The table below describes the significant fields shown in the display.

Table 1: 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 a backup path to use for that network. • x—The table entry has a best external route to use for the network.

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: • a—Path is selected as an additional path. • 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.
RPKI validation codes	If shown, the RPKI validation state for the network prefix, which is downloaded from the RPKI server. The codes are shown only if the bgp rpki server or neighbor announce rpki state command is configured.
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.

Examples

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)SX11, 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
```

Examples

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, 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
```

The table below describes the significant fields shown in the display.

Table 2: show ip bgp Field Descriptions

Field	Description
BGP routing table entry for	IP address or network number of the routing table entry.

Field	Description
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 load sharing 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.

Examples

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 ?
```

show ip bgp

```

*> 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

```

Examples

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          0 109 108 ?
*                   10.92.72.30     8896          0 109 108 ?
*> 10.92.1.0        10.92.72.30     8796          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 ?

```

Examples

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      0 200 ?

```

Examples

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      0 200 ?

```

Examples

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      0 200 ?

```

Examples

The following output indicates (for each neighbor) whether any of the additional path tags (group-best, all, best 2 or best 3) are applied to the path. A line of output indicates rx pathid (received from neighbor) and tx pathid (announcing to neighbors). Note that the "Path advertised to update-groups:" is now per-path when the BGP Additional Paths feature is enabled.

```

Device# show ip bgp 10.0.0.1 255.255.255.224

BGP routing table entry for 10.0.0.1/28, version 82
Paths: (10 available, best #5, table default)
  Path advertised to update-groups:
    21      25
  Refresh Epoch 1
  20 50, (Received from a RR-client)

```

```

192.0.2.1 from 192.0.2.1 (192.0.2.1)
  Origin IGP, metric 200, localpref 100, valid, internal, all
  Originator: 192.0.2.1, Cluster list: 2.2.2.2
  mpls labels in/out 16/nolabel
  rx pathid: 0, tx pathid: 0x9
Path advertised to update-groups:
  18      21
Refresh Epoch 1
30
192.0.2.2 from 192.0.2.2 (192.0.2.2)
  Origin IGP, metric 200, localpref 100, valid, internal, group-best, all
  Originator: 192.0.2.2, Cluster list: 4.4.4.4
  mpls labels in/out 16/nolabel
  rx pathid: 0x1, tx pathid: 0x8
Path advertised to update-groups:
  16      18      19      20      21      22      24
  25      27
Refresh Epoch 1
10
192.0.2.3 from 192.0.2.3 (192.0.2.3)
  Origin IGP, metric 200, localpref 100, valid, external, best2, all
  mpls labels in/out 16/nolabel
  rx pathid: 0, tx pathid: 0x7
Path advertised to update-groups:
  20      21      22      24      25
Refresh Epoch 1
10
192.0.2.4 from 192.0.2.4 (192.0.2.4)
  Origin IGP, metric 300, localpref 100, valid, external, best3, all
  mpls labels in/out 16/nolabel
  rx pathid: 0, tx pathid: 0x6
Path advertised to update-groups:
  10      13      17      18      19      20      21
  22      23      24      25      26      27      28
Refresh Epoch 1
10
192.0.2.5 from 192.0.2.5 (192.0.2.5)
  Origin IGP, metric 100, localpref 100, valid, external, best
  mpls labels in/out 16/nolabel
  rx pathid: 0, tx pathid: 0x0
Path advertised to update-groups:
  21
Refresh Epoch 1
30
192.0.2.6 from 192.0.2.6 (192.0.2.6)
  Origin IGP, metric 200, localpref 100, valid, internal, all
  Originator: 192.0.2.6, Cluster list: 5.5.5.5
  mpls labels in/out 16/nolabel
  rx pathid: 0x1, tx pathid: 0x5
Path advertised to update-groups:
  18      23      24      26      28
Refresh Epoch 1
60 40, (Received from a RR-client)
192.0.2.7 from 192.0.2.7 (192.0.2.7)
  Origin IGP, metric 250, localpref 100, valid, internal, group-best
  Originator: 192.0.2.7, Cluster list: 3.3.3.3
  mpls labels in/out 16/nolabel
  rx pathid: 0x2, tx pathid: 0x2
Path advertised to update-groups:
  25
Refresh Epoch 1
30 40, (Received from a RR-client)
192.0.2.8 from 192.0.2.8 (192.0.2.8)
  Origin IGP, metric 200, localpref 100, valid, internal, all
  Originator: 192.0.2.8, Cluster list: 2.2.2.2
  mpls labels in/out 16/nolabel
  rx pathid: 0x1, tx pathid: 0x3
Path advertised to update-groups:
  18      21      23      24      25      26      28
Refresh Epoch 1
20 40, (Received from a RR-client)
192.0.2.9 from 192.0.2.9 (192.0.2.9)
  Origin IGP, metric 200, localpref 100, valid, internal, group-best, all

```

```

    Originator: 192.0.2.9, Cluster list: 2.2.2.2
    mpls labels in/out 16/nolabel
    rx pathid: 0x1, tx pathid: 0x4
Path advertised to update-groups:
  21
Refresh Epoch 1
30 40
  192.0.2.9 from 192.0.2.9 (192.0.2.9)
    Origin IGP, metric 100, localpref 100, valid, internal, all
    Originator: 192.0.2.9, Cluster list: 4.4.4.4
    mpls labels in/out 16/nolabel
    rx pathid: 0x1, tx pathid: 0x1

```

Examples

The following is sample output from the **show ip bgp** command that displays unknown and discarded path attributes:

```

Router# show ip bgp 192.0.2.0/32

BGP routing table entry for 192.0.2.0/32, version 0
Paths: (1 available, no best path)
  Refresh Epoch 1
  Local
    192.168.101.2 from 192.168.101.2 (192.168.101.2)
      Origin IGP, localpref 100, valid, internal
      unknown transitive attribute: flag 0xE0 type 0x81 length 0x20
        value 0000 0000 0000 0000 0000 0000 0000 0000
              0000 0000 0000 0000 0000 0000 0000 0000

      unknown transitive attribute: flag 0xE0 type 0x83 length 0x20
        value 0000 0000 0000 0000 0000 0000 0000 0000
              0000 0000 0000 0000 0000 0000 0000 0000

      discarded unknown attribute: flag 0x40 type 0x63 length 0x64
        value 0000 0000 0000 0000 0000 0000 0000 0000
              0000 0000 0000 0000 0000 0000 0000 0000

```

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 bgp community new-format	Configures BGP to display communities in the format AA:NN.
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.
clear ip bgp	Resets BGP connections using hard or soft reconfiguration.

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** *route-distinguisher* | **vrf** *vrf-name*} | **mvpn** {**all** | **rd** *route-distinguisher* | **vrf** *vrf-name*} | **unicast** *prefix* | **multicast** *prefix* | **tunnel**}

Syntax Description

mdt	Displays entries for multicast distribution tree (MDT) sessions.
all	Displays all the entries in the routing table.
rd <i>route-distinguisher</i>	Displays information about the specified VPN route distinguisher.
vrf <i>vrf-name</i>	Displays information about the specified VRF.
mvpn	Displays entries for multicast VPN (MVPN) sessions.
unicast	Displays entries for unicast sessions.
<i>prefix</i>	Displays entries for the specified prefix.
multicast	Displays entries for multicast sessions.
tunnel	Displays entries for tunnel sessions.

Command Modes

Privileged EXEC (#)

Command History

Release	Modification
12.0(7)T	This command was introduced.
12.0(29)S	This command was modified. 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.

Release	Modification
12.4(20)T	This command was modified. The mdt keyword was added.
15.2(1)S	This command was modified. An RPKI validation code is displayed per network, if one applies.
Cisco IOS XE 3.5S	This command was modified. An RPKI validation code is displayed per network, if one applies.
Cisco IOS XE 3.7S	This command was modified. Imported paths from a VRF table to the global routing table are displayed, if any.
15.2(4)S	This command was implemented on the Cisco 7200 series routers.
Cisco IOS XE 3.8S	This command was modified. The mvpn 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
```

The table below describes the significant fields shown in the display.

Table 3: 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.

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 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 Border Gateway Protocol (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.

Field	Description
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

The following is sample output from the **show ip bgp ipv4 unicast** *prefix* command. The output indicates the imported path information from a VRF named vpn1.

Device# **show ip bgp ipv4 unicast 150.1.1.0**

BGP routing table entry for 150.1.1.0/24, version 2

Paths: (1 available, best #1, table default)

Not advertised to any peer

Refresh Epoch 1

65002, imported path from 1:1:150.1.1.0/24 (vpn1)

4.4.4.4 (metric 11) from 4.4.4.4 (4.4.4.4)

Origin IGP, metric 0, localpref 100, valid, internal, best

Extended Community: RT:1:1

mpls labels in/out nlabel/16

Related Commands

Command	Description
clear ip bgp ipv4 mdt	Resets MDT IPv4 BGP address-family sessions.
export map	Exports IP prefixes from a VRF table into the global table.
show ip bgp	Displays entries in the BGP routing table.

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**}|**vpn4 all**|**vpn6 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	(Optional) Displays peers in the IPv4 address family.
multicast	(Optional) Specifies IPv4 multicast address prefixes.
unicast	(Optional) Specifies IPv4 unicast address prefixes.
vpn4 all	(Optional) Displays peers in the VPNv4 address family.
vpn6 unicast all	(Optional) Displays peers in the VPNv6 address family.
slow	(Optional) Displays information about dynamically configured slow peers.
<i>ip-address</i>	(Optional) IP address of the IPv4 neighbor. If this argument is omitted, information about all neighbors is displayed.
<i>ipv6-address</i>	(Optional) IP address of 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 <i>reg-exp</i>	(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	Mainline and T Release	Modification
	10.0	This command was introduced.
	11.2	This command was modified. The received-routes keyword was added.
	12.2(4)T	This command was modified. The received and prefix-filter keywords were added.
	12.2(15)T	This command was modified. Support for the display of BGP graceful restart capability information was added.
	12.3(7)T	This command was modified. The command output was modified to support the BGP TTL Security Check feature and to display explicit-null label information.
	12.4(4)T	This command was modified. Support for the display of Bidirectional Forwarding Detection (BFD) information was added.
	12.4(11)T	This command was modified. Support for the policy and detail keywords was added.
	12.4(20)T	This command was modified. The output was modified to support BGP TCP path MTU discovery.

Mainline and T Release	Modification
12.4(24)T	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation was added.

Command History

S Release	Modification
12.0(18)S	This command was modified. The output was modified to display the no-prepend configuration option.
12.0(21)ST	This command was modified. The output was modified to display Multiprotocol Label Switching (MPLS) label information.
12.0(22)S	This command was modified. Support for the display of BGP graceful restart capability information was added. Support for the Cisco 12000 series routers (Engine 0 and Engine 2) was also added.
12.0(25)S	This command was modified. The policy and detail keywords were added.
12.0(27)S	This command was modified. The command output was modified to support the BGP TTL Security Check feature and to display explicit-null label information.
12.0(31)S	This command was modified. Support for the display of BFD information was added.
12.0(32)S12	This command was modified. Support for displaying 4-byte autonomous system numbers in asdot notation 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.
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	This command was modified. Support for the display of BFD information was added.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
12.2(33)SRA	This command was modified. The output was modified to support BGP TCP path Maximum Transmission Unit (MTU) discovery.
12.2(33)SRB	This command was modified. Support for the policy and detail keywords was added.

S Release	Modification
12.2(33)SXH	This command was modified. Support for displaying BGP dynamic neighbor information was added.
12.2(33)SRC	This command was modified. Support for displaying BGP graceful restart information was added.
12.2(33)SB	This command was modified. Support for displaying BFD and the BGP graceful restart per peer information was added, 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. Support for displaying BGP best external and BGP additional path features information was added. 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.
15.0(1)SY	This command was integrated into Cisco IOS Release 15.0(1)SY.
15.1(1)S	This command was modified. The Layer 2 VPN address family is displayed if graceful restart or nonstop forwarding (NSF) is enabled.
15.1(1)SG	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain notation was added and the default display format became asplain.
15.2(4)S	This command was modified and implemented on the Cisco 7200 series router. The configured discard and treat-as-withdraw attributes are displayed, along with counts of incoming Updates with a matching discard attribute or treat-as-withdraw attribute, and number of times a malformed Update is treat-as-withdraw. The capabilities of the neighbor to send and receive additional paths that are advertised or received are added.
15.1(2)SNG	This command was implemented on the Cisco ASR 901 Series Aggregation Services Routers.
15.2(1)E	This command was integrated into Cisco IOS Release 15.2(1)E.

Command History

Cisco IOS XE	Modification
Cisco IOS XE Release 2.1	This command was integrated into Cisco IOS XE Release 2.1.

Cisco IOS XE	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 became asplain.
Cisco IOS XE Release 3.1S	This command was modified. The slow keyword was added.
Cisco IOS XE Release 3.6S	This command was modified. Support for displaying BGP BFD multihop and C-bit information was added.
Cisco IOS XE Release 3.3SG	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 3.7S	This command was implemented on the Cisco ASR 903 router and the output modified. The configured discard and treat-as-withdraw attributes are displayed, along with counts of incoming Updates with a matching discard attribute or treat-as-withdraw attribute, and number of times a malformed Update is treat-as-withdraw. The capabilities of the neighbor to send and receive additional paths that are advertised or received are added.
Cisco IOS XE Release 3.8S	This command was modified. In support of the BGP Multi-Cluster ID feature, the cluster ID of a neighbor is displayed if the neighbor is assigned a cluster.

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, determining which policies are applied to the neighbor can be difficult.

In Cisco IOS Release 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.

Examples

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.

```
Device# 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 NLRI 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
The table below 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 4: 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.
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 a network administrator is 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 hh:mm:ss, that the underlying TCP connection has been in existence.
Last read	Time, in hh:mm:ss, since BGP last received a message from this neighbor.
last write	Time, in hh:mm:ss, since BGP last sent a message to this neighbor.

Field	Description
hold time	Time, in seconds, that BGP will maintain the session with this neighbor without receiving 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.
Revd	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.
For address family:	Address family to which the following fields refer.

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.
neighbor version	Number used by the software to track prefixes that have been sent and those that need to be sent.
1 update-group member	Number of the 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 a prefix has been withdrawn because it is no longer feasible.
Used as bestpath	Number of received prefixes installed as best paths.
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.

Field	Description
* 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 0.
* NEXT_HOP Martian	Displays outbound martian denials.
* NEXT_HOP non-local	Displays outbound nonlocal 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.
* 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 best path 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	Displays inbound denials because the best path came from an iBGP neighbor.
* Incorrect RIB for CE	Displays inbound denials due to RIB errors for a customer edge (CE) router.
* BGP distribute-list	Displays inbound denials due to a distribute list.

Field	Description
Number of NLRI...	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, in hh:mm:ss, since this peering session was last reset. The reason for the reset is displayed on this line.
External BGP neighbor may be...	Indicates that the BGP time to live (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.
unread input bytes	Number of bytes of packets still to be processed.
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 of times a packet is dropped due to no acknowledgment.
PmtuAger	Path MTU discovery timer.

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 higher 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:	Shortest recorded round-trip timeout (hard-wire value used for calculation).
maxRTT:	Longest 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.

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

Examples

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)SX11, 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
```

Examples

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

```
Device# 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
```

The table below describes the significant fields shown in the display.

Table 5: 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.

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 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.

Examples

The following is sample output from the **show ip bgp neighbors** command entered with the **check-control-plane-failure** option configured:

```
Device# show ip bgp neighbors 10.10.10.1

BGP neighbor is 10.10.10.1, remote AS 10, internal link
  Fall over configured for session
  BFD is configured. BFD peer is Up. Using BFD to detect fast fallover (single-hop) with
  c-bit check-control-plane-failure.
  Inherits from template cbit-tps for session parameters
  BGP version 4, remote router ID 10.7.7.7
  BGP state = Established, up for 00:03:55
  Last read 00:00:02, last write 00:00:21, 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
  Enhanced Refresh Capability: advertised and received
  Multisession Capability:
  Stateful switchover support enabled: NO for session 1

```

Examples

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

```
Device# show ip bgp neighbors 172.29.232.178 paths 10
```

```

Address      Refcount Metric Path
0x60E577B0      2      40 10 ?

```

The table below describes the significant fields shown in the display.

Table 6: 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.

Examples

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:

```
Device# 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

```

The table below describes the significant fields shown in the display.

Table 7: 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.

Examples

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 policies are policies that the neighbor inherits from a peer group or a peer-policy template.

```
Device# 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 policies:
  prefix-list NO-MARKETING in
  route-map ROUTE in
  weight 300
  maximum-prefix 10000
```

Examples

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:

```
Device# show ip bgp neighbors
BGP neighbor is 172.16.10.2, remote AS 45000, external link
.
.
.
Using BFD to detect fast fallover
```

Examples

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:

```
Device# 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
```

Examples

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:

```
Device# 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
.
.
.

```

Examples

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.

```

Device# 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 multiseession 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

```

Examples

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.

```

Device# 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

```

show ip bgp neighbors

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 multiseession 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)

Multiseession 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 NLRI 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 NLRI's 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
Minimum 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

```

Examples

The following is sample output from the **show ip bgp neighbors** command that indicates the discard attribute values and treat-as-withdraw attribute values configured. It also provides a count of received Updates matching a treat-as-withdraw attribute, a count of received Updates matching a discard attribute, and a count of received malformed Updates that are treat-as-withdraw.

```

Device# show ip bgp vpnv4 all neighbors 10.0.103.1

BGP neighbor is 10.0.103.1, remote AS 100, internal link
  Path-attribute treat-as-withdraw inbound
  Path-attribute treat-as-withdraw value 128
  Path-attribute treat-as-withdraw 128 in: count 2
  Path-attribute discard 128 inbound
  Path-attribute discard 128 in: count 2

      Outbound      Inbound
Local Policy Denied Prefixes:  -----  -----
MALFORM treat as withdraw:      0          1
Total:                          0          1

```

Examples

The following output indicates that the neighbor is capable of advertising additional paths and sending additional paths it receives. It is also capable of receiving additional paths and advertised paths.

```

Device# 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

```

show ip bgp neighbors

Last read 00:00:24, last write 00:00:24, hold time is 180, keepalive interval is 60 seconds

Neighbor capabilities:
 Additional paths Send: advertised and received
 Additional paths Receive: advertised and received
 Route refresh: advertised and received(old & new)
 Graceful Restart Capabilty: advertised and received
 Address family IPv4 Unicast: advertised and received

Examples

In the following output, the cluster ID of the neighbor is displayed. (The vertical bar and letter “i” for “include” cause the device to display only lines that include the user's input after the “i”, in this case, “cluster-id.”) The cluster ID displayed is the one directly configured through a neighbor or a template.

Device# **show ip bgp neighbors 192.168.2.2 | i cluster-id**

Configured with the cluster-id 192.168.15.6

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 enhanced-error	Restores the default behavior of treating Update messages that have a malformed attribute as withdrawn, or includes iBGP peers in the Enhanced Attribute Error Handling feature.
neighbor path-attribute discard	Configures the device to discard unwanted Update messages from the specified neighbor that contain a specified path attribute.
neighbor path-attribute treat-as-withdraw	Configures the device to withdraw from the specified neighbor unwanted Update messages that contain a specified attribute.
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 *regex*

Syntax Description

<i>regex</i>	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 ?
```

The table below describes the significant fields shown in the display.

Table 8: 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 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 all routing topology instances.
<i>routing-topology-instance-name</i>	(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	This command was modified. - 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.

Release	Modification
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.4(11)T	This command was modified. 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.
12.2(33)SX11	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.
15.2(1)S	This command was modified. It will show information about how many paths are in each RPKI state.
Cisco IOS XE Release 3.5S	This command was modified. It will show information about how many paths are in each RPKI state.
15.1(1)SG	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.

Release	Modification
Cisco IOS XE Release 3.3SG	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
15.2(4)S	This command was implemented on the Cisco 7200 series routers.

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)SX11, 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
```

The table below describes the significant fields shown in the display. Fields that are preceded by the asterisk (*) are not shown in the above output.

Table 9: 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 rinfo 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.

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.

Field	Description
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 peer group 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)SX11, 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)SX11, 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
```

The following example displays counts of prefix/AS pairs for each RPKI state. The fourth line of output indicates "Path RPKI states: x valid, x not found, x invalid." Of course the line of output indicating RPKI states can be displayed only if the **bgp rpki server** command or the **neighbor announce rpki state** command is configured.

```
Router> show ip bgp summary
```

```
For address family: IPv4 Unicast
BGP router identifier 10.0.96.2, local AS number 2
BGP table version is 8, main routing table version 8
Path RPKI states: 0 valid, 7 not found, 0 invalid
6 network entries using 888 bytes of memory
7 path entries using 448 bytes of memory
3/3 BGP path/bestpath attribute entries using 384 bytes of memory
2 BGP AS-PATH 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 1768 total bytes of memory
BGP activity 12/0 prefixes, 14/0 paths, scan interval 60 secs
```

Neighbor /PfxRcd	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State
10.0.0.3 3	4	3	6	9	8	0	0	00:01:04	
10.0.2.4 0	4	2	5	8	8	0	0	00:01:15	
10.0.3.5 3	4	4	6	7	8	0	0	00:01:14	
10.0.96.254	4	1	0	0	1	0	0	never	Idle

```
For address family: IPv6 Unicast
BGP router identifier 10.0.96.2, local AS number 2
BGP table version is 9, main routing table version 9
Path RPKI states: 3 valid, 4 not found, 0 invalid
6 network entries using 1032 bytes of memory
7 path entries using 616 bytes of memory
5/5 BGP path/bestpath attribute entries using 640 bytes of memory
2 BGP AS-PATH 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 2336 total bytes of memory
BGP activity 12/0 prefixes, 14/0 paths, scan interval 60 secs
```

Neighbor /PfxRcd	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State
2001::2 2	4	2	6	9	6	0	0	00:01:08	
2002::1 2	4	3	7	11	9	0	0	00:01:07	
2003::2	4	4	6	8	9	0	0	00:01:08	

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

<i>policy-template-name</i>	(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.

Release	Modification
12.2(33)SB	Support for the detail keyword was integrated into Cisco IOS Release 12.2(33)SB.
15.1(1)SY	This command was integrated into Cisco IOS Release 15.1(1)SY.

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.

```
Device# 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
```

The table below describes the significant fields shown in the display.

Table 10: 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.

Field	Description
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.

```
Device# 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 or privileged EXEC mode.

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

Syntax Description

<i>session-template-name</i>	(Optional) Name of a locally configured peer session template.
------------------------------	--

Command Default

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.
Cisco IOS XE Release 3.8S	This command was modified. The cluster ID for the template is displayed.

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.

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

Template:INTERNAL-BGP, index:1
Local policies:0x21, Inherited policies: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 policies: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
```

The table below describes the significant fields shown in the display.

Table 11: 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 policies:	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.

The following sample output displays the cluster ID assigned to the template:

```
Device# show ip bgp template peer-session TS1

Template:TS1, index:1
Local policies:0x10000000, Inherited policies:0x0
Locally configured session commands:
```

```
cluster-id 192.168.0.115
Inherited session commands:
```

Related Commands

Command	Description
bgp cluster-id	Sets the global cluster ID on a route reflector.
inherit peer-session	Configures a peer session template to inherit the configuration from another peer session template.
neighbor cluster-id	Sets the cluster ID for a neighbor.
template peer-session	Creates a peer session template and enters session-template configuration mode.

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

<i>community-list-number</i>	(Optional) A standard or expanded community list number in the range from 1 to 500.
<i>community-list-name</i>	(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
```

The Field Descriptions table below describes the significant fields shown in the display.

Table 12: 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

<i>list-number</i>	(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.
<i>list-name</i>	(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.

Release	Modification
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)SX11	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.
15.1(1)SG	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
Cisco IOS XE Release 3.3SG	This command was modified. Support for displaying 4-byte autonomous system numbers in asplain and asdot notation was added.
15.2(1)E	This command was integrated into Cisco IOS Release 15.2(1)E.

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)SX11, 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*

```

The table below describes the significant fields shown in the display.

Table 13: 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)SX11, 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)SX11, 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 route

To display contents 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

<i>ip-address</i>	(Optional) IP address for which routing information should be displayed.
repair-paths	(Optional) Displays the repair paths.
next-hop-override	(Optional) Displays the Next Hop Resolution Protocol (NHRP) next-hop overrides that are associated with a particular route and the corresponding default next hops.
dhcp	(Optional) Displays routes added by the Dynamic Host Configuration Protocol (DHCP) server.
<i>mask</i>	(Optional) Subnet mask.
longer-prefixes	(Optional) Displays output for longer prefix entries.
<i>protocol</i>	(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 , nhrp , or rip .
<i>process-id</i>	(Optional) Number used to identify a process of the specified protocol.
list	(Optional) Filters output by an access list name or number.
<i>access-list-number</i>	(Optional) Access list number.
<i>access-list-name</i>	(Optional) Access list name.
static	(Optional) Displays static routes.
download	(Optional) Displays routes 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	This command was modified. The “D—EIGRP, EX—EIGRP, N1—SPF NSSA external type 1 route” and “N2—OSPF NSSA external type 2 route” codes were included in the command output.
10.3	This command was modified. The <i>process-id</i> argument was added.
11.0	This command was modified. The longer-prefixes keyword was added.
11.1	This command was modified. The “U—per-user static route” code was included in the command output.
11.2	This command was modified. The “o—on-demand routing” code was included in the command output.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA, and the update-queue keyword was added.
11.3	This command was modified. The command output was enhanced to display the origin of an IP route in Intermediate System-to-Intermediate System (IS-IS) networks.
12.0(1)T	This command was modified. The “M—mobile” code was included in the command output.
12.0(3)T	This command was modified. The “P—periodic downloaded static route” code was included in the command output.
12.0(4)T	This command was modified. The “ia—IS-IS” code was included in the command output.
12.2(2)T	This command was modified. The command output was enhanced to display information on multipaths to the specified network.

Release	Modification
12.2(13)T	This command was modified. The <i>egp</i> and <i>igrp</i> arguments were removed because the Exterior Gateway Protocol (EGP) and the Interior Gateway Routing Protocol (IGRP) were no longer available in Cisco 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	This command was modified. The command output was enhanced to display route tag information.
12.3(8)T	This command was modified. The command 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(33)SRE	This command was modified. The dhcp and repair-paths keywords were 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. The next-hop-override and nhrp keywords were added.
15.2(2)S	This command was modified. The command output was enhanced to display route tag values in dotted decimal format.
Cisco IOS XE Release 3.6S	This command was modified. The command output was enhanced to display route tag values in dotted decimal format.
15.2(4)S	This command was implemented on the Cisco 7200 series router.
15.1(1)SY	This command was integrated into Cisco IOS Release 15.1(1)SY.

Examples

Examples

The following is sample output from the **show ip route** command when an IP address is not specified:

```
Device# 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,
```

show ip 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 sample output from the **show ip route** command includes routes learned from IS-IS Level 2:

```
Device# 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 not set
10.89.0.0 is subnetted (mask is 255.255.255.0), 3 subnets
C 10.89.64.0 255.255.255.0 is possibly down,
routing via 10.0.0.0, Ethernet0
i L2 10.89.67.0 [115/20] via 10.89.64.240, 0:00:12, Ethernet0
i L2 10.89.66.0 [115/20] via 10.89.64.240, 0:00:12, Ethernet0

```

The following is sample output from the **show ip route ip-address mask longer-prefixes** command. When this keyword is included, the address-mask pair becomes the prefix, and any address that matches that prefix is displayed. Therefore, multiple addresses are displayed. 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 with 10.0.0.0. Any destinations that fall into that range are displayed in the output.

```
Device# 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 sample outputs from the **show ip route** command display all downloaded static routes. A "p" indicates that these routes were installed using the AAA route download function.

Device# **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 172.16.17.1 to network 10.0.0.0

```

    172.31.0.0/32 is subnetted, 1 subnets
P    172.31.229.41 is directly connected, Dialer1 10.0.0.0/8 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

```

Device# **show ip route static**

```

    172.16.4.0/8 is variably subnetted, 2 subnets, 2 masks
P    172.16.1.1/32 is directly connected, BRI0
P    172.16.4.0/8 [1/0] via 10.1.1.1, BRI0
S    172.31.0.0/16 [1/0] via 172.16.114.65, Ethernet0
S    10.0.0.0/8 is directly connected, BRI0
P    10.0.0.0/8 is directly connected, BRI0
    172.16.0.0/16 is variably subnetted, 5 subnets, 2 masks
S    172.16.114.201/32 is directly connected, BRI0
S    172.16.114.205/32 is directly connected, BRI0
S    172.16.114.174/32 is directly connected, BRI0
S    172.16.114.12/32 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 172.16.114.65, Ethernet0
S    172.16.0.0/16 [1/0] via 172.16.114.65, Ethernet0

```

The following sample output from the **show ip route static download** command displays all active and inactive routes installed using the AAA route download function:

Device# **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 remotel
I    10.38.1.9 255.255.255.0 192.168.69.1

```

The following sample outputs from the **show ip route nhrp** command display shortcut switching on the tunnel interface:

Device# **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
```

Device# **show ip route nhrp**

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

The following are sample outputs from the **show ip route** command when the **next-hop-override** keyword is used. When this keyword is included, the NHRP next-hop overrides that are associated with a particular route and the corresponding default next hops are displayed.

```
=====
1) Initial configuration
=====
```

Device# **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
```

Device# **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

Device# **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
172.16.0.0/12	drop	

2) Add a next-hop override

address = 10.10.10.0
mask = 255.255.255.0
gateway = 10.1.1.1
interface = Tunnel0

Device# **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

Device# **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

Device# **show ip cef**

Prefix	Next Hop	Interface
.		
.		
.		
10.2.1.255/32	receive	Loopback1
10.10.10.0/24		

show ip route

```

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 next-hop override
   address = 10.10.10.0
   mask = 255.255.255.0
   gateway = 10.11.1.1
   interface = Tunnel0
=====

```

Device# **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

```

Device# **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

```

Device# **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		
.		
.		
.		

The table below describes the significant fields shown in the displays:

Table 14: show ip route Field Descriptions

Field	Description
Codes (Protocol)	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—On-demand routing • 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 • +—Replicated route
Codes (Type)	Type of route. It can be one of the following values: • *—Indicates the last path used when a packet was forwarded. This information is specific to nonfast-switched packets. • E1—OSPF external type 1 route • E2—OSPF external type 2 route • IA—OSPF interarea 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

Field	Description
10.110.0.0	Indicates the address of the remote network.
[160/5]	The first number in 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 device 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.

Examples

The following is sample output from the **show ip route** command when an IP address is specified:

```
Device# 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, the router includes one of its IP addresses to be used as the originator IP address. When other routers calculate IP routes, they 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 the origin of a particular IP route in your network. In the preceding example, the route to 10.0.0.1/32 was originated by a device with IP address 10.191.255.247.

The table below describes the significant fields shown in the display.

Table 15: 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.
Redistributing via...	Indicates the redistribution protocol.

Field	Description
Last update from 10.191.255.251	Indicates the IP address of the router that is the next hop to the remote network and the 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	Indicates the number of packets transmitted over various routes.

The following sample output from the **show ip route** command displays the tag applied to the route 10.22.0.0/16. You must specify an IP prefix to see the tag value. The fields in the display are self-explanatory.

```
Device# 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
```

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 172.16.2.2 [1/0] are static, and route 10.0.0.0/0 is a default route candidate. The fields in the display are self-explanatory.

```
Device# 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 172.16.2.2 [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 repair paths marked with the tag [RPR]. The fields in the display are self-explanatory:

Device# **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

```

Device# **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 interfaces tunnel	Displays tunnel interface information.
show ip route summary	Displays the current state of the routing table in summary format.

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 Description

<i>session-template-name</i>	Name or tag for the peer session template.
------------------------------	--

Command Default

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.
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

<i>keepalive</i>	Frequency (in seconds) with which the Cisco IOS software sends <i>keepalive</i> messages to its peer. The default is 60 seconds. The range is from 0 to 65535.
<i>holdtime</i>	Interval (in seconds) after not receiving a <i>keepalive</i> message that the software declares a peer dead. The default is 180 seconds. The range is from 0 to 65535.
<i>min-holdtime</i>	(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 <i>holdtime</i> argument. The range is from 0 to 65535.

Command Default

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 <i>min-holdtime</i> argument was added.
12.3(7)T	The <i>min-holdtime</i> argument was added.
12.2(22)S	The <i>min-holdtime</i> argument was added.
12.2(27)SBC	The <i>min-holdtime</i> argument was added and this command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(33)SRA	The <i>min-holdtime</i> argument was added and this command was integrated into Cisco IOS Release 12.2(33)SRA.

Release	Modification
12.2(33)SXH	The <i>min-holdtime</i> 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.