



show b

- [show banner](#), on page 2
- [show bfd drops](#), on page 3
- [show bfd map](#), on page 4
- [show bfd neighbors](#), on page 5
- [show bfd summary](#), on page 6
- [show bgp](#), on page 8
- [show bgp cidr-only](#), on page 14
- [show bgp community](#), on page 15
- [show bgp community-list](#), on page 16
- [show bgp filter-list](#), on page 18
- [show bgp injected-paths](#), on page 19
- [show bgp ipv4 unicast](#), on page 20
- [show bgp ipv6 unicast](#), on page 21
- [show bgp ipv4/ipv6 unicast community](#), on page 23
- [show bgp ipv4/ipv6 unicast community-list](#), on page 25
- [show bgp ipv4/ ipv6 unicast neighbors](#), on page 26
- [show bgp ipv4/ ipv6 unicast paths](#), on page 32
- [show bgp ipv4/ ipv6 unicast prefix-list](#), on page 34
- [show bgp ipv4/ ipv6 unicast regexp](#), on page 35
- [show bgp ipv4/ ipv6 unicast route-map](#), on page 36
- [show bgp ipv4/ ipv6 unicast summary](#), on page 37
- [show bgp neighbors](#), on page 39
- [show bgp paths](#), on page 48
- [show bgp prefix-list](#), on page 49
- [show bgp regexp](#), on page 50
- [show bgp rib-failure](#), on page 51
- [show bgp summary](#), on page 53
- [show bgp update-group](#), on page 57
- [show blocks](#), on page 60
- [show bootvar](#), on page 65
- [show bridge-group](#), on page 66

show banner

To display the configured banner message, enter the **show banner** command.

```
show banner [login]
```

Syntax Description

login	Displays the banner that has been set up for the password login prompt.
--------------	---

Command History

Release	Modification
6.1	This command was introduced.

Examples

```
> show banner
```

show bfd drops

To display the numbered of dropped packets in BFD, use the **show bfd drops** command.

show bfd drops

Command History	Release	Modification
	6.3	This command was introduced.

Examples

The following example displays the BFD dropped packets.

```
> show bfd drops
BFD Drop Statistics

```

	IPV4	IPV6	IPV4-M	IPV6-M
Invalid TTL	0	0	0	0
BFD Not Configured	0	0	0	0
No BFD Adjacency	0	0	0	0
Invalid Header Bits	0	0	0	0
Invalid Discriminator	0	0	0	0
Session AdminDown	0	0	0	0
Authen invalid BFD ver	0	0	0	0
Authen invalid len	0	0	0	0
Authen invalid seq	0	0	0	0
Authen failed	0	0	0	0

Related Commands	Command	Description
	clear bfd counters	Clears the BFD counters.
	show bfd map	Displays the configured BFD maps.
	show bfd neighbors	Displays a line-by-line listing of existing BFD adjacencies.
	show bfd summary	Displays summary information for BFD.

show bfd map

To display the configured BFD maps, use the **show bfd map** command.

show bfd map

Command History	Release	Modification
	6.3	This command was introduced.

Examples

The following example displays the BFD maps.

```
> show bfd map
Destination: 40.40.40.2/24
Source: 50.50.50.2/24
Template: mh
Authentication(Type): sha-1
```

Related Commands	Command	Description
	clear bfd counters	Clears the BFD counters.
	show bfd drops	Displays the numbered of dropped packets in BFD.
	show bfd neighbors	Displays a line-by-line listing of existing BFD adjacencies.
	show bfd summary	Displays summary information for BFD.

show bfd neighbors

To display a line-by-line listing of existing BFD adjacencies, use the **show bfd neighbors** command.

```
show bfd neighbors [client bgp] [ipv4 [ip_address] | ipv6 [ipv6_address] | multihop-ipv4
[ip_address] | multihop-ipv6 [ipv6_address]] [inactive] [detail]
```

Syntax	Description
client bgp	(Optional) Displays the neighbors of the BGP client.
ipv4 [ip_address]	(Optional) Displays single-hop IPv4 neighbors. You can optionally specify a particular neighbor address.
ipv6 [ipv6_address]	(Optional) Displays single-hop IPv6 neighbors. You can optionally specify a particular neighbor address.
multihop-ipv4 [ip_address]	(Optional) Displays multi-hop IPv4 neighbors. You can optionally specify a particular neighbor address.
multihop-ipv6 [ipv6_address]	(Optional) Displays multi-hop IPv6 neighbors. You can optionally specify a particular neighbor address.
inactive	(Optional) Displays the inactive adjacencies.
detail	(Optional) Displays all BFD protocol parameters and timers for each neighbor.

Command History	Release	Modification
	6.3	This command was introduced.

Examples

The following example displays the BFD neighbors.

```
> show bfd neighbors
OurAddr      NeighAddr    LD/RD  RH      Holdown(mult)  State Int
172.16.10.1  172.16.10.2  1/6    1       260 (3 )      Up    Fa0/1
```

Related Commands	Command	Description
	clear bfd counters	Clears the BFD counters.
	show bfd drops	Displays the numbered of dropped packets in BFD.
	show bfd map	Displays the configured BFD maps.
	show bfd summary	Displays summary information for BFD.

show bfd summary

To display summary information for BFD, use the **show bfd summary** command.

show bfd summary [**client** | **session**]

Syntax Description	client	(Optional) Displays the BFD summary for clients.
	session	(Optional) Displays the BFD summary for sessions.

Command History	Release	Modification
	6.3	This command was introduced.

Usage Guidelines Use this command to display summary information about BFD, BFD clients, or BFD sessions. When a BFD client launches a session with a peer, BFD sends periodic BFD control packets to the peer. Information about the following states of a session are included in the output of this command:

- **Up**—When another BFD interface acknowledges the BFD control packets, the session moves into an Up state.
- **Down**—The session and the data path are declared down if a data path failure occurs and BFD does not receive a control packet within the configured amount of time. When a session is down, BFD notifies the BFD client so that the client can perform necessary actions to reroute the traffic.

Examples

The following example displays the BFD summaries.

```
> show bfd summary
      Session      Up      Down
Total      1          1          0

> show bfd summary session
Protocol Session  Up  Down
IPV4      1          1    0
Total     1          1    0

> show bfd summary client
Client  Session  Up    Down
BGP     1          1     0
EIGRP   1          1     0
Total   2          2     0
```

Related Commands	Command	Description
	clear bfd counters	Clears the BFD counters.
	show bfd drops	Displays the numbered of dropped packets in BFD.
	show bfd map	Displays the configured BFD maps.

Command	Description
show bfd neighbors	Displays a line-by-line listing of existing BFD adjacencies.

show bgp

To display entries in the Border Gateway Protocol (BGP) routing table, use the **show bgp** command.

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

Syntax Description

<i>ip-address</i>	(Optional) Specifies the network in the BGP routing table to display.
<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 value for this argument is a number from 0 to 32.
bestpath	(Optional) Displays the bestpath for this prefix.
multipaths	(Optional) Displays multipaths for this prefix.
subnets	(Optional) Displays the subnet routes for the specified prefix.
all	(Optional) Displays all address family information in the BGP routing table.
prefix-list name	(Optional) Filters the output based on the specified prefix list.
pending-prefixes	(Optional) Displays prefixes that are pending deletion from the BGP routing table.
route-map name	(Optional) Filters the output based on the specified route map.
[vrf name all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command History

Release	Modification
6.1	This command was introduced.
6.6	The [vrf name all] keywords were added.

Usage Guidelines

The **show 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.

Examples

The following sample output shows the BGP routing table:

```
> show bgp
BGP table version is 22, local router ID is 10.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale, m multipath, b backup-path, x best-external
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop          Metric LocPrf Weight Path
*> 10.1.1.1/32    0.0.0.0           0         32768 i
*>i10.2.2.2/32    172.16.1.2        0         100     0 i
*bi10.9.9.9/32    192.168.3.2       0         100     0 10 10 i
*>                192.168.1.2       0         100     0 10 10 i
* i172.16.1.0/24 172.16.1.2        0         100     0 i
*>                0.0.0.0           0         32768 i
*> 192.168.1.0    0.0.0.0           0         32768 i
*>i192.168.3.0    172.16.1.2        0         100     0 i
*bi192.168.9.0    192.168.3.2       0         100     0 10 10 i
*>                192.168.1.2       0         100     0 10 10 i
*bi192.168.13.0   192.168.3.2       0         100     0 10 10 i
*>                192.168.1.2       0         100     0 10 10 i
```

The following table explains each field.

Table 1: show bgp Fields

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: <ul style="list-style-type: none"> • s—The table entry is suppressed. • d—The table entry is dampened. • h—The table entry history. • *—The table entry is valid. • >—The table entry is the best entry to use for that network. • i—The table entry was learned via an internal BGP (iBGP) session. • r—The table entry is a RIB-failure. • S—The table entry is stale. • m—The table entry has multipath to use for that network. • b—The table entry has backup path to use for that network. • x—The table entry has best external route to use for the network.

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: <ul style="list-style-type: none"> • i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised. • 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. 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.

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

```
> show 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:

```
> show bgp 10.3.3.3 255.255.255.255
BGP routing table entry for 10.3.3.3/32, version 35
Paths: (3 available, best #2, table default)
Multipath: eBGP
Flag: 0x860
  Advertised to update-groups:
    1
  200
```

```

10.71.8.165 from 10.71.8.165 (192.168.0.102)
  Origin incomplete, localpref 100, valid, external, backup/repair
  Only allowed to recurse through connected route
200
10.71.11.165 from 10.71.11.165 (192.168.0.102)
  Origin incomplete, localpref 100, weight 100, valid, external, best
  Only allowed to recurse through connected route
200
10.71.10.165 from 10.71.10.165 (192.168.0.104)
  Origin incomplete, localpref 100, valid, external,
  Only allowed to recurse through connected route

```

The following table explains each field.

Table 2: show bgp (4 byte autonomous system numbers) Fields

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

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

```

> show bgp all

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

   Network          Next Hop          Metric LocPrf Weight Path
*> 10.1.1.0/24      0.0.0.0            0         32768 ?
*> 10.13.13.0/24    0.0.0.0            0         32768 ?
*> 10.15.15.0/24    0.0.0.0            0         32768 ?

```

```

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

```

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

```
> show bgp 10.92.0.0 255.255.0.0 longer-prefixes
```

```

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

```

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

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

```

> show bgp 172.16.0.0/16 shorter-prefixes 8
*> 172.16.0.0      10.0.0.2          0 ?
*                  10.0.0.2          0 200 ?

```

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

```

> show bgp prefix-list ROUTE

BGP table version is 39, local router ID is 10.0.0.1

```

```
Status codes:s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes:i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 192.168.1.0	10.0.0.2			0	?
*	10.0.0.2	0		0	200 ?

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

```
> show bgp route-map LEARNED_PATH
```

```
BGP table version is 40, local router ID is 10.0.0.1
Status codes:s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes:i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 192.168.1.0	10.0.0.2			0	?
*	10.0.0.2	0		0	200 ?

show bgp cidr-only

To display routes with classless inter domain routing (CIDR), use the **show bgp cidr-only** command.

show bgp cidr-only [*vrf name* | **all**]

Syntax Description

[*vrf name* | **all**]

If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the **vrf name** keyword. If you want the command to affect all virtual routers, include the **all** keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command History

Release

Modification

6.1

This command was introduced.

6.6

The [*vrf name* | **all**] keywords were added.

Examples

The following is sample output from the **show bgp cidr-only** command. For an explanation of the output, see the **show bgp** command.

```
> show bgp cidr-only
```

```
BGP table version is 220, local router ID is 172.16.73.131
```

```
Status codes: s suppressed, * valid, > best, i - internal
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```

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

```

show bgp community

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

```
show bgp community [vrf name | all] [community-number] [exact-match] [no-advertise]
[no-export]
```

Syntax Description	
<i>community-number</i>	(Optional) Valid value is a community number in the range from 1 to 4294967295 or AA:NN (autonomous system-community number:2-byte number).
exact-match	(Optional) Displays only routes that have an exact match.
no-advertise	(Optional) Displays only routes that are not advertised to any peer (well-known community).
no-export	(Optional) Displays only routes that are not exported outside of the local autonomous system (well-known community).
[<i>vrf name</i> all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [vrf name all] keywords were added.

Examples

The following is sample output from the **show bgp community** command. For an explanation of the output, see the **show bgp** command.

```
> show bgp community 111:12345
BGP table version is 10, local router ID is 224.0.0.10
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 172.16.2.2/32    10.43.222.2        0             0 222 ?
*> 10.0.0.0         10.43.222.2        0             0 222 ?
*> 10.43.0.0        10.43.222.2        0             0 222 ?
*> 10.43.44.44/32   10.43.222.2        0             0 222 ?
* 10.43.222.0/24    10.43.222.2        0             0 222 i
*> 172.17.240.0/21  10.43.222.2        0             0 222 ?
*> 192.168.212.0    10.43.222.2        0             0 222 i
*> 172.31.1.0       10.43.222.2        0             0 222 ?
```

show bgp community-list

To display routes that are permitted by the Border Gateway Protocol (BGP) community list, use the **show bgp community-list** command.

```
show bgp community-list [vrf name | all] {community-list-number | community-list-name [exact-match] }
```

Syntax Description

<i>community-list-number</i>	A standard or expanded community list number in the range from 1 to 500.
<i>community-list-name</i>	Community list name. The community list name can be standard or expanded.
exact-match	(Optional) Displays only routes that have an exact match.
[vrf name all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command History

Release	Modification
6.1	This command was introduced.
6.6	The [vrf name all] keywords were added.

Examples

The following is sample output of the **show bgp community-list**. For an explanation of the output, see the **show bgp** command.

```
> show bgp community-list 20
BGP table version is 716977, local router ID is 192.168.32.1
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

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

```
*>i      10.0.16.1      0   100      0 1800 701 i
*         10.92.72.24 0 1878 704 701 i
```

show bgp filter-list

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

```
show bgp filter-list [vrf name | all] access-list-name
```

Syntax Description		
	<i>access-list-name</i>	Name of an autonomous system path access list. Valid values are from 1 to 500.
	[vrf <i>name</i> all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf <i>name</i> keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [vrf <i>name</i> all] keywords were added.

Examples

The following is sample output of the **show bgp filter-list** command. For an explanation of the output, see the **show bgp** command.

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

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

show bgp injected-paths

To display all the injected paths in the Border Gateway Protocol (BGP) routing table, use the **show bgp injected-paths** command.

```
show bgp injected-paths [vrf name | all]
```

Syntax Description	[vrf name all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.
---------------------------	------------------	---

Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [vrf name all] keywords were added.

Examples

The following is sample output from the **show bgp injected-paths** command. For an explanation of the output, see the **show bgp** command.

```
> show bgp injected-paths
BGP table version is 11, local router ID is 10.0.0.1
Status codes:s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes:i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf Weight Path
*> 172.16.0.0       10.0.0.2             0      0   0  ?
*> 172.17.0.0/16    10.0.0.2             0      0   0  ?
```

show bgp ipv4 unicast

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

```
show bgp ipv4 unicast [vrf name | all] [cidr-only]
```

Syntax Description	unicast	Specifies IPv4 unicast address prefixes.
	cidr-only	(Optional) Displays routes with non-natural netmasks.
	[vrf name all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [vrf name all] keywords were added.

Examples

The following is sample output from the **show bgp ipv4 unicast** command. For an explanation of the output, see the **show bgp** command.

```
> show 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
```

show bgp ipv6 unicast

To display entries in the IPv6 Border Gateway Protocol (BGP) routing table, use the **show bgp ipv6** command.

```
show bgp ipv6 unicast [vrf name | all] [ipv6-prefix/prefix-length] [longer-prefixes] [labels]
```

Syntax	Description
unicast	Specifies IPv6 unicast address prefixes.
<i>ipv6-prefix</i>	(Optional) IPv6 network number, entered to display a particular network in the IPv6 BGP routing table. This argument must be in the form documented in RFC 2373 where the address is specified in hexadecimal using 16-bit values between colons.
<i>/prefix-length</i>	(Optional) The length of the IPv6 prefix. A decimal value that indicates how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). A slash mark must precede the decimal value.
longer-prefixes	(Optional) Displays the route and more specific routes.
labels	(Optional) Displays the policies applied to this neighbor per address family.
[vrf name all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [vrf name all] keywords were added.

Examples

The following is sample output from the **show bgp ipv6 unicast** command, showing information for prefix 3FFE:500::/24. For an explanation of the output, see the **show bgp** command.

```
> show bgp ipv6 unicast 3FFE:500::/24
BGP routing table entry for 3FFE:500::/24, version 19421
Paths: (6 available, best #1)
 293 3425 2500
   3FFE:700:20:1::11 from 3FFE:700:20:1::11 (192.168.2.27)
     Origin IGP, localpref 100, valid, external, best
4554 293 3425 2500
   3FFE:C00:E:4::2 from 3FFE:C00:E:4::2 (192.168.1.1)
     Origin IGP, metric 1, localpref 100, valid, external
33 293 3425 2500
   3FFE:C00:E:5::2 from 3FFE:C00:E:5::2 (209.165.18.254)
     Origin IGP, localpref 100, valid, external
6175 7580 2500
```

show bgp ipv6 unicast

```

3FFE:C00:E:1::2 from 3FFE:C00:E:1::2 (209.165.223.204)
  Origin IGP, localpref 100, valid, external
1849 4697 2500, (suppressed due to dampening)
  3FFE:1100:0:CC00::1 from 3FFE:1100:0:CC00::1 (172.31.38.102)
    Origin IGP, localpref 100, valid, external
237 10566 4697 2500
  3FFE:C00:E:B::2 from 3FFE:C00:E:B::2 (172.31.0.3)
    Origin IGP, localpref 100, valid, external
> show bgp ipv6 unicast
BGP table version is 28, local router ID is 172.10.10.1
Status codes:s suppressed, 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
*>i4004::/64      ::FFFF:172.11.11.1
                                     0   100   0 ?
* i                ::FFFF:172.30.30.1
                                     0   100   0 ?

```

show bgp ipv4/ipv6 unicast community

To display entries in the IPv4 or IPv6 Border Gateway Protocol (BGP) routing table, use the **show bgp ipv4 unicast community** or **show bgp ipv6 unicast community** command respectively.

```
show bgp [vrf name | all] {ipv4 | ipv6} unicast community [community-number]
[exact-match] [local-as | no-advertise | no-export]
```

Syntax	Description
unicast	Specifies IPv4 or IPv6 unicast address prefixes.
<i>community-number</i>	(Optional) Valid value is a community number in the range from 1 to 4294967295 or AA:NN (autonomous system-community number:2-byte number).
exact-match	(Optional) Displays only routes that have an exact match.
local-as	(Optional) Displays only routes that are not sent outside of the local autonomous system (well-known community).
no-advertise	(Optional) Displays only routes that are not advertised to any peer (well-known community).
no-export	(Optional) Displays only routes that are not exported outside of the local autonomous system (well-known community).
[vrf name all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [vrf name all] keywords were added.

Examples

The following is sample output from the **show bgp ipv6 unicast community** command. For an explanation of the output, see the **show bgp** command.

```
BGP table version is 69, local router ID is 10.2.64.5
Status codes:s suppressed, h history, * valid, > best, i - internal
Origin codes:i - IGP, e - EGP, ? - incomplete
      Network                Next Hop                Metric LocPrf Weight Path
*> 2001:0DB8:0:1::1/64        ::                        0 32768 i
*> 2001:0DB8:0:1:1::/80       ::                        0 32768 ?
*> 2001:0DB8:0:2::/64         2001:0DB8:0:3::2        0 2 i
*> 2001:0DB8:0:2:1::/80       2001:0DB8:0:3::2        0 2 ?
* 2001:0DB8:0:3::1/64        2001:0DB8:0:3::2        0 2 ?
*>                             ::                        0 32768 ?
```

show bgp ipv4/ipv6 unicast community

```
*> 2001:0DB8:0:4::/64      2001:0DB8:0:3::2          0 2 ?
*> 2001:0DB8:0:5::1/64    ::                          0 32768 ?
*> 2001:0DB8:0:6::/64    2000:0:0:3::2            0 2 3 i
*> 2010::/64              ::                          0 32768 ?
*> 2020::/64              ::                          0 32768 ?
*> 2030::/64              ::                          0 32768 ?
*> 2040::/64              ::                          0 32768 ?
*> 2050::/64              ::                          0 32768 ?
```

show bgp ipv4/ipv6 unicast community-list

To display routes that are permitted by the IPv4 or IPv6 Border Gateway Protocol (BGP) community list, use the **show bgp ipv4 unicast community-list** or **show bgp ipv6 unicast community-list** command respectively.

```
show bgp [vrf name | all] {ipv4 | ipv6} unicast community-list {number | name}
[exact-match]
```

Syntax Description	unicast	Specifies IPv4 or IPv6 unicast address prefixes.
	<i>number</i>	Community list number in the range from 1 to 199.
	<i>name</i>	Community list name.
	exact-match	(Optional) Displays only routes that have an exact match.
	[<i>vrf name</i> all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [vrf name all] keywords were added.

Examples

The following is sample output of the **show bgp ipv6 unicast community-list** command for community list number 3. For an explanation of the output, see the **show bgp** command.

```
> show bgp ipv6 unicast community-list 3
BGP table version is 14, local router ID is 10.2.64.6
Status codes: s suppressed, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network                               Next Hop                               Metric LocPrf Weight Path
*> 2001:0DB8:0:1::/64                    2001:0DB8:0:3::1                       0 1 i
*> 2001:0DB8:0:1:1::/80                  2001:0DB8:0:3::1                       0 1 i
*> 2001:0DB8:0:2::1/64                   ::                                       0 32768 i
*> 2001:0DB8:0:2:1::/80                  ::                                       0 32768 ?
* 2001:0DB8:0:3::2/64                    2001:0DB8:0:3::1                       0 1 ?
*>                                       ::                                       0 32768 ?
*> 2001:0DB8:0:4::2/64                   ::                                       0 32768 ?
*> 2001:0DB8:0:5::/64                    2001:0DB8:0:3::1                       0 1 ?
*> 2010::/64                             2001:0DB8:0:3::1                       0 1 ?
*> 2020::/64                             2001:0DB8:0:3::1                       0 1 ?
*> 2030::/64                             2001:0DB8:0:3::1                       0 1 ?
*> 2040::/64                             2001:0DB8:0:3::1                       0 1 ?
*> 2050::/64                             2001:0DB8:0:3::1                       0 1 ?
```

show bgp ipv4/ ipv6 unicast neighbors

To display information about IPv4 or IPv6 Border Gateway Protocol (BGP) connections to neighbors, use the **show bgp ipv4 unicast neighbors** or **show bgp ipv6 neighbors** command.

```
show bgp [vrf name | all] {ipv4 | ipv6} unicast neighbors [ip-address] [received-routes
| routes | advertised-routes | paths regular-expression]
```

Syntax Description

unicast	Specifies IPv4 or IPv6 unicast address prefixes.
<i>ip-address</i>	(Optional) Address of the IPv4 or IPv6 BGP-speaking neighbor. If you omit this argument, all IPv4 or IPv6 neighbors are displayed. IPv6 addresses must be in the form documented in RFC 2373 where the address is specified in hexadecimal using 16-bit values between colons.
received-routes	(Optional) Displays all received routes (both accepted and rejected) from the specified neighbor.
routes	(Optional) Displays all routes received and accepted. This is a subset of the output from the received-routes keyword.
advertised-routes	(Optional) Displays all the routes the networking device advertised to the neighbor.
paths regular-expression	(Optional) Regular expression used to match the paths received.
[vrf name all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command History

Release	Modification
6.1	This command was introduced.
6.6	The [vrf name all] keywords were added.

Examples

The following is sample output from the **show bgp ipv6 unicast neighbors** command.

```
> show bgp ipv6 unicast neighbors
BGP neighbor is 3FFE:700:20:1::11, remote AS 65003, external link
  BGP version 4, remote router ID 192.168.2.27
  BGP state = Established, up for 13:40:17
  Last read 00:00:09, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 31306 messages, 20 notifications, 0 in queue
```

```

Sent 14298 messages, 1 notifications, 0 in queue
Default minimum time between advertisement runs is 30 seconds
For address family: IPv6 Unicast
BGP table version 21880, neighbor version 21880
Index 1, Offset 0, Mask 0x2
Route refresh request: received 0, sent 0
Community attribute sent to this neighbor
Outbound path policy configured
Incoming update prefix filter list is bgp-in
Outgoing update prefix filter list is aggregate
Route map for outgoing advertisements is uni-out
77 accepted prefixes consume 4928 bytes
Prefix advertised 4303, suppressed 0, withdrawn 1328
Number of NLRIs in the update sent: max 1, min 0
1 history paths consume 64 bytes
Connections established 22; dropped 21
Last reset 13:47:05, due to BGP Notification sent, hold time expired
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 3FFE:700:20:1::12, Local port: 55345
Foreign host: 3FFE:700:20:1::11, Foreign port: 179
Enqueued packets for retransmit: 0, input: 0 mis-ordered: 0 (0 bytes)
Event Timers (current time is 0x1A0D543C):
Timer           Starts    Wakeups      Next
Retrans         1218         5            0x0
TimeWait        0             0            0x0
AckHold         3327        3051         0x0
SendWnd         0             0            0x0
KeepAlive       0             0            0x0
GiveUp          0             0            0x0
PmtuAger        0             0            0x0
DeadWait        0             0            0x0
iss: 1805423033  snduna: 1805489354  sndnxt: 1805489354  sndwnd: 15531
irs: 821333727  rcvnxt: 821591465  rcvwnd: 15547  delrcvwnd: 837
SRTT: 300 ms, RTTO: 303 ms, RTV: 3 ms, KRTT: 0 ms
minRTT: 8 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: higher precedence, nagle
Datagrams (max data segment is 1420 bytes):
Rcvd: 4252 (out of order: 0), with data: 3328, total data bytes: 257737
Sent: 4445 (retransmit: 5), with data: 4445, total data bytes: 244128

```

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

Table 3: show bgp ipv4/ ipv6 unicast neighbors fields

Field	Description
BGP neighbor	IP address of the BGP neighbor and its autonomous system number. If the neighbor is in the same autonomous system as the router, then the link between them is internal; otherwise, it is considered external.
remote AS	Autonomous system of the neighbor.
internal link	Indicates that this peer is an interior Border Gateway Protocol (iBGP) peer.
BGP version	BGP version being used to communicate with the remote router; the router ID (an IP address) of the neighbor is also specified.
remote router ID	A 32-bit number written as 4 octets separated by periods (dotted-decimal format).

Field	Description
BGP state	Internal state of this BGP connection.
up for	Amount of time that the underlying TCP connection has been in existence.
Last read	Time that BGP last read a message from this neighbor.
hold time	Maximum amount of time that can elapse between messages from the peer.
keepalive interval	Time period between sending keepalive packets, which help ensure that the TCP connection is up.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor.
Route refresh	Indicates that the neighbor supports dynamic soft reset using the route refresh capability.
Address family IPv6 Unicast	Indicates that BGP peers are exchanging IPv6 reachability information.
Received	Number of total BGP messages received from this peer, including keepalives.
notifications	Number of error messages received from the peer.
Sent	Total number of BGP messages that have been sent to this peer, including keepalives.
notifications	Number of error messages the router has sent to this peer.
advertisement runs	Value of the minimum advertisement interval.
For address family	Address family to which the following fields refer.
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
neighbor version	Number used by the software to track the prefixes that have been sent and those that must be sent to this neighbor.
Route refresh request	Number of route refresh requests sent and received from this neighbor.
Community attribute (not shown in sample output)	Appears if the neighbor send-community command is configured for this neighbor.
Inbound path policy (not shown in sample output)	Indicates whether an inbound filter list or route map is configured.
Outbound path policy (not shown in sample output)	Indicates whether an outbound filter list, route map, or unsuppress map is configured.
bgp-in (not shown in sample output)	Name of the inbound update prefix filter list for the IPv6 unicast address family.
aggregate (not shown in sample output)	Name of the outbound update prefix filter list for the IPv6 unicast address family.

Field	Description
uni-out (not shown in sample output)	Name of the outbound route map for the IPv6 unicast address family.
accepted prefixes	Number of prefixes accepted.
Prefix advertised	Number of prefixes advertised.
suppressed	Number of prefixes suppressed
withdrawn	Number of prefixes withdrawn.
history paths (not shown in sample output)	Number of path entries held to remember history.
Connections established	Number of times the router has established a TCP connection and the two peers have agreed to speak BGP with each other.
dropped	Number of times that a good connection has failed or been taken down.
Last reset	Elapsed time (in hours:minutes:seconds) since this peering session was last reset.
Connection state	State of the BGP Peer
unread input bytes	Number of bytes of packets still to be processed.
Local host, Local port	Peering address of the local router, plus the port.
Foreign host, Foreign port	Peering address of the neighbor.
Event Timers	Table that displays the number of starts and wakeups for each timer.
snduna	Last send sequence number for which the local host sent but has not received an acknowledgment.
sndnxt	Sequence number the local host will send next.
sndwnd	TCP window size of the remote host.
irs	Initial receive sequence number.
rcvnxt	Last receive sequence number the local host has acknowledged.
rcvwnd	TCP window size of the local host.
delrcvwnd	Delayed receive window--data the local host has read from the connection, but has not yet subtracted from the receive window the host has advertised to the remote host. The value in this field gradually increases until it is larger than a full-sized packet, at which point it is applied to the rcvwnd field.
SRTT	A calculated smoothed round-trip timeout (in milliseconds).
RTTO	Round-trip timeout (in milliseconds).

Field	Description
RTV	Variance of the round-trip time (in milliseconds).
KRTT	New round-trip timeout (in milliseconds) using the Karn algorithm. This field separately tracks the round-trip time of packets that have been re-sent.
minRTT	Smallest recorded round-trip timeout (in milliseconds) with hard wire value used for calculation.
maxRTT	Largest recorded round-trip timeout (in milliseconds).
ACK hold	Time (in milliseconds) the local host will delay an acknowledgment in order to “piggyback” data on it.
Flags	IP precedence of the BGP packets.
Datagrams: Rcvd	Number of update packets received from neighbor.
with data	Number of update packets received with data.
total data bytes	Total number of bytes of data.
Sent	Number of update packets sent.
with data	Number of update packets with data sent.
total data bytes	Total number of data bytes.

The following is sample output from the **show bgp ipv6 unicast neighbors** command with the **advertised-routes** keyword. For an explanation of the output, see the **show bgp** command.

```
> show bgp ipv6 unicast neighbors 3FFE:700:20:1::11 advertised-routes
BGP table version is 21880, local router ID is 192.168.7.225
Status codes: s suppressed, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 2001:200::/35    3FFE:700:20:1::11      0 293 3425 2500 i
*> 2001:208::/35    3FFE:C00:E:B::2        0 237 7610 i
*> 2001:218::/35    3FFE:C00:E:C::2        0 3748 4697 i
```

The following is sample output from the **show bgp ipv6 unicast neighbors** command with the **routes** keyword:

```
> show bgp ipv6 unicast neighbors 3FFE:700:20:1::11 routes
BGP table version is 21885, local router ID is 192.168.7.225
Status codes: s suppressed, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 2001:200::/35    3FFE:700:20:1::11      0 293 3425 2500 i
* 2001:208::/35    3FFE:700:20:1::11      0 293 7610 i
* 2001:218::/35    3FFE:700:20:1::11      0 293 3425 4697 i
* 2001:230::/35    3FFE:700:20:1::11      0 293 1275 3748 i
```

The following is sample output from the **show bgp ipv6 neighbors** command with the **paths** keyword:

```
> show bgp ipv6 unicast neighbors 3FFE:700:20:1::11 paths ^293
Address      Refcount  Metric  Path
0x6131D7DC   2         0 293 3425 2500 i
0x6132861C   2         0 293 7610 i
0x6131AD18   2         0 293 3425 4697 i
0x61324084   2         0 293 1275 3748 i
0x61320E0C   1         0 293 3425 2500 2497 i
0x61326928   1         0 293 3425 2513 i
0x61327BC0   2         0 293 i
0x61321758   1         0 293 145 i
0x61320BEC   1         0 293 3425 6509 i
0x6131AAF8   2         0 293 1849 2914 ?
0x61320FE8   1         0 293 1849 1273 209 i
0x613260A8   2         0 293 1849 i
0x6132586C   1         0 293 1849 5539 i
0x6131BBF8   2         0 293 1849 1103 i
0x6132344C   1         0 293 4554 1103 1849 1752 i
0x61324150   2         0 293 1275 559 i
0x6131E5AC   2         0 293 1849 786 i
0x613235E4   1         0 293 1849 1273 i
0x6131D028   1         0 293 4554 5539 8627 i
0x613279E4   1         0 293 1275 3748 4697 3257 i
0x61320328   1         0 293 1849 1273 790 i
0x6131EC0C   2         0 293 1275 5409 i
```

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

Table 4: show bgp ipv6 neighbors paths fields

Field	Description
Address	Internal address where the 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.

The following sample output from the **show bgp ipv6 neighbors** command shows the **received routes** for IPv6 address 2000:0:0:4::2:

```
> show bgp ipv6 unicast neighbors 2000:0:0:4::2 received-routes
BGP table version is 2443, local router ID is 192.168.0.2
Status codes:s suppressed, h history, * valid, > best, i - internal
Origin codes:i - IGP, e - EGP, ? - incomplete

   Network                Next Hop                Metric LocPrf Weight Path
*> 2000:0:0:1::/64        2000:0:0:4::2                0 2 1 i
*> 2000:0:0:2::/64        2000:0:0:4::2                0 2 i
*> 2000:0:0:2:1::/80      2000:0:0:4::2                0 2 ?
*> 2000:0:0:3::/64        2000:0:0:4::2                0 2 ?
* 2000:0:0:4::1/64        2000:0:0:4::2                0 2 ?
```

show bgp ipv4/ ipv6 unicast paths

To display all the IPv4 or IPv6 Border Gateway Protocol (BGP) paths in the database, use the **show bgp ipv4 unicast paths** or **show bgp ipv6 unicast paths** command respectively.

```
show bgp [vrf name | all] {ipv4 | ipv6} unicast paths [regular-expression]
```

Syntax Description	<i>regular-expression</i>	(Optional) Regular expression used to match the paths received.
	[vrf name all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [vrf name all] keywords were added.

Examples

The following is sample output from the **show bgp ipv6 unicast paths** command:

```
> show bgp ipv6 unicast paths
Address      Hash Refcount Metric Path
0x61322A78   0      2      0 i
0x6131C214   3      2      0 6346 8664 786 i
0x6131D600   13     1      0 3748 1275 8319 1273 209 i
0x613229F0   17     1      0 3748 1275 8319 12853 i
0x61324AE0   18     1      1 4554 3748 4697 5408 i
0x61326818   32     1      1 4554 5609 i
0x61324728   34     1      0 6346 8664 9009 ?
0x61323804   35     1      0 3748 1275 8319 i
0x61327918   35     1      0 237 2839 8664 ?
0x61320504   38     2      0 3748 4697 1752 i
0x61320988   41     2      0 1849 786 i
0x6132245C   46     1      0 6346 8664 4927 i
```

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

Table 5: Show bgp ipv4/ ipv6 unicast path fields

Field	Description
Address	Internal address where the path is stored.
Refcount	Number of routes using that path.

Field	Description
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 bgp ipv4/ ipv6 unicast prefix-list

To display routes that match a prefix list, use the **show bgp ipv4 prefix-list** or **show bgp ipv6 prefix-list** command.

```
show bgp [vrf name | all] {ipv4 | ipv6} unicast prefix-list name
```

Syntax Description	prefix-list name	The specified prefix-list.
	[vrf name all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.
Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [vrf name all] keywords were added.

Examples

The following is sample output from the **show bgp ipv6 prefix-list** command:

```
> show bgp ipv6 unicast prefix-list pin
ipv6 prefix-list pin:
  count:4, range entries:3, sequences:5 - 20, refcount:2
  seq 5 permit 747::/16 (hit count:1, refcount:2)
  seq 10 permit 747:1::/32 ge 64 le 64 (hit count:2, refcount:2)
  seq 15 permit 747::/32 ge 33 (hit count:1, refcount:1)
  seq 20 permit 777::/16 le 124 (hit count:2, refcount:1)
The ipv6 prefix-list match the following prefixes:
seq 5: matches the exact match 747::/16
seq 10:first 32 bits in prefix must match with a prefixlen of /64
seq 15:first 32 bits in prefix must match with any prefixlen up to /128
seq 20:first 16 bits in prefix must match with any prefixlen up to /124
```

show bgp ipv4/ ipv6 unicast regexp

To display IPv4 or IPv6 Border Gateway Protocol (BGP) routes matching the autonomous system path regular expression, use the **show bgp ipv4 regexp** or **show bgp ipv6 regexp** command.

```
show bgp [vrf name | all] {ipv4 | ipv6} unicast regexp regular-expression
```

Syntax Description	regexp <i>regular-expression</i>	Regular expression that is used to match the BGP autonomous system paths
[vrf name all]		If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.
Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [vrf name all] keywords were added.

Examples

The following is sample output from the **show bgp ipv6 unicast regexp** command that shows paths beginning with 33 or containing 293. For an explanation of the output, see the **show bgp** command.

```
> show bgp ipv6 unicast regexp ^33|293
BGP table version is 69964, local router ID is 192.168.7.225
Status codes: s suppressed, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf Weight Path
*  2001:200::/35     3FFE:C00:E:4::2      1         0 4554 293 3425 2500 i
*                   2001:0DB8:0:F004::1  0 3320 293 3425 2500 i
*  2001:208::/35     3FFE:C00:E:4::2      1         0 4554 293 7610 i
*  2001:228::/35     3FFE:C00:E:F::2      0 6389 1849 293 2713 i
*  3FFE::/24         3FFE:C00:E:5::2      0 33 1849 4554 i
*  3FFE:100::/24     3FFE:C00:E:5::2      0 33 1849 3263 i
*  3FFE:300::/24     3FFE:C00:E:5::2      0 33 293 1275 1717 i
*                   3FFE:C00:E:F::2      0 6389 1849 293 1275
```

show bgp ipv4/ ipv6 unicast route-map

To display IPv4 or IPv6 Border Gateway Protocol (BGP) routes that failed to install in the routing table, use the **show bgp ipv4 unicast route-map** or **show bgp ipv6 unicast route-map** command.

```
show bgp [vrf name | all] {ipv4 | ipv6} unicast route-map name
```

Syntax Description	route-map name	A specified route map to match.
	[vrf name all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.
Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [vrf name all] keywords were added.

Examples

The following is sample output from the **show bgp ipv6 unicast route-map rmap** command for a route map named rmap. For an explanation of the output, see the **show bgp** command.

```
> show bgp ipv6 unicast route-map rmap
BGP table version is 16, local router ID is 172.30.242.1
Status codes:s suppressed, 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
*>i12:12::/64      2001:0DB8:101::1      0      100    50 ?
*>i12:13::/64      2001:0DB8:101::1      0      100    50 ?
*>i12:14::/64      2001:0DB8:101::1      0      100    50 ?
*>i543::/64        2001:0DB8:101::1      0      100    50 ?
```

show bgp ipv4/ ipv6 unicast summary

To display the status of all IPv4 or IPv6 Border Gateway Protocol (BGP) connections, use the **show bgp ipv4 unicast summary** or **show bgp ipv6 unicast summary** command respectively.

show bgp [*vrf name* | **all**] {**ipv4** | **ipv6**} **unicast summary**

Syntax Description	[vrf name all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.
--------------------	------------------	---

Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [<i>vrf name</i> all] keywords were added.

Examples

The following is sample output from the **show bgp ipv6 unicast summary** command:

```
> show bgp ipv6 unicast summary
BGP device identifier 172.30.4.4, local AS number 200
BGP table version is 1, main routing table version 1
Neighbor      V   AS  MsgRcvd  MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
2001:0DB8:101::2  4   200    6869    6882     0    0    0  06:25:24  Active
```

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

Table 6: show bgp ipv4/ ipv6 unicast summary fields

Field	Description
BGP device identifier	IP address of the networking device.
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
main routing table version	Last version of BGP database that was injected into the main routing table.
Neighbor	IPv6 address of a neighbor.
V	BGP version number spoken to that neighbor.
AS	Autonomous System
MsgRcvd	BGP messages received from that neighbor.

Field	Description
MsgSent	BGP messages sent to that neighbor
TblVer	Last version of the BGP database that was sent to that neighbor.
InQ	Number of messages from that neighbor waiting to be processed.
OutQ	Number of messages waiting to be sent to that neighbor.
Up/Down	The length of time that the BGP session has been in state Established, or the current state if it is not Established.
State/PfxRcd	<p>Current state of the BGP session/the number of prefixes the device has received from a neighbor. 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 Idle.</p> <p>An (Admin) entry with Idle status indicates that the connection has been shut down using the neighbor shutdown command.</p>

show bgp neighbors

To display information about Border Gateway Protocol (BGP) and TCP connections to neighbors, use the `show bgp neighbors` command.

```
show bgp neighbors [vrf name | all] [slow | ip-address [advertised-routes | paths [reg-exp]
| policy [detail] | received prefix-filter | received-routes | routes]]
```

Syntax	Description
slow	(Optional) Displays information about dynamically configured slow peers
<i>ip-address</i>	(Optional) Displays information about the IPv4 neighbor. If this argument is omitted, information about all neighbors is displayed.
advertised-routes	(Optional) Displays all routes that have been advertised to neighbors.
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.
[vrf <i>name</i> all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf <i>name</i> keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command Default The output of this command displays information for all neighbors.

Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [vrf <i>name</i> all] keywords were added.

Usage Guidelines

Use the **show 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.

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.

```
> show 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's 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 following table 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 7: show bgp neighbors Fields

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 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 hhhmss, that the underlying TCP connection has been in existence.
Last read	Time, in hhhmss, since BGP last received a message from this neighbor.
last write	Time, in hhhmss, 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 a messages.
keepalive interval	Time interval, in seconds, at which keepalive messages are transmitted to this neighbor.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor. "advertised and received" is displayed when a capability is successfully exchanged between two routers
Route Refresh	Status of the route refresh capability.
Graceful Restart Capability	Status of the graceful restart capability.
Address family IPv4 Unicast	IP Version 4 unicast-specific properties of this neighbor.
Message statistics	Statistics organized by message type.
InQ depth is	Number of messages in the input queue.
OutQ depth is	Number of messages in the output queue.
Sent	Total number of transmitted messages.
Received	Total number of received messages.
Opens	Number of open messages sent and received.
notifications	Number of notification (error) messages sent and received.
Updates	Number of update messages sent and received.
Keepalives	Number of keepalive messages sent and received.
Route Refresh	Number of route refresh request messages sent and received.
Total	Total number of messages sent and received.
Default minimum time between...	Time, in seconds, between advertisement transmissions.
For address family:	Address family to which the following fields refer.
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
neighbor version	Number used by the software to track prefixes that have been sent and those that need to be sent.
update-group	Number of update-group member for this address family
Prefix activity	Prefix statistics for this address family.

Field	Description
Prefixes current	Number of prefixes accepted for this address family.
Prefixes total	Total number of received prefixes.
Implicit Withdraw	Number of times that a prefix has been withdrawn and readvertised.
Explicit Withdraw	Number of times that prefix has been withdrawn because it is no longer feasible.
Used as bestpath	Number of received prefixes installed as bestpaths.
Used as multipath	Number of received prefixes installed as multipaths.
* Saved (soft-reconfig)	Number of soft resets performed with a neighbor that supports soft reconfiguration. This field is displayed only if the counter has a nonzero value.
* History paths	This field is displayed only if the counter has a nonzero value.
* Invalid paths	Number of invalid paths. This field is displayed only if the counter has a nonzero value.
Local Policy Denied Prefixes	Prefixes denied due to local policy configuration. Counters are updated for inbound and outbound policy denials. The fields under this heading are displayed only if the counter has a nonzero value.
* route-map	Displays inbound and outbound route-map policy denials.
* filter-list	Displays inbound and outbound filter-list policy denials.
* prefix-list	Displays inbound and outbound prefix-list policy denials.
* AS_PATH too long	Displays outbound AS-path length policy denials.
* AS_PATH loop	Displays outbound AS-path loop policy denials.
* AS_PATH confed info	Displays outbound confederation policy denials.
* AS_PATH contains AS 0	Displays outbound denials of autonomous system (AS) 0.
* NEXT_HOP Martian	Displays outbound martian denials.
* NEXT_HOP non-local	Displays outbound non-local next-hop denials.
* NEXT_HOP is us	Displays outbound next-hop-self denials.
* CLUSTER_LIST loop	Displays outbound cluster-list loop denials.
* ORIGINATOR loop	Displays outbound denials of local originated routes.
* unsuppress-map	Displays inbound denials due to an unsuppress-map.

Field	Description
* advertise-map	Displays inbound denials due to an advertise-map.
* Well-known Community	Displays inbound denials of well-known communities.
* SOO loop	Displays inbound denials due to site-of-origin.
* Bestpath from this peer	Displays inbound denials because the bestpath came from the local router.
* Suppressed due to dampening	Displays inbound denials because the neighbor or link is in a dampening state.
* Bestpath from iBGP peer	Displays inbound denials because the bestpath came from an iBGP neighbor.
* Incorrect RIB for CE	Displays inbound denials due to RIB errors for a CE router.
* BGP distribute-list	Displays inbound denials due to a distribute list.
Number of NLRIs...	Number of network layer reachability attributes in updates.
Connections established	Number of times a TCP and BGP connection has been successfully established.
dropped	Number of times that a valid session has failed or been taken down.
Last reset	Time since this peering session was last reset. The reason for the reset is displayed on this line.
External BGP neighbor may be... (not shown in the display)	Indicates that the BGP TTL security check is enabled. The maximum number of hops that can separate the local and remote peer is displayed on this line.
Connection state	Connection status of the BGP peer.
Connection is ECN Disabled	Explicit congestion notification status (enabled or disabled).
Local host: 10.108.50.1, Local port: 179	IP address of the local BGP speaker. BGP port number 179.
Foreign host: 10.108.50.2, Foreign port: 42698	Neighbor address and BGP destination port number.
Enqueued packets for retransmit:	Packets queued for retransmission by TCP.
Event Timers	TCP event timers. Counters are provided for starts and wakeups (expired timers).
Retrans	Number of times a packet has been retransmitted.

Field	Description
TimeWait	Time waiting for the retransmission timers to expire.
AckHold	Acknowledgment hold timer.
SendWnd	Transmission (send) window.
KeepAlive	Number of keepalive packets.
GiveUp	Number times a packet is dropped due to no acknowledgment.
PmtuAger	Path MTU discovery timer
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.
revnxt:	Last receive sequence number that has been locally acknowledged.
revwnd:	TCP window size of the local host.
delrcvwnd:	Delayed receive window—data the local host has read from the connection, but has not yet subtracted from the receive window the host has advertised to the remote host. The value in this field gradually increases until it is larger than a full-sized packet, at which point it is applied to the revwnd field.
SRTT:	A calculated smoothed round-trip timeout.
RTTO:	Round-trip timeout.
RTV:	Variance of the round-trip time.
KRTT:	New round-trip timeout (using the Karn algorithm). This field separately tracks the round-trip time of packets that have been re-sent.
minRTT:	Smallest recorded round-trip timeout (hard-wire value used for calculation).
maxRTT:	Largest recorded round-trip timeout.
ACK hold:	Length of time the local host will delay an acknowledgment to carry (piggyback) additional data.
IP Precedence value:	IP precedence of the BGP packets.
Datagrams	Number of update packets received from a neighbor.
Rcvd:	Number of received packets.

Field	Description
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 second retransmissions sent due to congestion.
Datagrams: Rcvd	Number of update packets received from a neighbor.
out of order:	Number of packets received out of sequence.
with data	Number of update packets received with data.
Last reset	Elapsed time since this peering session was last reset.
unread input bytes	Number of bytes of packets still to be processed.
retransmit	Number of packets retransmitted.
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).

The following example displays routes advertised for only the 172.16.232.178 neighbor. For an explanation of the output, see the **show bgp** command.

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

```
> show bgp neighbors 172.29.232.178 paths ^10
Address      Refcount Metric Path
0x60E577B0   2        40 10 ?
```

The following table explains each field.

Table 8: show bgp neighbors paths Fields

Field	Description
Address	Internal address where the path is stored.
Refcount	Number of routes using that path..

Field	Description
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..

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:

```
> show 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 following sample output shows the policies applied to the neighbor at 192.168.1.2. The output displays policies configured on the neighbor device.

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

The following is sample output from the **show 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:

```
> show 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, KRRTT: 0 ms
minRTT: 8 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: higher precedence, retransmission timeout, nagle, path mtu capable
```

show bgp paths

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

show bgp paths [*vrf name* | **all**] [*regexp*]

Syntax Description	<i>regexp</i>	Regular expression to match the BGP autonomous system paths.
	[<i>vrf name</i> all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [<i>vrf name</i> all] keywords were added.

Examples

The following is sample output from the **show bgp paths** command.

```
> show 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 following table explains each field.

Table 9: show bgp paths Fields

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 bgp prefix-list

To display information about a prefix list or prefix list entries, use the **show bgp prefix-list** command.

```
show bgp prefix-list [vrf name | all] [detail | summary] [prefix-list-name [seq
sequence-number | network/length [longer | first-match]]]
```

Syntax Description	detail summary	(Optional) Displays detailed or summarized information about all prefix lists.
	first-match	(Optional) Displays the first entry of the specified prefix list that matches the given network/length.
	longer	(Optional) Displays all entries of the specified prefix list that match or are more specific than the given network/length.
	<i>network/length</i>	(Optional) Displays all entries in the specified prefix list that use this network address and netmask length (in bits).
	<i>prefix-list-name</i>	(Optional) Displays the entries in a specific prefix list.
	seq <i>sequence-number</i>	(Optional) Displays only the prefix list entry with the specified sequence number in the specified prefix-list.
	[<i>vrf name</i> all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [vrf name all] keywords were added.

Examples

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

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

show bgp regexp

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

```
show bgp regexp [vrf name | all] regexp
```

Syntax Description	regex	Regular expression to match the BGP autonomous system paths.
	[vrf name all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [vrf name all] keywords were added.

Examples

The following is sample output from the **show bgp regexp** command.

```
> show bgp regexp 108$
BGP table version is 1738, local router ID is 172.16.72.24
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop           Metric LocPrf Weight Path
* 172.16.0.0      172.16.72.30          0 109 108 ?
* 172.16.1.0      172.16.72.30          0 109 108 ?
* 172.16.11.0     172.16.72.30          0 109 108 ?
* 172.16.14.0     172.16.72.30          0 109 108 ?
* 172.16.15.0     172.16.72.30          0 109 108 ?
* 172.16.16.0     172.16.72.30          0 109 108 ?
* 172.16.17.0     172.16.72.30          0 109 108 ?
* 172.16.18.0     172.16.72.30          0 109 108 ?
* 172.16.19.0     172.16.72.30          0 109 108 ?
* 172.16.24.0     172.16.72.30          0 109 108 ?
* 172.16.29.0     172.16.72.30          0 109 108 ?
* 172.16.30.0     172.16.72.30          0 109 108 ?
* 172.16.33.0     172.16.72.30          0 109 108 ?
* 172.16.35.0     172.16.72.30          0 109 108 ?
* 172.16.36.0     172.16.72.30          0 109 108 ?
* 172.16.37.0     172.16.72.30          0 109 108 ?
* 172.16.38.0     172.16.72.30          0 109 108 ?
* 172.16.39.0     172.16.72.30          0 109 108 ?
```

show bgp rib-failure

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

```
show bgp rib-failure [vrf name | all]
```

Syntax Description	[vrf name all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.
---------------------------	------------------	---

Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [vrf name all] keywords were added.

Examples

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

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

The following table explains each field.

Table 10: show bgp rib-failure Fields

Field	Description
Network	IP address of a network entity
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.
RIB-failure	Cause of RIB failure. Higher admin distance means that a route with a better (lower) administrative distance such as a static route already exists in the IP routing table.

Field	Description
RIB-NH Matches	<p>Route status that applies only when Higher admin distance appears in the RIB-failure column and bgp suppress-inactive is configured for the address family being used. There are three choices:</p> <ul style="list-style-type: none">• Yes—Means that the route in the RIB has the same next hop as the BGP route or next hop recurses down to the same adjacency as the BGP nexthop.• No—Means that the next hop in the RIB recurses down differently from the next hop of the BGP route.• n/a—Means that bgp suppress-inactive is not configured for the address family being used.

show bgp summary

To display the status of all Border Gateway Protocol (BGP) connections, use the **show bgp summary** command.

```
show bgp summary [vrf name | all]
```

Syntax Description	[<i>vrf name</i> all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf <i>name</i> keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.
---------------------------	----------------------------------	--

Command History	Release	Modification
	6.1	This command was introduced.
	6.6	The [<i>vrf name</i> all] keywords were added.

Usage Guidelines The **show 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.

Examples

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

```
> show 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 following table explains each field.

Table 11: show bgp summary Fields

Field	Description
BGP router identifier	In order of precedence and availability, the router identifier, 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	Number of advertised bitfield entries and the associated memory usage. A bitfield entry represents a piece of information (one bit) that is generated when a prefix is advertised to a peer. The advertised bit cache is built dynamically when required
...received paths for inbound soft reconfiguration	Number paths received and stored for inbound soft reconfiguration.
BGP using...	Total amount of memory, in bytes, used by the BGP process.
Dampening enabled...	Indicates that BGP dampening is enabled. The number of paths that carry an accumulated penalty and the number of dampened paths are displayed on this line.
BGP activity...	Displays the number of times that memory has been allocated or released for a path or prefix.
Neighbor	IP address of the neighbor.
V	BGP version number spoken to the neighbor.
AS	Autonomous system number.
MsgRcvd	Number of messages received from the neighbor.
MsgSent	Number of messages sent to the neighbor.
TblVer	Last version of the BGP database that was sent to the neighbor.
InQ	Number of messages queued to be processed from the neighbor.
OutQ	Number of messages queued to be sent to the neighbor.
Up/Down	The length of time that the BGP session has been in the Established state, or the current status if not in the Established state.
State/PfxRcd	Current state of the BGP session, and the number of prefixes that have been received from a neighbor or peer group. When the maximum number 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.

The following output from the **show 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.

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

Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
*192.168.3.2  4 50000     2      2        0    0    0 00:00:37      0
* Dynamically created based on a listen range command
```

Dynamically created neighbors: 1/(200 max), Subnet ranges: 1

BGP peergroup group192 listen range group members:
192.168.0.0/16

The following output from the **show 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.

```
> show 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 bgp summary** command shows the same two BGP neighbors, but the 4-byte autonomous system numbers are displayed in asdot notation format.

```
> show 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
```

show bgp update-group

To display information about BGP update-groups, use the **show bgp update-group** command.

```
show bgp update-group [vrf name | all] [index-group | ip-address] [summary]
```

Syntax Description

<i>index-group</i>	(Optional) Update group type with its corresponding index number. The range of update-group index numbers is from 1 to 4294967295.
<i>ip-address</i>	(Optional) IP address of a single neighbor who is a member of an update group.
summary	(Optional) Displays a summary of update-group member information. The output can be filtered to show information for a single index group or peer with the <i>index-group</i> or <i>ip-address</i> argument.
[vrf name all]	If you enable virtual routing and forwarding (VRF), also known as virtual routers, you can limit the command to a specific virtual router using the vrf name keyword. If you want the command to affect all virtual routers, include the all keyword. If you include neither of these VRF-related keywords, the command applies to the global VRF virtual router.

Command History

Release	Modification
6.1	This command was introduced.
6.6	The [vrf name all] keywords were added.

Usage Guidelines

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

Examples

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

```
> show bgp update-group
BGP version 4 update-group 1, internal, Address Family: IPv4 Unicast
  BGP Update version : 0, messages 0/0
  Route map for outgoing advertisements is COST1
  Update messages formatted 0, replicated 0
  Number of NLRIs in the update sent: max 0, min 0
  Minimum time between advertisement runs is 5 seconds
  Has 1 member:
  10.4.9.21
BGP version 4 update-group 2, internal, Address Family: IPv4 Unicast
  BGP Update version : 0, messages 0/0
  Update messages formatted 0, replicated 0
  Number of NLRIs in the update sent: max 0, min 0
```

```

Minimum time between advertisement runs is 5 seconds
Has 2 members:
10.4.9.5 10.4.9.8

```

The following table explains each field.

Table 12: show bgp update-group Fields

Field	Description
BGP version	BGP version.
update-group	Update-group number and type (internal or external).
update messages formatted..., replicated...	Number of update messages that have been formatted and replicated.
Number of NLRIs...	NLRI information sent in update.
.Minimum time between...	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	Number of advertised bitfield entries and the associated memory usage. A bitfield entry represents a piece of information (one bit) that is generated when a prefix is advertised to a peer. The advertised bit cache is built dynamically when required
...received paths for inbound soft reconfiguration	Number paths received and stored for inbound soft reconfiguration.
BGP using...	Total amount of memory, in bytes, used by the BGP process.
Dampening enabled...	Indicates that BGP dampening is enabled. The number of paths that carry an accumulated penalty and the number of dampened paths are displayed on this line.
BGP activity...	Displays the number of times that memory has been allocated or released for a path or prefix.
Neighbor	IP address of the neighbor.
V	BGP version number spoken to the neighbor.
AS	Autonomous system number.
MsgRcvd	Number of messages received from the neighbor.
MsgSent	Number of messages sent to the neighbor.
TblVer	Last version of the BGP database that was sent to the neighbor.
InQ	Number of messages queued to be processed from the neighbor.
OutQ	Number of messages queued to be sent to the neighbor.
Up/Down	The length of time that the BGP session has been in the Established state, or the current status if not in the Established state.
State/PfxRcd	Current state of the BGP session, and the number of prefixes that have been received from a neighbor or peer group. When the maximum number 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.

show blocks

To show the system buffer utilization, use the **show blocks** command.

```
show blocks [core | export-failed | interface]
show blocks address hex [diagnostics | dump | header | packet]
show blocks {all | assigned | free | old} [core-local [core-num] [diagnostics | dump |
header | packet]]
show blocks exhaustion {history [list | snapshot_num] | snapshot}
show blocks pool block-size
show blocks queue history [core-local [core-num]] [detail]
```

Syntax Description	
address <i>hex</i>	(Optional) Shows a block corresponding to this address, in hexadecimal.
all	(Optional) Shows all blocks.
assigned	(Optional) Shows blocks that are assigned and in use by an application.
core	(Optional) Shows core-specific buffers.
core-local [<i>core-num</i>]	(Optional) Shows system buffers for all cores. You can also specify a core number, for example, 1, to see the buffers for a specific core.
detail	(Optional) Shows a portion (128 bytes) of the first block for each unique queue type.
dump	(Optional) Shows the entire block contents, including the header and packet information. The difference between dump and packet is that dump includes additional information between the header and the packet.
diagnostics	(Optional) Shows block diagnostics.
exhaustion snapshot	(Optional) Prints the last x number (x is currently 10) of snapshots that were taken and the time stamp of the last snapshot). After a snapshot is taken, another snapshot is not taken if less than 5 minutes has passed.
exhaustion history [<i>list</i> <i>snapshot_num</i>]	(Optional) Shows the exhaustion snapshot history. You can specify a snapshot number to limit information to a single snapshot, or list to see a list of snapshots.
export-failed	(Optional) Show system buffer export failure counters.
free	(Optional) Shows blocks that are available for use.
header	(Optional) Shows the header of the block.
interface	(Optional) Show buffers attached to interfaces.
old	(Optional) Shows blocks that were assigned more than a minute ago.
packet	(Optional) Shows the header of the block as well as the packet contents.

pool <i>block-size</i>	(Optional) Shows blocks of a specific size.
queue history	(Optional) Shows where blocks are assigned when the threat defense device runs out of blocks. Sometimes, a block is allocated from the pool but never assigned to a queue. In that case, the location is the code address that allocated the block.

Command History

Release	Modification
6.1	This command was introduced.
7.0(1)	The output of this command was enhanced to include the failed count.

Usage Guidelines

The **show blocks** command helps you determine if the threat defense device is overloaded. This command lists preallocated system buffer utilization. A full memory condition is not a problem as long as traffic is moving through the threat defense device. You can use the **show conn** command to see if traffic is moving. If traffic is not moving and the memory is full, there may be a problem. You can also view this information using SNMP.

Examples

The following is sample output from the **show blocks** command.

```
> show blocks
  SIZE    MAX    LOW    CNT    FAILED
   0    1450   1450   1450     0
   4     100     99     99     0
   80    1996   1992   1992     0
  256    4148   4135   4142     0
 1550    6274   6270   6272     0
 2048     100     100     100     0
 2560     164     164     164     0
 4096     100     100     100     0
 8192     100     100     100     0
 9344     100     100     100     0
16384     100     100     100     0
65536     16      16      16      0
```

The following table explains each field.

Table 13: show blocks Fields

Field	Description
SIZE	Size, in bytes, of the block pool. Each size represents a particular type.
0	Used by dupb blocks.
4	Duplicates existing blocks in applications such as DNS, ISAKMP, URL filtering, uauth, TFTP, and TCP modules. Also, this sized block can be used normally by code to send packets to drivers, etc.
80	Used in TCP intercept to generate acknowledgment packets and for failover hello messages.

Field	Description
256	<p>Used for Stateful Failover updates, syslogging, and other TCP functions.</p> <p>These blocks are mainly used for Stateful Failover messages. The active threat defense device generates and sends packets to the standby threat defense device to update the translation and connection table. In bursty traffic, where high rates of connections are created or torn down, the number of available blocks might drop to 0. This situation indicates that one or more connections were not updated to the standby threat defense device. The Stateful Failover protocol catches the missing translation or connection the next time. If the CNT column for 256-byte blocks stays at or near 0 for extended periods of time, then the threat defense device is having trouble keeping the translation and connection tables synchronized because of the number of connections per second that the threat defense device is processing.</p> <p>Syslog messages sent out from the threat defense device also use the 256-byte blocks, but they are generally not released in such quantity to cause a depletion of the 256-byte block pool. If the CNT column shows that the number of 256-byte blocks is near 0, ensure that you are not logging at Debugging (level 7) to the syslog server. This is indicated by the logging trap line in the threat defense configuration. We recommend that you set logging at Notification (level 5) or lower, unless you require additional information for debugging purposes.</p>
1550	<p>Used to store Ethernet packets for processing through the threat defense device.</p> <p>When a packet enters an interface, it is placed on the input interface queue, passed up to the operating system, and placed in a block. The device determines whether the packet should be permitted or denied based on the security policy and processes the packet through to the output queue on the outbound interface. If the device is having trouble keeping up with the traffic load, the number of available blocks will hover close to 0 (as shown in the CNT column of the command output). When the CNT column is zero, the device attempts to allocate more blocks. The maximum can be greater than 8192 for 1550-byte blocks if you issue this command. If no more blocks are available, the device drops the packet.</p>
2048	Control or guided frames used for control updates.
16384	<p>Only used for the 64-bit, 66-MHz Gigabit Ethernet cards (i82543).</p> <p>See the description for 1550 for more information about Ethernet packets.</p>
MAX	Maximum number of blocks available for the specified byte block pool. The maximum number of blocks are carved out of memory at bootup. Typically, the maximum number of blocks does not change. The exception is for the 256- and 1550-byte blocks, where the device can dynamically create more when needed. The maximum can be greater than 8192 for 1550-byte blocks if you issue this command.
LOW	Low-water mark. This number indicates the lowest number of this size blocks available since the device was powered up, or since the last clearing of the blocks (with the clear blocks command). A zero in the LOW column indicates a previous event where memory was full.
CNT	Current number of blocks available for that specific size block pool. A zero in the CNT column means memory is full now.

Field	Description
FAILED	When the memory count for a block size is completely exhausted (LOW and CNT value is zero), the corresponding FAILED column is incremented with the number of allocation request for the same block size received thereafter. Eventually, when the memory space is freed, the current available blocks for allocation increments and the FAILED column value decreases. However, if CNT and FAILED values increase, it indicates an issue and must be resolved.

The following is sample output from the **show blocks all** command:

```
> show blocks all
Class 0, size 4
  Block   allocd_by   freed_by data size  alloccnt  dup_cnt  oper location
0x01799940 0x00000000 0x00101603    0      0        0 alloc not_specified
0x01798e80 0x00000000 0x00101603    0      0        0 alloc not_specified
0x017983c0 0x00000000 0x00101603    0      0        0 alloc not_specified
...
      Found 1000 of 1000 blocks
      Displaying 1000 of 1000 blocks
```

The following table explains each field.

Table 14: show blocks all Fields

Field	Description
Block	The block address.
allocd_by	The program address of the application that last used the block (0 if not used).
freed_by	The program address of the application that last released the block.
data size	The size of the application buffer/packet data that is inside the block.
alloccnt	The number of times this block has been used since the block came into existence.
dup_cnt	The current number of references to this block if used: 0 means 1 reference, 1 means 2 references.
oper	One of the four operations that was last performed on the block: alloc, get, put, or free.
location	The application that uses the block, or the program address of the application that last allocated the block (same as the allocd_by field).

The following is sample output from the **show blocks exhaustion history list** command:

```
> show blocks exhaustion history list
1 Snapshot created at 18:01:03 UTC Feb 19 2014:
  Snapshot created due to 16384 blocks running out

2 Snapshot created at 18:02:03 UTC Feb 19 2014:
  Snapshot created due to 16384 blocks running out

3 Snapshot created at 18:03:03 UTC Feb 19 2014:
  Snapshot created due to 16384 blocks running out
```

```
4 Snapshot created at 18:04:03 UTC Feb 19 2014:
  Snapshot created due to 16384 blocks running out
```

Related Commands

Command	Description
blocks	Increases the memory assigned to block diagnostics
clear blocks	Clears the system buffer statistics.
show conn	Shows active connections.

show bootvar

To show the boot file and configuration properties, use the **show bootvar** command.

show bootvar

Command History	Release	Modification
	6.1	This command was introduced.

Usage Guidelines

The BOOT variable specifies a list of bootable images on various devices. The CONFIG_FILE variable specifies the configuration file used during system initialization.

The output of this command is probably not meaningful for threat defense.

Examples

Following is an example of showing the boot variables for threat defense. Although the variables are empty, this example is from a functioning system.

```
> show bootvar
BOOT variable =
Current BOOT variable =
CONFIG_FILE variable =
Current CONFIG_FILE variable =
```

show bridge-group

To show bridge group information such as interfaces assigned, MAC addresses, and IP addresses, use the **show bridge-group** command.

show bridge-group [*bridge_group_number*]

Syntax Description	
<i>bridge_group_number</i>	Specifies the bridge group number as an integer between 1 and 250. If you do not specify a number, all bridge groups are shown.

Command History	Release	Modification
	6.1	This command was added.
	6.2	We added support in routed firewall mode when using Integrated Routing and Bridging.

Examples

The following is sample output from the **show bridge-group** command.

```
> show bridge-group
Static mac-address entries: 0 (in use), 16384 (max)
Dynamic mac-address entries: 0 (in use), 16384 (max)
Bridge Group: 1
Interfaces:
GigabitEthernet1/2
GigabitEthernet1/3
GigabitEthernet1/4
GigabitEthernet1/5
GigabitEthernet1/6
GigabitEthernet1/7
GigabitEthernet1/8
Management System IP Address: 192.168.1.1 255.255.255.0
Management Current IP Address: 192.168.1.1 255.255.255.0
Management IPv6 Global Unicast Address(es):
    2000:100::1, subnet is 2000:100::/64
Static mac-address entries: 0
Dynamic mac-address entries: 0
```

Related Commands	Command	Description
	show running-config interface bvi	Shows the bridge group interface configuration.