IP Unicast Routing

This chapter contains the following IP Unicast Routing commands:

ip cef load-sharing algorithm

To select a Cisco Express Forwarding load-balancing algorithm, use the**ip cef load-sharing algorithm** command in global configuration mode. To return to the default universal load-balancing algorithm, use the **no** form of this command.

ip cef load-sharing algorithm {**original** | [**universal** [*id*]]}
no ip cef load-sharing algorithm

Syntax Description

original	Sets the load-balancing algorithm to the original algorithm based on a source and destination hash.
universal	Sets the load-balancing algorithm to the universal algorithm that uses a source and destination and an ID hash.
id	(Optional) Fixed identifier.

Command Default

The universal load-balancing algorithm is selected by default. If you do not configure the fixed identifier for a load-balancing algorithm, the router automatically generates a unique ID.

Command Modes

Global configuration (config)

Command History

Release	Modification
	This command was introduced.

Usage Guidelines

The original Cisco Express Forwarding load-balancing algorithm produced distortions in load sharing across multiple devices because of the use of the same algorithm on every device. When the load-balancing algorithm is set to universal mode, each device on the network can make a different load sharing decision for each source-destination address pair, and that resolves load-balancing distortions.

Examples

The following example shows how to enable the Cisco Express Forwarding original load-balancing algorithm:

```
Device> enable
Device# configure terminal
Device(config)# ip cef load-sharing algorithm original
Device(config)# exit
```

Related Commands

Command	Description
ip load-sharing	Enables load balancing for Cisco Express Forwarding.

ip load-sharing

To enable load balancing for Cisco Express Forwarding on an interface, use the **ip load-sharing** command in interface configuration mode. To disable load balancing for Cisco Express Forwarding on the interface, use the **no** form of this command.

ip load-sharing {**per-packet** | **per-destination** }
no ip load-sharing per-packet

Syntax Description

per-packet	Enables per-packet load balancing for Cisco Express Forwarding on the interface. This functionality and keyword are not supported on all platforms. See "Usage Guidelines" for more information.
per-destination	Enables per-destination load balancing for Cisco Express Forwarding on the interface.

Command Default

Per-destination load balancing is enabled by default when you enable Cisco Express Forwarding.

Command Modes

Interface configuration (config-if)

Command History

Release	Modification
	This command was introduced.

Usage Guidelines

Per-packet load balancing allows the router to send data packets over successive equal-cost paths without regard to individual destination hosts or user sessions. Path utilization is good, but packets destined for a given destination host might take different paths and might arrive out of order.

Per-destination load balancing allows the device to use multiple, equal-cost paths to achieve load sharing. Packets for a given source-destination host pair are guaranteed to take the same path, even if multiple, equal-cost paths are available. Traffic for different source-destination host pairs tends to take different paths.

Note If you want to enable per-packet load sharing to a particular destination, then all interfaces that can forward traffic to the destination must be enabled for per-packet load sharing.

Examples

The following example shows how to enable per-packet load balancing:

```
Device> enable
Device# configure terminal
Device(config)# interface gigabitethernet 1/0/1
Device(config-if)# ip load-sharing per-packet
```

The following example shows how to enable per-destination load balancing:

```
Device> enable
Device# configure terminal
Device(config)# interface gigabitethernet 1/0/1
Device(config-if)# ip load-sharing per-destination
```

show ip bgp neighbors

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

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

Syntax Description

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

Command Default

The output of this command displays information for all neighbors.

Command Modes

User EXEC (>)

Privileged EXEC (#)

Command History

Command History

Release	Modification
Cisco IOS XE Denali 16.1.1	This command was introduced.
Cisco IOS XE Gibraltar 16.11.1	BGP Peak Prefix Watermark was added to the command output.

Usage Guidelines

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

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

Examples

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

show ip bgp neighbors: Example

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

```
Device#show ip bgp neighbors 10.108.50.2

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

```
  BGP advertise-best-external is enabled
  BGP table version 1, neighbor version 1/0
 Output queue size : 0
  Index 1, Offset 0, Mask 0x2
  1 update-group member
                               Sent       Rcvd
  Prefix activity:             ----       ----
    Prefixes Current:             0          0
    Prefixes Total:               0          0
    Implicit Withdraw:            0          0
    Explicit Withdraw:            0          0
    Used as bestpath:           n/a          0
    Used as multipath:          n/a          0
                                 Outbound    Inbound
  Local Policy Denied Prefixes:  --------    -------
    Total:                                0          0
  Number of NLRIs in the update sent: max 0, min 0
  Connections established 3; dropped 2
  Last reset 00:24:26, due to Peer closed the session
External BGP neighbor may be up to 2 hops away.
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled
Local host: 10.108.50.1, Local port: 179
Foreign host: 10.108.50.2, Foreign port: 42698
Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)
Event Timers (current time is 0x68B944):
Timer          Starts    Wakeups            Next
Retrans            27          0             0x0
TimeWait            0          0             0x0
AckHold            27         18             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0
iss: 3915509457  snduna: 3915510016  sndnxt: 3915510016     sndwnd:  15826
irs:  233567076  rcvnxt:  233567616  rcvwnd:      15845  delrcvwnd:    539
SRTT: 292 ms, RTTO: 359 ms, RTV: 67 ms, KRTT: 0 ms
minRTT: 12 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: passive open, nagle, gen tcbs
IP Precedence value : 6
Datagrams (max data segment is 1460 bytes):
Rcvd: 38 (out of order: 0), with data: 27, total data bytes: 539
Sent: 45 (retransmit: 0, fastretransmit: 0, partialack: 0, Second Congestion: 08
```

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

Table 1: show ip bgp neighbors Field Descriptions

Field	Description
BGP neighbor	IP address of the BGP neighbor and its autonomous system number.
remote AS	Autonomous system number of the neighbor.
local AS 300 no-prepend (not shown in display)	Verifies that the local autonomous system number is not prepended to received external routes. This output supports the hiding of the local autonomous systems when a network administrator is migrating autonomous systems.

Field	Description
internal link	“internal link” is displayed for iBGP neighbors; “external link” is displayed for external BGP (eBGP) neighbors.
BGP version	BGP version being used to communicate with the remote router.
remote router ID	IP address of the neighbor.
BGP state	Finite state machine (FSM) stage of session negotiation.
up for	Time, in hh:mm:ss, that the underlying TCP connection has been in existence.
Last read	Time, in hh:mm:ss, since BGP last received a message from this neighbor.
last write	Time, in hh:mm:ss, since BGP last sent a message to this neighbor.
hold time	Time, in seconds, that BGP will maintain the session with this neighbor without receiving messages.
keepalive interval	Time interval, in seconds, at which keepalive messages are transmitted to this neighbor.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor. “advertised and received” is displayed when a capability is successfully exchanged between two routers.
Route refresh	Status of the route refresh capability.
MPLS Label capability	Indicates that MPLS labels are both sent and received by the eBGP peer.
Graceful Restart Capability	Status of the graceful restart capability.
Address family IPv4 Unicast	IP Version 4 unicast-specific properties of this neighbor.
Message statistics	Statistics organized by message type.
InQ depth is	Number of messages in the input queue.
OutQ depth is	Number of messages in the output queue.
Sent	Total number of transmitted messages.
Revd	Total number of received messages.
Opens	Number of open messages sent and received.
Notifications	Number of notification (error) messages sent and received.
Updates	Number of update messages sent and received.
Keepalives	Number of keepalive messages sent and received.

Field	Description
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 is the primary routing table with which the neighbor has been updated. The number increments when the table changes.
neighbor version	Number used by the software to track prefixes that have been sent and those that need to be sent.
1 update-group member	Number of the update-group member for this address family.
Prefix activity	Prefix statistics for this address family.
Prefixes Current	Number of prefixes accepted for this address family.
Prefixes Total	Total number of received prefixes.
Implicit Withdraw	Number of times that a prefix has been withdrawn and readvertised.
Explicit Withdraw	Number of times that a prefix has been withdrawn because it is no longer feasible.
Used as bestpath	Number of received prefixes installed as best paths.
Used as multipath	Number of received prefixes installed as multipaths.
* Saved (soft-reconfig)	Number of soft resets performed with a neighbor that supports soft reconfiguration. This field is displayed only if the counter has a nonzero value.
* History paths	This field is displayed only if the counter has a nonzero value.
* Invalid paths	Number of invalid paths. This field is displayed only if the counter has a nonzero value.
Local Policy Denied Prefixes	Prefixes denied due to local policy configuration. Counters are updated for inbound and outbound policy denials. The fields under this heading are displayed only if the counter has a nonzero value.
* route-map	Displays inbound and outbound route-map policy denials.
* filter-list	Displays inbound and outbound filter-list policy denials.
* prefix-list	Displays inbound and outbound prefix-list policy denials.
* Ext Community	Displays only outbound extended community policy denials.
* AS_PATH too long	Displays outbound AS_PATH length policy denials.

Field	Description
* AS_PATH loop	Displays outbound AS_PATH loop policy denials.
* AS_PATH confed info	Displays outbound confederation policy denials.
* AS_PATH contains AS 0	Displays outbound denials of autonomous system 0.
* NEXT_HOP Martian	Displays outbound martian denials.
* NEXT_HOP non-local	Displays outbound nonlocal next-hop denials.
* NEXT_HOP is us	Displays outbound next-hop-self denials.
* CLUSTER_LIST loop	Displays outbound cluster-list loop denials.
* ORIGINATOR loop	Displays outbound denials of local originated routes.
* unsuppress-map	Displays inbound denials due to an unsuppress map.
* advertise-map	Displays inbound denials due to an advertise map.
* VPN Imported prefix	Displays inbound denials of VPN prefixes.
* Well-known Community	Displays inbound denials of well-known communities.
* SOO loop	Displays inbound denials due to site-of-origin.
* Bestpath from this peer	Displays inbound denials because the best path came from the local router.
* Suppressed due to dampening	Displays inbound denials because the neighbor or link is in a dampening state.
* Bestpath from iBGP peer	Deploys inbound denials because the best path came from an iBGP neighbor.
* Incorrect RIB for CE	Deploys inbound denials due to RIB errors for a customer edge (CE) router.
* BGP distribute-list	Displays inbound denials due to a distribute list.
Number of NLRIs...	Number of network layer reachability attributes in updates.
Current session network count peaked...	Displays the peak number of networks observed in the current session.
Highest network count observed at...	Displays the peak number of networks observed since startup.
Connections established	Number of times a TCP and BGP connection has been successfully established.
dropped	Number of times that a valid session has failed or been taken down.
Last reset	Time, in hh:mm:ss, since this peering session was last reset. The reason for the reset is displayed on this line.

Field	Description
External BGP neighbor may be...	Indicates that the BGP time to live (TTL) security check is enabled. The maximum number of hops that can separate the local and remote peer is displayed on this line.
Connection state	Connection status of the BGP peer.
unread input bytes	Number of bytes of packets still to be processed.
Connection is ECN Disabled	Explicit congestion notification status (enabled or disabled).
Local host: 10.108.50.1, Local port: 179	IP address of the local BGP speaker. BGP port number 179.
Foreign host: 10.108.50.2, Foreign port: 42698	Neighbor address and BGP destination port number.
Enqueued packets for retransmit:	Packets queued for retransmission by TCP.
Event Timers	TCP event timers. Counters are provided for starts and wakeups (expired timers).
Retrans	Number of times a packet has been retransmitted.
TimeWait	Time waiting for the retransmission timers to expire.
AckHold	Acknowledgment hold timer.
SendWnd	Transmission (send) window.
KeepAlive	Number of keepalive packets.
GiveUp	Number of times a packet is dropped due to no acknowledgment.
PmtuAger	Path MTU discovery timer.
DeadWait	Expiration timer for dead segments.
iss:	Initial packet transmission sequence number.
snduna:	Last transmission sequence number that has not been acknowledged.
sndnxt:	Next packet sequence number to be transmitted.
sndwnd:	TCP window size of the remote neighbor.
irs:	Initial packet receive sequence number.
rcvnxt:	Last receive sequence number that has been locally acknowledged.
rcvwnd:	TCP window size of the local host.

Field	Description
delrcvwnd:	Delayed receive window—data the local host has read from the connection, but has not yet subtracted from the receive window the host has advertised to the remote host. The value in this field gradually increases until it is higher than a full-sized packet, at which point it is applied to the rcvwnd field.
SRTT:	A calculated smoothed round-trip timeout.
RTTO:	Round-trip timeout.
RTV:	Variance of the round-trip time.
KRTT:	New round-trip timeout (using the Karn algorithm). This field separately tracks the round-trip time of packets that have been re-sent.
minRTT:	Shortest recorded round-trip timeout (hard-wire value used for calculation).
maxRTT:	Longest recorded round-trip timeout.
ACK hold:	Length of time the local host will delay an acknowledgment to carry (piggyback) additional data.
IP Precedence value:	IP precedence of the BGP packets.
Datagrams	Number of update packets received from a neighbor.
Rcvd:	Number of received packets.
out of order:	Number of packets received out of sequence.
with data	Number of update packets sent with data.
total data bytes	Total amount of data received, in bytes.
Sent	Number of update packets sent.
Second Congestion	Number of update packets with data sent.
Datagrams: Rcvd	Number of update packets received from a neighbor.
retransmit	Number of packets retransmitted.
fastretransmit	Number of duplicate acknowledgments retransmitted for an out of order segment before the retransmission timer expires.
partialack	Number of retransmissions for partial acknowledgments (transmissions before or without subsequent acknowledgments).
Second Congestion	Number of second retransmissions sent due to congestion.

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

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

```
Device#show ip bgp neighbors

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

show ip bgp neighbors advertised-routes

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

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

BGP table version is 27, local router ID is 172.16.232.181
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop          Metric LocPrf Weight Path
*>i10.0.0.0      172.16.232.179         0    100      0 ?
*> 10.20.2.0     10.0.0.0               0         32768 i
```

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

Table 2: show ip bgp neighbors advertised-routes Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This is the primary routing table with which the neighbor has been updated. The number increments when the table changes.
local router ID	IP address of the local BGP speaker.

Field	Description
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: • s—The table entry is suppressed. • d—The table entry is dampened and will not be advertised to BGP neighbors. • h—The table entry does not contain the best path based on historical information. • *—The table entry is valid. • >—The table entry is the best entry to use for that network. • i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: • i—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. • e—Entry originated from Exterior Gateway Protocol (EGP). • ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	IP address of a network entity.
Next Hop	IP address of the next system used to forward a packet to the destination network. An entry of 0.0.0.0 indicates that there are non-BGP routes in the path to the destination network.
Metric	If shown, this is the value of the interautonomous system metric. This field is not used frequently.
LocPrf	Local preference value as set with the **set local-preference** route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

show ip bgp neighbors check-control-plane-failure

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

```
Device#show ip bgp neighbors 10.10.10.1

BGP neighbor is 10.10.10.1,  remote AS 10, internal link
 Fall over configured for session
 BFD is configured. BFD peer is Up. Using BFD to detect fast fallover (single-hop) with
c-bit check-control-plane-failure.
```

```
 Inherits from template cbit-tps for session parameters
  BGP version 4, remote router ID 10.7.7.7
  BGP state = Established, up for 00:03:55
  Last read 00:00:02, last write 00:00:21, hold time is 180, keepalive interval is 60 seconds

  Neighbor sessions:
    1 active, is not multisession capable (disabled)
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    Enhanced Refresh Capability: advertised and received
    Multisession Capability:
    Stateful switchover support enabled: NO for session 1
```

show ip bgp neighbors paths

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

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

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

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

Table 3: show ip bgp neighbors paths Field Descriptions

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

show ip bgp neighbors received prefix-filter

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

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

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

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

Table 4: show ip bgp neighbors received prefix-filter Field Descriptions

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

show ip bgp neighbors policy

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

```
Device#show ip bgp neighbors 192.168.1.2 policy

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

BGP Attribute Filter and Enhanced Attribute Error Handling

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

```
Device#show ip bgp vpnv4 all neighbors 10.0.103.1

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

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

BGP Additional Paths

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

```
Device#show ip bgp neighbors 10.108.50.2
```

```
BGP neighbor is 10.108.50.2,  remote AS 1, internal link
  BGP version 4, remote router ID 192.168.252.252
  BGP state = Established, up for 00:24:25
  Last read 00:00:24, last write 00:00:24, hold time is 180, keepalive interval is 60 seconds

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

BGP—Multiple Cluster IDs

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

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

 Configured with the cluster-id 192.168.15.6
```

BGP Peak Prefix Watermark

The following sample output shows the peak watermarks and their timestamps displayed for the peak number of route entries per neighbor bases:

```
Device#show ip bgp ipv4 unicast neighbors 11.11.11.11

BGP neighbor is 11.11.11.11,  remote AS 1, internal link
  BGP version 4, remote router ID 0.0.0.0
  BGP state = Idle, down for 00:01:43
  Neighbor sessions:
    0 active, is not multisession capable (disabled)
    Stateful switchover support enabled: NO
  Do log neighbor state changes (via global configuration)
  Default minimum time between advertisement runs is 0 seconds

 For address family: IPv4 Unicast
  BGP table version 27, neighbor version 1/27
  Output queue size : 0
  Index 0, Advertise bit 0
  Slow-peer detection is disabled
  Slow-peer split-update-group dynamic is disabled
                               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
    Used as secondary:          n/a          0

                                 Outbound    Inbound
  Local Policy Denied Prefixes:  --------    -------
    Total:                              0          0
  Number of NLRIs in the update sent: max 2, min 0
```

```
  Current session network count peaked at 20 entries at 00:00:23 Aug 8 2018 PST (00:01:29.156
 ago).
  Highest network count observed at 20 entries at 23:55:32 Aug 7 2018 PST (00:06:20.156
ago).
  Last detected as dynamic slow peer: never
  Dynamic slow peer recovered: never
  Refresh Epoch: 1
  Last Sent Refresh Start-of-rib: never
  Last Sent Refresh End-of-rib: never
  Last Received Refresh Start-of-rib: never
  Last Received Refresh End-of-rib: never
                                       Sent       Rcvd
        Refresh activity:              ----       ----
          Refresh Start-of-RIB          0          0
          Refresh End-of-RIB            0          0
```

BGP Soft Inbound and Outbound Refresh Time

In the following example, the times of occurrence of the soft inbound and outbound refresh, to or from the given neighbour, are displayed:

```
Device#show ip bgp l2vpn evpn neighbors 11.11.11.11

BGP neighbor is 11.11.11.11,  remote AS 1, internal link
  BGP version 4, remote router ID 11.11.11.11
  BGP state = Established, up for 00:14:06
  Last read 00:00:21, last write 00:00:28, hold time is 180, keepalive
  ……
  Do log neighbor state changes (via global configuration)

  Default minimum time between advertisement runs is 0 seconds


 For address family: L2VPN E-VPN
  Session: 11.11.11.11
  BGP table version 30, neighbor version 30/0
  Output queue size : 0
  Index 1, Advertise bit 0
  1 update-group member
  Community attribute sent to this neighbor
  Extended-community attribute sent to this neighbor
  …….
  ……
  Last detected as dynamic slow peer: never
  Dynamic slow peer recovered: never
  Refresh Epoch: 2
  Last Sent Refresh Start-of-rib: never
  Last Sent Refresh End-of-rib: never
  Last Received Refresh Start-of-rib: 00:14:06
  Last Received Refresh End-of-rib: 00:14:06
  Refresh-In took 0 seconds
                                       Sent       Rcvd
        Refresh activity:              ----       ----
          Refresh Start-of-RIB          0          1
          Refresh End-of-RIB            0          1

  Address tracking is enabled, the RIB does have a route to 11.11.11.11
  Route to peer address reachability Up: 1; Down: 0
    Last notification 00:14:07
  Connections established 1; dropped 0
……
```

```
.............
Packets received in fast path: 0, fast processed: 0, slow path: 0
 fast lock acquisition failures: 0, slow path: 0
TCP Semaphore      0x7FA8A0AE7BA0  FREE
```

Related Commands

Command	Description
bgp asnotation dot	Changes the default display and the regular expression match format of BGP 4-byte autonomous system numbers from asplain (decimal values) to dot notation.
bgp enhanced-error	Restores the default behavior of treating Update messages that have a malformed attribute as withdrawn, or includes iBGP peers in the Enhanced Attribute Error Handling feature.
neighbor path-attribute discard	Configures the device to discard unwanted Update messages from the specified neighbor that contain a specified path attribute.
neighbor path-attribute treat-as-withdraw	Configures the device to withdraw from the specified neighbor unwanted Update messages that contain a specified attribute.
neighbor send-label	Enables a BGP router to send MPLS labels with BGP routes to a neighboring BGP router.
neighbor send-label explicit-null	Enables a BGP router to send MPLS labels with explicit-null information for a CSC-CE router and BGP routes to a neighboring CSC-PE router.
router bgp	Configures the BGP routing process.

show tech-support bgp

To automatically run show commands that display BGP related system information, use the **show tech-support bgp** command in the privileged EXEC mode.

show tech-support bgp [**address-family** {**all** | **ipv4** [**flowspec** | **multicast** | **unicast** | [**mdt** | **mvpn**] {**all** | **vrf** *vrf-instance-name*}] |**ipv6** [**flowspec** | **multicast** | **mvpn** {**all** | **vrf** *vrf-instance-name*} | **unicast**] | **l2vpn** [**evpn** | **vpls**] | **link-state** [**link-state**] | [**nsap** | **rtfilter**] [**unicast**] | [**vpnv4** | **vpnv6**] [**flowspec** | **multicast** | **unicast**] {**all** | **vrf** *vrf-instance-name*}}] [**detail**]

Syntax Description

address-family	(Optional) Displays the output for a specified address family.
address-family all	(Optional) Displays the output for all address families.
ipv4	(Optional) Displays the output for IPv4 address family.
ipv6	(Optional) Displays the output for IPv6 address family.
l2vpn	(Optional) Displays the output for L2VPN address family.
link-state	(Optional) Displays the output for Link State address family.
nsap	(Optional) Displays the output for NSAP address family.
rtfilter	(Optional) Displays the output for RT Filter address family.
vpnv4	(Optional) Displays the output for VPNv4 address family.
vpnv6	(Optional) Displays the output for VPNv6 address family.
flowspec	(Optional) Displays the flowspec related information for an address family.
multicast	(Optional) Displays the multicast related information for an address family.
unicast	(Optional) Displays the unicast related information for an address family.
mdt	(Optional) Displays the Multicast Distribution Tree (MDT) related information for an address family.

mvpn	(Optional) Displays the Multicast VPN (MVPN) related information for an address family.
vrf	Displays the information for a VPN Routing/Forwarding instance.
evpn	(Optional) Displays the Ethernet VPN (EVPN) related information for an address family.
vpls	(Optional) Displays the Virtual Private LAN Services (VPLS) related information for an address family.
vrf-instance-name	Specifies the name of the VPN Routing/Forwarding instance.
all	Displays the information about all VPN NLRIs.
detail	(Optional) Displays the detailed routes information.

Command Modes

User EXEC (>)

Privileged EXEC (#)

Command History

Release	Modification
Cisco IOS XE Gibraltar 16.11.1	This command was introduced.

Usage Guidelines

The **show tech-support bgp** command is used to display the outputs of various BGP show commands and log them to the show-tech file. The output from the **show tech-support bgp** com mand is very long. To better manage this output, you can redirect the output to a file (for example, **show tech-support** > **filename**) in the local writable storage file system or the remote file system. Redirecting the output to a file also makes sending the output to your Cisco Technical Assistance Center (TAC) representative easier.

You can use one of the following redirection methods:

- > filename - Redirects the output to a file.
- >> filename - Redirects the output to a file in append mode.

The following **show** commands run automatically when the **show tech-support bgp** command is used:

- **show clock**
- **show version**
- **show running-config**
- **show process cpu sorted**
- **show process cpu history**
- **show process memory sorted**

The following **show** commands for a specific address family run automnatically when the **show tech-support bgp address-family***address-family-name address-family-modifier* command is used:

- **show bgp** *address-family-name address-family-modifier* **summary**
- **show bgp** *address-family-name address-family-modifier* **detail**
- **show bgp** *address-family-name address-family-modifier* **internal**
- **show bgp** *address-family-name address-family-modifier* **neighbors**
- **show bgp** *address-family-name address-family-modifier* **update-group**
- **show bgp** *address-family-name address-family-modifier* **replication**
- **show bgp** *address-family-name address-family-modifier* **community**
- **show bgp** *address-family-name address-family-modifier* **dampening dampened-paths**
- **show bgp** *address-family-name address-family-modifier* **dampening flap-statistics**
- **show bgp** *address-family-name address-family-modifier* **dampening parameters**
- **show bgp** *address-family-name address-family-modifier* **injected-paths**
- **show bgp** *address-family-name address-family-modifier* **cluster-ids**
- **show bgp** *address-family-name address-family-modifier* **cluster-ids internal**
- **show bgp** *address-family-name address-family-modifier* **peer-group**
- **show bgp** *address-family-name address-family-modifier* **pending-prefixes**
- **show bgp** *address-family-name address-family-modifier* **rib-failure**

In addition to the above commands, the following segment routing specific **show** commands also run when the **show tech-support bgp** command is used:

- **show bgp all binding-sid**
- **show segment-routing client**
- **show segment-routing mpls state**
- **show segment-routing mpls gb**
- **show segment-routing mpls connected-prefix-sid-map protocol ipv4**
- **show segment-routing mpls connected-prefix-sid-map protocol backup ipv4**
- **show mpls traffic-eng tunnel auto-tunnel client bgp**