



Configuring BGP Support for TCP Path MTU Discovery per Session

This module describes configuration tasks to configure BGP support for TCP path MTU discovery. BGP is an interdomain routing protocol designed to provide loop-free routing between organizations. This module contains tasks that use BGP neighbor session commands to configure:

- Disabling TCP path MTU for all sessions
- Disabling TCP path for single neighbor session
- [Finding Feature Information, on page 1](#)
- [Restrictions for Configuring BGP Support for TCP Path MTU Discovery per Session, on page 2](#)
- [Information About BGP Support for TCP Path MTU Discovery per Session, on page 2](#)
- [Configuring BGP Support for TCP Path MTU Discovery per Session, on page 3](#)
- [Examples: Configuring BGP Support for TCP Path MTU Discovery per Session, on page 11](#)
- [Additional References, on page 19](#)
- [Feature Information for Configuring BGP Support for TCP Path MTU Discovery per Session, on page 20](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see [Bug Search Tool](#) and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for Configuring BGP Support for TCP Path MTU Discovery per Session

A router that runs Cisco software can be configured to run only one BGP routing process and to be a member of only one BGP autonomous system. However, a BGP routing process and autonomous system can support multiple address family configurations.

Information About BGP Support for TCP Path MTU Discovery per Session

Path MTU Discovery

The IP protocol family was designed to use a wide variety of transmission links. The maximum IP packet length is 65000 bytes. Most transmission links enforce a smaller maximum packet length limit, called the maximum transmission unit (MTU), which varies with the type of the transmission link. The design of IP accommodates link packet length limits by allowing intermediate routers to fragment IP packets as necessary for their outgoing links. The final destination of an IP packet is responsible for reassembling its fragments as necessary.

All TCP sessions are bounded by a limit on the number of bytes that can be transported in a single packet, and this limit is known as the maximum segment size (MSS). TCP breaks up packets into chunks in a transmit queue before passing packets down to the IP layer. A smaller MSS may not be fragmented at an IP device along the path to the destination device, but smaller packets increase the amount of bandwidth needed to transport the packets. The maximum TCP packet length is determined by both the MTU of the outbound interface on the source device and the MSS announced by the destination device during the TCP setup process.

Path MTU discovery (PMTUD) was developed as a solution to the problem of finding the optimal TCP packet length. PMTUD is an optimization (detailed in RFC 1191) wherein a TCP connection attempts to send the longest packets that will not be fragmented along the path from source to destination. It does this by using a flag, don't fragment (DF), in the IP packet. This flag is supposed to alter the behavior of an intermediate router that cannot send the packet across a link because it is too long. Normally the flag is off, and the router should fragment the packet and send the fragments. If a router tries to forward an IP datagram, with the DF bit set, to a link that has a lower MTU than the size of the packet, the router will drop the packet and return an ICMP Destination Unreachable message to the source of this IP datagram, with the code indicating "fragmentation needed and DF set." When the source device receives the ICMP message, it will lower the send MSS, and when TCP retransmits the segment, it will use the smaller segment size.

BGP Neighbor Session TCP PMTUD

TCP path MTU discovery is enabled by default for all BGP neighbor sessions, but there are situations when you may want to disable TCP path MTU discovery for one or all BGP neighbor sessions. Although PMTUD works well for larger transmission links (for example, Packet over Sonet links), a badly configured TCP implementation or a firewall may slow or stop the TCP connections from forwarding any packets. In this type of situation, you may need to disable TCP path MTU discovery.

In Cisco software, configuration options were introduced to permit TCP path MTU discovery to be disabled, or subsequently reenabled, either for a single BGP neighbor session or for all BGP sessions. To disable the TCP path MTU discovery globally for all BGP neighbors, use the **no bgp transport path-mtu-discovery** command in router configuration mode. To disable the TCP path MTU discovery for a single neighbor, use the **no neighbor transport path-mtu-discovery** command in router configuration mode or address family configuration mode. For more details, see the “Disabling TCP Path MTU Discovery Globally for All BGP Sessions” section or the “Disabling TCP Path MTU Discovery for a Single BGP Neighbor” section.

Configuring BGP Support for TCP Path MTU Discovery per Session

This section contains the following tasks:

Disabling TCP Path MTU Discovery Globally for All BGP Sessions

Perform this task to disable TCP path MTU discovery for all BGP sessions. TCP path MTU discovery is enabled by default when you configure BGP sessions, but we recommend that you enter the **show ip bgp neighbors** command to ensure that TCP path MTU discovery is enabled.

Before you begin

This task assumes that you have previously configured BGP neighbors with active TCP connections.

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	show ip bgp neighbors [<i>ip-address</i>] Example: Device# show ip bgp neighbors	(Optional) Displays information about the TCP and BGP connections to neighbors. <ul style="list-style-type: none"> • Use this command to determine whether BGP neighbors have TCP path MTU discovery enabled. <p>Note Only the syntax applicable to this task is used in this example. For more details, see the <i>Cisco IOS IP Routing: BGP Command Reference</i>.</p>
Step 3	configure terminal Example: Device# configure terminal	Enters global configuration mode.

	Command or Action	Purpose
Step 4	router bgp <i>autonomous-system-number</i> Example: <pre>Device(config)# router bgp 50000</pre>	Enters router configuration mode to create or configure a BGP routing process.
Step 5	no bgp transport path-mtu-discovery Example: <pre>Device(config-router)# no bgp transport path-mtu-discovery</pre>	Disables TCP path MTU discovery for all BGP sessions.
Step 6	end Example: <pre>Device(config-router)# end</pre>	Exits router configuration mode and returns to privileged EXEC mode.
Step 7	show ip bgp neighbors [<i>ip-address</i>] Example: <pre>Device# show ip bgp neighbors</pre>	(Optional) Displays information about the TCP and BGP connections to neighbors. <ul style="list-style-type: none"> In this example, the output from this command will not display that any neighbors have TCP path MTU enabled. <p>Note Only the syntax applicable to this task is used in this example. For more details, see the <i>Cisco IOS IP Routing: BGP Command Reference</i>.</p>

Examples

The following sample output from the **show ip bgp neighbors** command shows that TCP path MTU discovery is enabled for BGP neighbors. Two entries in the output—Transport(tcp) path-mtu-discovery is enabled and path mtu capable—show that TCP path MTU discovery is enabled.

```
Router# show ip bgp neighbors

BGP neighbor is 3.3.3.3, remote AS 65535, internal link
  BGP version 4, remote router ID 3.3.3.3
  .
  .
  .
For address family: VPNv4 Unicast
  Session: 3.3.3.3
  BGP table version 1, neighbor version 1/0
  .
  .
  .
Address tracking is enabled, the RIB does have a route to 3.3.3.3
Connections established 7; dropped 6
Last reset 00:00:12, due to Peer closed the session of session 1
Transport(tcp) path-mtu-discovery is enabled
```

```
.
.
.
SRTT: 487 ms, RTTO: 3168 ms, RTV: 2681 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 1000 ms, ACK hold: 200 ms
uptime: 5612 ms, Sent idletime: 4588 ms, Receive idletime: 4388 ms
Status Flags: active open
```

The following is sample output from the **show ip bgp neighbors** command after the **no bgp transport path-mtu-discovery** command has been entered. Note that the path mtu entries are missing.

```
Router# show ip bgp neighbors

BGP neighbor is 3.3.3.3, remote AS 65535, internal link
  BGP version 4, remote router ID 3.3.3.3
.
.
.
For address family: IPv4 Unicast
  Session: 3.3.3.3
  BGP table version 1, neighbor version 1/0
.
.
.
  Address tracking is enabled, the RIB does have a route to 3.3.3.3
  Connections established 8; dropped 7
  Last reset 00:00:13, due to Peer closed the session of session 1
  Transport(tcp) path-mtu-discovery is disabled
.
.
.
SRTT: 413 ms, RTTO: 3205 ms, RTV: 2792 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 1000 ms, ACK hold: 200 ms
uptime: 13472 ms, Sent idletime: 12248 ms, Receive idletime: 12448 ms
Status Flags: passive open, gen tcbs
```

Disabling TCP Path MTU Discovery for a Single BGP Neighbor

Perform this task to establish a peering session with an internal BGP (iBGP) neighbor and then disable TCP path MTU discovery for the BGP neighbor session. The **neighbor transport** command can be used in router configuration mode or address family configuration mode.

Before you begin

This task assumes that you know that TCP path MTU discovery is enabled by default for all your BGP neighbors.

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.

	Command or Action	Purpose
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	router bgp <i>autonomous-system-number</i> Example: Device(config)# router bgp 45000	Enters router configuration mode for the specified routing process.
Step 4	address-family {ipv4 [mdt multicast unicast [vrf <i>vrf-name</i>] vrf <i>vrf-name</i>] vpnv4 [unicast]} Example: Device(config-router)# address-family ipv4 unicast	Enters address family configuration mode to configure BGP peers to accept address-family-specific configurations. <ul style="list-style-type: none"> The example creates an IPv4 unicast address family session.
Step 5	neighbor {<i>ip-address</i> <i>peer-group-name</i>} remote-as <i>autonomous-system-number</i> Example: Device(config-router-af)# neighbor 192.168.1.1 remote-as 45000	Adds the IP address or peer group name of the neighbor in the specified autonomous system to the IPv4 multiprotocol BGP neighbor table of the local router.
Step 6	neighbor {<i>ip-address</i> <i>peer-group-name</i>} activate Example: Device(config-router-af)# neighbor 172.16.1.1 activate	Activates the neighbor under the IPv4 address family.
Step 7	neighbor {<i>ip-address</i> <i>peer-group-name</i>} transport {connection-mode path-mtu-discovery} { enable disable } Example: Device(config-router-af)# neighbor 172.16.1.1 transport path-mtu-discovery disable	Disables TCP path MTU discovery for a single BGP neighbor. <ul style="list-style-type: none"> In this example, TCP path MTU discovery is disabled for the neighbor at 172.16.1.1.
Step 8	end Example: Device(config-router-af)# end	Exits address family configuration mode and returns to privileged EXEC mode.
Step 9	show ip bgp neighbors Example:	(Optional) Displays information about the TCP and BGP connections to neighbors.

	Command or Action	Purpose
	<pre>Device# show ip bgp neighbors</pre>	<ul style="list-style-type: none"> In this example, the output from this command will not display that the neighbor has TCP path MTU discovery enabled. <p>Note Only the syntax applicable to this task is used in this example. For more details, see the <i>Cisco IOS IP Routing: BGP Command Reference</i>.</p>

Examples

The following sample output shows that TCP path MTU discovery has been disabled for BGP neighbor 172.16.1.1 but that it is still enabled for BGP neighbor 192.168.2.2. Two entries in the output—Transport(tcp) path-mtu-discovery is enabled and path mtu capable—show that TCP path MTU discovery is enabled.

```
Router# show ip bgp neighbors
BGP neighbor is 172.16.1.1, remote AS 45000, internal link
  BGP version 4, remote router ID 172.17.1.99
  .
  .
  .
  Address tracking is enabled, the RIB does have a route to 172.16.1.1
  Connections established 1; dropped 0
  Last reset never
  Transport(tcp) path-mtu-discovery is disabled
  .
  .
  .
  SRTT: 165 ms, RTTO: 1172 ms, RTV: 1007 ms, KRTT: 0 ms
  minRTT: 20 ms, maxRTT: 300 ms, ACK hold: 200 ms
  Flags: higher precedence, retransmission timeout, nagle
  .
  .
  .
  BGP neighbor is 192.168.2.2, remote AS 50000, external link
  BGP version 4, remote router ID 10.2.2.99
  .
  .
  .
  For address family: IPv4 Unicast
  BGP table version 4, neighbor version 4/0
  .
  .
  .
  Address tracking is enabled, the RIB does have a route to 192.168.2.2
  Connections established 2; dropped 1
  Last reset 00:05:11, due to User reset
  Transport(tcp) path-mtu-discovery is enabled
  .
  .
  .
  SRTT: 210 ms, RTTO: 904 ms, RTV: 694 ms, KRTT: 0 ms
  minRTT: 20 ms, maxRTT: 300 ms, ACK hold: 200 ms
  Flags: higher precedence, retransmission timeout, nagle, path mtu capable
```

Enabling TCP Path MTU Discovery Globally for All BGP Sessions

Perform this task to enable TCP path MTU discovery for all BGP sessions. TCP path MTU discovery is enabled by default when you configure BGP sessions, but if the BGP Support for TCP Path MTU Discovery per Session feature has been disabled, you can use this task to reenable it. To verify that TCP path MTU discovery is enabled, use the **show ip bgp neighbors** command.

Before you begin

This task assumes that you have previously configured BGP neighbors with active TCP connections.

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	router bgp <i>autonomous-system-number</i> Example: Device(config)# router bgp 45000	Enters router configuration mode to create or configure a BGP routing process.
Step 4	bgp transport path-mtu-discovery Example: Device(config-router)# bgp transport path-mtu-discovery	Enables TCP path MTU discovery for all BGP sessions.
Step 5	end Example: Device(config-router)# end	Exits router configuration mode and returns to privileged EXEC mode.
Step 6	show ip bgp neighbors Example: Device# show ip bgp neighbors	(Optional) Displays information about the TCP and BGP connections to neighbors. • In this example, the output from this command will show that all neighbors have TCP path MTU discovery enabled.

	Command or Action	Purpose
		Note Only the syntax applicable to this task is used in this example. For more details, see the <i>Cisco IOS IP Routing: BGP Command Reference</i> .

Examples

The following sample output from the **show ip bgp neighbors** command shows that TCP path MTU discovery is enabled for BGP neighbors. Two entries in the output—Transport(tcp) path-mtu-discovery is enabled and path mtu capable—show that TCP path MTU discovery is enabled.

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

Enabling TCP Path MTU Discovery for a Single BGP Neighbor

Perform this task to establish a peering session with an eBGP neighbor and then enable TCP path MTU discovery for the BGP neighbor session. The **neighbor transport** command can be used in router configuration mode or address family configuration mode.

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	configure terminal Example:	Enters global configuration mode.

	Command or Action	Purpose
	Device# configure terminal	
Step 3	router bgp <i>autonomous-system-number</i> Example: Device(config)# router bgp 45000	Enters router configuration mode for the specified routing process.
Step 4	address-family { ipv4 [mdt multicast unicast [vrf vrf-name] vrf vrf-name] vpn4 [unicast]} Example: Device(config-router)# address-family ipv4 unicast	Enters address family configuration mode to configure BGP peers to accept address-family-specific configurations. <ul style="list-style-type: none"> The example creates an IPv4 unicast address family session.
Step 5	neighbor { <i>ip-address</i> <i>peer-group-name</i> } remote-as <i>autonomous-system-number</i> Example: Device(config-router-af)# neighbor 192.168.2.2 remote-as 50000	Adds the IP address or peer group name of the neighbor in the specified autonomous system to the IPv4 multiprotocol BGP neighbor table of the local router.
Step 6	neighbor { <i>ip-address</i> <i>peer-group-name</i> } activate Example: Device(config-router-af)# neighbor 192.168.2.2 activate	Activates the neighbor under the IPv4 address family.
Step 7	neighbor { <i>ip-address</i> <i>peer-group-name</i> } transport { connection-mode path-mtu-discovery } Example: Device(config-router-af)# neighbor 192.168.2.2 transport path-mtu-discovery	Enables TCP path MTU discovery for a single BGP neighbor.
Step 8	end Example: Device(config-router-af)# end	Exits address family configuration mode and returns to privileged EXEC mode.
Step 9	show ip bgp neighbors [<i>ip-address</i>] Example: Device# show ip bgp neighbors 192.168.2.2	(Optional) Displays information about the TCP and BGP connections to neighbors. <p>Note Only the syntax applicable to this task is used in this example. For more details, see the <i>Cisco IOS IP Routing: BGP Command Reference</i>.</p>

Examples

The following sample output from the **show ip bgp neighbors** command shows that TCP path MTU discovery is enabled for the BGP neighbor at 192.168.2.2. Two entries in the output—Transport(tcp) path-mtu-discovery is enabled and path-mtu capable—show that TCP path MTU discovery is enabled.

```
Router# show ip bgp neighbors 192.168.2.2
BGP neighbor is 192.168.2.2, remote AS 50000, external link
  BGP version 4, remote router ID 10.2.2.99
  .
  .
  .
  For address family: IPv4 Unicast
    BGP table version 4, neighbor version 4/0
    .
    .
    .
    Address tracking is enabled, the RIB does have a route to 192.168.2.2
    Address tracking requires at least a /24 route to the peer
    Connections established 2; dropped 1
    Last reset 00:05:11, due to User reset
    Transport(tcp) path-mtu-discovery is enabled
    .
    .
    .
    SRTT: 210 ms, RTTO: 904 ms, RTV: 694 ms, KRTT: 0 ms
    minRTT: 20 ms, maxRTT: 300 ms, ACK hold: 200 ms
    Flags: higher precedence, retransmission timeout, nagle, path mtu capable
```

Examples: Configuring BGP Support for TCP Path MTU Discovery per Session

This section contains the following configuration examples:

Example: Configuring When TCP Path MTU Discovery is Enabled

The following example shows how to configure when TCP path MTU discovery is enabled. Use the **show ip bgp neighbors** command to verify that TCP path MTU discovery has been enabled.

```
router bgp 65535
  bgp log-neighbor-changes
  neighbor 3.3.3.3 remote-as 65535
  neighbor 3.3.3.3 update-source Loopback0
  !
  address-family vpnv4
  neighbor 3.3.3.3 activate
  neighbor 3.3.3.3 send-community extended
  exit-address-family

Router# show ip bgp neighbors
```

```
BGP neighbor is 3.3.3.3, remote AS 65535, internal link
BGP version 4, remote router ID 3.3.3.3
```

Example: Configuring When TCP Path MTU Discovery is Enabled

```

BGP state = Established, up for 00:00:05
Last read 00:00:05, last write 00:00:04, hold time is 180, keepalive interval is
 60 seconds
Neighbor sessions:
 1 active, is not multisession capable (disabled)
Neighbor capabilities:
Route refresh: advertised and received(new)
Four-octets ASN Capability: advertised and received
Address family IPv4 Unicast: advertised and received
Address family VPNv4 Unicast: advertised and received
Enhanced Refresh Capability: advertised and received
Multisession Capability:
Stateful switchover support enabled: NO for session 1

Message statistics:
InQ depth is 0
OutQ depth is 0

Sent Rcvd
Opens: 1 1
Notifications: 0 0
Updates: 2 2
Keepalives: 2 2
Route Refresh: 0 0
Total: 5 5
Default minimum time between advertisement runs is 0 seconds
For address family: IPv4 Unicast
Session: 3.3.3.3
BGP table version 1, neighbor version 1/0
Output queue size : 0
Index 7, Advertise bit 0
7 update-group member
Slow-peer detection is disabled
Slow-peer split-update-group dynamic is disabled

Interface associated: (none)
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
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

For address family: VPNv4 Unicast
Session: 3.3.3.3
BGP table version 1, neighbor version 1/0

```

```

Output queue size : 0
Index 6, Advertise bit 0
6 update-group member
Slow-peer detection is disabled
Slow-peer split-update-group dynamic is disabled
Interface associated: (none)

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
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
Address tracking is enabled, the RIB does have a route to 3.3.3.3
Connections established 7; dropped 6
Last reset 00:00:12, due to Peer closed the session of session 1
Transport(tcp) path-mtu-discovery is enabled
Graceful-Restart is disabled
SSO is disabled
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
Local host: 1.1.1.1, Local port: 19072
Foreign host: 3.3.3.3, Foreign port: 179
Connection tableid (VRF): 0
Maximum output segment queue size: 50

Enqueued packets for retransmit: 0, input: 0 mis-ordered: 0 (0 bytes)
Event Timers (current time is 0x4CB8D584):
Timer Starts Wakeups Next
Retrans 5 0 0x0
TimeWait 0 0 0x0
AckHold 4 1 0x0
SendWnd 0 0 0x0
KeepAlive 0 0 0x0
GiveUp 0 0 0x0
PmtuAger 1 0 0x4CC1E758
DeadWait 0 0 0x0
Linger 0 0 0x0
ProcessQ 0 0 0x0

iss: 3440615046 snduna: 3440615202 sndnxt: 3440615202
irs: 3709999990 rcvnxt: 371000146

sndwnd: 16229 scale: 0 maxrcvwnd: 16384
rcvwnd: 16229 scale: 0 delrcvwnd: 155

```

Example: Disabling Path MTU Discovery for All Sessions

```

SRTT: 487 ms, RTTO: 3168 ms, RTV: 2681 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 1000 ms, ACK hold: 200 ms
uptime: 5612 ms, Sent idletime: 4588 ms, Receive idletime: 4388 ms
Status Flags: active open
Option Flags: nagle, path mtu capable
IP Precedence value : 6

Datagrams (max data segment is 1436 bytes):
Rcvd: 9 (out of order: 0), with data: 5, total data bytes: 155
Sent: 10 (retransmit: 0, fastretransmit: 0, partialack: 0, Second Congestion: 0),
  with data: 5, total data bytes: 155
Packets received in fast path: 0, fast processed: 0, slow path: 0
fast lock acquisition failures: 0, slow path: 0
TCP Semaphore 0x3CCE2D54 FREE

```

Example: Disabling Path MTU Discovery for All Sessions

The following example shows how to disable TCP path MTU discovery for all sessions. Use the **show ip bgp neighbors** command to verify that TCP path MTU discovery has been disabled.

```

router bgp 65535
no bgp transport path-mtu-discovery
bgp log-neighbor-changes
neighbor 3.3.3.3 remote-as 65535
neighbor 3.3.3.3 update-source Loopback0
!
address-family vpnv4
neighbor 3.3.3.3 activate
neighbor 3.3.3.3 send-community extended
exit-address-family

```

Router# **show ip bgp neighbors**

```

BGP neighbor is 3.3.3.3, remote AS 65535, internal link
BGP version 4, remote router ID 3.3.3.3
BGP state = Established, up for 00:00:13
Last read 00:00:12, last write 00:00:13, hold time is 180, keepalive interval is
 60 seconds
Neighbor sessions:
 1 active, is not multisession capable (disabled)
Neighbor capabilities:
Route refresh: advertised and received(new)
Four-octets ASN Capability: advertised and received
Address family IPv4 Unicast: advertised and received
Address family VPNv4 Unicast: advertised and received
Enhanced Refresh Capability: advertised and received
Multisession Capability:
Stateful switchover support enabled: NO for session 1
Message statistics:
InQ depth is 0
OutQ depth is 0

Sent Rcvd
Opens: 1 1
Notifications: 0 0
Updates: 2 2
Keepalives: 2 2
Route Refresh: 0 0
Total: 5 5

```

```
Default minimum time between advertisement runs is 0 seconds
For address family: IPv4 Unicast
Session: 3.3.3.3
BGP table version 1, neighbor version 1/0
Output queue size : 0
Index 8, Advertise bit 0
8 update-group member
Slow-peer detection is disabled
Slow-peer split-update-group dynamic is disabled
Interface associated: (none)

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
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

For address family: VPNv4 Unicast
Session: 3.3.3.3
BGP table version 1, neighbor version 1/0
Output queue size : 0
Index 7, Advertise bit 0
7 update-group member
Slow-peer detection is disabled
Slow-peer split-update-group dynamic is disabled
Interface associated: (none)

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
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

Address tracking is enabled, the RIB does have a route to 3.3.3.3
Connections established 8; dropped 7
Last reset 00:00:13, due to Peer closed the session of session 1
Transport(tcp) path-mtu-discovery is disabled
Graceful-Restart is disabled
SSO is disabled
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
Local host: 1.1.1.1, Local port: 179
Foreign host: 3.3.3.3, Foreign port: 20089
Connection tableid (VRF): 0
Maximum output segment queue size: 50

Enqueued packets for retransmit: 0, input: 0 mis-ordered: 0 (0 bytes)
Event Timers (current time is 0x4CBC07BC):
Timer Starts Wakeups Next
Retrans 4 0 0x0
TimeWait 0 0 0x0
AckHold 4 1 0x0
SendWnd 0 0 0x0
KeepAlive 0 0 0x0
GiveUp 0 0 0x0
PmtuAger 0 0 0x0
DeadWait 0 0 0x0
Linger 0 0 0x0
ProcessQ 0 0 0x0

iss: 2825924659 snduna: 2825924815 sndnxt: 2825924815
irs: 683352676 rcvnxt: 683352832

sndwnd: 16229 scale: 0 maxrcvwnd: 16384
rcvwnd: 16229 scale: 0 delrcvwnd: 155

SRTT: 413 ms, RTTO: 3205 ms, RTV: 2792 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 1000 ms, ACK hold: 200 ms
uptime: 13472 ms, Sent idletime: 12248 ms, Receive idletime: 12448 ms
Status Flags: passive open, gen tcbs
Option Flags: nagle
IP Precedence value : 6

Datagrams (max data segment is 536 bytes):
Rcvd: 10 (out of order: 0), with data: 5, total data bytes: 155

Sent: 9 (retransmit: 0, fastretransmit: 0, partialack: 0, Second Congestion: 0),
with data: 5, total data bytes: 155
Packets received in fast path: 0, fast processed: 0, slow path: 0
fast lock acquisition failures: 0, slow path: 0
TCP Semaphore 0x3CCE2D54 FREE

```

Example: Disabling Path MTU Discovery Per Session or Neighbor

The following example shows how to disable TCP path MTU discovery per session or neighbor. Use the `show ip bgp neighbors` command to verify that TCP path MTU discovery has been disabled.


```

router bgp 65535
  bgp log-neighbor-changes
  neighbor 3.3.3.3 remote-as 65535
  neighbor 3.3.3.3 transport path-mtu-discovery disable
  neighbor 3.3.3.3 update-source Loopback0
  !
  address-family vpnv4
  neighbor 3.3.3.3 activate
  neighbor 3.3.3.3 send-community extended
  exit-address-family

```

Router# **show ip bgp neighbors**

```

rspl-2014#sh ip bgp neighbors 3.3.3.3
BGP neighbor is 3.3.3.3, remote AS 65535, internal link
BGP version 4, remote router ID 3.3.3.3
BGP state = Established, up for 00:00:37
Last read 00:00:37, last write 00:00:36, hold time is 180, keepalive interval is
 60 seconds
Neighbor sessions:
 1 active, is not multiseession capable (disabled)
Neighbor capabilities:
Route refresh: advertised and received(new)
Four-octets ASN Capability: advertised and received
Address family IPv4 Unicast: advertised and received
Address family VPNv4 Unicast: advertised and received
Enhanced Refresh Capability: advertised and received
Multiseession Capability:
Stateful switchover support enabled: NO for session 1
Message statistics:
InQ depth is 0
OutQ depth is 0

Sent Rcvd
Opens: 1 1
Notifications: 0 0
Updates: 2 2
Keepalives: 2 2
Route Refresh: 0 0
Total: 5 5
Default minimum time between advertisement runs is 0 seconds

For address family: IPv4 Unicast
Session: 3.3.3.3
BGP table version 1, neighbor version 1/0
Output queue size : 0
Index 10, Advertise bit 0
10 update-group member
Slow-peer detection is disabled
Slow-peer split-update-group dynamic is disabled
Interface associated: (none)

Sent Rcvd
Prefix activity: ---- ----
Prefixes Current: 0 0
Prefixes Total: 0 0
Implicit Withdraw: 0 0
Explicit Withdraw: 0 0
Used as bestpath: n/a 0
Used as multipath: n/a 0

Outbound Inbound
Local Policy Denied Prefixes: -----

```

```

Total: 0 0
Number of NLRI in the update sent: max 0, min 0
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

For address family: VPNv4 Unicast
Session: 3.3.3.3
BGP table version 1, neighbor version 1/0
Output queue size : 0
Index 9, Advertise bit 0
9 update-group member
Slow-peer detection is disabled
Slow-peer split-update-group dynamic is disabled
Interface associated: (none)

Sent Rcvd
Prefix activity: ---- ----
Prefixes Current: 0 0
Prefixes Total: 0 0
Implicit Withdraw: 0 0
Explicit Withdraw: 0 0
Used as bestpath: n/a 0
Used as multipath: n/a 0

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

Address tracking is enabled, the RIB does have a route to 3.3.3.3
Connections established 10; dropped 9
Last reset 00:00:45, due to Peer closed the session of session 1
Transport(tcp) path-mtu-discovery is disabled
Graceful-Restart is disabled
SSO is disabled
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255
Local host: 1.1.1.1, Local port: 63272
Foreign host: 3.3.3.3, Foreign port: 179
Connection tableid (VRF): 0
Maximum output segment queue size: 50

```

```

Enqueued packets for retransmit: 0, input: 0 mis-ordered: 0 (0 bytes)
Event Timers (current time is 0x4D1A8DF4):
Timer Starts Wakeups Next
Retrans 5 0 0x0
TimeWait 0 0 0x0
AckHold 4 1 0x0
SendWnd 0 0 0x0
KeepAlive 0 0 0x0
GiveUp 0 0 0x0
PmtuAger 0 0 0x0
DeadWait 0 0 0x0
Linger 0 0 0x0
ProcessQ 0 0 0x0

iss: 2572241784 snduna: 2572241940 sndnxt: 2572241940
irs: 3439105921 rcvnxt: 3439106077

sndwnd: 16229 scale: 0 maxrcvwnd: 16384
rcvwnd: 16229 scale: 0 delrcvwnd: 155

SRTT: 487 ms, RTTO: 3168 ms, RTV: 2681 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 1000 ms, ACK hold: 200 ms
uptime: 37108 ms, Sent idletime: 36088 ms, Receive idletime: 35888 ms
Status Flags: active open
Option Flags: nagle
IP Precedence value : 6

Datagrams (max data segment is 536 bytes):
Rcvd: 9 (out of order: 0), with data: 5, total data bytes: 155
Sent: 10 (retransmit: 0, fastretransmit: 0, partialack: 0, Second Congestion: 0),
with data: 5, total data bytes: 155

Packets received in fast path: 0, fast processed: 0, slow path: 0
fast lock acquisition failures: 0, slow path: 0
TCP Semaphore 0x3CCE2CE4 FREE

```

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mcl/allreleasemcl/all-book.html

Standards and RFCs

Standard/RFC	Title
No specific Standards and RFCs are supported by the features in this document.	—

MIBs

MB	MIBs Link
—	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

Technical Assistance

Description	Link
<p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p>	http://www.cisco.com/cisco/web/support/index.html

Feature Information for Configuring BGP Support for TCP Path MTU Discovery per Session

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for Configuring BGP Support for TCP Path MTU Discovery per Session

Feature Name	Releases	Feature Information
Configuring BGP Support for TCP Path MTU Discovery per Session	Cisco IOS XE Release 3.14.0S	This feature was introduced on the Cisco ASR 920 Series Aggregation Services Router (ASR-920-12CZ-A, ASR-920-12CZ-D, ASR-920-4SZ-A, ASR-920-4SZ-D, ASR-920-10SZ-PD, ASR-920-24SZ-IM, ASR-920-24SZ-M, ASR-920-24TZ-M) .