# 配置BGP最大前缀功能

## 目录


[简介](#)

[先决条件](#)

- [要求](#)
- [使用的组件](#)
- [规则](#)

[背景信息](#)

[配置](#)

- [网络图](#)
- [配置Maximum-Prefix仅警告行为](#)
- [配置Maximum-Prefix以终止BGP会话行为](#)

[验证与故障排除](#)

- [Maximum-Prefix Warning-Only](#)
- [最大前缀终止BGP会话](#)

[相关信息](#)


## 简介

本文档介绍有关边界网关协议(BGP)最大前缀功能的配置和故障排除信息。

---

## 先决条件

### 要求

Cisco 建议您了解以下主题：

- 基本 IP 路由
- BGP配置

---

### 使用的组件

本文档中的信息不限于特定的软件和硬件版本，但示例基于运行Cisco IOS XE软件版本17.12.x的Cisco Catalyst 8500系列边缘平台。

---

本文档中的信息都是基于特定实验室环境中的设备编写的。本文档中使用的所有设备最初均采用原始（默认）配置。如果您的网络处于活动状态，请确保您了解所有命令的潜在影响。

---

### 规则

有关文档规则的详细信息，请参阅 Cisco 技术提示规则。

---

## 背景信息

本文档提供有关BGP最大前缀功能的配置和故障排除信息。此功能允许您控制可以从邻居接收多少前缀。默认情况下，当从对等体收到的前缀数超过配置的最大前缀限制时，使用此功能可以使路由器关闭此对等体。这通常用于外部BGP对等体，但可以应用于内部BGP对等体。

---

当在远程对等站点的出站策略发生变化时，路由器开始接收超过路由器内存可采用的路由时，最大前缀功能非常有用。如果路由器还执行关键路由功能，则收到的BGP前缀意外增加可能会占用系统资源并影响内部网络连接。使用neighbor <neighbor-ip> maximum-prefix命令，可以保护路由器免受此情况的影响。

---

如果您打算使用此功能，请考虑以下要点：

- 确定远程 BGP 对等路由器通常发送的路由数。
- 将maximum-prefix限制设置为高于正常操作期间预期的前缀数。将警告阈值配置为最大前缀限制的百分比。
- 当接收的前缀计数超过配置的最大值时，选择所需的操作：
  - 终止BGP会话并要求手动重置。
  - 配置restart <restart-interval>以自动重新建立会话。
  - 配置warning-only以记录事件，但不终止会话。

---

---

注意：在最大前缀限制终止会话后，重新启动选项会自动尝试重新建立BGP会话。有关详

---

细的配置信息，请参阅达到最大前缀限制后重新启动邻居会话。

# 配置

本部分提供有关如何配置本文档所述功能的信息。

用于配置BGP Maximum-Prefix功能的命令语法为：

<#root>

```
neighbor {ip-address | peer-group-name} maximum-prefix <maximum> [threshold] [restart

      ] [warning-only]
```

其中：

- maximum - 表示允许从邻居接收的最大前缀数。
- threshold — 指定路由器生成警告消息的配置的最大前缀限制的百分比。有效范围为1至100%。
  - 默认值为75%。
  - 例如，如果配置的maximum-value为20，阈值为60，则当从邻居获取的BGP路由数超过20(12)条路由的60%时，路由器会生成警告消息。
- restart-interval — 指定路由器尝试重新建立BGP会话后的间隔（以分钟为单位）。有效范围为1至65535分钟，方法如下。
- warning-only（可选）— 允许路由器在超过最大前缀限制时生成日志消息，而不是终止对等会话。

为了更好地说明用法，请考虑以下示例：

<#root>

**neighbor 10.1.1.1 maximum-prefix 3000**

*!--- Drops the peering to 10.1.1.1 when*
*!--- more than 3000 prefixes are received.*

**neighbor 10.1.1.1 maximum-prefix 3000 warning-only**

*!--- Logs a warning message when the peer sends*
*!--- more than 3000 prefixes.*

**neighbor 10.1.1.1 maximum-prefix 3000 50**

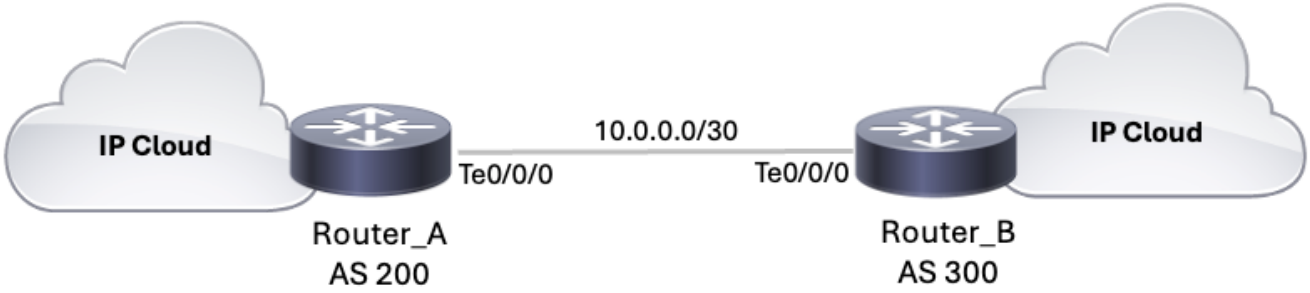
*!--- Logs a warning message at 1500 and drops the*
*!--- peering when over 3000 prefixes are sent.*

**neighbor 10.1.1.1 maximum-prefix 3000 50 warning-only**

*!--- Initially warns at 1500 and re-warns*
*!--- (different message) at 3000 prefixes received.*
*!--- However, the BGP Peer is not disconnected.*

---

## 网络图



自治系统200中的Router_A通过接口TenGigabitEthernet0/0/0直接连接到自治系统300中的Router_B。Router_A使用10.0.0.1/30,Router_B使用10.0.0.2/30。路由器通过此链路建立单跳外部边界网关协议(eBGP)会话。

## 配置Maximum-Prefix仅警告行为

在仅最大前缀警告配置中，按如下方式配置 Router_B：当从 Router_A 接收的前缀数超出所设置的阈值时，Router_B 仅记录一条警告消息。

下表显示了这两个路由器的配置。注意使用neighbor命令配置的warning-only关键字的存在：

| Router_A | Router_B |
| --- | --- |
| hostname Router_A<br>!<br>interface TenGigabitEthernet0/0/0<br>ip address 10.0.0.1 255.255.255.252<br>no negotiation auto<br>cdp enable<br>!<br>router bgp 200<br>bgp router-id 10.0.0.1<br>bgp log-neighbor-changes<br>neighbor 10.0.0.2 remote-as 300 | <#root><br><br>hostname Router_B<br>!<br>interface TenGigabitEthernet0/0/0<br>ip address 10.0.0.2 255.255.255.252<br>no negotiation auto<br>cdp enable<br>!<br>router bgp 300<br>bgp router-id 10.0.0.2<br>bgp log-neighbor-changes<br>neighbor 10.0.0.1 remote-as 200<br>**neighbor 10.0.0.1 maximum-prefix 10 80 warning-only** |

注意：在本示例中，当从邻居10.0.0.1接收的BGP前缀数量超过八个时，maximum-prefix命令会生成警告。

本文档的“验证和故障排除”部分中的show和debug命令输出报告当从Router_A接收的前缀数超过设定的阈值时，Router_B上会发生什么情况。

### 配置Maximum-Prefix以终止BGP会话行为

在本示例中，当received-prefix计数超过警告阈值时，Router_B会生成警告。当received-prefix计数超过最大前缀限制时，Router_B终止BGP会话。未配置warning-only关键字。当从邻居接收的前缀数超过10时，maximum-prefix命令终止BGP会话：

| Router_A | Router_B |
| --- | --- |
| hostname Router_A<br>!<br>interface TenGigabitEthernet0/0/0<br>ip address 10.0.0.1 255.255.255.252<br>no negotiation auto<br>cdp enable<br>!<br>router bgp 200<br>bgp router-id 10.0.0.1<br>bgp log-neighbor-changes<br>neighbor 10.0.0.2 remote-as 300 | <#root><br><br>hostname Router_B<br>!<br>interface TenGigabitEthernet0/0/0<br>ip address 10.0.0.2 255.255.255.252<br>no negotiation auto<br>cdp enable<br>!<br>router bgp 300<br>bgp router-id 10.0.0.2<br>bgp log-neighbor-changes<br>neighbor 10.0.0.1 remote-as 200<br>**neighbor 10.0.0.1 maximum-prefix 10 80** |

注意：在本示例中，当BGP从邻居获知的路由超过10时，maximum-prefix命令会强制断开邻居会话。

验证和故障排除部分中的show和debug命令输出报告当从Router_A接收的前缀数量超过设置的阈值时Router_B上发生的情况。

## 验证与故障排除

本部分所提供的信息可用于确认您的配置是否正常工作。本文档中使用的功能的命令语法和默认值可在BGP Command Page中找到。

注意：使用debug命令之前，请参阅了解有关Debug命令的重要信息。

- show ip bgp neighbor — 显示BGP邻居状态和前缀限制信息
- show ip bgp summary — 显示所有BGP连接的状态
- debug ip bgp updates in — 显示与BGP更新相关的信息

---

## Maximum-Prefix Warning-Only

请注意以下编号：

- 配置的最大前缀限制:10（10个前缀）
- 警告阈值：80%（八个前缀）

---



注意：省略用于测试前缀的确切路由生成和BGP通告配置。Router_A可以通过network语句或重新分发来生成前缀，或者从其他BGP邻居学习前缀并将其通告给Router_B。

---

只要接收的前缀数未超过设定的阈值，就不会记录任何消息。一旦从邻居10.0.0.1获知的BGP路由数量超过八个前缀的阈值限制，Router_B就会记录此消息。

---

当发送 9 个前缀时，系统将模拟此情况：

```
%BGP-4-MAXPFX: No. of prefix received from 10.0.0.1 (afi 0) reaches 9, max 10
```

如果情况恶化并超过Maximum-Prefix number set 10，路由器将记录此消息。当发送更多前缀时，将模拟此情况：

```
%BGP-3-MAXPFXEXCEED: No. of prefix received from 10.0.0.1 (afi 0): 11 exceed limit 10
```

<#root>

Router_B#

**show ip bgp neighbor 10.0.0.1**

```
BGP neighbor is 10.0.0.1,  remote AS 200, external link
  BGP version 4, remote router ID 10.0.0.1
  BGP state = Established, up for 00:17:22
  Last read 00:00:25, last write 00:00:22, hold time is 180, keepalive interval is 60 seconds
  Last update received: 00:04:04
  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
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

                         Sent       Rcvd
    Opens:                  1          1
    Notifications:          0          0
    Updates:                1          2
    Keepalives:            20         19
    Route Refresh:          0          0
    Total:                 22         22
  Do log neighbor state changes (via global configuration)
  Default minimum time between advertisement runs is 30 seconds

 For address family: IPv4 Unicast
  Session: 10.0.0.1
  BGP table version 12, neighbor version 12/0
  Output queue size : 0
  Index 1, Advertise bit 0
  1 update-group member
  Slow-peer detection is disabled
  Slow-peer split-update-group dynamic is disabled
                               Sent       Rcvd
  Prefix activity:             ----       ----


Prefixes Current: 0 11 (Consumes 1496 bytes)


    Prefixes Total:                0         11
    Implicit Withdraw:             0          0
    Explicit Withdraw:             0          0
    Used as bestpath:            n/a         11
    Used as multipath:           n/a          0
    Used as secondary:           n/a          0

                                  Outbound    Inbound
  Local Policy Denied Prefixes:   --------    -------
    Bestpath from this peer:            11        n/a
    Total:                              11          0


Maximum prefixes allowed 10 (warning-only) Threshold for warning message 80%


  Number of NLRIs in the update sent: max 0, min 0
  Current session network count peaked at 11 entries at 20:05:46 Aug 19 2026 UTC (00:04:05.075 ago)
```

**Highest network count observed at 11 entries at 20:05:46 Aug 19 2026 UTC (00:04:05.075 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

  Address tracking is enabled, the RIB does have a route to 10.0.0.1
  Route to peer address reachability Up: 1; Down: 0
    Last notification 00:17:27
  Connections established 1; dropped 0
  Last reset never
  External BGP neighbor configured for connected checks (single-hop no-disable-connected-check)
  Interface associated: TenGigabitEthernet0/0/0 (peering address in same link)
  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, Mininum incoming TTL 0, Outgoing TTL 1
Local host: 10.0.0.2, Local port: 179
Foreign host: 10.0.0.1, Foreign port: 48663
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 0x386BB365):
Timer          Starts    Wakeups            Next
Retrans            21          0             0x0
TimeWait            0          0             0x0
AckHold            22         21             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: 3438119007  snduna: 3438119468  sndnxt: 3438119468
irs: 2705427639  rcvnxt: 2705428185

sndwnd:  15924  scale:      0  maxrcvwnd:  16384
rcvwnd:  15839  scale:      0  delrcvwnd:    545

SRTT: 939 ms, RTTO: 1411 ms, RTV: 472 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 1000 ms, ACK hold: 120 ms
uptime: 1042758 ms, Sent idletime: 22095 ms, Receive idletime: 21895 ms
Status Flags: passive open, gen tcbs
Option Flags: nagle, path mtu capable
IP Precedence value : 6
Window update Optimisation : Enabled
ACK Optimisation : Dynamic ACK Tuning Enabled

Datagrams (max data segment is 1460 bytes):
```

```
Peer MSS:       1460
Rcvd: 44 (out of order: 0), with data: 22, total data bytes: 545
Sent: 45 (retransmit: 0, fastretransmit: 0, partialack: 0, Second Congestion: 0), with data: 22, total

 Packets received in fast path: 0, fast processed: 0, slow path: 0
 fast lock acquisition failures: 0, slow path: 0
TCP Semaphore      0x746BB5E1C7B0  FREE
```

<#root>

Router_B#

**show ip bgp summary**

```
BGP router identifier 10.0.0.2, local AS number 300
BGP table version is 12, main routing table version 12
11 network entries using 2728 bytes of memory
11 path entries using 1496 bytes of memory
1/1 BGP path/bestpath attribute entries using 296 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 4544 total bytes of memory
BGP activity 11/0 prefixes, 11/0 paths, scan interval 60 secs
11 networks peaked at 20:05:46 Aug 19 2026 UTC (00:08:51.371 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.0.1        4          200      27      27       12    0    0 00:22:09
```

**11**

---

Router_A向Router_B通告11个前缀。当received-prefix计数达到9时，Router_B生成警告，当计数达到11时，生成最大前缀超出消息。由于配置了warning-only，因此BGP会话保持已建立。

---

---

⚠ 警告：中的命令debug ip bgp updates可生成大量输出并影响设备性能。仅在受控故障排除窗口期间运行此命令，监控系统资源，使用命令过滤器，并在数据收集后禁用调试。

---

---

<#root>

Router_B#

**debug ip bgp updates in**

```
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd UPDATE w/ attr: nexthop 10.0.0.1, origin ?, metric 0, merge
```

```
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.0.0.0/30
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.10.1.0/30
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.1.1.1/32
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.2.2.2/32
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.3.3.3/32
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.4.4.4/32
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.5.5.5/32
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.6.6.6/32
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.7.7.7/32

*Aug 19 20:34:50.019: %BGP-4-MAXPFX: Number of prefixes received from 10.0.0.1 (afi 0) reaches 9, max 10

*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.9.9.9/32
*Aug 19 20:34:50.019: BGP(0): 10.0.0.1 rcvd 10.8.8.8/32

*Aug 19 20:34:50.019: %BGP-3-MAXPFXEXCEED: Number of prefixes received from 10.0.0.1 (afi 0): 11 exceeds

*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.0.0.0/30 -> 10.0.0.1(global)
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.1.1.1/32 -> 10.0.0.1(global)
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.2.2.2/32 -> 10.0.0.1(global)
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.3.3.3/32 -> 10.0.0.1(global)
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.4.4.4/32 -> 10.0.0.1(global)
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.5.5.5/32 -> 10.0.0.1(global)
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.6.6.6/32 -> 10.0.0.1(global)
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.7.7.7/32 -> 10.0.0.1(global)
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.8.8.8/32 -> 10.0.0.1(global)
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.9.9.9/32 -> 10.0.0.1(global)
*Aug 19 20:34:51.040: BGP(0): Revise route installing 1 of 1 routes for 10.10.1.0/30 -> 10.0.0.1(global)
```

---

在前面的示例中，即使相邻路由器发送的前缀数超过策略允许的数量，BGP邻居关系也会得到维护。Router_B会记录警告和maximum-prefix-exceeded消息，但会保持BGP会话的建立。Router_B继续接受来自邻居的前缀，因为配置了仅警告。

---

## 最大前缀终止BGP会话

此情况所需的初始条件是BGP邻居已启动且正在运行，并且具有Router_A向Router_B发送的六个前缀。如示例所示，当Router_A通告更多前缀（例如9）时，命令的输出准确反映了当Router_B配置为记录警告消息时已经看到的内容。

---

在Router_A通告第11个前缀后，接收的前缀计数超过了配置的限制10。Router_B发送已达到的最大前缀数通知并终止BGP会话。

<#root>

Router_B#

**debug ip bgp updates in**

*Aug 19 20:45:48.779: BGP(0): 10.0.0.1 rcvd UPDATE w/ attr: nexthop 10.0.0.1, origin ?, metric 0, merge
*Aug 19 20:45:48.779: BGP(0): 10.0.0.1 rcvd 10.7.7.7/32

***Aug 19 20:45:48.779: %BGP-4-MAXPFX: Number of prefixes received from 10.0.0.1 (afi 0) reaches 10, max 1**

*Aug 19 20:45:48.779: BGP(0): 10.0.0.1 rcvd 10.9.9.9/32

***Aug 19 20:45:48.779: %BGP-3-MAXPFXEXCEED: Number of prefixes received from 10.0.0.1 (afi 0): 11 exceeds**

***Aug 19 20:45:48.780: %BGP-5-NBR_RESET: Neighbor 10.0.0.1 reset (Peer over prefix limit)**

*Aug 19 20:45:48.780: BGP(0): no valid path for 10.0.0.0/30
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.1.1.1/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.2.2.2/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.3.3.3/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.4.4.4/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.5.5.5/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.6.6.6/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.7.7.7/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.8.8.8/32
*Aug 19 20:45:48.780: BGP(0): no valid path for 10.10.1.0/30

***Aug 19 20:45:48.780: %BGP-5-ADJCHANGE: neighbor 10.0.0.1 Down Peer over prefix limit *Aug 19 20:45:48.7**

*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.0.0.0/30
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.1.1.1/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.2.2.2/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.3.3.3/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.4.4.4/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.5.5.5/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.6.6.6/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.7.7.7/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.8.8.8/32
*Aug 19 20:45:48.780: BGP: topo global:IPv4 Unicast:base Remove_fwdroute for 10.10.1.0/30

<#root>

Router_B#

**show ip bgp summary**

BGP router identifier 10.0.0.2, local AS number 300
BGP table version is 25, main routing table version 25
17 networks peaked at 20:33:04 Aug 19 2026 UTC (00:13:00.072 ago)

| Neighbor | V | AS | MsgRcvd | MsgSent | TblVer | InQ | OutQ | Up/Down | State/PfxRcd |
|---|---|---|---|---|---|---|---|---|---|
| 10.0.0.1 | 4 | 200 | 0 | 0 | 1 | 0 | 0 | 00:00:15 | |

**Idle (PfxCt)**

注意：在此方案中，必须使用clear ip bgp <neighbor-ip>命令还原对等体会话。在重置会话之前，请减少对等体通告的前缀的数量或在容量验证后调整配置的maximum-prefix限制。此命令重置BGP会话并暂时删除从对等体获取的路由。

<#root>

Router_B#

**show ip bgp neighbors 10.0.0.1**

```
BGP neighbor is 10.0.0.1,  remote AS 200, external link
  BGP version 4, remote router ID 10.0.0.1
  BGP state = Idle, down for 00:00:39
  Last update received: n/a
  Neighbor sessions:
    0 active, is not multisession capable (disabled)
    Stateful switchover support enabled: NO for session 0
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

                         Sent       Rcvd
    Opens:                  0          1
    Notifications:          1          0
    Updates:                0          0
    Keepalives:             0          0
    Route Refresh:          0          0
    Total:                  1          1
  Do log neighbor state changes (via global configuration)
  Default minimum time between advertisement runs is 30 seconds

 For address family: IPv4 Unicast
  BGP table version 25, neighbor version 1/25
  Output queue size : 0
  Index 0, Advertise bit 0
  Address family not supported notification sent
  Slow-peer detection is disabled
  Slow-peer split-update-group dynamic is disabled
  Peer had exceeded the max. no. of prefixes configured.
  Maximum prefixes allowed 10
  Threshold for warning message 80%


 Reduce the no. of prefix and clear ip bgp 10.0.0.1 to restore peering


  Number of NLRIs in the update sent: max 0, min 0
```

```
Highest network count observed at 12 entries at 20:32:03 Aug 19 2026 UTC (00:14:25.012 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

  Address tracking is enabled, the RIB does have a route to 10.0.0.1
  Route to peer address reachability Up: 1; Down: 0
    Last notification 00:54:04
  Connections established 3; dropped 3
  Last reset 00:00:39, due to BGP protocol initialization
  External BGP neighbor configured for connected checks (single-hop no-disable-connected-check)
  Interface associated: TenGigabitEthernet0/0/0 (peering address in same link)
  Transport(tcp) path-mtu-discovery is enabled
  Graceful-Restart is disabled
  SSO is disabled


No active TCP connection
```

---

# 相关信息

- BGP 案例分析
- 排除 BGP 故障
- 思科技术支持和下载

关于此翻译
思科采用人工翻译与机器翻译相结合的方式将此文档翻译成不同语言，希望全球的用户都能通过各自的语言得到支持性的内容。

请注意：即使是最好的机器翻译，其准确度也不及专业翻译人员的水平。

Cisco Systems, Inc. 对于翻译的准确性不承担任何责任，并建议您总是参考英文原始文档（已提供链接）。