

配置 BGP 最大前缀特性

目录

[简介](#)

[先决条件](#)

[要求](#)

[使用的组件](#)

[规则](#)

[配置](#)

[网络图](#)

[配置](#)

[阈值超出阈值集时为 Warning-Only 配置 Maximum-Prefix](#)

[阈值超出阈值集时配置最大前缀以降低邻接关系](#)

[验证与故障排除](#)

[Maximum-Prefix Warning-Only](#)

[阈值超出阈值集时配置最大前缀以降低会话](#)

[相关信息](#)

简介

本文档提供有关边界网关协议 (BGP) 最大前缀功能的配置和故障排除信息。

使用 BGP 最大前缀功能，您可以控制可从邻居接收的前缀数目。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。此功能通常用于外部 BGP 对等体，但也可应用于内部 BGP 对等体。

当远程对等站点上的出站策略发生更改时，如果路由器收到的路由开始超出此路由器的存储器可容纳的路由，最大前缀功能则非常有用。如果上述同一路由器与 BGP 对等，并在网络内执行重要路由功能，此开销可能会导致不良后果。BGP 问题可能会中断内部网络连接。使用 `neighbor maximum-prefix` 命令，可以保护路由器以避免出现此种情况。

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。
- 设置阈值，使其略高于正常操作时预期收到的 BGP 前缀数。
- 确定当远程 BGP 对等体发送的前缀数超出预期前缀数时所采取的操作。可采取的操作可以是：
：在使用 `clear ip bgp x.x.x.x` 命令之前关闭会话并降低 BGP 邻居关系，或者仅记录警告消息。

注意： Cisco IOS® 软件版本 12.0(22)S^和12.2(15)T 中对此功能进行了增强。使用此增强功能，用户可以自动重建因超出配置的最大前缀限制而被关闭的对等会话。启用此功能后，网络操作员无需进行任何干预。有关进一步的信息，请参阅 [BGP 在超出最大前缀限制之后重新启动会话](#)。

先决条件

要求

思科建议本文档的读者基本了解[配置BGP网络](#)。

使用的组件

本文档中的信息基于以下软件和硬件版本：

Cisco IOS® 软件版本 12.2 (27) 上的 Cisco 2500 系列路由器

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。如果您使用的是真实网络，请确保您已经了解所有命令的潜在影响。

请访问 [Cisco Feature Navigator \(仅限注册用户 \)](#)，以便确定可在哪些 Cisco IOS 软件版本中使用此功能。

规则

有关文档规则的详细信息，请参阅 [Cisco 技术提示规则](#)。

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 最大前缀功能的命令语法是：

```
neighbor {ip-address | peer-group-name} maximum-prefix maximum [threshold] [restart restart-interval] [warning-only]
```

其中：


- **maximum** - 表示允许从邻居接收的最大前缀数。
- **threshold** - 一个可选整数值，用于指定配置的 **maximum-value** 所对应的百分比。如果超过此值，路由器将开始生成警告消息。其范围介于 1% 到 100% 之间，默认值为 75%。例如，如果配置的 **maximum-value** 为 20，并且阈值为 60，当来自邻居的 BGP 路由数超出 20 条的 60%（即 12 条路由）时，路由器将生成警告消息。**restart-interval** - 重建对等会话的可选时间间隔（以分钟为单位）。其范围介于 1 到 65535 分钟之间。**warning-only** -（可选）允许路由器在超出最大前缀限制时生成日志消息，而不是终止对等会话。

为了更好地演示此用法，请考虑以下示例：

```
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

本文档使用以下网络设置：



配置

本文档使用以下配置：

- [阈值超出阈值集时为 Warning-Only 配置 Maximum-Prefix](#)
- [阈值超出阈值集时配置最大前缀以降低邻接关系](#)

[阈值超出阈值集时为 Warning-Only 配置 Maximum-Prefix](#)

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。请注意，**neighbor** 命令配置了 *warning-only* 关键字。

Router_A	Router_B
<pre>hostname Router_A ! interface Loopback0 ip address 10.0.0.1 255.255.255.255 ! interface Serial0 ip address 192.168.1.1 255.255.255.252 ! interface Serial1 ip unnumbered Loopback0 ! router bgp 200 no synchronization bgp router-id 10.0.0.1 bgp log-neighbor- changes neighbor 192.168.1.2 local-as 100 neighbor 10.0.0.2 remote-as 300 neighbor 10.0.0.2 ebgp- multihop 2 neighbor 10.0.0.2 update-source Loopback0 neighbor 10.0.0.2 version 4 no auto-summary ! ip route 10.0.0.2 255.255.255.252 Serial1</pre>	<pre>hostname Router_B ! interface Loopback0 ip address 10.0.0.2 255.255.255.252 ! interface Ethernet0 ip address 10.0.1.1 255.255.255.0 ! interface Serial0 ip unnumbered Loopback0 ! router bgp 300 no synchronization bgp router-id 10.0.0.2 bgp log-neighbor-changes neighbor 10.0.0.1 remote-as 200 neighbor 10.0.0.1 ebgp- multihop 2 neighbor 10.0.0.1 update- source Loopback0 neighbor 10.0.0.1 version 4 neighbor 10.0.0.1 maximum- prefix 10 80 warning-only <i>!--- Enables warning message logging when the number !--- of BGP routes learned from neighbor !--- 10.0.0.1 exceeds eight.</i> no auto-summary ! ip route 10.0.0.1 255.255.255.252 Serial0</pre>

本文档的[验证与故障排除部分](#)中的 `show` 和 `debug` 命令输出报告了当从 Router_A 收到的前缀数超出所设置的阈值时，Router_B 上所发生的实际情况。

阈值超出阈值集时配置最大前缀以降低邻接关系

在配置为降低邻居关系的最大前缀配置中，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 应生成警告消息。同时，还对 Router_B 进行如下配置：当超出最大前缀限制时，Router_B 将减少 BGP 邻居。下表显示了这两个路由器的配置。请注意，`neighbor` 命令未设置 `warning-only` 关键字。

Router_A	Router_B
<pre>hostname Router_A ! interface Loopback0 ip address 10.0.0.1 255.255.255.255 ! interface Serial0 ip address 192.168.1.1 255.255.255.252 ! interface Serial1 ip unnumbered Loopback0 ! router bgp 200 no synchronization bgp router-id 10.0.0.1 bgp log-neighbor- changes neighbor 192.168.1.2 local-as 100 neighbor 10.0.0.2 remote-as 300 neighbor 10.0.0.2 ebgp- multihop 2 neighbor 10.0.0.2 update-source Loopback0 neighbor 10.0.0.2 version 4 no auto-summary ! ip route 10.0.0.2 255.255.255.252 Serial1</pre>	<pre>hostname Router_B ! interface Loopback0 ip address 10.0.0.2 255.255.255.252 ! interface Ethernet0 ip address 10.0.1.1 255.255.255.0 ! interface Serial0 ip unnumbered Loopback0 ! router bgp 300 no synchronization bgp router-id 10.0.0.2 bgp log-neighbor-changes neighbor 10.0.0.1 remote-as 200 neighbor 10.0.0.1 ebgp- multihop 2 neighbor 10.0.0.1 update- source Loopback0 neighbor 10.0.0.1 version 4 neighbor 10.0.0.1 maximum- prefix 10 80 <i>!--- This forces the neighbor session to tear down !--- when the BGP learned routes from !- -- the neighbor exceeds 10. no</i> auto-summary ! ip route 10.0.0.1 255.255.255.252 Serial0</pre>

[验证与故障排除部分](#)中的 `show` 和 `debug` 命令输出报告了当从 Router_A 收到的前缀数超出所设置的阈值时，Router_B 上所发生的实际情况。

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。

BGP命令页提供本文档中使用的命令语法和功能的[默认值](#)。

注意：在使用`debug`命令之前，请参[阅](#)有关Debug命令的重要信息。

- [show ip bgp neighbor](#) - 显示 BGP 邻居状态。
- [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：

- 允许的最大前缀数：10
- 警告阈值：80% (8)

只要收到的前缀数不超过所设置的阈值（即 8），则不会记录任何消息。只要来自邻居 10.0.0.1 的 BGP 路由数超出 8 条的阈值限值，Router_B 则会记录此消息。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
```

如果情况变得更糟，并且超过设置的最大前缀数 10，路由器会记录此消息。当发送 12 个前缀时，系统将模拟此情况：

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

激活 [debug ip bgp updates in](#) 之后，您可以更仔细地查看所发生的情况。然而，不要在具有数千个前缀的真实环境中使用此命令。所描述的情况是：Router_B 已建立对等体。Router_A 向 Router_B 通告 6 个前缀。现在，对等体 Router_A 通告另外三个前缀。

```
Router_B# debug ip bgp updates in
```

```
*Mar 12 07:31:18.944: BGP(0): 10.0.0.1 rcvd UPDATE w/ attr: nexthop 10.0.0.1, or
igin i, metric 0, path 200
*Mar 12 07:31:18.948: BGP(0): 10.0.0.1 rcvd 10.0.1.0/24...duplicate ignored
*Mar 12 07:31:18.952: BGP(0): 10.0.0.1 rcvd 10.0.2.0/24...duplicate ignored
*Mar 12 07:31:18.960: BGP(0): 10.0.0.1 rcvd 10.0.3.0/24...duplicate ignored
*Mar 12 07:32:20.224: BGP(0): 10.0.0.1 rcvd 10.0.4.0/24...duplicate ignored
*Mar 12 07:32:20.228: BGP(0): 10.0.0.1 rcvd 10.0.5.0/24...duplicate ignored
*Mar 12 07:32:20.232: BGP(0): 10.0.0.1 rcvd 10.0.6.0/24...duplicate ignored
*Mar 12 07:34:19.768: BGP(0): 10.0.0.1 rcvd 10.0.7.0/24
*Mar 12 07:34:19.772: BGP(0): 10.0.0.1 rcvd 10.0.8.0/24
*Mar 12 07:34:19.780: BGP(0): 10.0.0.1 rcvd 10.0.9.0/24
*Mar 12 07:34:19.780:
```

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

```
*Mar 12 07:34:19.792: BGP(0): Revise route installing 1 of 1 route for 10.0.7.0/
24 -> 10.0.0.1 to main IP table
*Mar 12 07:34:19.796: BGP(0): Revise route installing 1 of 1 route for 10.0.8.0/
24 -> 10.0.0.1 to main IP table
*Mar 12 07:34:19.804: BGP(0): Revise route installing 1 of 1 route for 10.0.9.0/
24 -> 10.0.0.1 to main IP table
```

```
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 00:13:22
  Last read 00:00:21, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
```

Route refresh: advertised and received(old & new)
Address family IPv4 Unicast: advertised and received
IPv4 MPLS Label capability:
Received 930 messages, 0 notifications, 0 in queue
Sent 919 messages, 1 notifications, 0 in queue
Default minimum time between advertisement runs is 30 seconds

For address family: IPv4 Unicast

BGP table version 30, neighbor version 30

Index 1, Offset 0, Mask 0x2

Route refresh request: received 0, sent 0

9 accepted prefixes consume 432 bytes

Prefix advertised 0, suppressed 0, withdrawn 0, **maximum limit 10 (warning-only**

)

Threshold for warning message 80%

Connections established 2; dropped 1

Last reset 00:29:13, due to BGP Notification sent, update malformed

Message received that caused BGP to send a Notification:

FFFFFFFF FFFFFFFFFF FFFFFFFFFF FFFFFFFFFF

003C0200 00001940 01010040 02040201

00C84003 040A0000 01800404 00000000

180A000A 180A000B 180A000C

External BGP neighbor can be up to 2 hops away.

Connection state is ESTAB, I/O status: 1, unread input bytes: 0

Local host: 10.0.0.2, Local port: 15668

Foreign host: 10.0.0.1, Foreign port: 179

Enqueued packets for retransmit: 0, input: 0 mis-ordered: 0 (0 bytes)

Event Timers (current time is 0x3A46EB54):

Timer	Starts	Wakeups	Next
Retrans	18	0	0x0
TimeWait	0	0	0x0
AckHold	22	9	0x0
SendWnd	0	0	0x0
KeepAlive	0	0	0x0
GiveUp	0	0	0x0
PmtuAger	0	0	0x0
DeadWait	0	0	0x0

iss: 2047376434 snduna: 2047376784 sndnxt: 2047376784 sndwnd: 16035

irs: 821061364 rcvnxt: 821062116 rcvwnd: 16188 delrcvwnd: 196

SRTT: 279 ms, RTTO: 500 ms, RTV: 221 ms, KRTT: 0 ms

minRTT: 24 ms, maxRTT: 384 ms, ACK hold: 200 ms

Flags: higher precedence, nagle

Datagrams (max data segment is 536 bytes):

Rcvd: 33 (out of order: 0), with data: 22, total data bytes: 751

Sent: 29 (retransmit: 0, fastretransmit: 0), with data: 17, total data bytes: 349

Router_B#show ip bgp summary

BGP router identifier 10.0.0.2, local AS number 300

BGP table version is 30, main routing table version 30

9 network entries and 9 paths using 1341 bytes of memory

1 BGP path attribute entries using 60 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 activity 36/101 prefixes, 36/27 paths, scan interval 60 secs

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.0.1	4	200	932	921	30	0	0	00:15:08	9

假定情况变得更糟，并且 Router_A 另外发送三个前缀，这会将总数增加至 12。

Router_B# **debug ip bgp updates in**

```
*Mar 12 07:39:21.192: BGP(0): 10.0.0.1 rcvd UPDATE w/ attr: nexthop 10.0.0.1, origin i, metric 0, path 200
*Mar 12 07:39:21.196: BGP(0): 10.0.0.1 rcvd 10.0.10.0/24
*Mar 12 07:39:21.200: %BGP-4-MAXPFX: No. of prefix received from 10.0.0.1 (afi 0) reaches 10, max 10
*Mar 12 07:39:21.208: BGP(0): 10.0.0.1 rcvd 10.0.11.0/24
*Mar 12 07:39:21.212: %BGP-3-MAXPFXEXCEED: No. of prefix received from 10.0.0.1 (afi 0): 11 exceed limit 10
*Mar 12 07:39:21.216: BGP(0): 10.0.0.1 rcvd 10.0.12.0/24
*Mar 12 07:39:21.228: BGP(0): Revise route installing 1 of 1 route for 10.0.10.0/24 -> 10.0.0.1 to main IP table
*Mar 12 07:39:21.236: BGP(0): Revise route installing 1 of 1 route for 10.0.11.0/24 -> 10.0.0.1 to main IP table
*Mar 12 07:39:21.240: BGP(0): Revise route installing 1 of 1 route for 10.0.12.0/24 -> 10.0.0.1 to main IP table
```

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 = Established, up for 00:19:56
  Last read 00:00:56, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    Address family IPv4 Unicast: advertised and received
    IPv4 MPLS Label capability:
  Received 937 messages, 0 notifications, 0 in queue
  Sent 925 messages, 1 notifications, 0 in queue
  Default minimum time between advertisement runs is 30 seconds
```

For address family: IPv4 Unicast

```
BGP table version 33, neighbor version 33
Index 1, Offset 0, Mask 0x2
Route refresh request: received 0, sent 0
12 accepted prefixes consume 576 bytes
Prefix advertised 0, suppressed 0, withdrawn 0, maximum limit 10 (warning-only)
Threshold for warning message 80%
```

```
Connections established 2; dropped 1
Last reset 00:35:47, due to BGP Notification sent, update malformed
Message received that caused BGP to send a Notification:
  FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF
  003C0200 00001940 01010040 02040201
  00C84003 040A0000 01800404 00000000
  180A000A 180A000B 180A000C
  External BGP neighbor can be up to 2 hops away.
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
```

Local host: 10.0.0.2, Local port: 15668
Foreign host: 10.0.0.1, Foreign port: 179

Enqueued packets for retransmit: 0, input: 0 mis-ordered: 0 (0 bytes)

Event Timers (current time is 0x3A4CEA98):

Timer	Starts	Wakeups	Next
Retrans	24	0	0x0
TimeWait	0	0	0x0
AckHold	29	16	0x0
SendWnd	0	0	0x0
KeepAlive	0	0	0x0
GiveUp	0	0	0x0
PmtuAger	0	0	0x0
DeadWait	0	0	0x0

iss: 2047376434 snduna: 2047376898 sndnxt: 2047376898 sndwnd: 15921
irs: 821061364 rcvnxt: 821062290 rcvwnd: 16014 delrcvwnd: 370

SRTT: 290 ms, RTTO: 376 ms, RTV: 86 ms, KRRTT: 0 ms
minRTT: 24 ms, maxRTT: 384 ms, ACK hold: 200 ms
Flags: higher precedence, nagle

Datagrams (max data segment is 536 bytes):

Rcvd: 40 (out of order: 0), with data: 29, total data bytes: 925
Sent: 42 (retransmit: 0, fastretransmit: 0), with data: 23, total data bytes: 463

Router_B#**show ip bgp summary**

BGP router identifier 10.0.0.2, local AS number 300
BGP table version is 33, main routing table version 33
12 network entries and 12 paths using 1788 bytes of memory
1 BGP path attribute entries using 60 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 activity 39/101 prefixes, 39/27 paths, scan interval 60 secs

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.0.1	4	200	939	927	33	0	0	00:21:28	12

正如您从所示的示例中看到的一样，即使邻接路由器发送的前缀多于策略所允许的前缀，也会保持 BGP 邻居关系。其结果是，Router_B 仅记录一条警告消息，而不会采取任何其他行动。

阈值超出阈值集时配置最大前缀以降低会话

此种情况的必需初始条件是：启动并运行 BGP 邻居，并且 Router_A 向 Router_B 发送 6 个前缀。正如您在示例中看到的一样，当 Router_A 通告更多前缀（如 9 个）时，命令输出反映的内容与将 Router_B 配置为仅记录警告消息时所显示的内容完全一致。如果增加发送的前缀数，并使 Router_A 通告 12 个前缀，Router_B 将终止与 Router_A 的邻居关系。

Router_B# debug ip bgp updates in

```
*Mar 12 08:03:27.864: BGP(0): 10.0.0.1 rcvd UPDATE w/ attr: nexthop 10.0.0.1, or  
igin i, metric 0, path 200  
*Mar 12 08:03:27.868: BGP(0): 10.0.0.1 rcvd 10.0.1.0/24...duplicate ignored  
*Mar 12 08:03:27.876: BGP(0): 10.0.0.1 rcvd 10.0.2.0/24...duplicate ignored
```



```

*Mar 12 08:03:27.880: BGP(0): 10.0.0.1 rcvd 10.0.3.0/24...duplicate ignored
*Mar 12 08:03:27.884: BGP(0): 10.0.0.1 rcvd 10.0.4.0/24...duplicate ignored
*Mar 12 08:03:27.892: BGP(0): 10.0.0.1 rcvd 10.0.5.0/24...duplicate ignored
*Mar 12 08:03:27.896: BGP(0): 10.0.0.1 rcvd 10.0.6.0/24...duplicate ignored
*Mar 12 08:03:27.900: BGP(0): 10.0.0.1 rcvd 10.0.7.0/24
*Mar 12 08:03:27.908: BGP(0): 10.0.0.1 rcvd 10.0.8.0/24
*Mar 12 08:03:27.912: BGP(0): 10.0.0.1 rcvd 10.0.9.0/24
*Mar 12 08:03:27.916: %BGP-4-MAXPFX: No. of prefix received from 10.0.0.1 (afi 0
) reaches 9, max 10
*Mar 12 08:03:27.924: BGP(0): 10.0.0.1 rcvd 10.0.10.0/24
*Mar 12 08:03:27.932: BGP(0): 10.0.0.1 rcvd 10.0.11.0/24
*Mar 12 08:03:27.932: %BGP-3-MAXPFXEXCEED: No. of prefix received from 10.0.0.1
(afi 0): 11 exceed limit 10
*Mar 12 08:03:27.940: %BGP-5-ADJCHANGE: neighbor 10.0.0.1 Down BGP Notification
sent
*Mar 12 08:03:27.940: %BGP-3-NOTIFICATION: sent to neighbor 10.0.0.1 3/1 (update
malformed) 0 bytes  FFFF FFFF FFFF FFFF FFFF FFFF FFFF FFFF 0060 0200 0000 1940
0101 0040 0204 0201 00C8 4003 040A 0000 0180 0404 0000 0000 180A 0001 180A 0002
180A 0003 180A 0004 180A 0005 180A 0006 180A 0007 180A 0008 180A 0009 180A 000A
180A 000B 180A 000C
*Mar 12 08:03:28.024: BGP(0): Revise route installing 1 of 1 route for 10.0.7.0/
24 -> 10.0.0.1 to main IP table
*Mar 12 08:03:28.032: BGP(0): Revise route installing 1 of 1 route for 10.0.8.0/
24 -> 10.0.0.1 to main IP table
*Mar 12 08:03:28.036: BGP(0): Revise route installing 1 of 1 route for 10.0.9.0/
24 -> 10.0.0.1 to main IP table
*Mar 12 08:03:28.044: BGP(0): Revise route installing 1 of 1 route for 10.0.10.0
/24 -> 10.0.0.1 to main IP table
*Mar 12 08:03:28.148: BGP(0): no valid path for 10.0.1.0/24
*Mar 12 08:03:28.152: BGP(0): no valid path for 10.0.2.0/24
*Mar 12 08:03:28.156: BGP(0): no valid path for 10.0.3.0/24
*Mar 12 08:03:28.156: BGP(0): no valid path for 10.0.4.0/24
*Mar 12 08:03:28.160: BGP(0): no valid path for 10.0.5.0/24
*Mar 12 08:03:28.164: BGP(0): no valid path for 10.0.6.0/24
*Mar 12 08:03:28.168: BGP(0): no valid path for 10.0.7.0/24
*Mar 12 08:03:28.168: BGP(0): no valid path for 10.0.8.0/24
*Mar 12 08:03:28.172: BGP(0): no valid path for 10.0.9.0/24
*Mar 12 08:03:28.176: BGP(0): no valid path for 10.0.10.0/24
*Mar 12 08:03:28.184: BGP(0): nettable_walker 10.0.1.0/24 no best path
*Mar 12 08:03:28.188: BGP(0): nettable_walker 10.0.2.0/24 no best path
*Mar 12 08:03:28.192: BGP(0): nettable_walker 10.0.3.0/24 no best path
*Mar 12 08:03:28.196: BGP(0): nettable_walker 10.0.4.0/24 no best path
*Mar 12 08:03:28.200: BGP(0): nettable_walker 10.0.5.0/24 no best path
*Mar 12 08:03:28.204: BGP(0): nettable_walker 10.0.6.0/24 no best path
*Mar 12 08:03:28.208: BGP(0): nettable_walker 10.0.7.0/24 no best path
*Mar 12 08:03:28.212: BGP(0): nettable_walker 10.0.8.0/24 no best path
*Mar 12 08:03:28.212: BGP(0): nettable_walker 10.0.9.0/24 no best path
*Mar 12 08:03:28.216: BGP(0): nettable_walker 10.0.10.0/24 no best path

```

```
Router_B# show ip bgp summary
```

```
BGP router identifier 10.0.0.2, local AS number 300
BGP table version is 87, main routing table version 87
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.0.1	4	200	965	948	0	0	0	00:02:24	Idle (PfxCt)

```
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 0.0.0.0
  BGP state = Idle
  Last read 00:02:43, hold time is 180, keepalive interval is 60 seconds
  Received 965 messages, 0 notifications, 0 in queue
  Sent 948 messages, 2 notifications, 0 in queue
  Default minimum time between advertisement runs is 30 seconds
```

```
For address family: IPv4 Unicast
  BGP table version 87, neighbor version 0
  Index 1, Offset 0, Mask 0x2
  Route refresh request: received 0, sent 0, maximum limit 10
Threshold for warning message 80%
```

```
Connections established 2; dropped 2
Last reset 00:02:43, due to BGP Notification sent, update malformed
Message received that caused BGP to send a Notification:
```

```
FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF
00600200 00001940 01010040 02040201
00C84003 040A0000 01800404 00000000
180A0001 180A0002 180A0003 180A0004
180A0005 180A0006 180A0007 180A0008
180A0009 180A000A 180A000B 180A000C
```

```
Peer had exceeded the max. no. of prefixes configured.
Reduce the no. of prefix and clear ip bgp 10.0.0.1 to restore peering
External BGP neighbor can be up to 2 hops away.
No active TCP connection
```

注意：使用此命令可恢复对等体功能：

```
Router_B# clear ip bgp 10.0.0.1
```

[相关信息](#)

- [BGP 在超出最大前缀限制之后重新启动会话](#)
- [BGP 故障排除](#)
- [BGP 案例分析](#)
- [技术支持和文档 - Cisco Systems](#)