

# 了解基本BGP故障排除

## 目录

---

[简介](#)

[先决条件](#)

[背景信息](#)

[确定为BGP对等体实施的配置](#)

[如何理解基本BGP输出](#)

[相关信息](#)

---

## 简介

本文档介绍执行边界网关协议(BGP)故障排除和了解Nexus基本输出的过程。

## 先决条件

### 要求

Cisco 建议您了解以下主题：

- Nexus交换机
- 调试输出中显示“BGP”

### 使用的组件

本文档不限于特定的软件和硬件版本。

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。

## 背景信息

BGP是大型网络中使用的外部网关协议，用于实现自治系统(ASes)之间的路由和可达性信息交换。它是支持Internet全局路由系统的核心路由协议。

## 确定为BGP对等体实施的配置

验证Nexus 9300上启用了BGP功能。

识别并理解BGP进程的配置。

```
switch# show running-config bgp
```


```
!Command: show running-config bgp
!Running configuration last done at: Tue Jul 18 19:45:05 2023
!Time: Tue Jul 18 19:45:44 2023
```

```
version 10.2(4) Bios:version 05.47
feature bgp
```

```
router bgp 64512
  router-id 172.17.255.255
  address-family ipv4 unicast
    network 10.100.1.0/24
  redistribute direct route-map permit_all
  neighbor 10.1.1.1
    remote-as 64512
    address-family ipv4 unicast
      prefix-list allow_in in
      prefix-list allow_out out
      soft-reconfiguration inbound always
  neighbor 172.18.255.255
    remote-as 65535
    update-source loopback10
    ebgp-multihop 3
    address-family ipv4 unicast
      route-map block_route in
      route-map no_local out
      soft-reconfiguration inbound always
```

-> The AS number of the local BGP speaker.  
-> BGP speaker identifier.  
-> Global address family configuration.  
-> Specifies a network as local to this autonomous system.  
-> Routes redistribution from other routing protocols (OSPF, EIGRP, etc.).  
-> IP address of the remote BGP peer.  
-> The AS number of the remote BGP peer.\*  
-> Local address family configuration.  
-> Prefix-list applied at the inbound of the BGP peer. \*\*  
-> Prefix-list applied at the outbound of the BGP peer. \*\*  
-> Store the inbound BGP route updates.  
-> Interface used to source BGP updates.  
-> Maximum hops to reach peer IP address, it modifies the ebgp-multihop value.  
-> Route-map applied at the inbound of the BGP peer. \*\*  
-> Route-map applied at the outbound of the BGP peer. \*\*

---

 注意：在本地和远程配置的不同ASN标识iBGP会话，而配置的其他ASN标识eBGP会话。路由映射的层次结构值比应用于同一对等体的前缀列表的层次结构值高。

---

## 如何理解基本BGP输出

### BGP对等体状态

```
switch# show ip bgp summary
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 172.17.255.255, local AS number 64512 -> Local BGP ID and ASN
BGP table version is 67, IPv4 Unicast config peers 2, capable peers 2
20 network entries and 19 paths using 5424 bytes of memory
BGP attribute entries [6/2112], BGP AS path entries [2/20]
BGP community entries [0/0], BGP clusterlist entries [0/0]
13 received paths for inbound soft reconfiguration
12 identical, 0 modified, 1 filtered received paths using 96 bytes
```

Neighbor	V AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.1.1	4 64512	346	334	67	0	0	05:25:12 8	-> BGP peer IP address, remote
172.18.255.255	4 65535	334	327	67	0	0	05:18:00 8	

 注：“打开/关闭”部分显示BGP对等会话已打开或关闭的时间。

State/PfxRcd部分显示会话当前处于的BGP状态。它显示建立状态时从其对等体获知的前缀的数量。

BGP状态	
空闲	这是BGP等待“开始事件”的第一个状态。当某人配置新的BGP邻居或重置已建立的BGP对等时，会发生启动事件。
连接	BGP正在等待TCP三次握手完成。成功后，它会继续进入OpenSent状态。如果它发生故障，我们继续进入“活动”状态。
主用	BGP尝试另一个TCP三次握手以建立与远程BGP邻居的连接。如果成功，则转到OpenSent状态。如果ConnectRetry计时器过期，则我们返回到连接状态。
OpenSent	在此状态下，BGP等待来自远程BGP邻居的Open消息。
OpenConfirm	BGP等待来自远程BGP邻居的keepalive消息。
已建立	BGP邻居邻接关系完成，BGP路由器发送更新数据包以交换路由信息。

了解BGP对等体信息。

```

switch# show ip bgp neighbors 10.1.1.1
BGP neighbor is 10.1.1.1, remote AS 64512, ibgp link, Peer index 3
  BGP version 4, remote router ID 172.16.255.255
  Neighbor previous state = OpenConfirm
  BGP state = Established, up for 00:05:29
  Neighbor vrf: default
  Peer is directly attached, interface Ethernet1/49
  Last read 00:00:28, hold time = 180, keepalive interval is 60 seconds
  Last written 00:00:28, keepalive timer expiry due 00:00:31
  Received 363 messages, 0 notifications, 0 bytes in queue
  Sent 354 messages, 1 notifications, 0(0) bytes in queue
  Enhanced error processing: On
    0 discarded attributes
  Connections established 2, dropped 1
  Last update recd 00:05:28, Last update sent = 00:05:28
    Last reset by us 00:06:21, due to holdtimer expired error
  Last error length sent: 0
  Reset error value sent: 0
  Reset error sent major: 4 minor: 0
  Notification data sent:
  Last reset by peer never, due to No error
  Last error length received: 0
  Reset error value received 0
  Reset error received major: 0 minor: 0
  Notification data received:

Neighbor capabilities:
  -> Peer IP address, remote
  -> BGP version, Peer Router ID
  -> Previous BGP state
  -> Current BGP state and up time
  -> VRF used for the peer
  -> Interface used to source
  -> Amount of time from last
  -> Amount of time from last
  -> Counters informing the
  -> Timer of the last major
  -> Last reset timer and

```

Dynamic capability: advertised (mp, refresh, gr) received (mp, refresh, gr)  
Dynamic capability (old): advertised received  
Route refresh capability (new): advertised received  
Route refresh capability (old): advertised received  
4-Byte AS capability: advertised received  
Address family IPv4 Unicast: advertised received  
Graceful Restart capability: advertised received

Graceful Restart Parameters:

Address families advertised to peer:  
  IPv4 Unicast  
Address families received from peer:  
  IPv4 Unicast  
Forwarding state preserved by peer for:  
Restart time advertised to peer: 120 seconds  
Stale time for routes advertised by peer: 300 seconds  
Restart time advertised by peer: 120 seconds  
Extended Next Hop Encoding Capability: advertised received  
Receive IPv6 next hop encoding Capability for AF:  
  IPv4 Unicast  VPNv4 Unicast

Message statistics:

	Sent	Rcvd
Opens:	2	2
Notifications:	1	0
Updates:	22	20
Keepalives:	340	339
Route Refresh:	1	0
Capability:	2	2
Total:	354	363
Total bytes:	7949	7524
Bytes in queue:	0	0

For address family: IPv4 Unicast  
BGP table version 88, neighbor version 88  
8 accepted prefixes (8 paths), consuming 2176 bytes of memory  
0 received prefixes treated as withdrawn  
11 sent prefixes (11 paths)  
Inbound soft reconfiguration allowed(always)  
Third-party Nexthop will not be computed.  
Inbound ip prefix-list configured is allow\_in, handle obtained  
Outbound ip prefix-list configured is allow\_out, handle obtained  
Last End-of-RIB received 00:00:01 after session start  
Last End-of-RIB sent 00:00:01 after session start  
First convergence 00:00:01 after session start with 11 routes sent

-> Amount of prefixes ad

Local host: 10.1.1.2, Local port: 28262  
Foreign host: 10.1.1.1, Foreign port: 179  
fd = 73

## 了解BGP表

此输出显示BGP表中获取的所有前缀的状态、下一跳、度量、本地优先级、权重和AS路径。

```
switch# show ip bgp  
BGP routing table information for VRF default, address family IPv4 Unicast
```

BGP table version is 88, Local Router ID is 172.17.255.255  
 Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, \*-valid, >-best  
 Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected  
 Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2

Network	Next Hop	Metric	LocPrf	Weight	Path
*>r10.1.1.0/30	0.0.0.0	0	100	32768	?
*>i10.100.1.0/24	10.1.1.1		100	0	i
*>i10.100.2.0/24	10.1.1.1		100	0	i
*>i10.100.3.0/24	10.1.1.1		150	0	i
*>i10.100.4.0/24	10.1.1.1	0	100	0	?
*>i10.100.5.0/24	10.1.1.1	0	100	0	?
*>i10.100.6.0/24	10.1.1.1	0	100	0	?
*>i10.100.7.0/24	10.1.1.1	0	100	0	?
*>i10.100.8.0/24	10.1.1.1	0	100	0	?
*>r172.17.255.255/32	0.0.0.0	0	100	32768	?
*>e172.30.1.0/24	172.18.255.255	0		0	65535 ?
*>e172.30.2.0/24	172.18.255.255	0		0	65535 ?
*>e172.30.3.0/24	172.18.255.255	0		0	65535 ?
*>e172.30.4.0/24	172.18.255.255	0		0	65535 ?
*>e172.30.5.0/24	172.18.255.255	0		0	65535 65534 65533 ?
*>e172.30.6.0/24	172.18.255.255	0		0	65535 65534 65533 ?
*>e172.30.7.0/24	172.18.255.255	0		0	65535 65534 65533 ?
*>e172.30.8.0/24	172.18.255.255	0		0	65535 65534 65533 ?
*>r192.168.1.0/30	0.0.0.0	0	100	32768	?

前缀会通告给特定BGP邻居。

switch# show ip bgp neighbors 172.18.255.255 advertised-routes

Peer 172.18.255.255 routes for address family IPv4 Unicast:  
 BGP table version is 88, Local Router ID is 172.17.255.255  
 Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, \*-valid, >-best  
 Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected  
 Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i10.100.1.0/24	10.1.1.1		100	0	i
*>i10.100.2.0/24	10.1.1.1		100	0	i
*>i10.100.3.0/24	10.1.1.1		150	0	i
*>i10.100.4.0/24	10.1.1.1	0	100	0	?
*>i10.100.5.0/24	10.1.1.1	0	100	0	?
*>i10.100.6.0/24	10.1.1.1	0	100	0	?
*>i10.100.7.0/24	10.1.1.1	0	100	0	?
*>i10.100.8.0/24	10.1.1.1	0	100	0	?


在任何过滤器 ( 前缀列表和/或路由映射 ) 之前从BGP对等体接收的前缀\*

switch# show ip bgp neighbors 172.18.255.255 received-routes

Peer 172.18.255.255 routes for address family IPv4 Unicast:  
 BGP table version is 88, Local Router ID is 172.17.255.255

Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, \*-valid, >-best  
 Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected  
 Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2

Network	Next Hop	Metric	LocPrf	Weight Path
* e172.18.255.255/32	172.18.255.255	0		0 65535 ?
*>e172.30.1.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.2.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.3.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.4.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.5.0/24	172.18.255.255	0		0 65535 65534 65533 ?
*>e172.30.6.0/24	172.18.255.255	0		0 65535 65534 65533 ?
*>e172.30.7.0/24	172.18.255.255	0		0 65535 65534 65533 ?
*>e172.30.8.0/24	172.18.255.255	0		0 65535 65534 65533 ?

 备注 必须在邻居上配置入站软重新配置

在过滤器 ( 前缀列表和/或路由映射 ) 后从BGP对等体接收的前缀

```
switch# show ip bgp neighbors 172.18.255.255 routes
```

```
Peer 172.18.255.255 routes for address family IPv4 Unicast:
BGP table version is 88, Local Router ID is 172.17.255.255
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2
```

Network	Next Hop	Metric	LocPrf	Weight Path
*>e172.30.1.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.2.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.3.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.4.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.5.0/24	172.18.255.255	0		0 65535 65534 65533 ?
*>e172.30.6.0/24	172.18.255.255	0		0 65535 65534 65533 ?
*>e172.30.7.0/24	172.18.255.255	0		0 65535 65534 65533 ?
*>e172.30.8.0/24	172.18.255.255	0		0 65535 65534 65533 ?

特定前缀的详细路径信息。

```
switch# show ip bgp 172.30.6.0
BGP routing table information for VRF default, address family IPv4 Unicast
BGP routing table entry for 172.30.6.0/24, version 28
Paths: (3 available, best #3)
Flags: (0x8000001a) (high32 00000000) on xmit-list, is in urib, is best urib route, is in HW
```

```
Path type: external, path is valid, not best reason: Router Id, no labeled nexthop
AS-Path: 65535 65534 65533 , path sourced external to AS
172.20.255.255 (metric 0) from 172.20.255.255 (172.20.255.255)
```

-> Prefi  
-> Numbe

-> As Pa  
-> Next

Origin incomplete, MED 0, localpref 100, weight 0

Path type: external, path is valid, not best reason: newer EBGP path, no labeled nexthop  
AS-Path: 65535 65534 65533 , path sourced external to AS  
172.19.255.255 (metric 0) from 172.19.255.255 (172.19.255.255)  
Origin incomplete, MED 0, localpref 100, weight 0

Advertised path-id 1


Path type: external, path is valid, is best path, no labeled nexthop, in rib  
AS-Path: 65535 65534 65533 , path sourced external to AS  
172.18.255.255 (metric 0) from 172.18.255.255 (172.18.255.255)  
Origin incomplete, MED 0, localpref 100, weight 0

-> Path s

Path-id 1 advertised to peers:  
10.1.1.2

-> BGP pee

---

 注:\*AS-Path表示到达前缀源位置所经过的ASN。  
\*\*AS路径从右到左读取。

---

要查看BGP中的最佳路径选择过程，请参阅[BGP最佳路径选择](#)。

## 相关信息

- [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）。