

在Nexus 9000上配置BGP路由反射器

目录

[简介](#)

[先决条件](#)

[要求](#)

[使用的组件](#)

[背景信息](#)

[配置](#)

[网络图](#)

[目标](#)

[初始验证](#)

[配置](#)

[验证](#)

[摘要](#)

[相关信息](#)

简介

本文档介绍在Nexus 9000系列上配置边界网关协议(BGP)路由反射器的过程。

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。如果您的网络处于生产状态，请确保您了解计划应用的任何命令的潜在影响。

背景信息

BGP用于避免路由环路的机制是通过将自己的自治系统编号(ASN)添加到AS PATH属性中。对于IBGP拓扑，更新不会修改AS-PATH属性，因为它们属于同一ASN，这可能会导致路由环路。

如果从iBGP对等体收到BGP前缀，则无法将此前缀通告到另一个iBGP邻居。BGP使用其水平分割规则来防止同一ASN内的环路。

此规则强制您使用全网状拓扑与每个iBGP邻居交换前缀。对于具有全网状场景的大型网络，由于它使用过多资源创建对等网络，因此不能进行扩展。

iBGP的全网状拓扑的替代方案包括：

- 路由反射
- 联盟


本文档仅重点介绍BGP路由反射器的用法。

使用路由反射器可以克服全网状限制。路由反射器是AS（自治系统）内的指定设备，它将iBGP前缀反射到同一AS中的其他路由器。

配置

要检查BGP路由反射器配置指南，请参阅[配置高级BGP](#)

网络图



在此图中，Nexus R2、Nexus R3、Nexus R4、Nexus R5和Nexus R6建立了iBGP邻居关系。而且，Nexus R1和Nexus R5建立了eBGP邻居关系。

目标

将Nexus R1生成的网络交换到同一ASN网络中的所有65000。

初始验证

Nexus R1执行前缀通告：

```
R1# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 78, Local Router ID is 192.168.1.1
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-i
njected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - b
```

est2

Network	Next Hop	Metric	LocPrf	Weight	Path
*>r192.168.1.0/30	0.0.0.0	0	100	32768	?
*>r192.168.2.0/30	0.0.0.0	0	100	32768	?
*>r192.168.3.0/30	0.0.0.0	0	100	32768	?
*>r192.168.4.0/30	0.0.0.0	0	100	32768	?
*>r192.168.5.0/30	0.0.0.0	0	100	32768	?

R1#

Nexus R5收到来自Nexus R1的通告，Nexus R1是eBGP会话。

Nexus R5

```
R5# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 123, Local Router ID is 172.16.1.2
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
*>r10.1.1.4/30	0.0.0.0	0	100	32768	?
*>r10.1.1.8/30	0.0.0.0	0	100	32768	?
*>r172.16.1.0/30	0.0.0.0	0	100	32768	?
*>e192.168.1.0/30	172.16.1.1	0		0	64512 ?
*>e192.168.2.0/30	172.16.1.1	0		0	64512 ?
*>e192.168.3.0/30	172.16.1.1	0		0	64512 ?
*>e192.168.4.0/30	172.16.1.1	0		0	64512 ?
*>e192.168.5.0/30	172.16.1.1	0		0	64512 ?

R5#

Nexus R5向其iBGP对等体Nexus R2和R6通告前缀。

```
R2# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 33, Local Router ID is 10.1.1.6
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
*>i10.1.1.4/30	10.1.1.5	0	100	0	?
*>i10.1.1.8/30	10.1.1.5	0	100	0	?
*>i172.16.1.0/30	10.1.1.5	0	100	0	?
*>i192.168.1.0/30	172.16.1.1	0	100	0	64512 ?
*>i192.168.2.0/30	172.16.1.1	0	100	0	64512 ?
*>i192.168.3.0/30	172.16.1.1	0	100	0	64512 ?
*>i192.168.4.0/30	172.16.1.1	0	100	0	64512 ?

```
*>i192.168.5.0/30    172.16.1.1          0          100          0 64512 ?
```

R2#

```
R6# show ip bgp
```

```
BGP routing table information for VRF default, address family IPv4 Unicast
```

```
BGP table version is 33, Local Router ID is 10.1.1.10
```

```
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
*>i10.1.1.4/30	10.1.1.9	0	100	0 ?
*>i10.1.1.8/30	10.1.1.9	0	100	0 ?
*>i172.16.1.0/30	10.1.1.9	0	100	0 ?
*>i192.168.1.0/30	172.16.1.1	0	100	0 64512 ?
*>i192.168.2.0/30	172.16.1.1	0	100	0 64512 ?
*>i192.168.3.0/30	172.16.1.1	0	100	0 64512 ?
*>i192.168.4.0/30	172.16.1.1	0	100	0 64512 ?
*>i192.168.5.0/30	172.16.1.1	0	100	0 64512 ?

R6#

根据环路避免规则，BGP不会向其他iBGP对等体通告iBGP对等体获取的任何前缀。

```
R2# show ip bgp neighbors 10.1.2.2 advertised-routes
```

```
Peer 10.1.2.2 routes for address family IPv4 Unicast:
```

```
BGP table version is 88, Local Router ID is 10.150.0.1
```

```
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
*>r10.1.1.4/30	0.0.0.0	0	100	32768 ?
*>r10.1.2.0/30	0.0.0.0	0	100	32768 ?
*>r10.1.2.4/30	0.0.0.0	0	100	32768 ?
*>r10.1.2.8/30	0.0.0.0	0	100	32768 ?
*>r10.150.0.0/30	0.0.0.0	0	100	32768 ?

R2#

Nexus R3和Nexus R4不会接收Nexus R1生成的前缀。

```
R3# show ip bgp
```

```
BGP routing table information for VRF default, address family IPv4 Unicast
```

```
BGP table version is 28, Local Router ID is 10.100.100.1
```

```
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
*>110.100.100.0/24	0.0.0.0		100	32768 i


```
R4# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 6, Local Router ID is 10.200.200.1
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
*>110.200.200.0/24	0.0.0.0		100	32768 i

R4#

配置

对于上述拓扑，要转换为路由反射器的选定设备是Nexus R2和Nexus R6，因为这些设备连接到同一ASN中的所有iBGP对等体。

 注意：路由反射器配置导致重置每个修改的bgp对等体。此配置必须在维护窗口下完成。

 注：通过拓扑设计，Nexus R2和R6成为ASN网络的路由反射65000。

Nexus R1

```
R1# show run bgp
```

```
!Command: show running-config bgp
!Running configuration last done at: Wed Jan 31 02:43:31 2024
!Time: Wed Jan 31 23:24:28 2024
```

```
version 10.2(5) Bios:version 05.47
feature bgp
```

```
router bgp 64512
  address-family ipv4 unicast
    redistribute direct route-map REDISTRIBUTE_LOCAL
  neighbor 172.16.1.2
    remote-as 65000
  address-family ipv4 unicast
```

R1#

Nexus R2

R2# show run bgp

```
!Command: show running-config bgp
!Running configuration last done at: Wed Jan 31 03:10:49 2024
!Time: Wed Jan 31 23:24:15 2024
```

```
version 10.2(5) Bios:version 05.47
feature bgp
```

```
router bgp 65000
  cluster-id 10.150.0.1
  address-family ipv4 unicast
    redistribute direct route-map ALLOW
  neighbor 10.1.1.5
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
  neighbor 10.1.2.2
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
  neighbor 10.1.2.6
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
  neighbor 10.1.2.10
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
```

R2#

Nexus R3

R3# show run bgp

```
!Command: show running-config bgp
!Running configuration last done at: Wed Jan 31 02:49:05 2024
!Time: Wed Jan 31 23:10:07 2024
```

```
version 10.2(5) Bios:version 05.47
feature bgp
```

```
router bgp 65000
  address-family ipv4 unicast
    network 10.100.100.0/24
  neighbor 10.1.2.5
    remote-as 65000
    address-family ipv4 unicast
  neighbor 10.1.3.5
    remote-as 65000
    address-family ipv4 unicast
```

R3#

Nexus R4

R4# show run bgp

```
!Command: show running-config bgp
!Running configuration last done at: Wed Jan 31 02:56:11 2024
!Time: Wed Jan 31 23:00:44 2024
```

```
version 10.2(5) Bios:version 05.47
feature bgp
```

```
router bgp 65000
  address-family ipv4 unicast
    network 10.200.200.0/24
  neighbor 10.1.2.1
    remote-as 65000
  address-family ipv4 unicast
  neighbor 10.1.3.1
    remote-as 65000
  address-family ipv4 unicast
```

R4#

Nexus R5

R5# show run bgp

```
!Command: show running-config bgp
!Running configuration last done at: Wed Jan 31 02:48:38 2024
!Time: Wed Jan 31 23:24:15 2024
```

```
version 10.2(5) Bios:version 05.47
feature bgp
```

```
router bgp 65000
  address-family ipv4 unicast
    redistribute direct route-map ALLOW
  neighbor 10.1.1.6
    remote-as 65000
  address-family ipv4 unicast
  neighbor 10.1.1.10
    remote-as 65000
  address-family ipv4 unicast
  neighbor 172.16.1.1
    remote-as 64512
  address-family ipv4 unicast
```


R5#

Nexus R6

R6# show run bgp

```
!Command: show running-config bgp
!Running configuration last done at: Wed Jan 31 03:12:40 2024
!Time: Wed Jan 31 23:24:18 2024
```

```
version 10.2(5) Bios:version 05.47
feature bgp

router bgp 65000
  cluster-id 10.160.0.1
  address-family ipv4 unicast
    redistribute direct route-map ALLOW
  neighbor 10.1.1.9
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
  neighbor 10.1.2.9
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
  neighbor 10.1.3.2
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
  neighbor 10.1.3.6
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
```

R6#

验证

将Nexus R2和Nexus R6配置为路由反射器后，前缀将通告给iBGP对等体。

R2# show ip bgp neighbors 10.1.2.2 advertised-routes

```
Peer 10.1.2.2 routes for address family IPv4 Unicast:
BGP table version is 22, Local Router ID is 10.150.0.1
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 - best
```

Network	Next Hop	Metric	LocPrf	Weight	Path
---------	----------	--------	--------	--------	------

```

*>r10.1.1.4/30      0.0.0.0      0      100      32768 ?
*>i10.1.1.8/30      10.1.1.5     0      100      0 ?
*>r10.1.2.0/30      0.0.0.0     0      100      32768 ?
*>r10.1.2.4/30      0.0.0.0     0      100      32768 ?
*>r10.1.2.8/30      0.0.0.0     0      100      32768 ?
*>i10.100.100.0/24  10.1.2.6    100      0 i
*>r10.150.0.0/30    0.0.0.0     0      100      32768 ?
*>i172.16.1.0/30    10.1.1.5     0      100      0 ?
*>i192.168.1.0/30   172.16.1.1   0      100      0 64512 ?
*>i192.168.2.0/30   172.16.1.1   0      100      0 64512 ?
*>i192.168.3.0/30   172.16.1.1   0      100      0 64512 ?
*>i192.168.4.0/30   172.16.1.1   0      100      0 64512 ?
*>i192.168.5.0/30   172.16.1.1   0      100      0 64512 ?

```

R2#

Nexus R3和Nexus R4接收Nexus R1生成的前缀。

```

R3# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 108, Local Router ID is 10.100.100.1
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
* i10.1.1.4/30	10.1.2.9	0	100	0	?
*>i	10.1.2.5	0	100	0	?
*>i10.1.1.8/30	10.1.3.5	0	100	0	?
* i	10.1.2.10	0	100	0	?
* i10.1.2.0/30	10.1.2.9	0	100	0	?
*>i	10.1.2.5	0	100	0	?
* i10.1.2.4/30	10.1.2.9	0	100	0	?
*>i	10.1.2.5	0	100	0	?
* i10.1.2.8/30	10.1.3.5	0	100	0	?
*>i	10.1.2.5	0	100	0	?
* i10.1.3.0/30	10.1.2.10	0	100	0	?
*>i	10.1.3.5	0	100	0	?
* i10.1.3.4/30	10.1.2.10	0	100	0	?
*>i	10.1.3.5	0	100	0	?
*>l10.100.100.0/24	0.0.0.0		100	32768	i
* i10.150.0.0/30	10.1.2.9	0	100	0	?
*>i	10.1.2.5	0	100	0	?
* i10.160.0.0/30	10.1.2.10	0	100	0	?
*>i	10.1.3.5	0	100	0	?
* i10.200.200.0/24	10.1.3.2		100	0	i
*>i	10.1.2.2		100	0	i
* i172.16.1.0/30	10.1.1.9	0	100	0	?
*>i	10.1.1.5	0	100	0	?
* i192.168.1.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.2.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.3.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.4.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.5.0/30	172.16.1.1	0	100	0	64512 ?

```
*>i 172.16.1.1 0 100 0 64512 ?
```

R3#

R4# show ip bgp

BGP routing table information for VRF default, address family IPv4 Unicast

BGP table version is 78, Local Router ID is 10.200.200.1

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
* i10.1.1.4/30	10.1.2.9	0	100	0	?
*>i	10.1.2.1	0	100	0	?
*>i10.1.1.8/30	10.1.3.1	0	100	0	?
* i	10.1.2.10	0	100	0	?
* i10.1.2.0/30	10.1.2.9	0	100	0	?
*>i	10.1.2.1	0	100	0	?
* i10.1.2.4/30	10.1.2.9	0	100	0	?
*>i	10.1.2.1	0	100	0	?
* i10.1.2.8/30	10.1.3.1	0	100	0	?
*>i	10.1.2.1	0	100	0	?
* i10.1.3.0/30	10.1.2.10	0	100	0	?
*>i	10.1.3.1	0	100	0	?
* i10.1.3.4/30	10.1.2.10	0	100	0	?
*>i	10.1.3.1	0	100	0	?
* i10.100.100.0/24	10.1.3.6		100	0	i
*>i	10.1.2.6		100	0	i
* i10.150.0.0/30	10.1.2.9	0	100	0	?
*>i	10.1.2.1	0	100	0	?
* i10.160.0.0/30	10.1.2.10	0	100	0	?
*>i	10.1.3.1	0	100	0	?
*>l10.200.200.0/24	0.0.0.0		100	32768	i
* i172.16.1.0/30	10.1.1.9	0	100	0	?
*>i	10.1.1.5	0	100	0	?
* i192.168.1.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.2.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.3.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.4.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.5.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?

R4#

路由反射器获取的前缀显示BGP对等体IP和集群ID。

R4# show ip bgp 192.168.2.0

BGP routing table information for VRF default, address family IPv4 Unicast

BGP routing table entry for 192.168.2.0/30, version 63

Paths: (2 available, best #2)

Flags: (0x8000001a) (high32 00000000) on xmit-list, is in urib, is best urib route, is in HW


```
Path type: internal, path is valid, not best reason: Neighbor Address, no labeled nexthop
AS-Path: 64512 , path sourced external to AS
 172.16.1.1 (metric 0) from 10.1.3.1 (10.160.0.1) >>>>>> Peer IP (Cluster ID)
  Origin incomplete, MED 0, localpref 100, weight 0
  Originator: 172.16.1.2 Cluster list: 10.160.0.1
```

Advertised path-id 1

```
Path type: internal, path is valid, is best path, no labeled nexthop, in rib
AS-Path: 64512 , path sourced external to AS
 172.16.1.1 (metric 0) from 10.1.2.1 (10.150.0.1)
  Origin incomplete, MED 0, localpref 100, weight 0
  Originator: 172.16.1.2 Cluster list: 10.150.0.1
```


Path-id 1 not advertised to any peer

R4#

 注意：路由反射器接收的Nexus R3和Nexus R4路由不会作为环路避免规则的一部分通告给其他iBGP对等体。

摘要

路由反射器配置不需要复杂的配置来将路由反射到其iBGP客户端，只需在地址系列下添加关键字“route-reflector-client”，即可通知系统充当该邻居的路由反射器。

 注意：要查看Nexus中的BGP联盟，请参阅[在Nexus 9000上配置BGP联盟](#)

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）。