

# 在Nexus 9000上配置BGP路由反射器

## 目錄

---

[簡介](#)

[必要條件](#)

[需求](#)

[採用元件](#)

[背景資訊](#)

[設定](#)

[網路圖表](#)

[目標](#)

[初始驗證](#)

[組態](#)

[驗證](#)

[摘要](#)

[相關資訊](#)

---

## 簡介

本檔案介紹在Nexus 9000系列上設定邊界閘道通訊協定(BGP)路由反射器的程式。

## 必要條件

### 需求

思科建議您瞭解以下主題：

- Nexus交換機
- BGP

### 採用元件

本文件所述內容不限於特定軟體和硬體版本。

本文中的資訊是根據特定實驗室環境內的裝置所建立。文中使用到的所有裝置皆從已清除（預設）的組態來啟動。如果您的網路正在作用，請確保您已瞭解計畫套用任何指令可能造成的影響。

## 背景資訊

BGP用來避免路由回圈的機制是將其自己的自治系統編號(ASN)新增到AS PATH屬性。對於IBGP拓撲，更新不會修改AS-PATH屬性，因為它們屬於同一個ASN，這可能會導致路由環路。

如果從iBGP對等點收到BGP首碼，則無法將這個首碼通告給另一個iBGP鄰居。BGP使用其水準分割規則來防止同一ASN中的環路。

此規則強制您使用全網狀拓撲與每個iBGP鄰居交換字首。對於具有全網狀方案的大型網路，由於它使用過多的資源來建立對等網路，因此不能進行擴展。

iBGP的全網狀拓撲的替代方案有：

- 路由反射
- 聯盟


本文僅重點介紹BGP路由反射器的使用情況。

使用路由反射器克服了全網格的限制。路由反射器是AS（自治系統）內將iBGP字首反射給同一AS中其他路由器的指定裝置。

## 設定

要檢查BGP路由反射器配置指南，請參閱[配置高級BGP](#)

## 網路圖表



在此圖中，Nexus R2、Nexus R3、Nexus 4、Nexus R5和Nexus R6建立了iBGP鄰居關係。而且，Nexus R1和Nexus R5建立了eBGP鄰居關係。

## 目標

將Nexus R1生成的網路交換到同一ASN網路中的所有65000。

## 初始驗證

Nexus R1執行字首通告：

```
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收到來自Nexus R1的通告，該通告是eBGP會話。

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對等體Nexus R2和R6通告字首。

```
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#

根據環路迴避規則，BGP不會向其他iBGP對等體通告iBGP對等體獲取的任何字首。

```
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不會接收Nexus R1生成的字首。

```
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#


## 組態

對於上述拓撲，要轉換為路由反射器的選定裝置是Nexus R2和Nexus R6，因為這些裝置與同一ASN中的所有iBGP對等體都有連線。

---

 注意：路由反射器配置會導致每個修改的bgp對等體都重置。此配置必須在維護視窗下完成。

---

 註：通過拓撲設計，Nexus R2和R6成為ASN網路的路由反65000器。

---

### 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#

## 驗證

將Nexus R2和Nexus R6轉換為路由反射器後，字首將通告給iBGP對等體。

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生成的字首。

```

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#
```

路由反射器學習的字首顯示BGP對等體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路由不會作為環路避免規則的一部分通告給其他iBGP對等體。

---

## 摘要

路由反射器組態不需要複雜組態即可將路由反射到其iBGP使用者端，只需在位址系列下新增關鍵字「route-reflector-client」，即可通知系統充當該鄰居的路由反射器。

---

 注意：要檢視Nexus中的BGP聯盟，請參閱[在Nexus 9000上配置BGP聯盟](#)

---

## 相關資訊

- [思科技術支援與下載](#)

## 關於此翻譯

思科已使用電腦和人工技術翻譯本文件，讓全世界的使用者能夠以自己的語言理解支援內容。請注意，即使是最佳機器翻譯，也不如專業譯者翻譯的內容準確。Cisco Systems, Inc. 對這些翻譯的準確度概不負責，並建議一律查看原始英文文件（提供連結）。