

# 在Nexus 9000上配置BGP聯盟

## 目錄

---

[簡介](#)  
[必要條件](#)  
[需求](#)  
[採用元件](#)  
[背景資訊](#)  
[配置示例](#)  
[網路圖表](#)  
[組態](#)  
[驗證](#)

---

## 簡介

本檔案介紹在Nexus 9000系列上設定邊界閘道通訊協定(BGP)聯盟的程式。

## 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

本文檔不限於特定的軟體和硬體版本。

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。如果您的網路運作中，請確保您瞭解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屬性中。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。

BGP使用其水準分割規則，該規則表明BGP裝置通過IBGP對等體接收字首，並且無法將該字首通告給另一個IBGP鄰居。該規則強制您使用全網狀拓撲與每個IBGP鄰居交換字首。對於使用大量資源建立對等體的大型網路，使用全網狀方案是不可擴展的。

iBGP的全網狀拓撲的替代方案有：

- 路由反射
- 聯盟


本文檔僅重點介紹BGP聯合體的用法。

BGP聯盟通過將一個ASN劃分為多個子自治系統並將它們分組到一個聯盟中，消除了全網狀拓撲的需要。

## 配置示例

要檢查BGP聯盟配置指南，請參閱[高級BGP配置指南Nexus 9000](#)。

## 網路圖表



附註：在拓撲中，Nexus 1、Nexus 2、Nexus 3和Nexus 4屬於同一個聯盟ASN 64513，將聯盟劃分為兩個較小的ASN(64512和64514)

附註：Nexus 1和Nexus 3建立iBGP鄰居關係。

附註：Nexus 2和Nexus 4建立iBGP鄰居關係。

附註：Nexus 1和Nexus 2建立了eBGP鄰居關係。

## 組態

 附註：聯盟配置導致重置到每個已建立的bgp對等體。此配置必須在維護視窗下完成。

### Nexus 1

```
Nexus1# show running-config bgp

!Command: show running-config bgp
!Running configuration last done at: Thu Jul 20 21:13:08 2023
!Time: Thu Jul 20 23:28:41 2023

version 10.2(5) Bios:version 05.47
feature bgp

router bgp 64512
    confederation identifier 64513
    confederation peers 64514
    log-neighbor-changes
    neighbor 10.0.5.2
        remote-as 64515
        address-family ipv4 unicast
    neighbor 172.16.255.2
        remote-as 64514
        update-source loopback0
        address-family ipv4 unicast
    neighbor 172.16.255.3
        remote-as 64512
        log-neighbor-changes
        update-source loopback0
        address-family ipv4 unicast
```

--> Local ASN  
--> BGP Confederation identifier  
--> ASN that belong to the Confederation

### Nexus 2

```
Nexus2# show running-config bgp

!Command: show running-config bgp
!Running configuration last done at: Thu Jul 20 21:12:54 2023
!Time: Thu Jul 20 23:31:04 2023

version 10.2(5) Bios:version 05.47
feature bgp

router bgp 64514
    confederation identifier 64513
    confederation peers 64512
    log-neighbor-changes
    neighbor 172.16.255.1
        remote-as 64512
        update-source loopback0
        address-family ipv4 unicast
    neighbor 172.16.255.4
        remote-as 64514
```

--> Local ASN  
--> BGP Confederation identifier  
--> ASN that belong to the Confederation

```
update-source loopback0
address-family ipv4 unicast
```

## Nexus 3

```
Nexus3# show running-config bgp

!Command: show running-config bgp
!Running configuration last done at: Thu Jul 20 21:14:54 2023
!Time: Thu Jul 20 23:31:45 2023

version 10.2(5) Bios:version 05.47
feature bgp

router bgp 64512
    confederation identifier 64513
    log-neighbor-changes
    neighbor 172.16.255.1
        remote-as 64512
        log-neighbor-changes
        update-source loopback0
        address-family ipv4 unicast
    --> Local ASN
    --> BGP Confederation identifier
```

## Nexus 4

```
Nexus4# show running-config bgp

!Command: show running-config bgp
!Running configuration last done at: Thu Jul 20 15:11:15 2023
!Time: Thu Jul 20 17:32:58 2023

version 10.2(5) Bios:version 05.47
feature bgp

router bgp 64514
    confederation identifier 64513
    log-neighbor-changes
    address-family ipv4 unicast
        network 172.30.1.0/24
        network 172.30.2.0/24
        network 172.30.3.0/24
        network 172.30.4.0/24
        network 172.30.5.0/24
        network 172.30.6.0/24
        network 172.30.7.0/24
        network 172.30.8.0/24
    neighbor 172.16.255.2
        remote-as 64514
        update-source loopback0
        address-family ipv4 unicast
    --> Local ASN
    --> BGP Confederation identifier
```

## Nexus 5

```
Nexus5# show running-config bgp

!Command: show running-config bgp
!Running configuration last done at: Thu Jul 20 21:13:32 2023
!Time: Thu Jul 20 23:33:34 2023

version 10.2(5) Bios:version 05.47
feature bgp

router bgp 64515                                --> Local ASN
  log-neighbor-changes
  address-family ipv4 unicast
    network 192.168.1.0/24
    network 192.168.2.0/24
    network 192.168.3.0/24
    network 192.168.4.0/24
    network 192.168.5.0/24
    network 192.168.6.0/24
    network 192.168.7.0/24
    network 192.168.8.0/24
  neighbor 10.0.5.1
    remote-as 64513
    address-family ipv4 unicast
```

## 驗證

### Nexus 1

步驟 1: 建立BGP鄰居並交換字首。

```
Nexus1# show ip bgp summary
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 172.16.255.1, local AS number 64512
BGP table version is 28, IPv4 Unicast config peers 3, capable peers 3
16 network entries and 16 paths using 4416 bytes of memory
BGP attribute entries [2/704], BGP AS path entries [2/12]
BGP community entries [0/0], BGP clusterlist entries [0/0]
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.5.2	4	64515	328	321	28	0	0	05:00:34	8
172.16.255.2	4	64514	270	264	28	0	0	04:16:51	8
172.16.255.3	4	64512	318	318	28	0	0	05:12:28	0

步驟 2: 聯盟對等體獲取的字首通過檢視狀態c-confed進行標識。

```
Nexus1# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
```

BGP table version is 28, Local Router ID is 172.16.255.1  
 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
*>c172.30.1.0/24	172.16.255.4		100	0	(64514) i
*>c172.30.2.0/24	172.16.255.4		100	0	(64514) i
*>c172.30.3.0/24	172.16.255.4		100	0	(64514) i
*>c172.30.4.0/24	172.16.255.4		100	0	(64514) i
*>c172.30.5.0/24	172.16.255.4		100	0	(64514) i
*>c172.30.6.0/24	172.16.255.4		100	0	(64514) i
*>c172.30.7.0/24	172.16.255.4		100	0	(64514) i
*>c172.30.8.0/24	172.16.255.4		100	0	(64514) i
*>e192.168.1.0/24	10.0.5.2			0	64515 i
*>e192.168.2.0/24	10.0.5.2			0	64515 i
*>e192.168.3.0/24	10.0.5.2			0	64515 i
*>e192.168.4.0/24	10.0.5.2			0	64515 i
*>e192.168.5.0/24	10.0.5.2			0	64515 i
*>e192.168.6.0/24	10.0.5.2			0	64515 i
*>e192.168.7.0/24	10.0.5.2			0	64515 i
*>e192.168.8.0/24	10.0.5.2			0	64515 i

## Nexus 2

步驟 1: 建立BGP鄰居並交換字首。

```
Nexus2# show ip bgp summary
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 172.16.255.2, local AS number 64514
BGP table version is 27, IPv4 Unicast config peers 2, capable peers 2
16 network entries and 16 paths using 3904 bytes of memory
BGP attribute entries [2/344], BGP AS path entries [1/12]
BGP community entries [0/0], BGP clusterlist entries [0/0]
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.1	4	64512	274	265	27	0	0	04:20:26	8
172.16.255.4	4	64514	269	262	27	0	0	04:16:00	8

步驟 2: 聯盟對等體獲取的字首通過檢視狀態c-confed進行標識。

```
Nexus2# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 27, Local Router ID is 172.16.255.2
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
*>i172.30.1.0/24	172.16.255.4		100	0	i
*>i172.30.2.0/24	172.16.255.4		100	0	i
*>i172.30.3.0/24	172.16.255.4		100	0	i

```

*>i172.30.4.0/24      172.16.255.4          100      0 i
*>i172.30.5.0/24      172.16.255.4          100      0 i
*>i172.30.6.0/24      172.16.255.4          100      0 i
*>i172.30.7.0/24      172.16.255.4          100      0 i
*>i172.30.8.0/24      172.16.255.4          100      0 i
*>c192.168.1.0/24     10.0.5.2            100      0 (64512) 64515 i
*>c192.168.2.0/24     10.0.5.2            100      0 (64512) 64515 i
*>c192.168.3.0/24     10.0.5.2            100      0 (64512) 64515 i
*>c192.168.4.0/24     10.0.5.2            100      0 (64512) 64515 i
*>c192.168.5.0/24     10.0.5.2            100      0 (64512) 64515 i
*>c192.168.6.0/24     10.0.5.2            100      0 (64512) 64515 i
*>c192.168.7.0/24     10.0.5.2            100      0 (64512) 64515 i
*>c192.168.8.0/24     10.0.5.2            100      0 (64512) 64515 i

```

## Nexus 3

步驟 1:已建立BGP鄰居並交換字首。

```

Nexus3# show ip bgp summary
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 172.16.255.3, local AS number 64512
BGP table version is 34, IPv4 Unicast config peers 1, capable peers 1
16 network entries and 16 paths using 4416 bytes of memory
BGP attribute entries [2/704], BGP AS path entries [2/12]
BGP community entries [0/0], BGP clusterlist entries [0/0]

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/
PfxRcd									
172.16.255.1	4	64512	332	323	34	0	0	05:17:52	16

第2步：從Nexus 3的角度來看，字首被學習為正常iBGP字首。從它的角度來看，它沒有任何聯盟鄰居，並且它唯一的對等體是iBGP對等體。

```

Nexus3# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 34, Local Router ID is 172.16.255.3
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
*>i172.30.1.0/24	172.16.255.4		100	0 (64514)	i
*>i172.30.2.0/24	172.16.255.4		100	0 (64514)	i
*>i172.30.3.0/24	172.16.255.4		100	0 (64514)	i
*>i172.30.4.0/24	172.16.255.4		100	0 (64514)	i
*>i172.30.5.0/24	172.16.255.4		100	0 (64514)	i
*>i172.30.6.0/24	172.16.255.4		100	0 (64514)	i
*>i172.30.7.0/24	172.16.255.4		100	0 (64514)	i
*>i172.30.8.0/24	172.16.255.4		100	0 (64514)	i
*>i192.168.1.0/24	10.0.5.2		100	0 64515	i
*>i192.168.2.0/24	10.0.5.2		100	0 64515	i

```

*>i192.168.3.0/24      10.0.5.2          100      0 64515 i
*>i192.168.4.0/24      10.0.5.2          100      0 64515 i
*>i192.168.5.0/24      10.0.5.2          100      0 64515 i
*>i192.168.6.0/24      10.0.5.2          100      0 64515 i
*>i192.168.7.0/24      10.0.5.2          100      0 64515 i
*>i192.168.8.0/24      10.0.5.2          100      0 64515 i

```

## Nexus 4

步驟 1:已建立BGP鄰居並交換字首。

```

Nexus4# show ip bgp summary
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 172.16.255.4, local AS number 64514
BGP table version is 18, IPv4 Unicast config peers 1, capable peers 1
16 network entries and 16 paths using 4416 bytes of memory
BGP attribute entries [2/704], BGP AS path entries [1/12]
BGP community entries [0/0], BGP clusterlist entries [0/0]

Neighbor      V   AS     MsgRcvd     MsgSent    TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.2  4  64514       268        267        18      0      0 04:21:40  8

```

步驟 2: 從Nexus 4的角度來看，字首是作為正常iBGP字首學習的。從它的角度來看，它沒有任何聯盟鄰居，並且它唯一的對等體是iBGP對等體。

```

Nexus4# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 18, Local Router ID is 172.16.255.4
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
*>1172.30.1.0/24      0.0.0.0          100        32768  i
*>1172.30.2.0/24      0.0.0.0          100        32768  i
*>1172.30.3.0/24      0.0.0.0          100        32768  i
*>1172.30.4.0/24      0.0.0.0          100        32768  i
*>1172.30.5.0/24      0.0.0.0          100        32768  i
*>1172.30.6.0/24      0.0.0.0          100        32768  i
*>1172.30.7.0/24      0.0.0.0          100        32768  i
*>1172.30.8.0/24      0.0.0.0          100        32768  i
*>i192.168.1.0/24     10.0.5.2          100      0 (64512) 64515 i
*>i192.168.2.0/24     10.0.5.2          100      0 (64512) 64515 i
*>i192.168.3.0/24     10.0.5.2          100      0 (64512) 64515 i
*>i192.168.4.0/24     10.0.5.2          100      0 (64512) 64515 i
*>i192.168.5.0/24     10.0.5.2          100      0 (64512) 64515 i
*>i192.168.6.0/24     10.0.5.2          100      0 (64512) 64515 i
*>i192.168.7.0/24     10.0.5.2          100      0 (64512) 64515 i
*>i192.168.8.0/24     10.0.5.2          100      0 (64512) 64515 i

```

## Nexus 5

步驟 1:已建立BGP鄰居並交換字首。在這種情況下，唯一的鄰居是Nexus 3，這是eBGP鄰居。

```
Nexus5# show ip bgp summary
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 172.16.255.5, local AS number 64515
BGP table version is 18, IPv4 Unicast config peers 1, capable peers 1
16 network entries and 16 paths using 3904 bytes of memory
BGP attribute entries [2/344], BGP AS path entries [1/6]
BGP community entries [0/0], BGP clusterlist entries [0/0]

Neighbor          V     AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.5.1          4    64513      317      317        18      0      0 05:11:26  8
```

步驟 2:從Nexus 5的角度來看，從Nexus 3收到的字首是使用聯盟ASN的AS PATH接收的，但作為正常eBGP字首接收。

```
Nexus5# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 18, Local Router ID is 172.16.255.5
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
*>e172.30.1.0/24      10.0.5.1                  0 64513 i
*>e172.30.2.0/24      10.0.5.1                  0 64513 i
*>e172.30.3.0/24      10.0.5.1                  0 64513 i
*>e172.30.4.0/24      10.0.5.1                  0 64513 i
*>e172.30.5.0/24      10.0.5.1                  0 64513 i
*>e172.30.6.0/24      10.0.5.1                  0 64513 i
*>e172.30.7.0/24      10.0.5.1                  0 64513 i
*>e172.30.8.0/24      10.0.5.1                  0 64513 i
*>l192.168.1.0/24    0.0.0.0                   100 32768 i
*>l192.168.2.0/24    0.0.0.0                   100 32768 i
*>l192.168.3.0/24    0.0.0.0                   100 32768 i
*>l192.168.4.0/24    0.0.0.0                   100 32768 i
*>l192.168.5.0/24    0.0.0.0                   100 32768 i
*>l192.168.6.0/24    0.0.0.0                   100 32768 i
*>l192.168.7.0/24    0.0.0.0                   100 32768 i
*>l192.168.8.0/24    0.0.0.0                   100 32768 i
```

---

附註：Nexus 5不知道Nexus 1 - Nexus 4之間使用的聯盟。從它的角度來看，這些是普通eBGP字首。

## 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）。