

使用BGP社群值配置上游提供商網路

目錄

[簡介](#)

[必要條件](#)

[需求](#)

[採用元件](#)

[背景資訊](#)

[慣例](#)

[配置和控制路由策略](#)

[網路圖表](#)

[組態](#)

[R3](#)

[R1](#)

[R2](#)

[驗證](#)

[相關資訊](#)

簡介

本檔案介紹如何使用BGP社群值控制上游提供者網路中的路由原則。

必要條件

需求

本檔案需要瞭解BGP路由通訊協定及其運作。

採用元件

本文件所述內容不限於特定軟體和硬體版本。但是本檔案中的資訊是根據以下軟體版本：

- Cisco IOS®軟體版本12.2

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

雖然社群本身不會改變[BGP最佳路徑](#)進程，但社群可用作標籤一組路由的標誌。然後，上游服務提供商路由器可以使用這些標誌在其網路中應用特定的路由策略（例如，本地首選項）。

提供商在提供商網路內的可配置社群值和相應的本地首選項值之間進行對映。可以具有要求修改提供商網路中的LOCAL_PREF的特定策略，以設定其路由更新中的相應團體值。

團體是一組字首，它們共用一些公共屬性，可以使用BGP團體屬性進行配置。BGP社群屬性是可選的可變長度傳遞屬性。該屬性由一組四個八位組值組成，它們指定了一個社群。社群屬性值在前兩個八位元中使用自治系統(AS)編號進行編碼，其他兩個八位元由AS定義。首碼可以有許多個社群屬性。看到字首中的多個團體屬性的BGP發言人可以基於一個、部分或全部屬性來操作。路由器可以先新增或修改社群屬性，然後再將屬性傳遞給其他對等體。若要詳細瞭解社群屬性，請參閱[BGP個案研究](#)。

本地優先屬性是向AS指示到達特定網路的首選路徑。當有多個路徑到達同一個目的地時，會選擇優先順序較高的路徑（本地優先順序屬性的預設值為100）。有關詳細資訊，請參閱案例研究。

慣例

如需文件慣例的詳細資訊，請參閱思科技術提示慣例。

配置和控制路由策略



注意：要查詢有關本文檔中使用的命令的其他資訊，請使用命令查詢工具。

為簡化起見，假定在上游服務提供商(AS 100)和您的裝置(AS 30)之間建立了社群屬性和本地首選項屬性對映。

本地優先使用程度	社群價值
130	100:300
125	100:250

如果通告字首的社群屬性等於100:300，則上游服務提供商會將這些路由的本地優先順序設定為130和125（如果社群屬性等於100:250）。

如果您更改向服務提供商通告的字首的團體值，則此操作將讓您控制服務提供商網路中的路由策略。

在網路圖中，AS 30希望將此路由策略與團體屬性結合使用。

- 從AS 100到網路10.0.10.0/24的入站流量通過R1-R3鏈路傳輸。如果R1-R3鏈路發生故障，所有流量都會通過R2-R3。
- 從AS 100到網路10.1.0.0/24的入站流量通過R2-R3鏈路傳輸。如果R2-R3鏈路發生故障，所有流量都會通過R1-R3。

為實現此路由策略，R3會以下列方式通告其字首：

到R1:

- 10.0.10.0/24，社群屬性為100:300
- 10.1.0.0/24，社群屬性為100:250

到R2:

- 10.0.10.0/24，社群屬性為100:250
- 10.1.0.0/24，社群屬性為100:300

一旦BGP鄰居R1和R2收到來自R3的字首，R1和R2就會根據團體與本地首選項屬性之間的對映應用已配置的策略（如上表所示），從而實現您指定的路由策略(AS 30)。R1在BGP表中安裝字首。


- 10.0.10.0/24，本地優先順序為130
- 10.1.0.0/24，本地優先順序為125

R2在其BGP表中安裝字首：

- 10.0.10.0/24，本地優先順序為125
- 10.1.0.0/24，本地優先順序為130

由於BGP路徑選擇標準中優先使用較高的本地優先順序，因此將本地優先順序為130（130大於125）的路徑選為AS 100中的最佳路徑，並將其安裝在R1和R2的IP路由表中。有關BGP路徑選取標準的詳細資訊，請參閱[BGP最佳路徑選取演演算法](#)。

網路圖表



組態

本檔案會使用以下設定：

- R3
- R1
- R2

R3

```
hostname R3
!
interface Loopback0
ip address 10.0.10.0 255.255.255.0
!
interface Ethernet0/0
ip address 10.1.0.0 255.255.255.1
!
interface Serial8/0
ip address 10.10.13.3 255.255.255.0

!--- Interface connected to R1

!
interface Serial9/0
ip address 10.10.23.3 255.255.255.0

!--- Interface connected to R2

!
router bgp 30
network 10.0.10.0 mask 255.255.255.0
network 10.1.0.0 mask 255.255.255.1

!--- Network commands announce prefix 10.0.10.0/24 and 10.1.0.0/24.

neighbor 10.10.13.1 remote-as 100

!--- Establishes peering with R1

neighbor 10.10.13.1 send-community

!--- Without this command, the community attributes are not sent to the neighbor

neighbor 10.10.13.1 route-map Peer-R1 out

!--- Configures outbound policy as defined by route-map "Peer-R1" when peering with R1

neighbor 10.10.23.2 remote-as 100

!--- Establishes peering with R2
```

```
neighbor 10.10.23.2 send-community

!--- Configures to send community attribute to R2

neighbor 10.10.23.2 route-map Peer-R2 out

!--- Configures outbound policy as defined by
!--- route-map "Peer-R2" when peering with R2.

no auto-summary
!
ip classless
ip bgp-community new-format

!--- Allows you to configure the BGP community
!--- attribute in AA:NN format.

!
access-list 101 permit ip host 10.0.10.0 host 255.255.255.0
access-list 102 permit ip host 10.1.1.0 host 255.255.255.1
!
!
route-map Peer-R1 permit 10
match ip address 101
set community 100:300

!--- Sets community 100:300 for routes matching access-list 101

!
route-map Peer-R1 permit 20
match ip address 102
set community 100:250

!--- Sets community 100:250 for routes matching access-list 102

!
route-map Peer-R2 permit 10
match ip address 101
set community 100:250

!--- Sets community 100:250 for routes matching access-list 101

!
route-map Peer-R2 permit 20
match ip address 102
set community 100:300

!--- Sets community 100:300 for routes matching access-list 102

!
end
```

R1

```
hostname R1
!
interface Loopback0
 ip address 200.200.10.1 255.255.255.0
!
interface Serial8/0
 ip address 10.10.13.1 255.255.255.1

!--- Connected to R3

!
interface Serial10/0
 ip address 10.10.12.1 255.255.255.0

!--- Connected to R2

!
router bgp 100
 no synchronization
 bgp log-neighbor-changes
 neighbor 10.10.12.2 remote-as 100

!--- Establishes peering with R2

neighbor 10.10.12.2 next-hop-self
neighbor 10.10.13.3 remote-as 30

!--- Establishes peering with R3

neighbor 10.10.13.3 route-map Peer-R3 in

!--- Configures the inbound policy as defined by route-map "Peer-R3" when peering with R3.

no auto-summary
!
ip bgp-community new-format

!--- Allows you to configure the BGP community attribute in AA:NN format.

ip community-list 1 permit 100:300
ip community-list 2 permit 100:250

!--- Defines community list 1 and 2.

!
route-map Peer-R3 permit 10
 match community 1
 set local-preference 130

!--- Sets local preference 130 for all routes matching community list 1.

!
route-map Peer-R3 permit 20
 match community 2
 set local-preference 125

!--- Sets local preference 125 for all routes matching community list 2.

!
```

```
route-map Peer-R3 permit 30
```

!--- Without this permit 30 statement, updates that do not match the permit 10 or permit 20 statements are dropped.

```
!  
end
```

R2

```
hostname R2  
!  
interface Loopback0  
ip address 10.0.10.0 255.255.255.0  
!  
interface Serial9/0  
ip address 10.10.23.2 255.255.255.1
```

!--- Connected to R3

```
!  
interface Serial10/0  
ip address 10.10.12.2 255.255.255.0
```

!--- Connected to R1

```
!  
router bgp 100  
no synchronization  
bgp log-neighbor-changes  
neighbor 10.10.12.1 remote-as 100
```

!--- Establishes iBGP peering with R1

```
neighbor 10.10.12.1 next-hop-self  
neighbor 10.10.23.3 remote-as 30
```

!--- Establishes peering with R3

```
neighbor 10.10.23.3 route-map Peer-R3 in
```

!--- Configures inbound policy as defined by route-map "Peer-R3" when peering with R3.

```
no auto-summary  
!  
ip bgp-community new-format
```

!--- Allows you to configure the BGP community attribute in AA:NN format.


```
!  
ip community-list 1 permit 100:300  
ip community-list 2 permit 100:250
```

!--- Defines community list 1 and 2.

```
!  
route-map Peer-R3 permit 10  
  match community 1  
  set local-preference 130  
  
!--- Sets local preference 130 for all routes matching community list 1.  
  
!  
route-map Peer-R3 permit 20  
  match community 2  
  set local-preference 125  
  
!--- Sets local preference 125 for all routes matching community list 2.  
  
!  
route-map Peer-R3 permit 30  
  
!--- Without this permit 30 statement, updates that do not match the permit 10 or permit 20 statements are dropped.  
  
!  
end
```

驗證

R1接收社群為100:300和100:250的字首10.0.10.0/24和10.1.0.0/24，如下圖所示 `show ip bgp` 命令輸出結果。

 注意：一旦根據配置的策略將這些路由安裝到BGP表中，則社群為100:300的字首會被分配本地優先順序130，社群為100:250的字首會被分配本地優先順序125。

```
<#root>
```

```
R1#
```

```
show ip bgp 10.0.10.0
```

```
BGP routing table entry for 10.0.10.0/24, version 2  
Paths: (1 available, best #1, table Default-IP-Routing-Table)  
  Advertised to non peer-group peers:  
    10.10.12.2  
    30  
    10.10.13.3 from 10.10.13.3 (10.0.10.0)  
      Origin IGP, metric 0, localpref 130, valid, external, best  
      Community: 100:300
```

```
!--- Prefix 10.0.10.0/24 with community 100:300 received from 10.10.13.3 (R3) is assigned local preference
```

```
<#root>
```

```
R1#
```

```
show ip bgp 10.1.0.0
```



```
BGP routing table entry for 10.1.0.0/24, version 4
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
    10.10.13.3
    30
    10.10.12.2 from 10.10.12.2 (10.1.0.0)
      Origin IGP, metric 0, localpref 130, valid, internal, best
```

```
!--- Received prefix 10.1.0.0/24 over iBGP from 10.10.12.2 (R2) with local preference 130
```

```
30 10.10.13.3 from 10.10.13.3 (198.50.100.0) Origin IGP, metric 0, localpref 125, valid, external Community: 100:250 !--- Prefix 10.1.0.0/24 with commu
```

<#root>


R1#

```
show ip bgp
```

```
BGP table version is 4, local router ID is 200.200.10.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.0.10.0/24	10.10.13.3	0	130	0	30 i
*>i 10.1.0.0/24	10.10.12.2	0	130	0	30 i
*	10.10.13.3	0	125	0	30 i

其 `show ip bgp` 命令，確認R1上選擇的最佳路徑具有本地優先順序(LocPrf)= 130。同樣，R2接收社群為100:250和100:300的字首10.0.10.0/24和10.1.0.0/24，如以下粗體所示 `show ip bgp` 命令輸出：

 注意：一旦這些路由安裝到BGP表中，根據配置的策略，將為community 100:300的字首分配本地優先順序130，為community 100:250的字首分配本地優先順序125。

<#root>

R2#

```
show ip bgp 10.0.10.0
```

```
BGP routing table entry for 10.0.10.0/24, version 2
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
    10.10.23.3
    30
    10.10.23.3 from 10.10.23.3 (10.0.10.0)
      Origin IGP, metric 0, localpref 125, valid, external
      Community: 100:250
```

```
!--- Prefix 10.0.10.0/24 with community 100:250 received from 10.10.23.3 (R3) is assigned local preference
```

```
30 10.10.12.1 from 10.10.12.1 (200.200.10.1) Origin IGP, metric 0, localpref 130, valid, internal, best !--- Received prefix 10.0.10.0/24 over iBGP from 10
```

<#root>

R2#

show ip bgp 10.1.0.0

```
BGP routing table entry for 10.1.0.0/24, version 3
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
    10.10.12.1
    30
    10.10.23.3 from 10.10.23.3 (10.1.0.0)
      Origin IGP, metric 0, localpref 130, valid, external, best
      Community: 100:300
```

!--- Prefix 10.1.0.0/24 with community 100:300 received from 10.10.23.3 (R3) is assigned local preferen

<#root>

R2#

show ip bgp

```
BGP table version is 3, local router ID is 192.168.50.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
* 10.0.10.0/24	10.10.23.3	0	125	0 30	i
*>i	10.10.12.1	0	130	0 30	i
*> 10.1.0.0/24	10.10.23.3	0	130	0 30	i

此 show ip bgp r2上的命令輸出確認R2上選擇的最佳路徑具有本地優先順序(localPrf)= 130。到達字首 10.0.10.0/24的IP路由優先於R1-R3鏈路從AS 100中退出，再從AS 30中退出。其 show ip route 命令 R1和R2確認此首選項。

<#root>

R1#

show ip route 10.0.10.0

```
Routing entry for 10.0.10.0/24
  Known via "bgp 100", distance 20, metric 0
  Tag 30, type external
  Last update from 10.10.13.3 3d21h ago
  Routing Descriptor Blocks:
  * 10.10.13.3, from 10.10.13.3, 3d21h ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
```

!--- On R1, the IP route to prefix 10.0.10.0/24 points to next hop 10.10.13.3 which is R3 serial 8/0 in

```
<#root>
```

```
R2#
```

```
show ip route 10.1.0.0
```

```
Routing entry for 10.1.0.0/24
  Known via "bgp 100", distance 200, metric 0
  Tag 30, type internal
  Last update from 10.10.12.1 3d21h ago
  Routing Descriptor Blocks:
  * 10.10.12.1, from 10.10.12.1, 3d21h ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
```

```
!--- On R2, IP route to prefix 10.1.0.0/24 points to next hop R1 (10.10.12.1) on its iBGP link
!--- Thus traffic to network 10.1.0.0/24 from R2 exits through R2-R1 and then R1-R3 link from AS 100 to
```

到達字首10.1.0.0/24的IP路由優先使用R2-R3鏈路從AS 100中退出，然後到達AS 30。其 `show ip route` 命令R1和R2確認此首選項。

```
<#root>
```

```
R2#
```

```
show ip route 10.1.0.0
```

```
Routing entry for 10.1.0.0/24
  Known via "bgp 100", distance 20, metric 0
  Tag 30, type external
  Last update from 10.10.23.3 3d22h ago
  Routing Descriptor Blocks:
  * 10.10.23.3, from 10.10.23.3, 3d22h ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
```

```
!--- On R2, IP route to prefix 10.1.0.0/24 points to next hop 10.10.23.3 which is R3 serial 9/0 interfa
```

```
<#root>
```

```
R1#
```

```
show ip route 10.1.0.0
```

```
Routing entry for 10.1.0.0/24
  Known via "bgp 100", distance 200, metric 0
  Tag 30, type internal
  Last update from 10.10.12.2 3d22h ago
  Routing Descriptor Blocks:
  * 10.10.12.2, from 10.10.12.2, 3d22h ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
```

```
!--- On R1, IP route to prefix 10.1.0.0/24 points to next hop R2 (10.10.12.2) on its iBGP link.
!--- Thus traffic to network 10.1.0.0/24 from R1 exits through R1-R2 and then R2-R3 link from AS 100 to
```

如果一個鏈路發生故障（例如R1-R3鏈路），所有流量都必須跟蹤R2-R3鏈路。如果關閉R1-R3之間的鏈路，則可以模擬此流量。

```
<#root>
```

```
R1#
```

```
configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
R1(config)#
```

```
interface serial8/0
```

```
R1(config-if)#
```

```
shut
```

```
R1(config-if)#
```

```
3d22h: %BGP-5-ADJCHANGE: neighbor 10.10.13.3 Down Interface flap
```

```
3d22h: %LINK-5-CHANGED: Interface Serial8/0, changed state to  
administratively down
```

```
3d22h: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial8/0,  
changed state to down
```

注意R1和R2上字首10.0.10.0/24和10.1.0.0/24的IP路由表。使用R2-R3鏈路退出AS 100。

```
<#root>
```

```
R1#
```

```
show ip route 10.0.10.0
```

```
Routing entry for 10.0.10.0/24
```

```
Known via "bgp 100", distance 200, metric 0
```

```
Tag 30, type internal
```

```
Last update from 10.10.12.2 00:01:47 ago
```

```
Routing Descriptor Blocks:
```

```
* 10.10.12.2, from 10.10.12.2, 00:01:47 ago
```

```
Route metric is 0, traffic share count is 1
```

```
AS Hops 1
```

```
<#root>
```

```
R1#
```

```
show ip route 10.1.0.0
```

```
Routing entry for 10.1.0.0/24
```

```
Known via "bgp 100", distance 200, metric 0
```

```
Tag 30, type internal
```

```
Last update from 10.10.12.2 3d22h ago
```

```
Routing Descriptor Blocks:
```

```
* 10.10.12.2, from 10.10.12.2, 3d22h ago
```

```
Route metric is 0, traffic share count is 1
AS Hops 1
```

此 `show` 命令輸出顯示，字首10.0.10.0/24和10.1.0.0/24的路由指向下一跳10.10.12.2(R2)，這是預期的。現在檢視R2上的IP路由表，檢查字首10.0.10.0/24和10.1.0.0/24的下一跳。下一跳必須為已配置策略的R3，才能成功工作。

```
<#root>
```

```
R2#
```

```
show ip route 10.0.10.0
```

```
Routing entry for 10.0.10.0/24
  Known via "bgp 100", distance 20, metric 0
  Tag 30, type external
  Last update from 10.10.23.3 00:04:10 ago
  Routing Descriptor Blocks:
  * 10.10.23.3, from 10.10.23.3, 00:04:10 ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
```

```
<#root>
```

```
R2#
```

```
show ip route 10.1.0.0
```

```
Routing entry for 10.1.0.0/24
  Known via "bgp 100", distance 20, metric 0
  Tag 30, type external
  Last update from 10.10.23.3 3d22h ago
  Routing Descriptor Blocks:
  * 10.10.23.3, from 10.10.23.3, 3d22h ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
```

下一跳10.10.23.3是R2-R3鏈路上的R3串列9/0介面。這確認配置的策略按預期工作。

相關資訊

- [RFC 1998](#)
- [疑難排解 BGP](#)
- [BGP：常見問題](#)
- [與BGP分擔負載](#)
- [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）。