

使用BGP社区值控制上行供应商网络的路由策略

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[背景理论](#)

[Conventions](#)

[Configure](#)

[控制路由策略](#)

[Network Diagram](#)

[配置](#)

[Verify](#)

[Related Information](#)

[Introduction](#)

本文档说明了如何使用边界网关协议 (BGP) 社区属性，以控制其上游服务提供商网络中的路由策略。

[Prerequisites](#)

[Requirements](#)

本文档要求对 BGP 路由协议及其操作有所了解。有关详细信息，请参阅 [BGP 案例分析](#)。

[Components Used](#)

This document is not restricted to specific software and hardware versions.但是，本文档中的信息基于以下软件和硬件版本：

- Cisco IOS软件版本12.2(27)
- Cisco 2500 系列路由器

本文档中的信息都是基于特定实验室环境中的设备创建的。All of the devices used in this document started with a cleared (default) configuration.如果您是在真实网络上操作，请确保您在使用任何命令前已经了解其潜在影响。

[背景理论](#)

尽管社区本身不会更改 [BGP 决策过程](#)，但社区可以用作标志以标记一组路由。然后，上游服务提

供商路由器便可使用这些标记在其网络中应用特定路由策略（例如，本地优先级）。

提供商会在客户可配置的社区值和提供商网络内对应的本地优先级值之间建立映射。其思路是，具有要求修改提供商网络中的 LOCAL_PREF 的特定策略的客户在其路由更新中设置对应的社区值。

社区是一组前缀，这些前缀共享某个公共属性，并且可以使用 BGP 社区属性进行配置。[BGP 社区属性](#)是一种长度可变的、可传递的可选属性。该属性由一组（四个）用于指定社区的八位组值组成。社区属性值用在前两个八位位组的自治系统(AS)编号编码，当依然是的两个八位位组定义的是由 AS。前缀可具有多个社区属性。检测到前缀中的多个社区属性的 BGP 扬声器可以基于一个、一些或所有属性进行操作。在路由器将社区属性传递到其他对等体之前，路由器可以添加或修改该属性。要了解有关社区属性的详细信息，请参阅 [BGP 案例分析](#)。

本地优先级属性指示 AS 以哪条路径作为首选来访问特定网络。如果有多条路径指向同一目标，将首选优先级较高的路径（本地优先级属性的默认值为 100）。有关详细信息，请参阅[本地优先级属性](#)。

Conventions

有关文档规则的详细信息，请参阅 [Cisco 技术提示规则](#)。

Configure

控制路由策略

本部分提供有关如何配置本文档所述功能的信息。

Note: 要查找本文档所用命令的其他信息，请使用[命令查找工具](#)（[仅限注册用户](#)）。

为进行简化，假设要在上游服务提供商 (AS 100) 和客户 (AS 30) 之间为社区属性和本地优先级属性建立以下映射。

本地首选	社区值
130	100:300
125	100:250

如果客户宣布社区属性等于 100:300 的前缀，则上游服务提供商会将这些路由的本地优先级设置为 130；如果社区属性等于 100:250，则设置为 125。

这将使您有可能可控制服务提供商网络中的路由策略（如果您更改向服务提供商宣布的前缀的社区值）。

在[网络图](#)中，客户 AS 30 希望通过社区属性实现此路由策略。

- 从 AS 100 发往网络 6.6.6.0/24 的入站流量通过 R1-R3 链路。如果 R1-R3 链路发生故障，所有流量将通过 R2-R3 传入。
- 从 AS 100 发往网络 7.7.7.0/24 的入站流量通过 R2-R3 链路。如果 R2-R3 链路发生故障，所有流量将通过 R1-R3 传入。

为了实现此路由策略，R3 按如下方式宣布其前缀：

到 R1：

- 6.6.6.0/24，社区属性为 100:300
- 7.7.7.0/24，社区属性为 100:250

到 R2：

- 6.6.6.0/24，社区属性为 100:250
- 7.7.7.0/24，社区属性为 100:300

当 BGP 邻居 R1 和 R2 从 R3 收到前缀后，R1 和 R2 便会基于社区属性和本地优先级属性之间的映射（如[上表](#)中所示）应用预先配置的策略，从而实现客户（AS 30）指定的路由策略。R1 在 BGP 表中安装以下前缀：

- 6.6.6.0/24，本地优先级为 130
- 7.7.7.0/24，本地优先级为 125


R2 在其 BGP 表中安装以下前缀：

- 6.6.6.0/24，本地优先级为 125
- 7.7.7.0/24，本地优先级为 130

由于在 BGP 路径选择标准中首选较高的本地优先级，因此将选择本地优先级为 130（130 大于 125）的路径作为 AS 100 中的最佳路径，并安装在 R1 和 R2 的 IP 路由表中。有关 BGP 路径选择标准的详细信息，请参阅[BGP 最佳路径选择算法](#)。

Network Diagram

本文档使用此图中所示的网络设置：



配置

本文档使用以下配置：

- [R3](#)
- [R1](#)
- [R2](#)

R3

```
Current configuration : 2037 bytes
!
version 12.2
!
hostname R3
!
interface Loopback0
 ip address 6.6.6.1 255.255.255.0
!
interface Ethernet0/0
 ip address 7.7.7.1 255.255.255.0
!
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 6.6.6.0 mask
255.255.255.0 network 7.7.7.0 mask 255.255.255.0 !---
Network commands announce prefix 6.6.6.0/24 !--- and
7.7.7.0/24. 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
6.6.6.0 host 255.255.255.0
access-list 102 permit ip host 7.7.7.0 host
255.255.255.0
!
!
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

```
Current configuration : 2037 bytes
!
version 12.2
!
hostname R3
!
interface Loopback0
 ip address 6.6.6.1 255.255.255.0
!
interface Ethernet0/0
 ip address 7.7.7.1 255.255.255.0
!
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 6.6.6.0 mask
255.255.255.0 network 7.7.7.0 mask 255.255.255.0 !---
Network commands announce prefix 6.6.6.0/24 !--- and
7.7.7.0/24. 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
6.6.6.0 host 255.255.255.0
access-list 102 permit ip host 7.7.7.0 host
255.255.255.0
!
!
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

```

R2

```

Current configuration : 2037 bytes
!
version 12.2
!
hostname R3
!
interface Loopback0
  ip address 6.6.6.1 255.255.255.0
!
interface Ethernet0/0
  ip address 7.7.7.1 255.255.255.0
!
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 6.6.6.0 mask
255.255.255.0 network 7.7.7.0 mask 255.255.255.0 !---
Network commands announce prefix 6.6.6.0/24 !--- and
7.7.7.0/24. 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
6.6.6.0 host 255.255.255.0
access-list 102 permit ip host 7.7.7.0 host
255.255.255.0
!
!
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

```

Verify

R1 将接收社区属性为 100:300 和 100:250 的前缀 6.6.6.0/24 和 7.7.7.0/24，如本部分的 **show ip bgp** 输出中的粗体所示。

Note: 将这些路由基于配置的策略安装到 BGP 表中后，系统会对社区属性为 100:300 的前缀分配本地优先级 130，对社区属性为 100:250 的前缀分配本地优先级 125。

```

R1# show ip bgp 6.6.6.0
BGP routing table entry for 6.6.6.0/24, 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 (6.6.6.1)
      Origin IGP, metric 0, localpref 130, valid, external, best
      Community: 100:300
!--- Prefix 6.6.6.0/24 with community 100:300 received from !--- 10.10.13.3 (R3) is assigned
local preference 130.

R1# show ip bgp 7.7.7.0
BGP routing table entry for 7.7.7.0/24, 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 (192.168.50.1)
      Origin IGP, metric 0, localpref 130, valid, internal, best
!--- Received prefix 7.7.7.0/24 over iBGP from 10.10.12.2 !--- (R2) with local preference 130.

    30
    10.10.13.3 from 10.10.13.3 (6.6.6.1)
      Origin IGP, metric 0, localpref 125, valid, external
      Community: 100:250
!--- Prefix 7.7.7.0/24 with community 100:250 received from !--- 10.10.13.3 (R3) is assigned

```

local preference 125.

```
R1# show ip bgp
BGP table version is 4, local router ID is 200.200.200.1
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
*> 6.6.6.0/24	10.10.13.3	0	130	0	30 i
*>i7.7.7.0/24	10.10.12.2	0	130	0	30 i
*	10.10.13.3	0	125	0	30 i

R1 上的 `show ip bgp` 命令将确认在 R1 上选择的最佳路径的本地优先级 (LocPrf) 为 130。

同样地，R2 也会接收社区属性为 100:250 和 100:300 的前缀 6.6.6.0/24 和 7.7.7.0/24，如本部分的 `show ip bgp` 命令输出中的粗体所示。

Note: 将这些路由基于配置的策略安装到 BGP 表中后，系统会对社区属性为 100:300 的前缀分配本地优先级 130，对社区属性为 100:250 的前缀分配本地优先级 125。

```
R2# show ip bgp 6.6.6.0
BGP routing table entry for 6.6.6.0/24, 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 (6.6.6.1)
      Origin IGP, metric 0, localpref 125, valid, external
      Community: 100:250
!--- Prefix 6.6.6.0/24 with community 100:250 received from !--- 10.10.23.3 (R3) is assigned local preference 125.
```

```
R2# show ip bgp 6.6.6.0
BGP routing table entry for 6.6.6.0/24, 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 (6.6.6.1)
      Origin IGP, metric 0, localpref 125, valid, external
      Community: 100:250
!--- Prefix 6.6.6.0/24 with community 100:250 received from !--- 10.10.23.3 (R3) is assigned local preference 125.
```

```
R2# show ip bgp 7.7.7.0
BGP routing table entry for 7.7.7.0/24, 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 (6.6.6.1)
      Origin IGP, metric 0, localpref 130, valid, external, best
      Community: 100:300
!--- Prefix 7.7.7.0/24 with community 100:300 received from !--- 10.10.23.3 (R3) is assigned local preference 130.
```

```
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
* 6.6.6.0/24	10.10.23.3	0	125	0	30 i
*>i	10.10.12.1	0	130	0	30 i
*> 7.7.7.0/24	10.10.23.3	0	130	0	30 i

R2 上的此 **show ip bgp** 命令输出将确认在 R2 上选择的最佳路径的本地优先级 (locPrf) 为 130。

指向前缀 6.6.6.0/24 的 IP 路由优先选择 R1-R3 链路退出朝向 AS 30 的 AS 100。R1 和 R2 上的 **show ip route** 命令可确认这一点。

```
R1# show ip route 6.6.6.0
Routing entry for 6.6.6.0/24
  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 6.6.6.0/24 points !--- to next hop 10.10.13.3 which is R3
  serial 8/0 !--- interface on the R1-R3 link.
```

```
R2# show ip route 6.6.6.0
Routing entry for 6.6.6.0/24
  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 6.6.6.0/24 points !--- to next hop R1 (10.10.12.1) on its iBGP
  link. !--- Thus traffic to network 6.6.6.0/24 from R2 !--- exits through R2-R1 and then R1-R3
  link from !--- AS 100 towards AS 30.
```

指向前缀 7.7.7.0/24 的 IP 路由优先选择 R2-R3 链路退出朝向 AS 30 的 AS 100。R1 和 R2 上的 **show ip route** 命令可确认这一点。

```
R2# show ip route 7.7.7.0
Routing entry for 7.7.7.0/24
  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 7.7.7.0/24 points !--- to next hop 10.10.23.3 which is R3 serial
  9/0 !--- interface on R2-R3 link.
```

```
R1# show ip route 7.7.7.0
Routing entry for 7.7.7.0/24
  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 7.7.7.0/24 points !--- to next hop R2 (10.10.12.2) on its iBGP
```

link. !--- Thus traffic to network 7.7.7.0/24 from R1 !--- exits through R1-R2 and then R2-R3 link !--- from AS 100 towards AS 30.

在一条链路 (如 R1-R3 链路) 出现故障的情况下 , 所有流量都必须经过 R2-R3 链路。如果关闭 R1-R3 之间的链路 , 则可以模拟这种情况。

```
R1# conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)#int s8/0
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
```

请注意 R1 和 R2 上的前缀 6.6.6.0/24 和 7.7.7.0/24 的 IP 路由表。使用 R2-R3 链路来退出 AS 100。

```
R1# show ip route 6.6.6.0
Routing entry for 6.6.6.0/24
  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
```

```
R1# show ip route 7.7.7.0
Routing entry for 7.7.7.0/24
  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** 命令输出显示到前缀 6.6.6.0/24 和 7.7.7.0/24 的路由指向下一跳 10.10.12.2 (R2) , 这符合预期。现在 , 请查看 R2 上的 IP 路由表以检查前缀 6.6.6.0/24 和 7.7.7.0/24 的下一跳。下一跳必须是 R3 才能成功运行所配置的策略。

```
R2# show ip route 6.6.6.0
Routing entry for 6.6.6.0/24
  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
```

```
R2# show ip route 7.7.7.0
Routing entry for 7.7.7.0/24
  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 链路上的 R3 系列 9/0 接口。这样便可确认所配置的策略是按预期方式运行的。

[Related Information](#)

- [RFC 1998](#)
- [排除BGP故障](#)
- [调试输出中显示“BGP:常见问题”](#)
- [负载分配的配置示例与在单宿主和多宿主环境的BGP](#)
- [BGP 支持页](#)
- [Technical Support & Documentation - Cisco Systems](#)