使用BGP社区值配置上游提供商网络

目录

[简介](#)

[先决条件](#)

- [要求](#)
- [使用的组件](#)

[背景信息](#)

[规则](#)

[配置和控制路由策略](#)

- [网络图](#)

[配置](#)

[R3](#)

[R1](#)

[R2](#)

[确认](#)

[相关信息](#)

简介

本文档介绍如何使用BGP社区值控制上游提供程序网络中的路由策略。

先决条件

要求

本文档需要了解边界网关协议(BGP)路由协议及其操作。

使用的组件

本文档不限于特定的软件和硬件版本。但是，本文档中的信息基于以下软件版本：

- Cisco IOS®软件版本12.2

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。如果您的网络处于活动状态，请确保您了解所有命令的潜在影响。

背景信息

虽然社区本身不会更改BGP最佳路径进程，但社区可用作标记一组路由的标志。然后，上游服务提供商路由器可以使用这些标志在其网络中应用特定路由策略（例如，本地优先级）。

提供商会在提供商网络内的可配置社区值和相应的本地首选项值之间进行映射。您可以具有需要修改提供商网络集中的LOCAL_PREF的特定策略及其路由更新中的相应社区值。

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。社区属性值使用自治系统(AS)编号在前两个二进制八位数中编码，其他两个二进制八位数由AS定义。前缀可具有多个社区属性。看到前缀中的多个社区属性的BGP发言人可以基于一个、部分或全部属性进行操作。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。

本地优先级属性是指AS首选哪条路径到达特定网络。当有多条路径到达同一个目标时，会选择优先级较高的路径（本地优先级属性的默认值为100）。 有关更多信息，请参阅案例研究。

规则

有关文档规则的详细信息，请参阅 Cisco 技术提示规则。

配置和控制路由策略

 注意：要寻找关于用于本文的指令的其他信息，请使用命令查找工具.

为简化起见，假定在上游服务提供商(AS 100)和您的设备(AS 30)之间建立了社区属性和本地首选项属性映射。

本地首选	社区值
130	100:300
125	100:250

如果通告前缀的社区属性等于100:300，则如果社区属性等于100:250，上游服务提供商将这些路由的本地优先级设置为130和125。

如果您更改向服务提供商通告的前缀的团体值，则这让您能够控制服务提供商网络内的路由策略。

在网络图中，AS 30希望将此路由策略与社区属性配合使用。

- 从AS 100到网络10.0.10.0/24的入站流量通过R1-R3链路传输。如果R1-R3链路发生故障，所有流量将通过R2-R3传输。
- 从AS 100到网络10.1.0.0/24的入站流量通过R2-R3链路传输。如果R2-R3链路发生故障，所有流量将通过R1-R3传输。

为实现此路由策略，R3通告其前缀的方式如下：

到 R1：

- 10.0.10.0/24，社区属性为 100:300
- 10.1.0.0/24，社区属性为 100:250

到 R2：

- 10.0.10.0/24，社区属性为 100:250
- 10.1.0.0/24，社区属性为 100:300

一旦BGP邻居R1和R2收到来自R3的前缀，R1和R2就会根据团体与本地首选项属性之间的映射应用配置的策略（如上表所示），从而获得您指定的路由策略(AS 30)。R1 在 BGP 表中安装以下前缀：

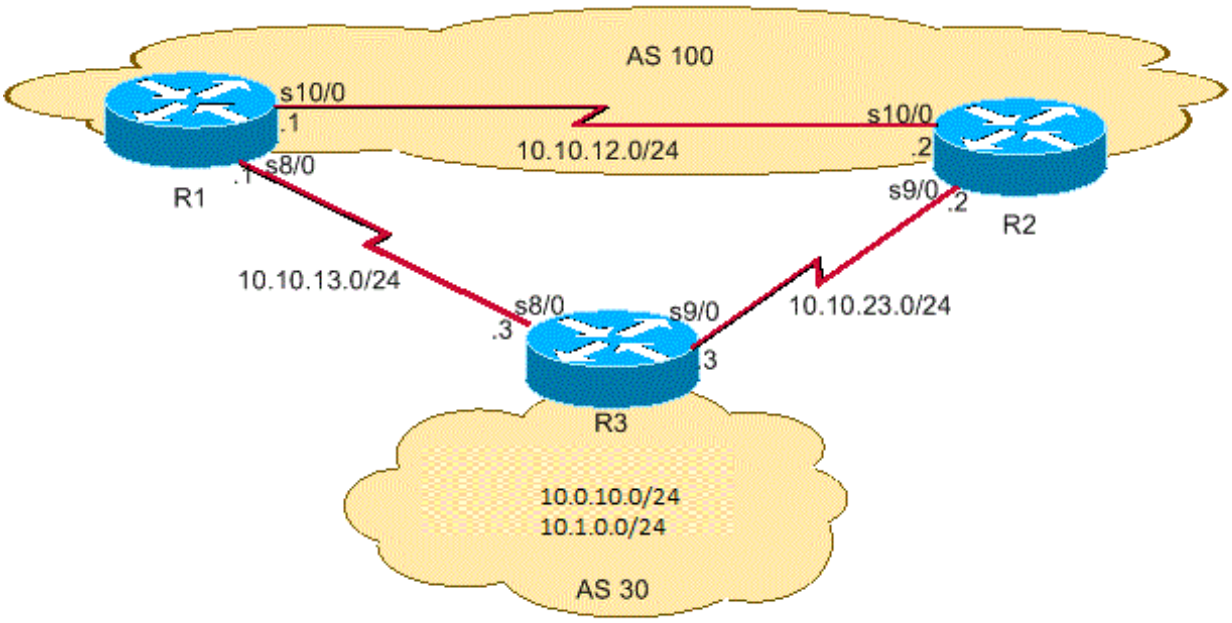
- 10.0.10.0/24，本地优先级为 130
- 10.1.0.0/24，本地优先级为 125

R2 在其 BGP 表中安装以下前缀：

- 10.0.10.0/24，本地优先级为 125
- 10.1.0.0/24，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，因此将本地优先级为130（130大于125）的路径选为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。

网络图



BGP网络

配置

本文档使用以下配置：

- 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 10.1.0.0 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 10.200.10.1 255.255.255.0
!
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
```

确认

R1接收社区为100:300和100:250的前缀10.0.10.0/24和10.1.0.0/24，如下一命令输出结 show ip bgp 果所示。

> 注意：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。

<#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 prefere
```

<#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 (198.51.100.1)
   Origin IGP, metric 0, localpref 125, valid, external
   Community: 100:250

!--- Prefix 10.1.0.0/24 with community 100:250 received from 10.10.13.3 (R3) is assigned local preference 125.
```

<#root>

R1#

show ip bgp

```
BGP table version is 4, local router ID is 10.200.10.1
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
```

show ip bgp R1上的命令确认R1上选择的最佳路径具有本地优先级(LoclPrf)= 130。同样，R2接收社区为100:250和100:300的前缀10.0.10.0/24和10.1.0.0/24，如下面的命令输出粗体所 show ip bgp 示：



注意：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。

<#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 prefere

 30
  10.10.12.1 from 10.10.12.1 (10.200.10.1)
   Origin IGP, metric 0, localpref 130, valid, internal, best

!--- Received prefix 10.0.10.0/24 over iBGP from 10.10.12.1 (R1) with local preference 130
```

<#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
```

show ip bgp R2上的此命令输出确认在R2上选择的最佳路径具有本地优先级(loclPrf)= 130。到前缀10.0.10.0/24的IP路由使R1-R3链路从AS 100退出，指向AS 30。在R1和R2上的命令确认此优 show ip route 先。

<#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
```

到前缀10.1.0.0/24的IP路由使R2-R3链路优先从AS 100退出，再到AS 30。 show ip route R1和R2上的命令可确认此首选项。

<#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
```

如果一条链路（例如R1-R3链路）发生故障，所有流量都必须跟踪R2-R3链路。如果关闭R1-R3之间的链路，则可以模拟此流量。

```
<#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上前缀10.0.10.0/24和10.1.0.0/24的IP路由表。使用R2-R3链路以退出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 命令输出显示到前缀10.0.10.0/24和10.1.0.0/24的路由指向预期中的下一跳10.10.12.2(R2)。现在，查看R2上的IP路由表，检查前缀10.0.10.0/24和10.1.0.0/24的下一跳。对于已配置的策略，下一跳必须是R3才能成功运行。

<#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 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。

相关信息

- RFC 1998
- 排除 BGP 故障
- BGP ：常见问题解答
- 与BGP共享负载
- 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）。