

IPv6 BGP路由反射器配置示例

目录

[简介](#)

[先决条件](#)

[要求](#)

[使用的组件](#)

[规则](#)

[配置](#)

[网络图](#)

[示例配置](#)

[验证](#)

[相关信息](#)

简介

本文档提供使用IPv6的示例配置，帮助您了解边界网关协议(BGP)中的路由反射器(RR)功能。默认情况下，从iBGP对等体接收的路由不会发送到另一个iBGP对等体，除非在AS内的所有BGP路由器之间形成全网状配置。这会导致可扩展性问题。使用BGP路由反射器可实现更高级别的可扩展性。

配置路由反射器允许路由器向其他iBGP扬声器通告或反映iBGP获知的路由。当使用neighbor route-reflector-client命令和[命令点是该RR客户端的邻居配置](#)时，路由器称为路由反射器。

先决条件

要求

尝试进行此配置之前，请确保满足以下要求：

- 了解 BGP 路由协议及其操作
- 了解 IPv6 编址方案

使用的组件

本文档不限于特定的软件和硬件版本。

本文档中的配置基于Cisco 3700系列路由器和Cisco IOS®^{软件}版本12.4(15)T1。

规则

有关文档规则的详细信息，请参阅 [Cisco 技术提示规则](#)。

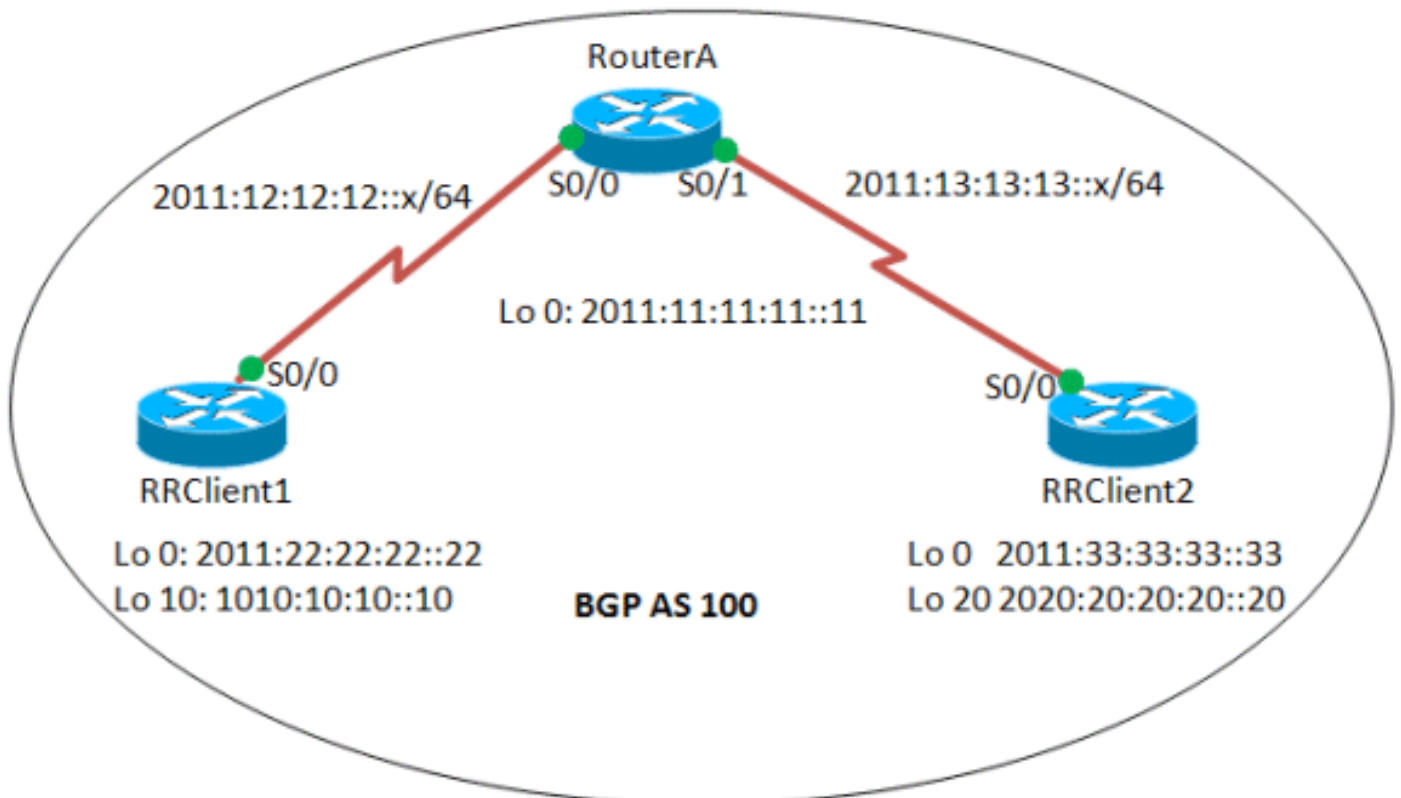
配置

在本例中，路由器A配置为RR，而路由器RRClient1和RRClient2是路由器A的客户端。所有路由器都配置为AS 100，但路由器没有全网状配置。相反，它使用BGP RR功能来相互通信。

注意：使用[命令查找工具](#)(仅限注册客户)可查找有关本文档中使用的命令的详细信息。

网络图

本文档使用以下网络设置：



示例配置

本文档使用以下配置：

- [Router A](#)
- [RRClient1](#)
- [RRClient2](#)

Router A

```
hostname Router-A
!
ip cef
!
ipv6 unicast-routing
!
interface Loopback0
 no ip address
 ipv6 address 2011:11:11:11::11/128
```

```

ipv6 ospf 10 area 0
!
interface Serial0/0
no ip address
ipv6 address 2011:12:12:12::1/64
ipv6 ospf 10 area 0
clock rate 2000000
!
interface Serial0/1
no ip address
ipv6 address 2011:13:13:13::1/64
ipv6 ospf 10 area 0
clock rate 2000000
!
router bgp 100
bgp router-id 1.1.1.1
no bgp default ipv4-unicast
bgp log-neighbor-changes
neighbor 2011:22:22:22::22 remote-as 100
neighbor 2011:22:22:22::22 update-source Loopback0
neighbor 2011:33:33:33::33 remote-as 100
neighbor 2011:33:33:33::33 update-source Loopback0
!
address-family ipv6
neighbor 2011:22:22:22::22 activate
neighbor 2011:22:22:22::22 route-reflector-client
!--- Configures the router RRClient1 as route reflector
client! neighbor 2011:33:33:33::33 activate neighbor
2011:33:33:33::33 route-reflector-client !--- Configures
the router RRClient2 as route reflector client! exit-
address-family ! ip forward-protocol nd ! ipv6 router
ospf 10 router-id 1.1.1.1 !--- Router ID of the route
reflector router A! log-adjacency-changes ! end

```

RRClient1

```

hostname RR-Client1
!
ip cef
!
ipv6 unicast-routing
!
interface Loopback0
no ip address
ipv6 address 2011:22:22:22::22/128
ipv6 ospf 10 area 0
!
interface Loopback10
no ip address
ipv6 address 1010:10:10:10::10/128
!
interface Serial0/0
no ip address
ipv6 address 2011:12:12:12::2/64
ipv6 ospf 10 area 0
clock rate 2000000
!
router bgp 100
bgp router-id 2.2.2.2
!--- Router ID of the RRClient1 no bgp default ipv4-
unicast bgp log-neighbor-changes neighbor
2011:11:11:11::11 remote-as 100 neighbor
2011:11:11:11::11 update-source Loopback0 ! address-

```

```
family ipv6 neighbor 2011:11:11:11::11 activate network
1010:10:10:10::10/128 exit-address-family ! ! ip
forward-protocol nd ! ipv6 router ospf 10 router-id
2.2.2.2 log-adjacency-changes ! ! end
```

RRClient2

```
hostname RR-Client2
!
ip cef
!
no ip domain lookup
ipv6 unicast-routing
!
!
interface Loopback0
no ip address
ipv6 address 2011:33:33:33::33/128
ipv6 ospf 10 area 0
!
interface Loopback20
no ip address
ipv6 address 2020:20:20:20::20/128
!
interface Serial0/0
no ip address
ipv6 address 2011:13:13:13::2/64
ipv6 ospf 10 area 0
clock rate 2000000
!
router bgp 100
bgp router-id 3.3.3.3
!--- Router ID of the RRClient2 no bgp default ipv4-
unicast bgp log-neighbor-changes neighbor
2011:11:11:11::11 remote-as 100 neighbor
2011:11:11:11::11 update-source Loopback0 ! address-
family ipv6 neighbor 2011:11:11:11::11 activate network
2020:20:20:20::20/128 exit-address-family ! ip forward-
protocol nd ! ipv6 router ospf 10 router-id 3.3.3.3 log-
adjacency-changes ! end
```

验证

使用本部分可确认配置能否正常运行。

[命令输出解释程序 \(仅限注册用户 \) \(OIT\) 支持某些 show 命令。](#) 使用 OIT 可查看对 show 命令输出的分析。

以下show命令用于验证配置：

- [show ipv6 route bgp](#)
- [show bgp ipv6 unicast](#)

在RR客户端中：

```
show ipv6 route bgp
```

在RRClient1中

```
RRClient1#sh ipv6 route bgp
```

```

IPv6 Routing Table - 9 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B
- BGP
    U - Per-user Static route, M - MIPv6
    I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea,
IS - ISIS summary
    O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext
1, OE2 - OSPF ext 2
    ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
    D - EIGRP, EX - EIGRP external
B 2020:20:20:20::20/128 [200/0]
via 2011:33:33:33::33
!--- The iBGP route from RRClient2 is reflected
RRClient1#ping 2011:33:33:33::33
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2011:33:33:33::33,
timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip
min/avg/max = 16/24/32 ms
!--- Ping to the RRClient2 from RRClient1 is successful
在RRClient2中

RRClient2#sh ipv6 route bgp
IPv6 Routing Table - 9 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B
- BGP
    U - Per-user Static route, M - MIPv6
    I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea,
IS - ISIS summary
    O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext
1, OE2 - OSPF ext 2
    ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
    D - EIGRP, EX - EIGRP external
B 1010:10:10:10::10/128 [200/0]
via 2011:22:22:22::22
!--- The iBGP route from RRClient1 is reflected
RRClient2#ping 1010:10:10:10::10
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1010:10:10:10::10,
timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip
min/avg/max = 40/50/76 ms
!--- Ping to the RRClient1 from RRClient2 is successful

```

在路由器A中：

```

show bgp ipv6 unicast IPv6 prefix

RouterA#sh bgp ipv6 unicast 1010:10:10:10::10/128
BGP routing table entry for 1010:10:10:10::10/128,
version 3
Paths: (1 available, best #1, table Global-IPv6-Table)
  Advertised to update-groups:
    1
Local, (Received from a RR-client)
!--- Indicates that the route was received from a route-
reflector client router RRClient1 2011:22:22:22::22
(metric 64) from 2011:22:22:22::22 (2.2.2.2) Origin IGP,

```

```
metric 0, localpref 100, valid, internal, best

RouterA#show bgp ipv6 unicast 2020:20:20:20::20/128
BGP routing table entry for 2020:20:20:20::20/128,
version 2
Paths: (1 available, best #1, table Global-IPv6-Table)
  Advertised to update-groups:
    1
  Local, (Received from a RR-client)
  !--- Indicates that the route was received from a route-
  reflector client router RRClient2 2011:33:33:33::33
  (metric 64) from 2011:33:33:33::33 (3.3.3.3) Origin IGP,
  metric 0, localpref 100, valid, internal, best
```

在RR客户端中：

每当反射iBGP路由（即传播到另一个iBGP对等体）时，实施路由反射器的路由器（在本例中为路由器A）会附加2个非传递属性：

- **呼叫方 ID:**这是非传递的可选BGP属性。反射的iBGP路由将具有从其接收路由的iBGP对等体的路由器ID作为其发起方ID。在本例中，从RRClient 2到RRClient1的路由2020:20:20:20::20/128由路由器A(RR)反映。因此，此路由将具有RRClient2(路由器ID:3.3.3.3)作为其发起方ID。
- **集群ID:**这是非传递的可选BGP属性。如果未配置cluster-id值，则反射的iBGP路由将RR的路由器ID作为其Cluster-ID。在本例中，路由2020:20:20:20::20/128未配置集群ID，因此路由器ID(路由器ID:路由器A(RR)的1.1.1.1)将是集群ID。

```
show bgp ipv6 unicast ipv6-prefix

在RRClient1中

show bgp ipv6 unicast 2020:20:20:20::20/128
BGP routing table entry for 2020:20:20:20::20/128,
version 3
Paths: (1 available, best #1, table Global-IPv6-Table)
  Not advertised to any peer
  Local
    2011:33:33:33::33 (metric 128) from
2011:11:11:11::11 (1.1.1.1)
    Origin IGP, metric 0, localpref 100, valid,
internal, best
    Originator: 3.3.3.3, Cluster list: 1.1.1.1
  !--- Originator ID 3.3.3.3 is the router id of the
  RRClient2 from which the route is received! !---
  Similarly, Cluster ID 1.1.1.1 is the router ID of the
  router reflector Router A!

在RRClient2中

show bgp ipv6 unicast 1010:10:10:10::10/128
BGP routing table entry for 1010:10:10:10::10/128,
version 3
Paths: (1 available, best #1, table Global-IPv6-Table)
  Not advertised to any peer
  Local
    2011:22:22:22::22 (metric 128) from
2011:11:11:11::11 (1.1.1.1)
    Origin IGP, metric 0, localpref 100, valid,
internal, best
    Originator: 2.2.2.2, Cluster list: 1.1.1.1
  !--- Originator ID 2.2.2.2 is the router ID of the
```

```
RRClient1 from which the route is received! !---  
Similarly, Cluster ID 1.1.1.1 is the router ID of the  
router reflector Router A!
```

相关信息

- [BGP 支持页](#)
- [IP 版本 6 支持页面](#)
- [BGP 案例分析](#)
- [技术支持和文档 - Cisco Systems](#)