使用IPv6 BGP配置IPV6远程触发黑洞

目录

简介

本文档介绍IPV6远程触发黑洞(RTBH)的行为。 它显示了使用路由映射故意使IPv6流量黑洞的场景。

先决条件

要求

Cisco 建议您了解以下主题：

- IPv6
- 边界网关协议 (BGP)

使用的组件

本文档中的信息基于Cisco IOS软件版本15.4。

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。如果您使用的是真实网络，请确保您已经了解所有命令的潜在影响。

背景信息

RTBH过滤是一种通常用于防止拒绝服务(DoS)攻击的技术。DoS攻击的一个常见问题是网络被大量有害/恶意流量泛洪。这会导致链路阻塞和其他问题，如CPU使用率过高等。这会清除合法流量，对网络造成严重影响。

根据RFC 2545，如果且仅当BGP发言者与“下一跳的网络地址”字段中承载的全局IPv6地址和路由通告到的对等体所标识的实体共享一个公共子网时，本地链路地址才应包括在“下一跳”字段中。在所有其他情况下，BGP发言者应在Network Address字段中仅向其对等体通告下一跳的全局IPv6地址。

它基本上意味着，如果您在直连子网上具有IPv6 EBGP邻居关系，则它将链路本地IP和全局IPv6地址作为下一跳。但是，命令请求(RFC)并未指定首选哪个。思科首选链路本地地址，因为它发送数据包时始终是最短距离。使用RTBH时，可能会出现问题，本文档将说明如何处理它。

配置

本文档以用例来解释RTBH的行为和使用的命令。

网络图

本图像用作本文档其余部分的示例拓扑。

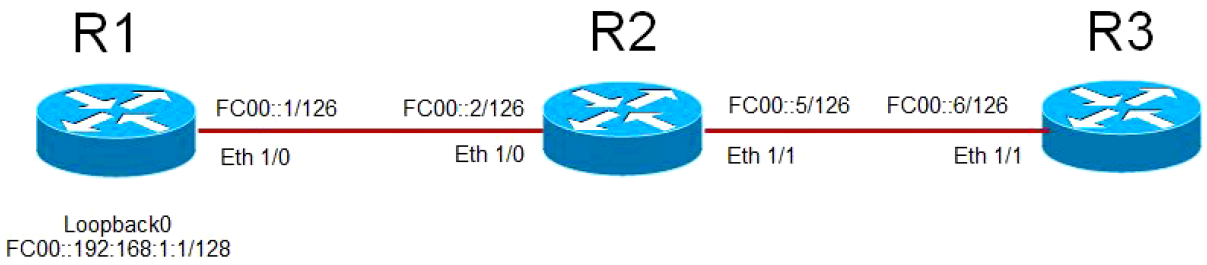

- R1与R2有EBGP邻居关系，R2与R3有EBGP邻居关系。
- 路由器R1通过BGP将其环回0(FC00::192:168:1:1/128)通告给R2，而R2将其通告给R3。
- R3使用路由映射将R1的环回前缀的下一跳设置为指向路由表中“NULL 0”的虚拟IPv6地址。

相关配置

此配置用于不同的路由器以模拟将使用RTBH的情况：

R1

```
interface Ethernet1/0
 no ip address
 ipv6 address FC00::1/126
end
!
interface Loopback0
 ip address 192.168.1.1 255.255.255.0
 ipv6 address FC00::192:168:1:1/128
 !
 router bgp 65500
 bgp router-id 192.168.1.1
 bgp log-neighbor-changes
 neighbor FC00::2 remote-as 65501
 !
 address-family ipv6
```

```
network FC00::/126
 network FC00::192:168:1:1/128
 neighbor FC00::2 activate
```

R2

```
interface Ethernet1/0
 no ip address
 ipv6 address FC00::2/126
end
!
interface Ethernet1/1
 no ip address
 ipv6 address FC00::5/126
!
router bgp 65501
 bgp router-id 192.168.1.2
 bgp log-neighbor-changes
 neighbor FC00::1 remote-as 65500
 neighbor FC00::6 remote-as 65502
 !
 address-family ipv6
 network FC00::/126
 network FC00::4/126
 neighbor FC00::1 activate
 neighbor FC00::6 activate
```

R3

```
interface Ethernet1/1
 no ip address
 ipv6 address FC00::6/126
end
!
ipv6 prefix-list BLACKHOLE-PREFIX seq 5 permit FC00::192:168:1:1/128
!
route-map BLACKHOLE-PBR permit 10
 match ipv6 address prefix-list BLACKHOLE-PREFIX
 set ipv6 next-hop FC00::192:168:1:3
route-map BLACKHOLE-PBR permit 20
!
router bgp 65502
 bgp router-id 192.168.1.3
 bgp log-neighbor-changes
 neighbor FC00::5 remote-as 65501
!
 address-family ipv6
 network FC00::4/126
 neighbor FC00::5 activate
 neighbor FC00::5 route-map BLACKHOLE-PBR in
```

验证

测试用例1

当R3上未配置基于策略的路由(PBR)时，在路由表中，到R3上R1的环回的路由指向R2的链路本地地址FE80::A8BB:CCFF:FE00:A211。

```
BGP Configuration

router bgp 65502
 bgp router-id 192.168.1.3
 bgp log-neighbor-changes
 neighbor FC00::5 remote-as 65501
 !
 address-family ipv6
 network FC00::4/126
 neighbor FC00::5 activate


BGP has both next-hops.

R3#show bgp ipv6 unicast FC00::192:168:1:1/128
BGP routing table entry for FC00::192:168:1:1/128, version 4
Paths: (1 available, best #1, table default)
 Not advertised to any peer
 Refresh Epoch 1
 65501 65500
   FC00::5 (FE80::A8BB:CCFF:FE00:A211) from FC00::5 (192.168.1.2)
     Origin IGP, localpref 100, valid, external, best
     rx pathid: 0, tx pathid: 0x0


Routing Table has Link Local address as the next-hop.

R3#show ipv6 route FC00::192:168:1:1
Routing entry for FC00::192:168:1:1/128
 Known via "bgp 65502", distance 20, metric 0, type external
 Route count is 1/1, share count 0
 Routing paths:
   FE80::A8BB:CCFF:FE00:A211, Ethernet1/1
     MPLS label: nolabel
     Last updated 00:02:45 ago


Destination is reachable

R3#ping ipv6 FC00::192:168:1:1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FC00::192:168:1:1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

测试用例2

当在R3上使用路由映射BLACKHOLE-PBR配置PBR时，观察到对于FC00::192:168:1:1/128（R1的环回），路由表中的下一跳仍然指向R2的链路本地地址FE80::A8BB:CCFF:FE00:A211。因此，流量从不黑洞，而是使用本地链路地址路由。

```
BGP Configuration

ipv6 prefix-list BLACKHOLE-PREFIX seq 5 permit FC00::192:168:1:1/128
!
route-map BLACKHOLE-PBR permit 10
 match ipv6 address prefix-list BLACKHOLE-PREFIX
 set ipv6 next-hop FC00::192:168:1:3
```

```
!
route-map BLACKHOLE-PBR permit 20
!
router bgp 65502
 bgp router-id 192.168.1.3
 bgp log-neighbor-changes
 neighbor FC00::5 remote-as 65501
 !
 address-family ipv4
 no neighbor FC00::5 activate
 exit-address-family
 !
 address-family ipv6
 network FC00::4/126
 neighbor FC00::5 activate
 neighbor FC00::5 route-map BLACKHOLE-PBR in
```

Next-hop in BGP changes to the one defined in route-map.

```
R3#show bgp ipv6 unicast FC00::192:168:1:1/128
BGP routing table entry for FC00::192:168:1:1/128, version 4
Paths: (1 available, best #1, table default)
 Not advertised to any peer
 Refresh Epoch 1
 65501 65500
   FC00::192:168:1:3 (FE80::A8BB:CCFF:FE00:A211) from FC00::5 (192.168.1.2)
     Origin IGP, localpref 100, valid, external, best
     rx pathid: 0, tx pathid: 0x0
```

New next-hop is not reachable and points to Null 0

```
R3#show ipv6 route FC00::192:168:1:3
Routing entry for FC00::192:168:1:3/128
 Known via "static", distance 1, metric 0
 Route count is 1/1, share count 0
 Routing paths:
   directly connected via Null0
     Last updated 00:19:23 ago
```

Routing table still uses Link Local address as next-hop.

```
R3#show ipv6 route FC00::192:168:1:1
Routing entry for FC00::192:168:1:1/128
 Known via "bgp 65502", distance 20, metric 0, type external
 Route count is 1/1, share count 0
 Routing paths:
FE80::A8BB:CCFF:FE00:A211, Ethernet1/1
     MPLS label: nolabel
     Last updated 00:00:41 ago
```

Destination is still reachable.

```
R3#ping ipv6 FC00::192:168:1:1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FC00::192:168:1:1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

测试用例3

为了克服此行为，请在R3上使用BGP邻居配置命令**disable-connected-check**。使用disable-connected-check来假设邻居的IPv6地址仅为一跳路。使用此命令的最常见情况是，在直连路由器的环回上建立EBGP邻居关系时。在这种情况下，该命令给人一种印象，即路由器正在构建EBGP邻居关系，并且不在公共子网中。邻居关系可以跨环回，因此，路由器在通告不传送链路本地地址但只传送全局IPv6地址的前缀时，会通告该前缀。

添加此命令后，您可以在R3的路由表中看到R1的环回**192:168:1:1/128**的路由，该路由根据FC00::192:**168:1:3的路由映射指向下一跳**。 现在，自**FC00::192:168:1:3有指**向Null 0的路由，因此流量会黑洞。

```
BGP Configuration

ipv6 prefix-list BLACKHOLE-PREFIX seq 5 permit FC00::192:168:1:1/128
!
route-map BLACKHOLE-PBR permit 10
 match ipv6 address prefix-list BLACKHOLE-PREFIX
 set ipv6 next-hop FC00::192:168:1:3
!
route-map BLACKHOLE-PBR permit 20
!
router bgp 65502
 bgp router-id 192.168.1.3
 bgp log-neighbor-changes
 neighbor FC00::5 remote-as 65501
neighbor FC00::5 disable-connected-check
 !
 address-family ipv4
 no neighbor FC00::5 activate
 exit-address-family
 !
 address-family ipv6
 network FC00::4/126
 neighbor FC00::5 activate
 neighbor FC00::5 route-map BLACKHOLE-PBR in


Next-hop in BGP changes to the one defined in route-map. There is no Link Local Address.

R3#show bgp ipv6 unicast FC00::192:168:1:1/128
BGP routing table entry for FC00::192:168:1:1/128, version 4
Paths: (1 available, best #1, table default)
 Not advertised to any peer
 Refresh Epoch 1
 65501 65500
   FC00::192:168:1:3 from FC00::5 (192.168.1.2)
     Origin IGP, localpref 100, valid, external, best
     rx pathid: 0, tx pathid: 0x0


Routing table uses the new next-hop.

R3#show ipv6 route FC00::192:168:1:1
Routing entry for FC00::192:168:1:1/128
 Known via "bgp 65502", distance 20, metric 0, type external
```

```
 Route count is 1/1, share count 0
 Routing paths:
FC00::192:168:1:3
     MPLS label: nolabel
     Last updated 00:00:37 ago


New next-hop is pointed to Null 0. Traffic will be dropped.

R3#show ipv6 route FC00::192:168:1:3
Routing entry for FC00::192:168:1:3/128
 Known via "static", distance 1, metric 0
 Route count is 1/1, share count 0
 Routing paths:
   directly connected via Null 0
     Last updated 02:18:03 ago


Destination is not reachable

R3#ping ipv6 FC00::192:168:1:1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to FC00::192:168:1:1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

注意： CSCuv60686disable-connected-check

故障排除

本文档目前没有特定的故障排除信息。