配置有IPv6 BGP的IPV6远程被触发的黑洞

目录

简介

本文描述在IPV6远程被触发的黑洞看到的行为(RTBH)。它显示方案使用路由映射的地方故意地黑色被钻孔的IPv6流量。

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

本文档中的信息根据Cisco IOS软件版本15.4版本。

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过滤是通常被使用的技术防止服务拒绝(DoS)攻击。在DOS攻击看到的常见问题是网络被巨大的音量不需要/恶意流量充斥。这导致链路堵塞和其他问题类似高CPU等。这使合法数据流挨饿并且导致在网络的严重的暗示。

根据RFC 2545，链路本地地址在下跳次字段将包括，如果，并且，只有当BGP扬声器共享与全局IPv6地址识别的实体的一普通的子网输入了下跳次字段网络地址，并且对等体路由通告。所有其他的案例BGP扬声器将通告给其对等体在网络地址地址字段下一跳的仅全局IPv6地址。

基本意味，如果直接地有在连接的子网的一IPv6 EBGP邻居关系，然后运载林克本地IP以及全局IPv6地址作为下一跳。然而，要求命令(RFC)不指定应该更喜欢哪个。思科更喜欢链路本地地址，因为，当发送数据包时它总是最短的距离。当您使用RTBH时，它可能是问题，并且本文解释如何处理它。

配置

本文采取用例解释用于的行为和命令获得RTBH工作。

网络图

此镜像使用，拓扑示例其余本文。

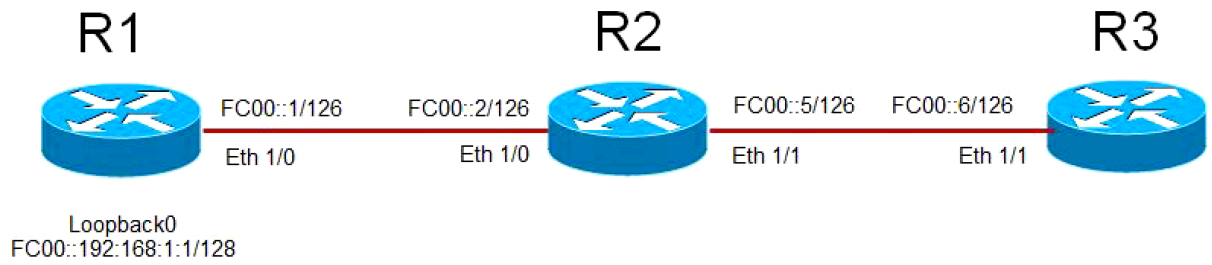

- R1有与R2和R2的EBGP邻居关系有与R3的EBGP邻居关系。
- 路由器R1通过BGP通告其loopback0 (FC00::192:168:1:1/128)对R2和R2通告它对R3。
- R3使用一route-map设置R1的环回前缀的下一跳为一个假的IPv6地址对“NULL 0"的该点在路由表里。

相关配置

此配置在另外路由器用于模拟将使用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

测试案例1

当没有时在R3配置的基于策略的路由(PBR)，在路由表里，路由对在R3的R1的环回指向R2's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
```

测试案例2

当有使用在R3时的route-map配置的PBR BLACKHOLE-PBR，注意到对于FC00::192:168:1:1/128 (R1的环回)， next-hop in路由表仍然指向R2's链路本地地址FE80::A8BB:CCFF:FE00:A211。所以，流量从未是黑色钻孔和路由使用链路本地地址。

```
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
```

测试案例3

为了解决此行为，请使用BGP邻居在R3的配置命令禁用连接**检查**。禁用连接检查用于假设，邻居的IPv6地址只是一种跳方式。此命令使用的最普通的scernario，当EBGP邻居关系在连接的路由器的时环回直接地被建立。在这种情况下，命令有印象路由器建立EBGP邻居关系并且不是在普通的子网。结邻可能是在环回间并且，路由器，当通告不运载链路本地地址，然而仅全局IPv6地址的前缀时。

一旦此命令被添加，您在R3里路由表能为R1的环回**192:168:1:1/128**看到路由，对是**FC00::192:168:1:3的**下一跳的点在符合route-map。 现在，因为**FC00::192:168:1:3**有指向Null0的一个路由，因此，黑色被钻孔的流量。

```
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)
```

Note:CSCuv60686route-map

故障排除

当前没有特定故障检修信息可用为本文。