使用IPv6 BGP配置IPV6遠端觸發黑洞

目錄

簡介

本檔案將說明IPV6遠端觸發的黑洞(RTBH)中出現的行為。 它顯示了這樣一個場景：IPv6流量有意使用路由對映進行黑洞。

必要條件

需求

思科建議您瞭解以下主題：

- IPv6
- 邊界閘道通訊協定(BGP)

採用元件

本檔案中的資訊是根據Cisco IOS軟體版本15.4。

本文中的資訊是根據特定實驗室環境內的裝置所建立。文中使用到的所有裝置皆從已清除（預設）的組態來啟動。如果您的網路正在作用，請確保您已瞭解任何指令可能造成的影響。

背景資訊

RTBH過濾技術通常用於防止拒絕服務(DoS)攻擊。DoS攻擊的一個常見問題是網路被大量不需要/惡意的流量淹沒。這會導致鏈路阻塞和其它問題，如高CPU等。這會導致合法流量不足並對網路造成嚴重影響。

根據RFC 2545，當BGP發言者與由「下一跳網路地址」欄位中攜帶的全域性IPv6地址標識的實體以及路由被通告到的對等體共用一個公共子網時，本地鏈路地址應包含在「下一跳」欄位中。在所有其他情況下，BGP發言人應只在Network Address欄位中向自己的同伴通告下一跳的全域性IPv6地址。

這基本上意味著，如果在直接連線的子網中存在IPv6 EBGP鄰居關係，則它會攜帶鏈路本地IP以及作為下一跳的全域性IPv6地址。但是，命令請求(RFC)未指定應首選哪一個。思科優先使用連結本地位址，因為當傳送封包時，它的距離總是最短。使用RTBH時，可能會出現問題，本文檔將介紹如何處理該問題。

設定

本文檔使用案例來解釋行為和用於使RTBH工作的命令。

網路圖表

此圖用作本文檔其餘部分的示例拓撲。

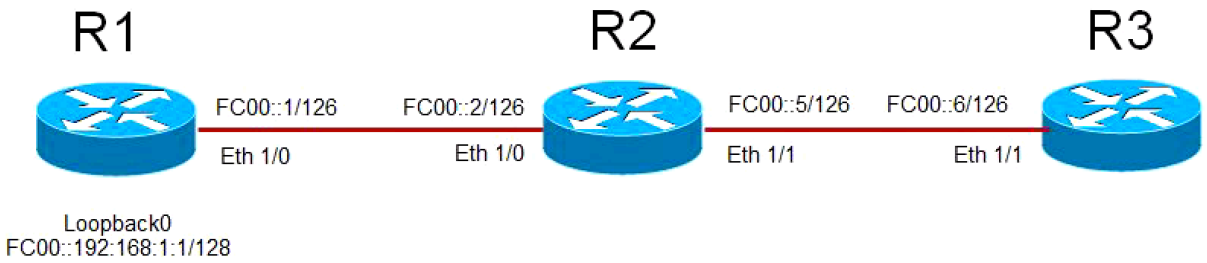

- R1與R2具有EBGP鄰居關係，R2與R3具有EBGP鄰居關係。
- 路由器R1通過BGP向R2通告其環回0(FC00::192:168:1:1/128),R2將其通告給R3。
- R3使用路由對映將R1的環回字首的下一跳設定為路由表中指向「NULL 0」的虛擬IPv6地址。

相關配置

不同路由器上使用以下配置來模擬使用RTBH的情況：

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
```

驗證

測試用例1

當R3上沒有配置基於策略的路由(PBR)時，在路由表中，到R1在R3上的環回的路由指向R2的鏈路本地地址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
```

測試用例2

當R3上使用路由對映BLACKHOLE-PBR配置PBR時，可以觀察到，對於FC00::192:168:1:1/128（R1的環回），路由表中的下一跳仍指向R2的鏈路本地地址FE80::A8BB:CFF:FE00:A211。因此，流量從未被黑遮蔽，而是使用鏈路本地地址進行路由。

```
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
```

測試用例3

為了克服此行為，請在R3上使用BGP neighbor配置命令**disable-connected-check**。Disable-connected-check用於假設鄰居的IPv6地址只有單跳方式（例如IPv6）。使用此命令的最常見情況是在直連路由器的環回上建立EBGP鄰居關係時。在這種情況下，該命令會產生這樣的印象：路由器正在建立EBGP鄰居關係，並且不在公共子網中。鄰居關係可以跨越環回，因此當路由器通告字首時，該字首不攜帶鏈路本地地址，而只攜帶全域性IPv6地址。

新增此命令後，您可以看到R1的環回**192:168:1:1/128的路由在R3的路由表**中，該路由根據路由對映(即**FC00::192:168:1:3)指向下一跳**。 現在，由於**FC00::192:168:1:3**具有指向Null 0的路由，因此，流量是黑洞。

```
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)
```

附註： CSCuv60686**disable-connected-check**

疑難排解

目前尚無適用於本文的特定疑難排解資訊。