採用iBGP作為PE到CE路由協定的VRF間路由洩漏

目錄

簡介

本文討論客戶邊緣(CE)和提供商邊緣(PE)運行內部BGP(iBGP)協定時，VRF間路由洩漏。討論了帶路由洩漏的電流限制及其解決辦法。

必要條件

需求

思科建議您瞭解BGP的基本知識。

採用元件

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。

設定

先前不支援將iBGP作為PE到CE協定。但是，現在已融入了iBGP，並且iBGP也可以被視為PE到CE路由的潛在候選者。此功能使客戶可以在所有站點中擁有單個自治系統。為實現此目的，引入了一個新屬性ATTR_SET ，它以透明的方式在服務提供商網路中承載VPN BGP屬性。此外，它需要將PE作為路由反射器與CE路由器進行iBGP會話。新引入的命令「neighbor x.x.x.x internal vpn-

client」可幫助實現這一目標。當配置此單個命令時，它會自動配置「neighbor x.x.x.x route-reflector-client」和「neighbor x.x.x.x next-hop-self」。

網路圖表

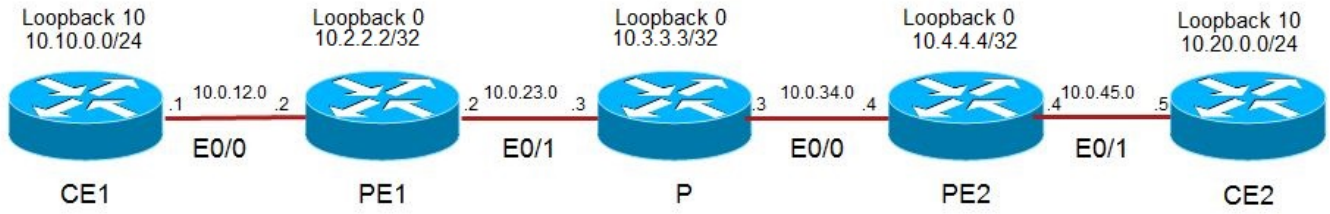

組態

CE1

```
interface Loopback10
ip address 10.10.0.1 255.255.255.0

interface Ethernet0/0
 ip address 10.0.12.1 255.255.255.0

router bgp 100
 bgp router-id 10.1.1.1
 bgp log-neighbor-changes
 neighbor 10.0.12.2 remote-as 100
 !
 address-family ipv4
 network 10.10.0.0 mask 255.255.255.0
 neighbor 10.0.12.2 activate
 exit-address-family
```

CE2

```
interface Loopback10
 ip address 10.20.0.1 255.255.255.0

interface Ethernet0/1
 ip address 10.0.45.5 255.255.255.0

router bgp 100
 bgp router-id 10.5.5.5
 bgp log-neighbor-changes
 neighbor 10.0.45.4 remote-as 100
 !
 address-family ipv4
  network 10.20.0.0 mask 255.255.255.0
  neighbor 10.0.45.4 activate
 exit-address-family
```

PE1

```
vrf definition A
 rd 10:10
 route-target export 100:100
 route-target import 100:100

!
 address-family ipv4
 exit-address-family
!
vrf definition B
 rd 20:20
 !
 address-family ipv4
 route-target import 50:50
 route-target import 100:100
 exit-address-family

interface Loopback0
 ip address 10.2.2.2 255.255.255.255
 ip ospf 100 area 0
!
interface Ethernet0/0
 vrf forwarding A
 ip address 10.0.12.2 255.255.255.0
!
interface Ethernet0/1
 ip address 10.0.23.2 255.255.255.0
 mpls ip

router bgp 100
 bgp router-id 10.2.2.2
 bgp log-neighbor-changes
 neighbor 10.4.4.4 remote-as 100
 neighbor 10.4.4.4 update-source Loopback0
 !
 address-family vpnv4
 neighbor 10.4.4.4 activate
 neighbor 10.4.4.4 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf A
 neighbor 10.0.12.1 remote-as 100
 neighbor 10.0.12.1 activate

 neighbor 10.0.12.1 internal-vpn-client // needed to exchange routes between PEs
 neighbor 10.0.12.1 next-hop-self
 exit-address-family
 !
 address-family ipv4 vrf B
 exit-address-family
```

PE2

```
vrf definition A
 rd 10:10
 route-target export 100:100
```

```
 route-target import 100:100

!
 address-family ipv4
 exit-address-family

interface Loopback0
 ip address 10.4.4.4 255.255.255.255
 ip ospf 100 area 0
!
interface Ethernet0/0
 ip address 10.0.34.4 255.255.255.0
 mpls ip
!
interface Ethernet0/1
 vrf forwarding A
 ip address 10.0.45.4 255.255.255.0

router bgp 100
 bgp router-id 10.4.4.4
 bgp log-neighbor-changes
 neighbor 10.2.2.2 remote-as 100
 neighbor 10.2.2.2 update-source Loopback0
 !
 address-family vpnv4
 neighbor 10.2.2.2 activate
 neighbor 10.2.2.2 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf A
 neighbor 10.0.45.5 remote-as 100
 neighbor 10.0.45.5 activate
 neighbor 10.0.45.5 internal-vpn-client //needed to exchange routes between PEs
 neighbor 10.0.45.5 route-reflector-client
 neighbor 10.0.45.5 next-hop-self
 exit-address-family
```

驗證

案例1:通過MP-BGP接受和交換客戶路由

如前所述，iBGP as PE to CE要求使用「neighbor x.x.x internal vpn-client」命令在VRF內部配置與客戶的BGP對等。 如果沒有此命令，本地PE在VRF中接受來自本地CE的路由，但是這些客戶路由不會通過MP-BGP與其他PR路由器共用。以下是已預配置「neighbor x.x.x.x internal vpn-client」的輸出。

以下輸出顯示了PE1和PE2上vrf A中的路由。

```
PE1#show ip route vrf A

Routing Table: A
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
```

```
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C        10.0.12.0/24 is directly connected, Ethernet0/0
L        10.0.12.2/32 is directly connected, Ethernet0/0
B        10.10.0.0/24 [200/0] via 10.0.12.1, 00:35:23
B        10.20.0.0/24 [200/0] via 10.4.4.4, 00:40:55
```

PE2#show ip route vrf A

```
Routing Table: A
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      a - application route
      + - replicated route, % - next hop override

Gateway of last resort is not set

     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C        10.0.45.0/24 is directly connected, Ethernet0/1
L        10.0.45.4/32 is directly connected, Ethernet0/1
B        10.10.0.0/24 [200/0] via 10.2.2.2, 00:00:08
B        10.20.0.0/24 [200/0] via 10.0.45.5, 00:41:55
```

CE1#show ip route bgp

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      a - application route
      + - replicated route, % - next hop override

Gateway of last resort is not set

     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
B        10.20.0.0/24 [200/0] via 10.0.12.2, 00:03:56
```

CE2#show ip route bgp

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      a - application route
      + - replicated route, % - next hop override

Gateway of last resort is not set
```

```
     10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
B        10.10.0.0/24 [200/0] via 10.0.45.4, 00:04:21
```

案例2:從一個VRF洩漏路由到另一個VRF。

案例1 ，成功演示了CE1和CE2之間的路由交換。現在考慮另一個vrf B ，需要將路由安裝到vrf A中。常規方法是使用VRF A中的export-map值並在VRF B中匯入相同的值，如下所示。

```
!
vrf definition A
 rd 10:10
 route-target export 100:100
 route-target import 100:100
 !
 address-family ipv4
 exit-address-family
!
vrf definition B
 rd 20:20

!
 address-family ipv4
 route-target import 100:100
 exit-address-family
!
```

完成上述配置後，VRF B無法安裝從本地CE接收的任何BGP路由。但是，通過MP-BGP從其他PE接收的路由已成功安裝，如下面的輸出所示。10.20.0.0/24屬於CE，它在VRF A中成功接收並匯出到VRF B。但是，從CE1本地接收的10.10.0.0/24無法進入VRF B。

```
PE1#show ip route vrf A bgp
Routing Table: A
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

     10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
B        10.10.0.0/24 [200/0] via 10.0.12.1, 00:12:35
B        10.20.0.0/24 [200/0] via 10.4.4.4, 00:54:22


PE1#show ip route vrf B

Routing Table: B
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
```

```
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      a - application route
      + - replicated route, % - next hop override

Gateway of last resort is not set

     10.0.0.0/24 is subnetted, 1 subnets
B        10.20.0.0 [200/0] via 10.4.4.4, 00:46:38
```

neighbor x.x.x.x internal vpn-clientVRF ABCEVRFPE1VRF BCE110.10.0.0/24

```
!
router bgp 100
 address-family ipv4 vrf A
 no  neighbor 10.0.12.1 internal-vpn-client
!
```

PE1#show ip route vrf B bgp

```
Routing Table: B
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
      D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
      ia - IS-IS inter area, * - candidate default, U - per-user static route
      o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
      a - application route
      + - replicated route, % - next hop override

Gateway of last resort is not set

     10.0.0.0/24 is subnetted, 2 subnets
B        10.10.0.0 [200/0] via 10.0.12.1 (A), 00:00:11
B        10.20.0.0 [200/0] via 10.4.4.4, 00:58:33
```

BAx.x.x.xvpn-client

PE2#show ip route vrf A bgp

```
Routing Table: A
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
B        10.20.0.0/24 [200/0] via 10.0.45.5, 01:04:21          // 10.10.0.0/24 is missing.
```

這是一個限制，且已安裝CSCuw43489增強錯誤來修正此問題。

因應措施

有一個解決方法可用於檢查上述問題。此解決方法允許在有命令「neighbor x.x.x.x internal vpn-client」的情況下將路由從VRF A匯入到VRF B。 從客戶匯入路由時，此解決方法需要設定虛擬社群（在下面的示例中為50:50）。將此虛擬擴展社群匯入vrf B。

```
!
route-map TEST, permit, sequence 10
  Match clauses:
  Set clauses:
    extended community RT:50:50
  Policy routing matches: 0 packets, 0 bytes
!
vrf definition B
 rd 20:20
 address-family ipv4
 route-target import 100:100
 route-target import 50:50                          // match dummy community
!
router bgp 100
 address-family ipv4 vrf A
 neighbor 10.0.12.1 route-map TEST in              // Set dummy community
!
```

PE1#show bgp vpnv4 uni vrf B 10.10.0.0

```
BGP routing table entry for 20:20:10.10.0.0/24, version 4
Paths: (1 available, best #1, table B)
 Not advertised to any peer
 Refresh Epoch 1
 Local, (Received from ibgp-pece RR-client), imported path from 10:10:10.10.0.0/24 (A)
   10.0.12.1 (via vrf A) (via A) from 10.0.12.1 (10.1.1.1)
     Origin IGP, metric 0, localpref 100, valid, internal, best
     Extended Community: RT:50:50
     rx pathid: 0, tx pathid: 0x0
```

PE1#show ip route vrf B

```
Routing Table: B
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

     10.0.0.0/24 is subnetted, 2 subnets
B        10.10.0.0 [200/0] via 10.0.12.1 (A), 00:00:25
B        10.20.0.0 [200/0] via 10.4.4.4, 00:00:25
```

如上所述，此解決方法使路由10.10.0.0/24在VRF A中安裝在VRF B中。