使用iBGP作为PE到CE路由协议的VRF间路由泄漏

目录

简介

本文讨论当客户边缘(CE)和提供商边缘(PE)运行内部BGP(iBGP)协议时，VRF间路由泄漏。讨论了当前路由泄漏的限制和解决方法。

先决条件

要求

建议本文的读者具备 BGP 的基础知识。

使用的组件

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。

配置

之前不支持将iBGP作为PE到CE协议。但是，现在已经纳入了此功能，iBGP也可以视为PE到CE路由的潜在候选。此功能允许客户在所有站点上拥有一个自治系统。为此，引入了新属性ATTR_SET，以透明方式跨服务提供商网络传输VPN BGP属性。此外，它还需要使PE成为与CE路由器的iBGP会话的路由反射器。新引入的命令neighbor x.x.x.x internal vpn-client有助于实现此目标

。当配置此单个命令时，它会自动配置“neighbor x.x.x.x route-reflector-client”和“neighbor x.x.x.x next-hop-self”。

网络图

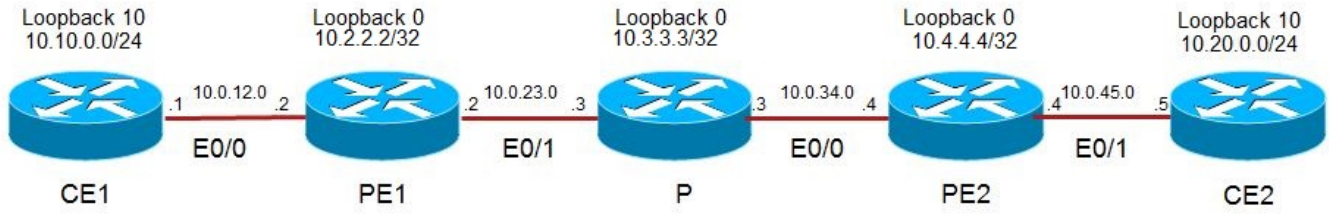

配置

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
```

验证

第 1 种情况：通过MP-BGP接受和交换客户路由

如前所述，iBGP作为PE到CE需要使用命令“neighbor x.x.x.x internal vpn-client”配置与VRF内部客户的BGP对等。 如果没有此命令，本地PE将接受来自VRF中本地CE的路由，但这些客户路由不通过MP-BGP与其他PR路由器共享。以下输出已采用“neighbor x.x.x.x internal vpn-client”预配置。

以下输出显示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
```

第 2 种情况：将路由从一个VRF泄露到另一个VRF。

案例1，成功演示了CE1和CE2之间的路由交换。现在考虑另一个vrf B，它需要将vrf A中的路由安装到自身中。常规方法是在VRF A中使用export-map值，在VRF B中导入相同值，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
!
```

完成上述配置后，VRF B无法安装从本地CE接收的任何BGP路由。但是，通过MP-BGP从其他PE接收的路由已成功安装，如下所示在输出中。10.20.0.0/24属于CE，在VRF A中成功接收并且也导出到VRF B。但本地从CE1接收的10.10.0.0/24无法进入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-client”VRF 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

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.
```

这是一个限制，已经提出了CSCuw43489增强版来解决此问题。

解决方法

有一种解决方法可用于检查上述问题。此解决方法允许在出现命令“neighbor x.x.x.x internal vpn-

client”时将路由从VRF A导入VRF B。 此解决方法要求在从客户导入路由时设置虚拟社区（以下示例中为50:50）。将此虚拟扩展社区导入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
```

如上所示，此解决方法使路由10.10.0.0/24出现在VRF A的VRF B中。