与iBGP的相互VRF路由泄漏作为对CE路由协议的PE

目录

简介

本文讨论相互VRF路由泄漏，当用户边缘(CE)和服务商边缘运行内部BGP (iBGP)时协议。它与路由泄漏讨论当前限制和它的一应急方案。

先决条件

要求

思科建议您有BGP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

作为对CE协议的PE不支持iBGP的支持前。然而，这当前合并，并且iBGP可以也考虑作为PE的潜在候选者对CE路由。此功能允许客户有在整个场地间的一单个自治系统。要达到运载在间服务提供商网络的VPN BGP属性以透明方式的此新团体ATTR_SET介绍。并且，它要求做PE作为iBGP会话的路由反射器有CE路由器的。新引入命令“邻接x.x.x.x内部vpn-client”帮助达到此。当此单个命令配

置时，自动地配置“邻接x.x.x.x route-reflector-client”和“邻接x.x.x.x next-hop-self”。

网络图

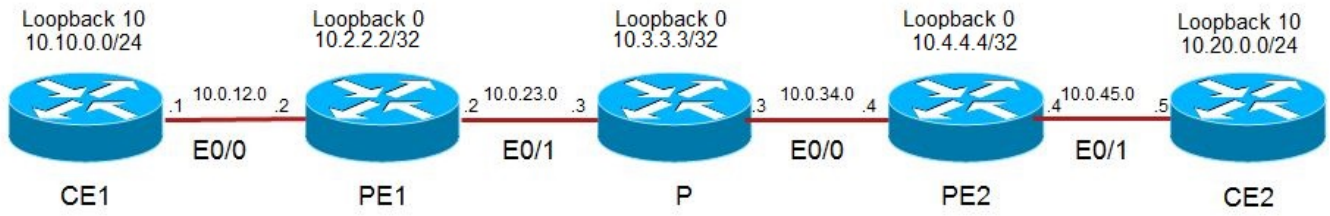

配置

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：接受和交换在MP-BGP的客户路由

如讨论前，iBGP作为PE对CE要求并列与客户的BGP的配置在与命令“邻接x.x.x.x内部vpn-client的”VRF里面。在没有此命令时，本地PE接受从本地CE的路由在VRF，然而这些客户路由没有通过MP-BGP共享用其他PRS路由器。在输出之下用“邻接”获得预先配置的x.x.x.x内部vpn-client。

在输出之下显示在VRF A的路由在PE1及PE2。

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：从一个VRF的漏的路由到另一个。

案例1，路由顺利地被展示的交换在CE1与CE2之间的。现在请考虑需要安装在VRF A的路由到本身的另一VRF B。正常方法将使用export map值在VRF A和导入在VRF B的同样值如下所示。

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
```

当在配置上执行， VRF B不能安装从本地CE接收的其中任一BGP路由。然而，从其他观点扫描器接收的路由通过MP-BGP在输出中顺利地安装和显示下面。10.20.0.0/24属于CE，并且那在VRF A顺利地接收和也导出对VRF B。但是从CE1接收的本地10.10.0.0/24不能输入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
```

“x.x.x.xvpn-client”VRF ACEVRFBPE1 VRF 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
```

BA(x.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.
```

这是限制，并且已经归档增强bug **CSCuw43489**解决此问题。

解决方法

有是可用在讨论问题上检查的应急方案。此应急方案允许到从VRF A的导入路由对VRF B在命令“邻接x.x.x.x内部vpn-client面前”。此应急方案要求设置一个假的社区(50:50完成在下面的示例)，当导入从客户时的路由。导入此假的扩展团体到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，此应急方案做在VRF A安装的路由10.10.0.0/24存在VRF B。