

Fuga de Ruta entre VRF con iBGP como protocolo de ruteo de PE a CE

Contenido

[Introducción](#)

[Prerequisites](#)

[Requirements](#)

[Componentes Utilizados](#)

[Configurar](#)

[Diagrama de la red](#)

[Configuraciones](#)

[CE1](#)

[CE2](#)

[PE1](#)

[PE2](#)

[Verificación](#)

[Caso 1: Aceptación e intercambio de rutas del cliente sobre MP-BGP](#)

[Caso 2: Fuga de rutas de un VRF a otro.](#)

[Solución Alternativa](#)

[Conversaciones relacionadas de la comunidad de soporte de Cisco](#)

Introducción

Este documento analiza la fuga de rutas entre VRF cuando el borde del cliente (CE) y el borde del proveedor (PE) ejecutan el protocolo BGP interno (iBGP). Describe la limitación actual con la fuga de rutas y una solución alternativa para ella también.

Prerequisites

Requirements

Cisco recomienda que tenga conocimiento básico de BGP.

Componentes Utilizados


The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Configurar

El soporte para iBGP como protocolo PE a CE no era soportado anteriormente. Sin embargo, esto se ha incorporado ahora y iBGP también puede considerarse como un candidato potencial para el

ruteo de PE a CE. Esta función permite a los clientes disponer de un único sistema autónomo en todos los sitios. Para lograr esto, se ha introducido un nuevo atributo ATTR_SET que transporta los atributos BGP VPN a través de la red del proveedor de servicios de una manera transparente. Además, requiere hacer el PE como reflector de ruta para la sesión iBGP con el router CE. El comando recientemente introducido " neighbor x.x.x.x internal vpn-client" ayuda a lograr esto. Cuando se configura este único comando, configura automáticamente "neighbor x.x.x.x route-reflector-client" y "neighbor x x.x.x.x next-hop-self".

Diagrama de la red



Configuraciones

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

```

Verificación

Caso 1: Aceptación e intercambio de rutas del cliente sobre MP-BGP

Como se ha mencionado anteriormente, iBGP como PE a CE requiere la configuración de peering BGP con el cliente dentro de VRF con el comando "neighbor x.x.x.x internal vpn-client". En ausencia de este comando, el PE local acepta las rutas desde el CE local en el VRF; sin embargo, estas rutas del cliente no se comparten a través de MP-BGP con otros routers PR. Los siguientes resultados se han tomado con "neighbor x.x.x.x internal vpn-client" preconfigurados.

A continuación se muestran las rutas en el vrf A en PE1 y 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
```

Caso 2: Fuga de rutas de un VRF a otro.

Caso 1, intercambio de rutas demostrado con éxito entre CE1 y CE2. Ahora considere otro vrf B que necesita instalar rutas en vrf A en sí mismo. El método regular es utilizar el valor export-map en el VRF A e importar el mismo valor en el VRF B como se muestra a continuación.

```
!
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
!
```

Cuando se realiza la configuración anterior, el VRF B no puede instalar ninguna de las rutas BGP que se recibieron del CE local. Sin embargo, las rutas recibidas de otros PE a través de MP-BGP se instalan correctamente como se muestra a continuación en la salida. 10.20.0.0/24 pertenece a CE y se recibe correctamente en VRF A y también se exporta a VRF B. Pero 10.10.0.0/24 recibido localmente de CE1 no puede ingresar el 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

Este problema de fuga de ruta VRF de la ruta CE local de VRF A a B sólo se ve hasta que se configura el punto "neighbor x.x.x.x internal vpn-client". Tan pronto como este comando se quita de PE1, el VRF B puede ver con éxito la ruta CE1 local 10.10.0.0/24, como se muestra a continuación.

```
!  
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

Y el sitio remoto B, deja de recibir las rutas del sitio A (como vecino x.x.x.x interno vpn-client ha sido eliminado).

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.

Esta es una limitación y un error de mejora **CSCuw43489** ya se ha registrado para solucionar este problema.

Solución Alternativa

Hay una solución alternativa disponible para comprobar el problema mencionado anteriormente. Esta solución temporal permite importar rutas del VRF A al VRF B en presencia del comando "neighbor x.x.x.x internal vpn-client". Esta solución temporal requiere establecer una comunidad ficticia (50:50 se realiza en el ejemplo siguiente) cuando se importan rutas del cliente. Importe esta comunidad extendida ficticia en 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

Como se muestra anteriormente, esta solución alternativa hace que la ruta 10.10.0.0/24 esté presente en la instalación de VRF A en el VRF B.