

路由目標約束

目錄

[簡介](#)

[路由目標約束的用途](#)

[不帶RTC的行為](#)

[RTC配置](#)

[PE配置](#)

[RR配置](#)

[RTC的行為](#)

[PE](#)

[RR](#)

[路由刷新處理](#)

[相關資訊](#)

簡介

本文描述一種機制，通過該機制，面向提供商邊緣(PE)路由器的VPNv4和VPNv6字首交換將減少到所需的最小數量。

路由目標約束的用途

使用多協定標籤交換(MPLS)VPN，內部邊界網關協定(iBGP)對等體或路由反射器(RR)將所有VPN4和/或VPN6字首傳送到PE路由器。PE路由器丟棄沒有匯入VPN路由和轉發(VRF)的VPN4/6字首。這種行為是RR將VPN4/6字首傳送到PE路由器，而它並不需要該字首。這是對RR和PE處理能力的浪費，也是對頻寬的浪費。

使用路由目標約束(RTC)時，RR僅將想要的VPN4/6字首傳送到PE。「Wanted」表示PE具有匯入特定字首的VRF。

RFC 4684指定RTC。支援通過適用於VPNv4和VPNv6的新地址系列rtfilter。

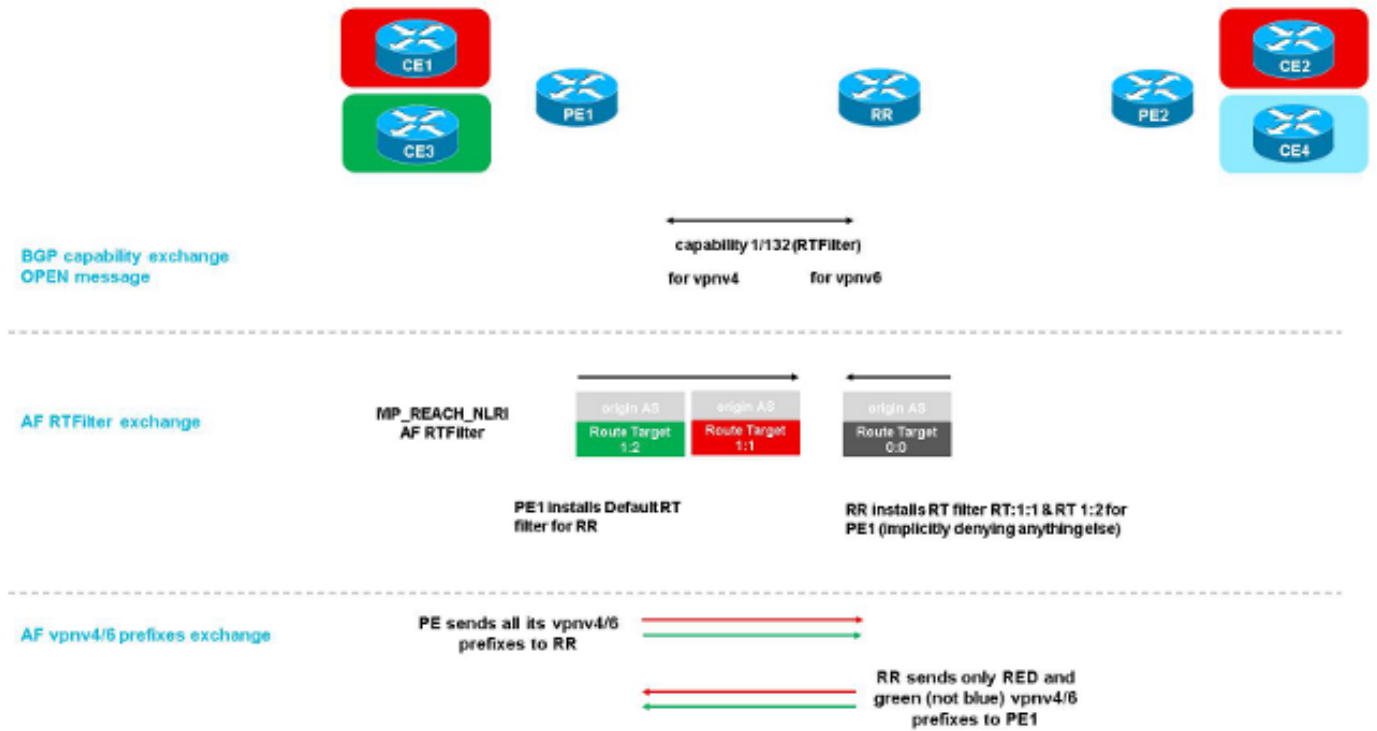
從VPN RT匯入清單中從PE路由器上的所有VRF獲取路由目標(RT)過濾資訊。PE路由器將此過濾資訊作為地址系列rtfilter中的BGP更新傳送到RR。此過濾資訊或RT成員資格在MP_REACH_NLRI和MP_UNREACH_NLRI屬性的網路層可達性資訊(NLRI)中編碼。

接收BGP對等體將此NLRI轉換為過濾器，並將此過濾器安裝在傳送對等體的出站。接收BGP對等體使用此過濾器來決定要傳送或不傳送哪個VPNv4/6字首，具體取決於連線的RT的存在。

若要使RTC正常運作，兩個BGP對等點都需要支援RTC。也就是說，RR和PE需要支援它。但是，部署可以是增量式，這意味著並非所有RR和PE路由器都需要一次支援它。RTC可以在網路中工

作，某些PE路由器支援它，而其它路由器不支援。在支援的路由器上，RTC將處於活動狀態。在尚不支援RTC的路由器上，通告將像以前一樣運行，而沒有RTC（因此沒有任何出站過濾）。

下圖顯示RTC的原理：



不帶RTC的行為

RR將所有VPN4/6字首傳送到PE。PE丟棄沒有匯入RT的PE。Debug BGP updates顯示丟棄的字首。消息「DENIED due to:提供了「不支援擴展社群」。

VPNv4單播的示例如下：

```
BGP(4): 10.100.1.3 rcvd UPDATE w/ att: nexthop 10.100.1.1, origin i, localpref 100,
metric 0, originator 10.100.1.1, clusterlist 10.100.1.3, merged path 65003,
AS_PATH , extended community RT:1:2
BGP(4): 10.100.1.3 rcvd 1:2:10.100.1.6/32, label 27 -- DENIED due to: extended
community not supported;
```

VPNv6單播的示例如下：

```
BGP(5): 10.100.1.3 rcvd UPDATE w/ attr: nexthop ::FFFF:10.100.1.1, origin i,
localpref 100, metric 0, originator 10.100.1.1, clusterlist 10.100.1.3,
merged path 65003, AS_PATH , extended community RT:1:2
BGP(5): 10.100.1.3 rcvd [1:2]2001:10:100:1::6/128, label 23 -- DENIED due to:
extended community not supported;
```

RTC配置

PE配置

```
vrf definition green
 rd 1:2
  route-target export 1:2
  route-target import 1:2
  !
  address-family ipv4
  exit-address-family
!
vrf definition red
 rd 1:1
  route-target export 1:1
  route-target import 1:1
  !
  address-family ipv4
  exit-address-family
  !
  address-family ipv6
  exit-address-family

router bgp 1
 bgp log-neighbor-changes
 neighbor 10.100.1.3 remote-as 1
 neighbor 10.100.1.3 update-source Loopback0
 neighbor 10.100.1.4 remote-as 1
 neighbor 10.100.1.4 update-source Loopback0
 !
 address-family vpnv4
 neighbor 10.100.1.3 activate
 neighbor 10.100.1.3 send-community both
 neighbor 10.100.1.4 activate
 neighbor 10.100.1.4 send-community both
 exit-address-family
 !
 address-family rtfiler unicast
  neighbor 10.100.1.3 activate
 neighbor 10.100.1.3 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf green
 neighbor 10.1.6.6 remote-as 65003
 neighbor 10.1.6.6 activate
 neighbor 10.1.6.6 send-community both
 exit-address-family
 !
 address-family ipv4 vrf red
 neighbor 10.1.5.5 remote-as 65001
 neighbor 10.1.5.5 activate
 neighbor 10.1.5.5 send-community both
 exit-address-family
```

RR配置

```
router bgp 1
```

```

bgp log-neighbor-changes
neighbor 10.100.1.1 remote-as 1
neighbor 10.100.1.1 update-source Loopback0
neighbor 10.100.1.2 remote-as 1
neighbor 10.100.1.2 update-source Loopback0
!
address-family vpnv4
neighbor 10.100.1.1 activate
neighbor 10.100.1.1 send-community both
neighbor 10.100.1.1 route-reflector-client
neighbor 10.100.1.2 activate
neighbor 10.100.1.2 send-community both
neighbor 10.100.1.2 route-reflector-client
exit-address-family
!
address-family rtfiler unicast
neighbor 10.100.1.1 activate
neighbor 10.100.1.1 send-community both
neighbor 10.100.1.1 route-reflector-client
neighbor 10.100.1.1 default-originate
exit-address-family

```

RTC的行為

當BGP對等建立時，對等體交換rtfilter的功能，即1/132（對於VPNv4和VPNv6）。

```

RR1# show bgp rtfiler unicast all neighbors 10.100.1.1
BGP neighbor is 10.100.1.1, remote AS 1, internal link
  BGP version 4, remote router ID 10.100.1.1
  BGP state = Established, up for 00:14:28
  Last read 00:00:01, last write 00:00:56, hold time is 180,
  keepalive interval is 60 seconds
  Neighbor sessions:
    1 active, is not multisession capable (disabled)
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: received
    Address family VPNv4 Unicast: advertised and received
    Address family VPNv6 Unicast: advertised and received
    Address family RT Filter: advertised and received
    Enhanced Refresh Capability: advertised and received
    Multisession Capability:
    Stateful switchover support enabled: NO for session 1
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

      Sent      Rcvd
  Opens:          1         1
  Notifications:  0         0
  Updates:        6         7
  Keepalives:     17        18
  Route Refresh:  0         0
  Total:          24        30
  Default minimum time between advertisement runs is 0 seconds

For address family: VPNv4 Unicast
Session: 10.100.1.1
BGP table version 65, neighbor version 65/0

```

Output queue size : 0
Index 19, Advertise bit 1
Route-Reflector Client
19 update-group member
RT Filter activate
Community attribute sent to this neighbor
Slow-peer detection is disabled
Slow-peer split-update-group dynamic is disabled

...

For address family: VPNv6 Unicast

Session: 10.100.1.1
BGP table version 5, neighbor version 5/0
Output queue size : 0
Index 3, Advertise bit 1
Route-Reflector Client
3 update-group member
RT Filter activate
Community attribute sent to this neighbor
Slow-peer detection is disabled
Slow-peer split-update-group dynamic is disabled

...

For address family: RT Filter

Session: 10.100.1.1
BGP table version 52, neighbor version 52/0
Output queue size : 0
Index 13, Advertise bit 0
Route-Reflector Client
13 update-group member
NEXT_HOP is always this router for eBGP paths
Community attribute sent to this neighbor
Default information originate, default sent
Slow-peer detection is disabled
Slow-peer split-update-group dynamic is disabled

	Sent	Rcvd
Prefix activity:	----	----
Prefixes Current:	1	2 (Consumes 160 bytes)
Prefixes Total:	1	2
Implicit Withdraw:	0	0
Explicit Withdraw:	0	0
Used as bestpath:	n/a	2
Used as multipath:	n/a	0

	Outbound	Inbound
Local Policy Denied Prefixes:	-----	-----
Bestpath from iBGP peer:	2	n/a
Total:	2	0

Number of NLRIs in the update sent: max 1, min 0
Last detected as dynamic slow peer: never
Dynamic slow peer recovered: never
Refresh Epoch: 1
Last Sent Refresh Start-of-rib: never
Last Sent Refresh End-of-rib: never
Last Received Refresh Start-of-rib: never
Last Received Refresh End-of-rib: never

	Sent	Rcvd
Refresh activity:	----	----
Refresh Start-of-RIB	0	0
Refresh End-of-RIB	0	0

Address tracking is enabled, the RIB does have a route to 10.100.1.1

```
Connections established 16; dropped 15
Last reset 00:14:28, due to Peer closed the session of session 1
Transport(tcp) path-mtu-discovery is enabled
Graceful-Restart is disabled
```

PE

```
debug bgp all
```

```
BGP: 10.100.1.3 active rcvd OPEN w/ optional parameter type 2 (Capability) len 6
BGP: 10.100.1.3 active OPEN has CAPABILITY code: 1, length 4
BGP: 10.100.1.3 active OPEN has MP_EXT CAP for afi/safi: 1/132
BGP: 10.100.1.3 accept RTC SAFI
```

```
PE1# show bgp rtfilter unicast rt 1:1
BGP routing table entry for 1:2:1:1, version 3
Paths: (1 available, best #1)
  Advertised to update-groups:
    13
  Refresh Epoch 1
  Local
    0.0.0.0 from 0.0.0.0 (10.100.1.1)
      Origin IGP, localpref 100, weight 32768, valid, sourced, local, best
      RT generation: import
      rx pathid: 0, tx pathid: 0x0
```

AF rtfilter還使用更新組 :

```
PE1# show bgp rtfilter unicast all update-group 13
BGP version 4 update-group 13, internal, Address Family: RT Filter
BGP Update version : 12/0, messages 0
Extended-community attribute sent to this neighbor
Topology: global, highest version: 12, tail marker: 12
Format state: Current working (OK, last not in list)
              Refresh blocked (not in list, last not in list)
Update messages formatted 1, replicated 1, current 0, refresh 0, limit 1000
Number of NLRIs in the update sent: max 2, min 0
Minimum time between advertisement runs is 0 seconds
Has 1 member:
  10.100.1.3
```

驗證PE傳送的RTFilter:

```
PE1# show bgp rtfilter unicast all neighbors 10.100.1.3 advertised-routes
BGP table version is 8, local router ID is 10.100.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	1:2:1:1	0.0.0.0			32768	i
*>	1:2:1:2	0.0.0.0			32768	i

```
Total number of prefixes 2
```

對於自治系統編號，路由目標成員資格字首的編碼為4個位元組，對於路由目標（擴展社群屬性）為8個位元組。在上方範例中，rtfilter首碼「1:2:1:1」解碼如下：

- 1是自治系統編號
- 2是延伸社群屬性的型別和子型別（以十進位制表示）（請參閱RFC 4360）
- 1:1是路由目標本身

RR將預設篩選條件傳送到PE(RR-client)。這是因為根據設計，RR需要所有VPNv4路由：

```
BGP(10): (base) 10.100.1.1 send UPDATE (format) 0:0:0:0, next 10.100.1.3,
metric 0, path Local
```

PE接收並安裝預設rt過濾器。例如，它會將所有內容傳送到RR:
(debug bgp rtfilter unicast updates)

```
BGP(10): 10.100.1.3 rcvd UPDATE w/ attr: nexthop 10.100.1.3, origin i,
localpref 100, metric 0, community no-export
BGP(10): 10.100.1.3 rcvd 0:0:0:0
BGP(4): Default RT filter installed for 10.100.1.3
```

RR從PE1接收並安裝rtfilter:
(debug bgp rtfilter unicast updates)

```
BGP(10): 10.100.1.1 rcvd UPDATE w/ attr: nexthop 10.100.1.1, origin i,
localpref 100, metric 0
BGP(10): 10.100.1.1 rcvd 1:2:1:1
BGP(4): 1:2:1:1 RT filter installed for 10.100.1.1
BGP: installing rt filter on 10.100.1.1
BGP: add installed RT filter 1:2:1:1 for 10.100.1.1
BGP(10): 10.100.1.1 rcvd 1:2:1:2
BGP(4): 1:2:1:2 RT filter installed for 10.100.1.1
BGP(4): 1:2:1:2 Initiating an incremental table walk for 10.100.1.1
BGP: installing rt filter on 10.100.1.1
BGP: add installed RT filter 1:2:1:2 for 10.100.1.1
```

檢查RR：上收到的過濾器

```
RR1# show bgp vpnv4 unicast all neighbors 10.100.1.1 received rtfilters
Address family: VPNv4 Unicast
Extended community filter has: 2 entries with default filtering disabled
Incremental refresh walk mode
Status codes: * valid, S Stale > installed
Route-Target Outbound Filter
*> Extended Community RT:1:2
*> Extended Community RT:1:1
```

PE不安裝具有特定RT的RT過濾器。PE從RR收到預設rt過濾器，因此PE傳送所有VPNv4/v6字首：

```
PE1# show bgp vpnv4 unicast all neighbors 10.100.1.3 received rtfilters
Address family: VPNv4 Unicast
Extended community filter has: 1 entries with default filtering enabled
Incremental refresh walk mode
```

要建立預設RT過濾器，請在AF rtfilter下配置「neighbor x.x.x.x default-originate」。

這將在RR上自動為RR客戶端對等建立此項。

RR

```
router bgp 1

address-family rtfilter unicast
neighbor 10.100.1.1 activate
neighbor 10.100.1.1 send-community both
neighbor 10.100.1.1 route-reflector-client
neighbor 10.100.1.1 default-originate
exit-address-family
```

路由刷新處理

當配置新的RT匯入或刪除RT匯入時，會從PE將路由更新傳送到地址系列VPNv4/6的RR。

當配置新的VRF時，PE向RR傳送路由刷新。

在這兩種情況下，RTC都處於活動狀態，RR不會將所有VPNv4/6字首傳送到PE。它只根據RT過濾器傳送集合。

相關資訊

- [技術支援與文件 - Cisco Systems](#)