

路由目标约束

目录

[简介](#)

[路由目标约束的用途](#)

[没有RTC的行为](#)

[RTC配置](#)

[PE配置](#)

[RR配置](#)

[RTC的行为](#)

[PE](#)

[RR](#)

[路由刷新处理](#)

[相关信息](#)

简介

本文档介绍一种机制，通过这种机制，可将VPNv4和VPNv6前缀交换减少到所需的最低程度。

路由目标约束的用途

使用多协议标签交换(MPLS)VPN时，内部边界网关协议(iBGP)对等体或路由反射器(RR)将所有VPN4和/或VPN6前缀发送到PE路由器。PE路由器丢弃没有导入VPN路由和转发(VRF)的VPN4/6前缀。这是RR向PE路由器发送VPN4/6前缀的行为，RR不需要这种行为。这浪费了RR和PE的处理能力，也浪费了带宽。

使用路由目标约束(RTC)时，RR只向PE发送想要的VPN4/6前缀。“想要”表示PE具有VRF导入特定前缀。

RFC 4684指定RTC。支持通过VPNv4和VPNv6的新地址系列rfilter。

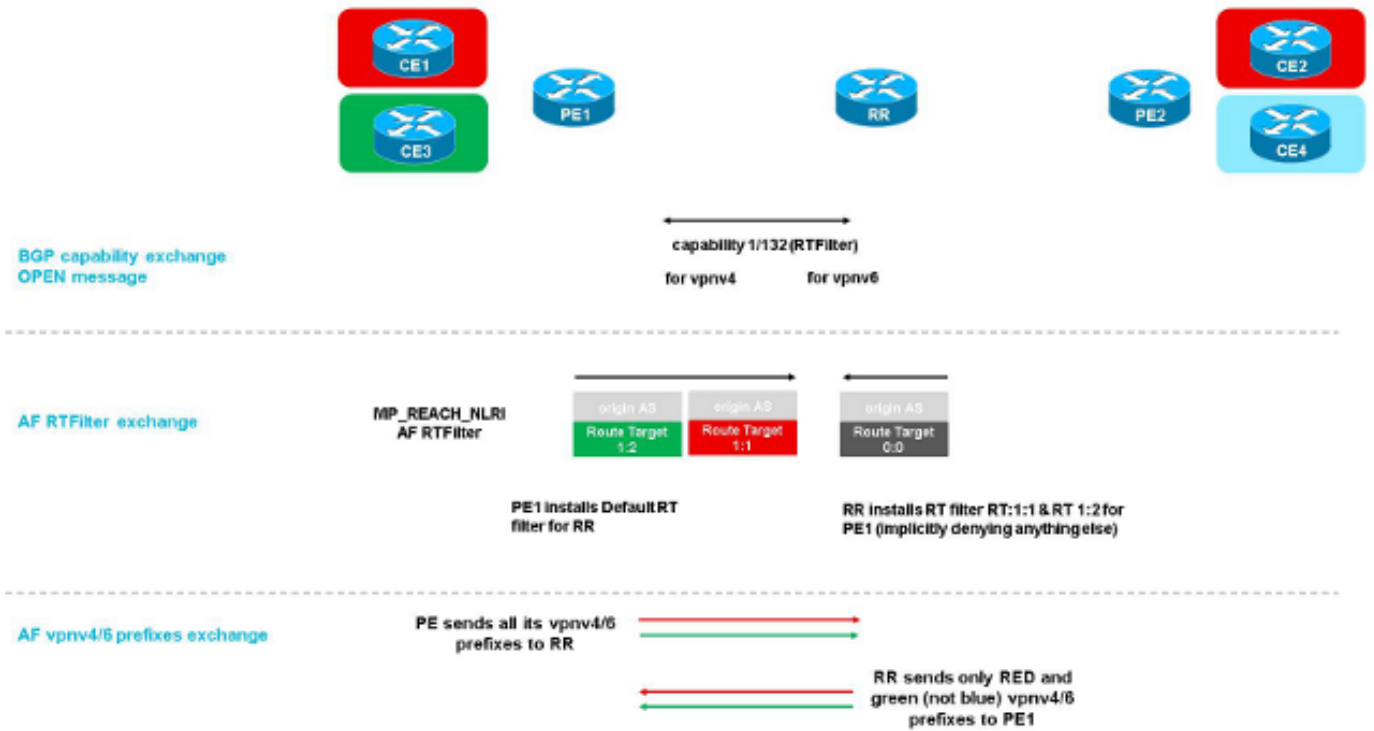
路由目标(RT)过滤信息从PE路由器上的所有VRF的VPN RT导入列表中获取。PE路由器将此过滤信息作为地址系列rfilter中的BGP更新发送到RR。此过滤信息或RT成员信息在MP_REACH_NLRI和MP_UNREACH_NLRI属性的网络层可达性信息(NLRI)中编码。

接收BGP对等体将此NLRI转换为过滤器，并将此过滤器安装到出站到发送对等体。接收BGP对等体使用此过滤器根据连接的RT的存在确定要发送或不发送的VPNv4/6前缀。

要使RTC工作，两个BGP对等体都需要支持RTC。即，RR和PE需要支持它。但是，部署可以是增量式的，这意味着并非所有路由反射器和PE路由器都需要一次性支持它。RTC可以在网络中工作，有些PE路由器支持RTC，而其他路由器不支持RTC。在支持它的路由器上，RTC已处于活动状态

。在尚不支持该功能的路由器上，通告将像以前一样工作，即没有RTC（因此没有任何出站过滤）。

此图显示了RTC的原理：



没有RTC的行为

RR将所有VPN4/6前缀发送到PE。PE会丢弃不导入RT的PE。调试BGP更新显示丢弃的前缀。消息“DENIED due to:提供不支持的扩展社区”。

VPNv4单播的示例如下：

```
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单播的示例如下：

```
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的行为

当BGP对等建立时，对等体交换rtfilter的功能，即1/132（对于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

	Sent	Rcvd
--	------	------

...

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还使用更新组：

```
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
```

验证PE发送的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

路由目标成员前缀的编码是自治系统编号的4个字节，路由目标的8个字节，这是扩展的社区属性。在上例中，rtfilter前缀“1:2:1:1”的解码如下：

- 1是自治系统编号
- 2是扩展社区属性的类型和子类型（十进制）（请参阅RFC 4360）
- 1:1是路由目标本身

RR将默认过滤器发送到PE(RR-client)。这是因为根据设计，路由反射器需要所有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接收并安装默认rt过滤器。例如，它将所有内容发送到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从PE1接收并安装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
```

检查RR上收到的过滤器：

```
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不安装具有特定RT的RT过滤器。PE从RR收到默认rt过滤器，因此PE发送所有VPNv4/v6前缀：

```
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
```

要创建默认RT过滤器，请在AF rtfilter下配置“neighbor x.x.x.x default-originate”。

这将在RR上为RR客户端项目自动创建。

RR

```
router bgp 1

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
```

路由刷新处理

当配置新RT导入或删除RT导入时，会从PE向RR发送地址系列VPNv4/6的路由刷新。

配置新VRF后，PE会向RR发送路由刷新。

在RTC处于活动状态的两种情况下，RR不会将所有VPNv4/6前缀发送到PE。它只根据RT过滤器发送集合。

相关信息

- [技术支持和文档 - Cisco Systems](#)