

路由目标限制条件

目录

[简介](#)

[路由目标限制条件目的](#)

[没有RTC的行为](#)

[RTC配置](#)

[PE配置](#)

[RR配置](#)

[RTC行为](#)

[PE](#)

[RR](#)

[路由刷新处理](#)

[相关信息](#)

简介

本文描述机制，藉以Vpvn4和VPNv6前缀交换往服务商边缘路由器减少到极小必要。

路由目标限制条件目的

使用多协议标签交换(MPLS) VPN，内部边界网关协议(iBGP)对等体或路由反射器(RR)发送所有VPN4和VPN6前缀到PE路由器。PE路由器下降没有导入的VPN路由与转发的VPN4/6前缀(VRF)。这是RR发送VPN4/6前缀对PE路由器，不需要的行为。这是浪费在RR和PE的处理功率和浪费带宽。

以路由目标限制条件(RTC)，RR发送仅希望的VPN4/6前缀对PE。‘希望’意味着PE有导入特定前缀的VRF。

RFC 4684指定RTC。支持是通过Vpvn4和VPNv6的一新的地址家族rfilter。

路由目标(RT)过滤信息从所有VRF的VPN RT海关进口货物分类表得到在PE路由器。PE路由器发送此过滤信息作为在地址家族rfilter的一次BGP更新对RR。此过滤信息或RT会员在MP_REACH_NLRI和MP_UNREACH_NLRI属性的网络层可达性信息(NLRI)编码。

接收的BGP对等体翻译此NLRI成过滤器并且安装此过滤器出站对发送的对等体。接收的BGP对等体使用此过滤器决定发送或发送的哪些VPNv4/6前缀，从属在附加的Rts出现。

为了使工作的RTC，两个BGP对等体需要支持RTC。即RR和PE需要支持它。然而，部署可以递增，含义不是所有的RR和PE路由器在一个需要支持它去。RTC在网络能工作，当一些PE路由器支持它和其他没有。在支持它的路由器上，RTC已经将是活跃的。在不支持它的路由器上广告将工作和

前面，是没有RTC (如此没有任何出局过滤)。

此图显示RTC的原理：

没有RTC的行为

RR发送所有VPN4/6前缀对PE。PE下降没有RT的导入的那个。调试BGP更新显示已丢失前缀。消息'拒绝由于：不支持的扩展团体'给。

Vpnv4单播的一示例如下：

```
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单播的一示例如下：

```
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行为

当BGP对等体设立时，对等体交换rtfilter的功能，是1/132 (VPNv4和VPNv6)。

```
RR1# show bgp rtfilter 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 NLRI 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

```

RR1# show bgp rtfilter 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 NLRI 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

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

也AF rtfilter用途更新组：

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 NLRI 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

路由目标成员关系前缀的编码是自治系统编号的4个字节和路由目标的8个字节，是扩展团体属性。在以上示例中，rtfilter前缀"1:2:1:1"解码如下：

- 1是自治系统编号
- 2是扩展团体属性的种类和子类型(在十进制) (参考的RFC 4360)
- 1:1是路由目标

RR发送默认过滤器对PE (RR客户端)。这是因为故意地，RR想要所有Vpnv4路由：

```
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

PE接收并且安装默认rt过滤器。例如，它发送一切对RR：
(调试bgp rtfilter单播更新)

```
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

RR接收并且安装从PE1的rtfilter：
(调试bgp rtfilter单播更新)

```
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
```

检查在RR的已接收过滤器：

```
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不安装有特定Rts的RT过滤器。PE接收从RR的默认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
```

为了创建默认RT过滤器，请配置“邻接x.x.x.x默认源”在AF rfilter下。

这在RR客户端对等互连的RR将自动地创建。

RR

```
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
```

路由刷新处理

当新的RT导入配置时或，当RT导入删除时，路由刷新从PE发送到地址家族的VPNv4/6 RR。

当新的VRF配置时，PE发送路由刷新对RR。

在两种情况下与RTC激活，RR不发送所有VPNv4/6前缀对PE。它根据RT过滤器只发送集。

相关信息

- [技术支持和文档 - Cisco Systems](#)