mVPN的严格RPF检查

目录

简介

本文档介绍VPN组播(mVPN)的严格反向路径转发(RPF)功能。 本文档使用Cisco IOS®中的示例和实施来说明该行为。

背景信息

RPF表示传入接口已检查到源。虽然会检查接口以确定它是指向源的正确接口，但不会检查它以确定它是该接口上的正确RPF邻居。在多路访问接口上，可能有多个邻居可以对其执行RPF。结果可能是路由器在该接口上收到两次相同的组播流并转发两者。

在多路访问接口上运行协议独立组播(PIM)的网络中，这不是问题，因为重复的组播流会导致运行断言机制，并且不再接收一个组播流。在某些情况下，PIM不在组播分布树(MDT)上运行，该树是多路访问接口。在这些情况下，边界网关协议(BGP)是重叠信令协议。

在使用分区MDT的配置文件中，即使PIM作为重叠协议运行，也不可能断言。原因是，在存在两个或多个入口PE路由器的情况下，一个入口提供商边缘(PE)不从另一个入口PE加入分区MDT。每个入口PE路由器可以将组播流转发到其分区MDT，而其他入口PE路由器不会看到组播流量。对于同一组播流，两个不同的出口PE路由器分别向不同的入口PE路由器加入MDT是有效方案：称为任播源。这允许不同的接收方加入同一组播流，但通过多协议标签交换(MPLS)核心中的不同路径。有关任播源的示例，请参阅图1。

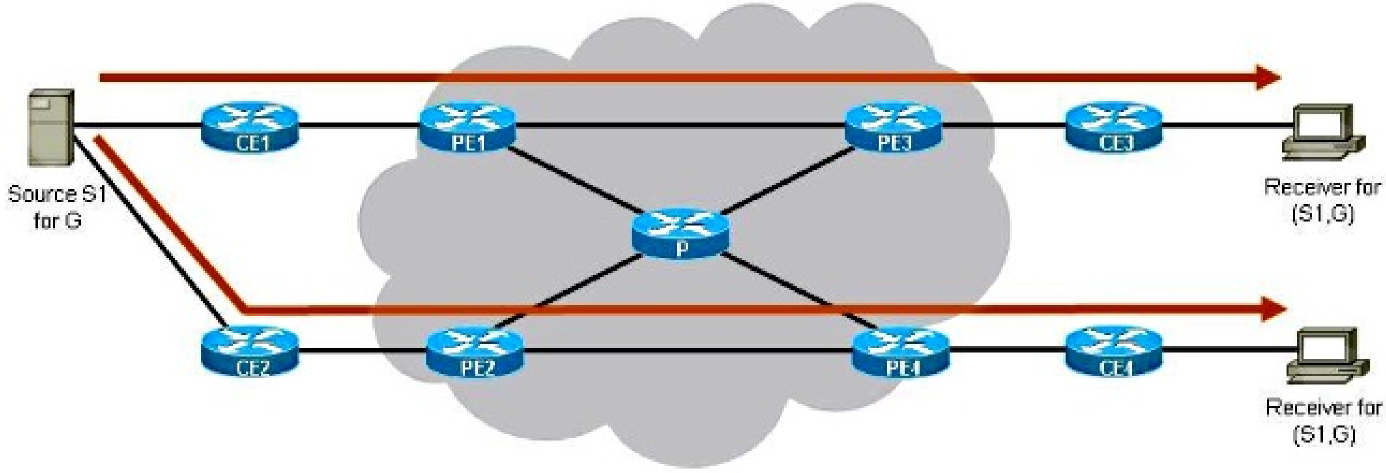

图 1

有两个入口PE路由器：PE1和PE2。有两个出口PE路由器：PE3和PE4。每个出口PE路由器都有一个不同的入口PE路由器作为其RPF邻居。PE3将PE1作为其RPF邻居。PE4将PE2作为其RPF邻居。出口PE路由器选择其最近的入口PE路由器作为其RPF邻居。

流(S1,G)将在顶部路径上从S1到接收器1，在底部路径上从S1到接收器2。两条路径（MPLS核心中的每条路径都是不同的分区MDT）上没有两个流的交集。

如果MDT是默认MDT（例如在默认MDT配置文件中），则这将不起作用，因为两个组播流将位于同一默认MDT上，并且断言机制将运行。如果MDT是默认MDT配置文件中的数据MDT，则所有入口PE路由器会从其他入口PE路由器加入数据MDT，因此会查看来自彼此的组播流量，并再次运行断言机制。如果重叠协议是BGP，则存在上游组播跳(UMH)选择，并且只选择一个入口PE路由器作为转发器，但这是按MDT。

任播源是运行分区MDT的一大优势。

问题

定期RPF检查确认数据包从正确的RPF接口到达路由器。不会检查是否从该接口上的正确RPF邻居接收了数据包。

请参阅图2。它显示了在分区MDT的场景中持续转发重复流量的问题。它表明，在分区MDT的情况下，常规RPF检查不足以避免重复的流量。

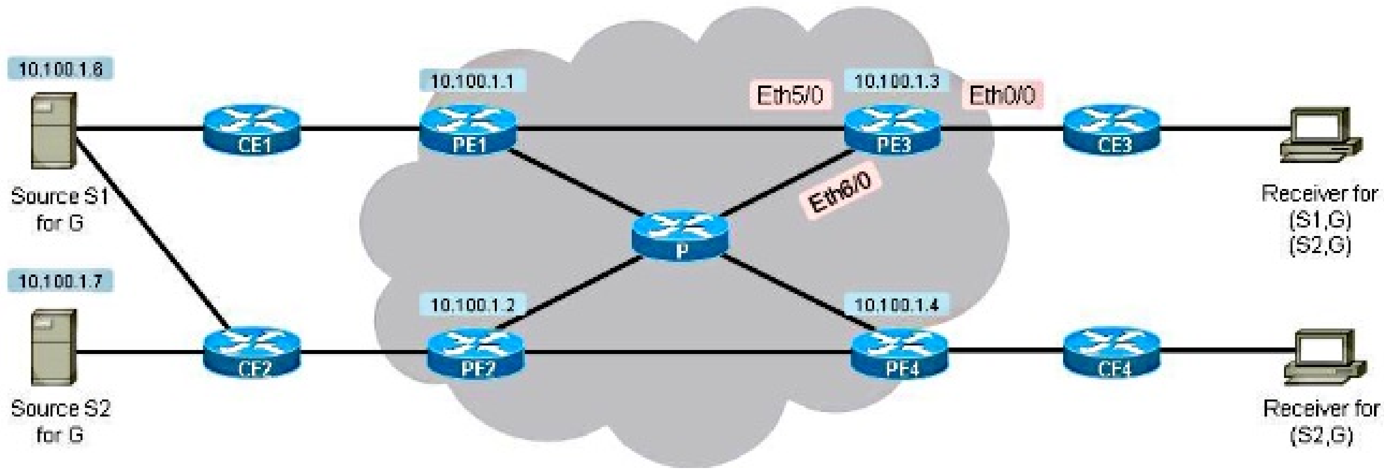

图 2

有两个接收器。第一接收器被设置为接收(S1,G)和(S2,G)的流量。 第二接收器被设置为仅接收(S2,G)的流量。有分区MDT，BGP是重叠信令协议。请注意，源S1可通过PE1和PE2访问。核心树协议是多点标签分发协议(mLDP)。

每个PE路由器通告第1类BGP IPv4 mVPN路由，该路由表示它是分区MDT的根候选。

```
PE3#show bgp ipv4 mvpn vrf one
BGP table version is 257, local router ID is 10.100.1.3
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-pah, 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
Route Distinguisher: 1:3 (default for vrf one)
*>i  [1][1:3][10.100.1.1]/12
                     10.100.1.1               0    100      0 ?
*>i  [1][1:3][10.100.1.2]/12
                      10.100.1.2               0    100      0 ?
*>  [1][1:3][10.100.1.3]/12
                      0.0.0.0                            32768 ?
*>i  [1][1:3][10.100.1.4]/12
                      10.100.1.4               0    100      0 ?
```

PE3在查找S1的单播路由后，发现PE1作为S1的RPF邻居。

```
PE3#show bgp vpnv4 unicast vrf one 10.100.1.6/32
BGP routing table entry for 1:3:10.100.1.6/32, version 16
Paths: (2 available, best #2, table one)
Advertised to update-groups:
    5
Refresh Epoch 2
65001, imported path from 1:2:10.100.1.6/32 (global)
  10.100.1.2 (metric 21) (via default) from 10.100.1.5 (10.100.1.5)
    Origin incomplete, metric 0, localpref 100, valid, internal
    Extended Community: RT:1:1 MVPN AS:1:0.0.0.0 MVPN VRF:10.100.1.2:1
    Originator: 10.100.1.2, Cluster list: 10.100.1.5
    mpls labels in/out nolabel/20
    rx pathid: 0, tx pathid: 0
Refresh Epoch 2
65001, imported path from 1:1:10.100.1.6/32 (global)
```

```
10.100.1.1 (metric 11) (via default) from 10.100.1.5 (10.100.1.5)
    Origin incomplete, metric 0, localpref 100, valid, internal, best
     Extended Community: RT:1:1 MVPN AS:1:0.0.0.0 MVPN VRF:10.100.1.1:1
    Originator: 10.100.1.1, Cluster list: 10.100.1.5
    mpls labels in/out nolabel/29
    rx pathid: 0, tx pathid: 0x0

PE3#show ip rpf vrf one 10.100.1.6
RPF information for ? (10.100.1.6)
  RPF interface: Lspvif0
  RPF neighbor: ? (10.100.1.1)
RPF route/mask: 10.100.1.6/32
RPF type: unicast (bgp 1)
Doing distance-preferred lookups across tables
RPF topology: ipv4 multicast base, originated from ipv4 unicast base
```

PE3选择PE1作为(S1,G)的RPF邻居，并将分区MDT与PE1作为根连接。PE3选择PE2作为(S2,G)的RPF邻居，并将分区MDT与PE2作为根连接。

```
PE3#show bgp vpnv4 unicast vrf one 10.100.1.7/32
BGP routing table entry for 1:3:10.100.1.7/32, version 18
Paths: (1 available, best #1, table one)
Advertised to update-groups:
    6
Refresh Epoch 2
65002, imported path from 1:2:10.100.1.7/32 (global)
     10.100.1.2 (metric 21) (via default) from 10.100.1.5 (10.100.1.5)
    Origin incomplete, metric 0, localpref 100, valid, internal, best
     Extended Community: RT:1:1 MVPN AS:1:0.0.0.0 MVPN VRF:10.100.1.2:1
    Originator: 10.100.1.2, Cluster list: 10.100.1.5
    mpls labels in/out nolabel/29
    rx pathid: 0, tx pathid: 0x0

PE3#show ip rpf vrf one 10.100.1.7
RPF information for ? (10.100.1.7)
  RPF interface: Lspvif0
  RPF neighbor: ? (10.100.1.2)
RPF route/mask: 10.100.1.7/32
RPF type: unicast (bgp 1)
Doing distance-preferred lookups across tables
RPF topology: ipv4 multicast base, originated from ipv4 unicast base
```

PE4选择PE2作为(S1,G)的RPF邻居，并将分区MDT与PE1作为根连接。

```
PE4#show bgp vpnv4 unicast vrf one 10.100.1.6/32
BGP routing table entry for 1:4:10.100.1.6/32, version 138
Paths: (2 available, best #1, table one)
Advertised to update-groups:
    2
Refresh Epoch 2
65001, imported path from 1:2:10.100.1.6/32 (global)
10.100.1.2 (metric 11) (via default) from 10.100.1.5 (10.100.1.5)
    Origin incomplete, metric 0, localpref 100, valid, internal, best
     Extended Community: RT:1:1 MVPN AS:1:0.0.0.0 MVPN VRF:10.100.1.2:1
    Originator: 10.100.1.2, Cluster list: 10.100.1.5
    mpls labels in/out nolabel/20
    rx pathid: 0, tx pathid: 0x0
Refresh Epoch 2
65001, imported path from 1:1:10.100.1.6/32 (global)
  10.100.1.1 (metric 21) (via default) from 10.100.1.5 (10.100.1.5)
    Origin incomplete, metric 0, localpref 100, valid, internal
    Extended Community: RT:1:1 MVPN AS:1:0.0.0.0 MVPN VRF:10.100.1.1:1
```

```
    Originator: 10.100.1.1, Cluster list: 10.100.1.5
    mpls labels in/out nolabel/29
    rx pathid: 0, tx pathid: 0

PE4#show ip rpf vrf one 10.100.1.6
RPF information for ? (10.100.1.6)
  RPF interface: Lspvif0
  RPF neighbor: ? (10.100.1.2)
RPF route/mask: 10.100.1.6/32
RPF type: unicast (bgp 1)
Doing distance-preferred lookups across tables
RPF topology: ipv4 multicast base, originated from ipv4 unicast base
```

注意，S1(10.100.1.6)和S2(10.100.1.7)的RPF接口均为Lspvif0。

PE3将(S2,G)的PE2分区MDT加入，PE4将(S1,G)的PE2分区MDT加入。PE1加入从PE1为(S1,G)划分的MDT。 通过在PE1和PE2上接收的第7类BGP IPv4 mVPN路由，您可以看到这一点。

```
PE1#show bgp ipv4 mvpn vrf one
BGP table version is 302, 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
Route Distinguisher: 1:1 (default for vrf one)
*>i  [7][1:1][1][10.100.1.6/32][232.1.1.1/32]/22
                      10.100.1.3               0    100      0 ?

PE2#show bgp ipv4 mvpn vrf one
BGP table version is 329, local router ID is 10.100.1.2
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
Route Distinguisher: 1:2 (default for vrf one)
*>i  [7][1:2][1][10.100.1.6/32][232.1.1.1/32]/22
                      10.100.1.4               0    100      0 ?
*>i  [7][1:2][1][10.100.1.7/32][232.1.1.1/32]/22
                      10.100.1.3               0    100      0 ?
```

PE3和PE4上的组播条目：

```
PE3#show ip mroute vrf one 232.1.1.1
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
      L - Local, P - Pruned, R - RP-bit set, F - Register flag,
      T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
      X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
      U - URD, I - Received Source Specific Host Report,
      Z - Multicast Tunnel, z - MDT-data group sender,
      Y - Joined MDT-data group, y - Sending to MDT-data group,
      G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
      N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
      Q - Received BGP S-A Route, q - Sent BGP S-A Route,
      V - RD & Vector, v - Vector, p - PIM Joins on route,
      x - VxLAN group
```

```
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(10.100.1.7, 232.1.1.1), 21:18:24/00:02:46, flags: sTg
Incoming interface: Lspvif0, RPF nbr 10.100.1.2
Outgoing interface list:
  Ethernet0/0, Forward/Sparse, 00:11:48/00:02:46

(10.100.1.6, 232.1.1.1), 21:18:27/00:03:17, flags: sTg
Incoming interface: Lspvif0, RPF nbr 10.100.1.1
Outgoing interface list:
  Ethernet0/0, Forward/Sparse, 00:11:48/00:03:17
```

```
PE4#show ip mroute vrf one 232.1.1.1
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
      L - Local, P - Pruned, R - RP-bit set, F - Register flag,
      T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
      X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
      U - URD, I - Received Source Specific Host Report,
      Z - Multicast Tunnel, z - MDT-data group sender,
      Y - Joined MDT-data group, y - Sending to MDT-data group,
      G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
      N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
      Q - Received BGP S-A Route, q - Sent BGP S-A Route,
      V - RD & Vector, v - Vector, p - PIM Joins on route,
      x - VxLAN group
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(10.100.1.6, 232.1.1.1), 20:50:13/00:02:37, flags: sTg
Incoming interface: Lspvif0, RPF nbr 10.100.1.2
Outgoing interface list:
  Ethernet0/0, Forward/Sparse, 20:50:13/00:02:37
```

这显示PE3加入根于PE1的点对多点(P2MP)树，以及根于PE2的树：

```
PE3#show mpls mldp database
* Indicates MLDP recursive forwarding is enabled

LSM ID : A   Type: P2MP   Uptime : 00:18:40
 FEC Root           : 10.100.1.1
 Opaque decoded     : [gid 65536 (0x00010000)]
Opaque length     : 4 bytes
Opaque value       : 01 0004 00010000
Upstream client(s) :
  10.100.1.1:0   [Active]
    Expires       : Never         Path Set ID : A
    Out Label (U) : None         Interface   : Ethernet5/0*
    Local Label (D): 29           Next Hop     : 10.1.5.1
Replication client(s):
  MDT (VRF one)
    Uptime         : 00:18:40     Path Set ID : None
    Interface     : Lspvif0

LSM ID : B   Type: P2MP   Uptime : 00:18:40
 FEC Root           : 10.100.1.2
 Opaque decoded     : [gid 65536 (0x00010000)]
Opaque length     : 4 bytes
Opaque value       : 01 0004 00010000
Upstream client(s) :
```

```
  10.100.1.5:0  [Active]
    Expires       : Never         Path Set ID : B
    Out Label (U) : None         Interface   : Ethernet6/0*
    Local Label (D): 30           Next Hop     : 10.1.3.5
Replication client(s):
  MDT (VRF one)
    Uptime        : 00:18:40     Path Set ID : None
    Interface     : Lspvif0
```

这显示PE4加入根于PE2的P2MP树：

```
PE4#show mpls mldp database
* Indicates MLDP recursive forwarding is enabled

LSM ID : 3   Type: P2MP   Uptime : 21:17:06
 FEC Root           : 10.100.1.2
 Opaque decoded     : [gid 65536 (0x00010000)]

Opaque value       : 01 0004 00010000
Upstream client(s) :
  10.100.1.2:0   [Active]
    Expires       : Never         Path Set ID : 3
    Out Label (U) : None         Interface   : Ethernet5/0*
    Local Label (D): 29           Next Hop     : 10.1.6.2
Replication client(s):
  MDT (VRF one)
    Uptime          : 21:17:06     Path Set ID : None
    Interface     : Lspvif0
```

组232.1.1.1的S1和S2数据流，10 pps。您可以在PE3和PE4上看到流。但是，在PE3上，您可以看到(S1,G)的速率为20 pps。

```
PE3#show ip mroute vrf one 232.1.1.1 count
Use "show ip mfib count" to get better response time for a large number of mroutes.

IP Multicast Statistics
3 routes using 1692 bytes of memory
2 groups, 1.00 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 232.1.1.1, Source count: 2, Packets forwarded: 1399687, Packets received:
2071455
Source: 10.100.1.7/32, Forwarding: 691517/10/28/2, Other: 691517/0/0
Source: 10.100.1.6/32, Forwarding: 708170/20/28/4, Other: 1379938/671768/0

PE4#show ip mroute vrf one 232.1.1.1 count
Use "show ip mfib count" to get better response time for a large number of mroutes.

IP Multicast Statistics
2 routes using 1246 bytes of memory
2 groups, 0.50 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 232.1.1.1, Source count: 1, Packets forwarded: 688820, Packets received:
688820
Source: 10.100.1.6/32, Forwarding: 688820/10/28/2, Other: 688820/0/0

PE3#show interfaces ethernet0/0 | include rate
Queueing strategy: fifo
30 second input rate 0 bits/sec, 0 packets/sec
```

```
30 second output rate 9000 bits/sec, 30 packets/sec
```
有重复的流。这种重复是流(S1,G)在PE1的分区MDT和PE2的分区MDT上存在的结果。第二个MDT（来自PE2）由PE3连接，以获得流(S2,G)。 但是，由于PE4从PE2加入分区MDT以获得(S1,G)，因此(S1,G)也存在于从PE2划分的MDT上。因此，PE3从其连接的两个分区MDT接收流(S1,G)。

PE3无法区分从PE1和PE2接收的(S1,G)的数据包。两个数据流都在正确的RPF接口上接收：Lspvif0。

```
PE3#show ip multicast vrf one mpls vif

Interface   Next-hop            Application     Ref-Count   Table / VRF name   Flags
 Lspvif0     0.0.0.0             MDT               N/A       1   (vrf one) 0x1
```
数据包可能到达PE3上或同一接口上的不同传入物理接口。无论如何，来自(S1,G)不同流的数据包在PE3上确实具有不同的MPLS标签：

```
PE3#show mpls forwarding-table vrf one
Local      Outgoing    Prefix            Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id      Switched      interface
29    [T] No Label   [gid 65536 (0x00010000)][V]    \
                                    768684        aggregate/one
30    [T] No Label   [gid 65536 (0x00010000)][V]    \
                                    1535940        aggregate/one

[T]     Forwarding through a LSP tunnel.
      View additional labelling info with the 'detail' option
```

解决方案

解决方案是制定更严格的RPF。使用严格RPF时，路由器会检查RPF接口上接收数据包的邻居。如果没有严格RPF，唯一的检查是确定传入接口是否是RPF接口，而不是确定是否从该接口上的正确RPF邻居接收数据包。

Cisco IOS注释

以下是有关使用Cisco IOS的RPF的一些重要说明。

- 在更改为严格RPF模式或从严格RPF模式更改时，请在配置分区MDT之前对其进行配置，或清除BGP。如果仅配置strict RPF命令，它不会立即创建另一个Lspvif接口。
- 默认情况下，Cisco IOS中未启用严格RPF。
- 不支持使用带有默认MDT配**置文件的**strict-rpf命令。

配置

您可以在PE3上为虚拟路由和转发(VRF)配置严格RPF。

```
vrf definition one
```

```
rd 1:3
!
address-family ipv4
mdt auto-discovery mldp
  mdt strict-rpf interface
  mdt partitioned mldp p2mp
mdt overlay use-bgp
route-target export 1:1
route-target import 1:1
exit-address-family
!
```

RPF信息已更改：

```
PE3#show ip rpf vrf one 10.100.1.6
RPF information for ? (10.100.1.6)
  RPF interface: Lspvif0
Strict-RPF interface: Lspvif1
  RPF neighbor: ? (10.100.1.1)
RPF route/mask: 10.100.1.6/32
RPF type: unicast (bgp 1)
Doing distance-preferred lookups across tables
RPF topology: ipv4 multicast base, originated from ipv4 unicast base

PE3#show ip rpf vrf one 10.100.1.7
RPF information for ? (10.100.1.7)
  RPF interface: Lspvif0
Strict-RPF interface: Lspvif2
  RPF neighbor: ? (10.100.1.2)
RPF route/mask: 10.100.1.7/32
RPF type: unicast (bgp 1)
Doing distance-preferred lookups across tables
RPF topology: ipv4 multicast base, originated from ipv4 unicast base
```

PE3为每个入口PE创建了Lspvif接口。Lspvif接口按入口PE、地址系列(AF)和VRF创建。10.100.1.6的RPF现在指向接口Lspvif1,10.100.1.7的RPF现在指向接口Lspvif2。

```
PE3#show ip multicast vrf one mpls vif

Interface   Next-hop          Application     Ref-Count   Table / VRF name   Flags
Lspvif0     0.0.0.0            MDT               N/A       1   (vrf one) 0x1
Lspvif1     10.100.1.1          MDT               N/A       1   (vrf one) 0x1
Lspvif2     10.100.1.2          MDT               N/A       1   (vrf one) 0x1
```

现在，PE1的数据包(S1,G)的RPF检查是针对RPF接口Lspvif1检查的。这些数据包带有MPLS标签29。来自PE2的数据包(S2,G)的RPF检查是针对RPF接口Lspvif2检查的。这些数据包带有MPLS标签30.流通过不同的传入接口到达PE3，但这也可能是同一接口。但是，由于mLDP从不使用Penultimate-Hop-Popping(PHP)，因此组播数据包上始终有一个常规MPLS标签。从PE1和PE2到达的(S1,G)数据包位于两个不同的分区MDT上，因此具有不同的MPLS标签。因此，PE3可以区分来自PE1的(S1,G)流和来自PE2的(S1,G)流。这样，PE3就可以分离数据包，并且RPF可以针对不同的入口PE路由器执行。

PE3上的mLDP数据库现在显示每个入口PE的不同Lspvif接口。

```
PE3#show mpls mldp database
* Indicates MLDP recursive forwarding is enabled

LSM ID : C   Type: P2MP   Uptime : 00:05:58
FEC Root           : 10.100.1.1
```

```
Opaque decoded     : [gid 65536 (0x00010000)]
Opaque length     : 4 bytes
Opaque value       : 01 0004 00010000
Upstream client(s) :
  10.100.1.1:0   [Active]
    Expires       : Never         Path Set ID : C
    Out Label (U) : None         Interface   : Ethernet5/0*
    Local Label (D): 29           Next Hop     : 10.1.5.1
Replication client(s):
  MDT (VRF one)
    Uptime         : 00:05:58     Path Set ID : None
    Interface     : Lspvif1

LSM ID : D   Type: P2MP   Uptime : 00:05:58
FEC Root           : 10.100.1.2
Opaque decoded     : [gid 65536 (0x00010000)]
Opaque length     : 4 bytes
Opaque value       : 01 0004 00010000
Upstream client(s) :
  10.100.1.5:0   [Active]
    Expires       : Never         Path Set ID : D
    Out Label (U) : None         Interface   : Ethernet6/0*
    Local Label (D): 30           Next Hop     : 10.1.3.5
Replication client(s):
  MDT (VRF one)
    Uptime         : 00:05:58     Path Set ID : None
    Interface     : Lspvif2
```

每个入口PE的严格RPF或RPF工作，因为组播流进入入口PE时，每个入口PE的MPLS标签不同：

```
PE3#show mpls forwarding-table vrf one
Local     Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label     Label      or Tunnel Id     Switched     interface
29   [T] No Label   [gid 65536 (0x00010000)][V]   \
                                      162708       aggregate/one
30   [T] No Label   [gid 65536 (0x00010000)][V]   \
                                      162750       aggregate/one

[T]     Forwarding through a LSP tunnel.
      View additional labelling info with the 'detail' option
```

严格RPF工作的证明是PE3上不再转发重复流(S1,G)。重复流仍然到达PE3，但由于RPF故障而被丢弃。RPF故障计数器为676255，并以10 pps的速率持续增加。

```
PE3#show ip mroute vrf one 232.1.1.1 count
Use "show ip mfib count" to get better response time for a large number of mroutes.

IP Multicast Statistics
3 routes using 1692 bytes of memory
2 groups, 1.00 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 232.1.1.1, Source count: 2, Packets forwarded: 1443260, Packets received:
2119515
Source: 10.100.1.7/32, Forwarding: 707523/10/28/2, Other: 707523/0/0
Source: 10.100.1.6/32, Forwarding: 735737/10/28/2, Other: 1411992/676255/0
```

PE3的输出速率现在为20 pps，即每个流(S1,G)和(S2,G)的10 pps:

```
PE3#show interfaces ethernet0/0 | include rate
```

```
Queueing strategy: fifo
30 second input rate 0 bits/sec, 0 packets/sec
30 second output rate 6000 bits/sec, 20 packets/sec
```

结论

必须使用分区MDT的mVPN部署模式必须使用严格RPF检查。

即使您不为使用分区MDT的mVPN部署模型配置严格RPF检查，情况也可能正常：组播流被传送到接收器。但是，当源连接到多个入口PE路由器时，可能存在重复的组播流量。这会导致网络带宽浪费，并可能对接收器上的组播应用造成负面影响。因此，必须为使用分区MDT的mVPN部署模型配置严格RPF检查。