適用於mVPN的嚴格RPF檢查

目錄

簡介

本檔案介紹透過VPN傳輸的多點傳送(mVPN)的嚴格反向路徑轉送(RPF)功能。 本檔案使用Cisco IOS®中的範例和實作來說明行為。

背景資訊

RPF意味著傳入介面被檢查到源。雖然會檢查介面以確定它是通向源的正確介面，但不會檢查介面以確定它是該介面上的正確RPF鄰居。在多路訪問介面上，可以有多個鄰居可供您使用RPF。結果可能是路由器在該介面上接收兩次相同的組播流，並轉發兩者。

在多路訪問介面上運行協定無關組播(PIM)的網路中，這不是問題，因為重複的多路傳播流會導致斷言機制運行，並且不再接收一個組播流。在某些情況下，PIM不會在組播分佈樹(MDT)上運行，該樹是多路訪問介面。在這些情況下，邊界閘道通訊協定(BGP)是重疊訊號通訊協定。

在具有Partitioned MDT的配置檔案中，即使PIM作為重疊協定運行，也無法斷言。原因是在存在兩個或多個輸入PE路由器的情況下，一個輸入提供程式邊緣(PE)不會從另一個輸入PE加入分割槽的MDT。每個輸入PE路由器可以將組播流轉發到其分割槽MDT上，而其他輸入PE路由器不會看到組播流量。對於同一組播流，兩個不同的輸出PE路由器各自將MDT加入不同的輸入PE路由器是一個有效的方案：它稱為任播源。這允許不同的接收者加入同一個多點傳播流，但通過多協定標籤交換(MPLS)核心中的不同路徑。有關任播源的示例，請參見圖1。

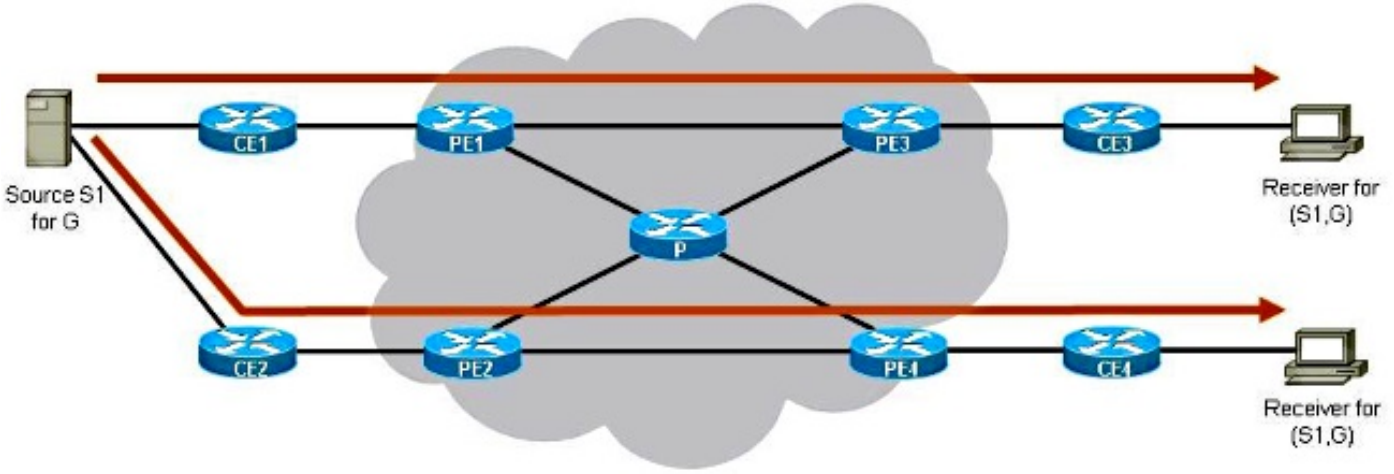

圖1

有兩台輸入PE路由器：PE1和PE2。有兩個出口PE路由器：PE3和PE4。每個出口PE路由器都有不同的輸入PE路由器作為其RPF鄰居。PE3將PE1作為其RPF鄰居。PE4將PE2作為其RPF鄰居。輸出PE路由器選擇其最近的輸入PE路由器作為其RPF鄰居。

流(S1,G)將通過頂部路徑從S1到接收器1，通過底部路徑從S1到接收器2。兩條路徑上的兩個流沒有交集（MPLS核心中的每個路徑都是不同的分割槽MDT）。

如果MDT是預設MDT（如在「預設MDT」配置檔案中），則此操作不會起作用，因為兩個組播流將位於同一個預設MDT上，並且斷言機制將運行。如果MDT是預設MDT配置檔案中的資料MDT，則所有輸入PE路由器都會從其他輸入PE路由器加入資料MDT，這樣就可以看到彼此之間的組播流量，並且斷言機制再次運行。如果重疊協定是BGP，則會選擇上游組播躍點(UMH)，且僅選擇一個輸入PE路由器作為轉發器，但這是每個MDT。

任播源是運行分割槽MDT的一大優勢。

問題

常規RPF檢查可確認資料包從正確的RPF介面到達路由器。沒有檢查可確認資料包是從該介面上的正確RPF鄰居接收的。

請參閱圖2。它顯示了在具有分割槽MDT的場景中持續轉發重複通訊的問題。它表明，對於分割槽的MDT，常規RPF檢查不足以避免重複流量。

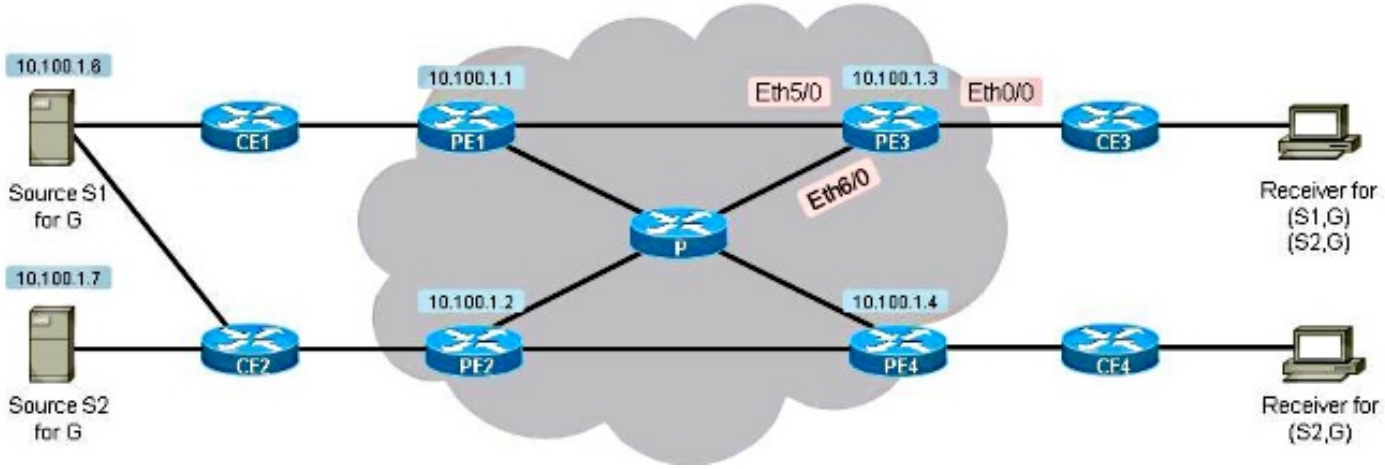

圖2

有兩個接收器。第一接收器設定為接收(S1,G)和(S2,G)的流量。 第二個接收器設定為僅接收(S2,G)的流量。存在分割槽MDT，BGP是重疊信令協定。請注意，源S1可以通過PE1和PE2訪問。核心樹協定是多點標籤分發協定(mLDP)。

每個PE路由器通告第1類BGP IPv4 mVPN路由，這表示該路由是作為分割槽MDT根的候選路由。

```
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 ?
```

查詢S1的單播路由後，PE3發現PE1是S1的RPF鄰居。

```
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選擇PE1作為(S1,G)的RPF鄰居，並以PE1作為根加入分割槽的MDT。PE3選擇PE2作為(S2,G)的RPF鄰居，並將分割槽的MDT與PE2作為根連線。

```
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選擇PE2作為(S1,G)的RPF鄰居，並以PE1作為根加入分割槽的MDT。

```
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
```

請注意，S1(10.100.1.6)和S2(10.100.1.7)的RPF介面都是Lspvif0。

對於(S2,G),PE3從PE2加入分割槽的MDT；對於(S1,G),PE4從PE2加入分割槽的MDT。對於(S1,G),PE1將分割槽的MDT與PE1連線。 您可通過在PE1和PE2上接收的型別7 BGP IPv4 mVPN路由看到這一點。

```
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上的組播條目：

```
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
```

這表明PE3加入了根在PE1的點對多點(P2MP)樹，也加入了根在PE2的樹：

```
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
```

這表明PE4加入根在PE2的P2MP樹：

```
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
```

組232.1.1.1的S1和S2流，帶10 pps。您可以在PE3和PE4看到資料流。但是，在PE3，您可以看到(S1,G)的速率是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
```

存在重複的流。該複製是來自PE1的分割槽MDT和來自PE2的分割槽MDT上存在流(S1,G)的結果。來自PE2的第二個分割槽MDT由PE3加入以獲得流(S2,G)。 但是，因為PE4從PE2加入分割槽的MDT以獲得(S1,G),(S1,G)也存在於PE2的分割槽MDT上。因此，PE3從它加入的兩個分割槽的MDT接收流(S1,G)。

PE3無法區分從PE1和PE2收到的(S1,G)資料包。兩個資料流均在正確的RPF介面上接收：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
```

資料包可能到達PE3上的不同傳入物理介面或同一介面上。在任何情況下，來自(S1,G)不同流的資料包到達時在PE3處確實帶有不同的MPLS標籤：

```
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
```

解決方案

解決方案是採用更嚴格的RPF。使用嚴格RPF時，路由器會檢查RPF介面上接收資料包的鄰居。如果沒有嚴格的RPF，唯一的檢查是確定傳入介面是否為RPF介面，而不是確定資料包是否從該介面上的正確RPF鄰居接收。

Cisco IOS說明

以下是有關使用Cisco IOS的RPF的一些重要說明。

- 當切換到/從嚴格RPF模式時，請在配置分割槽的MDT之前對其進行配置，或者清除BGP。如果只配置嚴格RPF命令，則不會立即建立另一個Lspvif介面。
- Cisco IOS預設不啟用嚴格RPF。
- 不支援將**strict-rpf** 命令與預設MDT配置檔案一起使用。

組態

您可以在PE3上為虛擬路由和轉發(VRF)配置嚴格的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資訊已更改：

```
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為每個輸入PE建立了一個Lspvif介面。Lspvif介面是根據入口PE、地址系列(AF)和VRF建立的。10.100.1.6的RPF現在指向介面Lspvif1,10.100.1.7的RPF現在指向介面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
```

現在，根據RPF介面Lspvif1檢查來自PE1的資料包(S1,G)。這些資料包帶有MPLS標籤29。根據RPF介面Lspvif2檢查來自PE2的資料包(S2,G)的RPF檢查。這些資料包帶有MPLS標籤30。流通過不同的傳入介面到達PE3，但此介面也可以是同一個介面。但是，由於mLDP從不使用倒數跳躍(Penultimate-Hop-Popping，PHP)，因此組播資料包的頂部始終有一個常規MPLS標籤。從PE1和PE2到達的(S1,G)資料包位於兩個不同的分割槽MDT上，因此具有不同的MPLS標籤。因此，PE3可以區分來自PE1的(S1,G)流和來自PE2的(S1,G)流。這樣，資料包可以被PE3分開，並且可以對不同的輸入PE路由器執行RPF。

PE3上的mLDP資料庫現在顯示每個輸入PE的不同Lspvif介面。

```
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
```

每個輸入PE的嚴格RPF或RPF工作正常，因為組播流傳入每個輸入PE具有不同MPLS標籤的輸入PE:

```
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
```

嚴格的RPF起作用的證明是，在PE3上不再有重複流(S1,G)轉發。重複流仍然到達PE3，但由於RPF失敗而被丟棄。RPF故障計數器處於676255狀態，並以10 pps的速率持續增加。

```
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處的輸出速率現在是20 pps，這是每個流(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
```

結論

對於使用分割槽MDT的mVPN部署模型，必須使用嚴格RPF檢查。

即使您沒有為具有分割槽MDT的mVPN部署模型配置嚴格RPF檢查，情況可能看起來仍然正常：將多播流傳送到接收器。但是，當源連線到多個輸入PE路由器時，可能存在重複的組播流量。這會導致網路中頻寬的浪費，並且可能會對接收器上的組播應用產生不利影響。因此，必須為使用分割槽MDT的mVPN部署模型配置嚴格的RPF檢查。