mVPN中的雙宿主源和資料MDT

目錄

簡介

本檔案介紹具有雙宿主來源和資料MDT（多點傳送發佈樹）的mVPN（多點傳送虛擬提供者網路）。 Cisco IOS®中的範例用於說明行為。

問題

如果mVPN世界中的一個來源是雙宿主到兩個輸入提供商邊緣(PE)路由器，則兩個輸入PE路由器可能都將一個(S，G)的流量轉發到多協定標籤交換(MPLS)雲中。例如，如果存在兩個輸出PE路由器，並且每個反向路徑轉發(RPF)到不同的輸入PE路由器，則可能會發生這種情況。如果兩個輸入PE路由器都轉發到預設MDT上，則斷言機制將啟動，其中一個輸入PE贏得斷言機制，另一個將丟失，因此一個且只有一個輸入PE繼續將客戶(C-)(S，G)轉發到MDT上。但是，如果由於任何原因斷言機制未在預設MDT上啟動，則兩個輸入PE路由器可能都開始將C-(S，G)組播流量傳輸到它們發起的一個Data-MDT上。由於流量不再位於預設MDT上，而是位於資料MDT上，因此兩個輸入PE路由器不會在MDT/隧道介面上相互接收C-(S，G)流量。這可能導致下游出現持續重複流量。本檔案將說明此問題的解決方案。

預設MDT上的斷言機制

無論核心樹協定如何，本節中的資訊對預設MDT都適用。選擇的核心樹協定是協定無關組播(PIM)。

範例中使用的是Cisco IOS，但前面提到的所有內容同樣適用於Cisco IOS-XR。使用的所有組播組都是源特定組播(SSM)組。

檢視圖1.雙宿主Source-1。有兩個輸入PE路由器（PE1和PE2）和兩個輸出PE路由器（PE3和PE4）。 源位於CE1,IP地址為10.100.1.6。CE1雙宿主到PE1和PE2。

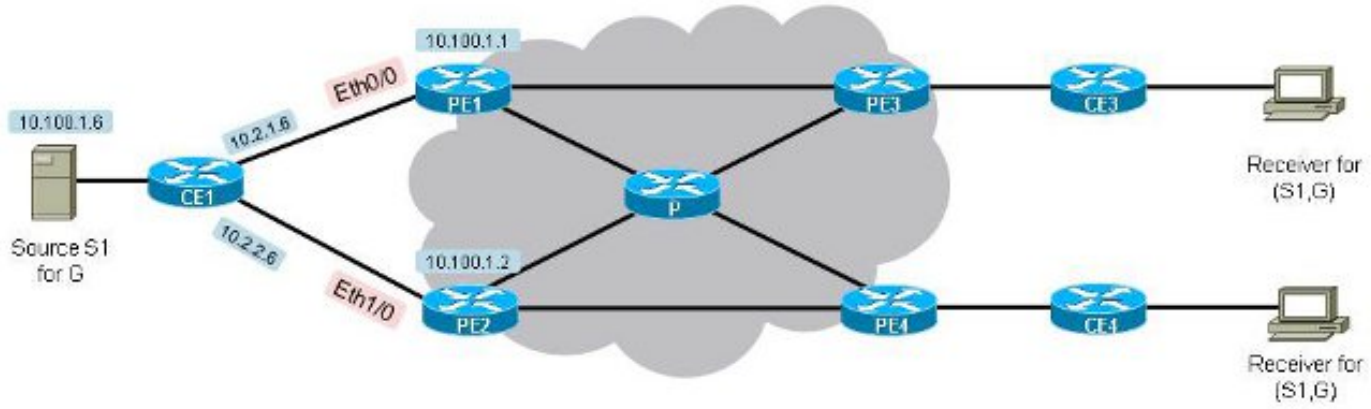

圖1.雙宿源1

所有PE路由器上的配置(PE路由器上的路由區分器(RD)可以不同)如下：

```
vrf definition one
 rd 1:1
 !
 address-family ipv4
 mdt default 232.10.10.10
 route-target export 1:1
 route-target import 1:1
 exit-address-family
!
```

為使兩個輸入PE路由器開始將組播流(10.100.1.6,232.1.1.1)轉發到預設MDT，它們都必須從輸出PE接收加入。檢視Figure1. Dual-Homed-Source-1中的拓撲。您可以看到，預設情況下，如果邊緣鏈路的所有成本都相同，且核心鏈路的所有成本都相同，則PE3將向PE1傳送RPF，PE4將向PE2傳送RPF(10.100.1.6,232.1.1.1)。 它們都將RPF設定為最接近的輸入PE。此輸出確認這點：

```
PE3#show ip rpf vrf one 10.100.1.6
RPF information for ? (10.100.1.6)
 RPF interface: Tunnel0
 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
 BGP originator: 10.100.1.1
 RPF topology: ipv4 multicast base, originated from ipv4 unicast base
```

PE3具有RPF到PE1。

```
PE4#show ip rpf vrf one 10.100.1.6
RPF information for ? (10.100.1.6)
 RPF interface: Tunnel0
 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
 BGP originator: 10.100.1.2
 RPF topology: ipv4 multicast base, originated from ipv4 unicast base
```

PE4與PE2之間具有RPF。PE3選擇PE1作為RPF鄰居的原因是，在VRF（虛擬路由/轉發）中，指向10.100.1.6/32的單播路由是通過PE1的最佳路由。PE3實際上從PE1和PE2都接收路由10.100.1.6/32。邊界網關協定(BGP)最佳路徑計算演算法中的所有條件都相同，但指向BGP下一跳地址的開銷除外。

```
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 333
Paths: (2 available, best #1, table one)
 Advertised to update-groups:
    21
 Refresh Epoch 1
 Local, 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 11, localpref 100, valid, internal,best
     Extended Community: RT:1:1 OSPF DOMAIN ID:0x0005:0x000000640200
       OSPF RT:0.0.0.0:2:0 OSPF ROUTER ID:10.2.4.1:0
     Originator: 10.100.1.1, Cluster list: 10.100.1.5
     Connector Attribute: count=1
      type 1 len 12 value 1:1:10.100.1.1
     mpls labels in/out nolabel/32
     rx pathid: 0, tx pathid: 0x0
 Refresh Epoch 1
 Local, 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 11, localpref 100, valid, internal
     Extended Community: RT:1:1 OSPF DOMAIN ID:0x0005:0x000000640200
       OSPF RT:0.0.0.0:2:0 OSPF ROUTER ID:10.2.2.2:0
     Originator: 10.100.1.2, Cluster list: 10.100.1.5
     Connector Attribute: count=1
      type 1 len 12 value 1:2:10.100.1.2
     mpls labels in/out nolabel/29
     rx pathid: 0, tx pathid: 0

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 1050
Paths: (2 available, best #2, table one)
 Advertised to update-groups:
    2
 Refresh Epoch 1
 Local, 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 11, localpref 100, valid, internal
     Extended Community: RT:1:1 OSPF DOMAIN ID:0x0005:0x000000640200
       OSPF RT:0.0.0.0:2:0 OSPF ROUTER ID:10.2.4.1:0
     Originator: 10.100.1.1, Cluster list: 10.100.1.5
     Connector Attribute: count=1
      type 1 len 12 value 1:1:10.100.1.1
     mpls labels in/out nolabel/32
     rx pathid: 0, tx pathid: 0
 Refresh Epoch 1
 Local, 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 11, localpref 100, valid, internal, best
     Extended Community: RT:1:1 OSPF DOMAIN ID:0x0005:0x000000640200
       OSPF RT:0.0.0.0:2:0 OSPF ROUTER ID:10.2.2.2:0
     Originator: 10.100.1.2, Cluster list: 10.100.1.5
```

```
     Connector Attribute: count=1
      type 1 len 12 value 1:2:10.100.1.2
     mpls labels in/out nolabel/29
     rx pathid: 0, tx pathid: 0x0
```

PE3選擇的最佳路徑是PE1通告的路徑，因為與PE2的IGP成本(21)相比，PE1具有最低的內部網關協定(IGP)成本(11)。對於PE4，則相反。拓撲結構顯示，從PE3到PE1隻有一跳，而從PE3到PE2隻有兩跳。由於所有鏈路的IGP成本都相同，因此PE3從PE1中選擇最佳路徑。

PE1和PE2上的組播路由資訊庫(MRIB)(10.100.1.6,232.1.1.1)如下所示（尚無組播流量）：

```
PE1#show ip mroute vrf one 232.1.1.1 10.100.1.6
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), 00:00:12/00:03:17, flags: sT
 Incoming interface: Ethernet0/0, RPF nbr 10.2.1.6
 Outgoing interface list:
   Tunnel0, Forward/Sparse, 00:00:12/00:03:17

PE2#show ip mroute vrf one 232.1.1.1 10.100.1.6
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), 00:00:47/00:02:55, flags: sT
 Incoming interface: Ethernet1/0, RPF nbr 10.2.2.6
 Outgoing interface list:
   Tunnel0, Forward/Sparse, 00:00:47/00:02:55
```

PE1和PE2都收到的PIM連線(10.100.1.6,232.1.1.1)。 Tunnel0介面位於兩台路由器上組播條目的傳出介面清單(OIL)中。

組播流量開始流向(10.100.1.6,232.1.1.1)。「Debug ip pim vrf one 232.1.1.1」和「debug ip mrouting vrf one 232.1.1.1」向我們表明，組播流量到達兩個輸入PE路由器的Tunnel0（在OIL中

）後，斷言機制將運行。

PE1

```
PIM(1): Send v2 Assert on Tunnel0 for 232.1.1.1, source 10.100.1.6, metric [110/11]
PIM(1): Assert metric to source 10.100.1.6 is [110/11]
MRT(1): not RPF interface, source address 10.100.1.6, group address 232.1.1.1
PIM(1): Received v2 Assert on Tunnel0 from 10.100.1.2
PIM(1): Assert metric to source 10.100.1.6 is [110/11]
PIM(1): We lose, our metric [110/11]
PIM(1): Prune Tunnel0/232.10.10.10 from (10.100.1.6/32, 232.1.1.1)
MRT(1): Delete Tunnel0/232.10.10.10 from the olist of (10.100.1.6, 232.1.1.1)
MRT(1): Reset the PIM interest flag for (10.100.1.6, 232.1.1.1)
MRT(1): set min mtu for (10.100.1.6, 232.1.1.1) 1500->18010 - deleted
PIM(1): Received v2 Join/Prune on Tunnel0 from 10.100.1.3, not to us
PIM(1): Join-list: (10.100.1.6/32, 232.1.1.1), S-bit set
```

PE2

```
PIM(1): Received v2 Assert on Tunnel0 from 10.100.1.1
PIM(1): Assert metric to source 10.100.1.6 is [110/11]
PIM(1): We win, our metric [110/11]
PIM(1): (10.100.1.6/32, 232.1.1.1) oif Tunnel0 in Forward state
PIM(1): Send v2 Assert on Tunnel0 for 232.1.1.1, source 10.100.1.6, metric [110/11]
PIM(1): Assert metric to source 10.100.1.6 is [110/11]
PIM(1): Received v2 Join/Prune on Tunnel0 from 10.100.1.3, to us
PIM(1): Join-list: (10.100.1.6/32, 232.1.1.1), S-bit set
PIM(1): Update Tunnel0/10.100.1.3 to (10.100.1.6, 232.1.1.1), Forward state, by PIM SG Join
```

如果兩個路由器到源10.100.1.6的度量和距離相同，則有一個連線中斷器來確定斷言獲勝者。分路器是Tunnel0（預設MDT）上PIM鄰居的最高IP地址。 在本例中，這是PE2:

```
PE1#show ip pim vrf one neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
10.100.1.4        Tunnel0                  06:27:57/00:01:29 v2    1 / DR S P G
10.100.1.3        Tunnel0                  06:28:56/00:01:24 v2    1 / S P G
10.100.1.2        Tunnel0                  06:29:00/00:01:41 v2    1 / S P G

PE1#show ip pim vrf one interface

Address          Interface                Ver/   Nbr    Query  DR         DR
                                          Mode   Count  Intvl  Prior
10.2.1.1         Ethernet0/0              v2/S   0      30     1          10.2.1.1
10.2.4.1         Ethernet1/0              v2/S   0      30     1          10.2.4.1
10.100.1.1       Lspvif1                  v2/S   0      30     1          10.100.1.1
10.100.1.1       Tunnel0                  v2/S   3      30     1          10.100.1.4
```

由於斷言，PE1從組播條目的OIL中刪除了Tunnel0。由於OIL變為空，組播條目將被剪除。

```
PE1#show ip mroute vrf one 232.1.1.1 10.100.1.6
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), 00:17:24/00:00:01, flags: sPT
 Incoming interface: Ethernet0/0, RPF nbr 10.2.1.6
 Outgoing interface list: Null
```

PE2在介面Tunnel0上設定了A標誌，因為它是斷言獲勝者。

```
PE2#show ip mroute vrf one 232.1.1.1 10.100.1.6
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), 00:17:20/00:02:54, flags: sT
 Incoming interface: Ethernet1/0, RPF nbr 10.2.2.6
 Outgoing interface list:
   Tunnel0, Forward/Sparse, 00:17:20/00:02:54, A
```

PE2定期在Tunnel0（預設MDT）上傳送一個斷言，就在斷言計時器到期之前。因此，PE2仍是絕對的贏家。

```
PE2#
PIM(1): Send v2 Assert on Tunnel0 for 232.1.1.1, source 10.100.1.6, metric [110/11]
PIM(1): Assert metric to source 10.100.1.6 is [110/11]
```

結論

斷言機制也與OIL中的隧道介面一起使用。當輸入PE路由器在OIL中的關聯隧道介面上接收C-(S，G)組播流量時，斷言將通過預設MDT交換。

具有資料MDT的斷言機制

在配置資料MDT的大部分時間，由於C-(S，G)流量僅在三秒後從預設MDT切換到資料MDT，因此斷言機制仍將在預設MDT上運行。然後如前所述發生同樣情況。請注意，**每個支援多點傳送的VRF只有一個通道介面**:預設MDT和所有資料MDT僅使用一個隧道介面。此隧道介面用於輸入PE路

由器的OIL或輸出PE路由器的RPF介面。

在某些情況下，在發出資料MDT訊號之前可能未觸發斷言機制。然後，C-(S，G)組播流量可能會在入口PE路由器PE1和PE2上的資料MDT上開始轉發。在這種情況下，可能會導致在MPLS核心網路中出現永久重複的C-(S，G)組播流量。為了避免這種情況，實施了以下解決方案：當入口PE路由器看到另一入口PE路由器通告資料MDT時，該PE路由器也是該資料MDT的入口PE路由器，該路由器將加入該資料MDT。原則上，只有輸出PE路由器（具有下游接收器）會加入資料MDT。由於輸入PE路由器加入其他輸入PE路由器通告的資料MDT，導致輸入PE路由器從OIL中存在的隧道介面接收組播流量，因此這會觸發斷言機制並導致其中一個輸入PE路由器停止將C-(S，G)組播流量轉發到其資料MDT（使用隧道介面），而另一個輸入PE（斷言獲勝者）可以繼續將C-(S，G)組播流量轉發到其資料MDT。

在下一個示例中，假設入口PE路由器PE1和PE2從未在預設MDT上看到彼此的C-(S，G)組播流量。流量在預設MDT上僅持續三秒，因此不難理解，例如，如果核心網路上臨時丟失流量，就會發生這種情況。

Data MDT的配置新增到所有PE路由器。所有PE路由器（PE路由器上的RD可以不同）的配置如下：

```
vrf definition one
 rd 1:1
 !
 address-family ipv4
 mdt default 232.10.10.10
mdt data 232.11.11.0 0.0.0.0
 route-target export 1:1
 route-target import 1:1
 exit-address-family
!
```

PE1和PE2一看到來自源的流量，就會建立一個C-(S，G)條目。兩個輸入PE路由器將C-(S，G)組播流量轉發到預設MDT上。出口PE路由器PE3和PE4接收組播流量並將其轉發。由於臨時問題，PE2在預設MDT上看不到來自PE1的流量，反之亦然。它們都傳送預設MDT上的資料MDT聯接型別長度值(TLV)。

如果沒有C-(S，G)流量，您會在輸入PE路由器上看到此組播狀態：

```
PE1#show ip mroute vrf one 232.1.1.1 10.100.1.6
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), 00:00:45/00:02:44, flags: sT
 Incoming interface: Ethernet0/0, RPF nbr 10.2.1.6
```

```
 Outgoing interface list:
   Tunnel0, Forward/Sparse, 00:00:45/00:02:42

PE2#show ip mroute vrf one 232.1.1.1 10.100.1.6
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), 00:02:18/00:03:28, flags: sT
 Incoming interface: Ethernet1/0, RPF nbr 10.2.2.6
 Outgoing interface list:
   Tunnel0, Forward/Sparse, 00:02:18/00:03:28
```

Y標誌尚未設定。兩台輸入PE路由器的OIL中都具有Tunnel0介面。這是因為PE3對PE1有RPF，而PE4對C-(S，G)對PE2有RPF。

當C-(S，G)的組播流量開始流動時，PE1和PE2都會轉發該流量。在入口PE路由器上跨越資料MDT的閾值，兩個都傳送資料MDT加入TLV，並在三秒鐘後開始轉發到其資料MDT。請注意，PE1加入源自PE2的資料MDT，PE2加入源自PE1的資料MDT。

```
PE1#show ip mroute vrf one 232.1.1.1 10.100.1.6
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), 00:01:26/00:03:02, flags: sTy
 Incoming interface: Ethernet0/0, RPF nbr 10.2.1.6
 Outgoing interface list:
   Tunnel0, Forward/Sparse, 00:01:26/00:03:02

PE2#show ip mroute vrf one 232.1.1.1 10.100.1.6
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), 00:00:41/00:02:48, flags: sTy
 Incoming interface: Ethernet1/0, RPF nbr 10.2.2.6
 Outgoing interface list:
   Tunnel0, Forward/Sparse, 00:00:41/00:02:48
```

PE1和PE都在Tunnel0介面上接收C-(S，G)的流量（但現在從Data MDT而不是預設MDT），並且啟動斷言機制。只有PE2在其資料MDT上繼續轉發C-(S，G)流量：

```
PE1#
PIM(1): Send v2 Assert on Tunnel0 for 232.1.1.1, source 10.100.1.6, metric [110/11]
PIM(1): Assert metric to source 10.100.1.6 is [110/11]
MRT(1): not RPF interface, source address 10.100.1.6, group address 232.1.1.1
PIM(1): Received v2 Assert on Tunnel0 from 10.100.1.2
PIM(1): Assert metric to source 10.100.1.6 is [110/11]
PIM(1): We lose, our metric [110/11]
PIM(1): Prune Tunnel0/232.11.11.0 from (10.100.1.6/32, 232.1.1.1)
MRT(1): Delete Tunnel0/232.11.11.0 from the olist of (10.100.1.6, 232.1.1.1)
MRT(1): Reset the PIM interest flag for (10.100.1.6, 232.1.1.1)
PIM(1): MDT Tunnel0 removed from (10.100.1.6,232.1.1.1)
MRT(1): Reset the y-flag for (10.100.1.6,232.1.1.1)
PIM(1): MDT next_hop change from: 232.11.11.0 to 232.10.10.10 for (10.100.1.6, 232.1.1.1)
Tunnel0
MRT(1): set min mtu for (10.100.1.6, 232.1.1.1) 1500->18010 - deleted
PIM(1): MDT threshold dropped for (10.100.1.6,232.1.1.1)
PIM(1): Receive MDT Packet (9889) from 10.100.1.2 (Tunnel0), length (ip: 44, udp: 24), ttl: 1
PIM(1): TLV type: 1 length: 16 MDT Packet length: 16

PE2#
PIM(1): Received v2 Assert on Tunnel0 from 10.100.1.1
PIM(1): Assert metric to source 10.100.1.6 is [110/11]
PIM(1): We win, our metric [110/11]
PIM(1): (10.100.1.6/32, 232.1.1.1) oif Tunnel0 in Forward state
PIM(1): Send v2 Assert on Tunnel0 for 232.1.1.1, source 10.100.1.6, metric [110/11]
PIM(1): Assert metric to source 10.100.1.6 is [110/11]
PE2#
PIM(1): Received v2 Join/Prune on Tunnel0 from 10.100.1.3, to us
PIM(1): Join-list: (10.100.1.6/32, 232.1.1.1), S-bit set
PIM(1): Update Tunnel0/10.100.1.3 to (10.100.1.6, 232.1.1.1), Forward state, by PIM SG Join
MRT(1): Update Tunnel0/232.10.10.10 in the olist of (10.100.1.6, 232.1.1.1), Forward state - MAC
built
MRT(1): Set the y-flag for (10.100.1.6,232.1.1.1)
PIM(1): MDT next_hop change from: 232.10.10.10 to 232.11.11.0 for (10.100.1.6, 232.1.1.1)
Tunnel0
```

PE1在OIL中不再具有隧道介面。

```
PE1#show ip mroute vrf one 232.1.1.1 10.100.1.6
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), 00:10:23/00:00:04, flags: sPT
 Incoming interface: Ethernet0/0, RPF nbr 10.2.1.6
 Outgoing interface list: Null
```

PE2在Tunnel0介面上設定了A標誌：

```
PE2#show ip mroute vrf one 232.1.1.1 10.100.1.6
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), 00:10:00/00:02:48, flags: sTy
 Incoming interface: Ethernet1/0, RPF nbr 10.2.2.6
 Outgoing interface list:
   Tunnel0, Forward/Sparse, 00:08:40/00:02:48, A
```

結論

使用資料MDT時，斷言機制也起作用。當輸入PE路由器在OIL中的關聯隧道介面上接收C-(S，G)組播流量時，斷言將通過預設MDT交換。