mVPN中的双宿主源和数据MDT

目录

简介

本文档介绍具有双宿主源和数据MDT（组播分布树）的mVPN（组播虚拟提供商网络）。 Cisco IOS®中的一个示例用于说明该行为。

问题

如果mVPN世界中的源是双宿主到两个入口提供商边缘(PE)路由器，则两个入口PE路由器可以同时将一个(S，G)的流量转发到多协议标签交换(MPLS)云。例如，如果有两个出口PE路由器，并且每个反向路径转发(RPF)都连接到不同的入口PE路由器，则可能出现这种情况。如果两个入口PE路由器都转发到默认MDT，则断言机制将启动，一个入口PE赢得断言机制，而另一个失去，从而一个和仅一个入口PE继续将客户(C-)(S，G)转发到MDT。但是，如果由于任何原因，断言机制未在默认MDT上启动，则两个入口PE路由器可能会开始将C-(S，G)组播流量传输到它们启动的一个Data-MDT上。由于流量不再在默认MDT上，而在数据MDT上，两个入口PE路由器在MDT/隧道接口上不会收到来自彼此的C-(S，G)流量。这可能导致持续的下游重复流量。本文档说明了此问题的解决方案。

默认MDT上的断言机制

无论核心树协议如何，本部分中的信息对于默认MDT都正确。所选核心树协议是协议无关组播(PIM)。

示例使用Cisco IOS，但所提及的所有内容同样适用于Cisco IOS-XR。使用的所有组播组都是源特定组播(SSM)组。

请看图1。双宿主源–1。有两个入口PE路由器（PE1和PE2）和两个出口PE路由器（PE3和PE4）。 源地址为CE1,IP地址为10.100.1.6。CE1双宿主到PE1和PE2。

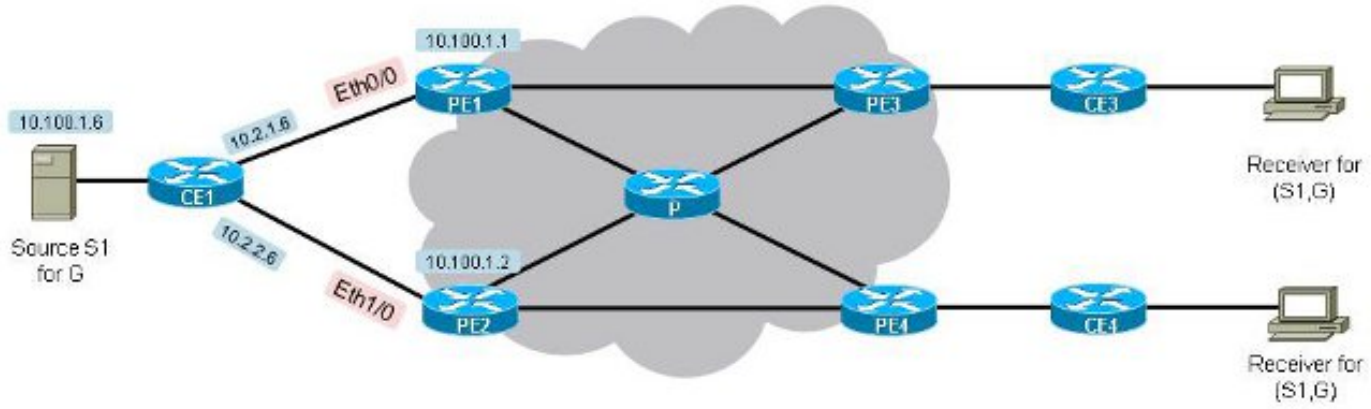

图1.双宿主源1

所有PE路由器(PE路由器上的路由识别器(RD)可以不同)上的配置如下：

```
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
!
```

为了使两个入口PE路由器开始将组播流(10.100.1.6,232.1.1.1)转发到默认MDT，它们必须都从出口PE接收加入。请看图1中的拓扑。双宿主源1。默认情况下，如果边缘链路的所有开销相同且核心链路的所有开销相同，则PE3将RPF用于PE1,PE4将RPF用于PE2(10.100.1.6,232.1.1.1)。 它们都通过RPF连接到最近的入口PE。此输出确认了这一点：

```
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具有RPF到PE2。PE3选择PE1作为RPF邻居的原因是虚拟路由/转发(VRF)中指向10.100.1.6/32的单播路由是通过PE1的最佳路由。PE3实际上从PE1和PE2接收路由10.100.1.6/32。边界网关协议(BGP)最佳路径计算算法相同，但BGP下一跳地址的开销除外。

```
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选择的最佳路径是PE1通告的路径，因为PE1的内部网关协议(IGP)开销(11)最低，而PE2的IGP开销(21)最低。对于PE4，则相反。拓扑显示，从PE3到PE1只有一跳，而从PE3到PE2有两跳。由于所有链路的IGP开销相同，因此PE3从PE1选择最佳路径。

(10.100.1.6,232.1.1.1)的组播路由信息库(MRIB)在PE1和PE2上类似于PE1和PE2，当尚未出现组播流量时：

```
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均收到(10.100.1.6、232.1.1.1)的PIM加入。 Tunnel0接口位于两台路由器上组播条目的传出接口列表(OIL)中。

组播流量开始流向(10.100.1.6,232.1.1.1)。“Debug ip pim vrf one 232.1.1”和“debug ip mrouting vrf one 232.1.1.1”显示组播流量到达两个P2P2P2P2P2P2.1.1.1PE路由器，使断言机制运行。

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
```

如果两台路由器到源10.100.1.6的度量和距离相同，则会进行一次平衡，以确定断言赢家。断路器是隧道0（默认MDT）上PIM邻居的最高IP地址。 在本例中，这是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
```

PE1已从组播条目的OIL中删除Tunnel0，因为该断言。由于OIL变为空，因此组播条目被修剪。

```
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在接口Tunnel0上设置了A标志，因为它是断言赢家。

```
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在断言计时器到期前定期在Tunnel0（默认MDT）上发送断言。因此，PE2仍然是坚定的赢家。

```
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]
```

结论

断言机制也适用于OIL中的隧道接口。当入口PE路由器在OIL中的关联隧道接口上接收C-(S，G)组播流量时，将通过默认MDT交换断言。

使用数据MDT的断言机制

在配置数据MDT的大多数时间中，断言机制仍将在默认MDT上运行，因为C-(S，G)流量仅在三秒后从默认MDT切换到数据MDT。然后，会发生与前面描述相同的情况。请注意，**每个启用组播的VRF只有一个隧道接口**:默认MDT和所有数据MDT仅使用一个隧道接口。此隧道接口用于入口PE路由器上的OIL或出口PE路由器上的RPF接口。

在某些情况下，在发出数据MDT信号之前，可能不会触发断言机制。因此，C-(S，G)组播流量可能

开始在入口PE路由器PE1和PE2上的数据MDT上转发。在这种情况下，这可能导致MPLS核心网络中永久重复的C-(S，G)组播流量。为避免这种情况，实施了以下解决方案：当入口PE路由器看到另一个入口PE路由器通告数据MDT时，PE路由器也是入口PE路由器，它会加入该数据MDT。原则上，只有出口PE路由器（具有下游接收器）才能加入数据MDT。由于入口PE路由器加入其他入口PE路由器通告的数据MDT，它导致入口PE路由器从OIL中存在的隧道接口接收组播流量，因此这会触发断言机制并导致其中一个入口PE路由器停止将C-(S，G)组播流量转发到其数据MDT（与Data MDT隧道接口），而其他入口PE（断言获胜者）可以继续将C-(S，G)组播流量转发到其数据MDT。

对于下一个示例，假设入口PE路由器PE1和PE2在默认MDT上从未看到来自彼此的C-(S，G)组播流量。流量在默认MDT上仅为三秒，因此不难理解，如果核心网络上出现临时流量丢失等情况，就会发生这种情况。

数据MDT的配置会添加到所有PE路由器。所有PE路由器（PE路由器上的RD可能不同）上的配置是：

```
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
!
```

一旦PE1和PE2看到来自源的流量，它们就会创建C-(S，G)条目。两个入口PE路由器将C-(S，G)组播流量转发到默认MDT。出口PE路由器PE3和PE4接收组播流量并转发。由于临时问题，PE2在默认MDT上看不到来自PE1的流量，反之亦然。它们都在默认MDT上发送数据MDT连接类型长度值(TLV)。

如果没有C-(S，G)流量，则在入口PE路由器上会看到以下组播状态：

```
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
```

尚未设置y标志。两个入口PE路由器在OIL中都有Tunnel0接口。这是由于PE3具有RPF，PE4具有RPF，PE2具有RPF，C-(S，G)。

当C-(S，G)的组播流量开始流动时，PE1和PE2都会转发流量。两个入口PE路由器上都超过了数据MDT的阈值，并且两个路由器都发送数据MDT加入TLV，在三秒后开始转发到其数据MDT。注意，PE1加入由PE2和PE2源的数据MDT，加入由PE1源的数据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接口（但现在从Data MDT，而不是Default MDT）上接收C-(S，G)的流量，断言机制将启动。只有PE2继续在其数据MDT上转发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中不再具有隧道接口。

```
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接口上设置了A标志：

```
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
```

结论

使用数据MDT时，断言机制也有效。当入口PE路由器在OIL中的关联隧道接口上接收C-(S，G)组播流量时，将通过默认MDT交换断言。