

Use Global Table Multicast (GTM) Non-Segmented for mVPN

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[Architecture](#)

[What are the Changes with RFC 6513/6514?](#)

[Configuration Changes](#)

[Route Advertisement](#)

[Examples](#)

[Example 1: Profile 12: Default MDT - mLDP - P2MP - BGP-AD - BGP C-mcast Signaling](#)

[Configuration](#)

[Troubleshooting](#)

[Ingress Border Router](#)

[Egress Border Router](#)

[Example 2: Profile 20 Default MDT - P2MP-TE - BGP-AD - PIM - C-mcast Signaling](#)

[Configuration](#)

[Troubleshooting](#)

[Ingress Border Router](#)

[Egress Border Router](#)

[Example 3: As in Example 1, but there is iBGP between PE and Border Router](#)

[Example 4: Seamless MPLS](#)

Introduction

This document describes the Global Table Multicast (GTM) Non-Segmented for mVPN.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

NG mVPN (RFC 6513/6514) has many profiles. Most of the profiles have Virtual Private Network (VPN) or Virtual Routing/Forwarding (VRF) at the PE routers. Some profiles (profiles 7 and are in the global context. For these profiles which are already global, BGP Auto Discovery (AD) is introduced with GTM. For the profiles that in the VRF context, the profiles are now available in the global context, with the Multicast Distribution Tree (MDT). These are all non-segmented GTM models. The procedures used by non-segmented GTM are described in IETF draft draft-ietf-bess-mvpn-global-table-mcast.

Both [RFC 7524](#) and draft draft-ietf-bess-mvpn-global-table-mcast ([RFC 7716](#)) require the GTM Source addresses to be reachable thru BGP unicast routes (either address-family ipv4 unicast or address-family ipv4 multicast).

The advantage of the draft draft-ietf-bess-mvpn-global-table-mcast over RFC 7524, is that the same procedures are kept as the ones used for regular NG mVPN (RFC 6514).

With GTM, the mVPN can be non-segmented or segmented.

Architecture

Within this article, the term border routers is used for an ABR, ASBR, or Aggregation router, that connects two segments of the network. Typically, the ABR is in Seamless MPLS networks. The ASBR is used when Inter-AS MPLS VPN is used. And, the Aggregation router is used when a GTM overlay non-segmented router connects the two parts of the core network, when either part runs a different multicast core tree protocol. For example, the Aggregation router can connect the PIM part of the core network with the mLDP part of the core network.

For any of the models, SAFI 2 can be used. The advantage is that SAFI 2 can have a different topology than SAFI 1. Hence, the RPF for multicast can be changed without changing the unicast forwarding.

A border router does not support dual-encapsulate. Meaning, that the router cannot forward multicast on two or more core-tree protocols at the same time. This can typically be used when you migrate from one core-tree to another. During migration, the ingress PE forwards onto both core-trees. This is not possible on border routers.

The GTM architecture supports non-segmented and segmented GTM. This document covers only the non-segmented GTM.

The procedures for GTM Overlay Non-Segmented are the ones described in draft-ietf-bess-mvpn-global-table-mcast. The same procedures are followed as in RFC 6513/6514 with a few changes.

What are the Changes with RFC 6513/6514?

With GTM, the next points applies. Some of these are the same as with RFC 6513/6514; some are different.

- Single Forwarding Selection (SFS) is not supported.
- AF IPv6 is supported.
- C-PIM and C-BGP signaling is supported.
- There is no VRF on the interfaces on the PE routers facing the edge. These interfaces are in global now. These routers are referred to as the Protocol Boundary Router (PBR) in the draft draft-ietf-bess-mvpn-global-table-mcast. These routers interface between an LSM core tree protocol and PIM. We call these routers the border routers.
- The core network runs a Label Switched Multicast (LSM) core tree protocol.
- mLDP, P2MP TE (static and dynamic), and IR are supported.
- Default, partitioned, and Data MDTs are supported
- Because there is no VPNv4/6 prefixes in GTM, the VRF Route-Import EC and Source-AS EC are attached to the IPv4 unicast (SAFI 1) or multicast (SAFI 2) prefixes.

Route-type 1, 3, and 5 have RTs. In Cisco IOS® XR these RTs must be present for GTM, even though this is not required as per the draft. You must configure the RTs under BGP for GTM to use. These RT are similar to the RTs used in the VRFs for regular mVPN, but apply now to the global context.

Route-type 4, 6, and 7 carry an RT that identifies the upstream PE router. The global administrator field is the IP address of the upstream PE. The local administrator field is set to 0 for GTM (it identifies the VRF in non-GTM or regular mVPN).

The PE routers become the interconnecting routers between a Label Switched Multicast (LSM) core tree protocol (mLDP, P2MP Traffic Engineering, Ingress Replication (IR)) and PIM. So, there is a part of the core network that runs LSM and we have a part of the core network that runs PIM. Lets call the core routers which act as the interface between the LSM part of the network and the PIM part of the network, the border routers. In some of the next examples, they are referred to as the C-PE routers (C for Core).


These border routers are the routers with the configuration needed for the GTM. None of the other routers are GTM-aware.

The configuration for GTM is similar to the configuration needed for the regular mVPN profiles. It is just that the interfaces towards the edge are not in a VRF.

There is no regular Route Distinguisher, because there are no VRFs. Since there are no regular Route Distinguishers (RDs), but RDs are used when signaling with BGP, all-zeros RD and all-ones RD are used for the signaling in GTM. To have this functionality, the BGP command global-table-multicast must be configured.


With GTM, the unicast routes are not in VPNv4/6. Hence, the unicast reachability must be provided in BGP in AF IPv4 or AF IPv6 and SAFI 1 or SAFI 2. That means that BGP must still be used between the border routers (PE routers without VRF).Refer to Image 1.

Image 1



Between the border and CE routers, there is no BGP. The border router adds the multicast attributes when it advertises the routes in iBGP towards the other border routers. It is important to note that BGP could be present between the CE and PE routers. Refer to Image 2.

Image 2



In this case, the PE router adds the multicast attributes when it forwards the unicast routes from eBGP into iBGP, towards the other PE routers. If the CE advertised the unicast routes with multicast attributes already to the PE router, then the PE router keeps the multicast attributes as they are and forward the unicast routes to the other PE routers. By default, for eBGP sessions, the multicast attributes are removed. So, when the PE routes advertises the unicast routes from iBGP into eBGP to the CE routes, there are no multicast attributes.

When the PE router advertises the unicast prefix through iBGP, it attaches the Extended Community (EC) VRF Route Import (VRF-RI) and the EC Source-AS. The other PE router strips these off before it propagates these routes in eBGP.

When the eBGP session is between two ASBRs, there is Inter-AS MPLS VPN and Inter-AS mVPN. In this case, the multicast attributes can be kept. Since the default behavior is to remove them on eBGP session, you would need to configure the command `send-multicast-attributes` on the eBGP session between the two ASBRs.

For the cases where we have a RR, there can be iBGP to iBGP propagation. This is the case on the inline-ABR (there is next-hop-self) of Seamless MPLS. Since the default behavior is to keep the multicast attributes for iBGP sessions, the inline-ABR needs to have the command `send-multicast-attributes-disable` in order to remove them.

Configuration Changes

You must configure `global-table-multicast` under the address family (AF) ipv4 mVPN under router BGP. This allows the operation of all-zeros RD and all-ones RD.

You must configure `import-rt` and `export-rt` under `multicast-routing` for AF ipv4 in global context. This is because there is no longer RTs configured for the VRFs, because GTM does not have VRFs. These RT must not overlap with any RTs used for regular mVPN.

The router pim commands (`rpf topology` and `mdt commands`) are now configured in global context.

The multicast-routing commands (`bgp auto-discovery` and `mdt commands`) are now configured in global context.

Route Advertisement

Between the border routers there is iBGP that advertises the Source prefixes. How can the ingress border router learn the Source prefix? There are three possibilities.

Image 3 shows these three possible scenarios.




Image 3

1. The border router received the prefixes from the PE as non-BGP prefixes. The border router needs to redistribute these prefixes into BGP. This border router adds the multicast attributes.
2. The border routers have an eBGP session towards the PE routers. The border router adds the multicast attributes before it propagates the prefix over iBGP to the other border routers. If the prefixes received over the eBGP session already have the multicast attributes, then they are kept and forwarded as is. The border router does not overwrite them.
3. The ingress border router learns the Source prefix from iBGP. In this case, the ingress border router is a RR. This scenario is used in Seamless MPLS, where the border routers are ABRs.

When the border router advertises a received iBGP prefix from another border router, it strips off the multicast attributes before it sends the prefix to the PE router. The border routers must have the command `send-multicast-attributes disable` under `router BGP` for this to happen.

Examples

Here are a few examples. The first example starts with a transformation of profile 12 into a GTM deployment.

Example 1: Profile 12: Default MDT - mLDP - P2MP - BGP-AD - BGP C-mcast Signaling

Image 4 shows this network. There is no VRF on the PE router towards the CE router.




Image 4

Note that the inner core network runs mLDP. The outer core network runs PIM. So, the border routers, that connect the PIM to mLDP core need to translate PIM to mLDP and vice versa.

The Source cannot be learned as an IGP route on the border router, router C-PE2. The IGP is ISIS here. If this is the case, then the RPF on the border router would use the ISIS route, that points to P1. If this is the case, then RPF fails as there is no PIM neighborship. You want the C-PE2 router to RPF for 10.2.1.8 and have it point to the MDT as the RPF interface. This could be an MDT based on mLDP, P2MP, or IR.

The solution is to use SAFI 2. It is used so that the Source is learned as AFI 2 route in BGP. So, the border router (C-PE2) have the Source as BGP SAFI 2 route (show route ipv4 multicast). The RPF for the Source points to the MDT interface.

Using SAFI 2 changes the RPF, and RPF for all sources now uses SAFI 2. This means that RPF for all sources in global uses SAFI 2, that includes RPF for the Ingress PE for example, for the VPN service. Once SAFI 2 is enabled, all RPF occurs through SAFI 2 only. Since only the sources are in SAFI 2, the RPF for the ingress PE routers fails. To make this work, you can configure **rump always-replicate** command under router rib. Because only RPF for the Source prefixes in global and the RPF for the PE routers must work, you can configure an access-list for the rump always-replicate command, and specify only the Sources in global and the ingress PE routers in the access-list. That way, if the border router already runs BGP for SAFI 1 and this SAFI 1 carried a large number of prefixes, these prefixes would not all be redistributed into the SAFI 2 RIB and use the memory unnecessarily.

Alternatively, you can configure distance bgp 20 20 20 for address-family ipv4 multicast under router BGP. This makes sure that if the Sources in global are also learned through AFI 2 of the IGP, that the BGP learned ones are preferred because of the lower distance of iBGP versus the distance of the IGP.

Configuration

This is the configuration of the border router.

```
hostname C-PE1

router rib
  address-family ipv4
  rump always-replicate
  !
route-policy global-one
  set core-tree mldp-default
end-policy
!
route-policy sources-in-ISIS
  if destination in (10.2.1.0/24) then
    pass
  endif
end-policy

!
router isis 1
  is-type level-1
  net 49.0001.0000.0000.0003.00
  address-family ipv4 unicast
  metric-style wide
  mpls traffic-eng level-1
  mpls traffic-eng router-id Loopback0
  !
  interface Loopback0
  address-family ipv4 unicast
  !
  address-family ipv4 multicast
  !
  !
  interface GigabitEthernet0/0/0/0
  address-family ipv4 unicast
  !
  address-family ipv4 multicast
  !
  !
  interface GigabitEthernet0/0/0/1
  address-family ipv4 unicast
  !
  address-family ipv4 multicast
  !
  !
!
router bgp 1
  address-family ipv4 unicast
  !
  address-family ipv4 multicast
  redistribute connected route-policy loopback
  redistribute isis 1 route-policy sources-in-ISIS
  !
  address-family ipv4 mvpn
  global-table-multicast
  !
  neighbor 10.100.1.5
  remote-as 1
```


You must have the next configuration on the border routers:

- BGP configured with AF ipv4 mvpn
- BGP AD enabled
- an MDT specified
- import-rt and export-rt configured under router bgp
- global-table-multicast configured under router bgp AF ipv4/6 mvpn

Optionally, the SAFI 2 must be enabled under router BGP

Troubleshooting

- First, the border routes must have the route-type 1 routes present.
- Verify the core tree in the inner core. Here, this is mLDP. So, is the mLDP signaling ok? Check the mLDP database entries for the Default MDT and possible Data MDTs.
- Check the source route in BGP.
- Check the RPF on the egress border router.
- Check the C-multicast signaling in BGP (route-type 6 & 7) on the border routers.

Ingress Border Router

The egress interface on the ingress border router is the Lmdt interface.

```
RP/0/0/CPU0:C-PE1#show mrib route 203.0.113.1 10.2.1.8
```

```
IP Multicast Routing Information Base
Entry flags: L - Domain-Local Source, E - External Source to the Domain,
             C - Directly-Connected Check, S - Signal, IA - Inherit Accept,
             IF - Inherit From, D - Drop, ME - MDT Encap, EID - Encap ID,
             MD - MDT Decap, MT - MDT Threshold Crossed, MH - MDT interface handle
             CD - Conditional Decap, MPLS - MPLS Decap, EX - Extranet
             MoFE - MoFRR Enabled, MoFS - MoFRR State, MoFP - MoFRR Primary
             MoFB - MoFRR Backup, RPFID - RPF ID Set, X - VXLAN
Interface flags: F - Forward, A - Accept, IC - Internal Copy,
                NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
                II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,
                LD - Local Disinterest, DI - Decapsulation Interface
                EI - Encapsulation Interface, MI - MDT Interface, LVIF - MPLS Encap,
                EX - Extranet, A2 - Secondary Accept, MT - MDT Threshold Crossed,
                MA - Data MDT Assigned, LMI - mLDP MDT Interface, TMI - P2MP-TE MDT Interface
                IRMI - IR MDT Interface

(10.2.1.8,203.0.113.1) RPF nbr: 10.1.2.2 Flags: RPF
Up: 00:08:58
Incoming Interface List
  GigabitEthernet0/0/0/1 Flags: A, Up: 00:08:58
Outgoing Interface List
  Lmdtdefault Flags: F LMI MA, Up: 00:08:58
```

```
RP/0/0/CPU0:C-PE1#show mfib route 203.0.113.1 10.2.1.8
```

```
IP Multicast Forwarding Information Base
```

Entry flags: C - Directly-Connected Check, S - Signal, D - Drop,
IA - Inherit Accept, IF - Inherit From, EID - Encap ID,
ME - MDT Encap, MD - MDT Decap, MT - MDT Threshold Crossed,
MH - MDT interface handle, CD - Conditional Decap,
DT - MDT Decap True, EX - Extranet, RPFID - RPF ID Set,
MoFE - MoFRR Enabled, MoFS - MoFRR State, X - VXLAN
Interface flags: F - Forward, A - Accept, IC - Internal Copy,
NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
EG - Egress, EI - Encapsulation Interface, MI - MDT Interface,
EX - Extranet, A2 - Secondary Accept
Forwarding/Replication Counts: Packets in/Packets out/Bytes out
Failure Counts: RPF / TTL / Empty Olist / Encap RL / Other

(10.2.1.8,203.0.113.1), Flags:
Up: 01:47:24
Last Used: 00:00:00
SW Forwarding Counts: 1197/1197/239400
SW Replication Counts: 1197/0/0
SW Failure Counts: 0/0/0/0/0
Lmtdefault Flags: F LMI, Up:01:47:24
GigabitEthernet0/0/0/1 Flags: A, Up:01:47:24

RP/0/0/CPU0:C-PE1#show route ipv4 multicast

Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
U - per-user static route, o - ODR, L - local, G - DAGR, l - LISP
A - access/subscriber, a - Application route
M - mobile route, r - RPL, (!) - FRR Backup path

Gateway of last resort is not set

```
i L1 10.1.1.0/24 [255/20] via 10.1.2.2, 1d21h, GigabitEthernet0/0/0/1
C 10.1.2.0/24 is directly connected, 1d21h, GigabitEthernet0/0/0/1
L 10.1.2.3/32 is directly connected, 3d19h, GigabitEthernet0/0/0/1
i L1 10.1.3.0/24 [115/20] via 10.1.3.4, 3d13h, GigabitEthernet0/0/0/0
L 10.1.3.3/32 is directly connected, 3d19h, GigabitEthernet0/0/0/0
i L1 10.1.4.0/24 [115/20] via 10.1.3.4, 3d13h, GigabitEthernet0/0/0/0
i L1 10.1.5.0/24 [115/30] via 10.1.3.4, 3d12h, GigabitEthernet0/0/0/0
i L1 10.1.6.0/24 [255/40] via 10.1.3.4, 1d21h, GigabitEthernet0/0/0/0
i L1 10.2.1.0/24 [255/30] via 10.1.2.2, 1d21h, GigabitEthernet0/0/0/1
i L1 10.2.2.0/24 [255/50] via 10.1.3.4, 1d21h, GigabitEthernet0/0/0/0
i L1 10.100.1.1/32 [255/30] via 10.1.2.2, 1d21h, GigabitEthernet0/0/0/1
i L1 10.100.1.2/32 [255/20] via 10.1.2.2, 1d21h, GigabitEthernet0/0/0/1
L 10.100.1.3/32 is directly connected, 1d21h, Loopback0
i L1 10.100.1.4/32 [115/20] via 10.1.3.4, 3d13h, GigabitEthernet0/0/0/0
i L1 10.100.1.5/32 [115/30] via 10.1.3.4, 3d12h, GigabitEthernet0/0/0/0
i L1 10.100.1.6/32 [255/40] via 10.1.3.4, 1d21h, GigabitEthernet0/0/0/0
i L1 10.100.1.7/32 [255/50] via 10.1.3.4, 1d21h, GigabitEthernet0/0/0/0
```

RP/0/0/CPU0:C-PE1#show pim rpf 10.2.1.8

Table: IPv4-Multicast-default
* 10.2.1.8/32 [255/30]
via GigabitEthernet0/0/0/1 with rpf neighbor 10.1.2.2

Egress Border Router

For the Source route, the VRF Route-Import EC and Source-AS EC are attached to the IPv4 unicast or multicast prefix. Here, it is an IPv4 multicast route:

```
RP/0/0/CPU0:C-PE2#show bgp ipv4 multicast 10.2.1.0/24
BGP routing table entry for 10.2.1.0/24
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          32        32
Last Modified: Sep 12 08:34:56.441 for 15:09:58
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
Local
  10.100.1.3 (metric 30) from 10.100.1.3 (10.100.1.3)
  Origin incomplete, metric 30, localpref 100, valid, internal, best, group-best
  Received Path ID 0, Local Path ID 1, version 32
  Extended community: VRF Route Import:10.100.1.3:0 Source AS:1:0
```

Note: If for whatever reason the VRF RI EC and Source AS EC are not there, then the RPF on the Egress border router fails.

An example when the route does not have these ECs:

```
RP/0/0/CPU0:C-PE2#show bgp ipv4 multicast 10.2.1.0/24
BGP routing table entry for 10.2.1.0/24
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          277       277
Last Modified: Sep 13 04:08:37.441 for 00:00:02
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
Local
  10.100.1.3 (metric 30) from 10.100.1.3 (10.100.1.1)
  Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best
  Received Path ID 0, Local Path ID 1, version 277
  Originator: 10.100.1.1, Cluster list: 10.100.1.3
```

Because of this, the RPF fails:

```
RP/0/0/CPU0:C-PE2#show pim rpf 10.2.1.8

Table: IPv4-Multicast-default
* 10.2.1.8/32 [200/30]
  via Null with rpf neighbor 0.0.0.0
```

```
RP/0/0/CPU0:C-PE2#show bgp ipv4 mvpn

BGP router identifier 10.100.1.5, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0   RD version: 0
```

```
BGP main routing table version 56
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
Global table multicast is enabled
BGP scan interval 60 secs
```

```
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 0:0:0					
*>i[1][10.100.1.3]/40	10.100.1.3			100	0 i
*> [1][10.100.1.5]/40	0.0.0.0				0 i
*>i[3][32][10.2.1.8][32][203.0.113.1][10.100.1.3]/120					
	10.100.1.3			100	0 i
*> [7][0:0:0][1][32][10.2.1.8][32][203.0.113.1]/184					
	0.0.0.0				0 i

```
Processed 4 prefixes, 4 paths
```

The command can be specified with the keywords `rd all-zero-rd`. It then shows all the entries with the all-zeros RD.

```
RP/0/0/CPU0:C-PE2#show bgp ipv4 mvpn rd all-zero-rd
```

```
BGP router identifier 10.100.1.5, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 56
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
Global table multicast is enabled
BGP scan interval 60 secs
```

```
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 0:0:0					
*>i[1][10.100.1.3]/40	10.100.1.3			100	0 i
*> [1][10.100.1.5]/40	0.0.0.0				0 i
*>i[3][32][10.2.1.8][32][203.0.113.1][10.100.1.3]/120					
	10.100.1.3			100	0 i
*> [7][0:0:0][1][32][10.2.1.8][32][203.0.113.1]/184					
	0.0.0.0				0 i

```
Processed 4 prefixes, 4 paths
```

The type 1 route:

```
RP/0/0/CPU0:C-PE2#show bgp ipv4 mvpn rd all-zero-rd [1][10.100.1.3]/40
```

```
BGP routing table entry for [1][10.100.1.3]/40, Route Distinguisher: 0:0:0
```

```
Versions:
```

Process	bRIB/RIB	SendTblVer
Speaker	43	43

```
Last Modified: Sep 8 07:42:43.786 for 1d17h
```

```
Paths: (1 available, best #1, not advertised to EBGP peer)
```

```
Not advertised to any peer
Path #1: Received by speaker 0
Not advertised to any peer
Local
 10.100.1.3 (metric 30) from 10.100.1.3 (10.100.1.3)
  Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate, imported
  Received Path ID 0, Local Path ID 1, version 43
  Community: no-export
  Extended community: RT:1:1
  PMSI: flags 0x00, type 2, label 0, ID 0x060001040a640103000701000400000001
  Source AFI: IPv4 MVPN, Source VRF: default, Source Route Distinguisher: 0:0:0
```

The PMSI decoded:

PMSI: flags 0x00, type 2, label 0, ID 0x060001040a640103000701000400000001

The decoded PMSI from the previous command is:

The PMSI Tunnel Type is : 2 : mLDP P2MP LSP The PMSI Tunnel ID is : 0x060001040a640103000701000400000001 FEC Element FEC Element Type : 6 : P2MP AF Type : 1 Address Length : 4 Root Node Address : 10.100.1.3 MP Opaque Length : 7 MP Opaque Value Element Opaque Type : 1 : LSP ID Global Opaque Length : 4 Global ID (Generic LSP Identifier) : 1

The Data MDT is signaled by a route-type 3 AD route from C-PE1.

```
RP/0/0/CPU0:C-PE2#show bgp ipv4 mvpn rd all-zero-rd [3][32][10.2.1.8]
[32][203.0.113.1][10.100.1.3]/120
```

BGP routing table entry for [3][32][10.2.1.8][32][203.0.113.1][10.100.1.3]/120, Route Distinguisher: 0:0:0

Versions:

Process	bRIB/RIB	SendTblVer
Speaker	56	56

Last Modified: Sep 10 00:51:52.786 for 00:04:57

Paths: (1 available, best #1, not advertised to EBGP peer)

Not advertised to any peer

Path #1: Received by speaker 0

Not advertised to any peer

Local

```
 10.100.1.3 (metric 30) from 10.100.1.3 (10.100.1.3)
  Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate, imported
  Received Path ID 0, Local Path ID 1, version 56
  Community: no-export
  Extended community: RT:1:1
  PMSI: flags 0x00, type 2, label 0, ID 0x060001040a640103000701000400000007
  Source AFI: IPv4 MVPN, Source VRF: default, Source Route Distinguisher: 0:0:0
```

The decoded PMSI shows that the Global LSP Identifier is 7. This is then used for the mLDP database entry for this Data MDT.

PMSI: flags 0x00, type 2, label 0, ID 0x060001040a640103000701000400000007

The decoded PMSI from the previous command is:

The PMSI Tunnel Type is : 2 : mLDP P2MP LSP The PMSI Tunnel ID is : 0x060001040a640103000701000400000007 FEC Element FEC Element Type : 6 : P2MP AF Type : 1 Address Length : 4 Root Node Address : 10.100.1.3 MP Opaque Length : 7 MP Opaque Value Element Opaque Type : 1 : LSP ID Global Opaque Length : 4 Global ID (Generic LSP Identifier) : 7

With the next commands, you can check what the Ingress PE advertises about the Data MDT. Notice that this is GTM, so there is no VRF in the next command.

```
RP/0/0/CPU0:C-PE2#show pim mdt mldp remote
```

Core Identifier	MDT Source	Cache Count	Max Agg	DIP	Local Entry	VRF Using	Routes Cache
[global-id 7]	10.100.1.3	1	255	N	N		1

```
RP/0/0/CPU0:C-PE2#show pim mdt mldp cache
```

Core Source	Cust (Source, Group)	Core Data	Expires
10.100.1.3	(10.2.1.8, 203.0.113.1)	[global-id 7]	never

Route-type 7 does not have a PMSI attached:

```
RP/0/0/CPU0:C-PE2#show bgp ipv4 mvpn rd all-zero-rd [7][0:0:0][1][32][10.2.1.8][32][203.0.113.1]/184
```

BGP routing table entry for [7][0:0:0][1][32][10.2.1.8][32][203.0.113.1]/184, Route Distinguisher: 0:0:0

Versions:

Process	bRIB/RIB	SendTblVer
Speaker	52	52

Last Modified: Sep 10 00:51:51.786 for 00:07:37

Paths: (1 available, best #1)

Advertised to peers (in unique update groups):

10.100.1.3

Path #1: Received by speaker 0

Advertised to peers (in unique update groups):

10.100.1.3

Local

0.0.0.0 from 0.0.0.0 (10.100.1.5)

Origin IGP, localpref 100, valid, redistributed, best, group-best, import-candidate

Received Path ID 0, Local Path ID 1, version 52

Extended community: RT:10.100.1.3:0

The RT identifies the upstream PE router. The global administrator field is the IP address of the upstream PE. The local administrator field is set to 0 for GTM.

```
RP/0/0/CPU0:C-PE2#show mrib route 203.0.113.1 10.2.1.8
```

IP Multicast Routing Information Base

Entry flags: L - Domain-Local Source, E - External Source to the Domain,

C - Directly-Connected Check, S - Signal, IA - Inherit Accept,

IF - Inherit From, D - Drop, ME - MDT Encap, EID - Encap ID,

MD - MDT Decap, MT - MDT Threshold Crossed, MH - MDT interface handle

CD - Conditional Decap, MPLS - MPLS Decap, EX - Extranet

MoFE - MoFRR Enabled, MoFS - MoFRR State, MoFP - MoFRR Primary

MoFB - MoFRR Backup, RPFID - RPF ID Set, X - VXLAN

Interface flags: F - Forward, A - Accept, IC - Internal Copy,

NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,

II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,
LD - Local Disinterest, DI - Decapsulation Interface
EI - Encapsulation Interface, MI - MDT Interface, LVIF - MPLS Encap,
EX - Extranet, A2 - Secondary Accept, MT - MDT Threshold Crossed,
MA - Data MDT Assigned, LMI - mLDP MDT Interface, TMI - P2MP-TE MDT Interface
IRMI - IR MDT Interface

```
(10.2.1.8,203.0.113.1) RPF nbr: 10.100.1.3 Flags: RPF
Up: 00:52:34
Incoming Interface List
  Lmdtdefault Flags: A LMI, Up: 00:52:34
Outgoing Interface List
  GigabitEthernet0/0/0/0 Flags: F NS, Up: 00:52:34
```

The incoming interface must be an Lmdt interface.

```
RP/0/0/CPU0:C-PE2#show mfib route 203.0.113.1 10.2.1.8
```

```
IP Multicast Forwarding Information Base
Entry flags: C - Directly-Connected Check, S - Signal, D - Drop,
  IA - Inherit Accept, IF - Inherit From, EID - Encap ID,
  ME - MDT Encap, MD - MDT Decap, MT - MDT Threshold Crossed,
  MH - MDT interface handle, CD - Conditional Decap,
  DT - MDT Decap True, EX - Extranet, RPFID - RPF ID Set,
  MoFE - MoFRR Enabled, MoFS - MoFRR State, X - VXLAN
Interface flags: F - Forward, A - Accept, IC - Internal Copy,
  NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
  EG - Egress, EI - Encapsulation Interface, MI - MDT Interface,
  EX - Extranet, A2 - Secondary Accept
Forwarding/Replication Counts: Packets in/Packets out/Bytes out
Failure Counts: RPF / TTL / Empty Olist / Encap RL / Other
```

```
(10.2.1.8,203.0.113.1),  Flags:
Up: 02:31:00
Last Used: never
SW Forwarding Counts: 0/2037/407400
SW Replication Counts: 0/2037/407400
SW Failure Counts: 0/0/0/0/0
Lmdtdefault Flags:  A LMI, Up:02:31:00
GigabitEthernet0/0/0/0 Flags:  NS EG, Up:02:31:00
```

Check the SAFI 2 routes:

```
RP/0/0/CPU0:C-PE2#show route ipv4 multicast
```

```
Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path
  D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
  N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
  E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
  i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2
  ia - IS-IS inter area, su - IS-IS summary null, * - candidate default
  U - per-user static route, o - ODR, L - local, G - DAGR, l - LISP
  A - access/subscriber, a - Application route
  M - mobile route, r - RPL, (!) - FRR Backup path
```

```
Gateway of last resort is not set
```

```
i L1 10.1.2.0/24 [115/30] via 10.1.4.4, 3d12h, GigabitEthernet0/0/0/1
```



```

i L1 10.1.3.0/24 [115/20] via 10.1.4.4, 3d12h, GigabitEthernet0/0/0/1
C    10.1.4.0/24 is directly connected, 1d21h, GigabitEthernet0/0/0/1
L    10.1.4.5/32 is directly connected, 3d12h, GigabitEthernet0/0/0/1
C    10.1.5.0/24 is directly connected, 1d21h, GigabitEthernet0/0/0/0
L    10.1.5.5/32 is directly connected, 3d12h, GigabitEthernet0/0/0/0
B    10.2.1.0/24 [200/30] via 10.100.1.3, 1d17h
i L1 10.100.1.3/32 [115/30] via 10.1.4.4, 3d12h, GigabitEthernet0/0/0/1
i L1 10.100.1.4/32 [115/20] via 10.1.4.4, 3d12h, GigabitEthernet0/0/0/1
L    10.100.1.5/32 is directly connected, 1d21h, Loopback0

```

Notice that the route for the Source is SAFI 2 (it is in AF IPv4 multicast), because it is in the RIB AF IPv4 multicast.

Notice that the next-hop is 10.100.1.3, the loopback of C-PE1, because that router has next-hop-self under AF ipv4 multicast under router BGP.

```

RP/0/0/CPU0:C-PE2#show bgp ipv4 multicast 10.2.1.0/24
BGP routing table entry for 10.2.1.0/24
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          34        34
Last Modified: Sep  8 07:42:18.786 for 1d17h
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
Local
  10.100.1.3 (metric 30) from 10.100.1.3 (10.100.1.3)
    Origin incomplete, metric 30, localpref 100, valid, internal, best, group-best
    Received Path ID 0, Local Path ID 1, version 34
    Extended community: VRF Route Import:10.100.1.3:0 Source AS:1:0

```

The RPF for the Source points to the Lmdt interface and the PIM neighbor across it. The RPF is performed in the IPv4 Multicast table.

```

RP/0/0/CPU0:C-PE2#show pim rpf 10.2.1.8

Table: IPv4-Multicast-default
* 10.2.1.8/32 [200/30]
  via Lmdtdefault with rpf neighbor 10.100.1.3

```

Check that the ingress border router is recognized as a PE router.

```

RP/0/0/CPU0:C-PE2#show pim pe

MVPN Provider Edge Router information

PE Address : 10.100.1.3 (0x1071da64)
RD: 0:0:0 (valid), RIB_HLI 0, RPF-ID 3, Remote RPF-ID 0, State: 1, S-PMSI: 2
PPMP_LABEL: 0, MS_PMSI_HLI: 0x00000, Bidir_PMSI_HLI: 0x00000, MLDP-added: [RD 0, ID 0, Bidir ID 0, Remote Bidir ID 0], Counts(SHR/SRC/DM/DEF-MD): 0, 1, 0, 0, Bidir: GRE RP Count 0, MPLS RP Count 0RSVP-TE added: [Leg 0, Ctrl Leg 0, Part tail 0 Def Tail 0, IR added: [Def Leg 0, Ctrl Leg 0, Part Leg 0, Part tail 0, Part IR Tail Label 0
  bgp_i_pmsi: 1,0/0 , bgp_ms_pmsi/Leaf-ad: 0/0, bgp_bidir_pmsi: 0, remote_bgp_bidir_pmsi: 0,
PMSIs: I 0x106a2d50, 0x0, MS 0x0, Bidir Local: 0x0, Remote: 0x0, BSR/Leaf-ad 0x0/0, Autorp-

```

```
disc/Leaf-ad 0x0/0, Autorp-ann/Leaf-ad 0x0/0
IIDs: I/6: 0x1/0x0, B/R: 0x0/0x0, MS: 0x0, B/A/A: 0x0/0x0/0x0
```

```
Bidir RPF-ID: 4, Remote Bidir RPF-ID: 0
I-PMSI: MLDP-P2MP, Opaque: [global-id 1] (0x106a2d50)
I-PMSI rem: (0x0)
MS-PMSI: (0x0)
Bidir-PMSI: (0x0)
Remote Bidir-PMSI: (0x0)
BSR-PMSI: (0x0)
A-Disc-PMSI: (0x0)
A-Ann-PMSI: (0x0)
RIB Dependency List: 0x1016446c
Bidir RIB Dependency List: 0x0
Sources: 1, RPs: 0, Bidir RPs: 0
```

The Inclusive PMSI (I-PMSI) is there.

You see the two P2MP mLDP entries that form the Default MDT between the two border routers in the mLDP database. There is also one P2MP mLDP entry with C-PE1 as root for the Data MDT.

```
RP/0/0/CPU0:C-PE2#show mpls mldp database brief
```

LSM ID	Type	Root	Up	Down	Decoded	Opaque Value
0x00007	P2MP	10.100.1.3	1	1	[global-id 1]	
0x00008	P2MP	10.100.1.5	0	2	[global-id 1]	
0x0000B	P2MP	10.100.1.3	1	1	[global-id 7]	

Example 2: Profile 20 Default MDT - P2MP-TE - BGP-AD - PIM - C-mcast Signaling

This is very similar as example 1. Now there is P2MP TE in the core. The tunnels are set up as auto tunnels. The tail end routers are discovered through BGP AD. Another difference with example 1, is that the overlay protocol is now PIM. Look at **Image 5**.




Image 5

Configuration

This is the configuration of the border router:

```
hostname C-PE1
logging console debugging
router rib
  address-family ipv4
  rump always-replicate
  !
!
line default
  timestamp disable
  exec-timeout 0 0
!
ipv4 unnumbered mpls traffic-eng Loopback0
interface Loopback0
  ipv4 address 10.100.1.3 255.255.255.255
!
interface MgmtEth0/0/CPU0/0
  shutdown
!
interface GigabitEthernet0/0/0/0
  ipv4 address 10.1.3.3 255.255.255.0
  load-interval 30
!
interface GigabitEthernet0/0/0/1
  ipv4 address 10.1.2.3 255.255.255.0
!
interface GigabitEthernet0/0/0/2
  shutdown
!
interface GigabitEthernet0/0/0/3
  shutdown
!
interface GigabitEthernet0/0/0/4
  shutdown
!
interface GigabitEthernet0/0/0/5
  shutdown
!
interface GigabitEthernet0/0/0/6
  shutdown
!
interface GigabitEthernet0/0/0/7
  shutdown
!
interface GigabitEthernet0/0/0/8
  shutdown
!
route-policy loopback
  if destination in (10.100.1.3/32) then
    pass
  endif
end-policy
!
route-policy global-one
  set core-tree p2mp-te-default
end-policy
```

```
!  
route-policy sources-in-ISIS  
  if destination in (10.2.1.0/24) then  
    pass  
  endif  
end-policy  
!  
router isis 1  
  is-type level-1  
  net 49.0001.0000.0000.0003.00  
  address-family ipv4 unicast  
metric-style wide  
mpls traffic-eng level-1  
mpls traffic-eng router-id Loopback0  
!  
interface Loopback0  
address-family ipv4 unicast  
!  
address-family ipv4 multicast  
!  
!  
interface GigabitEthernet0/0/0/0  
address-family ipv4 unicast  
!  
address-family ipv4 multicast  
!  
!  
interface GigabitEthernet0/0/0/1  
address-family ipv4 unicast  
!  
address-family ipv4 multicast  
!  
!  
!  
router bgp 1  
  address-family ipv4 unicast  
  !  
  address-family ipv4 multicast  
  redistribute connected route-policy loopback  
  redistribute ospf 1  
  redistribute isis 1 route-policy sources-in-ISIS  
  !  
  address-family ipv4 mvpn  
global-table-multicast  
!  
neighbor 10.100.1.5  
remote-as 1  
update-source Loopback0  
address-family ipv4 multicast  
  next-hop-self  
!  
address-family ipv4 mvpn  
!  
!  
!  
mpls oam  
!  
rsvp  
  interface GigabitEthernet0/0/0/0  
  bandwidth 1000000  
  !  
  interface GigabitEthernet0/0/0/1  
  bandwidth 1000000  
  !
```

```
!  
mpls traffic-eng  
  interface GigabitEthernet0/0/0/0  
  auto-tunnel backup  
  !  
  !  
  interface GigabitEthernet0/0/0/1  
  auto-tunnel backup  
  !  
  !  
  auto-tunnel p2mp  
  tunnel-id min 1000 max 2000  
  !  
!  
mpls ldp  
  log  
  neighbor  
  !  
  mldp  
  logging notifications  
  address-family ipv4  
    rib unicast-always  
  !  
  !  
  router-id 10.100.1.3  
  address-family ipv4  
  !  
  interface GigabitEthernet0/0/0/0  
  address-family ipv4  
  !  
  !  
  interface GigabitEthernet0/0/0/1  
  address-family ipv4  
  !  
  !  
!  
multicast-routing  
address-family ipv4  
  interface Loopback0  
  enable  
  !  
  interface GigabitEthernet0/0/0/1  
  enable  
  !  
mdt source Loopback0  
  export-rt 1:1  
  import-rt 1:1  
  bgp auto-discovery p2mp-te  
  !  
  mdt default p2mp-te  
  mdt data p2mp-te 100 immediate-switch  
  !  
!  
router pim  
  address-family ipv4  
rpf topology route-policy global-one  
  interface Loopback0  
  enable  
  !  
  interface GigabitEthernet0/0/0/1  
  !  
  !  
!
```

Troubleshooting

Ingress Border Router

Check that the RD all-zero are there. The route-type 1 routes must be there in order to build the P2MP TE based on P2MP TE tunnels.

```
RP/0/0/CPU0:C-PE1#show bgp ipv4 mvpn rd all-zero-rd
BGP router identifier 10.100.1.3, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0   RD version: 0
BGP main routing table version 140
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
Global table multicast is enabled
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network                Next Hop                Metric LocPrf Weight Path
Route Distinguisher: 0:0:0
*> [1][10.100.1.3]/40 0.0.0.0                                0 i
*>i[1][10.100.1.5]/40 10.100.1.5                          100   0 i

Processed 2 prefixes, 2 paths
```

Check the route-type 1 route in more detail:

```
RP/0/0/CPU0:C-PE1#show bgp ipv4 mvpn rd all-zero-rd [1][10.100.1.5]/40

BGP routing table entry for [1][10.100.1.5]/40, Route Distinguisher: 0:0:0
Versions:
  Process          bRIB/RIB   SendTblVer
  Speaker          135        135
Last Modified: Sep 12 08:21:42.207 for 00:20:14
Paths: (1 available, best #1, not advertised to EBGp peer)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
Local
  10.100.1.5 (metric 30) from 10.100.1.5 (10.100.1.5)
    Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate, imported
    Received Path ID 0, Local Path ID 1, version 135
    Community: no-export
    Extended community: RT:1:1
    PMSI: flags 0x00, type 1, label 0, ID 0x0000003e8000003e80a640105
    Source AFI: IPv4 MVPN, Source VRF: default, Source Route Distinguisher: 0:0:0
```

Check the PIM neighbors on the MDT Default:

```
RP/0/0/CPU0:C-PE1#show pim neighbor
```

PIM neighbors in VRF default

Flag: B - Bidir capable, P - Proxy capable, DR - Designated Router,
E - ECMP Redirect capable
* indicates the neighbor created for this router

Neighbor Address	Interface	Uptime	Expires	DR pri	Flags
10.1.2.2	GigabitEthernet0/0/0/1	6d02h	00:01:16	1	B
10.1.2.3*	GigabitEthernet0/0/0/1	6d02h	00:01:15	1 (DR)	B E
10.100.1.3*	Loopback0	6d02h	00:01:32	1 (DR)	B E
10.100.1.3*	Tmtddefault	00:36:21	00:01:40	1	
10.100.1.5	Tmtddefault	00:17:37	00:01:26	1 (DR)	

Check the MRIB route. The outgoing interface must be Tmtd:

```
RP/0/0/CPU0:C-PE1#show mrib route 203.0.113.1
```

IP Multicast Routing Information Base

Entry flags: L - Domain-Local Source, E - External Source to the Domain,
C - Directly-Connected Check, S - Signal, IA - Inherit Accept,
IF - Inherit From, D - Drop, ME - MDT Encap, EID - Encap ID,
MD - MDT Decap, MT - MDT Threshold Crossed, MH - MDT interface handle
CD - Conditional Decap, MPLS - MPLS Decap, EX - Extranet
MoFE - MoFRR Enabled, MoFS - MoFRR State, MoFP - MoFRR Primary
MoFB - MoFRR Backup, RPFID - RPF ID Set, X - VXLAN

Interface flags: F - Forward, A - Accept, IC - Internal Copy,
NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,
LD - Local Disinterest, DI - Decapsulation Interface
EI - Encapsulation Interface, MI - MDT Interface, LVIF - MPLS Encap,
EX - Extranet, A2 - Secondary Accept, MT - MDT Threshold Crossed,
MA - Data MDT Assigned, LMI - mLDP MDT Interface, TMI - P2MP-TE MDT Interface
IRMI - IR MDT Interface

```
(10.2.1.8,203.0.113.1) RPF nbr: 10.1.2.2 Flags: RPF  
Up: 00:09:10  
Incoming Interface List  
GigabitEthernet0/0/0/1 Flags: A, Up: 00:09:10  
Outgoing Interface List  
Tmtddefault Flags: F NS TMI, Up: 00:09:10
```

Check that there is one P2MP TE tunnel per border router as head end router:

```
RP/0/0/CPU0:C-PE1#show mpls traffic-eng tunnels tabular
```

Tunnel Name	LSP ID	Destination Address	Source Address	State	FRR State	LSP Role	Path Prot
^tunnel-mtel1001	10004	10.100.1.5	10.100.1.3	up	Inact	Head	
auto_C-PE2_mt1000	10005	10.100.1.3	10.100.1.5	up	Inact	Tail	

^ = automatically created P2MP tunnel

Once the Data MDT is triggered, we the route-type 3 & 4 routes:

```
RP/0/0/CPU0:C-PE1#show bgp ipv4 mvpn rd all-zero-rd
```

```
BGP router identifier 10.100.1.3, local AS number 1
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 143
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0
Global table multicast is enabled
BGP scan interval 60 secs
```

```
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale, N NextHop-discard
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 0:0:0					
*> [1][10.100.1.3]/40	0.0.0.0			0	i
*>i[1][10.100.1.5]/40	10.100.1.5			100	0 i
*> [3][32][10.2.1.8][32][203.0.113.1][10.100.1.3]/120	0.0.0.0			0	i
*>i[4][3][0:0:0][32][10.2.1.8][32][203.0.113.1][10.100.1.3][10.100.1.5]/224	10.100.1.5			100	0 i

```
Processed 4 prefixes, 4 paths
```

The route-type 3, announces to all tail end routers that a Data MDT is signalled:

```
RP/0/0/CPU0:C-PE1#show bgp ipv4 mvpn rd all-zero-rd
[3][32][10.2.1.8][32][203.0.113.1][10.100.1.3]/120
```

```
BGP routing table entry for [3][32][10.2.1.8][32][203.0.113.1][10.100.1.3]/120, Route
Distinguisher: 0:0:0
```

```
Versions:
```

Process	bRIB/RIB	SendTblVer
Speaker	141	141

```
Last Modified: Sep 12 08:46:17.207 for 00:00:41
```

```
Paths: (1 available, best #1, not advertised to EBGp peer)
```

```
Advertised to peers (in unique update groups):
```

```
10.100.1.5
```

```
Path #1: Received by speaker 0
```

```
Advertised to peers (in unique update groups):
```

```
10.100.1.5
```

```
Local
```

```
0.0.0.0 from 0.0.0.0 (10.100.1.3)
```

```
Origin IGP, localpref 100, valid, redistributed, best, group-best, import-candidate
```

```
Received Path ID 0, Local Path ID 1, version 141
```

```
Community: no-export
```

```
Extended community: RT:1:1
```

```
PMSI: flags 0x01, type 1, label 0, ID 0x000003ed000003ed0a640103
```

The PMSI decoded:

```
PMSI: flags 0x01, type 1, label 0, ID 0x000003ed000003ed0a640103
```

The decoded PMSI from the previous command is:

The PMSI Tunnel Type is : 1 : RSVP-TE P2MP LSP The PMSI Tunnel ID is : 0x000003ed000003ed0a640103 Extended Tunnel ID : 1005 Reserved part (should be zero): 0X0000 Tunnel ID : 1005 P2MP ID : 10.100.1.3

This can also be seen here:

```
RP/0/0/CPU0:C-PE1#show pim mdt cache
```

```
Core Source      Cust (Source, Group)      Core Data      Expires
10.100.1.3      (10.2.1.8, 203.0.113.1)  [p2mp 6]      never

Leaf AD: 10.100.1.5
```

The route-type 4 announces to the head end router which router is the tail end:

```
RP/0/0/CPU0:C-PE1#show bgp ipv4 mvpn rd all-zero-rd
```

```
[4][3][0:0:0][32][10.2.1.8][32][203.0.113.1][10.100.1.3][10.100.1.5]/224
```

BGP routing table entry for

```
[4][3][0:0:0][32][10.2.1.8][32][203.0.113.1][10.100.1.3][10.100.1.5]/224, Route Distinguisher: 0:0:0
```

Versions:

```
Process          bRIB/RIB  SendTblVer
Speaker          143      143
```

Last Modified: Sep 12 08:46:17.207 for 00:01:25

Paths: (1 available, best #1)

Not advertised to any peer

Path #1: Received by speaker 0

Not advertised to any peer

Local

10.100.1.5 (metric 30) from 10.100.1.5 (10.100.1.5)

Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate, imported

Received Path ID 0, Local Path ID 1, version 143

Extended community: SEG-NH:10.100.1.5:0 RT:10.100.1.3:0

Source AFI: IPv4 MVPN, Source VRF: default, Source Route Distinguisher: 0:0:0

Check that the Data MDT the P2MP TE tunnel- is set up:

```
RP/0/0/CPU0:C-PE1#show mpls traffic-eng tunnels tabular
```

Tunnel Name	LSP ID	Destination Address	Source Address	State	FRR State	LSP Role	Path Prot
^tunnel-mte1001	10004	10.100.1.5	10.100.1.3	up	Inact	Head	
^tunnel-mte1005	10002	10.100.1.5	10.100.1.3	up	Inact	Head	
auto_C-PE2_mt1000	10005	10.100.1.3	10.100.1.5	up	Inact	Tail	

^ = automatically created P2MP tunnel

Egress Border Router

Check that the incoming interface is the Tmdt interface:

```
RP/0/0/CPU0:C-PE2#show mrib route 203.0.113.1
```

IP Multicast Routing Information Base

Entry flags: L - Domain-Local Source, E - External Source to the Domain,

C - Directly-Connected Check, S - Signal, IA - Inherit Accept,
 IF - Inherit From, D - Drop, ME - MDT Encap, EID - Encap ID,
 MD - MDT Decap, MT - MDT Threshold Crossed, MH - MDT interface handle
 CD - Conditional Decap, MPLS - MPLS Decap, EX - Extranet
 MoFE - MoFRR Enabled, MoFS - MoFRR State, MoFP - MoFRR Primary
 MoFB - MoFRR Backup, RPFID - RPF ID Set, X - VXLAN

Interface flags: F - Forward, A - Accept, IC - Internal Copy,
 NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
 II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,
 LD - Local Disinterest, DI - Decapsulation Interface
 EI - Encapsulation Interface, MI - MDT Interface, LVIF - MPLS Encap,
 EX - Extranet, A2 - Secondary Accept, MT - MDT Threshold Crossed,
 MA - Data MDT Assigned, LMI - mLDP MDT Interface, **TMI - P2MP-TE MDT Interface**
 IRMI - IR MDT Interface

```
(10.2.1.8,203.0.113.1) RPF nbr: 10.100.1.3 Flags: RPF
Up: 00:18:03
Incoming Interface List
  Tmtddefault Flags: A TMI, Up: 00:18:00
Outgoing Interface List
  GigabitEthernet0/0/0/0 Flags: F NS, Up: 00:18:03
```

The RPF on the egress border router points to the ingress border router. The ingress interface is Tmtddefault. Notice the T for TE Tunnel:

```
RP/0/0/CPU0:C-PE2#show pim rpf 10.2.1.8
```

```
Table: IPv4-Multicast-default
* 10.2.1.8/32 [200/30]
  via Tmtddefault with rpf neighbor 10.100.1.3
```

Example 3: As in Example 1, but there is iBGP between PE and Border Router

Look at Image 6.

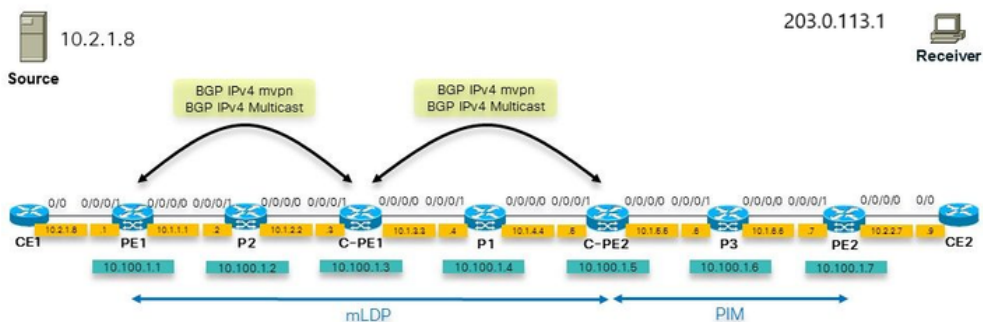


Image 6

We see an asymmetric setup where we have one core network with mLDP on one side and PIM still on the other side and GTM. This can occur during migration of core trees. The C-PE1 router must be a RR for BGP IPv4 multicast and BGP IPv4 mVPN. The configuration for PIM and multicast-routing that we had on C-PE1 in example 1 is now needed on PE1.

Example 4: Seamless MPLS

We deploy GTM over Seamless MPLS (Unified MPLS). The PE router must understand GTM, which only an Cisco IOS XR router can do and the PE router must originate the PIM RPF-Proxy vector in the PIM domain. This PIM RPF-Proxy vector is needed so that the P routers can RPF to the proxy IP address (the ABR). Since Cisco IOS XR 5.3.2, Cisco IOS XR can originate the RPF-Proxy Vector in global context. So, GTM can have the RPF-Proxy Vector.

To originate the PIM RPF-Proxy Vector, the PE router must have this configuration:

```
router pim
address-family [ipv4|ipv6]
  rpf-vector
!
```

Note: The support for interpreting the PIM RPF-Proxy Vector (this is what the P router must do) was introduced in Cisco IOS XR earlier releases.

This allows for the deployment of GTM over Seamless MPLS.