Migration of Core Tree Protocols on an IOS-XR PE Router in mVPN Networks

Contents

Introduction

This document describes the migration of Multicast VPN (mVPN) Protocol Independent Multicast (PIM) core tree-based Multicast Distribution Trees (MDTs) to Multipoint Label Distribution Protocol (mLDP) core tree-based MDTs. Also, how Data MDTs are signalled at the time of the migration in detail. This document describes the migration only for the Ingress Provider Edge (PE) router running Cisco IOS®-XR.

Migration of Core Tree Protocols

Dual-encap refers to an Ingress router that can forward a Customer (C)-multicast stream onto different types of core tree at the same time. For example, the Ingress PE router forwards one C-multicast stream onto a PIM-based core tree and an mLDP-based core tree at the same time. This is a requirement to successfully migrate mVPN from one core tree type to another.

Dual-encap is supported for PIM and mLDP.

Dual-encap is not supported for Multiprotocol Label Switching (MPLS) P2MP Traffic Engineering (TE).

Default MDT Generic Routing Encapsulation (GRE) and Default MDT mLDP migration or co-existence relies on the fact that the Ingress PE router forwards one C-multicast stream onto a PIM-based core tree and an mLDP-based core tree at the same time. While the Ingress PE forwards onto both MDTs, the Egress PE routers can be migrated one-by-one from one core tree type to another.

Typically, PE routes will migrate from the oldest mVPN deployment model using PIM-based core trees to an mVPN deployment model using mLDP-based trees. The oldest mVPN implementation is Profile 0, which is PIM-based core trees, no Border Gateway Protocol (BGP) Auto-Discovery (AD), and PIM in overlay signalling. However, migration can also occur in the opposite way.

Let's look at this migration scenario as this is the most common migration occurring: from GRE in the core (Profile 0) to a Default MDT mLDP Profile.

There are a few possible Default mLDP Profiles possible.

Let's look at these ones:

- mLDP without BGP AD
- mLDP with BGP AD and PIM C-signaling
- mLDP with BGP AD and BGP C-signaling

In the latter case, there is also a migration of the C-signaling protocol.

One of the things to keep in mind is that when BGP AD used, the Data MDT is signaled by BGP by default. If there is no BGP AD, then the Data MDT cannot be signaled by BGP.

In any case, the Ingress PE must have both Profile 0 and the mLDP profile configured. The Ingress PE will forward the C-multicast traffic onto both MDTs (Default or Data) of both core tree protocols. So, both Default MDTs must be configured on the Ingress PE.

If the Egress PE is capable of running core tree protocols PIM and mLDP, it can decide from which tree to pull the C-multicast traffic. This is done by configuring the Reverse Path Forwarding (RPF) policy on the Egress PE.

If the Egress PE router is capable of only Profile 0, then that PE will only join the PIM tree in the core and receive the C-multicast stream on the PIM-based tree.

> **Note**: If PIM Sparse Mode is used, then both the RP-PE and S-PE must be reachable across both the GRE-based and the mLDP-based MDT.

C-Multicast Protocol Migration

The C-multicast protocol can be migrated from PIM to BGP or vice-versa. This is done by configuring the Egress PE to choose either PIM or BGP as an overlay protocol. It is the Egress PE sending out a join either by PIM or BGP. The Ingress PE can receive and process both in a migration scenario.

This is a migration example of the C-multicast protocol, configured on the Egress PE:

```
router pim
 vrf one
  address-family ipv4
   rpf topology route-policy rpf-for-one
   mdt c-multicast-routing bgp
   !
   interface GigabitEthernet0/1/0/0
    enable
   !
  !
 !
!

route-policy rpf-for-one
```

```
  set core-tree mldp-default
end-policy
!
```

BGP is enabled as the overlay signaling protocol. The default is PIM.

Scenarios

Look at Figure 1. in order to see the setup used for the scenarios.
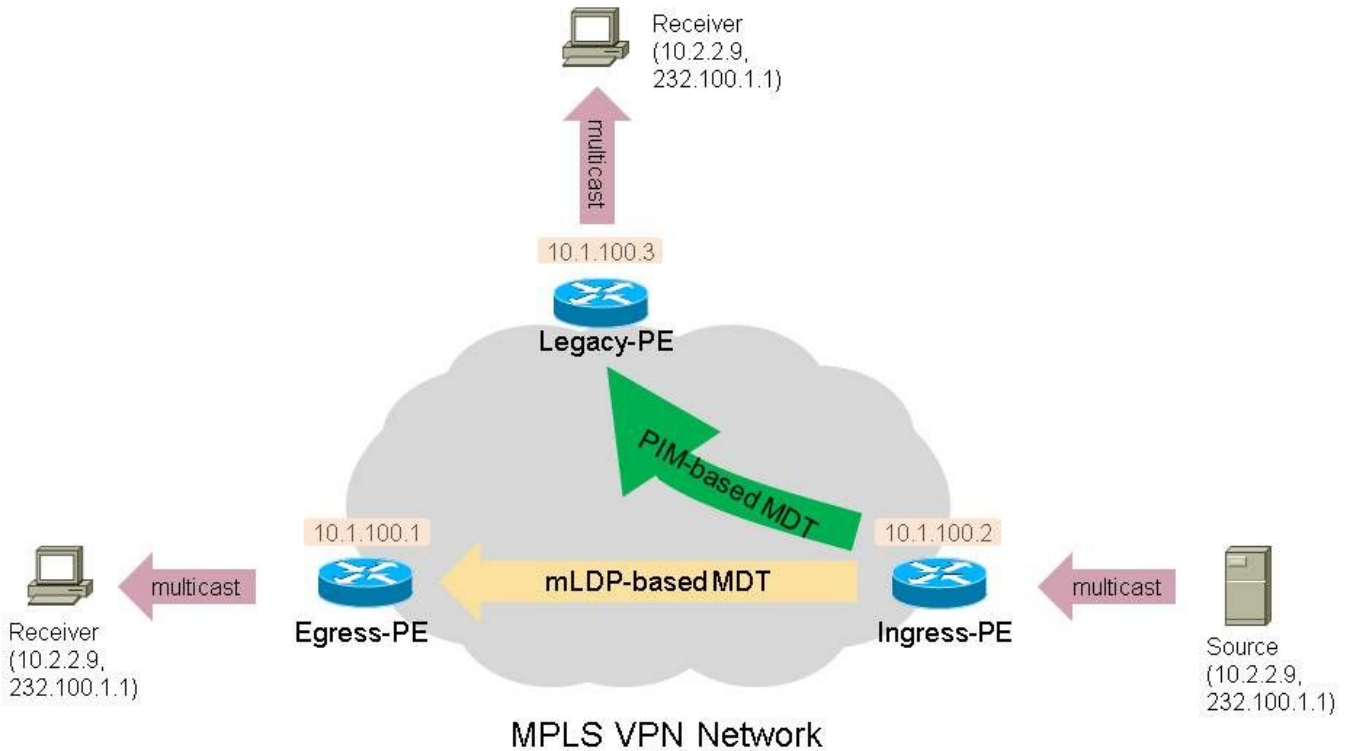

Figure 1.

In these scenarios, you have at least one legacy PE router as a Receiver PE router. This is a router that only runs Profile 0 (Default MDT - GRE - PIM C-mcast Signaling).

This router must have BGP IPv4 MDT configured.

There is at least one Receiver-PE router that runs an mLDP-based Profile. These are all Default MDT mLDP Profiles (1, 9, 13, 12, 17), all Partitioned MDT mLDP Profiles (2, 4, 5, 14 ,15), and Profile 7. Profile 8 for P2MP TE is supported as well.

The Ingress PE router is a dual-encap router: it runs profile 0 and an mLDP-based profile.

At all times this Ingress PE router must forward traffic on both the PIM-based MDT(s) and the mLDP-based MDT(s). These MDTs can be the Default and the Data MDTs.

As a legacy router, take a router running IOS, which can only run Profile 0. The configuration of the legacy router is this.

```
vrf definition one
rd 1:3
```

```
vpn id 1:1
route-target export 1:1
route-target import 1:1
!
address-family ipv4
mdt default 232.1.1.1
exit-address-family
```
BGP IPv4 MDT needs to be configured:

```
router bgp 1
…
address-family ipv4 mdt
neighbor 10.1.100.7 activate
neighbor 10.1.100.7 send-community extended
exit-address-family
!
…
```

Scenario 1.

There is one or more legacy PE router as a Receiver-PE router.

There is one or more PE router as Receiver-PE router that runs Profile 1 (Default MDT - mLDP MP2MP PIM C-mcast Signaling).

There is no BGP AD or BGP C-multicast signaling at all.

Configuration of the Receiver-PE router, running Profile 1:

```
vrf one
 vpn id 1:1
 address-family ipv4 unicast
  import route-target
   1:1
  !
  export route-target
   1:1
  !
 !

router pim
 vrf one
  address-family ipv4
   rpf topology route-policy rpf-for-one
   !
   interface GigabitEthernet0/1/0/0
    enable
   !
  !
 !
!

route-policy rpf-for-one
  set core-tree mldp-default
end-policy
!

multicast-routing
```

```
 vrf one
  address-family ipv4
   mdt source Loopback0
   mdt default mldp ipv4 10.1.100.7
   mdt data 100
   rate-per-route
   interface all enable
    !
   accounting per-prefix
  !
 !
!

mpls ldp
 mldp
  logging notifications
  address-family ipv4
   !
 !
!

route-policy rpf-for-one
 set core-tree mldp-default
```
Configuration of the Ingress PE router:

```
vrf one
vpn id 1:1
address-family ipv4 unicast
import route-target
1:1
!
export route-target
1:1
!
!

router pim
vrf one
address-family ipv4
!
interface GigabitEthernet0/1/0/0
enable
!
!
!

multicast-routing
vrf one
address-family ipv4
mdt source Loopback0
interface all enable
!
  mdt default ipv4 232.1.1.1
  mdt default mldp ipv4 10.1.100.7
  mdt data 255
  mdt data 232.1.2.0/24
!
!
!

mpls ldp
mldp
```

```
logging notifications
address-family ipv4
!
!
!
```

The ingress PE router must have BGP address family IPv4 MDT, matching what the legacy PE router has.

The Ingress PE must be forwarding onto both types of MDT:

```
Ingress-PE#show mrib vrf one route 232.100.1.1

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, MF - MPLS Encap, 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

(10.2.2.9,232.100.1.1) RPF nbr: 10.2.2.9 Flags: RPF MT
MT Slot: 0/1/CPU0
Up: 00:56:09
Incoming Interface List
  GigabitEthernet0/1/0/0 Flags: A, Up: 00:56:09
Outgoing Interface List
  mdtone Flags: F NS MI MT MA, Up: 00:22:59      <<< PIM-based tree
  Lmdtone Flags: F NS LMI MT MA, Up: 00:56:09    <<< mLDP-based tree
```

The ingress PE should see the legacy PE on the interface mdtone and the Profile 1 PE on interface Lmdtone as a PIM neighbor:

```
Ingress-PE#show pim vrf one neighbor

PIM neighbors in VRF one
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

10.1.100.1                Lmdtone 6w1d     00:01:29 1      P
10.1.100.2*               Lmdtone            6w1d     00:01:15 1 (DR) P
10.1.100.2*               mdtone             5w0d     00:01:30 1      P
10.1.100.3                mdtone 00:50:20 00:01:30 1 (DR) P
```

"debug pim vrf one mdt data" on the Ingress PE:

You see that a type 1 (PIM core tree) and a type 2 (mLDP core tree) PIM Join TLV are sent. The first one on mdtone and the second on Lmdtone.

```
pim[1140]: [13] MDT Grp lookup: Return match for grp 232.1.2.4 src 10.1.100.2 in local list (-)
pim[1140]: [13] In mdt timers process...
pim[1140]: [13] Processing MDT JOIN SEND timer for MDT null core mldp pointer in one
pim[1140]: [13] In join_send_update_timer: route->mt_head 50c53b44
pim[1140]: [13] Create new MDT tlv buffer for one for type 0x1
pim[1140]: [13] Buffer allocated for one mtu 1348 size 0
pim[1140]: [13] TLV type set to 0x1
pim[1140]: [13] TLV added for one mtu 1348 size 16
pim[1140]: [13] MDT cache upd: pe 0.0.0.0, (10.2.2.9,232.100.1.1), mdt_type 0x1, core
(10.1.100.2,232.1.2.4), for vrf one [local, -], mt_lc 0x11, mdt_if 'mdtone', cache NULL
pim[1140]: [13] Looked up cache pe 0.0.0.0(10.2.2.9,232.100.1.1) mdt_type 0x1 in one (found) -
No error
pim[1140]: [13] Cache get: Found entry for 0.0.0.0(10.2.2.9,232.100.1.1) mdt_type 0x1 in one
pim[1140]: [13] pim_mvrf_mdt_cache_update:946, mt_lc 0x11, copied mt_mdt_ifname 'mdtone'
pim[1140]: [13] Create new MDT tlv buffer for one for type 0x2
pim[1140]: [13] Buffer allocated for one mtu 1348 size 0
pim[1140]: [13] TLV type set to 0x2, o_type 0x2
pim[1140]: [13] TLV added for one mtu 1348 size 36
pim[1140]: [13] MDT cache upd: pe 0.0.0.0, (10.2.2.9,232.100.1.1), mdt_type 0x2, core src
10.1.100.2, id [mdt 1:1 1], for vrf one [local, -], mt_lc 0x11, mdt_if 'Lmdtone', cache NULL
pim[1140]: [13] Looked up cache pe 0.0.0.0(10.2.2.9,232.100.1.1) mdt_type 0x2 in one (found) -
No error
pim[1140]: [13] Cache get: Found entry for 0.0.0.0(10.2.2.9,232.100.1.1) mdt_type 0x2 in one
pim[1140]: [13] pim_mvrf_mdt_cache_update:946, mt_lc 0x11, copied mt_mdt_ifname 'Lmdtone'
pim[1140]: [13] Set next send time for core type (0x0/0x2) (v: 10.2.2.9,232.100.1.1) in one
pim[1140]: [13] 2. Flush MDT Join for one on Lmdtone(10.1.100.2) 6 (Cnt:1, Reached size 36 MTU
1348)
pim[1140]: [13] 2. Flush MDT Join for one (Lo0) 10.1.100.2
pim[1140]: [13] 2. Flush MDT Join for one on mdtone(10.1.100.2) 6 (Cnt:1, Reached size 16 MTU
1348)
pim[1140]: [13] 2. Flush MDT Join for one (Lo0) 10.1.100.2


Ingress-PE#show pim vrf one mdt cache

Core Source      Cust (Source, Group)           Core Data       Expires
10.1.100.2       (10.2.2.9, 232.100.1.1)          232.1.2.4       00:02:36
10.1.100.2       (10.2.2.9, 232.100.1.1)          [mdt 1:1 1]     00:02:36
```

> **Note**: The PIM Join Type Length Value (TLV) is a PIM message sent over the Default MDT
> and is used to signal the Data MDT. It is sent periodically, once every minute.

The legacy Egress PE:

"debug ip pim vrf one 232.100.1.1":

```
PIM(1): Receive MDT Packet (55759) from 10.1.100.2 (Tunnel3), length (ip: 44, udp: 24), ttl:
1PIM(1): TLV type: 1 length: 16 MDT Packet length: 16
```
The legacy PE caches the PIM Join TLV:

```
Legacy-PE#show ip pim vrf one mdt receive

Joined MDT-data [group/mdt number : source] uptime/expires for VRF: one
[232.1.2.4 : 10.1.100.2] 00:01:10/00:02:45
```
The legacy PE joins the Data MDT in the core:

```
Legacy-PE#show ip mroute vrf one 232.100.1.1
```

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

(10.2.2.9, 232.100.1.1), 00:08:48/00:02:34, flags: sT**Y**
Incoming interface: Tunnel3, RPF nbr 10.1.100.2, **MDT:[10.1.100.2,232.1.2.4]**/00:02:46
Outgoing interface list:
GigabitEthernet1/1, Forward/Sparse, 00:08:48/00:02:34
```

The Profile 1 Receiver-PE receives the PIM Join TLV as well, but for the mLDP-based Data MDT:

```
Egress-PE#**debug pim vrf one mdt data**

pim[1161]: [13] Received MDT Packet on Lmdtone (vrf:one) from 10.1.100.2, len 36
pim[1161]: [13] Processing type 2 tlv
pim[1161]: [13] Received MDT Join TLV from 10.1.100.2 for cust route 10.2.2.9,232.100.1.1
MDT number 1 len 36
pim[1161]: [13] Looked up cache pe 10.1.100.2(10.2.2.9,232.100.1.1) mdt_type 0x2 in one
(found) - No error
pim[1161]: [13] MDT cache upd: pe 10.1.100.2, (10.2.2.9,232.100.1.1), **mdt_type 0x2**, core
**src 10.1.100.2**, **id [mdt 1:1 1]**, for vrf one [remote, -], mt_lc 0xffffffff, mdt_if 'xxx',
cache NULL
pim[1161]: [13] Looked up cache pe 10.1.100.2(10.2.2.9,232.100.1.1) mdt_type 0x2 in one
(found) - No error
pim[1161]: [13] Cache get: Found entry for 10.1.100.2(10.2.2.9,232.100.1.1) mdt_type 0x2
in one
RP/0/RP1/CPU0:Nov 27 16:04:02.726 : Return match for [mdt 1:1 1] src 10.1.100.2 in remote
list (one)
pim[1161]: [13] Remote join: MDT [mdt 1:1 1] known in one. Refcount (1, 1)


Egress-PE#**show pim vrf one mdt cache**

Core Source     Cust (Source, Group)            Core Data       Expires
10.1.100.2       (10.2.2.9, 232.100.1.1)         [mdt 1:1 1]    00:02:12


Egress-PE#**show mrib vrf one route 232.100.1.1**

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, MF - MPLS Encap, 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

(10.2.2.9,232.100.1.1) RPF nbr: 10.1.100.2 Flags: RPF
Up: 00:45:20
Incoming Interface List
  Lmdtone Flags: A LMI, Up: 00:45:20
Outgoing Interface List
  GigabitEthernet0/0/0/9 Flags: F NS LI, Up: 00:45:20
```

Scenario 2.

There are one or more legacy PE routers as Receiver-PE routers.

There are one or more PE routers as Receiver-PE router that run Profile 9 (Default MDT - mLDP MP2MP BGP-AD PIM C-mcast Signaling).

There is BGP AD involved, but no BGP C-multicast signaling.

Configuration of the Receiver-PE router, running Profile 9:

```
vrf one
 vpn id 1:1
 address-family ipv4 unicast
  import route-target
   1:1
  !
  export route-target
   1:1
  !
 !

router pim
 vrf one
  address-family ipv4
   rpf topology route-policy rpf-for-one
   !
   interface GigabitEthernet0/1/0/0
    enable
   !
  !
 !
!

route-policy rpf-for-one
  set core-tree mldp-default
end-policy
!

multicast-routing
vrf one
address-family ipv4
  mdt source Loopback0
  rate-per-route
  interface all enable
  accounting per-prefix
  bgp auto-discovery mldp
  !
```

```
 mdt default mldp ipv4 10.1.100.7
!
!
!

router bgp 1
!
address-family vpnv4 unicast
!
!
 address-family ipv4 mvpn
!
!
neighbor 10.1.100.7      <<< iBGP neighbor
remote-as 1
update-source Loopback0

address-family vpnv4 unicast
!
 address-family ipv4 mvpn
!
!
vrf one
rd 1:1
address-family ipv4 unicast
  redistribute connected
!
 address-family ipv4 mvpn
!
!

mpls ldp
 mldp
  logging notifications
  address-family ipv4
   !
 !
!
```
The ingress PE router must have BGP address family IPv4 MDT, matching what the legacy PE router has. The ingress PE router must have BGP address family IPv4 MVPN, matching what the Profile 9 Egress PE router has.

Configuration of the Ingress PE router:


```
vrf one
vpn id 1:1
address-family ipv4 unicast
import route-target
  1:1
!
export route-target
  1:1
!
!
address-family ipv6 unicast
!
!

router pim
vrf one
address-family ipv4
```

```
      mdt c-multicast-routing pim
      announce-pim-join-tlv
        !
      interface GigabitEthernet0/1/0/0
      enable
        !
  !
  !
  !


multicast-routing
vrf one
address-family ipv4
  mdt source Loopback0
  interface all enable
    bgp auto-discovery mldp
!
    mdt default ipv4 232.1.1.1
    mdt default mldp ipv4 10.1.100.7
    mdt data 255
    mdt data 232.1.2.0/24
!
!
!


router bgp 1
address-family vpnv4 unicast
!
 address-family ipv4 mdt
!
 address-family ipv4 mvpn
!
neighbor 10.1.100.7      <<< iBGP neighbor
remote-as 1
update-source Loopback0
address-family vpnv4 unicast
!
 address-family ipv4 mdt
!
 address-family ipv4 mvpn
!
!
vrf one
rd 1:2
address-family ipv4 unicast
  redistribute connected
!
 address-family ipv4 mvpn
!

mpls ldp
 mldp
  logging notifications
  address-family ipv4
   !
  !
!
```
Without the command "announce-pim-join-tlv", the Ingress PE router does not send out the PIM Join TLV messages over the Default MDTs, if BGP Auto-Discovery (AD) is enabled. Without this command, the Ingress PE router only sends out a BGP IPv4 mvpn route-type 3 update. The Profile 9 Egress PE router does receive the BGP update and installs the Data MDT message in his cache. The legacy PE router does not run BGP AD and hence does not learn the Data MDT

Join message through BGP.

The Ingress PE must be forwarding the C-multicast traffic onto both types of MDT:

```
Ingress-PE#show mrib vrf one route 232.100.1.1

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, MF - MPLS Encap, 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

(10.2.2.9,232.100.1.1) RPF nbr: 10.2.2.9 Flags: RPF MT
MT Slot: 0/1/CPU0
Up: 05:03:56
Incoming Interface List
  GigabitEthernet0/1/0/0 Flags: A, Up: 05:03:56
Outgoing Interface List
  mdtone Flags: F NS MI MT MA, Up: 05:03:56
  Lmdtone Flags: F NS LMI MT MA, Up: 05:03:12
```

The ingress PE should see the legacy PE on the interface mdtone and the Profile 9 PE on interface Lmdtone as PIM neighbor:

```
Ingress-PE#show pim vrf one neighbor

PIM neighbors in VRF one
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

10.1.100.1                Lmdtone            6w1d      00:01:18 1       P
10.1.100.2*               Lmdtone            6w1d      00:01:34 1 (DR)  P
10.1.100.2*               mdtone             5w0d      00:01:18 1       P
10.1.100.3                mdtone             06:00:03 00:01:21 1 (DR)
```

The Profile 9 Egress PE receives the Data MDT message as a BGP update for a route-type 3 in address family IPv4 MVPN:

```
Egress-PE#show bgp ipv4 mvpn vrf one
BGP router identifier 10.1.100.1, local AS number 1
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0x0   RD version: 1367879340
BGP main routing table version 92
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: 1:1 (default for vrf one)
*> [1][10.1.100.1]/40 0.0.0.0                              0 i
*>i[1][10.1.100.2]/40 10.1.100.2                    100    0 i
*>i[3][32][10.2.2.9][32][232.100.1.1][10.1.100.2]/120
                      10.1.100.2                    100    0 i


Processed 3 prefixes, 3 paths


Egress-PE#show bgp ipv4 mvpn vrf one [3][32][10.2.2.9][32][232.100.1.1][10.1.100.2]/120
BGP routing table entry for [3][32][10.2.2.9][32][232.100.1.1][10.1.100.2]/120, Route
Distinguisher: 1:1
Versions:
Process            bRIB/RIB SendTblVer
Speaker              92        92
Last Modified: Nov 27 20:25:32.474 for 00:44:22
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
  10.1.100.2 (metric 12) from 10.1.100.7 (10.1.100.2)
    Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
imported
    Received Path ID 0, Local Path ID 1, version 92
    Community: no-export
    Extended community: RT:1:1
    Originator: 10.1.100.2, Cluster list: 10.1.100.7
    PMSI: flags 0x00, type 2, label 0, ID
0x060001040a016402000e02000b000001000000010000001
    Source VRF: default, Source Route Distinguisher: 1:2
```

This BGP route is a route-type 3, for protocol tunnel type 2, which is mLDP P2MP LSP (the Data MDT built on a P2MP mLSP LSP). There is no BGP route-type 3 entry for any PIM tree, as BGP AD is not enabled for PIM.

"debug pim vrf one mdt data" on the Ingress PE:

```
pim[1140]: [13] In mdt timers process...
pim[1140]: [13] Processing MDT JOIN SEND timer for MDT null core mldp pointer in one
pim[1140]: [13] In join_send_update_timer: route->mt_head 50c53b44
pim[1140]: [13] Create new MDT tlv buffer for one for type 0x1
pim[1140]: [13] Buffer allocated for one mtu 1348 size 0
pim[1140]: [13] TLV type set to 0x1
pim[1140]: [13] TLV added for one mtu 1348 size 16
pim[1140]: [13] MDT cache upd: pe 0.0.0.0, (10.2.2.9,232.100.1.1), mdt_type 0x1, core
(10.1.100.2,232.1.2.5), for vrf one [local, -], mt_lc 0x11, mdt_if 'mdtone', cache NULL
pim[1140]: [13] Looked up cache pe 0.0.0.0(10.2.2.9,232.100.1.1) mdt_type 0x1 in one
(found) - No error
pim[1140]: [13] Cache get: Found entry for 0.0.0.0(10.2.2.9,232.100.1.1) mdt_type 0x1 in
one
pim[1140]: [13] pim_mvrf_mdt_cache_update:946, mt_lc 0x11, copied mt_mdt_ifname 'mdtone'
pim[1140]: [13] Create new MDT tlv buffer for one for type 0x2
pim[1140]: [13] Buffer allocated for one mtu 1348 size 0
pim[1140]: [13] TLV type set to 0x2, o_type 0x2
pim[1140]: [13] TLV added for one mtu 1348 size 36
pim[1140]: [13] MDT cache upd: pe 0.0.0.0, (10.2.2.9,232.100.1.1), mdt_type 0x2, core src
10.1.100.2, id [mdt 1:1 1], for vrf one [local, -], mt_lc 0x11, mdt_if 'Lmdtone', cache
NULL
```

```
: pim[1140]: [13] Looked up cache pe 0.0.0.0(10.2.2.9,232.100.1.1) mdt_type 0x2 in one
(found) - No error
pim[1140]: [13] Cache get: Found entry for 0.0.0.0(10.2.2.9,232.100.1.1) mdt_type 0x2 in
one
pim[1140]: [13] pim_mvrf_mdt_cache_update:946, mt_lc 0x11, copied mt_mdt_ifname 'Lmdtone'
pim[1140]: [13] Set next send time for core type (0x0/0x2) (v: 10.2.2.9,232.100.1.1) in
one
pim[1140]: [13] 2. Flush MDT Join for one on Lmdtone(10.1.100.2) 6 (Cnt:1, Reached size
36 MTU 1348)
pim[1140]: [13] 2. Flush MDT Join for one (Lo0) 10.1.100.2
pim[1140]: [13] 2. Flush MDT Join for one on mdtone(10.1.100.2) 6 (Cnt:1, Reached size 16
MTU 1348)
pim[1140]: [13] 2. Flush MDT Join for one (Lo0) 10.1.100.2
```

The Ingress PE sends a PIM Join TLV for both the PIM-based and the mLDP-based Data MDT.

On the legacy PE:

 "debug ip pim vrf one 232.100.1.1":

```
PIM(1): Receive MDT Packet (56333) from 10.1.100.2 (Tunnel3), length (ip: 44, udp: 24), ttl: 1
PIM(1): TLV type: 1 length: 16 MDT Packet length: 16
```

The legacy PE receives and caches the PIM Join TLV:

```
Legacy-PE#show ip pim vrf one mdt receive

Joined MDT-data [group/mdt number : source] uptime/expires for VRF: one
[232.1.2.5 : 10.1.100.2] 00:23:30/00:02:33
```

The legacy PE joins the Data MDT in the core:

```
Legacy-PE#show ip mroute vrf one 232.100.1.1
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(10.2.2.9, 232.100.1.1), 05:13:35/00:03:02, flags: sTY
Incoming interface: Tunnel3, RPF nbr 10.1.100.2, MDT:[10.1.100.2,232.1.2.5]/00:02:37
Outgoing interface list:
  GigabitEthernet1/1, Forward/Sparse, 05:13:35/00:03:02
```

The Profile 9 Receiver-PE.

"debug pim vrf one mdt data" on the Profile 9 Egress PE:

```
pim[1161]: [13] Received MDT Packet on Lmdtone (vrf:one) from 10.1.100.2, len 36
```

```
pim[1161]: [13] Processing type 2 tlv
pim[1161]: [13] Received MDT Join TLV from 10.1.100.2 for cust route 10.2.2.9,232.100.1.1
MDT number 1 len 36
pim[1161]: [13] Looked up cache pe 10.1.100.2(10.2.2.9,232.100.1.1) mdt_type 0x2 in one
(found) - No error
pim[1161]: [13] MDT cache upd: pe 10.1.100.2, (10.2.2.9,232.100.1.1), mdt_type 0x2, core
src 10.1.100.2, id [mdt 1:1 1], for vrf one [remote, -], mt_lc 0xffffffff, mdt_if 'xxx',
cache NULL
pim[1161]: [13] Looked up cache pe 10.1.100.2(10.2.2.9,232.100.1.1) mdt_type 0x2 in one
(found) - No error
pim[1161]: [13] Cache get: Found entry for 10.1.100.2(10.2.2.9,232.100.1.1) mdt_type 0x2
in one
pim[1161]: [13] MDT lookup: Return match for [mdt 1:1 1] src 10.1.100.2 in remote list
(one)
pim[1161]: [13] Remote join: MDT [mdt 1:1 1] known in one. Refcount (1, 1)
```

The Profile 9 Receiver-PE receives and caches the PIM Join TLV. The Profile 9 Receiver-PE also learned of the Data MDT because of receiving the BGP update message for a route-type 3 from the Ingress PE. The PIM Join TLV and the BGP update message route-type are equivalent and hold the same information with regards to the core tree tunnel for the Data MDT.

```
Egress-PE#show pim vrf one mdt cache

Core Source     Cust (Source, Group)              Core Data      Expires
10.1.100.2       (10.2.2.9, 232.100.1.1)           [mdt 1:1 1]     00:02:35


Egress-PE#show mrib vrf one route 232.100.1.1

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, MF - MPLS Encap, 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

(10.2.2.9,232.100.1.1) RPF nbr: 10.1.100.2 Flags: RPF
Up: 05:10:22
Incoming Interface List
  Lmdtone Flags: A LMI, Up: 05:10:22
Outgoing Interface List
  GigabitEthernet0/0/0/9 Flags: F NS LI, Up: 05:10:22
```

Scenario 3.

There is one or more legacy PE router as a Receiver-PE router.
There is one or more PE router as Receiver-PE router that runs Profile 13 (Default MDT - mLDP MP2MP BGP-AD BGP C-mcast Signaling).
There is BGP AD involved and BGP C-multicast signaling.

Configuration of the Receiver-PE router, running Profile 13:

```
vrf one
 vpn id 1:1
 address-family ipv4 unicast
 import route-target
  1:1
 !
 export route-target
  1:1
 !
 !

router pim
 vrf one
 address-family ipv4
  rpf topology route-policy rpf-for-one
  mdt c-multicast-routing bgp
  !
  interface GigabitEthernet0/1/0/0
   enable
  !
 !
 !
!

route-policy rpf-for-one
 set core-tree mldp-default
end-policy
!

multicast-routing
 vrf one
 address-family ipv4
  mdt source Loopback0
  rate-per-route
  interface all enable
  accounting per-prefix
  bgp auto-discovery mldp
  !
  mdt default mldp ipv4 10.1.100.7
 !
 !
!

router bgp 1
!
 address-family vpnv4 unicast
 !
!
address-family ipv4 mvpn
 !
!
 neighbor 10.1.100.7      <<< iBGP neighbor
 remote-as 1
 update-source Loopback0
 !
 address-family vpnv4 unicast
 !
 address-family ipv4 mvpn
 !
 !
```

```
vrf one
 rd 1:1
 address-family ipv4 unicast
  redistribute connected
  !
 address-family ipv4 mvpn
  !
  !

mpls ldp
 mldp
 logging notifications
 address-family ipv4
  !
  !
!
```

Configuration of the Ingress PE router:

```
vrf one
 vpn id 1:1
 address-family ipv4 unicast
 import route-target
  1:1
 !
 export route-target
  1:1
 !
 !
 address-family ipv6 unicast
 !
!

router pim
 vrf one
 address-family ipv4
   mdt c-multicast-routing bgp
    announce-pim-join-tlv
  !
  interface GigabitEthernet0/1/0/0
   enable
  !
 !
 !
!

multicast-routing
 vrf one
 address-family ipv4
  mdt source Loopback0
  interface all enable
   mdt default ipv4 232.1.1.1
   mdt default mldp ipv4 10.1.100.7
   mdt data 255
   mdt data 232.1.2.0/24
 !
 !
!

router bgp 1
 address-family vpnv4 unicast
 !
 address-family ipv4 mdt
```

```
 !
 address-family ipv4 mvpn
 !
 neighbor 10.1.100.7      <<< iBGP neighbor
 remote-as 1
 update-source Loopback0
 address-family vpnv4 unicast
  !
  address-family ipv4 mdt
  !
  address-family ipv4 mvpn
  !
  !
 vrf one
 rd 1:2
 address-family ipv4 unicast
  redistribute connected
  !
  address-family ipv4 mvpn
  !

mpls ldp
 mldp
 logging notifications
 address-family ipv4
  !
  !
 !
```

Without the command announce-pim-join-tlv, the Ingress PE router does not send out the PIM Join TLV messages over the Default MDT, if BGP AD is enabled. Without this command, the Ingress PE router only sends out a BGP IPv4 mvpn route-type 3 update. The Profile 13 Egress PE router does receive the BGP update and installs the Data MDT message in his cache. The legacy PE router does not run BGP AD and hence does not learn the Data MDT Join message through BGP.

The Ingress PE must be forwarding onto both types of MDT:

```
Ingress-PE#show mrib vrf one route 232.100.1.1

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, MF - MPLS Encap, 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

(10.2.2.9,232.100.1.1) RPF nbr: 10.2.2.9 Flags: RPF MT
MT Slot: 0/1/CPU0
Up: 19:49:27
```

```
Incoming Interface List
  GigabitEthernet0/1/0/0 Flags: A, Up: 19:49:27
Outgoing Interface List
  mdtone Flags: F MI MT MA, Up: 19:49:27
  Lmdtone Flags: F LMI MT MA, Up: 01:10:15
```

The ingress PE should see the legacy PE on the interface mdtone as a PIM neighbor. However, it is not a must to have the Profile 13 PE on interface Lmdtone as a PIM neighbor, because BGP is now used as a C-multicast signaling protocol.

"debug pim vrf one mdt data" on the Ingress PE:

```
pim[1140]: [13] In mdt timers process...
pim[1140]: [13] Processing MDT JOIN SEND timer for MDT null core mldp pointer in one
pim[1140]: [13] In join_send_update_timer: route->mt_head 50c53b44
pim[1140]: [13] Create new MDT tlv buffer for one for type 0x1
pim[1140]: [13] Buffer allocated for one mtu 1348 size 0
pim[1140]: [13] TLV type set to 0x1
pim[1140]: [13] TLV added for one mtu 1348 size 16
pim[1140]: [13] MDT cache upd: pe 0.0.0.0, (10.2.2.9,232.100.1.1), mdt_type 0x1, core
(10.1.100.2,232.1.2.5), for vrf one [local, -], mt_lc 0x11, mdt_if 'mdtone', cache NULL
pim[1140]: [13] Looked up cache pe 0.0.0.0(10.2.2.9,232.100.1.1) mdt_type 0x1 in one (found) -
No error
pim[1140]: [13] Cache get: Found entry for 0.0.0.0(10.2.2.9,232.100.1.1) mdt_type 0x1 in one
pim[1140]: [13] pim_mvrf_mdt_cache_update:946, mt_lc 0x11, copied mt_mdt_ifname 'mdtone'
pim[1140]: [13] Create new MDT tlv buffer for one for type 0x2
pim[1140]: [13] Buffer allocated for one mtu 1348 size 0
pim[1140]: [13] TLV type set to 0x2, o_type 0x2
pim[1140]: [13] TLV added for one mtu 1348 size 36
pim[1140]: [13] MDT cache upd: pe 0.0.0.0, (10.2.2.9,232.100.1.1), mdt_type 0x2, core src
10.1.100.2, id [mdt 1:1 1], for vrf one [local, -], mt_lc 0x11, mdt_if 'Lmdtone', cache NULL
pim[1140]: [13] Looked up cache pe 0.0.0.0(10.2.2.9,232.100.1.1) mdt_type 0x2 in one (found) -
No error
pim[1140]: [13] Cache get: Found entry for 0.0.0.0(10.2.2.9,232.100.1.1) mdt_type 0x2 in one
pim[1140]: [13] pim_mvrf_mdt_cache_update:946, mt_lc 0x11, copied mt_mdt_ifname 'Lmdtone'
pim[1140]: [13] Set next send time for core type (0x0/0x2) (v: 10.2.2.9,232.100.1.1) in one
pim[1140]: [13] 2. Flush MDT Join for one on Lmdtone(10.1.100.2) 6 (Cnt:1, Reached size 36 MTU
1348)
pim[1140]: [13] 2. Flush MDT Join for one (Lo0) 10.1.100.2
pim[1140]: [13] 2. Flush MDT Join for one on mdtone(10.1.100.2) 6 (Cnt:1, Reached size 16 MTU
1348)
pim[1140]: [13] 2. Flush MDT Join for one (Lo0) 10.1.100.2
pim[1140]: [13] MDT Grp lookup: Return match for grp 232.1.2.5 src 10.1.100.2 in local list (-)
```

The Ingress PE sends PIM Join TLV for both the PIM-based and the mLDP-based Data MDT.

"debug ip pim vrf one 232.100.1.1" on the legacy PE:

```
PIM(1): Receive MDT Packet (57957) from 10.1.100.2 (Tunnel3), length (ip: 44, udp: 24), ttl: 1
PIM(1): TLV type: 1 length: 16 MDT Packet length: 16
```

The legacy PE caches the PIM Join TLV:

```
Legacy-PE#show ip pim vrf one mdt receive

Joined MDT-data [group/mdt number : source]  uptime/expires for VRF: one
  [232.1.2.5 : 10.1.100.2]  00:03:36/00:02:24
```

The legacy PE joins the Data MDT in the core:

```
Legacy-PE#show ip mroute vrf one 232.100.1.1
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode


(10.2.2.9, 232.100.1.1), 18:53:53/00:02:50, flags: sTY
 Incoming interface: Tunnel3, RPF nbr 10.1.100.2, MDT:[10.1.100.2,232.1.2.5]/00:02:02
 Outgoing interface list:
   GigabitEthernet1/1, Forward/Sparse, 18:53:53/00:02:50
```
The Profile 13 Receiver-PE:
"debug pim vrf one mdt data" on the Profile 13 Egress PE:


```
pim[1161]: [13] Received MDT Packet on Lmdtone (vrf:one) from 10.1.100.2, len 36
pim[1161]: [13] Processing type 2 tlv
pim[1161]: [13] Received MDT Join TLV from 10.1.100.2 for cust route 10.2.2.9,232.100.1.1 MDT
number 1 len 36
pim[1161]: [13] Looked up cache pe 10.1.100.2(10.2.2.9,232.100.1.1) mdt_type 0x2 in one (found)
- No error
pim[1161]: [13] MDT cache upd: pe 10.1.100.2, (10.2.2.9,232.100.1.1), mdt_type 0x2, core src
10.1.100.2, id [mdt 1:1 1], for vrf one [remote, -], mt_lc 0xffffffff, mdt_if 'xxx', cache NULL
pim[1161]: [13] Looked up cache pe 10.1.100.2(10.2.2.9,232.100.1.1) mdt_type 0x2 in one (found)
- No error
pim[1161]: [13] Cache get: Found entry for 10.1.100.2(10.2.2.9,232.100.1.1) mdt_type 0x2 in one
pim[1161]: [13] MDT lookup: Return match for [mdt 1:1 1] src 10.1.100.2 in remote list (one)
pim[1161]: [13] Remote join: MDT [mdt 1:1 1] known in one. Refcount (1, 1)


RP/0/RP1/CPU0:Legacy-PE#show pim vrf one mdt cache


Core Source      Cust (Source, Group)          Core Data        Expires
10.1.100.2       (10.2.2.9, 232.100.1.1)       [mdt 1:1 1]      00:02:21
```
The Profile 13 Receiver-PE receives and caches the PIM Join TLV for the mLDP-based MDT. The
Profile 13 Receiver-PE also learned of the Data MDT because of receiving the BGP update
message for a route-type 3 from the Ingress PE. The PIM Join TLV and the BGP update message
route-type are equivalent and hold the same information with regards to the core tree tunnel for
the Data MDT.


```
Ingress-PE#show bgp ipv4 mvpn vrf one
BGP router identifier 10.1.100.1, local AS number 1
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0x0   RD version: 1367879340
BGP main routing table version 93
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: 1:1 (default for vrf one)
*> [1][10.1.100.1]/40 0.0.0.0                            0 i
*>i[1][10.1.100.2]/40 10.1.100.2               100       0 i
*>i[3][32][10.2.2.9][32][232.100.1.1][10.1.100.2]/120
                     10.1.100.2                100       0 i
*> [7][1:2][1][32][10.2.2.9][32][232.100.1.1]/184
                     0.0.0.0                             0 i

Processed 4 prefixes, 4 paths


Egress-PE#show bgp ipv4 mvpn vrf one [3][32][10.2.2.9][32][232.100.1.1][10.1.100.2]/120
BGP routing table entry for [3][32][10.2.2.9][32][232.100.1.1][10.1.100.2]/120, Route
Distinguisher: 1:1
Versions:
 Process          bRIB/RIB  SendTblVer
 Speaker              92        92
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
   10.1.100.2 (metric 12) from 10.1.100.7 (10.1.100.2)
     Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate, imported
     Received Path ID 0, Local Path ID 1, version 92
     Community: no-export
     Extended community: RT:1:1
     Originator: 10.1.100.2, Cluster list: 10.1.100.7
     PMSI: flags 0x00, type 2, label 0, ID 0x060001040a016402000e02000b0000010000000100000001
     Source VRF: default, Source Route Distinguisher: 1:2
```

This BGP route is a route-type 3, for protocol tunnel type 2, which is mLDP P2MP LSP (the Data MDT built on a P2MP mLSP LSP). There is no BGP route-type 3 for any PIM tree, as BGP AD is not enabled for PIM.

There is also a route-type 7 because C-multicast signaling is turned on between the Profile 13 Egress PE and the Ingress PE. The route-type 7 BGP update is sent from the Profile 13 Egress PE to the Ingress PE.

Scenario 4.

There is PIM Sparse mode in the VPN context in this scenario.

There is one or more legacy PE router as Source-PE router.

There is one or more PE router as Receiver-PE router that runs Profile 13 (Default MDT - mLDP MP2MP BGP-AD BGP C-mcast Signaling). There is BGP AD involved and BGP C-multicast signaling. Because these PE routers will need to be able to receive traffic directly from the Source-PE - the legacy PE router - they also need to run Profile 0.

The RP-PE is a PE router that runs Profile 13 (Default MDT - mLDP MP2MP BGP-AD BGP C-mcast Signaling). There is BGP AD involved and BGP C-multicast signaling. Because the RP-PE router will need to be able to receive traffic directly from the Source-PE - the legacy PE router - they also need to run Profile 0.

The multicast routing worked in scenario 3, but this might only work for Source-Specific Multicasting (SSM). If the C-signaling is Sparse Mode, then multicast might fail. This can depend on where the Rendez-Vous Point (RP) is placed. If the signaling in the overlay is only (S, G), then the multicast routing will work as in scenario 3. This occurs if the RP is located at the Receiver site. If the RP is at the site of a Receiver, then the Receiver-PE will not send out a (*, G) Join in overlay, either by PIM or BGP. If the RP is however located at the Source-PE, or another PE, then there will be (*, G) and (S, G) signaling in the overlay. The multicast routing may fail if this is done with the configuration as in scenario 3.

Look at Figure 2. It shows a network with a Source-PE (Legacy-PE), an RP-PE (PE2) and a Receiver-PE (PE1).
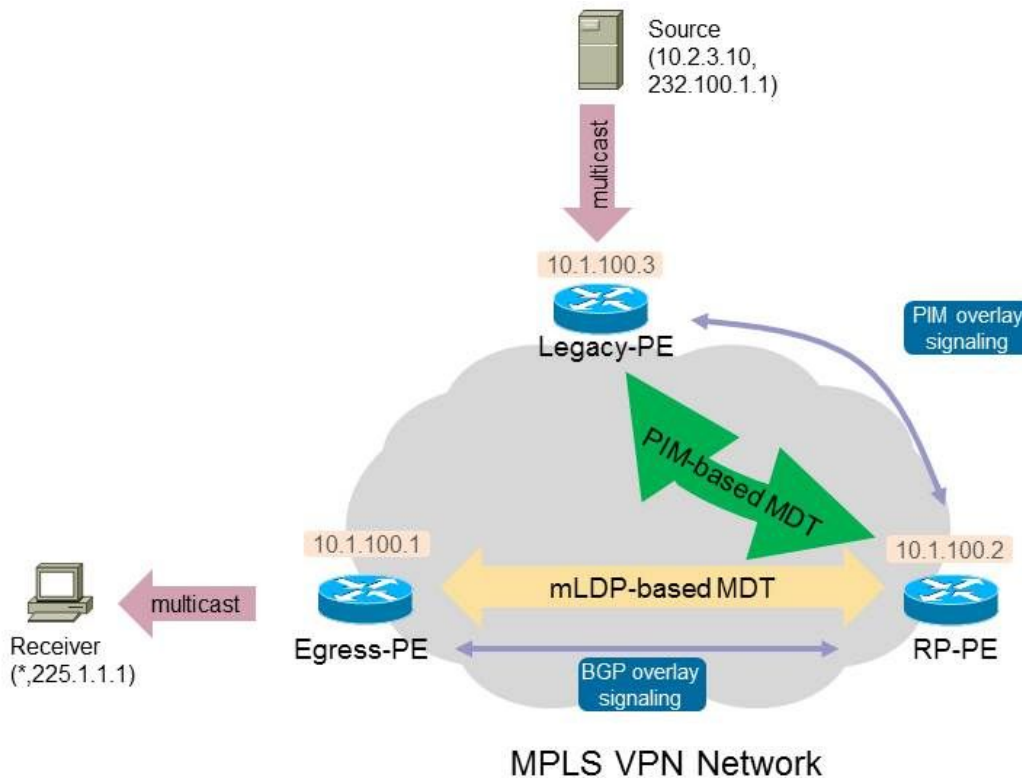

Figure 2.

The Egress PE routers need to send out Joins for (*,G). Which protocol they will use is determined by the configuration. The Egress-PE will use BGP, and the Legacy-Source-PE router will use PIM if it has a Receiver as well. The shared tree will hence be signaled fine. There will be an issue when the Source starts sending: the Source tree will not be signaled.

The Issue

Once the Source starts sending, the RP will receive the register packets from the PIM First Hop Router (FHR). This could be the Legacy-Source-PE router here. The RP-PE would then need to send out a PIM (S, G) Join towards the Legacy-Source-PE, since the Legacy-Source-PE does not run BGP as an overlay signaling protocol. However, the RP-PE has BGP configured as the overlay signaling protocol. So, the Legacy-Source-PE will never receive a PIM (S, G) Join message from the RP-PE and hence the Source tree from Source to RP cannot be signaled. The setup is stuck in the phase of registering. The Outgoing Interface List (OIL) on the Legacy-Source-PE will be empty:

```
Legacy-PE#show ip mroute vrf one 225.1.1.1
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.1.1.1), 00:05:47/stopped, RP 10.2.100.9, flags: SPF
Incoming interface: Tunnel3, RPF nbr 10.1.100.2
Outgoing interface list: Null

(10.2.3.10, 225.1.1.1), 00:05:47/00:02:42, flags: PFT
Incoming interface: GigabitEthernet1/1, RPF nbr 10.2.3.10
Outgoing interface list: Null
```

In order to fix this, you need to have the RP-PE send out a PIM Join for (S, G) to the Legacy-Source-PE, while RP-PE still has BGP enabled as an overlay signaling protocol for the non-legacy routers. If a Source comes online behind a non-legacy router, then the RP-PE needs to send out a route-type 7 BGP update message towards that non-legacy router.

The RP-PE can use both PIM and BGP as overlay signaling. The choice of either one will be determined by a route policy. You need to have the migration command under router PIM for the VRF. For the network depicted in Figure 2, this is the needed configuration on the RP-PE:

```
router pim
vrf one
address-family ipv4
   rpf topology route-policy rpf-for-one
  mdt c-multicast-routing bgp
    migration route-policy PIM-to-BGP
    announce-pim-join-tlv
  !
  interface GigabitEthernet0/1/0/0
  enable
  !
!
!
!
route-policy rpf-for-one
 if next-hop in (10.1.100.3/32) then
  set core-tree pim-default
else
  set core-tree mldp-default
endif
end-policy
!

route-policy PIM-to-BGP
 if next-hop in (10.1.100.3/32) then
  set c-multicast-routing pim
 else
```

```
   set c-multicast-routing bgp
endif
end-policy
!

multicast-routing
vrf one
address-family ipv4
  mdt source Loopback0
  rate-per-route
  interface all enable
  accounting per-prefix
  bgp auto-discovery mldp
  !
   mdt default ipv4 232.1.1.1
  mdt default mldp ipv4 10.1.100.7
  !
!
!
```

The route-policy PIM-to-BGP specifies that if the remote PE router is 10.1.100.3 (Legacy-Source-PE), then use PIM as an overlay signaling protocol. Else (so for the non-legacy PE router), BGP is used as the overlay signaling protocol. So, the RP-PE now sends out a PIM (S, G) Join towards the Legacy-Source-PE on the PIM-based Default MDT. The Legacy-Source-PE now has the (S, G) entry:

```
Legacy-PE#show ip mroute vrf one 225.1.1.1
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
     N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.1.1.1), 00:11:56/stopped, RP 10.2.100.9, flags: SPF
Incoming interface: Tunnel3, RPF nbr 10.1.100.2
Outgoing interface list: Null

(10.2.3.10, 225.1.1.1), 00:11:56/00:03:22, flags: FT
Incoming interface: GigabitEthernet1/1, RPF nbr 10.2.3.10
Outgoing interface list:
  Tunnel3, Forward/Sparse, 00:00:11/00:03:18
```

The Receiver can receive the multicast packets if the RP-PE U-turns the packets: it forwards the multicast packets received from the MDT onto the Lmdt tree.

> **Note**: Check if the RP-PE router has support for the PE turnaround feature on that platform and software.

```
RP/0/3/CPU1:PE2#show mrib vrf one route 225.1.1.1
```

```
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, MF - MPLS Encap, 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

(*,225.1.1.1) RPF nbr: 10.2.2.9 Flags: C RPF
Up: 00:53:59
Incoming Interface List
  GigabitEthernet0/1/0/0 Flags: A, Up: 00:53:59
Outgoing Interface List
  Lmdtone Flags: F LMI, Up: 00:53:59

(10.2.3.10,225.1.1.1) RPF nbr: 10.1.100.3 Flags: RPF
Up: 00:03:00
Incoming Interface List
  mdtone Flags: A MI, Up: 00:03:00
Outgoing Interface List
  Lmdtone Flags: F NS LMI, Up: 00:03:00
```

No matter if the Last Hop Router (LHR) has SPT-switchover configured or not, the multicast traffic continues to be forwarded over the shared tree, towards the RP-PE. Look at Figure 3. in order to see how the multicast traffic is forwarded.
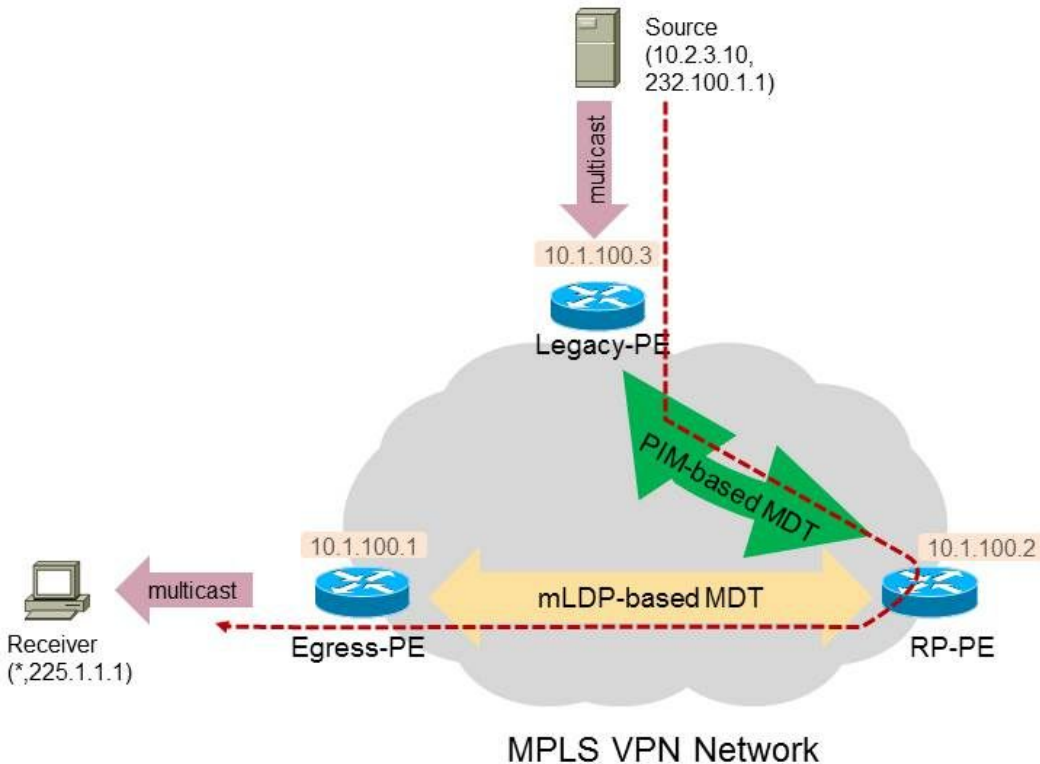

Figure 3.

The Egress-PE has no (S, G) entry:

```
RP/0/RP1/CPU0:PE1#show mrib vrf one route 225.1.1.1
IP Multicast Routing Information Bas
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, MF - MPLS Encap, 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


(*,225.1.1.1) RPF nbr: 10.1.100.2 Flags: C RPF
Up: 04:35:36
Incoming Interface List
  Lmdtone Flags: A LMI, Up: 03:00:24
Outgoing Interface List
  GigabitEthernet0/0/0/9 Flags: F NS, Up: 04:35:36
```

If the Egress-PE is the LHR, then it will not have an (S, G) entry. The reason that the Egress-PE cannot switchover to the (S, G) entry, is that it did not receive a BGP Source Active route from a PE router. The multicast traffic is forwarded as in Figure 3.

However, it is possible that the Egress-PE is not the LHR, but a CE router at the Egress-PE site- is the LHR. If that CE router does switchover to the Source tree, then the Egress-PE will receive a PIM (S, G) Join and install the (S, G) entry.

```
RP/0/RP1/CPU0:PE1#show mrib vrf one route 225.1.1.1

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, MF - MPLS Encap, 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


(*,225.1.1.1) RPF nbr: 10.1.100.2 Flags: C RPF
Up: 00:04:51
Incoming Interface List
  Lmdtone Flags: A LMI, Up: 00:04:51
Outgoing Interface List
```

```
  GigabitEthernet0/0/0/9 Flags: F NS, Up: 00:04:51

(10.2.3.10,225.1.1.1) RPF nbr: 10.1.100.3 Flags: RPF
Up: 00:00:27
Incoming Interface List
  Lmdtone Flags: A LMI, Up: 00:00:27
Outgoing Interface List
  GigabitEthernet0/0/0/9 Flags: F NS, Up: 00:00:27
```
But, the Egress-PE will now RPF to the Source and find the router Legacy-Source-PE as the RPF neighbor:

```
RP/0/RP1/CPU0:PE1#show pim vrf one rpf 10.2.3.10

Table: IPv4-Unicast-default
* 10.2.3.10/32 [200/0]
  via Lmdtone with rpf neighbor 10.1.100.3
   Connector: 1:3:10.1.100.3, Nexthop: 10.1.100.3
```
Since there is no MDT between the Egress-PE and the Legacy-Source-PE, the Egress-PE cannot send a Join to the Legacy-Source-PE. Remember that the Egress-PE only build mLDP trees and does BGP customer signaling. Remember that the Legacy-Source-PE only build PIM-based trees and only does PIM customer signaling.

However, since the Egress-PE has RPF info pointing to incoming interface Lmdt and the multicast traffic arrives still on that MDT from the RP-PE, the multicast traffic will be forwarded towards the receiver and will not fail RPF. The reason is that the RPF does not do a strict RPF check to check if the multicast traffic actually arrives from RPF neighbor 10.1.100.3, the Legacy-PE router. Note that there is no PIM adjacency for 10.1.100.3 on PE1 on Lmdt, because the Legacy-PE cannot have Lmdt because it only runs PIM as core tree protocol (Profile 0):

```
RP/0/RP1/CPU0:PE1#show pim vrf one neighbor
PIM neighbors in VRF one
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

10.1.100.1*                Lmdtone             01:32:46 00:01:32 100 (DR) P
10.1.100.2                 Lmdtone             01:30:46 00:01:16 1      P
10.1.100.4                 Lmdtone             01:30:38 00:01:24 1      P
10.1.100.1*                mdtone              01:32:46 00:01:34 100 (DR) P
10.1.100.2                 mdtone              01:32:45 00:01:29 1      P
10.1.100.3                 mdtone              01:32:17 00:01:29 1      P
10.1.100.4                 mdtone              01:32:43 00:01:20 1      P
10.2.1.1*                  GigabitEthernet0/0/0/9 01:32:46 00:01:18 100    B P E
10.2.1.8                   GigabitEthernet0/0/0/9 01:32:39 00:01:16 100 (DR)
```
The reason that PE1 picks Lmdt as the incoming interface is that this is the info received from the RPF topology command on PE1:

```
route-policy rpf-for-one
set core-tree mldp-default
end-policy
!
```
If the RPF is still ok on PE1, then the multicast traffic can reach the Receiver behind PE1. But, the traffic does not take the shortest path Legacy-PE to PE1 in the core.

The Solution

In order to fix this, the Egress-PE (PE1) must be configured to signal PIM-based MDT and BGP as overlay signaling as well. This configuration is needed on the Egress-PE in that case:

```
router pim
vrf one
address-family ipv4
   rpf topology route-policy rpf-for-one
  mdt c-multicast-routing bgp
  migration route-policy PIM-to-BGP
    announce-pim-join-tlv
  !
  rp-address 10.2.100.9 override
  !
  interface GigabitEthernet0/0/0/9
   enable
  !
!
!
!

route-policy rpf-for-one
 if next-hop in (10.1.100.3/32) then
   set core-tree pim-default
else
  set core-tree mldp-default
endif
end-policy
!

route-policy PIM-to-BGP
 if next-hop in (10.1.100.3/32) then
  set c-multicast-routing pim
 else
  set c-multicast-routing bgp
endif
end-policy
!

multicast-routing
vrf one
address-family ipv4
  mdt source Loopback0
  rate-per-route
  interface all enable
  accounting per-prefix
  bgp auto-discovery mldp
  !
  mdt default ipv4 232.1.1.1
   mdt default mldp ipv4 10.1.100.7
!
!
!
```

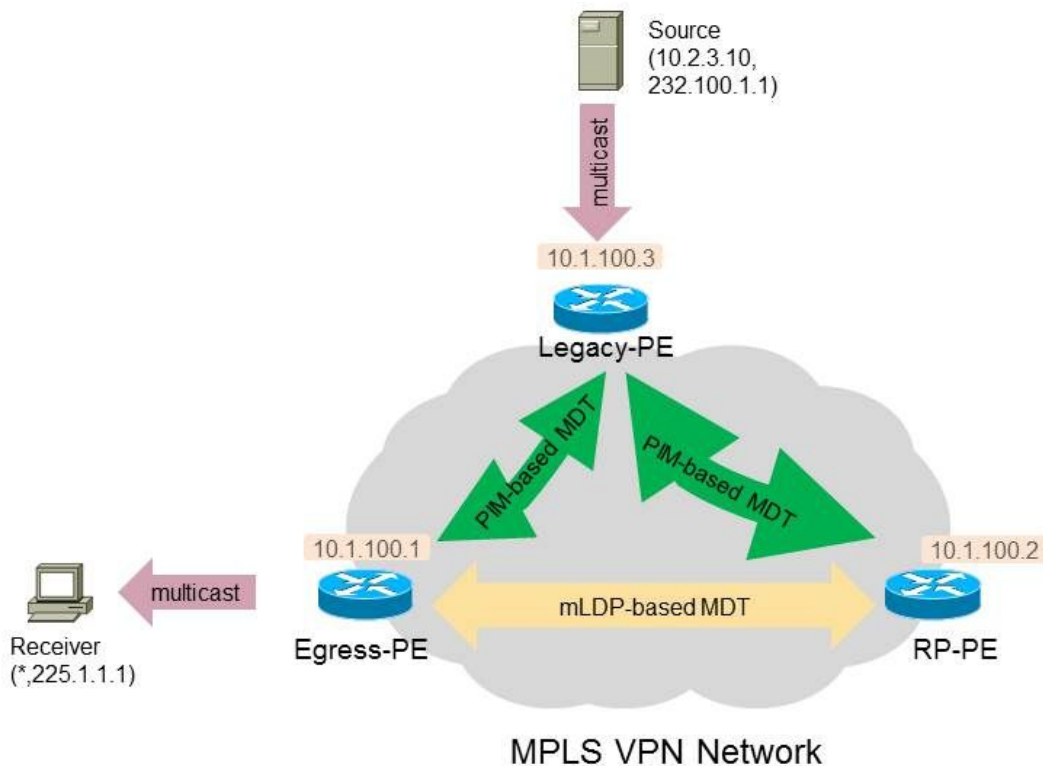Look at Figure 4. There is now a PIM-based MDT between the Legacy-PE and the Egress-PE.

Figure 4.

The Egress-PE sends PIM Join messages across the PIM-based MDT towards the Legacy-Source-PE for (S, G) after SPT switchover. The incoming interface on the Egress-PE is now mdtone. The RP-PE is no longer a turnaround router for multicast traffic.

```
RP/0/RP1/CPU0:PE1#show mrib vrf one route 225.1.1.1

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, MF - MPLS Encap, 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

(*,225.1.1.1) RPF nbr: 10.1.100.2 Flags: C RPF
Up: 00:09:59
Incoming Interface List
  Lmdtone Flags: A LMI, Up: 00:09:59
Outgoing Interface List
  GigabitEthernet0/0/0/9 Flags: F NS, Up: 00:09:59

(10.2.3.10,225.1.1.1) RPF nbr: 10.1.100.3 Flags: RPF
Up: 00:14:29
Incoming Interface List
  mdtone Flags: A MI, Up: 00:14:29
```

```
Outgoing Interface List
   GigabitEthernet0/0/0/9 Flags: F NS, Up: 00:14:29
```
And PE1 has this PIM RPF info for the Source:

```
RP/0/RP1/CPU0:PE1#show pim vrf one rpf 10.2.3.10

Table: IPv4-Unicast-default
* 10.2.3.10/32 [200/0]
  via mdtone with rpf neighbor 10.1.100.3
  RT:1:1 ,Connector: 1:3:10.1.100.3, Nexthop: 10.1.100.3
```
This means that the traffic now flows directly from Legacy-Source-PE to Egress-PE in the core network across the PIM-based MDT. See Figure 5.
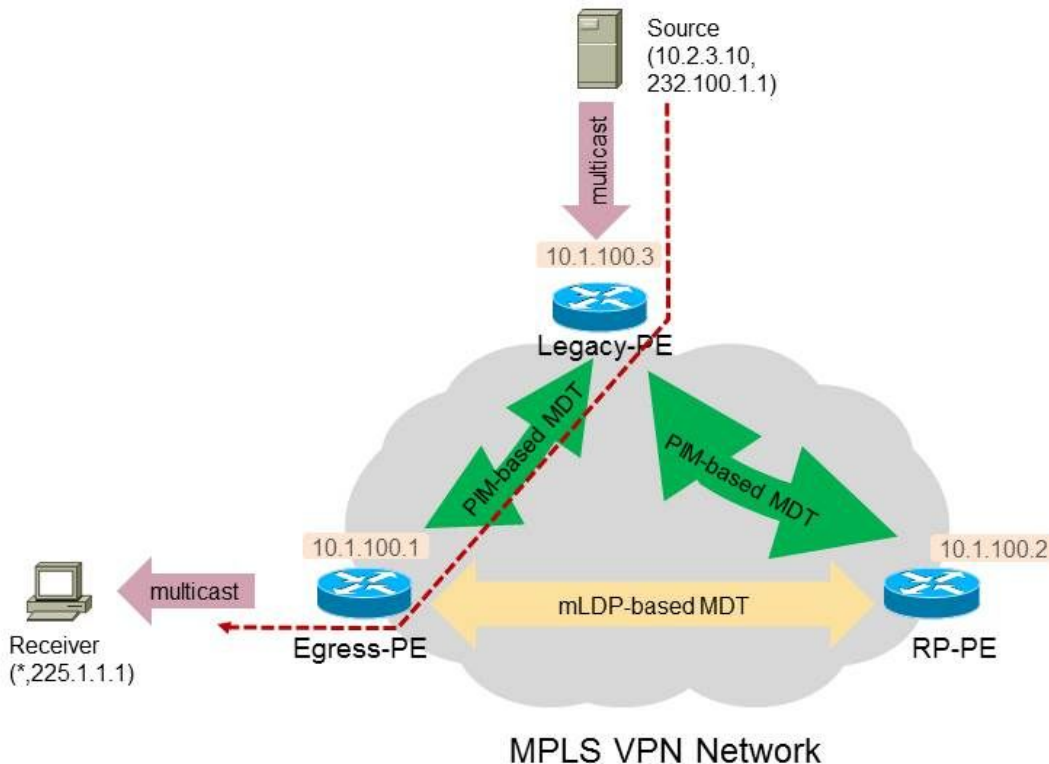

Figure 5.

Conclusion

All non-legacy PE routers, which are Receiver-PE or RP-PE routers, must have the configuration in place for migrating the core-tree protocols and migrating the C-signaling protocols.

Alternatively, a workaround is to ensure that the SPT-switchover does not occur, but then the routing of the multicast traffic might not be over the shortest path in the core of the network.