

# mVPN Profile Migration with Turnaround PE Router

## Contents

[Introduction](#)

[Solution](#)

[Configuration](#)

[Configuration of the RR](#)

[Configuration of the Source PE](#)

[Configuration of the TA PE](#)

[Configuration of the egress PE](#)

[Verification](#)

[Profile 6 PE – PE3](#)

[TA PE](#)

[Profile 0 PE – PE2](#)

[RR](#)

[Exit Strategy](#)

[Conclusion](#)

## Introduction

This document describes a migration strategy from mVPN (multicast Virtual Private Network) profile 0 to an mLDP-based (multipoint Label Distribution Protocol) profile in Cisco IOS<sup>®</sup>-XR using a turnaround router.

A network is running mVPN profile 0, which is the profile with PIM (Protocol Independent Multicast) in the core network and PIM in the overlay. The network will migrate to a profile using mLDP in the core. Here, the migration is to profile 6: In-band signalling with mLDP and using VRF (Virtual Routing/Forwarding) on the PE (Provider Edge) routers.

The migration solution works for SSM (Source Specific Multicast) and for ASM (Any Source Multicast) traffic.

Look at image 1.




Image 1

Image 1 shows the network running profile 0. The sources are behind PE1. The migration is towards profile 6, but the issue is that the ingress PE (Provider Edge) router, PE1, is a legacy router that cannot immediately migrate to profile 6. The solution is to keep using profile 0 on PE1 and use a TurnAround (TA) PE router in the network, that takes care of switching the multicast traffic from profile 0 to 6. The solution using a turnaround router is a temporary solution until the source PE can be migrated to profile 6 as well. The arrows on the topology show the multicast traffic flow.

## Solution

The solution needs the following:

- A turnaround router able of running profile 0 and profile 6.
- SAFI (Subsequent Address Family Identifier) 2 routing must be enabled on the non-legacy routers. This is key to this solution. The RPF (Reverse Path Forwarding) towards the source (or the RP (Rendez-Vous Point) for ASM) needs to be towards the TA PE router. For this, there must be static routes for sources and RP (if ASM is used) in the VRF in SAFI 2 on the turnaround router. These static routes are advertised on the TA PE router by BGP in SAFI 129 (vpng4 multicast). The static routes are in SAFI 2 as to not override the routes in SAFI 1 (unicast) and not override the unicast forwarding decision on the TA PE router and on the PE routers receiving the SAFI 129 routes.
- SAFI 129 in BGP is used on the profile 6 PE and RR (Route Reflector) routers. SAFI 2 in BGP is used on the profile 6 PE routers. This SAFI 2 carries the multicast routes in BGP used for multicast RPF, overriding the unicast routes. SAFI 2 is the multicast routes in the VRF, and the SAFI 129 routes are the multicast routes for vpng4.

The turnaround PE router is responsible for attracting the multicast traffic in profile 0 and sending it back out into the core over profile 6. This makes it the turnaround router. The turnaround does not

need a locally connected receiver, but it could have one.

## Configuration

### Configuration of the RR

```
router bgp 65001
  bgp router-id 10.0.0.5
  address-family ipv4 unicast
  !
  address-family vpnv4 unicast
  !
address-family ipv4 mdt    ## for profile 0
  !
  address-family ipv4 mvpn
  !
address-family vpnv4 multicast  ## SAFI 129
  !
  neighbor 10.0.0.1          ## profile 0 peer
  remote-as 65001
  update-source Loopback0
  address-family vpnv4 unicast
    route-reflector-client
  !
  address-family ipv4 mdt
    route-reflector-client
  !
  !
  neighbor 10.0.0.2          ## profile 0 peer
  remote-as 65001
  update-source Loopback0
  address-family vpnv4 unicast
    route-reflector-client
  !
  address-family ipv4 mdt
    route-reflector-client
  !
  !
  neighbor 10.0.0.3          ## TA peer
  remote-as 65001
  update-source Loopback0
  address-family vpnv4 unicast
    route-reflector-client
  !
  address-family ipv4 mvpn
    route-reflector-client
  !
address-family vpnv4 multicast  ## SAFI 129
  route-reflector-client
  !
  !
  neighbor 10.0.0.4          ## profile 6 peer
  remote-as 65001
  update-source Loopback0
  address-family vpnv4 unicast
    route-reflector-client
  !
  address-family ipv4 mdt
    route-reflector-client
  !
  !
```

```

address-family ipv4 mvpn
  route-reflector-client
!
address-family vpnv4 multicast
  route-reflector-client

```

The RR is reflecting routes for profile 0 (AF ipv4 mdt). MDT stands for Multicast Distribution Tree.

The RR needs SAFI 129. This is address family *vpnv4 multicast*. BGP sessions for this AF between the RR and every router running profile 6 need to be there.

## Configuration of the Source PE

**Note:** Source PE configuration must be added to any other profile 0 PE that is part of migration.

```

vrf one
  address-family ipv4 unicast
  import route-target
    65001:1
  !
  export route-target
    65001:1

router bgp 65001
  bgp router-id 10.0.0.1
  address-family vpnv4 unicast
  !
address-family ipv4 mdt
  !
  neighbor 10.0.0.5
  remote-as 65001
  update-source Loopback0
  address-family vpnv4 unicast
  !
address-family ipv4 mdt
  !
  !
  vrf one
  rd 1:2
  address-family ipv4 unicast
  redistribute onnedcted

multicast-routing
  address-family ipv4
  interface Loopback0
  enable
  !
  interface GigabitEthernet0/0/0/0
  enable
  !
  !
  vrf one
  address-family ipv4
  interface GigabitEthernet0/0/0/1
  enable
  !
  mdt source Loopback0

```

```
rate-per-route
mdt default ipv4 232.1.1.1  ## profile 0 Default MDT
```

The source PE router has the configuration for profile 0 only. There is no SAFI 129 or SAFI 2 configured. There is no profile 6 configuration.

## Configuration of the TA PE

```
vrf one
address-family ipv4 unicast
import route-target
65001:1
!
export route-target
65001:1
!
!
address-family ipv4 multicast  ## SAFI 2
import route-target
65001:1
!
export route-target
65001:1

router bgp 65001
bgp router-id 10.0.0.4
address-family ipv4 unicast
!
address-family ipv4 multicast  ## this is needed to have the static route in SAFI 2
!
address-family vpnv4 unicast
!
address-family ipv4 mdt  ## for profile 0
!
address-family ipv4 mvpn
!
address-family vpnv4 multicast  ## SAFI 129
!
neighbor 10.0.0.5  ## RR peer
remote-as 65001
update-source Loopback0
address-family vpnv4 unicast
!
address-family ipv4 mdt
!
address-family ipv4 mvpn
!
address-family vpnv4 multicast  ## SAFI 129
!
!
vrf one
rd 1:4
address-family ipv4 unicast
redistribute connected
redistribute static
!
address-family ipv4 multicast  ## SAFI 2
redistribute connected
redistribute static  ## redistribute SAFI 2 static routes
```

```

!
!

router static
 vrf one
  address-family ipv4 multicast
    10.1.6.0/24 vrf default 10.0.0.1  ## SAFI 2 static route

route-policy rpf-PE-TA
 set core-tree pim-default
end-polic

multicast-routing
 address-family ipv4
 interface Loopback0
  enable
!
 interface GigabitEthernet0/0/0/0
  enable
!
!
 vrf one
 address-family ipv4
  mdt source Loopback0
  rate-per-route
    mdt default ipv4 232.1.1.1  ## profile 0
    mdt mldp in-band-signaling ipv4  ## profile 6
!
!
!
router pim
 vrf one
 address-family ipv4
  rpf topology route-policy rpf-PE-TA

```

The TA PE needs SAFI 129. This is address family vpnv4 multicast towards the RR.

SAFI 2 is needed under the VRF and BGP.

The static route in the VRF towards the source (or RP for ASM) is needed, pointing to the ingress PE router. This static SAFI 2 route must be redistributed, as a SAFI 129 route in BGP. This SAFI 129 route is received by profile 6 PE routes in BGP as a SAFI 129 route and installed as a SAFI 2 route in the VRF.

Both profile 0 and 6 configuration is needed. The RPF topology command is configured for profile 0, as this is where the source (or RP) is present.

**Note:** The turnaround router needs to have a physical (sub)interface (not a loopback interface) enabled for multicast routing in the VRF. If this is not the case, then the multicast routes are not installed in the linecards, and the multicast traffic will not be turned around.

## Configuration of the egress PE

```

vrf one
 address-family ipv4 unicast
 import route-target

```

```

65001:1
export route-target
65001:1
!
address-family ipv4 multicast    ## SAFI 2
import route-target
65001:1
!
export route-target
65001:1

router bgp 65001
  bgp router-id 10.0.0.3
  address-family ipv4 unicast
  !
  address-family vpnv4 unicast
  !
  address-family ipv4 mvpn
  !
address-family vpnv4 multicast    ## SAFI 129
  !
  neighbor 10.0.0.5    ## RR peer
  remote-as 65001
  update-source Loopback0
  address-family vpnv4 unicast
  !
  address-family ipv4 mvpn
  !
address-family vpnv4 multicast    ## SAFI 129
  !
  !
  vrf one
  rd 1:3
  address-family ipv4 unicast
  redistribute connected
  !
address-family ipv4 multicast    ## SAFI 2
  redistribute connected
  redistribute static

route-policy in-band-mldp
  set core-tree mldp-inband    ## profile 6
end-polic

multicast-routing
  address-family ipv4
  interface Loopback0
  enable
  !
  !
  vrf one
  address-family ipv4
  interface GigabitEthernet0/0/0/1
  enable
  !
  mdt source Loopback0
  rate-per-route
  mdt mldp in-band-signaling ipv4    ## profile 6
  !
  !
!
router pim
  address-family ipv4
  interface Loopback0

```

```

enable
!
!
vrf one
address-family ipv4
  rpf topology route-policy in-band-mldp  ## profile 6

```

The egress PE router has the configuration for profile 6. In addition to that: in order for the egress PE router to successfully RPF towards the TA PE router for the source (or the RP for ASM), it needs the configuration for SAFI 2 and SAFI 129.

## Verification

### Profile 6 PE – PE3

```

RP/0/RP0/CPU0:PE3#show bgp vpnv4 multicast rd 1:3 10.1.6.0/24
BGP routing table entry for 10.1.6.0/24, Route Distinguisher: 1:3
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          136      136
Last Modified: Jul  7 12:02:27.278 for 00:49:22
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Not advertised to any peer
Local
  10.0.0.4 (metric 30) from 10.0.0.5 (10.0.0.4)
    Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best, import-
candidate, imported
    Received Path ID 0, Local Path ID 1, version 136
    Extended community: RT:65001:1
    Originator: 10.0.0.4, Cluster list: 10.0.0.5
    Connector: type: 1, Value:1:4:10.0.0.4
    Source AFI: VPNv4 Multicast, Source VRF: default, Source Route Distinguisher: 1:4

```

The next-hop is 10.0.0.4 which is the TA PE router.

```

RP/0/RP0/CPU0:PE3#show route vrf one ipv4 multicast 10.1.6.0/24
Routing entry for 10.1.6.0/24
  Known via "bgp 65001", distance 200, metric 0, type internal
  Installed Jul  7 12:02:27.236 for 00:50:44
Routing Descriptor Blocks
  10.0.0.4, from 10.0.0.5
    Nexthop in Vrf: "default", Table: "default", IPv4 Unicast, Table Id: 0xe0000000
    Route metric is 0
  No advertising protos.

```

```

RP/0/RP0/CPU0:PE3#show pim vrf one rpf 10.1.6.6
Table: IPv4-Multicast-default
* 10.1.6.6/32 [200/0]
  via Imdtone with rpf neighbor 10.0.0.4

```



Connector: 1:4:10.0.0.4, Nexthop: 10.0.0.4

The RPF is towards the TA PE router.

```
RP/0/RP0/CPU0:PE3#show mrib vrf one route 232.1.1.1
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, 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, TRMI - TREE SID MDT Interface, MH - Multihome Interface
(10.1.6.6,232.1.1.1) RPF nbr: 10.0.0.4 Flags: RPF
Up: 09:29:38
Incoming Interface List
  Imdtone Flags: A LMI, Up: 00:47:04
Outgoing Interface List
  GigabitEthernet0/0/0/1 Flags: F NS, Up: 09:29:38
```

The ingress interface is profile 6.

## TA PE

```
RP/0/RP0/CPU0:TA-PE#show bgp vpnv4 multicast rd 1:4 10.1.6.0/24
BGP routing table entry for 10.1.6.0/24, Route Distinguisher: 1:4
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          80        80
Last Modified: Jul  7 12:02:27.317 for 01:04:42
Paths: (1 available, best #1)
  Advertised to peers (in unique update groups):
    10.0.0.5
  Path #1: Received by speaker 0
Advertised to peers (in unique update groups):
  10.0.0.5
Local
  10.0.0.1 (metric 30) from 0.0.0.0 (10.0.0.4)
  Origin incomplete, metric 0, localpref 100, weight 32768, valid, redistributed, best,
  group-best, import-candidate
  Received Path ID 0, Local Path ID 1, version 80
  Extended community: RT:65001:1
```

This route is local, but the next-hop is the source PE (10.0.0.1). The route is advertised to the RR (10.0.0.5).

```
RP/0/RP0/CPU0:TA-PE#show route vrf one ipv4 multicast 10.1.6.0/24
Routing entry for 10.1.6.0/24
  Known via "static", distance 1, metric 0
  Installed Jul  7 12:02:27.234 for 01:07:01
  Routing Descriptor Blocks
    10.0.0.1
      Nexthop in Vrf: "default", Table: "default", IPv4 Multicast, Table Id: 0xe0100000
      Route metric is 0, Wt is 1
  No advertising protos.
```

```
RP/0/RP0/CPU0:PE-TA#show pim vrf one rpf 10.1.6.6
Table: IPv4-Multicast-default
* 10.1.6.6/32 [1/0]
  via mdtone with rpf neighbor 10.0.0.1
```

The RPF is towards the source router using profile 0.

```
RP/0/RP0/CPU0:TA-PE#show mrrib vrf one route 232.1.1.1
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, 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, TRMI - TREE SID MDT Interface, MH - Multihome Interface
(10.1.6.6,232.1.1.1) RPF nbr: 10.0.0.1 Flags: RPF
Up: 01:13:28
Incoming Interface List
  mdtone Flags: A MI, Up: 01:13:28
Outgoing Interface List
  Imdtone Flags: F LMI, Up: 01:13:28
```

The incoming interface is MDT of profile 0 and the outgoing interface is MDT of profile 6. This is the turnaround.

```
RP/0/RP0/CPU0:TA-PE#show mfib vrf one route 232.1.1.1 detail
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  
 (10.1.6.6,232.1.1.1), Flags: EID RPFID  
 Up: 01:15:01  
 Last Used: never  
 SW Forwarding Counts: 0/0/0  
 SW Replication Counts: 0/0/0  
 SW Failure Counts: 0/0/0/0/0  
 Route ver: 0xd672  
 MVPN Info :-  
 Associated Table ID : 0xe0000000  
 MDT Handle: 0x0, MDT Probe:N [N], Rate:Y, Acc:N  
 MDT SW Ingress Encap V4/V6, Egress decap: 0 / 0, 0  
 Encap ID: 262146, RPF ID: 3  
 Local Receiver: False, Turnaround: True  
**mdtone** Flags: **A** MI, Up:01:15:01  
**lmdtone** Flags: **F** LMI, Up:01:15:01

The incoming interface is MDT of profile 0 and the outgoing interface is MDT of profile 6. This is the turnaround.

## Profile 0 PE – PE2

```

RP/0/RP0/CPU0:PE2#show pim vrf one rpf 10.1.6.6
Table: IPv4-Unicast-default
* 10.1.6.6/32 [200/0]
  via mdtone with rpf neighbor 10.0.0.1
  Connector: 1:1:10.0.0.1, Nexthop: 10.0.0.1
  
```

The RPF is towards the profile 0 ingress PE router.

```

RP/0/RP0/CPU0:PE2#show mrib vrf one route 232.1.1.1
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, 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, TRMI - TREE SID MDT Interface, MH - Multihome Interface
(10.1.6.6,232.1.1.1) RPF nbr: 10.0.0.1 Flags: RPF
Up: 1d22h
Incoming Interface List
  mdtone Flags: A MI, Up: 02:49:35
  
```

Outgoing Interface List  
GigabitEthernet0/0/0/1 Flags: F NS, Up: 1d22h

The ingress interface is profile 0.

## RR

```
RP/0/RP0/CPU0:P#show bgp vpnv4 multicast rd 1:4 10.1.6.0/24
BGP routing table entry for 10.1.6.0/24, Route Distinguisher: 1:4
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          84        84
Last Modified: Jul  7 12:02:27.979 for 00:54:33
Paths: (1 available, best #1)
  Advertised to update-groups (with more than one peer):
    0.2
  Path #1: Received by speaker 0
Advertised to update-groups (with more than one peer):
  0.2
Local, (Received from a RR-client)
  10.0.0.4 (metric 20) from 10.0.0.4 (10.0.0.4)
  Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best, import-
candidate, not-in-vrf
  Received Path ID 0, Local Path ID 1, version 84
  Extended community: RT:65001:1
  Connector: type: 1, Value:1:4:10.0.0.4
```

The route towards the source is advertised to profile 6 PE routers and is received from the TA router (10.0.0.4).

## Exit Strategy

The migration solution using a turnaround router is a temporary solution. The migration should be concluded by migrating every PE router to profile 6. This can be done with the following steps:

- Add a new source PE router
- Add a route policy on PE routers to either RPF to the legacy source PE router (Profile 0), or to the TA PE router (profile 6), or to the new source PE router (profile 6) Specify a source and/or a group in the route policy
- Move the multicast source to the new source PE router
- Remove the old source PE router once all the multicast groups are migrated to the new source PE router

## Conclusion

Using a turnaround router for mVPN can be an easy way to facilitate the migration from profile 0 to a new mVPN profile as a temporary solution while waiting for a newer source PE router which is capable of running the new mVPN profile.