

# Troubleshoot mVPN Profile 14

## Contents

- [Introduction](#)
- [Prerequisites](#)
- [Requirements](#)
- [Components Used](#)
- [mLDP Concepts](#)
- [Core Tree - mLDP](#)
- [MP2MP LSP vs P2MP LSP](#)
- [Recursive Forwarding](#)
- [Configuration](#)
- [Label Switch Multicast \(LSM\)/NG mVPN](#)
- [BGP](#)
- [AF IPv4 mVPN](#)
- [BGP mVPN Route Types](#)
- [Route Types in AF IPv4 MVPN](#)
- [Intra-AS I-PMSI A-D Route \(Type 1\)](#)
- [Inter-AS I-PMSI A-D Route \(Type 2\)](#)
- [S-PMSI A-D Route \(Type 3\)](#)
- [Leaf A-D Route \(Type 4\)](#)
- [Source Active A-D Route \(Type 5\)](#)
- [Shared Tree Join Route \(Type 6\)](#)
- [Source Tree Join Route \(Type 7\)](#)
- [AF IPv4 MVPN NLRI Encoding](#)
- [Configuration](#)
- [Topology](#)
- [Steps to Troubleshoot](#)
- [Tail End/Receiver PE Node](#)
- [IGMP Join Check](#)
- [Towards Core](#)
- [MID Router](#)
- [Towards Receivers](#)
- [Head End Router \(Multicast Source PE\)](#)
- [Towards Core](#)
- [Towards Multicast Source](#)
- [Log Collection](#)

## Introduction

This document describes how to troubleshoot Multicast VPN (mVPN) profile 14 within the Cisco IOS® XR.

## Prerequisites

You must have knowledge of Multicast profile14 and its behaviour.

## Requirements

Cisco recommends that MVPN BGP auto-discovery must be configured.

Also, it is recommended that you know these topics:

- Profiles “ 14
- Core Tree “ mLDP
- Service “ IPv4 VPN
- C-Overlay Signaling “ BGP
- Auto-Discovery “ BGP
- Model “ Partition MDT P2MP with BGP AD
- Encapsulation/forwarding “ MPLS

## Components Used

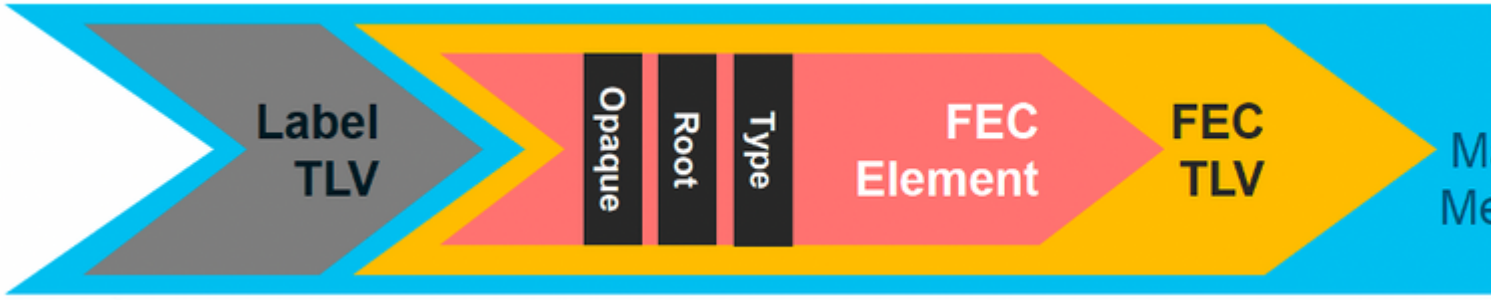
This document is not restricted to specific software and hardware versions.

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.

## mLDP Concepts

### Core Tree - mLDP

- Multipoint LDP = LDP + extensions
- P2MP tree -Receiver driven “Root learned from routing
- MP2MP tree “Configuration driven “Root configured
- Protection by MPLS TE or Loop-Free Alternate (LFA)
- No PHP “Top label identifies the tree
- Replication of mcast on the core routers
- FEC elements hold: Type of tree + Root + Opaque value: (S, G), MDT number, LSP ID



mLDP

## MP2MP LSP vs P2MP LSP

We have two ways to send traffic between the PEs.

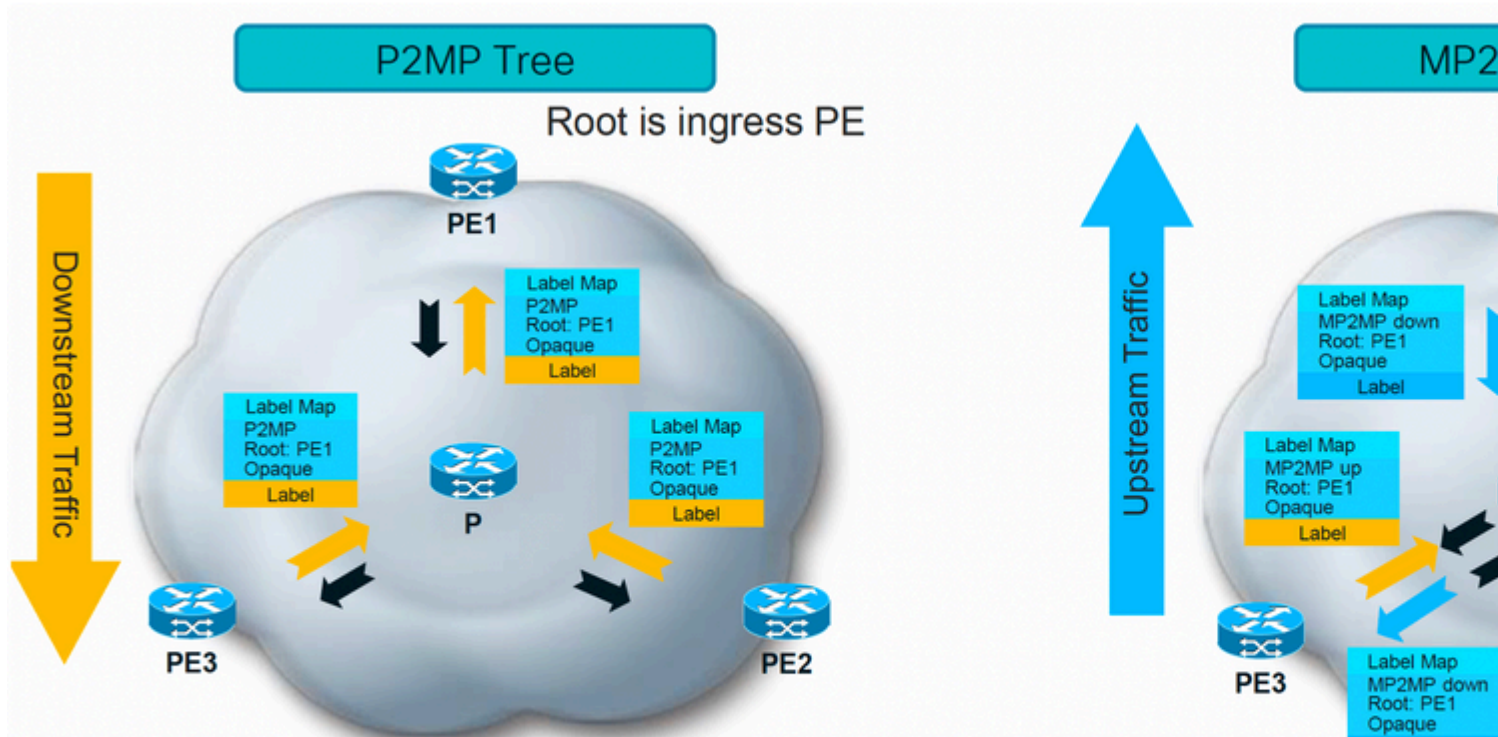
1. MP2MP - a static default MDT is created among all PEs for that particular user, this works as if all PEs for that user were connected in the same LAN segment.
2. P2MP - when a configurable threshold value is reached, the sender PE can create a new LSP so only the PEs that are involved in that communication see the traffic, this LSP is created to avoid unnecessary waste of bandwidth sending traffic to PEs that did not join the stream. This LSP only creates for an (S, G) entry in the user network, if you send traffic that exceeds the configured threshold through an (\*, G) entry, no new P2MP is created (the concept of PHP can not be used in this scenario when the leaf acts as a root, otherwise it does not know which user that traffic belongs to).

mLDP work with a concept as root, it is somewhat similar to a Rendezvous Point in the PIM-SM scenario, but remember, in mLDP, a PIM-free MPLS core is mentioned, and PIM is only enabled in the user-facing interfaces.

In the MP2MP tree, the roots are responsible to receive the traffic from a leaf and replicate it to all directly connected neighbors.

Every PE can have more than one root configured, they choose its active root based on metric, if the metric is the same the highest IP address is elected as the root for that particular PE.

In the P2MP tree, the root is always the device sending the traffic (In this document the term root is used more to reference the MP2MP tree unless specified).



```
log
hello-adjacency
neighbor
nsr
graceful-restart
!
igmp sync delay on-session-up 25
mldp
address-family ipv4
make-before-break delay 60 60
forwarding recursive
recursive-fec
```

This concept is explained in the troubleshoot section. This configuration is required in the MID and BUD router.

Head/Root = This router has a Multicast source connected

Tail/Leaf = This router has Multicast receivers

Bud = This router is in transit and has some receivers connected

Mid = This router is in transit

## **Label Switch Multicast (LSM)/NG mVPN**

- Leverage MPLS encapsulation > Share control and data plane with unicast
- Leverage new core tree protocols > mLDP, P2MP TE, BIER
- Fast ReRoute (FRR)
- Leverage proven functionality: Default and Data MDT
- More flexible designs per VPN
- Manageability: no need to track Multicast Groups per VPN/Default/Data MDT

## **BGP**

### **AF IPv4 mVPN**

AF IPv4 mVPN is added to be able to support the NG profiles. It is referred to as the MCAST-VPN NLRI. There are two main functions of this address family:

- Auto-Discovery (AD)
- Customer Signaling in overlay

Auto-Discovery means that the PE routers learn of each other in an automatic manner. PE routers learn which other PE routers are enabled for multicast for a VPN. This avoids the need to manually configure this on every PE router.

The user signaling in overlay is the signaling of (\*, G) or (S, G) information, but now in BGP instead of PIM.

Look at Figure A to see the overview of BGP AD IPv4 MVPN.

# BGP Auto Discovery + C-Signalling

BGP update

PMSI Tunnel Attribute (PTA)

mcast vprn

Flags (1 octet)	0-6 – Reserved 7 – L – Leaf information required
Tunnel type (1 octet)	0 No tunnel information present 1 RSVP-TE P2MP LSP 2 mLDP P2MP LSP 3 PIM-SSM Tree 4 PIM-SM Tree 5 PIM-Bidir Tree 6 Ingress Replication 7 mLDP MP2MP LSP
MPLS label (3 octets)	MPLS label
Tunnel Identifier (variable)	1 RSVP-TE P2MP LSP - <Extended Tunnel ID, Reserved, Tunnel ID, P2MP ID> 2 mLDP P2MP LSP - <P2MP FEC Element> 3 PIM-SSM Tree - <P- Root Node Address, P-Multicast Group> 4 PIM-SM Tree - <Sender Address, P-Multicast Group> 5 PIM-Bidir Tree - <Sender Address, P-Multicast Group> 6 Ingress replication - <unicast tunnel endpoint IP address of the local PE that is to be this PE's receiving endpoint address for the tunnel> 7 mLDP MP2MP LSP - <MP2MP FEC Element>

Route type (1 octet)	1 Intra-AS 2 Inter-AS 3 S-PMSI 4 Leaf A-D 5 Source 6 Shared 7 Source
Length (1 octet)	Length in field of MC
Route type specific (variable length)	One or mo  RD (8 octe MCAST se MCAST se MCAST g MCAST g Originating

Figure A - BGP IPv4 AD MVPN

There are two important pieces of information carried out by BGP in the AF IPv4 MVPN. The first part of information sits in the NLRI itself. There is a number from 1-7 that represents the Route Type. Each Route Type has a specific function. The first 5 are used for BGP-AD. Note that Route Type 3 is actually also used for Data MDT (S-PMSI) signaling.

The last two types, 6 and 7, are used for C-multicast signaling. Next to the Route Type, there is also specific Route Type information embedded which pertains to the RD (if VPN context is used), specific C-multicast info (Source, Group) and the IP address (PE) of the originator.

The second part is the PMSI Tunnel Attribute (PTA). This PTA holds information that pertains to the core tree tunnels that are used. The tunnel type indicates a number of which core tree protocol is used to build the tunnel. The tunnel identifier is a set of fields which uniquely describes the tunnel by specific core tree protocol fields.

## BGP mVPN Route Types

Route Type	Name	AD or C-mcast	Usage
1	Intra-AS I-PMSI A-D route	AD	AD Advertise member PE

2	Inter-AS I-PMSI A-D route	AD	Same as type 1, but for inter-AS
3	S-PMSI A-D route	AD	Advertised by Source PE to signal Data MDT
4	Leaf A-D route	AD	Advertised by Receiver-PE as a response to type 3 route with leaf info required flag set
5	Source Active A-D route	AD	Advertised by Source PE with an active source to facilitate SPT switchover
6	Shared Tree Join route	C-mcast	Advertise (*, G) by Receiver PE
7	Source Tree Join route	C-mcast	Advertise (S, G) by Receiver PE

## Route Types in AF IPv4 MVPN

This section provides an overview of the different route types and their purpose.

### Intra-AS I-PMSI A-D Route (Type 1)

Purpose: To advertise the PE as to participate in MVPN for a certain VPN. This is BGP-AD.

Encoding:

```

+-----+
|      RD      (8 octets)      |
+-----+
| Originating Router's IP Addr |
+-----+

```

### Inter-AS I-PMSI A-D Route (Type 2)

#### Purpose

Advertised by an ASBR, so that Inter-Autonomous mVPN can be signaled.

#### Encoding

```

+-----+
|      RD      (8 octets)      |
+-----+
|      Source AS (4 octets)    |
+-----+

```

### S-PMSI A-D Route (Type 3)

#### Purpose

1. Advertised by the PE router, in order for it to signal to other PE routers that it is the Ingress PE for a certain C-(S, G) and that it switches over to a Data MDT. The core tree information needed to receive PE routers joins this stream is encoded in the PTA.
2. Advertised by a PE router, it indicates that it wants to participate in a Partitioned MDT model. In that case, the Source and Group fields are wildcard (\*,\*) elements.

#### Encoding

```

+-----+
|      RD      (8 octets)      |
+-----+
| Multicast Source Length (1 octet) |
+-----+
| Multicast Source (variable)      |
+-----+
| Multicast Group Length (1 octet) |
+-----+
| Multicast Group (variable)       |
+-----+
| Originating Router's IP Addr     |
+-----+

```

The Source and Group fields relate to C-multicast state.

### Leaf A-D Route (Type 4)

#### Purpose

This message is sent in response to obtain an Inter-AS I-PMSI A-D route or a S-PMSI A-D route, only if the Leaf Information (LI) Required bit is set.

## Encoding

```
+-----+
|   Route Key (variable)   |
+-----+
| Originating Router's IP Addr |
+-----+
```

The Route Key is set to the NLRI of the received route with the Leaf Information Required bit set.

If the Leaf A-D is generated that corresponds to inter-as I-PMSI route, then its NLRI looks like this:

```
+-----+
|   Type (1 octet): 4     |
+-----+
|   Length (1 octet)     |
+-----+
|   Type (1 octet): 2     |
+-----+
|   Length (1 octet)     |
+-----+
|   RD (8 octet)         |
+-----+
| Source AS (4 octet)    |
+-----+
| Originating Router's IP Addr |
+-----+
```

If the Leaf A-D route is generated that corresponds to the S-PMSI route, then its NLRI looks like this:

```
+-----+
|   Type (1 octet): 4     |
```



```

+-----+
| Length (1 octet) |
+-----+
| Type (1 octet): 3 |
+-----+
| Length (1 octet) |
+-----+
| RD (8 octet) |
+-----+
| Multicast Source Length (1 octet) |
+-----+
| Multicast Source (Variable) |
+-----+
| Multicast Group Length (1 octet) |
+-----+
| Multicast Group (Variable) |
+-----+
| Origin Routers IP addr for SPMSI |
+-----+
|Origin Router's IP Addr for Leaf-ad|
+-----+

```

## Source Active A-D Route (Type 5)

### Purpose

The usage of the Source Active Route is related to PIM Sparse Mode in overlay signaling. Refer to that section for more information.

### Encoding

```

+-----+
| RD (8 octets) |
+-----+

```

```

| Multicast Source Length (1 octet) |
+-----+
| Multicast Source (variable)      |
+-----+
| Multicast Group Length (1 octet) |
+-----+
| Multicast Group (variable)       |
+-----+

```

## Shared Tree Join Route (Type 6)

### Purpose

The BGP Update or Withdraw message is the equivalent of a PIM (\*, G) Join or Prune message.

### Encoding

Type 6 & 7 Routes have the same NLRI encoding.

```

+-----+
| RD (8 octets) |
+-----+
| Source AS (4 octets) |
+-----+
| Multicast Source Length (1 octet) |
+-----+
| Multicast Source (variable) |
+-----+
| Multicast Group Length (1 octet) |
+-----+
| Multicast Group (variable) |
+-----+

```

## Source Tree Join Route (Type 7)

### Purpose

The BGP Update or Withdraw message is the equivalent of a PIM (S,G) Join or Prune message.

### Encoding

Type 6 & 7 Routes have the same NLRI encoding.

```

+-----+
| RD (8 octets) |
+-----+
| Source AS (4 octets) |
+-----+
| Multicast Source Length (1 octet) |
+-----+

```

```

+-----+
| Multicast Source (variable) |
+-----+
| Multicast Group Length (1 octet) |
+-----+
| Multicast Group (variable) |
+-----+

```

## AF IPv4 MVPN NLRI Encoding

This section lists the encoding of the network part for IPv4 MVPN routes per Route Type. Note that the RD is not included for Route Types 1, 2, and 3. The network part is listed with the number in () and it represents the number of octets for each part. The first number is always the Route Type (1-7). The Prefix-length part is listed in the number of bits.

### Intra-AS A\_D Route

```

Network:          [1] [PE address (4)]
Prefix-length:    40

```

### Inter-as A-D route

```

Network:          [2] [AS (4)]
Prefix-length:    40

```

### S-PMSI A-D route

```

Network:          [3] [srclen (1)] [src (4)] [1 grplen] [grp (4)] [PE address (4)]
Prefix-length:    120

```

### Leaf A-D route

```

Network:          [4] [2] [RD] (8)] [AS (4)] [ PE address (4)]
Prefix-length:    144

```

or

Network: [4] [3] [RD (8)] [srclen (1)] [src (4)] [grplen (1)] [grp (4)] [PE address (4)]  
Prefix-length: 224

### Source-Active A-D route

Network: [5] [srclen (1)] [src (4)] [grplen (1)] [grp (4)]  
Prefix-length: 88

### C-Multicast Route - Shared-tree Join

Network: [6] [RD (8)] [AS (4)] [srclen (1)] [src (4)] [grplen (1)] [grp (4)]  
Prefix-length: 184

### C-Multicast Route Source-tree Join

Network: [7] [RD (8)] [AS (4)] [srclen (1)] [src (4)] [grplen (1)] [4 grp]  
Prefix-length: 184

The network and prefix length for each route type can be found in the output of `show bgp ipv4 mvpn` command in Cisco IOS or Cisco IOS XR. Here is an example of the output for this command from Cisco IOS XR:

```
DRP/0/3/CPU1:Router#show bgp ipv4 mvpn vrf one
Route Distinguisher: 1:2 (default for vrf one)

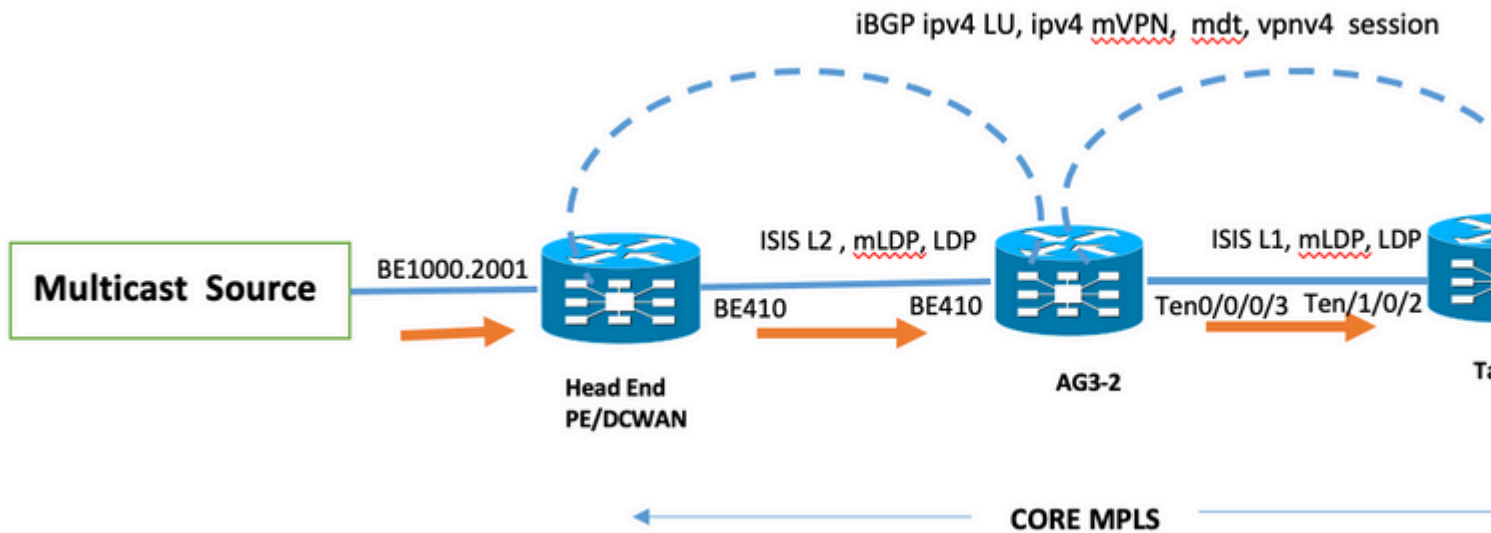
*>i[1][10.1.100.1]/40 10.1.100.1 100 0 i
*>i [3][1:1][*][*][10.100.1.1]/14
10.100.1.1 0 100 0 ?
*>i[3][3][1:2][32][10.2.2.9][32][225.1.1.1][10.1.100.2][10.1.100.5]/224
10.1.100.5 100 0 i
*>i[4][3][1:2][32][10.2.100.9][32][225.1.1.1][10.1.100.2][10.1.100.1]/224
10.1.100.1 100 0 i
*>i[5][32][10.2.2.9][32][225.1.1.1]/88
10.1.100.2 100 0 i
*>i[6][1:5][1][32][10.2.100.12][32][225.1.1.1]/184
10.1.100.1 100 0 i
*>i[7][1:2][1][32][10.2.2.9][32][232.1.1.1]/184
10.1.100.1 100 0 i
```

# Configuration

Refer to this [document](#) to know how to configure mVPN with Cisco IOS XR.

# Topology

Partitioned MDT - mLDP P2MP - BGP-AD - BGP C-Mcast Signaling - SSM



Multicast Profile 14 topology

Platform: ASR9K

## Steps to Troubleshoot

1. Check if there are any system-level alarms on the reported node (Source PE/Receiver PE/MID Router).

```
Show pfm location all
```

```
Show logging (Check for critical or major logs)
```

```
admin show alarm
```

2. Check the MPLS core connectivity between the Headend router and Tail End Router.

3. Check the local connectivity of the Multicast Source/Receiver to connected PE routers.

4. Check the End-to-End connectivity between Multicast Source and receiver.

5. Check the configuration as per the configuration guide mentioned in the previously mentioned document.

6. Basic verification from system health check:

- Check if any system alarm (**show pfm location all**)
- Check memory and CPU

```
show memory summary location all
```

```
show shmwin summary location all
```

```
show processes memory detail location all
```

```
show process cpu location <>
```

- Check interface and QoS Stats
- Check show logging, if there is any critical notification or errors
- Check for any configuration changes due to which impact or issue observed
- Platform drops

Fri Mar 12 02:43:14.720 EST

232.11.11.11 Bundle-Ether5000.1 15:26:11 never 10.1.0.2

232.11.11.11 Bundle-Ether5000.2 15:26:03 not used 10.1.1.2

<#root>

RP/0/RSP0/CPU0:AG2-1#sh igmp vrf MLDP-14 groups 232.11.11.11 detail

Fri Mar 12 02:44:19.393 EST

Interface: Bundle-Ether5000.1

Group: 232.11.11.11

Uptime: 15:27:15 >>>>>>>

**Check the IGMP join timestamp**

Router mode: INCLUDE

Host mode: INCLUDE

Last reporter: 10.1.0.2

Suppress: 0

Group source list:

Source Address Uptime Expires Fwd Flags

10.1.0.2 15:27:15 00:03:46 Yes Remote Local e

Interface: Bundle-Ether5000.2

Group: 232.11.11.11

>>> **Multicast Group**

Uptime: 15:27:07

Router mode: INCLUDE

Host mode: INCLUDE

Last reporter: 10.1.1.2

>>> **Receiver IP**

Suppress: 0

Group source list:

Source Address Uptime Expires Fwd Flags

10.1.0.2 15:27:07 00:03:59 Yes Remote 4

>>> Source is 10.1.0.2 & uptime

RP/0/RSP0/CPU0:AG2-1#

Check the interface status with bit rates and errors:

RP/0/RSP0/CPU0:AG2-1#sh int be5000

Fri Mar 12 02:45:32.033 EST

Bundle-Ether5000 is up, line protocol is up

Interface state transitions: 1

Hardware is Aggregated Ethernet interface(s), address is 0026.9813.b986

Internet address is Unknown

MTU 1514 bytes, BW 10000000 Kbit (Max: 10000000 Kbit)

reliability 255/255, txload 20/255, rxload 0/255

Encapsulation ARPA,

Full-duplex, 10000Mb/s

loopback not set,

Last link flapped 15:28:35

No. of members in this bundle: 1

TenGigE0/0/0/17 Full-duplex 10000Mb/s Active

Last input 00:00:00, output 00:00:00

Last clearing of "show interface" counters never

30 second input rate 1000 bits/sec, 2 packets/sec

30 second output rate 793646000 bits/sec, 200031 packets/sec

118212 packets input, 8341734 bytes, 0 total input drops

0 drops for unrecognized upper-level protocol

Received 6 broadcast packets, 6819 multicast packets

0 runts, 0 giants, 0 throttles, 0 parity

0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort

2857201595 packets output, 1416646430161 bytes, 0 total output drops



Output 242 broadcast packets, 2857201353 multicast packets  
0 output errors, 0 underruns, 0 applique, 0 resets  
0 output buffer failures, 0 output buffers swapped out  
0 carrier transitions

Check if any LPTS drops:

```
show lpts pifib hardware static-police location <> | i "IGMP|PIM|Flow"
```

```
RP/0/RSP0/CPU0:AG2-1#show lpts pifib hardware static-police location 0/0/CPU0 | i "IGMP|PIM|Flow"
```

Fri Mar 12 02:45:56.259 EST

Punt Reason SID Flow Rate Burst Rate Accepted Dropped Destination

PUNT\_IGMP\_SNOOP NETIO\_MED 4000 2000 0 0 0x0030 (0/RSP0/CPU0)

PUNT\_IFIB\_IGMP\_OPT NETIO\_LOW 2000 500 0 0 Local

PUNT\_IFIB\_PIM\_OPT NETIO\_LOW 2000 500 36 0 Local

PIM Entry Check:

```
RP/0/RSP0/CPU0:AG2-1#sh pim vrf MLDP-14 topology 232.11.11.11
```

Fri Mar 12 02:47:57.462 EST

IP PIM Multicast Topology Table

Entry state: (\*S,G)[RPT/SPT] Protocol Uptime Info

Entry flags: KAT - Keep Alive Timer, AA - Assume Alive, PA - Probe Alive

RA - Really Alive, IA - Inherit Alive, LH - Last Hop

DSS - Don't Signal Sources, RR - Register Received

SR - Sending Registers, SNR - Sending Null Registers

E - MSDP External, EX - Extranet

MFA - Mofrr Active, MFP - Mofrr Primary, MFB - Mofrr Backup

DCC - Don't Check Connected, ME - MDT Encap, MD - MDT Decap

MT - Crossed Data MDT threshold, MA - Data MDT Assigned

SAJ - BGP Source Active Joined, SAR - BGP Source Active Received,

SAS - BGP Source Active Sent, IM - Inband mLDP, X - VxLAN

Interface state: Name, Uptime, Fwd, Info

Interface flags: LI - Local Interest, LD - Local Dissinterest,

II - Internal Interest, ID - Internal Dissinterest,

LH - Last Hop, AS - Assert, AB - Admin Boundary, EX - Extranet,

BGP - BGP C-Multicast Join, BP - BGP Source Active Prune,

MVS - MVPN Safi Learned, MV6S - MVPN IPv6 Safi Learned

(10.1.0.2,232.11.11.11)SPT SSM Up: 15:30:53

JP: Join(BGP) RPF: LmdtMLDP-14,10.44.0.10 Flags:

Bundle-Ether5000.1 15:30:53 fwd LI LH >>>> UPTIME of PIM & all egress Interfaces last hop & local inte

Bundle-Ether5000.2 15:30:45 fwd LI LH

RP/0/RSP0/CPU0:AG2-1#

<#root>

RP/0/RSP0/CPU0:AG2-1#sh pim vrf MLDP-14 interface be5000.1 detail

Fri Mar 12 02:49:36.698 EST

PIM interfaces in VRF MLDP-14

IP PIM Multicast Interface State

Flag: B - Bidir enabled, NB - Bidir disabled

P - PIM Proxy enabled, NP - PIM Proxy disabled

V - Virtual Interface

BFD State - State/Interval/Multiplier

Interface PIM Nbr Hello DR

Count Intvl Prior

Bundle-Ether5000.1 on 2 30 1

>>> PIM ON, Hello Interval 30 sec

Primary Address : 10.1.0.1

Flags : NB NP V

```
BFD : On/500 ms/3 >>> BFD ON
DR : this system >>> This is DR
Propagation delay : 500
Override Interval : 2500
Hello Timer : 00:00:27
Neighbor Filter : -
```

<#root>

```
RP/0/RSP0/CPU0:AG2-1#sh pim vrf MLDP-14 neighbor bundle-ether 5000.1 detail
Fri Mar 12 03:03:43.892 EST
PIM neighbors in VRF MLDP-14
```

```
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
10.1.0.1* Bundle-Ether5000.1 15:46:40 00:01:38 1 (DR) B E
```

>>>> **Verify DR**

```
Expiry Timer: 00:00:07
```

```
10.1.0.2 Bundle-Ether5000.1 15:46:27 00:01:25 0
```

```
BFD State: enabled
```

<#root>

```
RP/0/RSP0/CPU0:AG2-1#sh pim vrf MLDP-14 mdt cache | utili egrep -w "(10.1.0.2, 232.11.11.11)"
```

```
Fri Mar 12 03:04:19.783 EST
```

```
10.44.0.10 (10.1.0.2, 232.11.11.11) [global-id 4] never >>>>
```

```
Verify Global ID for specific S,G
```

---

**Note:** This global id remains the same/unique from receiver to Source for this specific S, G.

---

PIM Errors/Stats:

```
show pim traffic
```

PIM Traffic Counters

Elapsed time since counters cleared: 1d01h

	Received	Sent	
Valid PIM Packets	15759217	15214426	>>> verify the pim packets increment
Hello	9207	12336	>>> verify the pim hello Negotiations
Join-Prune	1076805	531981	
Data Register	14673205	0	
Null Register	73205	0	
Register Stop	0	14673205	
Assert	0	0	
Batched Assert	0	0	
Bidir DF Election	0	0	
BSR Message	0	0	>>> If you have RP check this stats
Candidate-RP Adv.	0	0	>>> If you have RP check this stats
Join groups sent		0	
Prune groups sent		0	
Output JP bytes		0	
Output hello bytes		410	

Errors: >>>> Check below paramters for errors or drops

Malformed Packets	0
Bad Checksums	0
Socket Errors	0
Subnet Errors	0
Packets dropped since send queue was full	0
Packets dropped due to invalid socket	0
Packets which couldn't be accessed	0
Packets sent on Loopback Errors	6
Packets received on PIM-disabled Interface	0

MRIB Check:

RP/0/RSP0/CPU0:AG2-1#sh mrib vrf MLDP-14 route 10.1.0.2 232.11.11.11 detail

Fri Mar 12 03:05:38.300 EST

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

(10.1.0.2,232.11.11.11) Ver: 0x7023 RPF nbr: 10.44.0.10 Flags: RPF RPFID,

PD: Slotmask: 0x4 >>>>Slot Mask denotes which slot it will be flooded (binary conversion : 0100) Slot 2

MGID: 5489 >> Note MGID

Up: 15:48:34

RPF-ID: 639, Encap-ID: 0 >>Note RFID & ENCAP-ID. EnCAP ID 0 is Default (Partition MDT)

Incoming Interface List

LmdtMLDP-14 Flags: A LMI, Up: 00:36:35

Outgoing Interface List

Bundle-Ether5000.1 (0/0/CPU0) Flags: F NS LI, Up: 15:48:34 >>Match this slot

Bundle-Ether5000.2 (0/0/CPU0) Flags: F NS LI, Up: 15:48:26

RP/0/RSP0/CPU0:AG2-1#

Check bundle member interface active:

RP/0/RSP0/CPU0:AG2-1#sh int bundle-ether 5000 | in Active

Fri Mar 12 03:09:52.186 EST

TenGigE0/0/0/17 Full-duplex 10000Mb/s Active >>> It belong to LC0 , NP5, FIA 2

RP/0/RSP0/CPU0:AG2-1#

Check NP and FIA mapping of TenGigE0/0/0/17:

RP/0/RSP0/CPU0:AG2-1#sh controllers np ports all location 0/0/CPU0

Fri Mar 12 03:11:53.080 EST

Node: 0/0/CPU0:

-----

NP Bridge Fia Ports

-- -----

0 -- 0 TenGigE0/0/0/0 - TenGigE0/0/0/2  
1 -- 0 TenGigE0/0/0/3 - TenGigE0/0/0/5  
2 -- 1 TenGigE0/0/0/6 - TenGigE0/0/0/8  
3 -- 1 TenGigE0/0/0/9 - TenGigE0/0/0/11  
4 -- 2 TenGigE0/0/0/12 - TenGigE0/0/0/14  
5 -- 2 TenGigE0/0/0/15 - TenGigE0/0/0/17  
6 -- 3 TenGigE0/0/0/18 - TenGigE0/0/0/20  
7 -- 3 TenGigE0/0/0/21 - TenGigE0/0/0/23

BGP mVPN check:

RP/0/RSP0/CPU0:AG2-1#sh bgp ipv4 mvpn vrf MLDP-14

Fri Mar 12 03:14:31.903 EST

BGP router identifier 10.44.0.5, local AS number 55836

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 13952

BGP NSR Initial initsync version 207 (Reached)

BGP NSR/ISSU Sync-Group versions 13952/0

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: 1000:1001 (default for vrf MLDP-14)

\*> [1][10.44.0.5]/40 0.0.0.0 0 i

\*>i[1][10.44.0.10]/40 10.44.0.10 100 0 i

\*>i[1][10.44.0.110]/40

10.44.0.110 100 0 i

\*> [3][0][0.0.0.0][0][0.0.0.0][10.44.0.5]/120

0.0.0.0 0 i

\*>i[3][0][0.0.0.0][0][0.0.0.0][10.44.0.10]/120

10.44.0.10 100 0 i

\*>i[3][0][0.0.0.0][0][0.0.0.0][10.44.0.110]/120

10.44.0.110 100 0 i

\*>i[3][32][10.1.0.2][32][232.11.11.11][10.44.0.10]/120 >>> Type 3 is received from Source PE

10.44.0.10 100 0 i

\*> [7][1000:1001][55836][32][10.1.0.2][32][232.11.11.11]/184 >>> Type 7 is advertised from receiver PE

0.0.0.0 0 i

Processed 8 prefixes, 8 paths

LC mFIB Programming Check:

<#root>

RP/0/RSP0/CPU0:AG2-1#sh mfib vrf MLDP-14 platform route 10.1.0.2 232.11.11.11 location 0/0/CPU0

Fri Mar 12 03:19:41.712 EST

LC Type: A9K-24x10GE-SE

-----

Legend:

Common SW:

Route Information

- B: BAACL check S: RPF Interface signal
- DC: Directly connected PL: Punt to LC CPU
- PR: Punt to RP PF: Punt if forwarded
- DR: Drop all RI: RPF interface
- MF: MoFRR enabled TR: Tunnel route
- TE: Tunnel encap CD: Conditional decap
- MI: MVET index NI: Not Installed
- BD: Bidir

MVPN Information (Common SW)

- MLI: MVPN table index VTID: VRF Table ID
- RPFID: RPF Identifier

MoFRR Information

- A: Active RPF interface RI: RPF interface (or RPF ID)
- WDI: Watchdog counter index

-----

Source: 10.1.0.2 Group: 232.11.11.11 Mask: 64 RPF ID : 639

>> Verify RFID

Common SW Information

-----

MGID: 5489 MLI: 2132 Fabric Slotmask: 0x4 FGID: 0x4

>>> Verify MGID & Slotmask

-----

B S DC PL PR PF DR BD RI MF TR TE CD MI

-----  
F F F F F F F F 639 F F F F 0x0  
-----

MVPN/MLDP Decap Information:

RPF ID: 639



RPF ID MGID: Primary:7737 Backup:7738

Last refcount bitmask (RPF ID): 0x20, Sets: 21, Clears: 20

MNP Olist: [NP0:0 NP1:0 NP2:0 NP3:0 NP4:0 NP5:2 NP6:0 NP7:0 ]

>>> Verify NP as mentioned above it is NP5

Turnaround flag:

Check forwarding rate in bytes and packets for specific S, G on platform/PD:

RP/0/RSP0/CPU0:AG2-1#sh mfib vrf MLDP-14 route rate detail

Fri Mar 12 03:17:06.998 EST

IP Multicast Forwarding Rates

(Source Address, Group Address)

Incoming rate: (Incoming interface)

Node: (Incoming node) : pps/bps

HW Incoming count : (in packets)

HW Drop count : (in packets)

Outgoing rate:

Node: (Outgoing node) : pps/bps

HW Forwarding count: (in packets)

HW Drop count: (in packets)

Interfaces: (Outgoing interface list)

(10.1.0.2,232.11.11.11) >>>> Forwarding rate for specific S,G

Incoming rate : MDT Lm14

Node : 0/0/CPU0 : 200565 / 766963242 >>>> Incoming rate bps/pps

HW Incoming count : 3233699818 packets >>> Incoming rate in packets

HW Drop count : 0 packets >>> Incoming drops

Outgoing rate :

Node : 0/0/CPU0 : 99013 / 378628394 >>>> Outgoing rate bps/pps

HW Forwarding count: 1616848961 packets >>>> Outgoing rate in packets

HW Drop count: 0 packets >>>> Outgoing drops

Interfaces: BE5000.1 BE5000.2

PD statistics for S, G:

RP/0/RSP0/CPU0:AG2-1#sh mfib vrf MLDP-14 platform route statistics 10.1.0.2 232.11.11.11 location 0/0/CP

Fri Mar 12 03:21:57.222 EST

LC Type: A9K-24x10GE-SE

-----  
Legend:

N: NP ID R: Received

F: Forwarded P: Punted to CPU

ID: Ingress Drop ED: Egress Drop

Source: 10.1.0.2 Group: 232.11.11.11 Mask:64

-----  
NP R(packets:bytes)/F(packets:bytes)/P(packets)/ID(packets)/ED(packets)  
-----

0 1645909137:786744567486 / 0:0 / 0 / 0 / 1645909136

1 0:0 / 0:0 / 0 / 0 / 0

2 0:0 / 0:0 / 0 / 0 / 0

3 0:0 / 0:0 / 0 / 0 / 0

4 0:0 / 0:0 / 0 / 0 / 0

5 1645819072:786701516416 / 1645914052:786746916856 / 0 / 0 / 0 >> Forwarding rate on NP5 in bytes & pa

6 0:0 / 0:0 / 0 / 0 / 0

7 0:0 / 0:0 / 0 / 0 / 0  
-----

Interface Statistics:

-----  
C Interface F/P/D (packets:bytes)  
-----

5 BE5000.1 1645819072:786701516416 / 0:0 / 0:0 >> Forwarding rate on interface

5 BE5000.2 1645914051:786746916378 / 0:0 / 0:0  
-----

Check forwarding rate in bps and pps for specific S, G on platform/PD:

```
RP/0/RSP0/CPU0:AG2-1#sh mfib vrf MLDP-14 route stati 10.1.0.2 232.11.11.11 location 0/0/CPU0
```

```
Fri Mar 12 03:25:44.237 EST
```

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, RPFID - RPF ID Set

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

SW/HW Forwarding/Replication Counts: Packets in/Packets out/Bytes out

SW Failure Counts: RPF / TTL / Empty Olist / Encap RL / Other

HW Drop Counts: Ingress / Egress

HW Forwarding Rates: bps In/pps In/bps Out/pps Out

(10.1.0.2,232.11.11.11), Flags:

Up: 16:08:40

Last Used: never

SW Forwarding Counts: 0/0/0

SW Replication Counts: 0/0/0

SW Failure Counts: 0/0/0/0/0

HW Forwarding Counts: 3337154990/1668576693/797579659254

HW Replication Counts: 3337154990/3337166638/1595165652964

HW Drop Counts: 0/0 >>>> DROP counter if it is incrementing then it is a problem (Check np DROP counter)

HW Forwarding Rates: 766982645/200570/388346650/101555 -à It is in current bps & pps on LC (it should be)

LmdtMLDP-14 Flags: A LMI, Up:00:56:41 >>> MDT uptime

Bundle-Ether5000.1 Flags: NS EG, Up:16:08:40 >>>> uptime in PD

Bundle-Ether5000.2 Flags: NS EG, Up:16:08:32

RP/0/RSP0/CPU0:AG2-1#

Use MGID Index (noted previously) to check the clients and mapping of respective FIA and XBAR:

RP/0/RSP0/CPU0:AG2-1#sh controllers mgidprgm mgidindex 5489 location 0/0/CPU0

Fri Mar 12 03:28:22.682 EST

Device MGID-Bits Client-Last-Modified

=====

XBAR-0 100 MFIBV4 >> X-BAR instance 0

FIA-0 0 MFIBV4

FIA-1 0 MFIBV4

FIA-2 10 MFIBV4 >> Ten-0/0/0/17 is part of FIA 2

FIA-3 0 MFIBV4

=====

Client Mask

=====

MFIBV4 0x20

MFIBV6 0x0

L2FIB 0x0

sRP-pseudo-mc 0x0

UT 0x0

Prgm-Svr 0x0

P2MP 0x0

xbar 0x0

UT1 0x0

UT2 0x0

punt\_lib 0x0 punt\_lib 0x0

Platform programming for S, G and OLIST state:

<#root>

RP/0/RSP0/CPU0:AG2-1#sh mfib vrf MLDP-14 platform route olist detail 10.1.0.2 232.11.11.11 location 0/0

Fri Mar 12 03:31:17.067 EST

LC Type: A9K-24x10GE-SE

-----

Legend:

Common SW:

Route Information

NP: NP ID B: BAACL check

S: RPF Interface signal DC: Directly connected

PL: Punt to LC CPU PR: Punt to RP

PF: Punt if forwarded DR: Drop all

RI: RPF interface FS: Fabric slotmask

G: Multicast group ID M: Multicast Leaf Index

T: Table ID for lookup OC: Count of OLIST members

Base: Base of the statistics pointer NI: Not Installed

BD: Bidir

Interface Information

NP: NP ID Intf: Interface

U: uIDB index OT: OLE Type

T: Table ID IC: HW IC flag

B: HW BAACL bit EU: Interface uIDB index

IB: Bundle interface EH: In HW OLIST table

OIDX: OLIST index on NP PT: Punt table entry

Base: Statistics Ptr base RM: Remote FGID (Pri/Back)

BD\_ID: Bridge Domain ID if BVI RH: Raw hash

Software OLIST Information

SW OC: Software OLIST counts HW OC: Hardware OLIST counts

T: Table ID SD: Send direct flag

-----

Engine: Olist Information:

NP: NP ID Intf: Interface

U: uIDB index OT: OLE Type

T: Table ID IC: HW IC flag

B: HW BAACL bit EU: Interface uIDB index

IB: Bundle interface EH: In HW OLIST table

OIDX: OLIST index on NP PT: Punt table entry

Base: Statistics Ptr base RM: Remote FGID (Pri/Back)

BD\_ID: Bridge Domain ID if BVI RH: Raw hash

-----  
Source: 10.1.0.2 Group: 232.11.11.11 Mask: 64 RPF ID : 639

Common SW Information

-----

-----  
NP B S DC PL PR PF DR BD RI FS G M T OC Base

-----  
0 F F F F F F F F 639 0x4 5489 2132 0 0 0x12c88

1 F F F F F F F F 639 0x4 5489 2132 0 0 0x12c88

2 F F F F F F F F 639 0x4 5489 2132 0 0 0x12c88

3 F F F F F F F F 639 0x4 5489 2132 0 0 0x12c88

4 F F F F F F F F 639 0x4 5489 2132 0 0 0x12cb0

5 F F F F F F F F 639 0x4 5489 2132 0 2 0x12cb0

6 F F F F F F F F 639 0x4 5489 2132 0 0 0x12cc8

7 F F F F F F F F 639 0x4 5489 2132 0 0 0x12cb8

-----  
Interface Information

-----

NP Intf OT U T IC B EU IB EH OIDX PT Base RH

-----

5 BE5000.1 REG 931 0 F F 931 T T/T 0 T 0x12cb8 0xef8edd84

>> verify UIDB 931

5 BE5000.2 REG 930 0 F F 930 T T/T 1 T 0x12cbb 0xef8edd84

-----

Software OLIST Information

-----

NP SW OC HW OC T SD

-----

5 2 2 0 F

-----

Virtual Interface Local Presence

-----

NP Intf UL Intf Bundle Parent

-----

5 Bundle-Ether5000.1 Te0/0/0/17 Bundle-Ether5000

>> Check the member interface this S,G is forwarding

5 Bundle-Ether5000.2 Te0/0/0/17 Bundle-Ether5000

-----

Engine Information

-----

-----

NP Intf OT U T IC B VRF Base RH

-----

5 0xa03e0200 REG 41731 0 F F 0xb82c0100 0x84dd8eef

5 0xe03e0200 REG 41475 0 F F 63747 0xbb2c0100 0x84dd8eef

-----

Verify UIDB programming:

RP/0/RSP0/CPU0:AG2-1#sh uidb index location 0/0/CPU0 | in 931

Fri Mar 12 03:41:21.876 EST

0/0/CPU0 0x00023ea0 Bundle-Ether5000.1 Sub-interface 931

Check Bundle Hashing for S, G:

RP/0/RSP0/CPU0:AG2-1#sh mrib platform bundle hash ip 232.11.11.11 10.1.0.2 bundle-ether5000

Fri Mar 12 03:31:19.067 EST

=====

(Source, Group): (10.1.0.2,232.11.11.11)

Bundle interface: Bundle-Ether500 (0x20002a0)

Bundle member: TenGigE0/0/0/17 (0x580)

Raw Hash (32b): 0x8febeed6

Lag Hash (12b): 0x00d6

Member index: 0

=====

RP/0/RSP0/CPU0:AG2-1#sh mrib platform bundle hash ip 232.11.21.11 10.1.0.2 bundle-ether5000

Fri Mar 12 03:31:21.067 EST

=====

(Source, Group): (10.1.0.2,232.11.11.11)

Bundle interface: Bundle-Ether5000 (0x20002a0)

Bundle member: TenGigE0/1/0/17 (0x5c0)

Raw Hash (32b): 0x8014609b

Lag Hash (12b): 0x009b

Member index: 1

Check BCDL client status:

RP/0/RSP0/CPU0:AG2-1#sh mrib ipv4 client

Tue Aug 31 06:35:37.639 EDT

IP MRIB client-connections

bcdl\_agent:1 (client id 28)

igmp:7467 (client id 23)

ldp:17503 (client id 52)

lmrib\_bcdl:1 (client id 30) >>> Expected

lmrib\_bcdl:2 (client id 60) >>>> Not expected if more than 1 >> fix provided Cisco bug ID [CSCvy5896](#)

pim:7471 (client id 26)

pim6\_lbl:7644 (client id 20)

pim\_lbl:7471 (client id 17)

xtc:7458 (client id 33)



## Towards Core

Verify MLDP status:

```
RP/0/RSP0/CPU0:AG2-1#sh mpls mldp status
```

```
Fri Mar 12 02:56:00.727 EST
```

```
mLDP statistics
```

```
Process status : Active, Running and Ready
```

```
Logging notifications : Disabled
```

```
Label release scan in : never
```

```
LSM ID scan in : never
```

```
RIB connection status : Connected
```

```
RIB connection open : Yes
```

Check the MLDP database with Global-id:

The behaviour of recursive-fec.

Here for Tail/leaf end to reach root/Head end, there is next-hop change.

On Tail End, Head end loopback is learned via AG3-2 (MID Router). So here recursive-fec plays an important role

Configuration:

```
<#root>
```

```
mldp
```

```
address-family ipv4
```

```
recursive-fec
```

```
!
```

- Enable mLDP recursive fec makea Leaf (Ingress PE) to build Partitioned MDT towards the upstream router and insert the inline-RR (AG3-2) (that is, next-hop for Ingress PE/Original Root prefix) as the root.

```
<#root>
```

```
RP/0/RSP0/CPU0:AG2-1#sh mpls mldp database root 10.44.0.10 opaquetype global-id 4 details
```

Fri Mar 12 03:42:20.594 EST

mLDP database

LSM-ID: 0x000D6 (LSP-ID: 0x8005D) Type: P2MP Uptime: 01:13:02

FEC Root : 10.44.0.10 >>>>

**This is the FEC root where Multicast Source is connected**

FEC Length : 12 bytes

FEC Value internal : 0201000400000004312C000A

Opaque length : 4 bytes

Opaque value : 01 0004 00000004

Opaque decoded : [global-id 4]

Features : MBB RFEC RFWD Trace

Upstream neighbor(s) :

Recursive encode LSM-ID: 0x000D7 >>>

**10.44.0.10 is reachable via 10.44.0.4 so Recursive LSM check the next command**

Downstream client(s):

PIM MDT Uptime: 01:13:02

Egress intf : LmdtMLDP-14 >>>>

**Pointing towards mdt tunnel VRF**

Table ID : IPv4: 0xe00003f9 IPv6: 0xe08003f9

RPF ID : 639 >>>> Check the RPFID it should match with mfib & mrrib o/p

RD : 1000:1001 >> Check RD value

Verify the LSM ID state in the MLDP database:

<#root>

RP/0/RSP0/CPU0:AG2-1#sh mpls mldp database 0x000D7 details >>>>

**Recursive LSM ID (Outer FEC)**

Fri Mar 12 03:44:54.350 EST

mLDP database

LSM-ID: 0x000D7 (LSP-ID: 0x8005D) Type: P2MP Uptime: 01:15:36

FEC Root : 10.44.0.4

FEC Length : 25 bytes

FEC Value internal : 020700110600010312C000A000701000400000004312C0004

Opaque length : 17 bytes

Opaque value : 07 0011 0600010312C000A000701000400000004

Opaque decoded : [recursive] 10.44.0.10:[global-id 4] >>>>

inner fec which is pointing towards actual source 10.44.0.10

Features : MBB RFEC RFWD Trace

Upstream neighbor(s) :

Is CSI accepting : N

10.44.0.4:0 [Active] [MBB] Uptime: 01:15:36 >>>

upstream neighbor. Check ldp neighbor with hop

Local Label (D) : 29233 >>>> local label

Downstream client(s):

Recursive 0x000D6 Uptime: 01:15:36 >>>

This is outer fec which is pointing towards VRF interface actual receiver (check above)

PIM MDT Uptime: 01:15:36 (from 0x000D6) >>

MDT tunnel uptime

Egress intf : LmdtMLDP-14

Table ID : IPv4: 0xe00003f9 IPv6: 0xe08003f9

RPF ID : 639

RD : 1000:1001

Verify MRIB parameters for local labels for S, G:

<#root>

RP/0/RSP0/CPU0:AG2-1#sh mrib mpls forwarding labels 29233 detail

Fri Mar 12 03:47:45.907 EST

LSP information (mLDP) :

LSM-ID: 0x8005D, Role: Tail

>>>> Verify the LSM ID & RPF -ID with MLDP database

RPF-ID: 0x0027F, Assoc-TIDs: 0xe00003f9/0xe08003f9, MDT: LmdtMLDP-14

>> RFID is in hex convert into DEC & verify it

Incoming Label : 29233

Transported Protocol : <unknown>

Explicit Null : None

IP lookup : enabled

Platform information :

Slotmask: Primary: 0x4, Backup: 0x4 >> Verify slotmask

MGID Primary: 7751, Backup: 7752 >>> Verify the same with mfib platform route location command

Label Node Role: LMRIB\_ROUTER\_ROLE\_EGRESS, Tunnel Type: MVPN Tunnel

Outsegment Info #1 [T/Pop]:

No info.

PlatInfo:

Primary: [Main IF: NULL (IFH: 0), UL IF: NULL (IFH: 0)]

Backup: [Main IF: NULL (IFH: 0), UL IF: NULL (IFH: 0)]

Check the RPFID and respective MGID:

RP/0/RSP0/CPU0:AG2-1#show mrib rpf-id | b 639

Fri Mar 12 03:52:43.375 EST

Addr: 127.2.0.0, RPF-ID: 639 (0x27f)

Alloc: F, IPv6 Intr: F, Stale: F, Chkpt Obj-ID: 0x0

Routes - Non-core: 1, Core: 0, Label: 2

In PD Retry: F

Platform Annot: Primary MGID: 7737, Backup MGID: 7738

Verify MPLS forwarding:

```
RP/0/RSP0/CPU0:AG2-1#sh mpls forwarding labels 29233
```

```
Fri Mar 12 04:00:00.695 EST
```

```
Local Outgoing Prefix Outgoing Next Hop Bytes
```

```
Label Label or ID Interface Switched
```

```
-----  
29233 Unlabelled mLDP/IR: 0x8005d >> verify with MPLS MLDP database LSP -ID
```

```
RP/0/RSP0/CPU0:AG2-1#sh mpls forwarding labels 29233 location 0/0/CPU0
```

```
Fri Mar 12 04:00:05.962 EST
```

```
Local Outgoing Prefix Outgoing Next Hop Bytes
```

```
Label Label or ID Interface Switched
```

```
-----  
29233 Unlabelled mLDP/IR: 0x8005d 289 > For LSP card, MPLS forwarding stats is not incrementing fix pro
```

```
RP/0/RSP0/CPU0:AG2-1#
```

Check CEF forwarding:

```
<#root>
```

```
RP/0/RSP0/CPU0:AG2-1#sh cef 10.44.0.10
```

```
>> FEC Root IP
```

```
Fri Mar 12 05:48:01.838 EST
```

```
10.44.0.10/32, version 149189, internal 0x1000001 0x0 (ptr 0x7a11058) [1], 0x0 (0x79c18c08), 0xa08 (0x79c18c08)
```

```
Updated Mar 12 02:29:03.109
```

```
Prefix Len 32, traffic index 0, precedence n/a, priority 4
```

```
via
```

```
10.44.0.4/32
```

```
, 6 dependencies, recursive [flags 0x6000]
```

```
path-idx 0 NHID 0x0 [0x727a0730 0x0]
```

```
recursion-via-/32
```

```
next hop 10.44.0.4/32 via
```

24215/0/21

local label 29229

next hop 10.16.26.10/32 Te0/1/0/2 labels imposed {ImplNull 25509}

RP/0/RSP0/CPU0:AG2-1#sh cef 10.44.0.4 >> Next Hop for 10.44.0.10

Fri Mar 12 05:48:07.914 EST

10.44.0.4/32, version 117384, internal 0x1000001 0x0 (ptr 0x721ee2a8) [1], 0x0 (0x79c14ae4), 0xa20 (0x79c14ae4)

Updated Mar 11 20:41:39.060

remote adjacency to TenGigE0/1/0/2

Prefix Len 32, traffic index 0, precedence n/a, priority 1

via 10.16.26.0/32, Bundle-Ether45, 3 dependencies, weight 0, class 0, backup (Local-LFA) [flags 0x300]

path-idx 0 NHID 0x0 [0x79a002a0 0x79a00508]

next hop 10.16.26.0/32

local adjacency

local label 24215 labels imposed {ImplNull}

via 10.16.26.10/32, TenGigE0/1/0/2, 4 dependencies, weight 0, class 0, protected [flags 0x400]

path-idx 1 bkup-idx 0 NHID 0x0 [0x78eeb420 0x78ef0b50]

next hop 10.16.26.10/32

local label 24215 labels imposed {ImplNull}

RP/0/RSP0/CPU0:AG2-1#

Check the NP drops counters:

RP/0/RSP0/CPU0:ASR9006-L#sh controllers np counters np0 location 0/2/CPU0 | util egrep -i "DROP|OLE|MDF"

Fri Oct 30 03:31:18.619 UTC

16 MDF\_TX\_LC\_CPU 131216974 491

17 MDF\_TX\_WIRE 1233084179074 393586. >>> Packets towards egress PHY from NP

21 MDF\_TX\_FABRIC 1297085738439 297619 >>> Packets towards Fabric from NP

105 RSV\_DROP\_EGR\_UIDB\_DOWN 1460 0

117 RSV\_DROP\_MPLS\_NRLDI\_NO\_MATCH 413748715 0 >>

```

161 RSV_DROP_EGR_LAG_NO_MATCH 716 0
232 RSV_DROP_MPLS_LEAF_NO_MATCH 860077083 0
438 RSV_LSM_NO_OLE 2714 0. >>>>due to no OLE mapped or underlay Join for LSM ID
544 MDF_PIPE_LPBK 537269848146 297618
564 MDF_PUNT_POLICE_DROP 39231 0
573 MODIFY_PUNT_REASON_MISS_DROP 2 0
633 MDF_LB_TM0_LPBK0 144115523350 75159 >>> TM0 loopback utilization (Tomhawk/Typh case)
634 MDF_LB_TM0_LPBK1 799575201453 74405
635 MDF_LB_TM1_LPBK0 143959136311 74405 >>> TM1 loopback utilization (in Tomhawk case)
636 MDF_LB_TM1_LPBK1 143733432720 73661
637 MDF_TM_LP_Q_P1 6321879 10
639 MDF_TM_LP_Q_PN 1231376971969 297618
1945 PARSE_DROP_IN_UIDB_DOWN 696210 0
1948 PARSE_DROP_IPV4_DISABLED 2866 0
1979 PARSE_DROP_IPV6_DISABLED 269 0

```

Verify NP or FIA Drops on LC:

```

RP/0/RSP0/CPU0:B#sh drops all ongoing location 0/2/CPU0
=====

```

```

Checking for ongoing drops on 0/2/CPU0
=====

```

```

[np:NP0] RSV_DROP_IPV4_LEAF_NO_MATCH: +323
[np:NP0] RSV_DROP_MPLS_TXADJ_DROP: +393205
[np:NP0] RSV_EGR_LAG_NOT_LOCAL_DROP_CNT: +36
[np:NP0] MDF_RPF_FAIL_DROP: +24058109. >>>> Check use case TZ to understand this
[np:NP1] RSV_DROP_IPV4_LEAF_NO_MATCH: +246
[np:NP1] RSV_DROP_MPLS_TXADJ_DROP: +393216
[np:NP1] RSV_EGR_LAG_NOT_LOCAL_DROP_CNT: +304
[np:NP1] MDF_RPF_FAIL_DROP: +24116093

```

```

show controller np counters:

```

[np:NP0] RSV\_DROP\_IPV4\_LEAF\_NO\_MATCH: +29  
[np:NP0] RSV\_DROP\_MPLS\_TXADJ\_DROP: +16076  
[np:NP0] MDF\_RPF\_FAIL\_DROP: +979311  
[np:NP1] RSV\_DROP\_IPV4\_LEAF\_NO\_MATCH: +9  
[np:NP1] RSV\_DROP\_MPLS\_TXADJ\_DROP: +16080  
[np:NP1] RSV\_EGR\_LAG\_NOT\_LOCAL\_DROP\_CNT: +11  
[np:NP1] MDF\_RPF\_FAIL\_DROP: +984429  
[np:NP2] UNKNOWN\_L2\_ON\_L3\_DISCARD: +2

---

**Note:** MDF\_RPF\_FAIL\_DROP is due to RPF fail is when the RPF interface in the route does not match with the incoming intf of the packet.

---

RP/0/RSP0/CPU0:PUNEKLCLEAR001#sh drops all ongoing location 0/0/CPU0 | include "RSV\_DROP\_ACL\_DENY|TM"

Thu Dec 24 20:33:30.086 IST

[np:NP3] RSV\_DROP\_ACL\_DENY: +109

[np:NP5] RSV\_DROP\_ACL\_DENY: +154

[np:NP6] RSV\_DROP\_ACL\_DENY: +74

[np:NP7] RSV\_DROP\_ACL\_DENY: +29

[np:NP 3 TM 0] drop paks: +5909 >>> Check TM drops if it increasing rapidly check the QOS stats & RFD bu

[np:NP 5 TM 0] drop paks: +909

[np:NP 6 TM 0] drop paks: +8925

[np:NP 7 TM 0] drop paks: +5491

If you observe that TM drops increase more rapidly, then collect these logs:

Show policy-map interface <> out >>>> egress stats >>>>Check if there is any drops in queue

Collect np datalogs

Get into LC shell (run attach 0/x/cpu0)

Execute **show\_np** "e <np\_id> -d npdatalog

Also, collect multiple(4-5) iterations of the pipeline from the LC shell.



<#root>

```
show_np -e <np_id> -d -pipeline >>>>
```

Only applicable for Tomhawk LC

Since Multicast traffic uses TM loopback, check if there are any drops in qshal (example NP4 Location 4).

Sample o/p:

```
Tailend#show qshal loopback-queue np 4 location 0/0/CPU0
Tue Jun 9 10:51:12.908 IST
TY Options argc:6
nphal_show_chk -p 2049loopback-queue-n 0x4
Done
Loopback Queues : NP 4
=====
Multi-cast Loopback Queue NP 4
Ingress: QID 0x40 Entity: 4/0/0/4/8/0 Priority: Priority 1 Qdepth: 0
StatIDs: commit/fast_commit/drop: 0x5f0140/0x0/0x5f0141
Statistics(Pkts/Bytes):
Tx_To_TM 0/0
Total Xmt 0/0 Dropped 0/0 >>> Check Transmit & Drop packets in P1
Ingress: QID 0x41 Entity: 4/0/0/4/8/1 Priority: Priority 2 Qdepth: 0
StatIDs: commit/fast_commit/drop: 0x5f0145/0x0/0x5f0146
Statistics(Pkts/Bytes):
Tx_To_TM 0/0
Total Xmt 0/0 Dropped 0/0 >>> Check Transmit & Drop packets in P2
Ingress: QID 0x43 Entity: 4/0/0/4/8/3 Priority: Priority 3 Qdepth: 0
StatIDs: commit/fast_commit/drop: 0x5f014f/0x0/0x5f0150
Statistics(Pkts/Bytes):
Tx_To_TM 0/0
Total Xmt 0/0 Dropped 0/0 >>> Check Transmit & Drop packets in P3
Ingress: QID 0x42 Entity: 4/0/0/4/8/2 Priority: Priority Normal Qdepth: 0
StatIDs: commit/fast_commit/drop: 0x5f014a/0x0/0x5f014b
Statistics(Pkts/Bytes):
Tx_To_TM 6607491182505/9627383230841487
Total Xmt 349820789/509702036782 Dropped 6607141361716/9626873528804705 >>>> Check Transmit & Drop packets
```

To catch any bursty traffic:

```
np_perf -e 0 -R -o 4 -D 10 -n 5000
```

R is for rfd

O is for special ports (loopback comes under that)

D is duration; in the above example it is 10 secs; max you can specify is 2 days or something like that

## MID Router

### Towards Receivers

Check MLDP status:

```
RP/0/RSP1/CPU0:AG3-2#sh mpls mldp status
```

```
Fri Mar 12 08:08:37.366 UTC
```

```
mLDP statistics
```

```
Process status : Active, Running and Ready >>>> Active
```

```
Logging notifications : Disabled
```

```
Label release scan in : never
```

```
LSM ID scan in : never
```

```
RIB connection status : Connected
```

```
RIB connection open : Yes
```

```
RP/0/RSP1/CPU0:AG3-2#
```

Check the MLDP database:

Behaviour of recursive-fec

Here for the mid-end (inline RR) to reach the root/head-end, there is next-hop as root.

Configuration:

```
mldp
```

```
address-family ipv4
```

```
recursive-fec
```

```
!
```

Sample o/p:

```
<#root>
```

```
RP/0/RSP1/CPU0:AG3-2#sh mpls mldp database opaquetype recursive global-id 4
```

```
Fri Mar 12 07:54:05.599 UTC
```

```
mLDP database
```

```
LSM-ID: 0x00023 (LSP-ID: 0x80023) Type: P2MP Uptime: 01:26:35
```

```
FEC Root : 10.44.0.4 (we are the root) >>>
```

**In this router, outer FEC get remove**

Opaque decoded : [recursive] 10.44.0.10:[global-id 4]

Features : MBB RFWD

Upstream neighbor(s) : >>>>>>>>>>>>>>>>>>

**Upstream client towards HEADEND**

Recursive decode LSM-ID: 0x00024 >> Inner FEC to reach 10.44.0.10

Downstream client(s): >>>>>>

**Downstream client towards Receiver**

LDP 10.44.0.5:0 Uptime: 01:26:35

Rec Next Hop : 10.44.0.5 >> Receiver end IP

Remote label (D) : 29233 >>> Remote label to reach 10.44.0.5

<#root>

RP/0/RSP1/CPU0:AG3-2#sh mpls mldp database 0x00024 >>>

**Inner FEC (Actual LSM to reach Source)**

Fri Mar 12 07:55:20.346 UTC

mLDP database

LSM-ID: 0x00024 (LSP-ID: 0x80023) Type: P2MP Uptime: 01:27:50

FEC Root : 10.44.0.10

Opaque decoded : [global-id 4]

Features : MBB RFWD

Upstream neighbor(s) :

Is CSI accepting : N

10.44.0.10:0 [Active] [MBB] Uptime: 01:27:50

Local Label (D) : 25515

Downstream client(s):

Recursive 0x00023 Uptime: 01:27:50

Check multicast PI > mrIB entry/state for S, G with the help of the label:

RP/0/RSP1/CPU0:AG3-2#sh mrib mpls forwarding labels 25515 detail

Fri Mar 12 08:01:03.353 UTC

LSP information (mLDP) :

LSM-ID: 0x80023, Role: Mid >> Note LSM & Role

Incoming Label : 25515

Transported Protocol : <unknown>

Explicit Null : None

IP lookup : disabled

Platform information :

Slotmask: Primary: 0x4, Backup: 0xc

MGID Primary: 34318, Backup: 34319

Label Node Role: LMrib\_ROUTER\_ROLE\_MID, Tunnel Type: MVPN Tunnel

Outsegment Info #1 [M/Swap, Recursive]:

OutLabel: 29233, NH: 10.44.0.5, ID: 0x23, Sel IF: TenGigE0/0/0/3 >> verify out label with local label

UL IF: TenGigE0/0/0/3, Node-ID: 0x8220

Backup Tunnel: Un:0x0 Backup State: Ready, NH: 0.0.0.0, MP Label: 0

Backup Sel IF: TenGigE0/1/0/0/2, UL IF: TenGigE0/1/0/0/2, Node-ID: 0x8320

PlatInfo:

Primary: [Main IF: TenGigE0/0/0/3 (IFH: 0x4000180), UL IF: TenGigE0/0/0/3 (IFH: 0x4000180)] >> Egress

Backup: [Main IF: Bundle-Ether45 (IFH: 0x20002e0), UL IF: TenGigE0/1/0/0/2 (IFH: 0x60020c0)]

RP/0/RSP1/CPU0:AG3-2#

Check MGID with this command:

show mpls forwarding labels 25515 hardware egress location 0/0/CPU0

This MGID is in hex value and then converts to decimal.

Verify MGID:

RP/0/RSP1/CPU0:AG3-2#sh controllers mgidprgm mgidindex 34318 location 0/0/CPU0 >>> mgid value you got

Fri Mar 12 08:02:02.873 UTC

Device MGID-Bits Client-Last-Modified

=====

XBAR-0 1 P2MP

FIA-0 1 P2MP

FIA-1 0 None

=====

Client Mask

=====

MFIBV4 0x0

MFIBV6 0x0

L2FIB 0x0

sRP-pseudo-mc 0x0

UT 0x0

Prgm-Svr 0x0

P2MP 0x1

xbar 0x0

UT1 0x0

UT2 0x0

punt\_lib 0x0

RP/0/RSP1/CPU0:AG3-2#

Check forwarding rate on RSP and LC:

RP/0/RSP1/CPU0:AG3-2#sh mpls forwarding labels 25515 location 0/0/CPU0

Fri Mar 12 08:04:35.608 UTC

Local Outgoing Prefix Outgoing Next Hop Bytes

Label Label or ID Interface Switched

-----

25515 29233 mLDP/IR: 0x00023 10.44.0.5 279353664958 >> Forwarding rate should increase

RP/0/RSP1/CPU0:AG3-2#sh mpls forwarding labels 25515 location 0/0/CPU0

Fri Mar 12 08:04:38.798 UTC

Local Outgoing Prefix Outgoing Next Hop Bytes

Label Label or ID Interface Switched

-----

25515 29233 mLDP/IR: 0x00023 10.44.0.5 279353664958 >>> Forwarding rate should increase

Verify MPLS label programming related to S, G:

<#root>

RP/0/RSP1/CPU0:AG3-2#sh mpls forwarding labels 25515 detail location 0/0/CPU0

Fri Mar 12 08:04:56.793 UTC

Local Outgoing Prefix Outgoing Next Hop Bytes

Label Label or ID Interface Switched

-----

25515 mLDP/IR: 0x00023 (0x80023) >>> Verify LSM/LSP ID as per MLDP database

Updated Mar 12 06:27:29.419

mLDP/IR LSM-ID: 0x80023, MDT: 0x0

Flags:IP Lookup:not-set, Expnullv4:not-set, Expnullv6:not-set

Payload Type v4:not-set, Payload Type v6:not-set, l2vpn:not-set

Head:not-set, Tail:not-set, Bud:not-set, Peek:not-set, inclusive:not-set

Ingress Drop:not-set, Egress Drop:not-set

RPF-ID:0, Encap-ID:0

Disp-Tun:[ifh:0x0, label:-]

Platform Data [72]:

{ 0 0 0 4 0 0 0 4

0 0 0 0 0 0 0 0

0 0 0 0 0 0 0 0

0 0 0 12 0 0 0 12

0 0 0 0 0 0 0 0

0 0 0 0 0 0 0 0

0 0 134 14 0 0 134 15

0 0 134 14 0 0 134 15

```
0 0 0 1 0 0 0 2
}
mpls paths: 1, local mpls paths: 1, protected mpls paths: 1
29233 mLDP/IR: 0x00023 (0x80023) \
>>> verify LSM ID

10.44.0.5 280073746930
Updated: Mar 12 06:27:29.420
My Nodeid:0x8220
Interface Nodeids:
[ 0x8220 - - - - - ]
Interface Handles:
[ 0x4000180 - - - - - ]
Backup Interface Nodeids:
[ 0x8320 - - - - - ]
Backup Interface Handles:
[ 0x60020c0 - - - - - ]
via-label:25502, mpi-flags:0x200 tos_masks:[ primary:0x0 backup:0x0]
MPI resolution time Mar 12 06:27:29.420
Packets Switched: 581065865
```

Check BGP mVPN:

You can see Type 3 from the Source and Type 7 from the receiver.

```
RP/0/RSP1/CPU0:AG3-2#sh bgp ipv4 mvpn rd 1000:1001
Fri Mar 12 08:08:09.413 UTC
BGP router identifier 10.44.0.4, local AS number 55836
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 11597
```

```
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 11597/0
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: 1000:1001
*>i[1][10.44.0.5]/40 10.44.0.5 100 0 i
*>i[1][10.44.0.10]/40 10.44.0.10 100 0 i
*>i[1][10.44.0.110]/40
10.44.0.110 100 0 i
*>i[3][0][0.0.0.0][0][0.0.0.0][10.44.0.5]/120
10.44.0.5 100 0 i
*>i[3][0][0.0.0.0][0][0.0.0.0][10.44.0.10]/120
10.44.0.10 100 0 i
*>i[3][0][0.0.0.0][0][0.0.0.0][10.44.0.110]/120
10.44.0.110 100 0 i
*>i[3][32][10.1.0.2][32][232.11.11.11][10.44.0.10]/120
10.44.0.10 100 0 i
*>i[7][1000:1001][55836][32][10.1.0.2][32][232.11.11.11]/184
10.44.0.5 100 0 i
Processed 8 prefixes, 8 paths
```

## Head End Router (Multicast Source PE)

### Towards Core

Check MPLS MLDP Database:

```
RP/0/RSP0/CPU0:DC1-DCWAN#sh mpls mldp database root 10.44.0.10 opaquetype global-id 4
```

Fri Mar 12 05:27:57.124 EST



mLDP database

LSM-ID: 0x00009 Type: P2MP Uptime: 03:05:01 >> Verify LSM ID

FEC Root : 10.44.0.10 (we are the root)

Opaque decoded : [global-id 4]

Features : MBB RFWD

Upstream neighbor(s) : >>> Upstream is under VRF so it is none

None

Downstream client(s):

LDP 10.44.0.4:0 Uptime: 03:05:01 >>> Downstream client 10.44.0.4

Rec Next Hop : 10.44.0.4

Remote label (D) : 25515 >>> Verify this label with Mid router local label

PIM MDT Uptime: 03:05:01

Egress intf : LmdtMLDP-14

Table ID : IPv4: 0xe00003f8 IPv6: 0xe08003f8

HLI : 0x00009

Ingress : Yes

Local Label : 33031 (internal)

RP/0/RSP0/CPU0:DC1-DCWAN#

Check the egress core interface with the use of mrrib:

RP/0/RSP0/CPU0:DC1-DCWAN#sh mrrib mpls forwarding labels 33031 detail

Fri Mar 12 05:29:57.603 EST

LSP information (mLDP) :

LSM-ID: 0x00009, Role: Head, Head LSM-ID: 0x00009

Incoming Label : (33031) >>> Local Label

Transported Protocol : <unknown>

Explicit Null : None

IP lookup : disabled

Platform information :

Slotmask: Primary: 0x1, Backup: 0x1>>>>>> It gives egress LC0 as slotmask

MGID Primary: 26096, Backup: 26097 >>> Don't consider this MGID to verify the egress forwarding

Label Node Role: LMRIB\_ROUTER\_ROLE\_INGRESS, Tunnel Type: MVPN Tunnel

Outsegment Info #1 [H/Push, Recursive]:

OutLabel: 25515, NH: 10.44.0.4, ID: 0x9, Sel IF: TenGigE0/0/1/2 >> Hash to one bundle member interface

UL IF: TenGigE0/0/1/2, Node-ID: 0x801

PlatInfo:

Primary: [Main IF: Bundle-Ether410 (IFH: 0x20009e0), UL IF: TenGigE0/0/1/2 (IFH: 0x140)]

Backup: [Main IF: NULL (IFH: 0x0), UL IF: NULL (IFH: 0x0)]

Alternate command:

```
RP/0/RSP0/CPU0:DC1-DCWAN#sh mrib mpls forwarding head-lsm-id 0x00009 detail
```

Fri Mar 12 05:44:46.639 EST

LSP information (mLDP) :

LSM-ID: 0x00009, Role: Head, Head LSM-ID: 0x00009

Incoming Label : (33031)

Transported Protocol : <unknown>

Explicit Null : None

IP lookup : disabled

Platform information :

Slotmask: Primary: 0x1, Backup: 0x1

MGID Primary: 26096, Backup: 26097

Label Node Role: LMRIB\_ROUTER\_ROLE\_INGRESS, Tunnel Type: MVPN Tunnel

Outsegment Info #1 [H/Push, Recursive]:

OutLabel: 25515, NH: 10.44.0.4, ID: 0x9, Sel IF: TenGigE0/0/1/2

UL IF: TenGigE0/0/1/2, Node-ID: 0x801

PlatInfo:

Primary: [Main IF: Bundle-Ether410 (IFH: 0x20009e0), UL IF: TenGigE0/0/1/2 (IFH: 0x140)]

Backup: [Main IF: NULL (IFH: 0x0), UL IF: NULL (IFH: 0x0)]

Check mpls forwarding rate on RSP and LC - egress from DCWAN is MPLS forwarding:

```
RP/0/RSP0/CPU0:DC1-DCWAN#sh mpls forwarding labels 33031
```

```
Fri Mar 12 05:32:09.026 EST
```

```
Local Outgoing Prefix Outgoing Next Hop Bytes
```

```
Label Label or ID Interface Switched
```

```
-----  
33031 25515 mLDP/IR: 0x00009 10.44.0.4 545528292634 >> Verify forwarding rate & LSM
```

```
RP/0/RSP0/CPU0:DC1-DCWAN#sh mpls forwarding labels 33031 detail location 0/0/CPU0
```

```
Fri Mar 12 05:32:22.904 EST
```

```
Local Outgoing Prefix Outgoing Next Hop Bytes
```

```
Label Label or ID Interface Switched
```

```
-----  
33031 mLDP/IR: 0x00009 (0x00009)
```

```
Updated Mar 12 02:22:56.152
```

```
mLDP/IR LSM-ID: 0x00009, MDT: 0x0, Head LSM-ID: 0x00009
```

```
IPv4 Tableid: 0xe0000000, IPv6 Tableid: 0xe0800000
```

```
Flags:IP Lookup:not-set, Expnulv4:not-set, Expnulv6:not-set
```

```
Payload Type v4:not-set, Payload Type v6:not-set, l2vpn:not-set
```

```
Head:set, Tail:not-set, Bud:not-set, Peek:not-set, inclusive:not-set
```

```
Ingress Drop:not-set, Egress Drop:not-set
```

```
RPF-ID:0, Encap-ID:0
```

```
Disp-Tun:[ifh:0x0, label:-]
```

```
Platform Data [64]:
```

```
{ 0 0 0 1 0 0 0 1
```

```
0 0 0 0 0 0 0 0
```

```
0 0 0 0 0 0 0 0
```

```
0 0 0 1 0 0 0 1
```

```
0 0 0 0 0 0 0 0
```

```
0 0 0 0 0 0 0 0
```

0 0 101 240 0 0 101 241

0 0 0 1 0 0 0 1

}

mpls paths: 1, local mpls paths: 1, protected mpls paths: 0

25515 mLDP/IR: 0x00009 (0x00009) \

10.44.0.4 546248376534

Updated: Mar 12 02:22:56.150

My Nodeid:0x801

Interface Nodeids:

[ 0x801 - - - - - ]

Interface Handles:

[ 0x140 - - - - - ]

Backup Interface Nodeids:

[ - - - - - ]

Backup Interface Handles:

[ - - - - - ]

via-label:33009, mpi-flags:0x200 tos\_masks:[ primary:0x0 backup:0x0]

MPI resolution time Mar 12 02:22:56.151

Packets Switched: 1133295387 >>> verify Traffic rate increase

RP/0/RSP0/CPU0:DC1-DCWAN#

Check CEF for next-hop shown in MLDP downstream:

RP/0/RSP0/CPU0:DC1-DCWAN#sh cef 10.44.0.4 internal

Fri Mar 12 05:39:40.543 EST

IPv4:default:0xe0000000[0x718462d8]:lsd:10.44.0.4/32[ref:1 proto:ipv4 flags:0x1000001 f2:0x0 src:lsd ver

LSPA => [ref:3 count:1 src:lsd flags:[pd create pending, updated, has-pd-li] walk-idx:0 flags:0xa20][0x

0={

LABEL\_INFO => [default [o-label:pop l-label:33009 type:0 flags:[linked to lsw, ipv4, recycle, retain] e

LABEL\_INFO => [default [o-label:pop l-label:33009 type:0 flags:[owner locked, eos li, linked to lsw, ip

LDI => [ref:9 pl:75fe41a8 proto:mpls type:imp-eos1 lvl:1 buckets:1 slots:1 fixup:0 flags:[owner locked,

slots:

```
{0,p0,s0,f:[]}=
```

```
NHINFO => [type:tx, linkt:link_IPv4, nh_proto:IPv4, ifh:Bundle-Ether410(0x20009e0), nh:10.16.74.0, adj
```

```
IDB => [ref:1 ifh:Bundle-Ether410(0x20009e0) type:28/EtherBundle parent-type:0/none flags:[owner locked
```

```
IDB_EXT:{MPLS IPv4:[ref:1 MPLS:[0x71a6c058] NH:2 MTU:1500 flags:[forwarding is enabled][0x75659058]],IF
```

```
}
```

buckets:

```
{0,p0,s0}={[0x73c97890]}
```

```
LDI_LW => [type:imp loadinfo:[0x72811554] flags: timestamp:Mar 12 02:22:40.904][0x8aa0a3ac]
```

```
PATHLIST => [ref:12 count:1 depth:1 ldi:72811554 type:Shared flags:[mpls ext, imp, lsw] extn_flags:[] t
```

```
0={
```

```
[nh:10.16.74.0 ifh:Bundle-Ether410(0x20009e0) tbl:0xe0000000 type:implicit-null flags:[none]]
```

```
[depth:1 flags:[resolved,ldi-preferred] resolves-via: ###
```

```
NHINFO => [type:tx, linkt:link_IPv4, nh_proto:IPv4, ifh:Bundle-Ether410(0x20009e0), nh:10.16.74.0, adj
```

```
IDB => [ref:1 ifh:Bundle-Ether410(0x20009e0) type:28/EtherBundle parent-type:0/none flags:[owner locked
```

```
IDB_EXT:{MPLS IPv4:[ref:1 MPLS:[0x71a6c058] NH:2 MTU:1500 flags:[forwarding is enabled][0x75659058]],IF
```

```
NHINFO => [type:tx, linkt:link_MPLS, nh_proto:IPv4, ifh:Bundle-Ether410(0x20009e0), nh:10.16.74.0, adj
```

```
IDB => [ref:1 ifh:Bundle-Ether410(0x20009e0) type:28/EtherBundle parent-type:0/none flags:[owner locked
```

```
IDB_EXT:{MPLS IPv4:[ref:1 MPLS:[0x71a6c058] NH:2 MTU:1500 flags:[forwarding is enabled][0x75659058]],IF
```

```
RP/0/RSP0/CPU0:DC1-DCWAN#
```

Verify Bundle member interfaces:

```
RP/0/RSP0/CPU0:DC1-DCWAN#sh int be310 | in Active
```

```
Wed Mar 10 06:37:16.756 EST
```

```
TenGigE0/0/0/0 Full-duplex 10000Mb/s Active
```

```
TenGigE0/0/0/1 Full-duplex 10000Mb/s Active
```

```
TenGigE0/0/0/4 Full-duplex 10000Mb/s Active
```

```
TenGigE0/0/0/5 Full-duplex 10000Mb/s Active
```

Check BCDL client status:

RP/0/RSP0/CPU0:AG2-1#sh mrib ipv4 client

Tue Aug 31 06:35:37.639 EDT

IP MRIB client-connections

bcdl\_agent:1 (client id 28)

igmp:7467 (client id 23)

ldp:17503 (client id 52)

lmrib\_bcdl:1 (client id 30) >>> Expected

lmrib\_bcdl:2 (client id 60) >>>> Not expected if more than 1 >> fix provided Cisco bug ID [CSCvy5890](#)

pim:7471 (client id 26)

pim6\_lbl:7644 (client id 20)

pim\_lbl:7471 (client id 17)

xtc:7458 (client id 33)

Check mFIB:

RP/0/RSP0/CPU0:DC1-DCWAN#sh mrib vrf MLDP-14 platform route 10.1.0.2 232.11.11.11/32 location 0/0/CPU0

Fri Mar 12 09:39:34.651 EST

LC Type: A9K-MOD400-SE

-----  
Legend:

Common SW:

Route Information

B: BAACL check S: RPF Interface signal

DC: Directly connected PL: Punt to LC CPU

PR: Punt to RP PF: Punt if forwarded

DR: Drop all RI: RPF interface

MF: MoFRR enabled TR: Tunnel route

TE: Tunnel encap CD: Conditional decap

MI: MVET index NI: Not Installed

BD: Bidir

MVPN Information (Common SW)

MLI: MVPN table index VTID: VRF Table ID

RPFID: RPF Identifier

MoFRR Information

A: Active RPF interface RI: RPF interface (or RPF ID)

WDI: Watchdog counter index

-----

Source: 10.1.0.2 Group: 232.11.11.11 Mask: 64 RPF Int: BE1000.2001

Common SW Information

-----

MGID: 26081 MLI: 75537 Fabric Slotmask: 0x0 FGID: 0x0 >> Here no slotmask programmed as traffic is ing

-----

B S DC PL PR PF DR BD RI MF TR TE CD MI

-----

F F F F F F F F 0x2013120 F F T F 0x2

-----

MVPN/MLDP Encap Information:

Encap flag: T, Encap ID: 2 -->>>> Note this Encap ID

Last refcount bitmask (Enc ID): 0x0, Sets: 0, Clears: 0

Local receiver flag: F

Turnaround flag: F

Verify Encap ID:

RP/0/RSP0/CPU0:DC1-DCWAN#sh mfib encap identifier

Fri Mar 12 09:45:01.228 EST

Encap Id: 2

#Route: 1 #Encap: 1 Stale: N IP Lookup:N (MRIB Update VRF Lite(v4/v6):N/N #LocalEG:0 V6:N)

LSM ID: 9 Leaf Turnaround: True

RP/0/RSP0/CPU0:DC1-DCWAN#sh mfib encap identifier location 0/0/CPU0

Fri Mar 12 09:45:11.745 EST

Encap Id: 2

#Route: 1 #Encap: 1 Stale: N IP Lookup:N (MRIB Update VRF Lite(v4/v6):N/N #LocalEG:0 V6:N)

LSM ID: 9 Leaf Turnaround: True

Check BGP mVPN routes:

Check if BGP Type 7 is received from the receiver and Type 3 is advertised from Source to Receiver. BGP Type 3 is only advertised when data MDT is triggered.

<#root>

RP/0/RSP0/CPU0:DC1-DCWAN#sh bgp ipv4 mvpn rd 1000:1001

Wed Mar 10 07:39:26.898 EST

BGP router identifier 10.44.0.10, local AS number 55836

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 351943

BGP NSR Initial initsync version 31 (Reached)

BGP NSR/ISSU Sync-Group versions 351943/0

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: 1000:1001 (default for vrf MLDP-14 )

\*>i[1][10.44.0.5]/40 10.44.0.5 100 0 i

\* i 10.44.0.5 100 0 i

\*> [1][10.44.0.10]/40 0.0.0.0 0 i

\*>i[1][10.44.0.110]/40

10.44.0.110 100 0 i

\* i 10.44.0.110 100 0 i



```

*>i[3][0][0.0.0.0][0][0.0.0.0][10.44.0.5]/120
  10.44.0.5 100 0 i
* i 10.44.0.5 100 0 i
*> [3][0][0.0.0.0][0][0.0.0.0][10.44.0.10]/120
  0.0.0.0 0 i
*>i[3][0][0.0.0.0][0][0.0.0.0][10.44.0.110]/120
  10.44.0.110 100 0 i
* i 10.44.0.110 100 0 i
*> [3][32][10.1.0.2][32][232.11.11.11][10.44.0.10]/120
>>>> Type 3 route advertised towards receiver

  0.0.0.0 0 i
*>i[7][1000:1001][55836][32][10.1.0.2][32][232.11.11.11]/184
>>> Type 7 route received from receiver

  10.44.0.5 100 0 i
* i 10.44.0.5 100 0 i

```

Processed 8 prefixes, 13 paths.

## **Towards Multicast Source**

Check mRIB VRF:

```

RP/0/RSP0/CPU0:DC1-DCWAN#sh mrib vrf MLDP-14 route 10.1.0.2 232.11.11.11/32 detail
Fri Mar 12 05:17:51.828 EST

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

(10.1.0.2,232.11.11.11) Ver: 0x2c12 RPF nbr: 10.1.0.2 Flags: RPF MT EID,

PD: Slotmask: 0x0 >> No Slotmask as it is incoming traffic . It will get programmed when traffic is eg

MGID: 26081

MT Slot: 0/0/CPU0 I>>>> Ingress LC

Up: 02:55:10

RPF-ID: 0, Encap-ID: 2

Incoming Interface List

Bundle-Ether1000.2001 Flags: A, Up: 02:55:10 >> Multicast Source connected interface

Outgoing Interface List

LmdtMLDP-14 Flags: F LMI MT MA TR, Up: 02:55:10, Head LSM-ID: 0x00009 >>Verify LSM ID with MLDP dB & ch

RP/0/RSP0/CPU0:DC1-DCWAN#

Verify MGID with ENCAP-ID:

RP/0/RSP0/CPU0:DC1-DCWAN#sh mrib encap-id 2

Wed Mar 17 12:50:58.421 EST

Encap ID Key : 000001100000002001000201000000000000

Encap ID Length : 19

Encap ID Value : 2

Platform Annotation:

Slotmask: Primary: 0x1, Backup: 0x1 >>> Verify the Slot-mask point to LC 0

MGID: Primary: 43570, Backup: 43571 >>> Check this MGID

Flags (Vrflite(v4/v6),Stale,v6): N/N, N, N

Oles:

[1] type: 0x5, len: 12

LSM-ID: 0x00010 MDT: 0x2000120 Turnaround: TRUE  
Primary: 0/0/CPU0[1]  
Backup:  
TableId: 0xe00003f8[1]

Redist History:  
client id 22 redist time: 1d01h redist flags 0x0

MGID programming for egress LC (towards mdt):

RP/0/RSP0/CPU0:DC1-DCWAN#sh controllers mgidprgm mgidindex 43570 location 0/0/CPU0

Fri Mar 12 09:41:36.319 EST

Device MGID-Bits Client-Last-Modified

=====  
XBAR-0 10 P2MP

FIA-0 1 P2MP

FIA-1 1 P2MP

=====

Client Mask

=====

MFIBV4 0x0

MFIBV6 0x0

L2FIB 0x0

sRP-pseudo-mc 0x0

UT 0x0

Prgm-Svr 0x0

P2MP 0x3

xbar 0x0

UT1 0x0

UT2 0x0

punt\_lib 0x0

Check mFIB VRF:

RP/0/RSP0/CPU0:DC1-DCWAN#sh mfib vrf MLDP-14 platform route 10.1.0.2 232.11.11.11/32 location 0/0/CPU0

Fri Mar 12 05:20:24.584 EST

LC Type: A9K-MOD400-SE

-----

Legend:

Common SW:

Route Information

- B: BAACL check S: RPF Interface signal
- DC: Directly connected PL: Punt to LC CPU
- PR: Punt to RP PF: Punt if forwarded
- DR: Drop all RI: RPF interface
- MF: MoFRR enabled TR: Tunnel route
- TE: Tunnel encap CD: Conditional decap
- MI: MVET index NI: Not Installed
- BD: Bidir

MVPN Information (Common SW)

- MLI: MVPN table index VTID: VRF Table ID
- RPFID: RPF Identifier

MoFRR Information

- A: Active RPF interface RI: RPF interface (or RPF ID)
- WDI: Watchdog counter index

-----

Source: 10.1.0.2 Group: 232.11.11.11 Mask: 64 RPF Int: BE1000.2001 -- Ingress interface

Common SW Information

-----

MGID: 26081 MLI: 75537 Fabric Slotmask: 0x0 FGID: 0x0

-----

B S DC PL PR PF DR BD RI MF TR TE CD MI

-----

F F F F F F F F 0x2013120 F F T F 0x2

-----

MVPN/MLDP Encap Information:

Encap flag: T, Encap ID: 2 >> Verify EnCap ID with mrib

Last refcount bitmask (Enc ID): 0x0, Sets: 0, Clears: 0

Local receiver flag: F

Turnaround flag: F

Check PIM:

```
RP/0/RSP0/CPU0:DC1-DCWAN#sh pim vrf MLDP-14 neighbor
```

PIM neighbors in VRF MLDP-14

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

10.44.0.10\* LmdtMLDP-14 1d18h 00:01:18 1 (DR)

10.1.0.1\* Bundle-Ether1000.2001 19:06:32 00:01:30 1 (DR) B E >> This router is DR

```
RP/0/RSP0/CPU0:DC1-DCWAN#sh pim vrf MLDP-14 topology 232.11.11.11 detail
```

Fri Mar 12 05:23:32.981 EST

IP PIM Multicast Topology Table

Entry state: (\*S,G)[RPT/SPT] Protocol Uptime Info

Entry flags: KAT - Keep Alive Timer, AA - Assume Alive, PA - Probe Alive

RA - Really Alive, IA - Inherit Alive, LH - Last Hop

DSS - Don't Signal Sources, RR - Register Received

SR - Sending Registers, SNR - Sending Null Registers

E - MSDP External, EX - Extranet

MFA - Mofrr Active, MFP - Mofrr Primary, MFB - Mofrr Backup

DCC - Don't Check Connected, ME - MDT Encap, MD - MDT Decap

MT - Crossed Data MDT threshold, MA - Data MDT Assigned

SAJ - BGP Source Active Joined, SAR - BGP Source Active Received,

SAS - BGP Source Active Sent, IM - Inband mLDP, X - VxLAN

Interface state: Name, Uptime, Fwd, Info

Interface flags: LI - Local Interest, LD - Local Dissinterest,

II - Internal Interest, ID - Internal Dissinterest,

LH - Last Hop, AS - Assert, AB - Admin Boundary, EX - Extranet,

BGP - BGP C-Multicast Join, BP - BGP Source Active Prune,

MVS - MVPN Safi Learned, MV6S - MVPN IPv6 Safi Learned

(10.1.0.2,232.11.11.11)SPT SSM Up: 03:00:51

JP: Join(now) RPF: Bundle-Ether1000.2001,10.1.0.2\* Flags: MT MA >>>> MA Data MDT tunnel formed ..Ingress

Up: MT set (03:00:37) MDT: JoinSend Y, Cache N/N, Misc (0x0,0/0)

Cache: Add 03:00:37, Rem 1d18h. MT Cnt: Set 6, Unset 0. Joins sent 181

MDT-ifh 0x0/0x0, MT Slot none/ none

RPF-redirect BW usage: 0, Flags: 0x0, ObjID: 0x0

c-multicast-routing: PIM BGPJP: 1d18h

RPF Table: IPv4-Unicast-default

LmdtMLDP-14 03:00:51 fwd BGP >>> Signaling via BGP

Check PIM interface parameters:

RP/0/RSP0/CPU0:DC1-DCWAN#sh pim vrf MLDP-14 interface Bundle-Ether1000.2001 detail

Fri Mar 12 05:25:43.812 EST

PIM interfaces in VRF MLDP-14

IP PIM Multicast Interface State

Flag: B - Bidir enabled, NB - Bidir disabled

P - PIM Proxy enabled, NP - PIM Proxy disabled

V - Virtual Interface

BFD State - State/Interval/Multiplier

Interface PIM Nbr Hello DR

Count Intvl Prior

Bundle-Ether1000.2001 on 1 30 1 >> PIM ON Hello interval 30 sec

Primary Address : 10.1.0.1 >>> Self Address

Flags : B NP V  
BFD : Off/150 ms/3  
DR : this system >>> This DR  
Propagation delay : 500  
Override Interval : 2500  
Hello Timer : 00:00:15  
Neighbor Filter : -

### PIM Errors/Stats:

<#root>

**show**  
**pim**  
**traffic**

### PIM Traffic Counters

Elapsed time since counters cleared: 1d01h

	Received	Sent	
Valid PIM Packets	15759217	15214426	>>> verify the pim packets increment
Hello	9207	12336	>>> verify the pim hello Negotiations
Join-Prune	1076805	531981	
Data Register	14673205	0	
Null Register	73205	0	
Register Stop	0	14673205	
Assert	0	0	
Batched Assert	0	0	
Bidir DF Election	0	0	
BSR Message	0	0	>>> If you have RP check this stats
Candidate-RP Adv.	0	0	>>> If you have RP check this stats
Join groups sent		0	
Prune groups sent		0	
Output JP bytes		0	

Output hello bytes 410

Errors: >>>> Check below paramters for errors or drops

Malformed Packets	0
Bad Checksums	0
Socket Errors	0
Subnet Errors	0
Packets dropped since send queue was full	0
Packets dropped due to invalid socket	0
Packets which couldn't be accessed	0
Packets sent on Loopback Errors	6
Packets received on PIM-disabled Interface	0

## Log Collection

Collect these show techs related to multicast:

```
show tech-support multicast address-family ipv4
show tech-support multicast address-family ipv4 hardware
show tech-support multicast vrf <> address-family ipv4
show tech-support multicast vrf <> address-family ipv4 hardware
```

If you observe any NP for fabric issues or drops, capture this CLI:

```
show tech np
show tech fabric
```

Collect np datalogs:

1. Get into LC shell (run `attach 0/x/cpu0` )
2. Execute this command: `show_np <np_id> -d npdatalog`

If you observe any process crash:

1. Upload core dump for that process file end with `core.tgz` and `core.txt` . The default file generated gets saved in `harddisk:/`



2. Upload show tech for that process. The Default file generated gets saved in `harddisk:/showtech`

If you observe any process blocked through CLI, use this command: `show process blocked location all`

Take:

- Follow process `<JID>` iteration 5 location `<>`
- Dumpcore running `<JID/Process name>` location `<>`
  - This CLI generates dump core and uploads core dump for that process, files end with `core.tgz` and `core.txt`.
  - The default file generated gets saved in `harddisk:/`
- Upload show tech for that process.