MLDP-Based MVPN

MLDP-Based MVPN

The MLDP-based MVPN feature provides extensions to Label Distribution Protocol (LDP) for the setup of point-to-multipoint (P2MP) and multipoint-to-multipoint (MP2MP) label switched paths (LSPs) for transport in the Multicast Virtual Private Network (MVPN) core network.

Prerequisites for MLDP-Based MVPN

- You must be familiar with IPv4 multicast routing configuration tasks and concepts.

- Cisco Express Forwarding (CEF) must be enabled on the router for label switching.

- Unicast routing must be operational.

- To enable MLDP-based multicast VPN, you must configure a VPN routing and forwarding (VRF) instance.

Restrictions for MLDP-Based VPN

- Only MLDP profiles 1, 13, and 14 are supported.

- MLDP extranet is not supported.

- GRE tunnel in core is not supported for MLDP.

- MLDP FRR is not supported.

- Supported content group modes are Protocol Independent Multicast (PIM) sparse mode (PIM-SM), and Source Specific Multicast (SSM). Bidirectional PIM (PIM-Bidir) traffic is supported only on Profile 1.

- PIM dense mode (PIM-DM) is not supported.

- RSVP-TE-based LSM is not supported.

- The PIM-sparse content group mode is supported if the RP is configured behind the PE router (on CE) or on the source PE router.

- IGP MLDP ECMP is not supported. You must configure **no mpls mldp forwarding recursive** to use MLDP multipath.

- Dual Homing of L2 PEs are not supported for any MVPN profiles.

- MLDP in Seamless MPLS architecture is not supported.

Information About MLDP-Based MVPN

Overview of MLDP-Based MVPN

MVPN allows a service provider to configure and support multicast traffic in an MPLS VPN environment. This feature supports routing and forwarding of multicast packets for each individual VPN routing and forwarding (VRF) instance, and it also provides a mechanism to transport VPN multicast packets across the service provider backbone.

A VPN is network connectivity across a shared infrastructure, such as an Internet service provider (ISP). Its function is to provide the same policies and performance as a private network, at a reduced cost of ownership, thus creating many opportunities for cost savings through operations and infrastructure.

An MVPN allows an enterprise to transparently interconnect its private network across the network backbone of a service provider. The use of an MVPN to interconnect an enterprise network in this way does not change the way that the enterprise network is administered, nor does it change general enterprise connectivity.

As shown in the figure, in an MLDP-based MVPN, a static default multicast distribution tree (MDT) is established for each multicast domain. The default MDT defines the path used by provider edge (PE) devices to send multicast data and control messages to every other PE device in the multicast domain. A default MDT is created in the core network using a single MP2MP LSP. The default MDT behaves like a virtual LAN.

Figure 1: MLDP with the Default MDT Scenario

As shown in the figure, an MLDP-based MVPN also supports the dynamic creation of data MDTs for high-bandwidth transmission. For high-rate data sources, a data MDT is created using P2MP LSPs to off-load traffic from the default MDT to avoid unnecessary waste of bandwidth to PEs that did not join the stream. The creation of the data MDT is signaled dynamically using MDT Join TLV messages. Data MDTs are a feature unique to Cisco IOS software. Data MDTs are intended for high-bandwidth sources such as full-motion video inside the VPN to ensure optimal traffic forwarding in the MPLS VPN core. The threshold at which the data MDT is created can be configured on a per-device or a per-VRF basis. When the multicast transmission exceeds the defined threshold, the sending PE device creates the data MDT and sends a User Datagram Protocol (UDP) message, which contains information about the data MDT to all devices on the default MDT.

Figure 2: MLDP with the Data MDT Scenario

Data MDTs are created only for (S, G) multicast route entries within the VRF multicast routing table. They are not created for (*, G) entries regardless of the value of the individual source data rate.

The only transport mechanism previously available was Protocol Independent Multicast (PIM) with Multipoint Generic Routing Encapsulation (mGRE) over an IP core network. The introduction of Multicast Label Distribution Protocol (MLDP) provides transport by using MLDP with label encapsulation over an MPLS core network.

MLDP creates the MDTs as follows:

- The default MDT uses MP2MP LSPs.

    - Supports low bandwidth and control traffic between VRFs.

- The data MDT uses P2MP LSPs.

    - Supports a single high-bandwidth source stream from a VRF.

All other operations of MVPN remain the same regardless of the tunneling mechanism:

- PIM neighbors in a VRF are seen across a Label Switched Path virtual interface (LSP-VIF).

- The VPN multicast state is signaled by PIM.

The only other difference when using MLDP is that the MDT group address used in the mGRE solution is replaced with a VPN ID.

MLDP-based MVPN provides the following benefits:

- Enables the use of a single MPLS forwarding plane for both unicast and multicast traffic.

- Enables existing MPLS protection (for example, MPLS Traffic Engineering/Resource Reservation Protocol (TE/RSVP link protection) and MPLS Operations Administration and Maintenance (OAM) mechanisms to be used for multicast traffic.

- Reduces operational complexity due to the elimination of the need for PIM in the MPLS core network.

Initial Deployment of an MLDP-Based MVPN

Initial deployment of an MLDP-based MVPN involves the configuration of a default MDT and one or more data MDTs.

A static default MDT is established for each multicast domain. The default MDT defines the path used by PE devices to send multicast data and control messages to every other PE device in the multicast domain. A default MDT is created in the core network using a single MP2MP LSP.

An MLDP-based MVPN also supports the dynamic creation of data MDTs for high-bandwidth transmission.

Default MDT Creation

The figure shows the default MDT scenario. The Opaque value used to signal a default MDT consists of two parameters: the VPN ID and the MDT number for the VPN in the format (vpn-id, 0) where vpn-id is a manually configured 7-byte number that uniquely identifies this VPN. The default MDT is set to zero.

In this scenario, each of the three PE devices belong to the VRF called VRF and they have the same VPN ID. Each PE device with the same VPN ID will join the same MP2MP tree. The PE devices have created a primary MP2MP tree rooted at P-Central (Root 1) and a backup MP2MP tree rooted at PE-North (Root 2). There are two sources at PE-West and interested receivers at both PE-North and PE-East. PE-West will choose one of

the MP2MP trees to transmit the customer VPN traffic, but all PE devices can receive traffic on either of the MP2MP trees.

Figure 3: Default MDT Scenario

LSP Downstream Default MDT Creation

The figures show the downstream tree creation for each of the roots. Each PE device configured with VPN ID 100:2 creates the same Forwarding Equivalence Class (FEC) Type Length Value (TLV), but with a different root and downstream labels per MP2MP tree. The FEC type will be MP2MP Down, which prompts the receiving Label Switched Route (LSR) to respond with an upstream label mapping message to create the upstream path.

Figure 4: Default MDT Downstream--Root 1
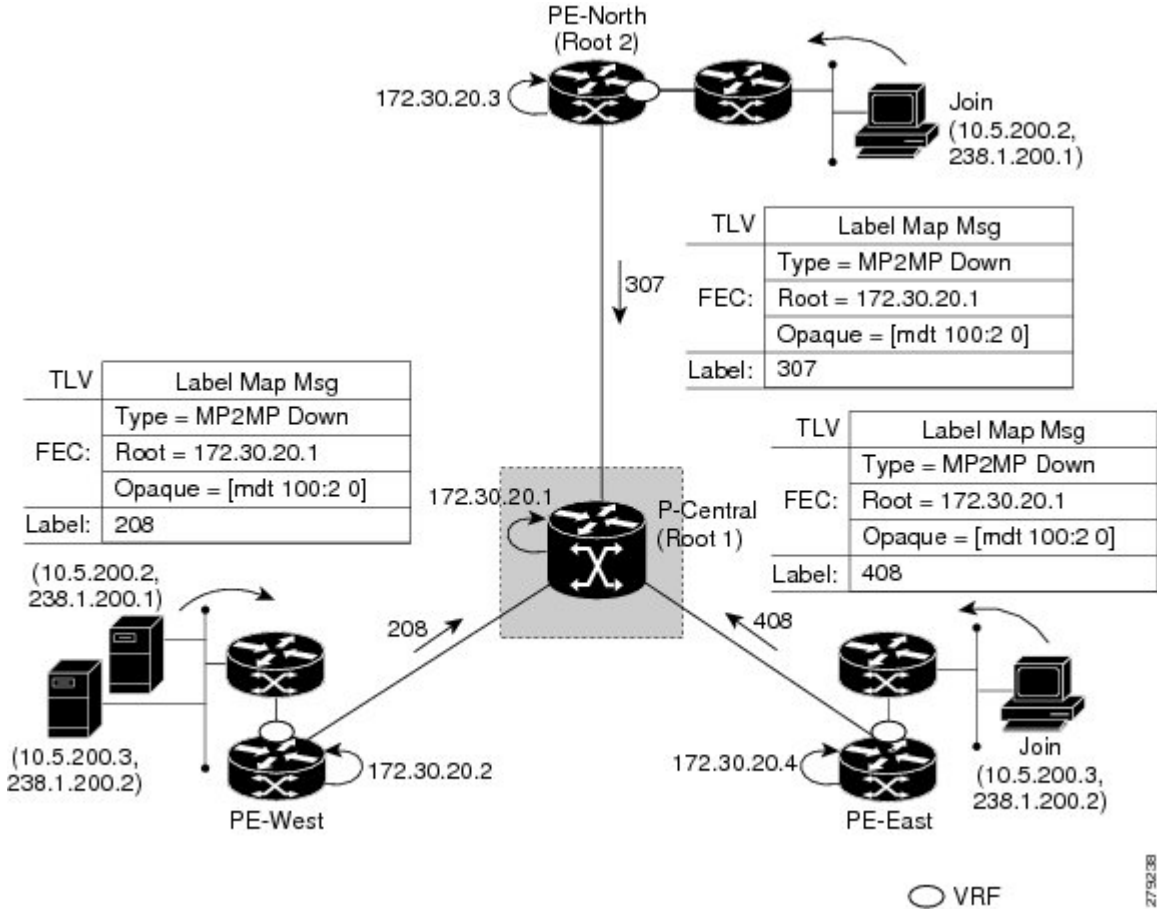
Figure 5: Default MDT Downstream--Root 2

LSP Upstream Default MDT Creation

The figures show the upstream LSP creation for the default MDTs. For each downstream label received, a corresponding upstream label is sent. In the first figure, P-Central sends out three upstream labels (111, 109, and 105) to each downstream directly connected neighbor (downstream is away from the root). The process for PE-North is the same except that it only sends a single upstream label (313) as there is only one directly connected downstream neighbor, as shown in the second figure.

Figure 6: Default MDT Upstream--Root 1

Figure 7: Default MDT Upstream--Root 2

PIM Overlay Signaling of VPN Multicast State

The signaling of the multicast state within a VPN is via PIM. It is called overlay signaling because the PIM session runs over the multipoint LSP and maps the VPN multicast flow to the LSP. In an MVPN, the operation of PIM is independent of the underlying tunnel technology. In the MVPN solution, a PIM adjacency is created between PE devices, and the multicast states within a VRF are populated over the PIM sessions. When using MLDP, the PIM session runs over an LSP-VIF interface. The figure shows PIM signaling running over the default MDT MP2MP LSP. Access to the MP2MP LSP is via the LSP-VIF, which can see all the leaf PE devices at the end of branches, much like a LAN interface. In the figure, PE-East sends a downstream label mapping message to the root, P-Central, which in turn sends an upstream label mapping message to PE-West. These messages result in the creation of the LSP between the two leaf PE devices. A PIM session can then be activated over the top of the LSP allowing the (S, G) states and control messages to be signaled between PE-West and PE-East. In this case, PE-East receives a Join TLV message for (10.5.200.3, 238.1.200.2) within VRF, which it inserts into the mroute table. The Join TLV message is then sent via the PIM session to PE-West (BGP next-hop of 10.5.200.3), which populates its VRF mroute table. This procedure is identical to the procedure using an mGRE tunnel.

Figure 8: PIM Signaling over LSP

Data MDT Scenario

In an MVPN, traffic that exceeds a certain threshold can move off the default MDT onto a data MDT.

The figure shows the data MDT scenario. The Opaque value used to signal a data MDT consists of two parameters: the VPN ID and the MDT number in the format (vpn-id, MDT# > 0) where vpn-id is a manually configured 7-byte number that uniquely identifies this VPN. The second parameter is the unique data MDT number for this VPN, which is a number greater than zero.

In the scenario, two receivers at PE-North and PE-East are interested in two sources at PE-West. If the source 10.5.200.3 exceeds the threshold on the default MDT, PE-West will issue an MDT Join TLV message over the default MDT MP2MP LSP advising all PE devices that a new data MDT is being created.

Because PE-East has an interested receiver in VRF, it will build a multipoint LSP using P2MP back to PE-West, which will be the root of the tree. PE-North does not have a receiver for 10.5.200.3, therefore it will just cache the Join TLV message.

Figure 9: Data MDT Scenario

P2MP and MP2MP Label Switched Paths

MLDP is an application that sets up Multipoint Label Switched Paths (MP LSPs) in MPLS networks without requiring multicast routing protocols in the MPLS core. MLDP constructs the P2MP or MP2MP LSPs without interacting with or relying upon any other multicast tree construction protocol. Using LDP extensions for MP LSPs and Unicast IP routing, MLDP can setup MP LSPs. The two types of MP LSPs that can be setup are Point-to-Multipoint (P2MP) and Multipoint-to-Multipoint (MP2MP) type LSPs.

A P2MP LSP allows traffic from a single root (ingress node) to be delivered to a number of leaves (egress nodes), where each P2MP tree is uniquely identified with a 2-tuple (root node address, P2MP LSP identifier). A P2MP LSP consists of a single root node, zero or more transit nodes, and one or more leaf nodes, where typically root and leaf nodes are PEs and transit nodes are P routers. A P2MP LSP setup is receiver-driven and is signaled using MLDP P2MP FEC, where LSP identifier is represented by the MP Opaque Value element. MP Opaque Value carries information that is known to ingress LSRs and Leaf LSRs, but need not be interpreted by transit LSRs. There can be several MP LSPs rooted at a given ingress node, each with its own identifier.

A MP2MP LSP allows traffic from multiple ingress nodes to be delivered to multiple egress nodes, where a MP2MP tree is uniquely identified with a 2-tuple (root node address, MP2MP LSP identifier). For a MP2MP LSP, all egress nodes, except the sending node, receive a packet sent from an ingress node.

A MP2MP LSP is similar to a P2MP LSP, but each leaf node acts as both an ingress and egress node. To build an MP2MP LSP, you can setup a downstream path and an upstream path so that:

- Downstream path is setup just like a normal P2MP LSP.

- Upstream path is setup like a P2P LSP towards the upstream router, but inherits the downstream labels from the downstream P2MP LSP.

Note We recommend that you configure one P2MP MDT tree per prefix. For example, if 500 multicast routes are needed, then you should configure at least 500 P2MP MDT trees.

Packet Flow in MLDP-based MVPN

For each packet coming in, MPLS creates multiple out-labels. Packets from the source network are replicated along the path to the receiver network. The CE1 router sends out the native IP multicast traffic. The PE1 router imposes a label on the incoming multicast packet and replicates the labeled packet towards the MPLS core network. When the packet reaches the core router (P), the packet is replicated with the appropriate labels for the MP2MP default MDT or the P2MP data MDT and transported to all the egress PEs. Once the packet reaches the egress PE, the label is removed and the IP multicast packet is replicated onto the VRF interface.

Realizing an MLDP-based MVPN

There are different ways a Label Switched Path (LSP) built by MLDP can be used depending on the requirement and nature of application such as:

- P2MP LSPs for global table transit Multicast using in-band signaling.

- P2MP/MP2MP LSPs for MVPN based on MI-PMSI or Multidirectional Inclusive Provider Multicast Service Instance (Rosen Draft).

- P2MP/MP2MP LSPs for MVPN based on MS-PMSI or Multidirectional Selective Provider Multicast Service Instance (Partitioned E-LAN).

The device performs the following important functions for the implementation of MLDP:

1. Encapsulating VRF multicast IP packet with GRE/Label and replicating to core interfaces (Imposition node).

2. Replicating multicast label packets to different interfaces with different labels (Mid node).

3. Decapsulate and replicate label packets into VRF interfaces (Disposition node).

Overview of MVPN MLDP Partitioned MDT

MVPN allows a service provider to configure and support multicast traffic in an MPLS VPN environment. This type supports routing and forwarding of multicast packets for each individual VPN routing and forwarding (VRF) instance, and it also provides a mechanism to transport VPN multicast packets across the service provider backbone. In the MLDP case, the regular label switch path forwarding is used, so core does not need

to run PIM protocol. In this scenario, the c-packets are encapsulated in the MPLS labels and forwarding is based on the MPLS Label Switched Paths (LSPs).

The MVPN MLDP service allows you to build a Protocol Independent Multicast (PIM) domain that has sources and receivers located in different sites.

To provide Layer 3 multicast services to customers with multiple distributed sites, service providers look for a secure and scalable mechanism to transmit customer multicast traffic across the provider network. Multicast VPN (MVPN) provides such services over a shared service provider backbone, using native multicast technology similar to BGP/MPLS VPN.

MVPN emulates MPLS VPN technology in its adoption of the multicast domain (MD) concept, in which provider edge (PE) routers establish virtual PIM neighbor connections with other PE routers that are connected to the same customer VPN. These PE routers thereby form a secure, virtual multicast domain over the provider network. Multicast traffic is then transmitted across the core network from one site to another, as if the traffic were going through a dedicated provider network.

Separate multicast routing and forwarding tables are maintained for each VPN routing and forwarding (VRF) instance, with traffic being sent through VPN tunnels across the service provider backbone.

In the Rosen MVPN MLDP solution, a multipoint-to-multipoint (MP2MP) default MDT is setup to carry control plane and data traffic. A disadvantage with this solution is that all PE routers that are part of the MVPN need to join this default MDT tree. Setting up a MP2MP tree between all PE routers of a MVPN is equivalent to creating N P2MP trees rooted at each PE (Where N is the number of PE routers). In an Inter-AS (Option A) solution this problem is exacerbated since all PE routers across all AS'es need to join the default MDT. Another disadvantage of this solution is that any packet sent through a default MDT reaches all the PE routers even if there is no requirement.

In the partitioned MDT approach, only those egress PE routers that receive traffic requests from a particular ingress PE join the PMSI configured at that ingress PE. This makes the number of ingress PE routers in a network to be low resulting in a limited number of trees in the core.

Supported MLDP Profiles

Profile Name	Supported on MLDP
Profile 1 Default MDT - MLDP MP2MP - PIM C-mcast Signaling	Yes
Profile 2 Partitioned MDT - MLDP MP2MP - PIM C-mcast Signaling	No
Profile 4 Partitioned MDT - MLDP MP2MP - BGP-AD - PIM C-mcast Signaling	No
Profile 5 Partitioned MDT - MLDP P2MP - BGP-AD - PIM C-mcast Signaling	No
Profile 6 VRF MLDP - In-band Signaling	No
Profile 7 Global MLDP In-band Signaling	No
Profile 9 Default MDT - MLDP - MP2MP - BGP-AD - PIM C-mcast Signaling	No

Profile Name	Supported on MLDP
Profile 12 Default MDT - MLDP - P2MP - BGP-AD - BGP C-mcast Signaling	No
Profile 13 Default MDT - MLDP - MP2MP - BGP-AD - BGP C-mcast Signaling	Yes
Profile 14 Partitioned MDT - MLDP P2MP - BGP-AD - BGP C-mast Signaling	Yes
Profile 15 Partitioned MDT - MLDP MP2MP - BGP-AD - BGP C-mast Signaling	No
Profile 17 Default MDT - MLDP - P2MP - BGP-AD - PIM C-mcast Signaling	No

How to Configure MLDP-Based MVPN

Configuring Initial MLDP Settings

Perform this task to configure the initial MLDP settings.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**mpls mldp logging notifications** **Example:** `Device(config)# mpls mldp logging notifications`	Enables MLDP logging notifications.
Step 4	**end** **Example:** `Device(config)# end`	Ends the current configuration session and returns to privileged EXEC mode.

Configuring an MLDP-Based MVPN

Perform this task to configure an MLDP-based MVPN.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**ip multicast-routing** **Example:** `Device(config)# ip multicast-routing`	Enables IP multicast routing.
Step 4	**ip multicast-routing vrf** *vrf-name* **Example:** `Device(config)# ip multicast-routing vrf VRF`	Enables IP multicast routing for the MVPN VRF specified for the *vrf-name* argument.
Step 5	**vrf definition** *vrf-name* **Example:** `Device(config)# vrf definition VRF`	Enters VRF configuration mode and defines the VPN routing instance by assigning a VRF name.
Step 6	**rd** *route-distinguisher* **Example:** `Device(config-vrf)# rd 50:11`	Creates a route distinguisher (RD) (in order to make the VRF functional). Creates the routing and forwarding tables, associates the RD with the VRF instance, and specifies the default RD for a VPN.
Step 7	**vpn id** *oui* **:** *vpn-index* **Example:** `Device(config-vrf)# vpn id 50:10`	Sets or updates the VPN ID on a VRF instance.
Step 8	**address family ipv4** **Example:** `Device(config-vrf)# address family ipv4`	Enters VRF address family configuration mode to specify an address family for a VRF. • The **ipv4** keyword specifies an IPv4 address family for a VRF.

	Command or Action	**Purpose**

Step 9	**mdt preference** { **mldp** | **pim** }
	
	Example:
	
	`Device(config-vrf-af)# mdt preference`
	`mldp`
		Specifies a preference for a particular MDT type (MLDP or PIM).

Step 10	**mdt default mpls mldp** *group-address*
	
	Example:
	
	`Device(config-vrf-af)# mdt default mpls`
	`mldp 172.30.20.1`
		Configures a default MDT group for a VPN VRF instance.

Step 11	**mdt data mpls mldp** *number-of-data-mdt*
	
	Example:
	
	`Device(config-vrf-af)# mdt data mpls`
	`mldp 255`
		Specifies a range of addresses to be used in the data MDT pool.

Step 12	**mdt data threshold** *kb/s* **list** *access-list*
	
	Example:
	
	`Device(config-vrf-af)# mdt data`
	`threshold 40 list 1`
		Defines the bandwidth threshold value in kilobits per second.

Step 13	**route target export**
	route-target-ext-community
	
	Example:
	
	`Device(config-vrf-af)# route target`
	`export 100:100`
		Creates an export route target extended community for the specified VRF.

Step 14	**route target import**
	route-target-ext-community
	
	Example:
	
	`Device(config-vrf-af)# route target`
	`import 100:100`
		Creates an import route target extended community for the specified VRF.

Step 15	**end** **Example:** `Device(config-vrf-af)# end`	Ends the current configuration session and returns to privileged EXEC mode.

Verifying the Configuration of an MLDP-Based MVPN

Perform this task in privileged EXEC mode to verify the configuration of an MLDP-based MVPN.

Procedure

Step 1 show mpls mldp database

Enter the **show mpls mldp database** command to display information in the MLDP database. It shows the FEC, the Opaque value of the FEC decoded, and the replication clients associated with it:

Example:

```
Device# show mpls mldp database
  * For interface indicates MLDP recursive forwarding is enabled
  * For RPF-ID indicates wildcard value
  > Indicates it is a Primary MLDP MDT Branch

LSM ID : CB (RNR LSM ID: CC)   Type: MP2MP   Uptime : 00:01:38
  FEC Root          : 2.2.2.2 (we are the root)
  Opaque decoded    : [mdt 3001:1 0]
  Opaque length     : 11 bytes
  Opaque value      : 02 000B 0030010000000100000000
  RNR active LSP     : (this entry)
  Upstream client(s) :
    None
      Expires       : N/A           Path Set ID  : D5
  Replication client(s):
>   MDT  (VRF vrf3001)
      Uptime        : 00:01:38      Path Set ID  : D6
      Interface     : Lspvif101     RPF-ID       : *
    33.33.33.33:0
      Uptime        : 00:01:22      Path Set ID  : D7
      Out label (D) : 2343          Interface    : Vlan2222*
      Local label (U): 466          Next Hop     : 26.1.3.2
```

Step 2 **show ip pim neighbor** [**vrf** *vrf-name*] **neighbor** [*interface-type interface-number*]

Enter the **show ip pim neighbor** command to display PIM adjacencies information:

Example:

```
Device# show ip pim vrf vrf3001 neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
192.168.1.2       Port-channel122.3001     3d19h/00:01:30    v2    1 / DR B S P G
5.5.5.5           Lspvif101                00:01:48/00:01:25 v2    1 / B S P G
7.7.7.7           Lspvif101                00:01:48/00:01:25 v2    1 / DR S P G
```

Step 3 **show ip mroute** [**vrf** *vrf-name*] [[**active** [*kbps*] [**interface** *type number*] | **bidirectional** | **count** [**terse**] | **dense** | **interface** *type number* | **proxy** | **pruned** | **sparse** | **ssm** | **static** | **summary**] | [*group-address* [*source-address*]] [**count** [**terse**] | **interface** *type number* | **proxy** | **pruned** | **summary**] | [*source-address group-address*] [**count** [**terse**] | **interface** *type number* | **proxy** | **pruned** | **summary**] | [*group-address*] **active** [*kbps*] [**interface** *type number* | **verbose**]]

Enter the **show ip mroute** command to display the contents of the multicast routing (mroute) table:

Example:

```
Device# show ip mroute vrf vrf3001 225.1.1.1 30.22.1.10
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(30.22.1.10, 225.1.1.1), 00:31:08/00:02:14, flags: JTY
  Incoming interface: Lspvif101, RPF nbr 2.2.2.2, MDT: [2, 2.2.2.2]/00:02:51
  Outgoing interface list:
    Vlan3001, Forward/Sparse, 00:31:08/00:02:35
```

Step 4   **show mpls forwarding-table** [*network* {*mask* | *length*} | **labels** *label* [**-** *label*] | **interface** *interface* | **next-hop** *address* | **lsp-tunnel** [*tunnel-id*]] [**vrf** *vrf-name*] [**detail**]

Enter the **show mpls forwarding-table**command to display the contents of the MPLS Label Forwarding Information Base (LFIB):

Example:

```
Device# show mpls forwarding-table vrf vrf3001
Local      Outgoing    Prefix           Bytes Label    Outgoing    Next Hop
Label      Label       or Tunnel Id     Switched       interface
150        No Label    192.168.1.0/24[V]   \
                                        0               aggregate/vrf3001
356        No Label    30.1.30.2/32[V]  0              Po122.3001  192.168.1.2
357        No Label    30.1.30.1/32[V]  0              Po122.3001  192.168.1.2
358        No Label    30.22.1.0/24[V]  0              Po122.3001  192.168.1.2
466   [T]  No Label    [mdt 3001:1 0][V]   \
                                        65660           aggregate/vrf3001

[T]     Forwarding through a LSP tunnel.
        View additional labelling info with the 'detail' option
```

Configuration Examples for MLDP-Based MVPN

Example: Initial Deployment of an MLDP-Based MVPN

Initial deployment of an MLDP-based MVPN involves the configuration of a default MDT and one or more data MDTs.

Default MDT Configuration

The following example shows how to configure the default MDT for an MLDP-based MVPN. This configuration is based on the sample topology illustrated in the figure.

Figure 10: Default MDT Example

This configuration is consistent for every PE device participating in the same VPN ID. The **vpn id 100:2** command replaces the MDT group address used with the mGRE transport method. To provide redundancy, two default MDT trees are statically configured, rooted at P-Central and PE-North. The selection as to which MP2MP tree the default MDT will use at a particular PE device is determined by Interior Gateway Protocol (IGP) metrics. An MP2MP LSP is implicit for the default MDT.

```
ip pim mpls source Loopback0
ip multicast-routing
ip multicast-routing vrf VRF
!
ip vrf VRF
 rd 100:2
 vpn id 100:2
 route-target export 200:2
 route-target import 200:2
 mdt default mpls mldp 172.30.20.1 (P-Central)
 mdt default mpls mldp 172.30.20.3 (PE-North)
```

PIM Adjacencies

PIM operates over the LSP-VIF as if it were a regular tunnel interface. That means PIM hellos are exchanged over the LSP-VIF to establish PIM adjacencies over the default MDT. The sample output in this section displays the three PIM adjacencies in VRF of PE-East. The listed entried here are the adjacencies to PE-West and PE-North over the MP2MP LSP via LSP-VIF interface 101.

```
PE-East# show ip pim vrf vrf3001 neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor        Interface            Uptime/Expires     Ver   DR
Address                                                       Prio/Mode
5.5.5.5         Lspvif0              00:18:54/00:01:33 v2    1 / S P G
2.2.2.2         Lspvif0              1d00h/00:01:34    v2    1 / S P G
22.22.22.22     Lspvif0              1d00h/00:01:34    v2    1 / DR S P G
```

The output from the **show ip mroute** command also shows the (S, G) entry for VRF. The stream 225.1.1.1 has the Reverse Path Forwarding (RPF) interface of LSP-VIF interface 101 and the neighbor 2.2.2.2, which is PE-West.

```
PE-East# show ip mroute vrf vrf3001 225.1.1.1 30.22.1.10
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(30.22.1.10, 225.1.1.1), 00:31:08/00:02:14, flags: JTY
  Incoming interface: Lspvif101, RPF nbr 2.2.2.2, MDT: [2, 2.2.2.2]/00:02:51
  Outgoing interface list:
    Vlan3001, Forward/Sparse, 00:31:08/00:02:35
```

MLDP Database Entry--PE-East

The sample output in this section displays the database entries for the MP2MP trees supporting the default MDT at PE-East. The database is searched by Opaque value MDT 3001:1, which results in information for two MP2MP trees (one for each root) being returned. Both trees have different system IDs and use the same Opaque value ([mdt 3001:1 1]), but with different roots. Entry 3E0 shows it is the primary MP2MP tree, therefore PE-East will transmit all source multicast traffic on this LSP, and 21C will be the backup root. The interface LSP-VIF interface 101 represents both MP2MP LSPs. The Local Label (D) is the downstream label allocated by PE-East for this tree. In other words, traffic from the root will be received with either the abel of the primary tree or the backup tree. The Out Label (U) is the label that PE-East will use to send traffic into

the tree; upstream towards the root, either 361 for the Primary Tree or 363 for the Backup Tree. Both these labels were received from P-Central.

```
PE-East# show mpls mldp database opaque_type mdt 3001:1
LSM ID : 3E0   Type: P2MP   Uptime : 00:34:24
 FEC Root          : 2.2.2.2
 Opaque decoded    : [mdt 3001:1 1]
 Opaque length     : 11 bytes
 Opaque value      : 02 000B 003001000000100000001
 Upstream client(s) :
   33.33.33.33:0    [Active]
     Expires      : Never        Path Set ID  : 1C0
     Out Label (U)  : None        Interface    : Port-channel23*
     Local Label (D): 361        Next Hop     : 104.2.3.2
 Replication client(s):
   MDT  (VRF vrf3001)
     Uptime       : 00:34:24     Path Set ID  : None
     Interface    : Lspvif101    RPF-ID       : *

LSM ID : 21C   Type: P2MP   Uptime : 00:34:16
 FEC Root          : 2.2.2.2
 Opaque decoded    : [mdt 3001:1 2]
 Opaque length     : 11 bytes
 Opaque value      : 02 000B 003001000000100000002
 Upstream client(s) :
   33.33.33.33:0    [Active]
     Expires      : Never        Path Set ID  : 17D
     Out Label (U)  : None        Interface    : Port-channel23*
     Local Label (D): 363        Next Hop     : 104.2.3.2
 Replication client(s):
   MDT  (VRF vrf3001)
     Uptime       : 00:34:16     Path Set ID  : None
     Interface    : Lspvif101    RPF-ID       : *
```

Label Forwarding Entry--P-Central (Root 1)

The sample output shown in this section displays the VRF (MDT 3001:1) MLDP database entry 7035A for the primary MP2MP LSP, which is P-Central. Because the local device P-Central is the root, there is no upstream peer ID, therefore no labels are allocated locally. However there are three replication clients, representing each of the three PE devices: PE-North, PE-West, and PE-East. These replication clients are the downstream nodes of the P2MP LSP. These clients receive multipoint replicated traffic.

In the replication entry looking from the perspective of the root, there are two types of labels:

- Out label (D)--These are labels received from remote peers that are downstream to the root (remember traffic flows downstream away from the root).

- Local label (U)--These are labels provided by P-Central to its neighbors to be used as upstream labels (sending traffic to the root). It is easy to identify these labels as they all start in the 100 range, which we have configured for P-Central to use. P-Central sends these labels out when it receives a FEC with the type as MP2MP Down.

From the labels received and sent in the replication entries, the Label Forwarding Information Base (LFIB) is created. The LFIB has one entry per upstream path and one entry per downstream path. In this case because P-Central is the root, there are only upstream entries in the LFIB that have been merged with the corresponding downstream labels. For example, label 105 is the label P-Central sent to PE-East to send source traffic upstream. Traffic received from PE-East will then be replicated using the downstream labels 307 to PE-West and 208 to PE-North.

```
P-Central# show mpls mldp database opaque_type mdt 3001:1
LSM ID : 7035A   Type: P2MP   Uptime : 00:01:13
  FEC Root          : 2.2.2.2
  Opaque decoded    : [mdt 3001:1 1]
  Opaque length     : 11 bytes
  Opaque value      : 02 000B 0030010000000100000001
  Upstream client(s) :
    33.33.33.33:0    [Active]
      Expires       : Never         Path Set ID  : 501A2
      Out Label (U) : None          Interface    : Vlan31*
      Local Label (D): 997          Next Hop     : 104.3.1.2
  Replication client(s):
    MDT  (VRF vrf3001)
      Uptime        : 00:01:13      Path Set ID  : None
      Interface     : Lspvif1       RPF-ID       : *
```

The sample output shown in this section displays the entry on P-Central for the P2MP LSP rooted at PE-North (backup root). In this tree P-Central is a branch of the tree, not a root, therefore there are some minor differences to note:

- The upstream peer ID is PE-North, therefore P-Central has allocated label 915 in the downstream direction towards PE-North and subsequently PE-North has responded with an upstream label.

- Two replication entries representing PE-East and PE-West are displayed.

- The merged LFIB shows three entries:

  - One downstream entry label 915 receiving traffic from Root 2 (PE-North), which is then directed further downstream using out labels of PE-West and PE-East.

  - Two upstream entries receiving traffic from the leaves and directing it either downstream or upstream using label out label.

```
Central_P# show mpls mldp database opaque_type mdt 3001:1
LSM ID : 3024C (RNR LSM ID: 1026F)  Type: MP2MP  Uptime : 2w3d
FEC Root               : 2.2.2.2
Opaque decoded         : [mdt 3001:1 0]
Opaque length          : 11 bytes
Opaque value           : 02 000B 0030010000000100000000
RNR active LSP         : 101F6 (root: 22.22.22.22)
Upstream client(s)     :
33.33.33.33:0 [Active]
Expires                : Never         Path Set ID  : D0157
Out Label (U)          : 4069          Interface    : Port-channel31*
Local Label (D)        : 915           Next Hop     : 104.3.1.2
Replication client(s)   :
> MDT (VRF vrf3001)
Uptime                 : 2w3d          Path Set ID  : F0036
Interface              : Lspvif1       RPF-ID       : *
7.7.7.7:0
Uptime                 : 1d20h         Path Set ID  : B01ED
Out label (D)          : 25            Interface    : Port-channel71.1*
Local label (U)        : 941           Next Hop     : 104.71.1.1

LSM ID : 101F6 (RNR LSM ID: 1026F)  Type: MP2MP  Uptime : 21:17:45
FEC Root               : 22.22.22.22 (we are the root)
Opaque decoded         : [mdt 3001:1 0]
Opaque length          : 11 bytes
Opaque value           : 02 000B 0030010000000100000000
```

```
RNR active LSP          : (this entry)
Candidate RNR ID(s)     : 3024C
Upstream client(s)      :
None
Expires                 : N/A          Path Set ID   : F007B
Replication client(s)   :
> MDT (VRF vrf3001)
Uptime                  : 20:51:46     Path Set ID   : C001F
Interface               : Lspvif1      RPF-ID        : *
7.7.7.7:0
Uptime                  : 20:51:43     Path Set ID   : C0020
Out label (D)           : 44           Interface     : Port-channel71.1*
Local label (U)         : 1191         Next Hop      : 104.71.1.1
33.33.33.33:0
Uptime                  : 00:00:34     Path Set ID   : 100049
Out label (D)           : 3109         Interface     : Port-channel31*
Local label (U)         : 1340         Next Hop      : 104.3.1.2
```

Data MDT Configuration

The following example shows how to configure the data MDT for an MLDP-based MVPN. This configuration is based on the sample topology illustrated in the figure.

Figure 11: Data MDT Example

The sample output in this section displays the data MDT configuration for all the PE devices. The **mdt data** commands are the only additional commands necessary. The first **mdt data**command allows a maximum of 60 data MDTs to be created, and the second **mdt data**command sets the threshold. If the number of data MDTs exceeds 60, then the data MDTs will be reused in the same way as they are for the mGRE tunnel method (the one with the lowest reference count).

```
ip pim vrf VRF mpls source Loopback0
!
ip vrf VRF
 rd 100:2
 vpn id 100:2
 route-target export 200:2
 route-target import 200:2
 mdt default mpls mldp 172.30.20.1 (P-Central)
 mdt default mpls mldp 172.30.20.3 (PE-North)
 mdt data mpls mldp 60
 mdt data threshold 1
```

VRF mroute Table--PE-West

The sample output in this section displays the VRF mroute table on PE-West before the high-bandwidth source exceeds the threshold. At this point there are two streams, representing each of the two VPN sources at PE-West, on a single MP2MP LSP (System ID 2). The LSP represents the default MDT accessed via LSP-VIF interface 0.

```
PE-West# show ip mroute vrf vrf3001 verbose.
.
.

(30.0.5.10, 228.1.1.1), 16:08:00/00:02:21, flags: FTAp
Incoming interface: Vlan3001, RPF nbr 0.0.0.0
Outgoing interface list:
Lspvif0, LSM MDT: 2 (default), Forward/Sparse, 16:08:00/00:03:25, Pkts:0, p

.
.
.

(30.0.5.10, 228.1.1.3), 15:55:20/00:01:38, flags: FTAp
Incoming interface: Vlan3001, RPF nbr 0.0.0.0
Outgoing interface list:
Lspvif0, LSM MDT: 2 (default), Forward/Sparse, 15:55:13/00:02:44, Pkts:0, p
```

The sample output in this section displays the output after the source transmission exceeds the threshold. PE-West sends an MDT Join TLV message to signal the creation of a data MDT. In this case, the data MDT number is 8, therefore PE-East will send a label mapping message back to PE-West with a FEC TLV containing root=PE-West, Opaque value=(mdt vpn-id 8). The System ID is now changed to D signaling a different LSP; however, the LSP-VIF is still LSP-VIF interface 0. The (S, G) entry also has the "y" flag set indicating this stream has switched to a data MDT.

```
PE-West# show ip mroute vrf vrf3001 228.1.1.3 30.0.5.10 verbose
.
.
.
(30.0.5.10, 228.1.1.3), 16:00:17/00:02:49, flags: FTAyp
Incoming interface: Vlan3001, RPF nbr 0.0.0.0
MDT TX nr: 8 LSM-ID: 0xD
Outgoing interface list:
```

```
             Lspvif0, LSM MDT: D (data), Forward/Sparse, 16:00:10/00:02:43, Pkts:0, p
```

MLDP Database Entries

The sample output in this section displays the MLDP entry for the data MDT (F) on the ingress device PE-West. The following points about this entry should be noted:

- The tree type is P2MP with PE-West (5.5.5.5) as the root.

- The Opaque value is [mdt 3001:1 10] denoting the first data MDT.

- There are no labels allocated as it is the root.

- There is one replication client entry on this tree.

- The MDT entry is an internal construct.

```
PE-West# show mpls mldp database id F
LSM ID : F   Type: P2MP   Uptime : 00:02:37
  FEC Root          : 5.5.5.5 (we are the root)
  Opaque decoded    : [mdt 3001:1 10]
  Opaque length     : 11 bytes
  Opaque value      : 02 000B 003001000000010000000A
  Upstream client(s) :
    None
      Expires       : N/A          Path Set ID  : 10
  Replication client(s):
>   MDT  (VRF vrf3001)
      Uptime        : 00:02:37     Path Set ID  : None
      Interface     : Lspvif0      RPF-ID       : *
    33.33.33.33:0
      Uptime        : 00:02:37     Path Set ID  : None
      Out label (D) : 3326         Interface    : Port-channel23*
      Local label (U): None        Next Hop     : 104.2.3.2
```

The sample output in this section displays the database entry for the data MDT on PE-East, the egress device. Also shown is the MDT Join TLV message that was sent from PE-West over the default MDT. The MDT Join TLV message contains all the necessary information to allow PE-East to create a label mapping message P2MP LSP back to the root of PE-West.

```
PE-East# show mpls mldp database opaque_type mdt 3001:1
LSM ID : CD   Type: P2MP   Uptime : 00:33:46
  FEC Root          : 2.2.2.2 (we are the root)
  Opaque decoded    : [mdt 3001:1 1]
  Opaque length     : 11 bytes
  Opaque value      : 02 000B 003001000000010000001
  Upstream client(s) :
    None
      Expires       : N/A          Path Set ID  : D8
  Replication client(s):
>   MDT  (VRF vrf3001)
      Uptime        : 00:33:46     Path Set ID  : None
      Interface     : Lspvif101    RPF-ID       : *
    33.33.33.33:0
      Uptime        : 00:33:46     Path Set ID  : None
      Out label (D) : 348          Interface    : Vlan2222*
      Local label (U): None        Next Hop     : 26.1.3.2

LSM ID : CE   Type: P2MP   Uptime : 00:33:38
```

```
FEC Root          : 2.2.2.2 (we are the root)
Opaque decoded    : [mdt 3001:1 2]
Opaque length     : 11 bytes
Opaque value      : 02 000B 00300100000000100000002
Upstream client(s) :
  None
    Expires      : N/A          Path Set ID  : D9
  Replication client(s):
>   MDT  (VRF vrf3001)
      Uptime      : 00:33:38     Path Set ID  : None
      Interface   : Lspvif101    RPF-ID       : *
    33.33.33.33:0
      Uptime      : 00:33:38     Path Set ID  : None
      Out label (D) : 2399       Interface    : Vlan2222*
      Local label (U): None      Next Hop     : 26.1.3.2
```

LFIB Entry for the Data MDT

The sample output in this section displays the LFIB entry for the data MDT as it passes through P-Central and PE-East. The Tunnel ID used for the LSP is the Opaque value [mdt 3001:1 0].

```
P-Central# show mpls forwarding-table labels 1191
Local      Outgoing   Prefix         Bytes Label   Outgoing    Next Hop
Label      Label      or Tunnel Id   Switched      interface
1191       2602       [mdt 3001:1 0][V]   \
                                     156663076     Po31        104.3.1.2
      [T]  No Label   [mdt 3001:1 0][V]   \
                                     45279264      aggregate/vrf3001

[T]    Forwarding through a LSP tunnel.
       View additional labelling info with the 'detail' option


PE-East# show mpls forwarding-table vrf vrf3001
 Local      Outgoing   Prefix         Bytes Label   Outgoing    Next Hop
Label      Label      or Tunnel Id   Switched      interface
132        No Label   30.0.1.0/24[V]  0             drop
133        Pop Label  30.30.1.1/32[V] 0             aggregate/vrf3001
137        Pop Label  30.1.30.1/32[V] 0             aggregate/vrf3001
138        No Label   30.0.5.0/24[V]  0             aggregate/vrf3001
142   [T]  No Label   [mdt 3001:1 0][V]   \
                                     905056        aggregate/vrf3001
145   [T]  No Label   [mdt 3001:1 0][V]   \
                                     7448          aggregate/vrf3001

[T]    Forwarding through a LSP tunnel.
       View additional labelling info with the 'detail' option
```

Example: Configuring MVPN Profile 1 - Default MDT - MLDP MP2MP - PIM C-mcast Signaling

The following example shows how to configure MVPN Profile 1:

```
vrf definition one
 rd 1:2
 vpn id 1000:2000
!
 address-family ipv4
  mdt default mpls mldp 10.100.1.1
  route-target export 1:1
```

```
  route-target import 1:1
 exit-address-family
!

ip multicast-routing vrf one

mpls mldp logging notifications

router bgp 1
 bgp log-neighbor-changes
 neighbor 10.100.1.7 remote-as 1
 neighbor 10.100.1.7 update-source Loopback0
 !
 address-family vpnv4
  neighbor 10.100.1.7 activate
  neighbor 10.100.1.7 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf one
  redistribute connected
  neighbor 10.2.2.9 remote-as 65002
  neighbor 10.2.2.9 activate
 exit-address-family
```

Example: Configuring MVPN Profile 13 - Default MDT - MLDP - MP2MP - BGP-AD - BGP C-mcast Signaling

The following example shows how to configure MVPN Profile 13:

```
vrf definition one
 rd 1:1
 vpn id 1000:2000
 !
 address-family ipv4
  mdt auto-discovery mldp
  mdt default mpls mldp 10.100.1.3
  mdt overlay use-bgp
  route-target export 1:1
  route-target import 1:1
 exit-address-family
!

interface Ethernet2/0
 vrf forwarding one
 ip address 10.2.1.1 255.255.255.0
 ip pim sparse-mode

router bgp 1
 neighbor 10.100.1.7 remote-as 1
 neighbor 10.100.1.7 update-source Loopback0
 !
 address-family ipv4 mvpn
  neighbor 10.100.1.7 activate
  neighbor 10.100.1.7 send-community extended
 exit-address-family
 !
 address-family vpnv4
  neighbor 10.100.1.7 activate
  neighbor 10.100.1.7 send-community extended
 exit-address-family
```

```
!
```

Example: Configuring MVPN Profile 14 - Partitioned MDT - MLDP P2MP - BGP-AD - BGP C-mast Signaling

The following example shows how to configure MVPN Profile 14:

```
vrf definition one
 rd 1:1
 !
 address-family ipv4
  mdt auto-discovery mldp
  mdt strict-rpf interface
  mdt partitioned mldp p2mp
  mdt overlay use-bgp
  route-target export 1:1
  route-target import 1:1
 exit-address-family

!
interface Ethernet2/0
 vrf forwarding one
 ip address 10.2.1.1 255.255.255.0
 ip pim sparse-mode
!

router bgp 1
 neighbor 10.100.1.7 remote-as 1
 neighbor 10.100.1.7 update-source Loopback0
 !
 address-family ipv4 mvpn
  neighbor 10.100.1.7 activate
  neighbor 10.100.1.7 send-community extended
 exit-address-family
 !
 address-family vpnv4
  neighbor 10.100.1.7 activate
  neighbor 10.100.1.7 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf one
  redistribute connected
  neighbor 10.2.1.8 remote-as 65001
  neighbor 10.2.1.8 activate
 exit-address-family
!
```

Feature History for MLDP-Based MVPN

This table provides release and related information for features explained in this module.

These features are available on all releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Amsterdam 17.3.3	MLDP-Based MVPN	The MLDP-based MVPN feature provides extensions to Label Distribution Protocol (LDP) for the setup of point-to-multipoint (P2MP) and multipoint-to-multipoint (MP2MP) label switched paths (LSPs) for transport in the Multicast Virtual Private Network (MVPN) core network.

Use Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to http://www.cisco.com/go/cfn.