Multicast VPNs - mLDP Core Tree

« Multicast VPNs - PIM Core Tree (/workbook/view/service-provider-v4/task/multicast-vpns-pim-core-tree-Mjg5MA%3D%3D) | Unified MPLS (/workbook/view/service-provider-v4/task/unified-mpls-Mjk1MA%3D%3D) »

Last updated: April 22, 2016

CONTENTS ›

Note:

Initial Configuration & Diagrams: Load the initial configuration files for the section named **Multicast**, which can be found in CCIE SPv4 Topology Diagrams & Initial Configurations (http://labs.ine.com/workbook/view/service-provider-v4/task/ccie-spv4-topology-diagrams-initial-configs). Refer to the **Multicast Diagram** in order to complete this task.

A Note Multicast Platform Support: Multicast data-plane forwarding is not supported in the current implementations of XRv. The control-plane portion is fully functional however. Throughout the Multicast sections of this workbook, we will configure the different variations of Multicast on the XRv platform to explore IOS-XR specific syntax, and to validate the protocol mechanics via the control-plane state.

Task

- R2 and XR1 are preconfigured as PE routers for Unicast MPLS L3VPN service to VRF VPN_A. Configure the network as defined below to also provide Multicast transport services to this customer.
- Configure Multicast Routing in the Service Provider Core as follows:
  - Enable mLDP in the core network.
  - R2 and XR1 should use R6 as the root node.
  - Configure VRF VPN_A on R2 and XR1 to use the default MDT built by mLDP.
  - After traffic rate increases, it should be transferred to a Data MDT. Set the threshold to 30 Kbps and ensure only up to 110 Data Trees are creates.
- Configure the Multicast Routing in the Customer sites as follows:
  - Enable PIM in Sparse Mode on all transit interfaces of R1 & XR2, and on the CE facings links of R2 and XR1.
  - Enable PIM in Sparse Mode on the Loopback0 interface of XR2.
  - Configure XR2 to announce itself as the BSR and RP Candidate for all groups in the range 224.0.0.0/4.
  - Configure R7's link to R1 to generate an IGMP Report message for the group 227.7.7.7.
  - Configure R8's link to XR2 to generate an IGMP Report message for the group 228.8.8.8.
- When complete, you should be able to perform the following verifications:
  - R1 should have built a (*,G) tree for the group (*,227.7.7.7) back to the Rendezvous Point XR2.
  - XR2 should have a (*,G) tree built for the group (*,228.8.8.8) to its attached receiver, R8.
  - R7 should be able to ping the multicast group address 228.8.8.8, and have packets forwarded out to the receiver R8.
  - R8 should be able to ping the multicast group address 227.7.7.7, and have packets forwarded

out to the receiver R7.
- High volumes of multicast traffic generated by R7 or R8 for these group addresses should cause new MDT data tunnels to form between R2 and XR1.

Configuration **Click to *collapse***

CONTENTS

```
R1:
ip multicast-routing distributed
!
interface GigabitEthernet1.12
 ip pim sparse-mode
!
interface GigabitEthernet1.17
 ip pim sparse-mode

R2:
ip multicast-routing vrf VPN_A distributed
!
vrf definition VPN_A
 vpn id 100:100
 address-family ipv4
 mdt default mpls mldp 6.6.6.6
 mdt data mpls mldp 110
 mdt data threshold 30
!
interface Loopback0
 ip pim sparse-mode
!
interface GigabitEthernet1.12
 ip pim sparse-mode
!


R7:
interface GigabitEthernet1.17
 ip igmp join-group 227.7.7.7

R8:
interface GigabitEthernet1.820
 ip igmp join-group 228.8.8.8

XR1:
vrf VPN_A
 vpn id 100:1
  !
 !
!
mpls ldp
 mldp
 !
!
multicast-routing
 address-family ipv4
  interface Loopback0
   enable
  !
  mdt source Loopback0
 !
 vrf VPN_A
```

CONTENTS

```
  address-family ipv4
   mdt data 110 threshold 30
   mdt default mldp ipv4 6.6.6.6
   interface all enable
  !
 !
!
router pim
 vrf VPN_A
  address-family ipv4
   rpf topology route-policy ROSEN_mLDP
  !
 !
!
route-policy ROSEN_mLDP
  set core-tree mldp-default
end-policy
!



XR2:
multicast-routing
 address-family ipv4
  interface all enable
 !
!
router pim
 address-family ipv4
  bsr candidate-bsr 20.20.20.20 hash-mask-len 30 priority 1
  bsr candidate-rp 20.20.20.20 group-list 224/4 priority 192 interval 60
```


Verification

This example shows one of the new methods for deploying Multicast VPNs. Instead of using PIM in the core, mLDP (Multipoint LDP) is used as a replacement to built trees yielding similar functionality as those that PIM would have built. Specifically, the trees built in the core - from now on referred to as "Core Trees" - are the Default and Data MDTs.

An mLDP built Default MDT consists of a Multipoint-to-Muiltipoint, or MP2MP, tree. This is a bidirectional LSP tree which allows all joined members to send and receive traffic. In the previous example, this tree was built using PIM. All PEs learned about all other PEs via the MDT BGP SAFI, and multicast SSM trees were built between all the PEs using the multicast group address configured under the VRF. mLDP allows the core network to use the MPLS forwarding tables to build a tree with similar functionality. The connectivity model of this tree resembles a LAN segment, where all PEs can communicate bidirectionally over the same LSP (tree).

mLDP re-uses the existing LDP adjacencies in the core network, and removes PIM from the core. In unicast forwarding, FECs are formed by binding labels to prefixes. mLDP binds labels to a 3 tuple consisting of (tree_type, root_node, opaque_value). In our example the tree-type can be either the Default or Data tree. The root-node is statically configured for the Default MDT. For the Data MDT, the the PE with the multicast source becomes the root-node. The opaque value is set to the VPN-ID configured under the VRF, and the MDT# where the Default MDT always has an MDT# of 0, and all Data MDTs have MDT#s greater than 0.

Like in the previous example, tunnels are automatically created in the VRF of each PE router which is what allows the PIM domain to be extended over the core network. The tunnels in the previous example were mGRE tunnels which used the underlying PIM based Default MDT. In this example, the tunnels automatically generated are a logical construct in each PE which map to the MP2MP LSP created in the core. These logical interfaces are unnumbered to the Loopback used for VPNv4 peering and have PIM enabled on them. In IOS this is seen as the "Lspvif0" interface, and in IOS-XR as "LmdtVPN/A".

In the previous example, the Customer PIM signaling was begin tunneled ontop of the provider PIM infrastructure. In the example, the customer PIM signaling is simply carried ontop of an LSP, similar to how regular unicast VPN traffic is forwarded. The end result is that PIM is extended over the provider network just like in the previous example. The only thing that has changed is the underlying tree in the core which the Customer PIM is signaled over.

Notice that mLDP neighbors are established with core devices without any special config in the core. mLDP is enabled by default - the only pre-requisite is that LDP is enabled, and the version of LDP code on the routers supports mLDP.

```
R2#show mpls mldp neighbors

 MLDP peer ID    : 3.3.3.3:0, uptime 00:36:28 Up,
  Target Adj     : No
  Session hndl   : 6
  Upstream count : 1
  Branch count   : 0
  Path count     : 1
  Path(s)        : 20.2.3.3          LDP GigabitEthernet1.23
  Nhop count     : 1
  Nhop list      : 20.2.3.3

 MLDP peer ID    : 4.4.4.4:0, uptime 00:36:14 Up,
  Target Adj     : No
  Session hndl   : 7
  Upstream count : 0
  Branch count   : 0
  Path count     : 1
  Path(s)        : 20.2.4.4          LDP GigabitEthernet1.24
  Nhop count     : 1
  Nhop list      : 20.2.4.4


R5#show mpls mldp neighbors

 MLDP peer ID    : 6.6.6.6:0, uptime 00:13:13 Up,
  Target Adj     : No
  Session hndl   : 6
  Upstream count : 0
  Branch count   : 0
  Path count     : 1
  Path(s)        : 20.5.6.6          LDP GigabitEthernet1.56
  Nhop count     : 0

 MLDP peer ID    : 4.4.4.4:0, uptime 00:12:59 Up,
  Target Adj     : No
  Session hndl   : 7
  Upstream count : 0
  Branch count   : 0
  Path count     : 1
  Path(s)        : 20.4.5.4          LDP GigabitEthernet1.45
  Nhop count     : 0

 MLDP peer ID    : 19.19.19.19:0, uptime 00:12:53 Up,
  Target Adj     : No
  Session hndl   : 8
  Upstream count : 0
  Branch count   : 0
  Path count     : 1
  Path(s)        : 20.5.19.19        LDP GigabitEthernet1.519
  Nhop count     : 0

RP/0/0/CPU0:XR1#show mpls mldp neighbors
```

```
Sun Jun  7 18:01:32.699 UTC

 mLDP neighbor database

  MLDP peer ID      : 5.5.5.5:0, uptime 00:26:13 Up ,

   Capabilities     : Typed Wildcard FEC, P2MP, MP2MP

   Target Adj       : No

   Upstream count   : 0

   Branch count     : 0

   Label map timer  : never

   Policy filter in : None

   Path count       : 1

   Path(s)          : 20.5.19.5         GigabitEthernet0/0/0/0.519 LDP

   Adj list         : 20.5.19.5         GigabitEthernet0/0/0/0.519

   Peer addr list   : 20.4.5.5

                    : 20.5.6.5

                    : 20.5.19.5

                    : 5.5.5.5
```

```
  MLDP peer ID      : 6.6.6.6:0, uptime 00:26:13 Up,

   Capabilities     : Typed Wildcard FEC, P2MP, MP2MP

   Target Adj       : No

   Upstream count   : 1

   Branch count     : 0

   Label map timer  : never

   Policy filter in : None

   Path count       : 1

   Path(s)          : 20.6.19.6         GigabitEthernet0/0/0/0.619 LDP

   Adj list         : 20.6.19.6         GigabitEthernet0/0/0/0.619

   Peer addr list   : 20.3.6.6

                    : 20.4.6.6

                    : 20.5.6.6

                    : 20.6.19.6

                    : 6.6.6.6
```

CONTENTS ›

mLDP is exchanged as a capability and can be disabled with the 'no mpls mldp' command:

```
R2#show mpls ldp capabilities


LDP Capabilities - [<description> (<type>)]
--------------------------------------------------------
    [ICCP (type 0x0405) MajVer 1 MinVer 0]
    [Dynamic Announcement (0x0506)]
    [ mLDP Point-to-Multipoint (0x0508) ]
    [ mLDP Multipoint-to-Multipoint (0x0509) ]
    [Typed Wildcard (0x050B)]


R2(config)#no mpls mldp
R2(config)#end


%PIM-5-NBRCHG: VRF VPN_A: neighbor 19.19.19.19 DOWN on interface Lspvif0 DR
%SYS-5-CONFIG_I: Configured from console by console
%LINK-5-CHANGED: Interface Lspvif0, changed state to administratively down
%LINEPROTO-5-UPDOWN: Line protocol on Interface Lspvif0, changed state to down


R2#show mpls mldp neighbors
R2#




R2(config)#mpls mldp
R2(config)#end
LDP Peer 3.3.3.3:0 re-announced
LDP Peer 4.4.4.4:0 re-announced
MLDP: Reevaluating peers for nhop: 20.2.4.4
MLDP: Reevaluating peers for nhop: 20.2.3.3
%SYS-5-CONFIG_I: Configured from console by console
%PIM-5-DRCHG: VRF VPN_A: DR change from neighbor 0.0.0.0 to 2.2.2.2 on interface Lspvif0
%LINK-3-UPDOWN: Interface Lspvif0, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface Lspvif0, changed state to up
```

Before validating what the forwarding entries look like in the core, lets start by validating that a PIM adjacency has been established via the virtual interface between the R2 and XR1. Recall that this adjacency is being tunneled over the MP2MP LSP created by mLDP. Due to its LAN-like forwarding behavior, if any other PE joins the tree by advertising the same 3-tuple via mLDP, it would establish a PIM adjacency with R2 and XR1 as if they were connected to a LAN.

CONTENTS

```
R2#show ip pim vrf VPN_A interface

Address          Interface                Ver/   Nbr    Query  DR         DR
                                          Mode   Count  Intvl  Prior
 2.2.2.2          Lspvif0                  v2/S   1      30     1          19.19.19.19
10.1.2.2         GigabitEthernet1.12      v2/S   1      30     1          10.1.2.2
R2#

R2#show ip pim vrf VPN_A 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
 19.19.19.19       Lspvif0                  00:00:39/00:01:35 v2    1 / DR P G
10.1.2.1          GigabitEthernet1.12      00:51:57/00:01:32 v2    1 / S P G




RP/0/0/CPU0:XR1#show pim vrf VPN_A interface
Sun Jun  7 18:29:06.685 UTC

PIM interfaces in VRF VPN_A
Address               Interface                     PIM  Nbr   Hello  DR    DR
                                                         Count Intvl  Prior

 19.19.19.19           LmdtVPN/A                     on   2     30     1     this system
10.19.20.19           GigabitEthernet0/0/0/0.1920   on   2     30     1     10.19.20.20

RP/0/0/CPU0:XR1#show pim vrf VPN_A neighbor
Sun Jun  7 18:29:16.425 UTC

PIM neighbors in VRF VPN_A
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

10.19.20.19*                 GigabitEthernet0/0/0/0.1920 00:11:19  00:01:44 1      B P E
10.19.20.20                  GigabitEthernet0/0/0/0.1920 00:11:09  00:01:40 1 (DR) B P
 2.2.2.2                      LmdtVPN/A              00:04:16  00:01:43 1      P
 19.19.19.19*                 LmdtVPN/A              00:04:19  00:01:15 1 (DR) P
```

Now that PIM has been extended over the core for the VPN, the PIM signaling taking place at the customer site should be end to end. XR2's BSR messages should have been heard by R2 and R1.

CONTENTS >

```
R2#show ip pim vrf VPN_A rp mapping
PIM Group-to-RP Mappings


Group(s) 224.0.0.0/4
  RP 20.20.20.20 (?), v2
    Info source: 20.20.20.20 (?), via bootstrap, priority 192, holdtime 150
         Uptime: 01:50:15, expires: 00:02:11



R1#show ip pim rp mapping
PIM Group-to-RP Mappings


Group(s) 224.0.0.0/4
  RP 20.20.20.20 (?), v2
    Info source: 20.20.20.20 (?), via bootstrap, priority 192, holdtime 150
         Uptime: 01:49:59, expires: 00:01:27
```

R7 is sending an IGMP Membership Report for 227.7.7.7, which R1 is receiving and turning into a PIM Join.

```
R7#show ip igmp groups
IGMP Connected Group Membership
Group Address    Interface                Uptime    Expires   Last Reporter   Group Accounted
 227.7.7.7        GigabitEthernet1.17      02:40:19  never     10.1.7.7


R1#show ip mroute 227.7.7.7
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode


( *, 227.7.7.7 ), 02:41:16/00:02:37, RP 20.20.20.20, flags: SJC
  Incoming interface: GigabitEthernet1.12, RPF nbr 10.1.2.2
  Outgoing interface list:
    GigabitEthernet1.17 , Forward/Sparse, 02:41:16/00:02:37
```

R1 now forwards this PIM join hop by hop towards XR2 - the RP. R2 receives the join, installs local state for the (*, G). Notice that the incoming interface for this (*, G) is the Lspvif0 interface.

CONTENTS

```
R2#show ip mroute vrf VPN_A 227.7.7.7
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

( *, 227.7.7.7 ), 01:53:58/00:02:48, RP 20.20.20.20, flags: S
   Incoming interface: Lspvif0, RPF nbr 19.19.19.19
  Outgoing interface list:
    GigabitEthernet1.12, Forward/Sparse, 01:53:58/00:02:48
```

R2 performs an RPF check towards 20.20.20.20 and finds that it needs to forward the Join out of interface Lspvif0, the virtual interface using the Default Core Tree created by mLDP, towards PIM neighbor 19.19.19.19.

```
R2#show ip rpf vrf VPN_A 20.20.20.20
RPF information for ? ( 20.20.20.20 )
   RPF interface: Lspvif0
   RPF neighbor: ? (19.19.19.19)
  RPF route/mask: 20.20.20.20/32
  RPF type: unicast ( bgp 100 )
  Doing distance-preferred lookups across tables
  BGP originator: 19.19.19.19
  RPF topology: ipv4 multicast base, originated from ipv4 unicast base
```

XR1 receives the Join and forwards it to XR2, creating (*, G) state in its MRIB table. The outgoing interface is set to LmdtVPN/A, the virtual interface created by mLDP.

CONTENTS

```
RP/0/0/CPU0:XR1#show pim vrf VPN_A topology 227.7.7.7
Sun Jun  7 20:24:57.359 UTC


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


( *,227.7.7.7 ) SM Up: 00:02:54 RP: 20.20.20.20
JP: Join(now) RPF: GigabitEthernet0/0/0/0.1920,10.19.20.20 Flags:
  LmdtVPN/A                  00:02:54  fwd Join(00:03:15)
```


Finally, we notice that XR2 does in fact receive the Join from R7 - properly building the shared tree by tunneling the PIM Signaling over the Default Core Tree.

```
RP/0/0/CPU0:XR2#show pim topology 227.7.7.7
Sun Jun  7 20:25:49.506 UTC

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

( *,227.7.7.7 ) SM Up: 00:03:46 RP: 20.20.20.20*
JP: Join(never)  RPF: Decapstunnel1,20.20.20.20  Flags:
  GigabitEthernet0/0/0/0.1920 00:03:46  fwd Join(00:02:46)
```


Now that we have validated that the customer signaling is working, lets explore how mLDP building these trees in the core network.

```
R2#show mpls mldp database
  * Indicates MLDP recursive forwarding is enabled

LSM ID : 7 (RNR LSM ID: 8)   Type: MP2MP   Uptime : 02:00:51
  FEC Root           : 6.6.6.6
  Opaque decoded     : [mdt 100:1 0]
  Opaque length      : 11 bytes
  Opaque value       : 02 000B 0001000000000100000000
  RNR active LSP     : (this entry)
  Upstream client(s) :
    3.3.3.3:0    [Active]
      Expires        : Never         Path Set ID  : 7
      Out Label (U)  : 23            Interface    : GigabitEthernet1.23*
      Local Label (D): 18            Next Hop     : 20.2.3.3
  Replication client(s):
    MDT  (VRF VPN_A)
      Uptime         : 02:00:51      Path Set ID  : 8
      Interface      : Lspvif0
```

R2 has an mLDP database entry for the Default MDT. Notice that the 3-tuple used for binding the label is present in the output above.

- Type : MP2MP - Multipoint-to-Multipoint
- Root : 6.6.6.6
- Opaque Value : VPN-ID:1, MDT#: 0

The Type is set to "MP2MP", which is the type of LSP used by the Default MDT - any PE connected to a Mulipoint-to-Multipoint tree can send a packet in and it be received by all other routers joined to the tree, emulating a LAN. The same PE router can also receive packets over this same LSP. The root node is also specified, which in this case is 6.6.6.6. Additionally, the Opaque value is set to "mdt 100:1 0". The "100:1" value represents the VPN-ID, and "0" at the end represents the MDT#, which is always #0 for the Default MDT. We also see that R2 advertised label 18 towards R3 - as seen by the "Local Label", and R3 advertised label 23 to R2 - as seen by the "Out Label". This LSP allows bidirectional communications between the PEs, unlike a normal LDP signaled LSP. R2 can send labeled packets into the LSP "upstream" towards to the root using label 23, but it also advertises a label in the "upstream" direction to be used for packets being sent back TO R2, coming in the "downstream" direction from the root towards R2.

Notice that the bottom of this output shows Lspvif0 as the "Replication Clients". This means that when R2 receives packets IN via this LSP - meaning packets flowing in the downstream direction from the root towards R2, it will decapsulate them and send them to the Lspvif0 interface for further processing. This is what "stitches" the Core Tree with the Customer network.

R3 also has the database entry. When it receives packets going "downstream", from the root towards the leaf R2, it has the labels in place to forward the packets towards R2. Notice that the 3-tuple is also present in this entry.

```
R3#show mpls mldp database
  * Indicates MLDP recursive forwarding is enabled


LSM ID : 4   Type: MP2MP   Uptime : 02:16:57
   FEC Root           : 6.6.6.6
   Opaque decoded     : [mdt 100:1 0]
  Opaque length      : 11 bytes
  Opaque value       : 02 000B 0001000000000100000000
  Upstream client(s) :
    6.6.6.6:0    [Active]
      Expires        : Never         Path Set ID  : 7
       Out Label (U)  : 29            Interface    : GigabitEthernet1.36*
       Local Label (D): 31            Next Hop     : 20.3.6.6
  Replication client(s):
    2.2.2.2:0
      Uptime         : 02:16:57      Path Set ID  : 8
      Out label (D)  : 18            Interface    : GigabitEthernet1.23*
      Local label (U): 23            Next Hop     : 20.2.3.2
```

This creates forwarding entries in the LFIB:

```
R3#show mpls forwarding-table
Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
16         Pop Label  20.2.4.0/24      0             Gi1.23     20.2.3.2
           Pop Label  20.2.4.0/24      0             Gi1.34     20.3.4.4
17         31         5.5.5.5/32       0             Gi1.34     20.3.4.4
           17         5.5.5.5/32       0             Gi1.36     20.3.6.6
18         Pop Label  4.4.4.4/32       0             Gi1.34     20.3.4.4
19         Pop Label  2.2.2.2/32       0             Gi1.23     20.2.3.2
20         Pop Label  6.6.6.6/32       0             Gi1.36     20.3.6.6
 23         29         [mdt 100:1 0]    42264         Gi1.36     20.3.6.6
 31         18         [mdt 100:1 0]    34032         Gi1.23     20.2.3.2
32         Pop Label  20.6.19.0/24     0             Gi1.36     20.3.6.6
33         33         20.5.19.0/24     0             Gi1.34     20.3.4.4
           20         20.5.19.0/24     0             Gi1.36     20.3.6.6
34         Pop Label  20.4.6.0/24      0             Gi1.34     20.3.4.4
           Pop Label  20.4.6.0/24      0             Gi1.36     20.3.6.6
35         Pop Label  20.5.6.0/24      0             Gi1.36     20.3.6.6
36         Pop Label  20.4.5.0/24      0             Gi1.34     20.3.4.4
37         35         19.19.19.19/32   31478         Gi1.36     20.3.6.6
```

R6 is the root of the tree, so it does not have any upstream clients. The same 3-tuple is preset in this entry, and the labels are in place to receive labeled packets from either branch (Gig1.36 or Gig1.619), and forward them towards the all other branches. This allows any leaf device rooted at 6.6.6.6 for this MP2MP tree to send and receive packets over this same tree - something which is extremely efficient from a resource consumption point of view. No PIM state is kept in the core, and very few labels are allocated in order to accomplish this.

```
R6#show mpls mldp database

  * Indicates MLDP recursive forwarding is enabled


LSM ID : 2   Type: MP2MP   Uptime : 03:14:09

   FEC Root           : 6.6.6.6 (we are the root)

   Opaque decoded     : [mdt 100:1 0]

  Opaque length      : 11 bytes

  Opaque value       : 02 000B 0001000000000100000000

  Upstream client(s) :

    None

      Expires        : N/A           Path Set ID  : 3

  Replication client(s):

     19.19.19.19:0

      Uptime         : 02:24:32      Path Set ID  : 9

       Out label (D)  : 16017         Interface    : GigabitEthernet1.619*

       Local label (U): 30            Next Hop     : 20.6.19.19

     3.3.3.3:0

      Uptime         : 02:21:47      Path Set ID  : A

       Out label (D)  : 31            Interface    : GigabitEthernet1.36*

       Local label (U): 29            Next Hop     : 20.3.6.3



R6#show mpls forwarding-table

Local      Outgoing    Prefix           Bytes Label   Outgoing   Next Hop

Label      Label       or Tunnel Id     Switched      interface

16         Pop Label  4.4.4.4/32       0             Gi1.46     20.4.6.4

17         Pop Label  5.5.5.5/32       0             Gi1.56     20.5.6.5

18         19         2.2.2.2/32       0             Gi1.36     20.3.6.3

           27         2.2.2.2/32       0             Gi1.46     20.4.6.4

19         Pop Label  3.3.3.3/32       0             Gi1.36     20.3.6.3

20         Pop Label  20.5.19.0/24     0             Gi1.56     20.5.6.5

           Pop Label  20.5.19.0/24     0             Gi1.619    20.6.19.19

 29         16017      [mdt 100:1 0]    44440         Gi1.619    20.6.19.19

 30         31         [mdt 100:1 0]    35754         Gi1.36     20.3.6.3

31         Pop Label  20.3.4.0/24      0             Gi1.36     20.3.6.3

           Pop Label  20.3.4.0/24      0             Gi1.46     20.4.6.4

32         Pop Label  20.2.4.0/24      0             Gi1.46     20.4.6.4

33         Pop Label  20.4.5.0/24      0             Gi1.46     20.4.6.4

           Pop Label  20.4.5.0/24      0             Gi1.56     20.5.6.5

34         Pop Label  20.2.3.0/24      0             Gi1.36     20.3.6.3

35         Pop Label  19.19.19.19/32   30850         Gi1.619    20.6.19.19
```

XR1 has similar output to R2 - showing that packets received on this LSP are picked up by the virtual interface for further processing.

```
RP/0/0/CPU0:XR1#show mpls mldp database

Sun Jun  7 21:04:38.696 UTC

mLDP database

LSM-ID: 0x00003 (RNR LSM-ID: 0x00004)   Type: MP2MP   Uptime: 02:39:41

  FEC Root           : 6.6.6.6

  Opaque decoded     : [mdt 100:1 0]

  RNR active LSP     : (this entry)

  Candidate RNR ID(s):

  Upstream neighbor(s) :

    6.6.6.6:0 [Active] Uptime: 02:39:41

      Next Hop       : 20.6.19.6

      Interface      : GigabitEthernet0/0/0/0.619

      Local Label (D) : 16017             Remote Label (U): 30

  Downstream  client(s):

    PIM MDT            Uptime: 02:39:41

      Egress intf     : LmdtVPN/A

      Table ID       : IPv4: 0xe000001c IPv6: 0xe080001c

      HLI            : 0x00004

      RPF ID         : 1

      Local Label    : 16001 (internal)
```

Data MDTs will be created in the core after the threshold configured is breached. This can be tested by sending traffic at a rate of over 30 Kbps on R7, as shown below. Note that due to a limitation in the implementation of XRv, only the control-plane portion of the multicast features will work. Due to this reason, pings sent to the multicast groups between R7 and R8 will not work - as they are hitting the limitation when crossing XR1 and XR2.

```
R7#ping 228.8.8.8 repeat 100000 timeout 0 size 1400


Type escape sequence to abort.

Sending 100000, 1400-byte ICMP Echos to 228.8.8.8, timeout is 0 seconds:

......................................................................

......................................................................

......................................................................

<snip>
```

R2 installs (S,G) state for this entry after the RP builds the SPT towards R7. Notice that the flags on this entry are 'Ty', meaning the SPT bit is set and that a Data MDT is being used. Also notice that the outgoing interface for the (S,G) is the Lspvif0 interface.

CONTENTS

```
R2#show ip mroute vrf VPN_A 228.8.8.8
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set , J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group,  y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 228.8.8.8), 00:00:33/stopped, RP 20.20.20.20, flags: SP
  Incoming interface: Lspvif0, RPF nbr 19.19.19.19
  Outgoing interface list: Null

( 10.1.7.7, 228.8.8.8 ), 00:00:33/00:02:26,  flags: Ty
  Incoming interface: GigabitEthernet1.12, RPF nbr 10.1.2.1
  Outgoing interface list:
    Lspvif0, Forward/Sparse, 00:00:33/00:02:56
```

This same command was ran in the previous example, and it showed the Data MDT being created with a new multicast group. Since we are not using multicast in the core, the Data MDT is differentiated by the MDT# - in this case #1.

```
R2#show ip pim vrf VPN_A mdt send


MDT-data send list for VRF: VPN_A
  (source, group)                     MDT-data group/num   ref_count
  ( 10.1.7.7, 228.8.8.8 )              1                    1
```

R2 generated a new mLDP database entry, advertising a new 3-tuple for this Data MDT. The 3-tuple in this case is as follows:

- Type : P2MP - Point-to-Multipoint
- Root : 2.2.2.2
- Opaque Value : VPN-ID:1, MDT#: 1

```
R2#show mpls mldp database
  * Indicates MLDP recursive forwarding is enabled


LSM ID : 9   Type: P2MP   Uptime : 00:01:00
   FEC Root           : 2.2.2.2 (we are the root)
   Opaque decoded     : [mdt 100:1 1]
  Opaque length      : 11 bytes
  Opaque value       : 02 000B 0001000000000100000001
  Upstream client(s) :
    None
      Expires        : N/A          Path Set ID  : 9
  Replication client(s):
     MDT  (VRF VPN_A)
      Uptime         : 00:01:00      Path Set ID  : None
       Interface      : Lspvif0
    4.4.4.4:0
      Uptime         : 00:01:00      Path Set ID  : None
       Out label (D)  : 21            Interface    : GigabitEthernet1.24*
      Local label (U): None          Next Hop     : 20.2.4.4


LSM ID : 7 (RNR LSM ID: 8)   Type: MP2MP   Uptime : 04:46:13
  FEC Root           : 6.6.6.6
  Opaque decoded     : [mdt 100:1 0]
  Opaque length      : 11 bytes
  Opaque value       : 02 000B 0001000000000100000000
  RNR active LSP     : (this entry)
  Upstream client(s) :
    3.3.3.3:0    [Active]
      Expires        : Never         Path Set ID  : 7
      Out Label (U)  : 28            Interface    : GigabitEthernet1.23*
      Local Label (D): 18            Next Hop     : 20.2.3.3
  Replication client(s):
    MDT  (VRF VPN_A)
      Uptime         : 04:46:13      Path Set ID  : 8
      Interface      : Lspvif0
```

Two notable differences between the Default and Data MDT are that Data MDTs are P2MP (Point-to-Multipoint) instead of MP2MP - meaning that R2 can send traffic into the LSP and multiple PEs can receive it, however, those same PEs cannot send traffic back to R2 using the same LSP, unlike in a MP2MP LSP. In order words, this is an unidirectional LSP. Additionally, the tree is rooted at R2 instead of at a central node in the core. Since this tree is purpose built to transport packets between the source behind R2 and any interested receivers at other customer sites, there is no need to have a LAN like tree where all PEs receive all packets with a central root node in the core.

The database entry is seen throughout the network until it reaches XR1 - the tail of the tree. Notice that the 3-tuple is visible by all routers and this is how entries are differentiated. This is analogous to a destination prefix in regular unicast label bindings.

```
R4#sh mpls mldp database
  * Indicates MLDP recursive forwarding is enabled

LSM ID : 7   Type: P2MP   Uptime : 00:01:57
   FEC Root           : 2.2.2.2
   Opaque decoded     : [mdt 100:1 1]
  Opaque length      : 11 bytes
  Opaque value       : 02 000B 0001000000000100000001
  Upstream client(s) :
    2.2.2.2:0    [Active]
      Expires        : Never        Path Set ID  : D
      Out Label (U)  : None         Interface    : GigabitEthernet1.24*
      Local Label (D): 21           Next Hop     : 20.2.4.2
  Replication client(s):
    6.6.6.6:0
      Uptime         : 00:01:57     Path Set ID  : None
      Out label (D)  : 22           Interface    : GigabitEthernet1.46*
      Local label (U): None         Next Hop     : 20.4.6.6


R4#show mpls forwarding-table
Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
17         23         19.19.19.19/32   21675         Gi1.45     20.4.5.5
           21         19.19.19.19/32   0             Gi1.46     20.4.6.6
21         22         [mdt 100:1 1]    703890        Gi1.46     20.4.6.6
26         Pop Label  3.3.3.3/32       0             Gi1.34     20.3.4.3
27         Pop Label  2.2.2.2/32       23945         Gi1.24     20.2.4.2
28         Pop Label  20.3.6.0/24      0             Gi1.34     20.3.4.3
           Pop Label  20.3.6.0/24      0             Gi1.46     20.4.6.6
29         Pop Label  20.2.3.0/24      0             Gi1.24     20.2.4.2
           Pop Label  20.2.3.0/24      0             Gi1.34     20.3.4.3
30         Pop Label  6.6.6.6/32       0             Gi1.46     20.4.6.6
31         Pop Label  5.5.5.5/32       0             Gi1.45     20.4.5.5
32         Pop Label  20.6.19.0/24     0             Gi1.46     20.4.6.6
33         Pop Label  20.5.19.0/24     0             Gi1.45     20.4.5.5
34         Pop Label  20.5.6.0/24      0             Gi1.45     20.4.5.5
           Pop Label  20.5.6.0/24      0             Gi1.46     20.4.6.6


R6#show mpls mldp database
  * Indicates MLDP recursive forwarding is enabled

LSM ID : 9   Type: P2MP   Uptime : 00:02:16
   FEC Root           : 2.2.2.2
   Opaque decoded     : [mdt 100:1 1]
  Opaque length      : 11 bytes
  Opaque value       : 02 000B 0001000000000100000001
  Upstream client(s) :
    4.4.4.4:0    [Active]
      Expires        : Never        Path Set ID  : 1C
      Out Label (U)  : None         Interface    : GigabitEthernet1.46*
      Local Label (D): 22           Next Hop     : 20.4.6.4
```

CONTENTS

```
  Replication client(s):
    19.19.19.19:0
      Uptime         : 00:02:16      Path Set ID  : None
      Out label (D)  : 16017         Interface    : GigabitEthernet1.619*
      Local label (U): None          Next Hop     : 20.6.19.19

LSM ID : 8   Type: MP2MP   Uptime : 01:52:34
  FEC Root           : 6.6.6.6 (we are the root)
  Opaque decoded     : [mdt 100:1 0]
  Opaque length      : 11 bytes
  Opaque value       : 02 000B 0001000000000100000000
  Upstream client(s) :
    None
      Expires        : N/A           Path Set ID  : 18
  Replication client(s):
    3.3.3.3:0
      Uptime         : 01:52:34      Path Set ID  : 19
      Out label (D)  : 31            Interface    : GigabitEthernet1.36*
      Local label (U): 25            Next Hop     : 20.3.6.3
    19.19.19.19:0
      Uptime         : 01:25:08      Path Set ID  : 1B
      Out label (D)  : 16000         Interface    : GigabitEthernet1.619*
      Local label (U): 24            Next Hop     : 20.6.19.19



R6#show mpls forwarding-table
Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
16         Pop Label  4.4.4.4/32       0             Gi1.46     20.4.6.4
17         Pop Label  5.5.5.5/32       0             Gi1.56     20.5.6.5
18         19         2.2.2.2/32       630           Gi1.36     20.3.6.3
           27         2.2.2.2/32       0             Gi1.46     20.4.6.4
19         Pop Label  3.3.3.3/32       0             Gi1.36     20.3.6.3
20         Pop Label  20.5.19.0/24     0             Gi1.56     20.5.6.5
           Pop Label  20.5.19.0/24     0             Gi1.619    20.6.19.19
21         Pop Label  19.19.19.19/32   2311980       Gi1.619    20.6.19.19
22         16017      [mdt 100:1 1]    792054        Gi1.619    20.6.19.19
24         31         [mdt 100:1 0]    33016         Gi1.36     20.3.6.3
25         16000      [mdt 100:1 0]    221360        Gi1.619    20.6.19.19
31         Pop Label  20.3.4.0/24      0             Gi1.36     20.3.6.3
           Pop Label  20.3.4.0/24      0             Gi1.46     20.4.6.4
32         Pop Label  20.2.4.0/24      0             Gi1.46     20.4.6.4
33         Pop Label  20.4.5.0/24      0             Gi1.46     20.4.6.4
           Pop Label  20.4.5.0/24      0             Gi1.56     20.5.6.5
34         Pop Label  20.2.3.0/24      0             Gi1.36     20.3.6.3



RP/0/0/CPU0:XR1#show mpls mldp database
Sun Jun  7 23:16:23.230 UTC
mLDP database
LSM-ID: 0x00003   Type: P2MP   Uptime: 00:02:28
  FEC Root           : 2.2.2.2
  Opaque decoded     : [mdt 100:1 1]
  Upstream neighbor(s) :
```

CONTENTS

```
    6.6.6.6:0 [Active] Uptime: 00:02:28
      Local Label (D) : 16017
  Downstream  client(s):
    PIM MDT            Uptime: 00:02:28
      Egress intf     : LmdtVPN/A
      Table ID        : IPv4: 0xe0000011 IPv6: 0xe0800011
      RPF ID          : 1

LSM-ID: 0x00001 (RNR LSM-ID: 0x00002)   Type: MP2MP   Uptime: 01:26:47
  FEC Root           : 6.6.6.6
  Opaque decoded     : [mdt 100:1 0]
  RNR active LSP     : (this entry)
  Candidate RNR ID(s):
  Upstream neighbor(s) :
    6.6.6.6:0 [Active] Uptime: 01:26:33
      Next Hop        : 20.6.19.6
      Interface       : GigabitEthernet0/0/0/0.619
      Local Label (D) : 16000             Remote Label (U): 24
  Downstream  client(s):
    PIM MDT            Uptime: 01:26:47
      Egress intf     : LmdtVPN/A
      Table ID        : IPv4: 0xe0000011 IPv6: 0xe0800011
      HLI             : 0x00002
      RPF ID          : 1
      Local Label     : 16001 (internal)
```


As with the previous example of Multicast VPN using the Rosen model with a PIM Core Tree, the RPF check rules change on the PEs. In the Next-Generation Multicast model, IOS-XR uses a routing-policy to specify how to perform this RPF check:

```
RP/0/0/CPU0:XR1#sh run router pim
Mon Jun  8 00:10:39.417 UTC
router pim
 vrf VPN_A
  address-family ipv4
   rpf topology route-policy ROSEN_mLDP
  !
 !
!
RP/0/0/CPU0:XR1#show running-config rpl
Mon Jun  8 00:10:53.616 UTC
route-policy ROSEN_mLDP
  set core-tree mldp-default
end-policy
!
```

In this example we are using the default mLDP Core-Tree, however there are several other Core-Tree models that have been implemented in IOS-XR. The default is "pim-default", which is why we did not have to change the Core-Tree type in our previous example using PIM in the core.

```
RP/0/0/CPU0:XR1(config)#route-policy TEST
RP/0/0/CPU0:XR1(config-rpl)#set core-tree ?
  ingress-replication-default      Ingress Replication Default MDT core
  ingress-replication-partitioned  Ingress Replication Partitioned MDT core
  mldp-default                     MLDP Default MDT core
  mldp-inband                      MLDP Inband core
  mldp-partitioned-mp2mp           MLDP Partitioned MP2MP MDT core
  mldp-partitioned-p2mp            MLDP Partitioned P2MP MDT core
  p2mp-te-default                  P2MP TE Default MDT core
  p2mp-te-partitioned              P2MP TE Partitioned MDT core
  parameter                        Identifier specified in the format: '$' followed by alphanumeric characters
  pim-default                      PIM Default MDT core
```

Nevertheless, the RPF neighbor/interface on R2/XR1 is still found by looking at the VPNv4 route and its next-hop in the VRF.

```
R2#show ip route vrf VPN_A 10.8.20.8


Routing Table: VPN_A

Routing entry for 10.8.20.0/24

  Known via "bgp 100", distance 200, metric 2, type internal

  Redistributing via ospf 100

  Advertised by ospf 100 subnets

  Last update from 19.19.19.19 02:23:16 ago

  Routing Descriptor Blocks:

  * 19.19.19.19 (default), from 19.19.19.19, 02:23:16 ago

      Route metric is 2, traffic share count is 1

      AS Hops 0

      MPLS label: 16014

      MPLS Flags: MPLS Required


R2#show ip rpf vrf VPN_A 10.8.20.8

RPF information for ? (10.8.20.8)

   RPF interface: Lspvif0

   RPF neighbor: ? (19.19.19.19)

  RPF route/mask: 10.8.20.0/24

  RPF type: unicast (bgp 100)

  Doing distance-preferred lookups across tables

  BGP originator: 19.19.19.19

  RPF topology: ipv4 multicast base, originated from ipv4 unicast base



RP/0/0/CPU0:XR1#show route vrf VPN_A 10.1.7.7

Mon Jun  8 00:15:27.517 UTC


Routing entry for 10.1.7.0/24

   Known via "bgp 100", distance 200, metric 2, type internal

  Installed Jun  7 23:02:48.666 for 01:12:39

  Routing Descriptor Blocks

    2.2.2.2, from 2.2.2.2

      Nexthop in Vrf: "default", Table: "default", IPv4 Unicast, Table Id: 0xe0000000

      Route metric is 2

  No advertising protos.



RP/0/0/CPU0:XR1#show pim vrf VPN_A rpf 10.1.7.7

Mon Jun  8 00:15:47.656 UTC


Table: IPv4-Unicast-default

* 10.1.7.7/32 [200/2]

    via LmdtVPN/A with rpf neighbor 2.2.2.2
```

« Multicast VPNs - PIM Core Tree (/workbook/view/service-provider-v4/task/multicast-vpns-pim-core-tree-Mjg5MA%3D%3D) | Unified MPLS (/workbook/view/service-provider-v4/task/unified-mpls-Mjk1MA%3D%3D) »

CONTENTS ›