Overlay Multicast in SDA: Effortless Deployment for Multicast-based Applications

Vineetha Potu
Technical Consulting Engineer, @vipotu

Ivan Lagunes
Technical Consulting Engineer, @ilagunes

Session ID: TACENT-2028

Agenda

01 Overview

02 Headend Replication

03 Native Multicast

04 Headend Vs Native Multicast

05 L2 Flooding

06 L2 Flooding Vs Overlay Multicast

07 Common Application Considerations

Overview

Common Applications

Application	Description
ASHRAE BACnet®	IoT, Manufacturing and Building Automation Industries (HVAC, Security, Access Control)
CobraNet®	Professional Audio Industry (Digital audio networking and control systems)
mDNS	Networking/IT Industry (Service discovery e.g., printers, devices), in conference rooms and screen sharing
Dante®	Audio/Video devices (Digital media networking)
IPTV	Entertainment Industry, Broadcasting, Telecommunications (TV over internet)

Which Feature should be used?

In SD-Access, there are several features available to handle multicast traffic.

How can we determine which feature to use for a specific application?

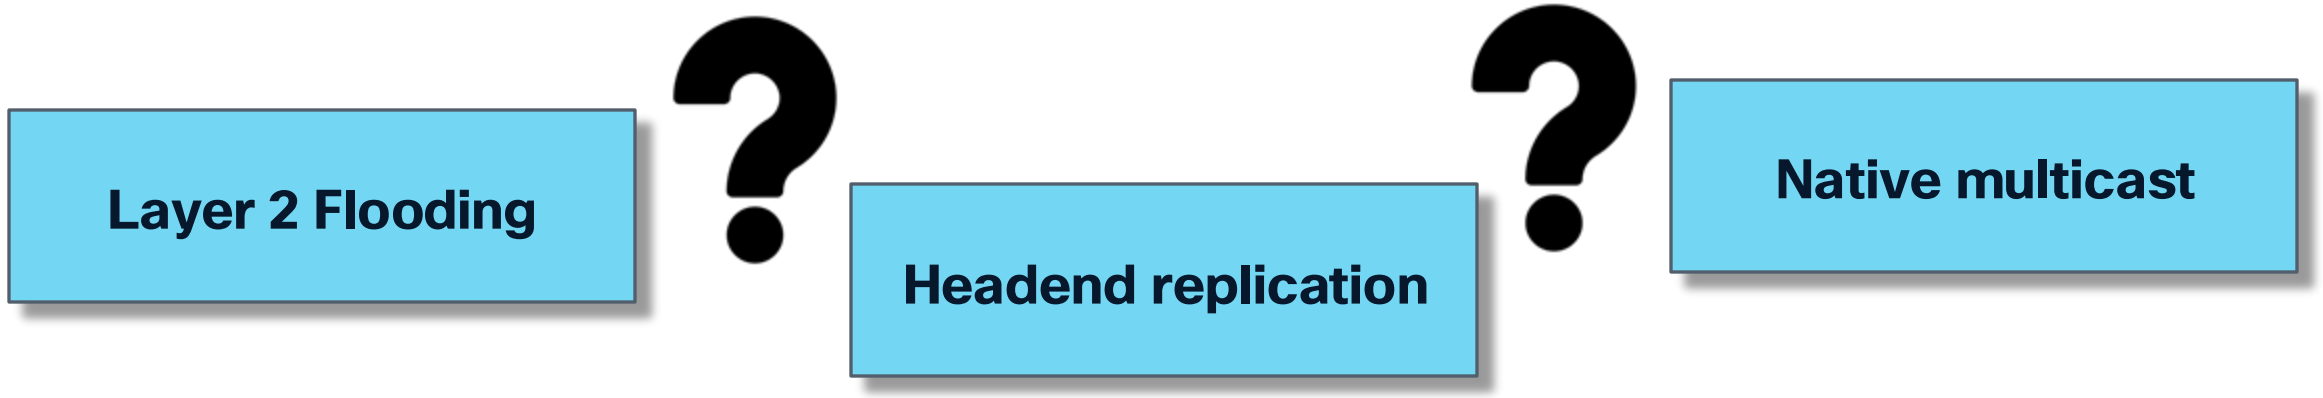

Fields within the application packet can be analyzed to determine the appropriate feature for your application.

Destination MAC

Time to Live (TTL)

Is it B-U-M traffic?

Overlay Multicast in SDA

SD-Access work with an underlay and overlay network separation.

Overlay multicast delivers traffic using the overlay network running LISP and VXLAN.

Overlay multicast enables efficient one-to-many communication for applications.

Multicast can be configured using:

- Any source multicast (Rendezvous Point)
- Source specific multicast (IGMPv3)
- Using Both, SSM and ASM

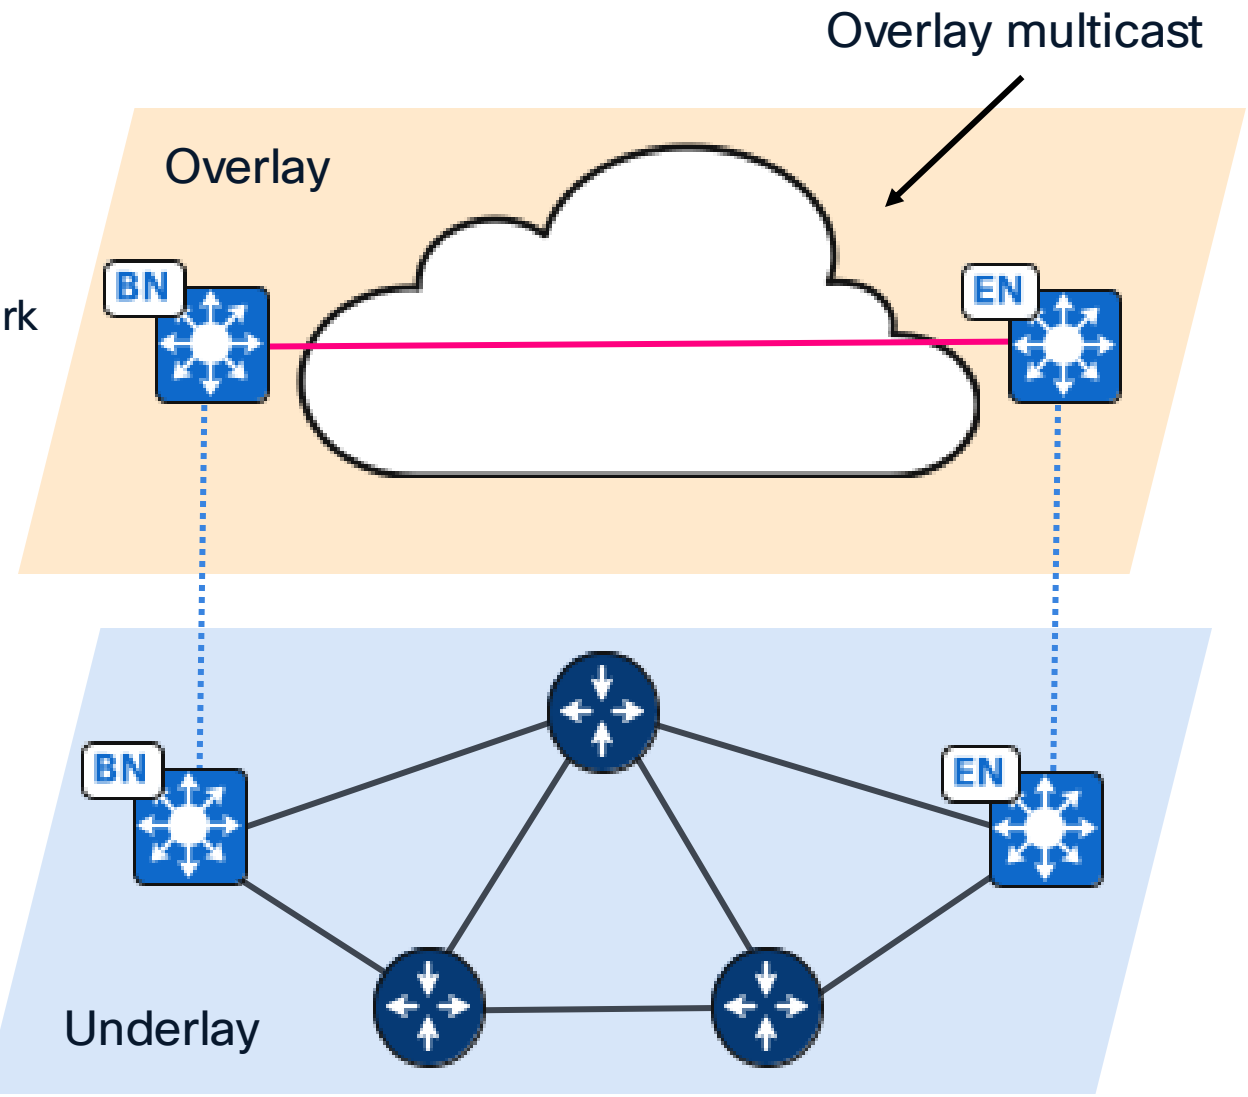

Multicast Modes in SDA

Criteria	PIM-ASM (Any source multicast)	PIM-SSM (Source specific multicast)
Rendezvous Point	Uses a Rendezvous Point	It does not use a Rendezvous Point
IGMP (Internet Group Management Protocol)	It supports IGMPv2 and v3	Every receiver needs to know the source using IGMPv3
MSDP (Multicast Source Discovery Protocol)	MSDP for Anycast-RP to enhance redundancy and load sharing.	MSDP is not needed, it does not use a Rendezvous Point
When to use	Suitable for one-to-many applications	Deployment with known source applications

Pre-requisites and Considerations

TAC Tip

You cannot change the multicast pool on the fly; you need to plan the required size of the pool in advance.

Multicast Deployment

- Overlay multicast require the use of an IP pool to configure a unique multicast IP address to each device
- Deployment for both methods can be configured using SSM, ASM or both
- Overlay multicast cannot be deployed per VLAN, it is deployed per VN (virtual network)
- You can only use one flavor per site, meaning, you cannot have headend replication for one VN and native multicast for another

Deployment

Deployment for overlay multicast can be done using Cisco Catalyst Center GUI: go to **Provision > Fabric Sites > Site_Name > Site Actions > Configure multicast**

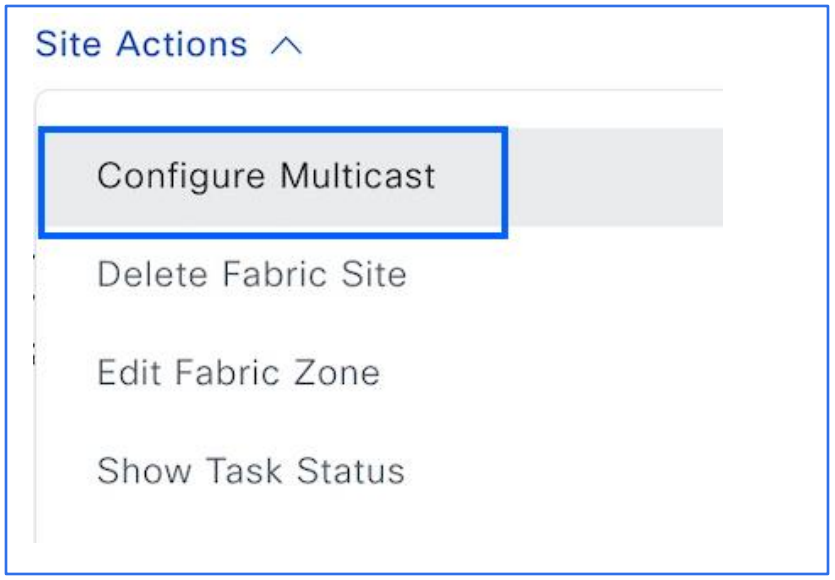

You need to choose the replication mode; there are two flavors for overlay multicast; Native multicast and Headend replication

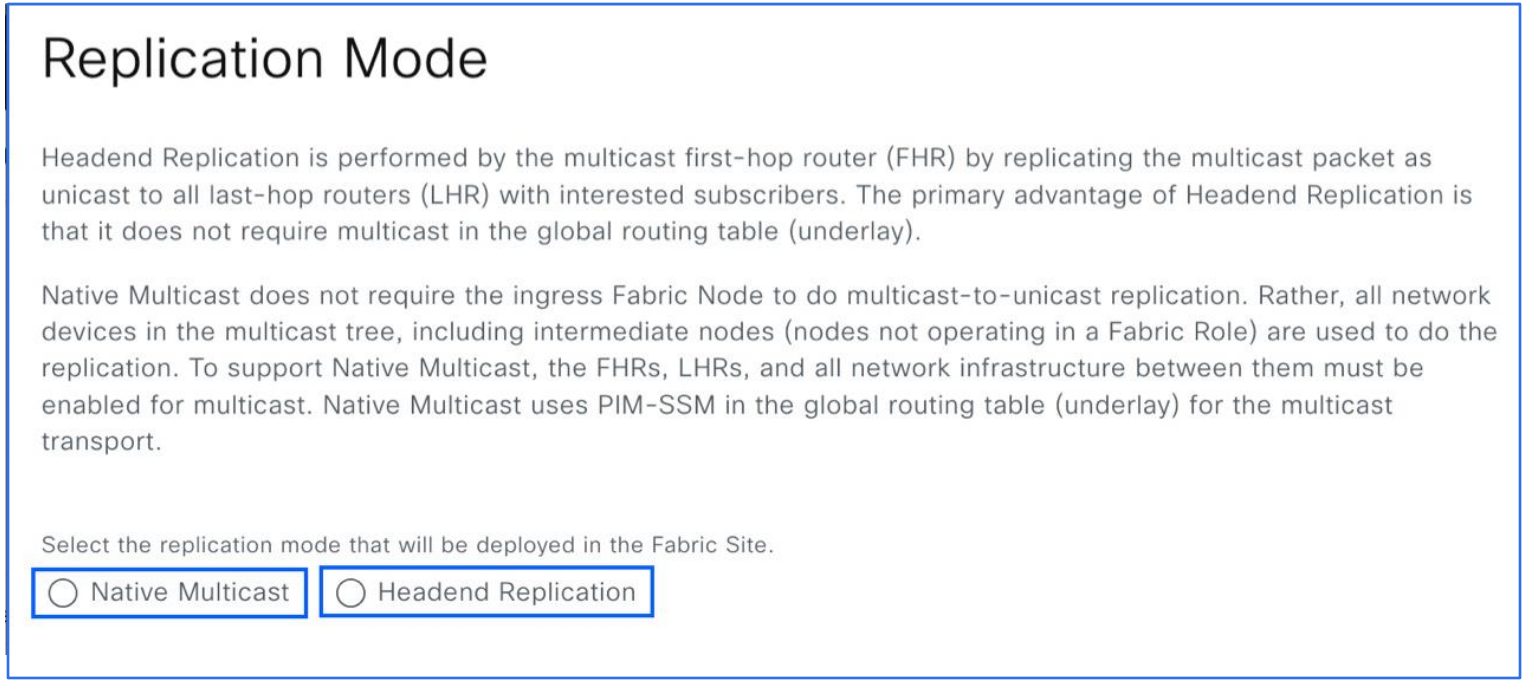

Headend Replication

Headend Replication

- Multicast packets are replicated by the First-Hop Router (FHR) and sent as unicast VXLAN packets to each Last-Hop Router (LHR).
- Multicast operates entirely in the overlay, independent of underlay multicast protocols.
- Utilizes the underlay routing table using unicast encapsulation from RLOC to RLOC (Routing Locator).

4 LHRs = Packet replicated **4** times

Source MAC: aaaa.aaaa.aaaa	**Destination MAC:** 0100.5e01.0101
Source IP: **172.13.233.1**	Destination IP: **239.1.1.1**
UDP / Data Multicast	RTP / RSTP / PRP / etc

Source MAC: source mac of underlay interface	**Destination MAC:** next hop of underlay
Source IP: **172.16.0.5**	Destination IP: **172.16.0.1,2,3,4**
UDP (VXLAN)	DST: 4789 VNI LISP IID (L3)
Source MAC: aaaa.aaaa.aaaa	Destination MAC: 0100.5e01.0101
Source IP: **172.13.233.1**	Destination IP: **239.1.1.1**
UDP / Data multicast	RTP / RSTP / PTP / etc

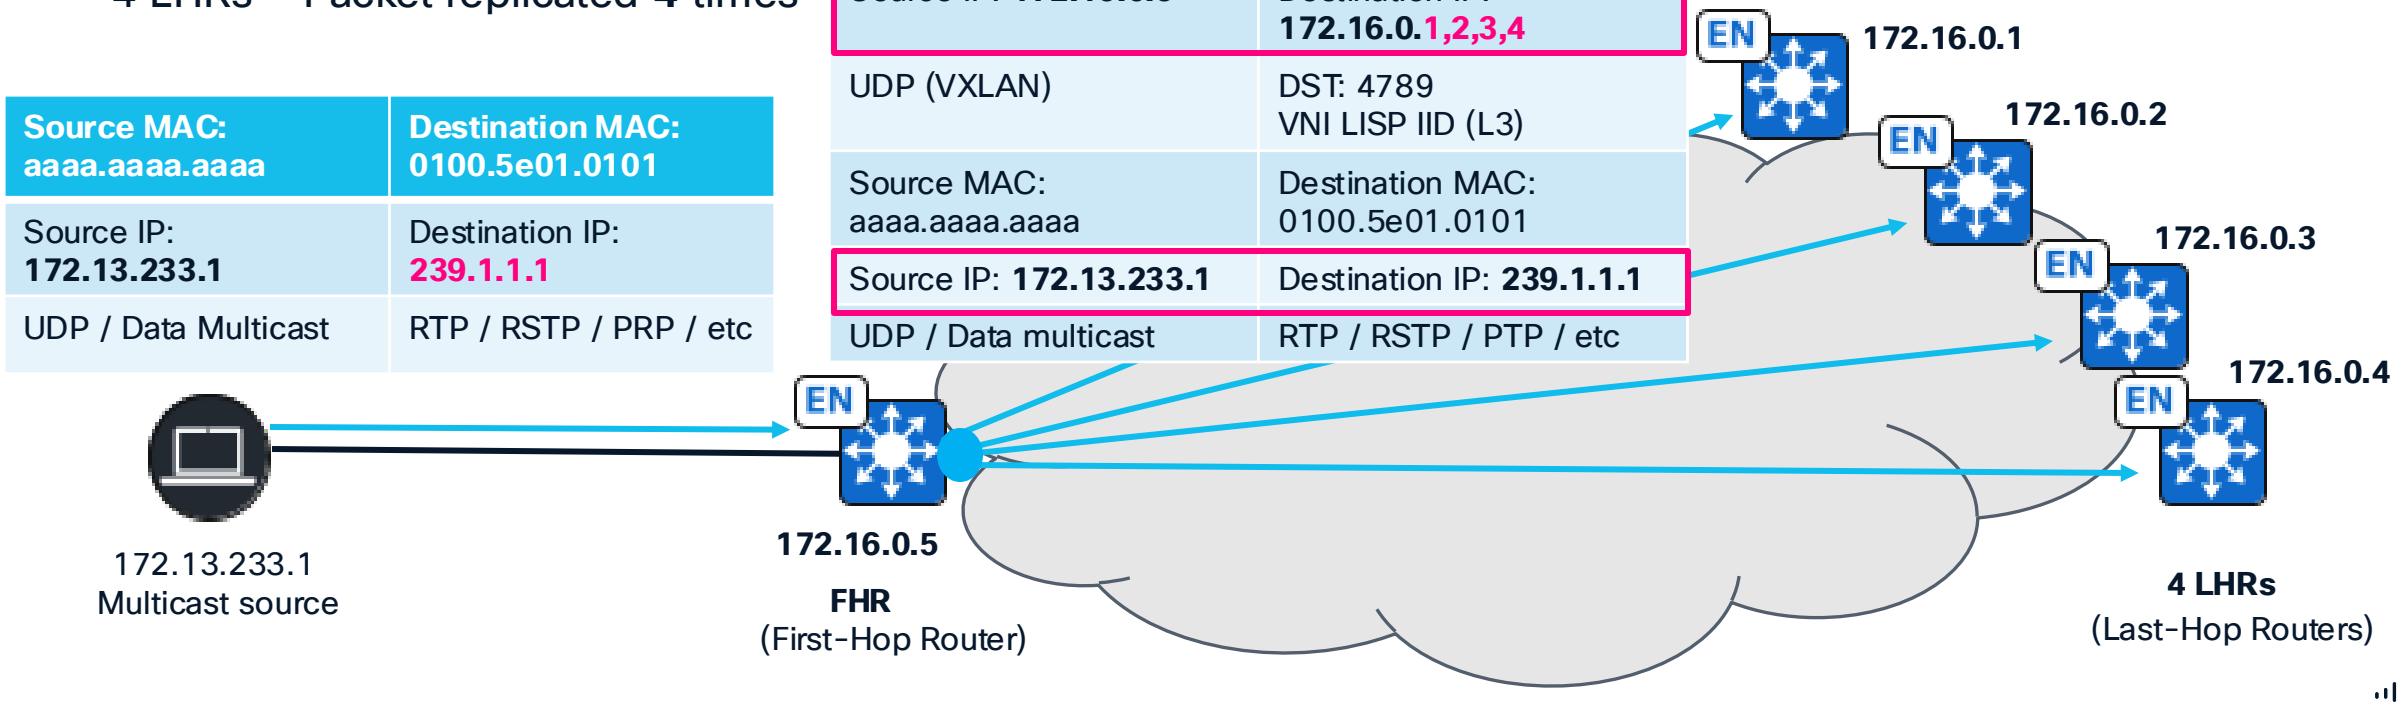

Headend Replication

Headend Replication:

- First-Hop Router (FHR) replicates multicast packets to Last-Hop Routers (LHRs) via unicast n-times.
- No underlay multicast is required, simplifying deployment, using unicast for encapsulation instead.

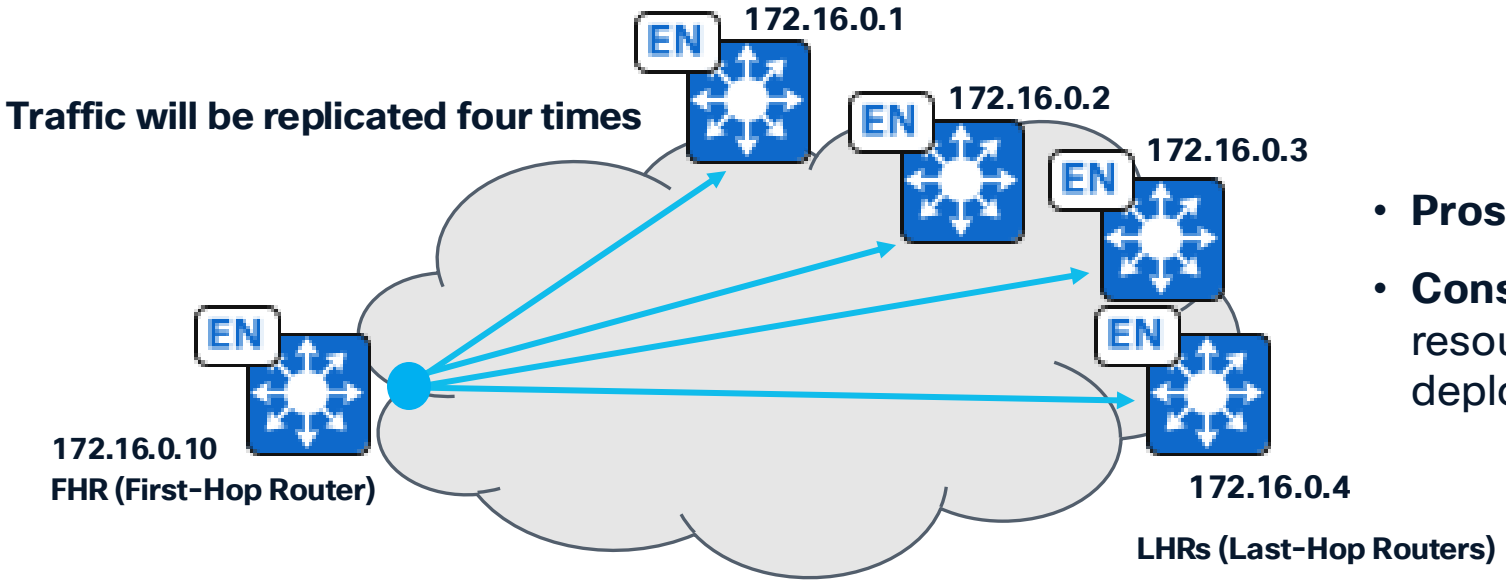

- **Pros:** Simple configuration and implementation.
- **Cons:** Scale issues for large fabrics. Higher resource/bandwidth usage on FHR for large-scale deployments.

Native Multicast

Native Multicast

- A single VXLAN-encapsulated multicast stream is sent through the underlay to reach multiple Last-Hop Routers.
- Uses underlay multicast to distribute overlay multicast traffic efficiently across the SDA fabric.
- LISP maps overlay multicast groups to underlay multicast groups.
- VXLAN encapsulation carries overlay traffic; underlay multicast ensures efficient bandwidth usage.

Overlay multicast group is mapped to an underlay multicast group in range of 232.x.x.x/8

239.1.1.1 mapped to underlay group **232.0.2.255**

Source MAC: aaaa.aaaa.aaaa	**Destination MAC:** 0100.5e01.0101
Source IP: **172.13.233.1**	Destination IP: **239.1.1.1**
UDP / Data Multicast	RTP / RSTP / PRP / etc

Source MAC: source mac of underlay interface	**Destination:** multicast MAC for the SSM group
Source IP: **172.16.0.5**	Destination IP: **232.0.2.255**
UDP (VXLAN)	DST: 4789 VNI LISP IID (L3)
Source MAC: aaaa.aaaa.aaaa	Destination MAC: 0100.5e01.0101
Source IP: **172.13.233.1**	Destination IP: **239.1.1.1**
UDP / Data multicast	RTP / RSTP / PTP / etc

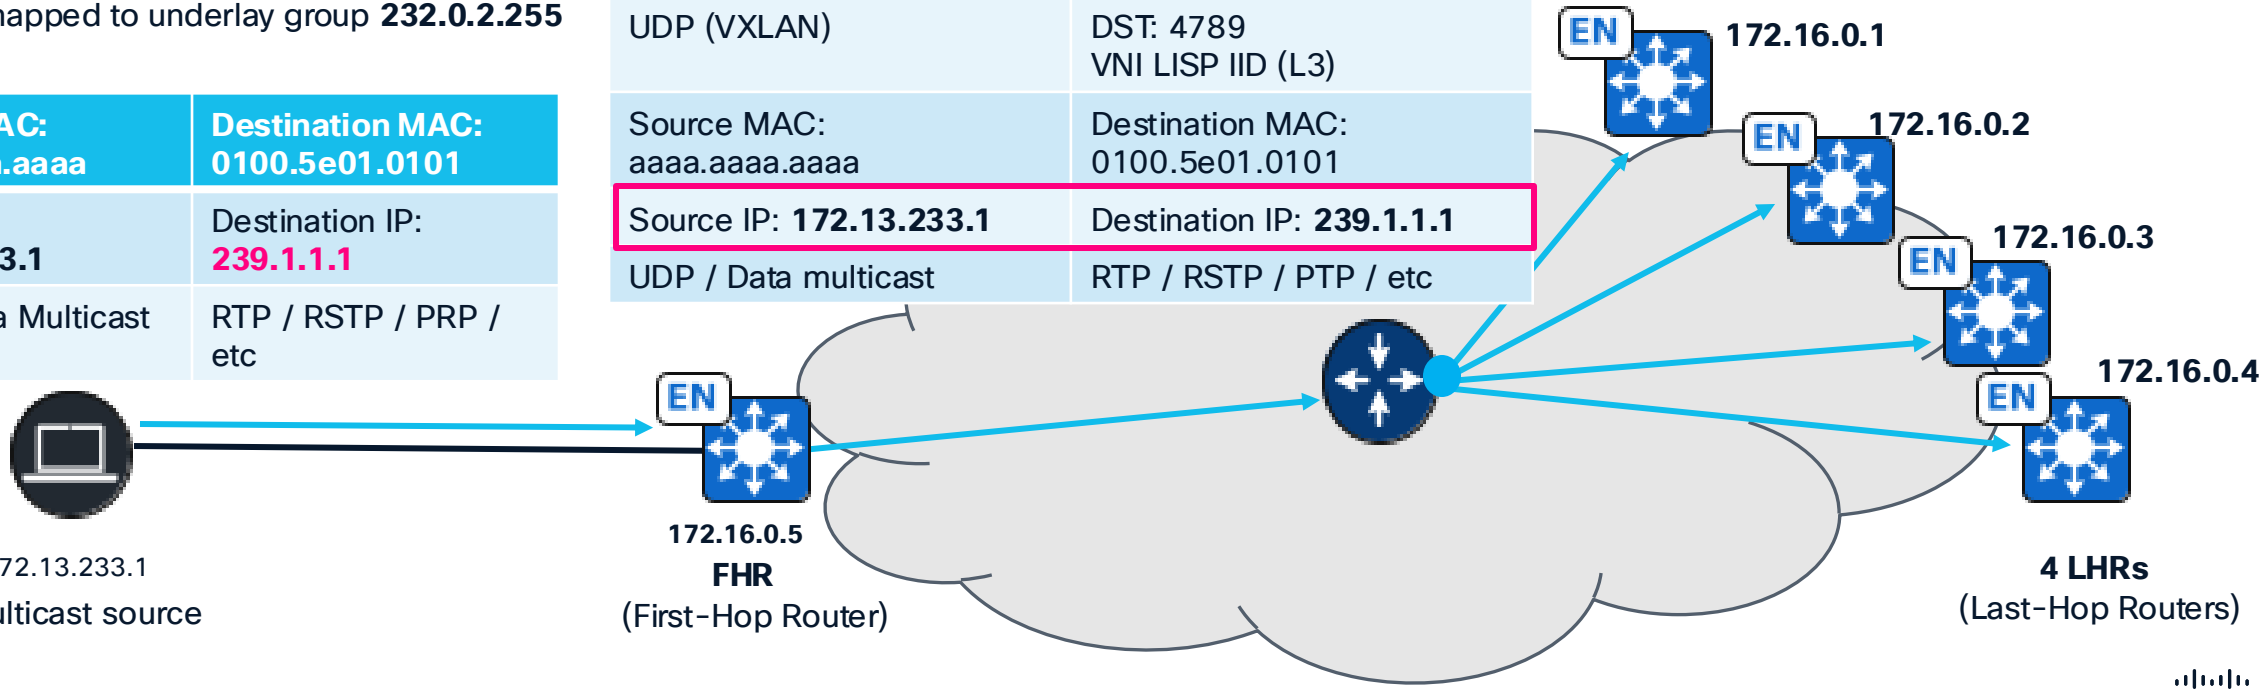

Verification

You can confirm which overlay multicast mode configured on our network by reviewing the configuration on layer 3 LISP Interface:

”ip pim sparse-mode” configured – Headend Replication

Headend Replication

```
Edge#show run interface lisp0.4104
interface LISP0.4104
 vrf forwarding water
 ip pim sparse-mode
```

”no ip pim sparse-mode” configured – Native Multicast

Native Multicast

```
Edge#show run interface lisp0.4104
interface LISP0.4104
vrf forwarding water
ip pim lisp transport multicast
ip pim lisp core-group-range 232.0.0.1 1000
end
```

SSM core group 232.0.0.0 range is defined, used for underlay multicast mapping

Another way to confirm native multicast mode is by checking the mroute form on the overlay, you will notice that the overlay group 239.1.1.1 is mapped to the underlay group 232.0.2.255.

Native Multicast

```
Edge#show ip mroute vrf water 239.1.1.1 verbose
(*, 239.1.1.1), 00:00:28/stopped, RP 172.93.94.1, flags: SPFl
  Incoming interface: LISP0.4104, RPF nbr 172.40.1.1, LISP: [172.40.1.1, 232.0.2.255]
  Outgoing interface list: Null
```

Verification

To confirm which flavors of overlay multicast is configured from Cisco Catalyst Center GUI, go to **Provision > Fabric Sites > Site_Name > Site Actions > Edit Multicast Configuration > Summary page > Look under Replication mode**

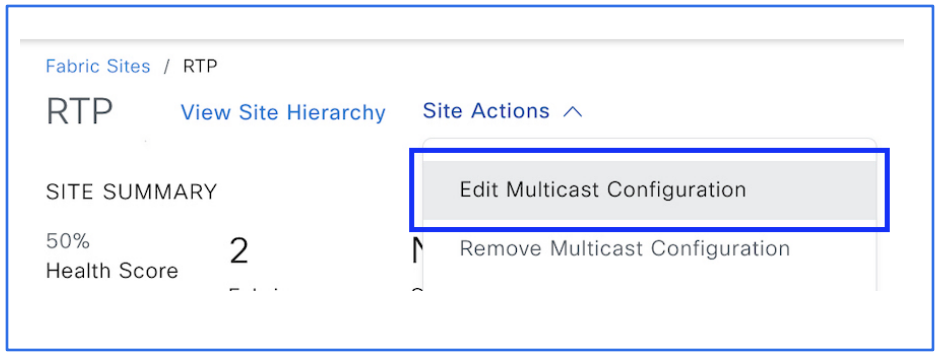

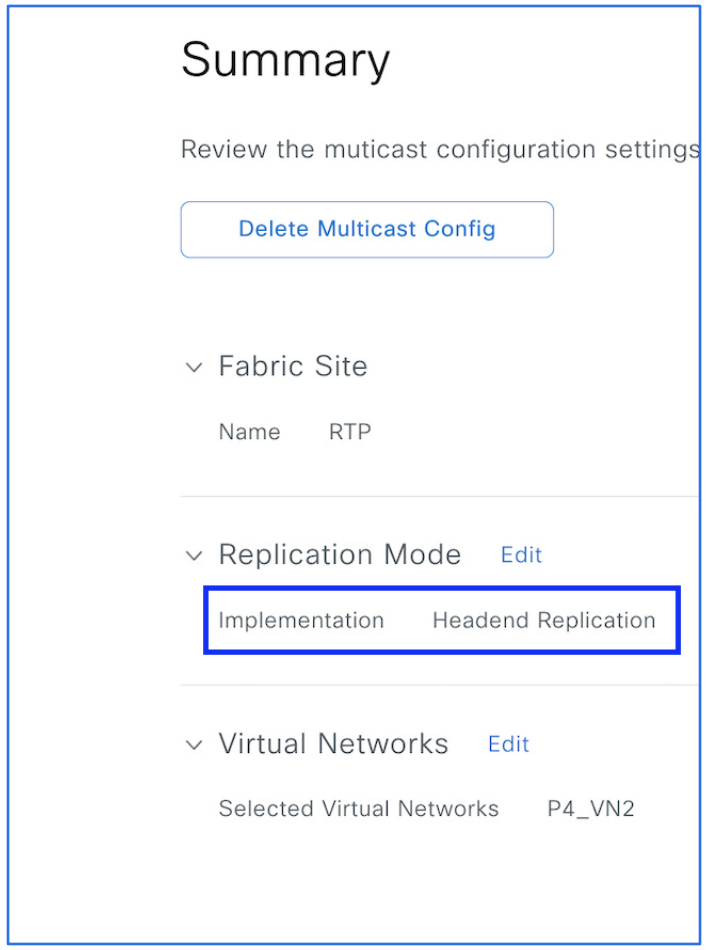

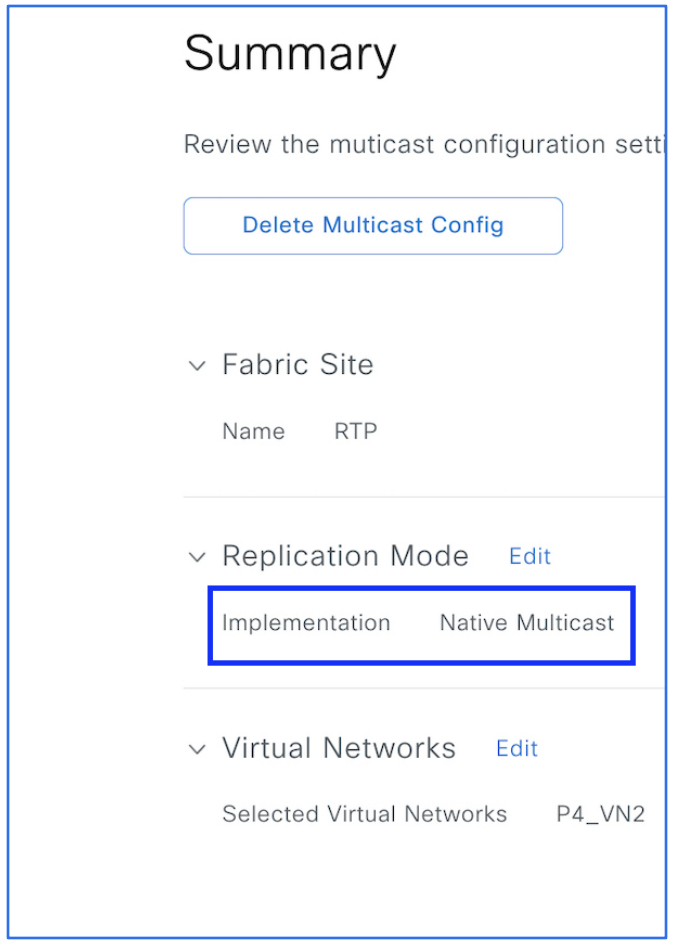

Headend vs Native Multicast

Table of Comparison

Criteria	Headend Replication	Native Multicast
Definition	FHR replicates multicast packets and sends them as unicast VXLAN packets to LHR's (Multicast into Unicast)	Uses underlay multicast (PIM) to deliver a single VXLAN-encapsulated multicast stream to multiple LHR's (Multicast into Multicast)
Underlay multicast requirement	No underlay multicast required; operates entirely in the overlay using underlay unicast	Requires PIM (ASM or SSM) configured in the underlay network
Bandwidth Efficiency	**Inefficient;** Multiple unicast copies increase bandwidth usage as receivers grow	**Highly efficient;** Single multicast stream in underlay serves all receivers ✓
Scalability	**Limited;** Scales poorly for large receiver counts (e.g., <100 receivers are recommended). Ideal for small network	**Excellent;** scales to thousands of receivers, ideal for large fabrics. ✓
Supportability	**Supported** for all device types ✓	**Not supported** for Nexus 7k, Cat 6k, and Cat 4K's
Use cases	**Non multicast underlays** Simpler to deploy and troubleshoot ✓	**Multicast capable underlays only** Complexity in deploying and troubleshooting

Challenges in SDA

- SDA is a routed access network; every switch is separated by a layer 3 routed links. There is no switching between the fabric nodes.
- The overlay multicast solution in SDA routes the packet, so the packet that can be routed is handled by overlay multicast.

What if we have a packet that meant to be switched/flooded but cannot be routed?

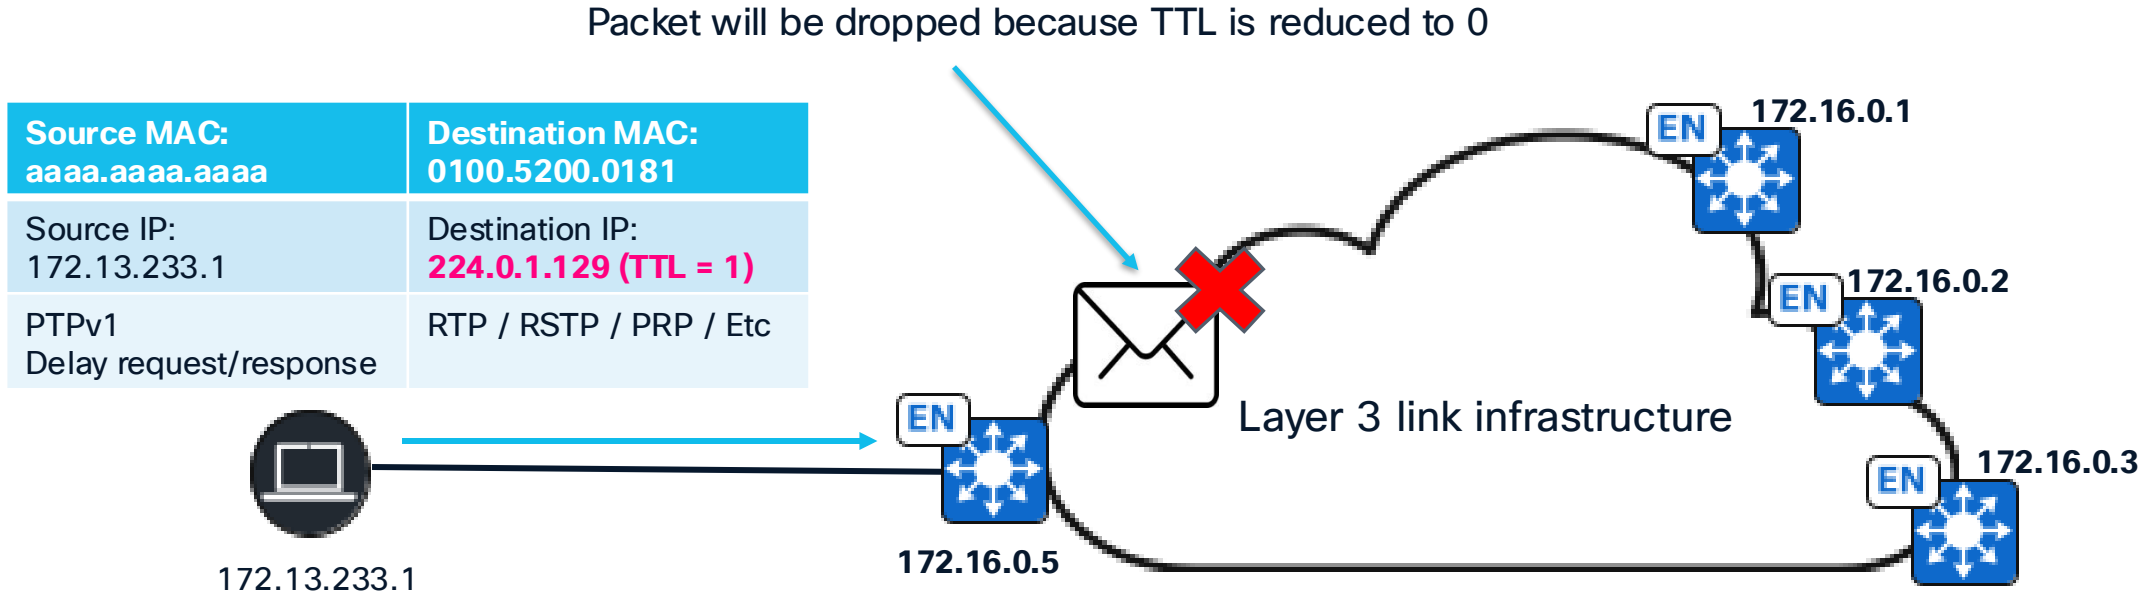

L2 Flooding

L2 Flooding

- To distribute packets that cannot be routed across SDA network, layer 2 flooding is required.
- Layer 2 flooding achieves this by creating a multicast tree in the underlay, where each edge node joins the multicast tree and receives a copy of the packet.
- All L2-VNI enabled fabric nodes will join a common multicast group (239.0.17.x) in the underlay network.
- Layer 2 flooding is a feature that can be enabled per VLAN.

Source MAC: aaaa.aaaa.aaaa	**Destination MAC: 0100.5200.0181**
Source IP: **172.13.233.1**	Destination IP: **224.0.1.129 (TTL = 1)**
PTPv1 Delay request/response	RTP / RSTP / PRP / etc

Source Mac: Underlay Mac	**Destination Mac: 0100.5e00.1101**
Source IP: **172.16.0.5**	Destination IP: **239.0.17.1**
UDP SRC: X (VXLAN)	UDP DST: 4789 VNI LISP IID (L3)
Source MAC: aaaa.aaaa.aaaa	Destination MAC: 0100.5200.0181
Source IP: **172.13.233.1**	Destination IP: **224.0.1.129**
PTPv1 request/response	RTP / RSTP / PRP / etc

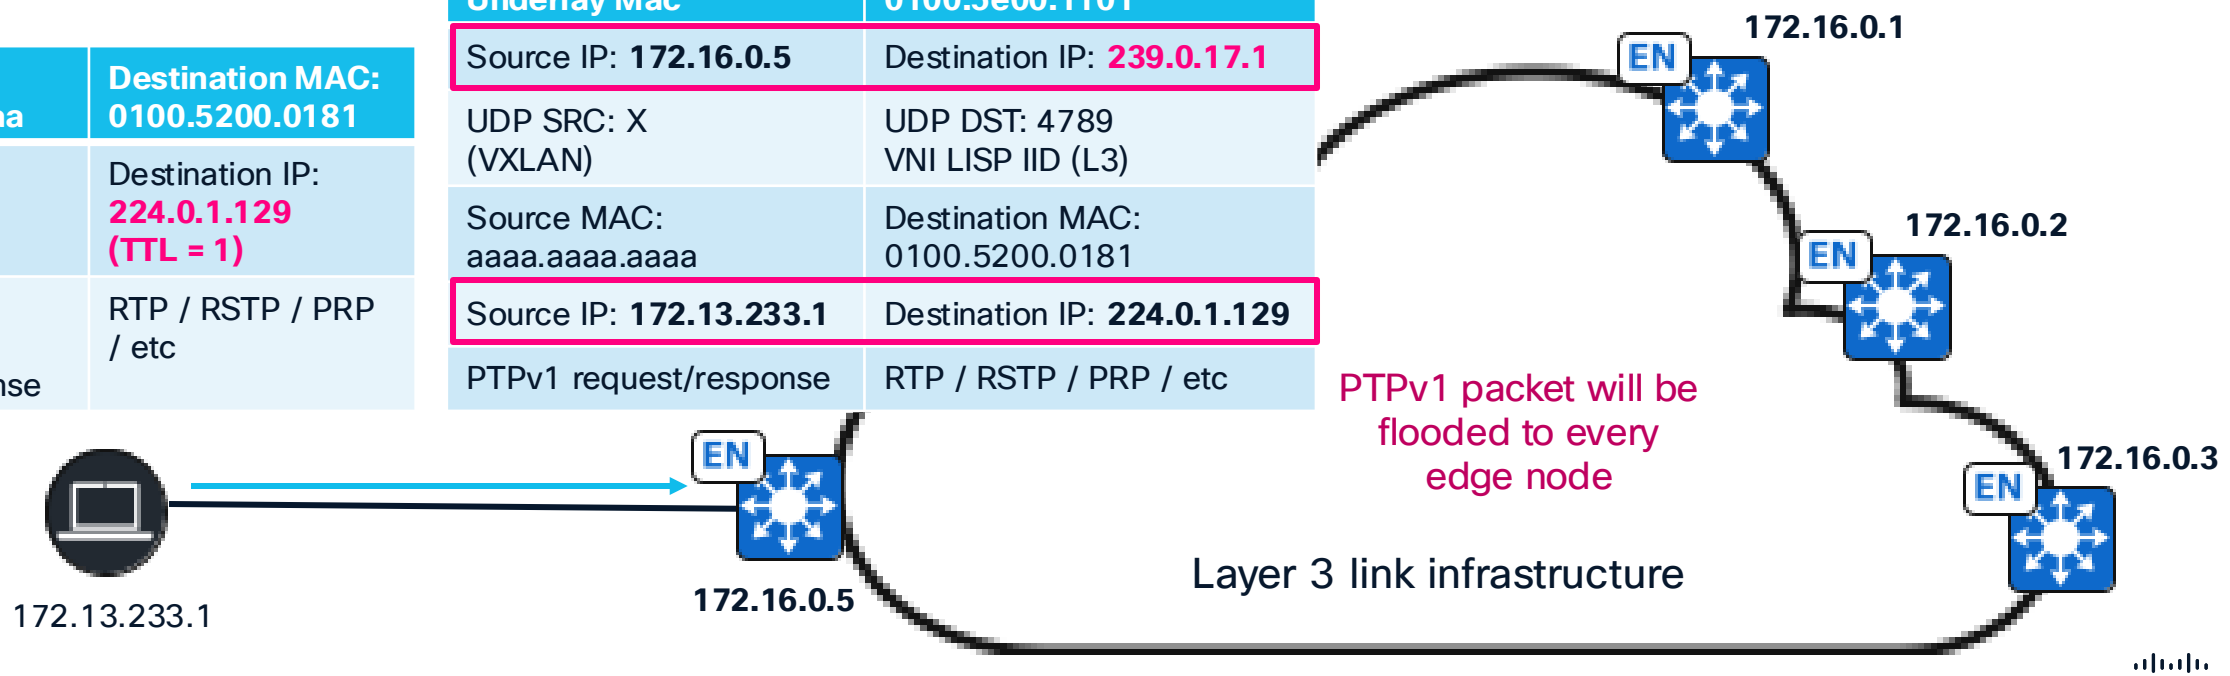

B-U-M traffic

B-U-M Stands for Broadcast, Unknown Unicast and Multicast.

Broadcast is any frame that has a destination mac address of ffff.ffff.ffff

Source MAC: aaaa.aaaa.aaaa	Destination MAC: ffff.ffff.ffff
Source IP: 172.13.233.1	Destination IP: 172.13.233.9
ICMP Payload	

L2 flooding is required to handle this type of Traffic

Some frames with ethernet destination of ffff.ffff.ffff:

- ARP request (L2 flooding enabled vlan)
- Gratuitous ARP
- DHCP packets (in L2 only vlan)
- IP directed broadcast

Unknown Unicast is a unicast packet with a mac address unknown to the switch

Source MAC: aaaa.aaaa.aaaa	Destination MAC: aaaa.bbbb.cccc
Source IP: 172.13.233.1	Destination IP: 172.13.233.9
ICMP Payload	

Destination MAC aaaa.bbbb.cccc is unknown to the switch

L2 flooding is required to handle this type of Traffic

Unknown unicast MAC addresses can be anything except:

- ffff.ffff.ffff broadcast mac address
- Any MAC address with the broadcast/multicast bit set (01xx.xxxx.xxxx or 3333.xxxx.xxx)

B-U-M traffic

Multicast traffic (Link local) is any packet within the reserved range of 224.0.0.0/24

Source MAC: aaaa.aaaa.aaaa	Destination MAC: 0100.5e00.000D
Source IP: 172.13.233.1	Destination IP: 224.0.0.13
ICMP Payload	

L2 flooding is required to handle this type of Traffic

Some link local traffic is:

224.0.0.5 = 0100.5E00.0005 – All OSPF routers

224.0.0.13 = 0100.5E00.000D – PIM

224.0.0.233 = 0100.5E00.00E9 – Dante discovery

224.0.0.251 = 0100.5E00.00fb – mDNS

L2 only multicast = 0160.2BFF.FF00 – Cobranet

Multicast traffic (Non-link local) is everything within the private range 224.0.1.0 to 239.255.255.255

Source MAC: aaaa.aaaa.aaaa	Destination MAC: 0100.5e7f.fffa
Source IP: 172.13.233.1	Destination IP: 239.255.255.250
ICMP Payload	

Overlay multicast (of any flavor) can handle this type of Traffic

Some non-link local traffic is:

224.0.1.129 – 0100.5e00.0181 – PTP (Based on TTL value)

239.255.255.250 – 0100.5e7f.fffa – SSDP

Deployment

- L2 flooding can be deployed using Cisco Catalyst Center GUI, go to **Provision > Fabric sites > Select_site > Anycast gateways** in here select the anycast gateway on which you want to enable layer 2 flooding and click on edit

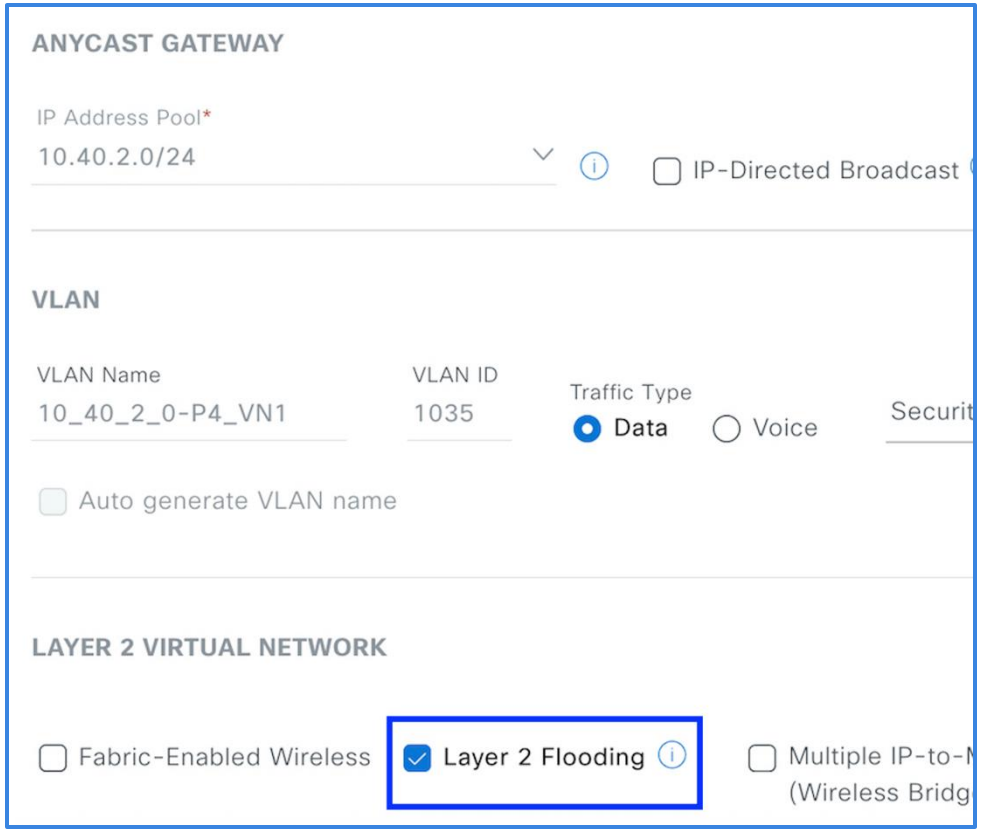

Verification from CLI

```
!
interface L2LISP0.8189
 instance-id 8189
  remote-rloc-probe on-route-change
  service ethernet
   eid-table vlan 1035
   broadcast-underlay 239.0.17.X
   flood arp-nd
   flood unknown-unicast
   database-mapping mac locator-set <rloc_id>
   exit-service-ethernet
  !
  exit-instance-id
```

Top 5 Reasons for Malfunctioning

Reason	Overlay Multicast	L2 Flooding
Reason 1	TTL reducing to zero before reaching to destinations	Not enabling multicast in the underlay
Reason 2	If we have an external RP, then make sure to configure MSDP between internal and external RP's	Misconfigurations (Not enabling underlay multicast on intermediate nodes, PIM on interfaces, not enabling MSDP, configuring wrong RP IP on L3 nodes)
Reason 3	L2-only Vlan is a bad selection for high dense multicast applications like IPTV	ARP source silent hosts with no device-tracking entry, can cause RPF to point to wrong interface on FHR
Reason 4	Not enabling multicast over SD-Access transit on border nodes (Multicast in SDA transit networks)	For L2 flooding use /24 or smaller address pool to limit the number of broadcasts
Reason 5	With overlay multicast enabled, SSDP traffic can cause high CPU utilizations on the RP's	With L2 flooding enabled SSDP traffic gets flooded and causes high bandwidth issues

Underlay Multicast

Deployment

- Layer 2 flooding as well as Native multicast require underlay multicast to be configured on every node that is in the data path.
- Underlay Multicast can only be configured by LAN Automation or manually configured using CLI.

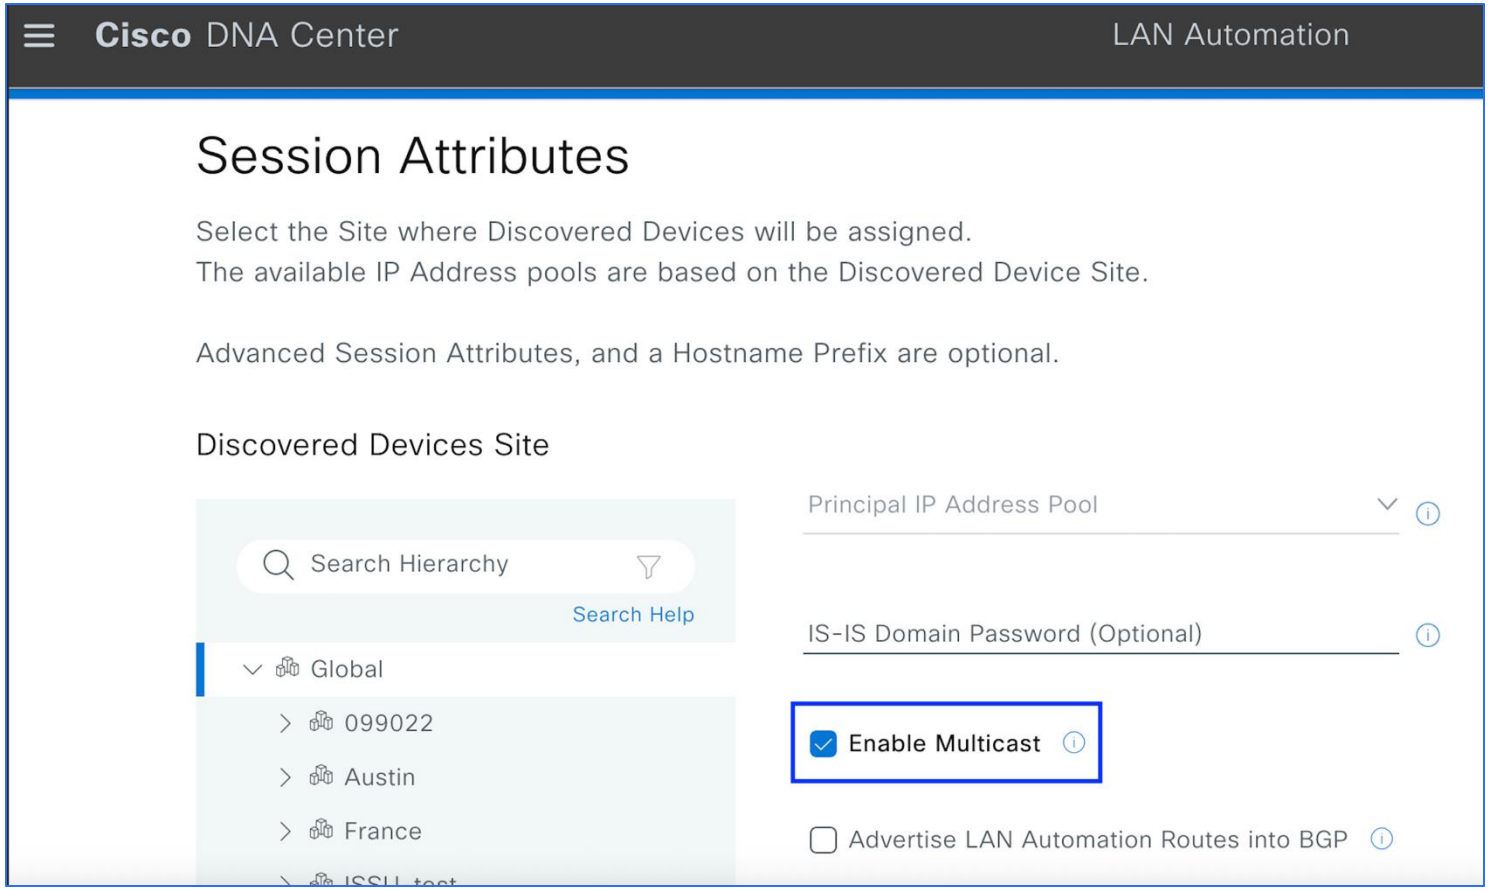

Deployment

- Underlay multicast configuration using CLI is as follows:
- This CLI configurations can also be deployed as CLI templates using Catalyst Center for scalability.

Designated Anycast RPs (Rendezvous Point)

```
interface loopback60000 <<< Create a Loopback60000 to configure RP
ip address x.x.x.x x.x.x.x <<< Configure RP IP address
ip router isis (or the equivalent in other IGP)
ip pim sparse-mode <<< Enable ip pim sparse-mode

If we have Anycast RP setup, then we need to configure msdp for
redundancy
ip msdp peer (Loopback0 of the other RP) connect-source loopback 0
ip msdp originator-id Loopback0
ip msdp cache-sa-state
```

All Layer 3 Nodes

```
IP multicast-routing <<< Enables Multicast Routing

Interface (All IGP/L3 interfaces part of the underlay)
ip pim sparse-mode <<< Enable ip pim sparse-mode
ip router isis

Interface Loopback0 (RLOC (Routing Locator))of Fabric devices)
ip pim sparse-mode <<< Enable ip pim sparse-mode
ip router isis (or the equivalent in other IGP)

ip pim register-source Loopback0 <<< Source it from reachable/Unique interface
ip pim rp-address x.x.x.x (The same RP for all nodes)
```

L2 Flooding Vs Overlay Multicast

Table of Comparison

The main difference between overlay multicast and layer 2 flooding is that overlay multicast routes the packet, while layer 2 flooding switches/floods the packet.

Aspect	L2 Flooding	Overlay Multicast
Definition	Forwards Layer 2 Frames (Broadcast, Unknown unicast, multicast) by flooding to all ports in a vlan.	Delivers IP multicast traffic by routing to specific receivers in a virtual network.
Traffic types	Broadcast, Unknown unicast, multicast (BUM traffic)	IP multicast traffic (Non link Local)
Configuration	Enabled per VLAN using Cisco Catalyst Center	Enabled per VN for specific multicast groups using Cisco Catalyst Center
Dependency	Underlay multicast is needed	Optimal with underlay multicast; falls back to headend replication
Control plane	LISP for known endpoint resolution; floods for unresolved or broadcast traffic	LISP + multicast protocols (PIM, IGMP) for precise delivery

Common Application Considerations

BACnet

- BACnet will use only Layer 2 flooding because it is a protocol that uses broadcast mac address ffff.ffff.ffff to deliver the "who-is" packet to discover other BACnet devices on the network

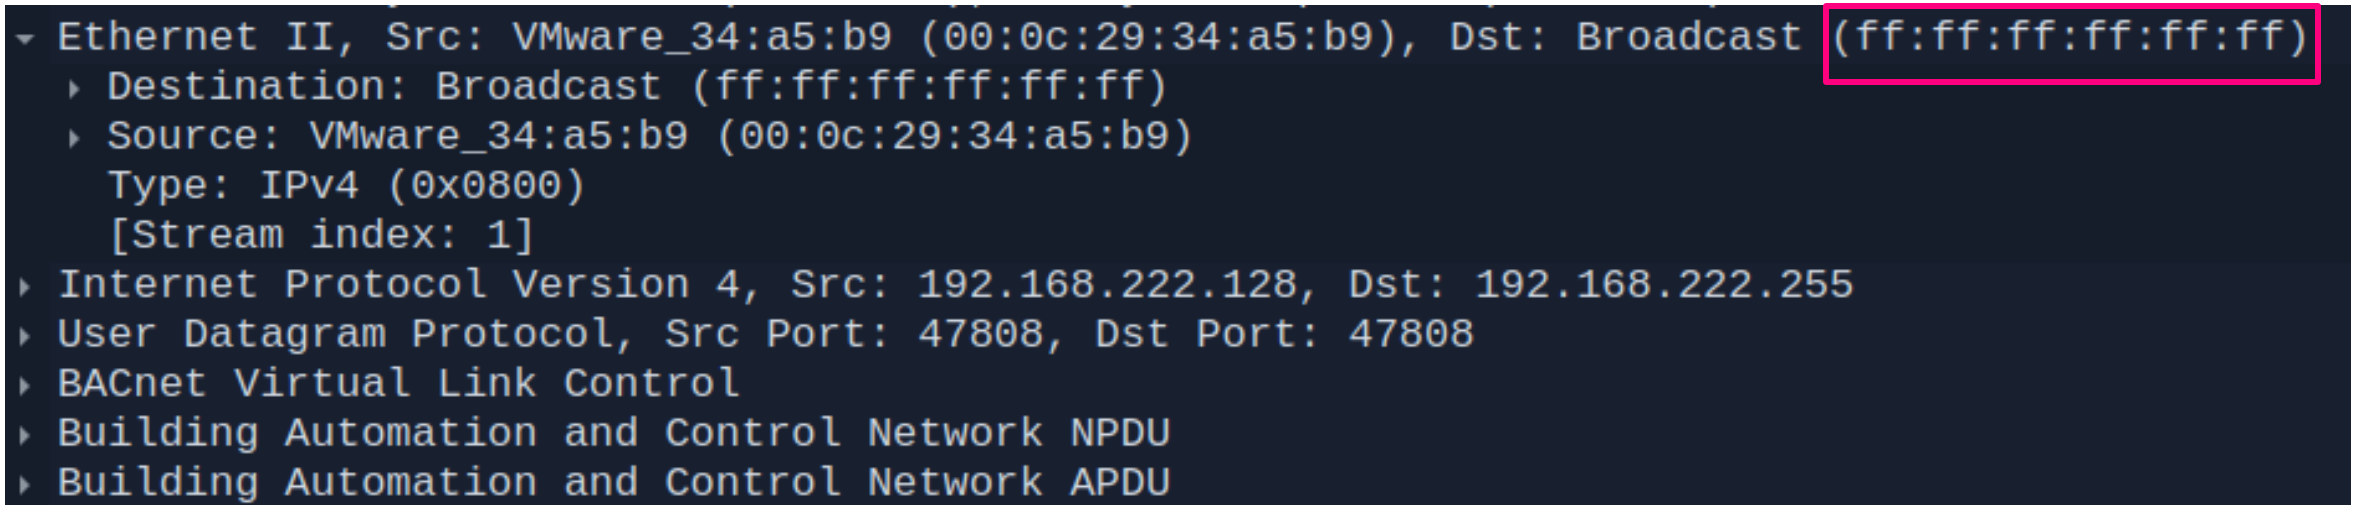

Destination MAC for the packet of ffff.ffff.ffff which is classified as BUM traffic, a type of traffic that Layer 2 flooding can handle.

CobraNET

CobraNET does not use IP headers, it uses only an ethernet header where the destination MAC is a multicast MAC address in the reserve range:

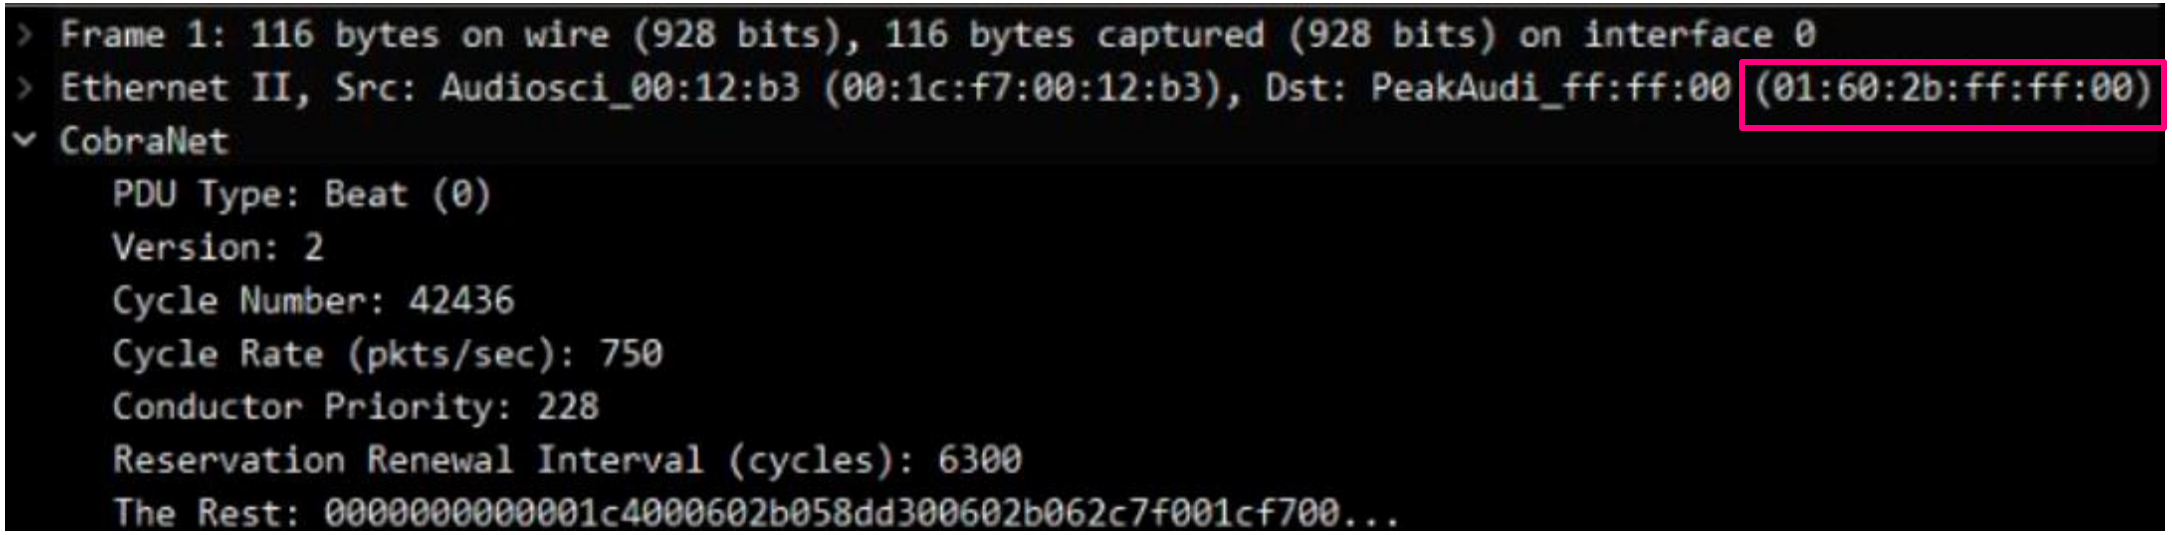

The destination multicast MAC address 01:60:2b:ff:ff:00 is a multicast MAC address that classifies as BUM traffic, this packet using a multicast L2 MAC address will be handle by Layer 2 flooding.

mDNS

- mDNS is a protocol that uses Link-Local multicast address 224.0.0.251 on reserved range, and the destination MAC is a Multicast MAC (01:00:5E:00:00:FB) in reserved range.

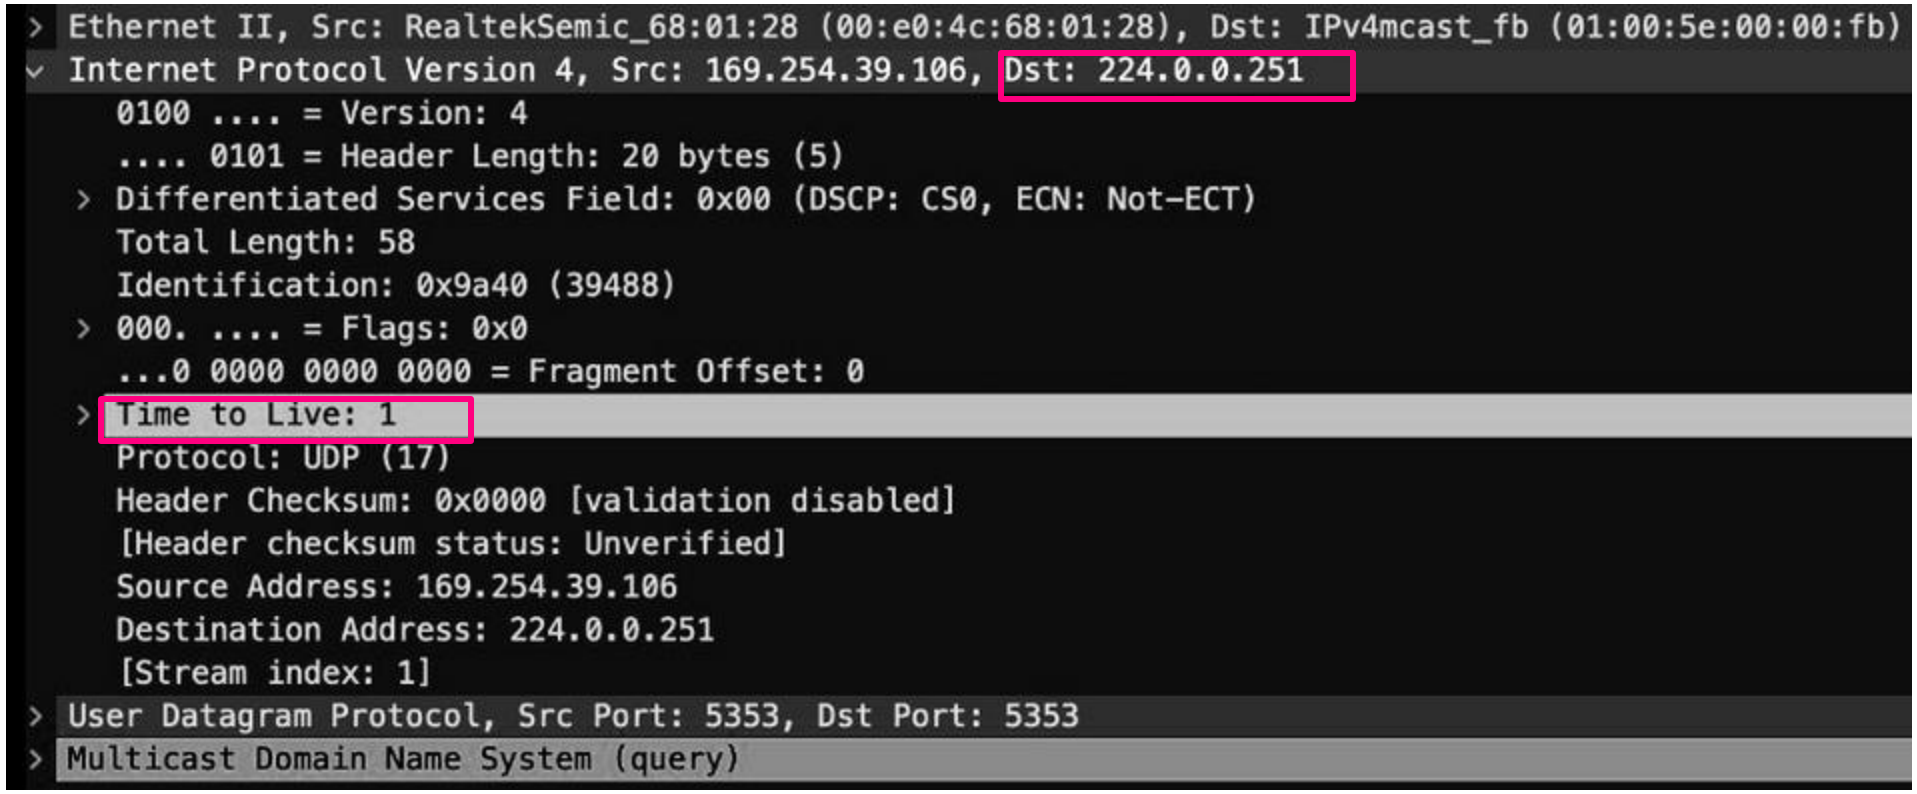

The destination MAC address 01:00:5e:00:00:fb is a multicast MAC address that classifies as BUM traffic, and based on TTL of this packet, this type of traffic will be handle by Layer 2 flooding.

PTPv1

- PTPv1 requires the use of Layer 2 flooding because of how the packets are constructed, for example, let's look at a "delay_response" message sent by a master clock:

```
▸ Ethernet II, Src: FujitsuS_1d:1e:27 (00:30:05:1d:1e:27), Dst: IPv4mcast_01:81 (01:00:5e:00:01:81)
▾ Internet Protocol Version 4, Src: 10.10.100.5, Dst: 224.0.1.129
    0100 .... = Version: 4
    .... 0101 = Header Length: 20 bytes (5)
  ▸ Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
    Total Length: 88
    Identification: 0x005e (94)
  ▸ 010. .... = Flags: 0x2, Don't fragment
    ...0 0000 0000 0000 = Fragment Offset: 0
    Time to Live: 1
    Protocol: UDP (17)
    Header Checksum: 0x29a7 [validation disabled]
    [Header checksum status: Unverified]
    Source Address: 10.10.100.5
    Destination Address: 224.0.1.129
    [Stream index: 0]
▸ User Datagram Protocol, Src Port: 320, Dst Port: 320
▸ Precision Time Protocol (IEEE1588)
```

Based on the TTL that the packets uses for PTPv1, you will need to use Layer 2 flooding to deliver this packet across the network

PTPv2

- PTPv2 requires the use overlay multicast (any flavor) because the TTL for this application can be greater than 1, let's examine the same "delay_response" packet sent by the master clock:

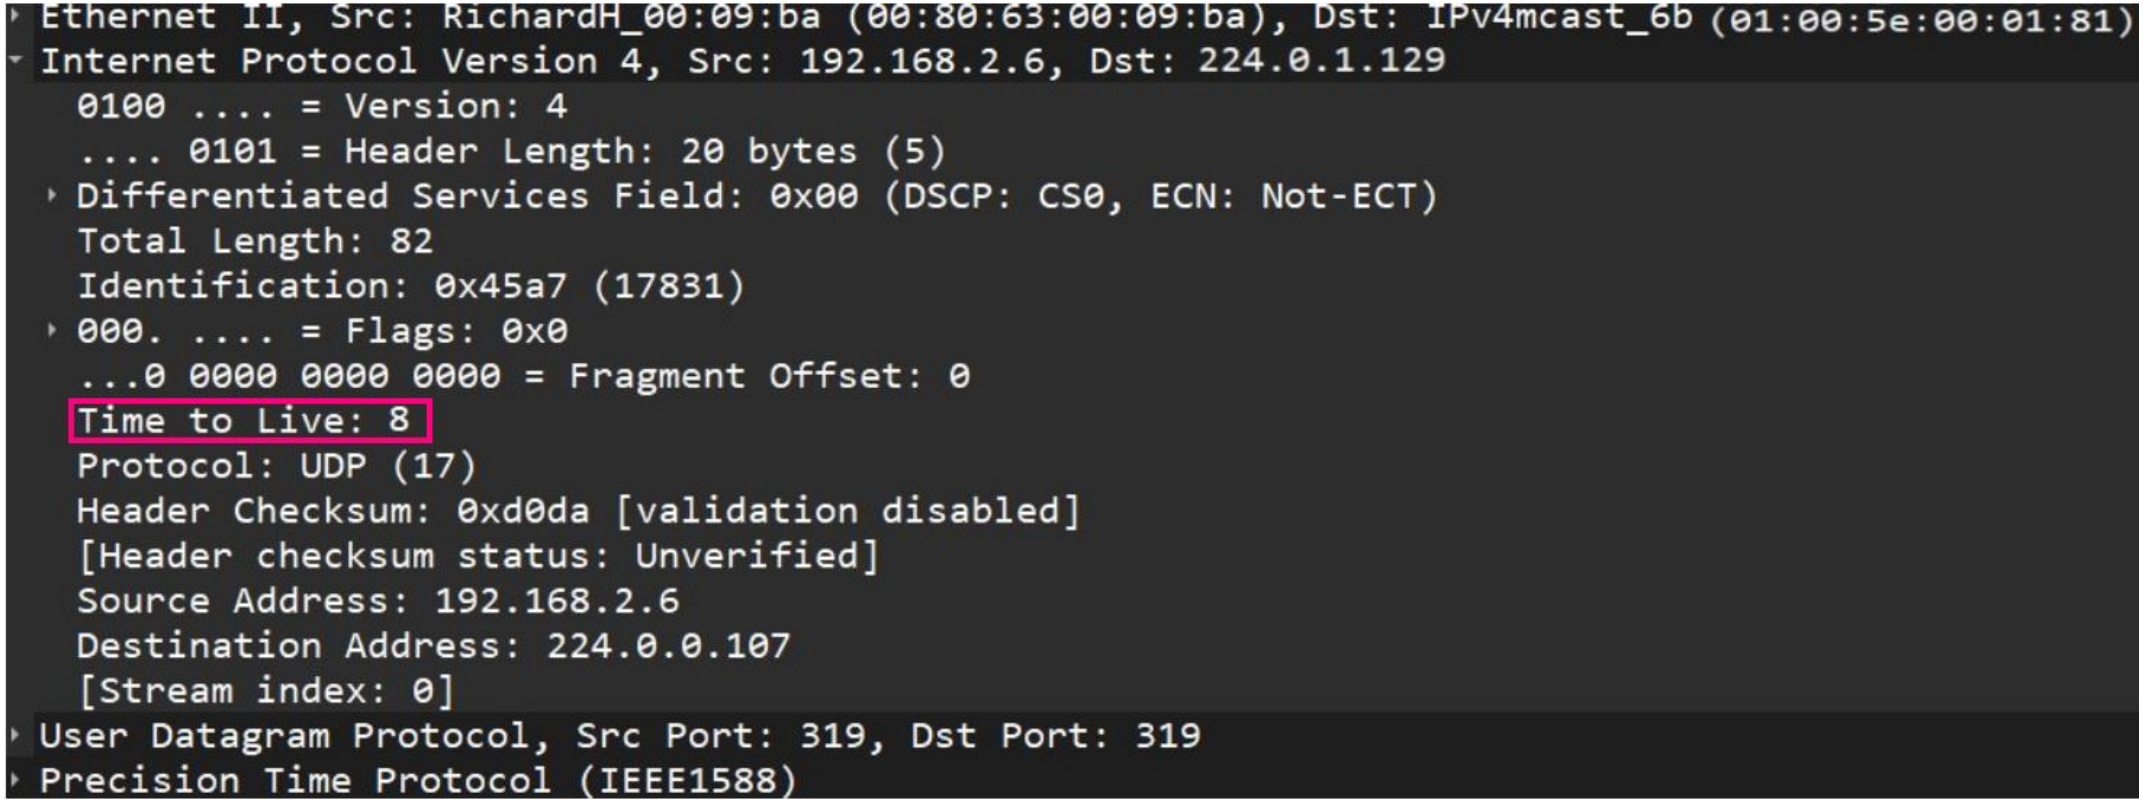

Based on the TTL that the packets for PTPv2 can use you will need to use either; overlay multicast or Layer 2 flooding if TTL=1 as we previously observed on PTPv1

Dante

- Uses Precision time protocol (PTP)
- PTPv1 or PTPv2 is used depending on the capabilities of the Dante devices in the network.
- PTPv1 is used for mixed networks where some devices do not support PTPv2.
- PTPv2 is used in networks where all devices can support it.
- **Dante discovery** is a packet that uses the reserved link local multicast range 224.0.0.233/24
- If Dante discovery is used along with PTPv2 Layer 2 flooding and overlay multicast will be required.

IPTV

IPTV multicast is a method of delivering video content from a source to multiple receivers on a network, IPTV will use only overlay multicast because of the TTL it utilizes on its packets, which is greater than 1:

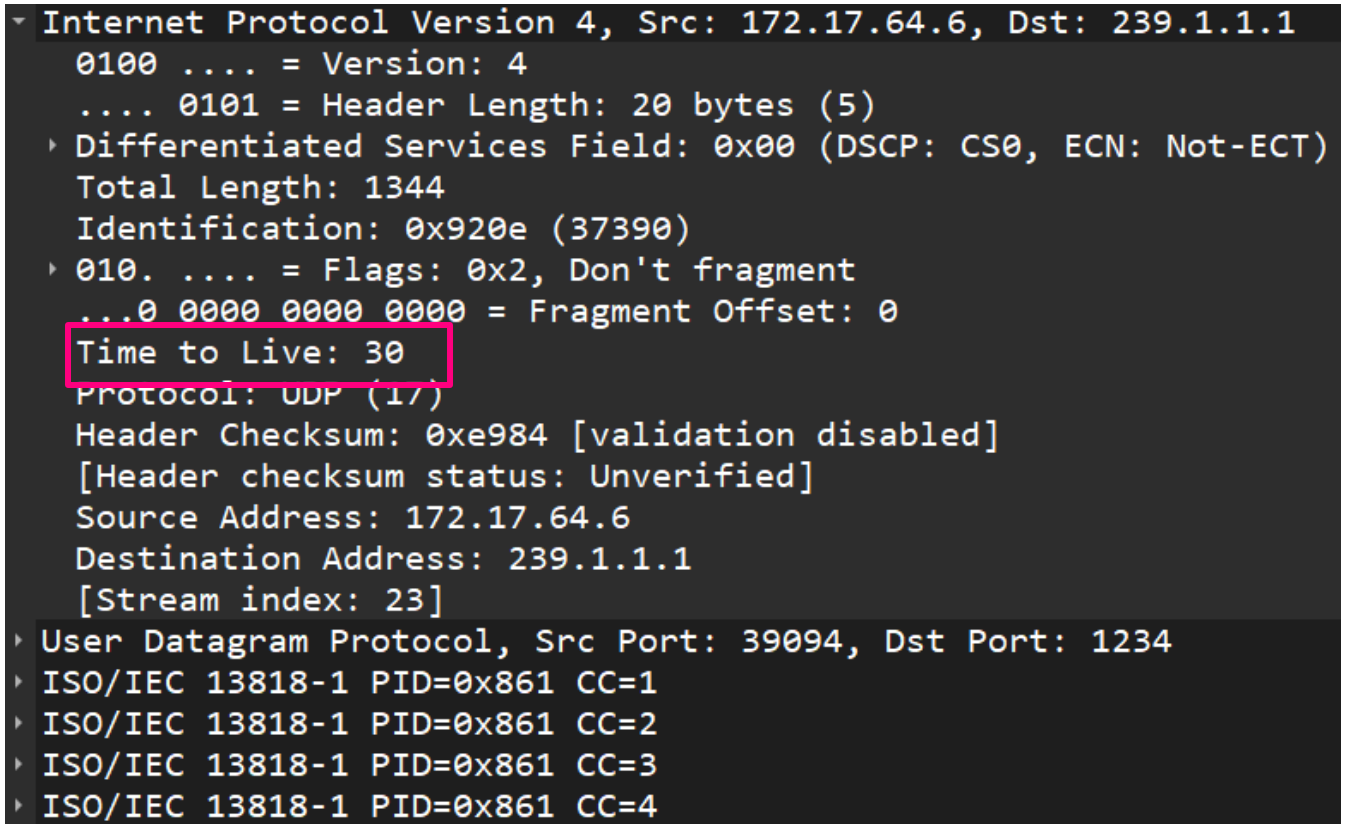

Destination IP for the packet is 239.1.1.1 and TTL=30, therefore it will be handled by overlay multicast (any flavor)

Summary Table

Technologies	Industry	Packet type	Traffic type	Solution
BACnet	IoT, Manufacturing & Building Automation Industries (HVAC, Security, Access Control)	IP broadcast	Destination MAC is broadcast (FF:FF:FF:FF:FF:FF) Classified as BUM traffic	L2 Flooding
CobraNet	Professional Audio Industry (Digital audio networking and control systems)	No IP header; L2 Multicast on reserved range	Destination MAC is multicast MAC (01:60:2B:FF:FF:00) in reserved range Classified as BUM traffic	L2 Flooding
mDNS	Networking/IT Industry Service discovery, conference rooms and screen sharing	Link Local Multicast on reserved range	Destination MAC is multicast MAC (01:00:5E:00:00:FB) in reserved range (224.0.0.251) Classified as BUM traffic	L2 Flooding
Dante	Audio/Video devices (Digital media networking)	PTPv1 uses TTL=1 PTPv2 uses TTL>1	Uses reserved range 224.0.0.233 for Dante discovery; Requires L2 flooding Uses multicast range 224.0.1.129 to 224.0.1.132 for PTPv2; Requires Overlay multicast	L2 Flooding & Overlay Multicast
IPTV	Entertainment Industry, Broadcasting/Telecommunications (TV over internet)	Uses TTL>1	Uses Multicast range 224.0.1.0 – 239.255.255.255	Overlay Multicast

Q&A

Complete Your Session Evaluations

Complete a minimum of 4 session surveys and the Overall Event Survey to be entered in a drawing to win 1 of 5 full conference passes to Cisco Live 2026.

Earn 100 points per survey completed and compete on the Cisco Live Challenge leaderboard.

Level up and earn exclusive prizes!

Complete your surveys in the Cisco Live mobile app.

Continue Your Education

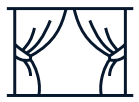

Visit the Cisco Showcase for related demos

Book your one-on-one Meet the Engineer meeting

Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs

Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

Contact us at: ilagunes@cisco.com, vipotu@cisco.com

Thank you

cisco Live !


CISCO