

# Troubleshoot DHCP in Layer 2 Only VLAN - Wired

## Contents

---

### [Introduction](#)

### [Prerequisites](#)

[Requirements](#)

### [L2 Only Overview](#)

[Overview](#)

[DHCP Behavior Change in L2 Only VLANs](#)

[Underlay Multicast](#)

[DHCP Server Inside the SD-Access Fabric](#)

### [Topology](#)

### [L2 Only VLAN Configuration](#)

[L2 Only VLAN Deployment from Catalyst Center](#)

### [L2 Only VLAN Configuration - Fabric Edges](#)

### [L2 Hand-off Configuration - Fabric Border](#)

### [DHCP Traffic Flow](#)

[DHCP Discover and Request - Edge](#)

[MAC Learning and Endpoint Registration](#)

[DHCP Broadcast Bridged in L2 Flooding](#)

[Packet Captures](#)

[DHCP Discover and Request - L2 Border](#)

[Packet Captures](#)

[DHCP Offer and ACK - Broadcast - L2 Border](#)

[MAC Learning and Gateway Registration](#)

[DHCP Broadcast Bridged in L2 Flooding](#)

[DHCP Offer and ACK - Broadcast - Edge](#)

[DHCP Offer and ACK - Unicast - L2 Border](#)

[DHCP Offer and ACK - Unicast - Edge](#)

---

## Introduction

This document describes how to troubleshoot DHCP for wired endpoints in a Layer-2 Only network in SD-Access (SDA) fabric.

## Prerequisites

### Requirements

Cisco recommends that you have knowledge of these topics:

- Internet Protocol (IP) Forwarding

- Locator/ID Separation Protocol (LISP)
- Protocol Independent Multicast (PIM) Sparse-Mode

## Hardware & Software Requirements

- Catalyst 9000 series switches
- Catalyst Center Version 2.3.7.9
- Cisco IOS® XE 17.12 and later

## Limitations

- Only one L2 Border can handoff a unique VLAN/VNI concurrently, unless robust loop prevention mechanisms, such as FlexLink+ or EEM scripts to disable links, are properly configured.

# L2 Only Overview

## Overview

In typical SD-Access deployments, the L2/L3 boundary resides at the Fabric Edge (FE), where the FE hosts the client's gateway in the form of an SVI, which is often called "Anycast Gateway". L3 VNIs (Routed) are established for inter-subnet traffic, while L2 VNIs (Switched) manage intra-subnet traffic. Consistent configuration across all FEs enables seamless client roaming. Forwarding is optimized: intra-subnet (L2) traffic is directly bridged between FEs, and inter-subnet (L3) traffic is routed either between FEs or between an FE and a Border Node.

For endpoints in SDA Fabrics that require a strict network entry point outside the fabric, the SDA Fabric must provide an L2 channel from the Edge to an external gateway.

This concept is analogous to traditional Ethernet campus deployments where a Layer 2 access network connects to a Layer 3 router. Intra-VLAN traffic remains within the L2 network, while inter-VLAN traffic is routed by the L3 device, often returns to a different VLAN on the L2 network.

Within a LISP context, the Site Control Plane primarily tracks MAC addresses and their corresponding MAC-to-IP bindings, much like traditional ARP entries. L2 VNI/L2-only pools are designed to facilitate registration, resolution, and forwarding exclusively based on these two EID types. Therefore, any LISP-based forwarding in an L2-only environment relies solely on MAC and MAC-to-IP information, it completely disregards IPv4 or IPv6 EIDs. To complement LISP EIDs, L2-only pools heavily depend on flood-and-learn mechanisms, similar to the behavior of traditional switches. Consequently, L2 Flooding becomes a critical component for handling Broadcast, Unknown Unicast, and Multicast (BUM) traffic within this solution, requires the use of Underlay Multicast. Conversely, normal unicast traffic is forwarded using standard LISP forwarding processes, primarily via Map-Caches.

Both Fabric Edges and the "L2 Border" (L2B) maintain L2 VNIs, which map to local VLANs (this mapping is locally device-significant within SDA, allowing different VLANs to map to the same L2 VNI across nodes). In this specific use case, no SVI is configured on these VLANs at these nodes, meaning there is no corresponding L3 VNI.

## DHCP Behavior Change in L2 Only VLANs

In Anycast Gateway pools, DHCP presents a challenge because every Fabric Edge acts as the gateway for its directly connected endpoints, with the same gateway IP across all FEs. To properly identify the original source of a DHCP relayed packet, FEs must insert DHCP Option 82 and its sub-options, including the LISP RLOC information. This is achieved with DHCP Snooping on the client VLAN at the Fabric Edge. DHCP Snooping serves a dual purpose in this context: it facilitates the insertion of Option 82 and, crucially,

prevents the flood of DHCP broadcast packets across the bridge-domain (VLAN/VNI). Even when Layer-2 Flooding is enabled for an Anycast Gateway, DHCP Snooping effectively suppresses the broadcast packet to be forwarded out of the Fabric Edge as a broadcast.

In contrast, a Layer 2 Only VLAN lacks a gateway, which simplifies DHCP source identification. Since packets are not relayed by any Fabric Edges, complex mechanisms for source identification are unnecessary. Without DHCP Snooping on the L2 Only VLAN, the flood-control mechanism for DHCP packets is effectively bypassed. This allows DHCP broadcasts to be forwarded via L2 Flooding to their final destination, which could be a DHCP server directly connected to a Fabric Node or a Layer 3 device that provides DHCP relay functionality.

---



**Warning:** The "Multiple IP to MAC" functionality within an L2 Only pool automatically activates DHCP Snooping in Bridge VM mode, which enforces DHCP flood control. Consequently, this renders the L2 VNI pool incapable to support DHCP for its endpoints.

---

## Underlay Multicast

Given DHCP's heavy reliance on broadcast traffic, Layer 2 flooding must be leveraged to support this protocol. As with any other L2 Flooding-enabled pool, the underlay network must be configured for

multicast traffic, specifically Any-Source-Multicast utilizing PIM Sparse-Mode. While underlay multicast configuration is automated via the LAN Automation workflow, if this step was omitted, additional configuration is required (manual or template).

- Enable IP Multicast Routing on all nodes (Borders, Edges, Intermediate Nodes, etc.).
- Configure PIM Sparse-Mode on the Loopback0 interface of each Border and Edge node.
- Enable PIM Sparse-Mode on each IGP (underlay routing protocol) interface.
- Configure the PIM Rendezvous Point (RP) on all nodes (Borders, Edges, Intermediate Nodes), RP placement on Borders is encouraged.
- Verify PIM Neighbors, PIM RP, and PIM Tunnel status.

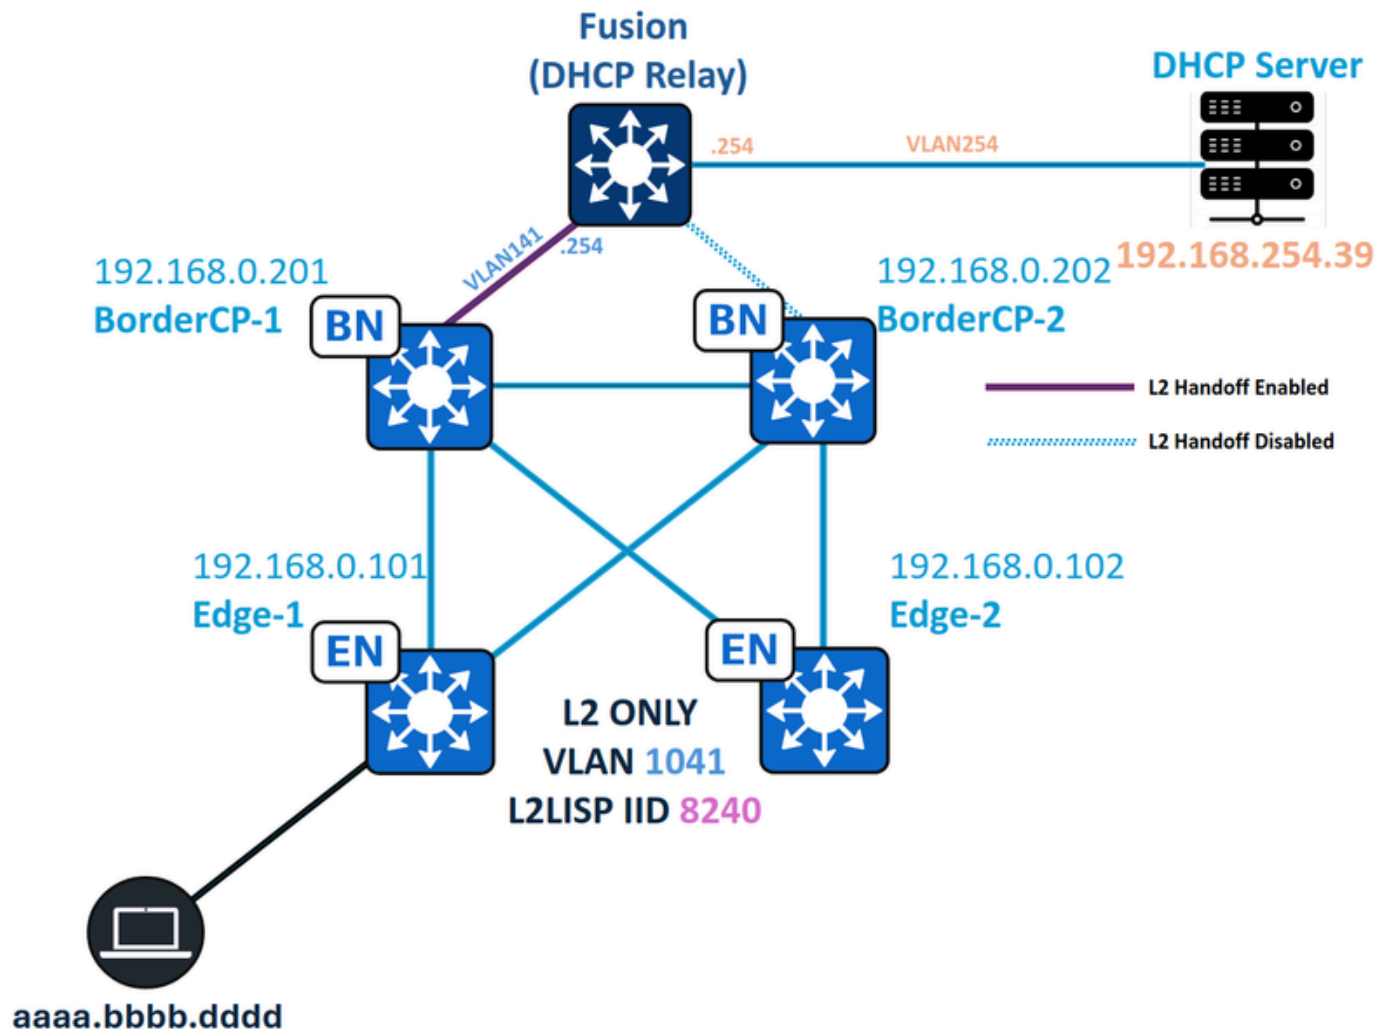
## **DHCP Server Inside the SD-Access Fabric**

A common design question is whether a DHCP Server can be deployed within an SD-Access fabric. The answer, in essence, is both yes and no.

The official [Cisco Validated Design](#) recommends that the DHCP Server should be placed outside the fabric, typically within the Shared Services block. However, if circumstances necessitate the DHCP server's physical attachment to a Fabric Node (e.g., an Edge or Border), the only supported method is via an L2 Only network. This is due to the inherent behavior of Anycast Gateway pools, where DHCP Snooping is enabled by default. This not only blocks DHCP Offers and Acknowledges from the server but also prevents DHCP Discover and Request packets, even when encapsulated in VXLAN, from being forwarded. While "DHCP Snooping Trust" on DHCP server ports allow Offers and Acknowledges, Discover and Request packets are not forwarded using the same method. Furthermore, the removal of DHCP Snooping in an Anycast Gateway pool is not a supported option, as Catalyst Center flags such a configuration deviation during compliance validation.

Conversely, when the DHCP server is placed within an L2 Only network, DHCP Snooping is not enforced, allowing all DHCP packets to pass without policy-based inspection or blockage. The network device upstream of the SD-Access fabric (e.g., a Fusion Router) is configured as the gateway for the L2 Only network, enabling traffic from multiple VRFs to access the same DHCP server within that L2 Only segment.

## **Topology**



Network Topology

In this topology:

- 192.168.0.201 and 192.168.0.202 are Collocated Borders for the Fabric Site, BorderCP-1 is the only Border with the Layer 2 Hand-off feature enabled.
- 192.168.0.101 and 192.168.0.102 are Fabric Edge Nodes
- 192.168.254.39 is the DHCP Server
- aaaa.bbbb.dddd is the DHCP-enabled endpoint
- The Fusion device acts as DHCP Relay for the fabric subnets.

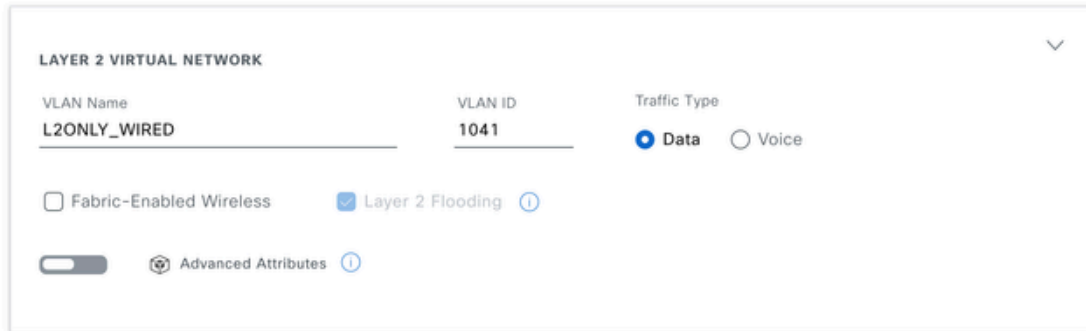
## L2 Only VLAN Configuration

### L2 Only VLAN Deployment from Catalyst Center

Path: Catalyst Center / Provision / Fabric Site / Layer 2 Virtual Networks / Edit Layer 2 Virtual Networks

## Configuration Attributes

Provide a name for each Layer 2 Virtual Network and define its attributes.



**LAYER 2 VIRTUAL NETWORK**

VLAN Name: L2ONLY\_WIRED VLAN ID: 1041 Traffic Type: ☒ Data ☐ Voice

☐ Fabric-Enabled Wireless ☒ Layer 2 Flooding ⓘ

☐ Advanced Attributes ⓘ

### L2VNI Configuration

## L2 Only VLAN Configuration - Fabric Edges

Fabric Edge nodes have the VLAN configured with CTS enabled, IGMP and IPv6 MLD disabled, and the required L2 LISP configuration. This L2 Only pool is **not** a Wireless pool; therefore, features typically found in L2 Only Wireless Pools, such as RA-Guard, DHCPGuard, and Flood Access Tunnel, are not configured. Instead, the flooding of ARP packets is explicitly enabled with "**flood arp-nd**"

### Fabric Edge (192.168.0.101) Configuration

<#root>

```
cts role-based enforcement vlan-list
1041
```

```
vlan
1041
```

```
name L2ONLY_WIRED
```

```
no ip igmp snooping vlan 1041 querier
```

```
no ip igmp snooping vlan 1041
```

```
no ipv6 mld snooping vlan 1041
```

```
router lisp
```

```
instance-id
8240

remote-rloc-probe on-route-change
service ethernet

eid-table vlan

1041

broadcast-underlay
239.0.17.1

flood arp-nd

flood unknown-unicast

database-mapping mac locator-set rloc_91947dad-3621-42bd-ab6b-379ecebb5a2b
exit-service-ethernet
```

## L2 Hand-off Configuration - Fabric Border

From an operational perspective, the DHCP server (or Router/Relay) is allowed to be connected to any Fabric Node, including both Borders and Edges.

Using Border nodes to connect the DHCP server is the recommended approach, however, requires careful design consideration. This is because the Border must be configured for L2 Hand-Off on a per-interface basis. This allows the Fabric Pool to be handed off to either the same VLAN as within the Fabric or a different one. This flexibility in VLAN IDs between Fabric Edges and Borders is possible because both are mapped to the same L2 LISP Instance-ID. L2 Hand-off physical ports must not be simultaneously enabled with the same VLAN to prevent Layer 2 loops within the SD-Access network. For redundancy, methods such as StackWise Virtual, FlexLink+, or EEM scripts are required.

In contrast, connecting the DHCP Server or Gateway Router to a Fabric Edge requires no additional configuration.

**BorderCP-1.DNA2.local**

Interface: TenGigabitEthernet1/0/44

VLAN Name	Enable Layer-2 Handoff	External VLAN
L2_Only_Wireless	<input checked="" type="checkbox"/>	31
L2_Only_Wireless_2	<input checked="" type="checkbox"/>	40
L2ONLY_WIRED	<input checked="" type="checkbox"/>	141

3 Record(s) | Show Records: 25 | 1 - 3

*L2 Hand-off Configuration*

## Fabric Border (192.168.0.201) Configuration

<#root>

```
cts role-based enforcement vlan-list
```

```
141
```

```
vlan
```

```
141
```

```
name L2ONLY_WIRED
```

```
no ip igmp snooping vlan 141 querier
```

```
no ip igmp snooping vlan 141
```

```
no ipv6 mld snooping vlan 141
```

```
router lisp
instance-id
```

```
8240
```



```
remote-rloc-probe on-route-change  
service ethernet
```

```
eid-table
```

```
vlan 141
```

```
broadcast-underlay 239.0.17.1
```

```
flood arp-nd
```

```
flood unknown-unicast
```

```
database-mapping mac locator-set rloc_91947dad-3621-42bd-ab6b-379ecebb5a2b  
exit-service-ethernet
```

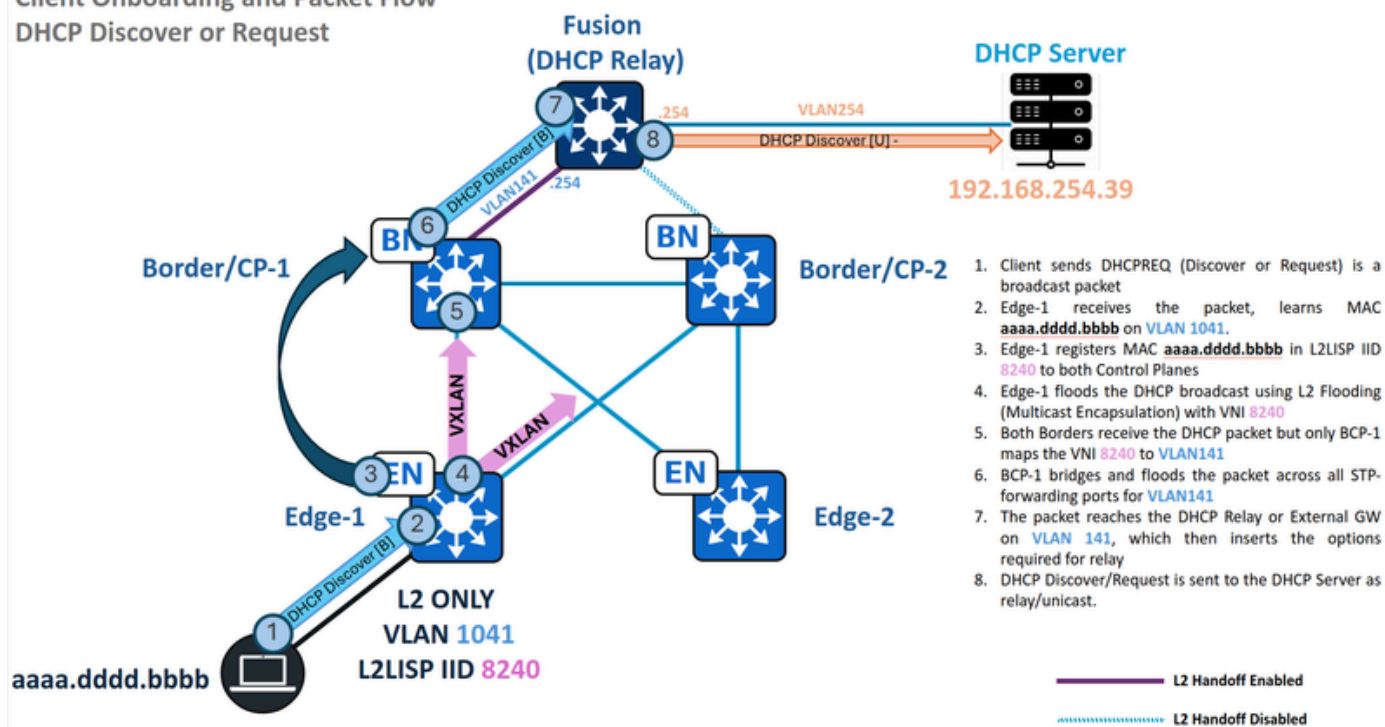
```
interface TenGigabitEthernet1/0/44
```

```
switchport mode trunk
```

## **DHCP Traffic Flow**

### **DHCP Discover and Request - Edge**

## Client Onboarding and Packet Flow DHCP Discover or Request



Traffic Flow - DHCP Discover and Request in L2 Only

## MAC Learning and Endpoint Registration

When endpoint **aaaa.dddd.bbbb** sends a DHCP Discover or Request (a broadcast packet), the Edge node must learn the endpoint's MAC address, add it to its MAC address table, then to the L2/MAC SISF table, and finally to the L2LISP Database for **VLAN 1041**, mapped to L2LISP Instance **8240**.

<#root>

Edge-1#

```
show mac address-table interface te1/0/2
```

Mac Address Table			
Vlan	Mac Address	Type	Ports
-----	-----	-----	-----
1041			
	aaaa.dddd.bbbb	DYNAMIC	Te1/0/2

Edge-1#

```
show vlan id 1041
```

VLAN Name	Status	Ports
-----------	--------	-------

-----

1041 L2ONLY\_WIRED

active

L2LI0:

8240

, Te1/0/2, Te1/0/17, Te1/0/18, Te1/0/19, Te1/0/20, Ac2, Po1

Edge-1#

show device-tracking database mac | i aaaa.ddd.bbb|vlan

MAC	Interface	vlan	prlv	state	Time left	Policy
aaaa.ddd.bbb	Te1/0/2	1041	NO TRUST	MAC-REACHABLE	123 s	IPDT_POLICY

Edge-1#

show lisp instance-id 8240 dynamic-eid summary | i Name|aaaa.ddd.bbb

Dyn-EID Name	Dynamic-EID	Interface	Uptime	Last	Pending
Auto-L2-group-					
8240					

aaaa.ddd.bbb

N/A	6d04h	never
-----	-------	-------

0

Edge-1#

show lisp instance-id 8240 ethernet database aaaa.ddd.bbb

LISP ETR MAC Mapping Database for LISP 0 EID-table

Vlan 1041 (IID 8240)

, LSBs: 0x1

Entries total 1, no-route 0, inactive 0, do-not-register 0

aaaa.ddd.bbb/48

,

dynamic-eid Auto-L2-group-8240

, inherited from default locator-set rloc\_91947dad-3621-42bd-ab6b-379ecebb5a2b

Uptime: 6d04h, Last-change: 6d04h

Domain-ID: local

Service-Insertion: N/A

Locator	Pri/Wgt	Source	State
---------	---------	--------	-------

192.168.0.101

10/10	cfg-intf	site-self,	reachable
-------	----------	------------	-----------

Map-server	Uptime	ACK	Domain-ID
192.168.0.201	6d04h	Yes	0
192.168.0.202	6d04h	Yes	0

If the MAC address of the endpoint is correctly learned and the ACK flag has been marked as "Yes" for the Fabric Control planes, this stage is considered completed.

## DHCP Broadcast Bridged in L2 Flooding

When DHCP Snooping is disabled, DHCP Broadcasts are not blocked; instead, they are encapsulated in multicast for Layer 2 Flooding. Conversely, enabling DHCP Snooping prevents the flooding of these broadcast packets.

<#root>

Edge-1#

```
show ip dhcp snooping
```

```
Switch DHCP snooping is enabled
```

```
Switch DHCP gleaning is disabled
```

```
DHCP snooping is configured on following VLANs:
```

```
12-13,50,52-53,333,1021-1026
```

```
DHCP snooping is operational on following VLANs:
```

```
12-13,50,52-53,333,1021-1026
```

<--

VLAN1041 should not be listed, as DHCP snooping must be disabled in L2 Only pools.

```
Proxy bridge is configured on following VLANs:
```

```
1024
```

```
Proxy bridge is operational on following VLANs:
```

```
1024
```

<snip>

Since DHCP Snooping is disabled, the DHCP Discover/Request utilizes the L2LISP0 interface, bridging traffic via L2 Flooding. Depending on the Catalyst Center version and applied Fabric Banners, the L2LISP0 interface may have access-lists configured in both directions; therefore, ensure DHCP traffic (UDP ports 67 and 68) is not explicitly denied by any Access Control Entries (ACEs).

```
<#root>
```

```
interface L2LISP0
```

```
ip access-group
```

```
SDA-FABRIC-LISP
```

```
in
```

```
ip access-group
```

```
SDA-FABRIC-LISP out
```

```
Edge-1#
```

```
show access-list SDA-FABRIC-LISP
```

```
Extended IP access list SDA-FABRIC-LISP
```

```
10 deny ip any host 224.0.0.22
```

```
20 deny ip any host 224.0.0.13
```

```
30 deny ip any host 224.0.0.1
```

```
40 permit ip any any
```

Utilize the configured broadcast-underlay group for the L2LISP instance and the Fabric Edge's Loopback0 IP address to verify the L2 Flooding (S,G) entry that bridges this packet to other Fabric Nodes. Consult the mroute and mfib tables to validate parameters such as the incoming interface, outgoing interface list, and forwarding counters.

```
<#root>
```

```
Edge-1#
```

```
show ip interface loopback 0 | i Internet
```

```
Internet address is
```

```
192.168.0.101/32
```

```
Edge-1#
```

```
show running-config | se 8240
```

```
interface L2LISP0.8240
```

```
instance-id 8240
```

```
remote-rloc-probe on-route-change  
service ethernet
```

```
eid-table vlan 1041
```

```
broadcast-underlay 239.0.17.1
```

```
Edge-1#
```

```
show ip mroute 239.0.17.1 192.168.0.101 | be \((
```

```
(192.168.0.101, 239.0.17.1)
```

```
, 00:00:19/00:03:17, flags: FT  
Incoming interface:
```

```
Null0
```

```
, RPF nbr 0.0.0.0
```

```
<--
```

```
Local S,G IIF must be Null0
```

```
Outgoing interface list:
```

```
TenGigabitEthernet1/1/2
```

```
,
```

```
Forward
```

```
/Sparse, 00:00:19/00:03:10, flags:
```

```
<--
```

```
1st OIF = Tel1/1/2 = Border2 Uplink
```

```
TenGigabitEthernet1/1/1
```

```
,
```

```
Forward
```

```
/Sparse, 00:00:19/00:03:13, flags:
```

```
<--
```

```
2nd OIF = Tel1/1/1 = Border1 Uplink
```

```
Edge-1#
```

```
show ip mfib 239.0.17.1 192.168.0.101 count
```

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second

Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Default

13 routes, 6 (\*,G)s, 3 (\*,G/m)s

Group:

239.0.17.1

Source:

192.168.0.101

,

SW Forwarding: 1/0/392/0, Other: 1/1/0

HW Forwarding:


7

/0/231/0, Other: 0/0/0

<--

HW Forwarding counters (First counter = Pkt Count) must increase

Totals - Source count: 1, Packet count: 8



**Tip:** If an (S,G) entry is not found or the Outgoing Interface List (OIL) contains no Outgoing Interfaces (OIFs), it indicates an issue with the underlay multicast configuration or operation.

---

## Packet Captures

Configure a simultaneous embedded packet capture on the switch to record both the incoming DHCP packet from the endpoint and the corresponding egress packet for L2 Flooding. Upon packet capture, two distinct packets should be observed: the original DHCP Discover/Request and its VXLAN-encapsulated counterpart, destined for the Underlay Group (239.0.17.1).

### Fabric Edge (192.168.0.101) packet captures

<#root>

```
monitor capture cap interface TenGigabitEthernet1/0/2 IN      <-- Endpoint Interface
```



```
monitor capture cap interface TenGigabitEthernet1/1/1 OUT    <-- One of the OIFs from the multicast route
```

```
monitor capture cap match any
monitor capture cap buffer size 100
monitor capture cap limit pps 1000
monitor capture cap start
monitor capture cap stop
```

Edge-1#

```
show monitor capture cap buffer display-filter "bootp and dhcp.hw.mac_addr==aaaa.ddd.ddd.bbbb"
```

<-- aaaa.ddd.ddd.bbbb is the endpoint MAC

Starting the packet display ..... Press Ctrl + Shift + 6 to exit

```
22  2.486991      0.0.0.0 -> 255.255.255.255 DHCP
```

356 DHCP Discover

- Transaction ID 0xf8e

<--

356 is the Length of the original packet

```
23  2.487037      0.0.0.0 -> 255.255.255.255 DHCP
```

406 DHCP Discover

- Transaction ID 0xf8e

<--

406 is the Length of the VXLAN encapsulated packet

Edge-1#

```
show monitor capture cap buffer display-filter "bootp and dhcp.hw.mac_addr==aaaa.ddd.ddd.bbbb and vxlan"
```

Starting the packet display ..... Press Ctrl + Shift + 6 to exit

```
23  2.487037      0.0.0.0 -> 255.255.255.255 DHCP
```

406 DHCP Discover

- Transaction ID 0xf8e

Edge-1#

```
show monitor capture cap buffer display-filter "bootp and dhcp.hw.mac_addr==aaaa.ddd.ddd.bbbb and vxlan" de
```

Internet Protocol Version 4, Src:

192.168.0.101, Dst: 239.0.17.1 <-- DHCP Discover is encapsulated for Layer 2 Flooding

Internet Protocol Version 4, Src:  
0.0.0.0, Dst: 255.255.255.255

## DHCP Discover and Request - L2 Border

After the Edge sends the DHCP Discover and Request packets via Layer 2 Flooding, encapsulated with the Broadcast-Underlay group 239.0.17.1, these packets are received by the L2 Hand-off Border, specifically Border/CP-1 in this scenario.

For this to occur, Border/CP-1 must possess a multicast route with the (S,G) of the Edge, and its outgoing interface list must include the L2LISP instance of the L2 Handoff VLAN. It's important to note that L2 Hand-off Borders share the same L2LISP Instance-ID, even if they utilize different VLANs for the Hand-off.

<#root>

BorderCP-1#

show vlan id 141

VLAN Name	Status	Ports
141 L2ONLY_WIRED		

active

L2LI0:

8240

, Te1/0/44

BorderCP-1#

show ip mroute 239.0.17.1 192.168.0.101 | be \((

(192.168.0.101, 239.0.17.1)

, 00:03:20/00:00:48, flags: MTA

Incoming interface:

TenGigabitEthernet1/0/42

, RPF nbr 192.168.98.3

<--

Incoming Interface Te1/0/42 is the RPF interface for 192.168.0.101 (Edge RLOC)

Outgoing interface list:

TenGigabitEthernet1/0/26, Forward/Sparse, 00:03:20/00:03:24, flags:

L2LISP0.

8240

, Forward/Sparse-Dense, 00:03:20/00:02:39, flags:

BorderCP-1#

show ip mfib 239.0.17.1 192.168.0.101 count

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second

Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Default

13 routes, 6 (\*,G)s, 3 (\*,G/m)s

Group:

239.0.17.1

Source:

192.168.0.101

,

SW Forwarding: 1/0/392/0, Other: 0/0/0


HW Forwarding:

3

/0/317/0, Other: 0/0/0

<-- HW Forwarding counters (First counter = Pkt Count) must increase

Totals - Source count: 1, Packet count: 4



**Tip:** If an (S,G) entry is not found, it indicates an issue with the underlay multicast configuration or operation. If the L2LISP for the required instance is not present as OIF, it indicates an issue with the operation UP/DOWN status of the L2LISP sub-interface or the IGMP enablement status of the L2LISP interface.

---

Similar to the Fabric Edge node, ensure no Access Control Entry denies the ingress DHCP packet on the L2LISP0 interface.

<#root>

BorderCP-1#

**show access-list SDA-FABRIC-LISP**

```
Extended IP access list SDA-FABRIC-LISP
 10 deny ip any host 224.0.0.22
```

```
20 deny ip any host 224.0.0.13
30 deny ip any host 224.0.0.1
```

```
40 permit ip any any
```

After the packet is de-encapsulated and placed on the VLAN matching VNI 8240, its broadcast nature dictates that it is flooded out all Spanning Tree Protocol forwarding ports for hand-off VLAN 141.

<#root>

BorderCP-1#

```
show spanning-tree vlan 141 | be Interface
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
-----					
Te1/0/44					
	Desg				
FWD					
2000	128.56	P2p			

The Device-Tracking table confirms that interface Te1/0/44, which connects to the Gateway/DHCP Relay, must be an STP-forwarding port.

<#root>

BorderCP-1#

```
show device-tracking database address 172.16.141.254 | be Network
```

Network Layer Address	Link Layer Address	Interface	vlan	prlv1	age
ARP					
172.16.141.254					
f87b.2003.7fc0					
Te1/0/44					
141					
0005	133s	REACHABLE	112 s	try 0	

## Packet Captures

Configure a simultaneous embedded packet capture on the switch to record both the incoming DHCP packet from L2 Flooding (S,G incoming interface) and the corresponding egress packet to the DHCP Relay. Upon packet capture, two distinct packets should be observed: the VXLAN encapsulated packet from Edge-1, and the de-encapsulated packet that goes to the DHCP Relay.

### Fabric Border/CP (192.168.0.201) packet captures

```
<#root>
```

```
monitor capture cap interface TenGigabitEthernet1/0/42 IN    <-- Incoming interface for Edge's S,G Mrou
```

```
monitor capture cap interface TenGigabitEthernet1/0/44 OUT    <-- Interface that connects to the DHCP Rel
```

```
monitor capture cap match any
```

```
monitor capture cap buffer size 100
```

```
monitor capture cap start
```

```
monitor capture cap stop
```

```
BorderCP-1#
```

```
show monitor capture cap buffer display-filter "bootp and dhcp.hw.mac_addr==aaaa.ddd.ddd"
```

```
Starting the packet display ..... Press Ctrl + Shift + 6 to exit
```

```
427  16.695022      0.0.0.0 -> 255.255.255.255 DHCP
```

```
406
```

```
DHCP Discover - Transaction ID 0x2030
```

```
<-- 406 is the Lenght of the VXLAN encapsulated packet
```

```
428  16.695053      0.0.0.0 -> 255.255.255.255 DHCP
```

```
364
```

```
DHCP Discover - Transaction ID 0x2030
```

```
<-- 364 is the Lenght of the VXLAN encapsulated packet
```

```
Packet 427: VXLAN Encapsulated
```

BorderCP-1#

```
show monitor capture cap buffer display-filter "bootp and dhcp.hw.mac_addr==aaaaa.ddddd.bbbb and vxlan"
```

Internet Protocol Version 4, Src:

192.168.0.101, Dst: 239.0.17.1

Internet Protocol Version 4, Src:

0.0.0.0, Dst: 255.255.255.255

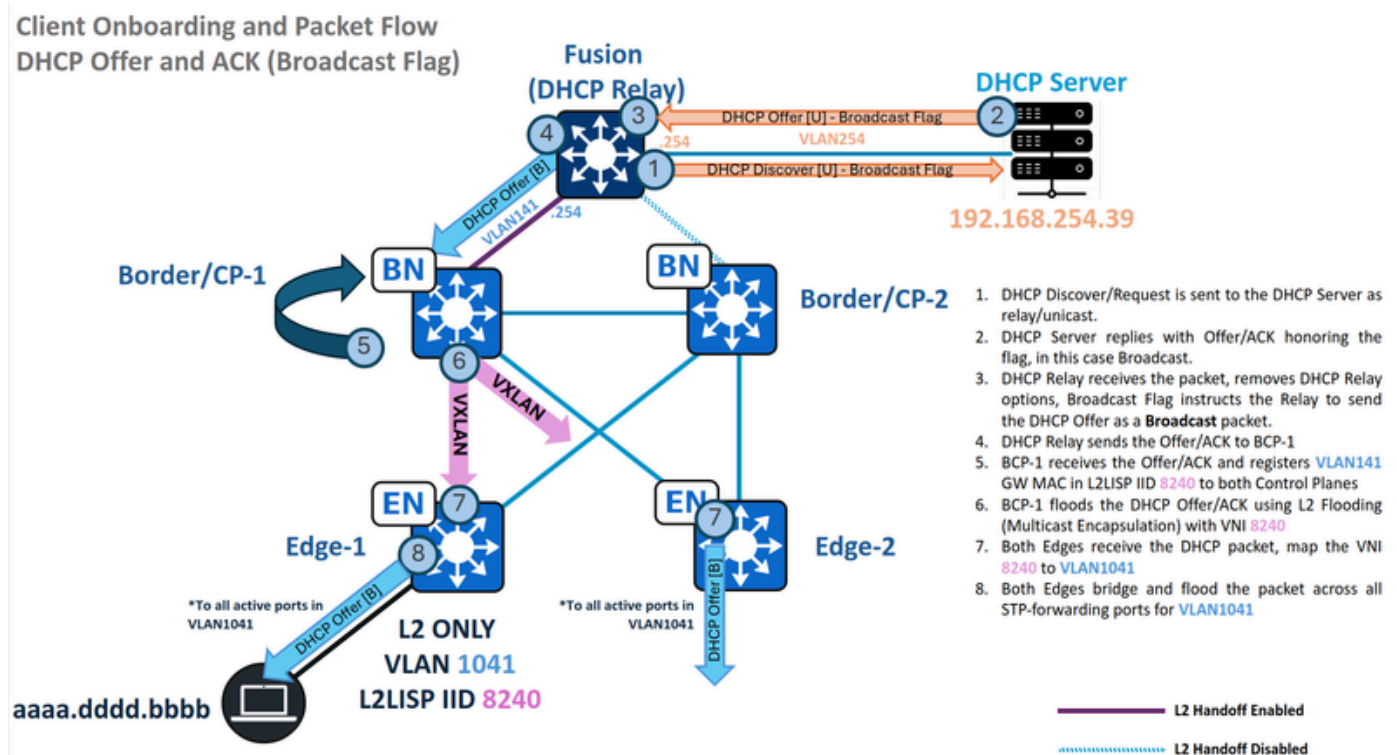
Packet 428: Plain (dot1Q cannot be captured at egress direction)

BorderCP-1#

```
show monitor capture cap buffer display-filter "bootp and dhcp.hw.mac_addr==aaaaa.ddddd.bbbb and not vxlan"
```

Internet Protocol Version 4, Src: 0.0.0.0, Dst: 255.255.255.255

## DHCP Offer and ACK - Broadcast - L2 Border



Traffic Flow - Broadcast DHCP Offer and ACK in L2 Only

Now that the DHCP Discover has exited the SD-Access fabric, the DHCP relay inserts traditional DHCP Relay Options (e.g., GiAddr/GatewayIPAddress) and forwards the packet as a unicast transmission to the DHCP Server. In this flow, the SD-Access fabric does not append any special DHCP options.

Upon the arrival of a DHCP Discover/Request to the server, the server honors the embedded **Broadcast** or

**Unicast** flag. This flag dictates whether the DHCP Relay Agent forwards the DHCP Offer to the downstream device (our Borders) as a broadcast or unicast frame. For this demonstration, a broadcast scenario is assumed.

**MAC Learning and Gateway Registration**

When the DHCP relay sends a DHCP Offer or ACK, the L2BN node must learn the gateway's MAC address, add it to its MAC address table, then to the L2/MAC SISF table, and finally to the L2LISP Database for VLAN **141**, mapped to L2LISP Instance **8240**.

<#root>

BorderCP-1#

show mac address-table interface te1/0/44

Mac Address Table			
Vlan	Mac Address	Type	Ports
141	f87b.2003.7fc0	DYNAMIC	Te1/0/44

BorderCP-1#

show vlan id 141

VLAN Name	Status	Ports
141	L2ONLY_WIRED	active L2LI0:
8240		,
Te1/0/44		

BorderCP-1#

show device-tracking database mac | i 7fc0|vlan



MAC	Interface	vlan	prlv1	state	Time left	Policy
f87b.2003.7fc0						
Tel/0/44 141						
NO TRUST						
MAC-REACHABLE						
61 s	LISP-DT-GLEAN-VLAN 64					

BorderCP-1#

show lisp ins 8240 dynamic-eid summary | i Name|f87b.2003.7fc0

Dyn-EID Name	Dynamic-EID	Interface	Uptime	Last	Pending
Auto-L2-group-8240					
f87b.2003.7fc0					
N/A	6d06h	never			
0					

BorderCP-1#

show lisp instance-id 8240 ethernet database f87b.2003.7fc0

LISP ETR MAC Mapping Database for LISP 0 EID-table Vlan

141

(IID

8240

), LSBs: 0x1

Entries total 1, no-route 0, inactive 0, do-not-register 0

f87b.2003.7fc0/48

, dynamic-eid Auto-L2-group-8240, inherited from default locator-set rloc\_0f43c5d8-f48d-48a5-a5a8-094b8

Uptime: 6d06h, Last-change: 6d06h

Domain-ID: local

Service-Insertion: N/A

Locator	Pri/Wgt	Source	State
192.168.0.201			
10/10	cfg-intf	site-self,	reachable
Map-server	Uptime	ACK	Domain-ID
192.168.0.201			

6d06h

Yes

0

192.168.0.202

6d06h

Yes

0

If the MAC address of the gateway is correctly learned and the ACK flag has been marked as "**Yes**" for the Fabric Control planes, this stage is considered completed.

### **DHCP Broadcast Bridged in L2 Flooding**

Without DHCP Snooping enabled, DHCP Broadcasts are not blocked and are encapsulated in multicast for Layer 2 Flooding. Conversely, if DHCP Snooping is enabled, the flood of DHCP Broadcast packets is prevented.

<#root>

BorderCP-1#


**show ip dhcp snooping**

Switch DHCP snooping is enabled  
Switch DHCP gleaning is disabled  
DHCP snooping is configured on following VLANs:  
1001

**DHCP snooping is operational on following VLANs:**

1001      <-- VLAN141 should not be listed, as DHCP snooping must be disabled in L2 Only pools.

Proxy bridge is configured on following VLANs:  
none  
Proxy bridge is operational on following VLANs:  
none  
<snip>



**Tip:** Because DHCP Snooping is not enabled in the L2Border, DHCP Snooping Trust configuration is not needed.

---

At this stage, L2LISP ACL validation is already done in both devices.

Utilize the configured broadcast-underlay group for the L2LISP instance and the L2Border Loopback0 IP address to verify the L2 Flooding (S,G) entry that bridges this packet to other Fabric Nodes. Consult the mroute and mfib tables to validate parameters such as the incoming interface, outgoing interface list, and forwarding counters.

<#root>

BorderCP-1#

```
show ip int loopback 0 | i Internet
```

Internet address is  
192.168.0.201/32

BorderCP-1#

show run | se 8240

interface L2LISP0.8240

instance-id 8240

remote-rloc-probe on-route-change  
service ethernet  
eid-table vlan 1041

broadcast-underlay 239.0.17.1

BorderCP-1#

show ip mroute 239.0.17.1 192.168.0.201 | be \((

(  
192.168.0.201, 239.0.17.1  
) , 1w5d/00:02:52, flags: FTA  
Incoming interface:

Null0

, RPF nbr 0.0.0.0

<-- Local S,G IIF must be Null0

Outgoing interface list:

TenGigabitEthernet1/0/42

, Forward/Sparse, 1w3d/00:02:52, flags:

<-- Edge1 Downlink

TenGigabitEthernet1/0/43

, Forward/Sparse, 1w3d/00:02:52, flags:

<-- Edge2 Downlink

BorderCP-1#

show ip mfib 239.0.17.1 192.168.0.201 count

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second  
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)  
Default  
13 routes, 6 (\*,G)s, 3 (\*,G/m)s  
Group:

239.0.17.1

Source:

192.168.0.201


,  
SW Forwarding: 1/0/392/0, Other: 1/1/0  
HW Forwarding:

92071

/0/102/0, Other: 0/0/0

<-- HW Forwarding counters (First counter = Pkt Count) must increase

Totals - Source count: 1, Packet count: 92071



**Tip:** If an (S,G) entry is not found or the Outgoing Interface List (OIL) contains no Outgoing Interfaces (OIFs), it indicates an issue with the underlay multicast configuration or operation.

---

With these validations, along packet captures similar to the previous steps, this section is concluded, as the DHCP Offer is forwarded as a broadcast to all Fabric Edges using the outgoing interface list contents, in this case, out of interface TenGig1/0/42 and TenGig1/0/43.

## **DHCP Offer and ACK - Broadcast - Edge**

Exactly as the previous flow, verify the L2Border S,G in the Fabric Edge, where the incoming interface points towards the L2BN and the OIL contains the L2LISP instance mapped to VLAN 1041.

<#root>

Edge-1#

show vlan id 1041

VLAN Name	Status	Ports
-----------	--------	-------

1041

L2ONLY\_WIRED

active

L2L10:

8240

,

Te1/0/2

, Te1/0/17, Te1/0/18, Te1/0/19, Te1/0/20, Ac2, Po1

Edge-1#

show ip mroute 239.0.17.1 192.168.0.201 | be \

(

192.168.0.201

,

239.0.17.1

), 1w3d/00:01:52, flags: JT

Incoming interface:

TenGigabitEthernet1/1/2

, RPF nbr 192.168.98.2

<-- IIF Te1/1/2 is the RPF interface for 192.168.0.201 (L2BN RLOC)

Outgoing interface list:

L2LISP0.8240,

Forward/Sparse-Dense

,

1w3d/00:02:23, flags:

Edge-1#

show ip mfib 239.0.17.1 192.168.0.201 count

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second

Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Default

13 routes, 6 (\*,G)s, 3 (\*,G/m)s  
Group:

239.0.17.1

Source:

192.168.0.201,

SW Forwarding: 1/0/96/0, Other: 0/0/0

HW Forwarding:


76236

/0/114/0, Other: 0/0/0

<-- HW Forwarding counters (First counter = Pkt Count) must increase

Totals - Source count: 1, Packet count: 4





**Tip:** If an (S,G) entry is not found, it indicates an issue with the underlay multicast configuration or operation. If the L2LISP for the required instance is not present as OIF, it indicates an issue with the operation UP/DOWN status of the L2LISP sub-interface or the IGMP enablement status of the L2LISP interface.

---

L2LISP ACL validation is already done in both devices.

After the packet is de-encapsulated and placed on the VLAN matching VNI 8240, its broadcast nature dictates that it is flooded out all Spanning Tree Protocol forwarding ports for VLAN1041.

<#root>

Edge-1#

```
show spanning-tree vlan 1041 | be Interface
```

Interface	Role	Sts	Cost	Prio.	Nbr	Type
-----						
Te1/0/2						
Desg						
FWD						
20000	128.2	P2p	Edge			
Te1/0/17		Desg				
FWD						
2000	128.17	P2p				
Te1/0/18		Back				
BLK						
2000	128.18	P2p				
Te1/0/19		Desg				
FWD						
2000	128.19	P2p				
Te1/0/20		Back				
BLK						
2000	128.20	P2p				

The MAC address table identifies port Te1/0/2 as the endpoint port, which is in FWD state by STP, the packet is flooded out to the endpoint.

<#root>

Edge-1#

show mac address-table interface te1/0/2

Mac Address Table			
Vlan	Mac Address	Type	Ports
-----			
1041			
	aaaa.ddd.d.bbbb		
	DYNAMIC		
Te1/0/2			

The DHCP Offer and ACK process remains consistent. Without DHCP Snooping enabled, no entries are created in the DHCP Snooping table. Consequently, the Device-Tracking entry for the DHCP-enabled endpoint is generated by the glean of ARP packets. It is also expected that commands like "show platform

dhcpsnooping client stats" display no data, as DHCP snooping is disabled.

<#root>

Edge-1#

show device-tracking database interface te1/0/2 | be Network

Network Layer Address	Link Layer Address	Interface	vlan	prlv1	ag
ARP					
172.16.141.1					
aaaa.dddd.bbbb					
	Te1/0/2				
1041					
0005	45s	REACHABLE	207 s	try 0	

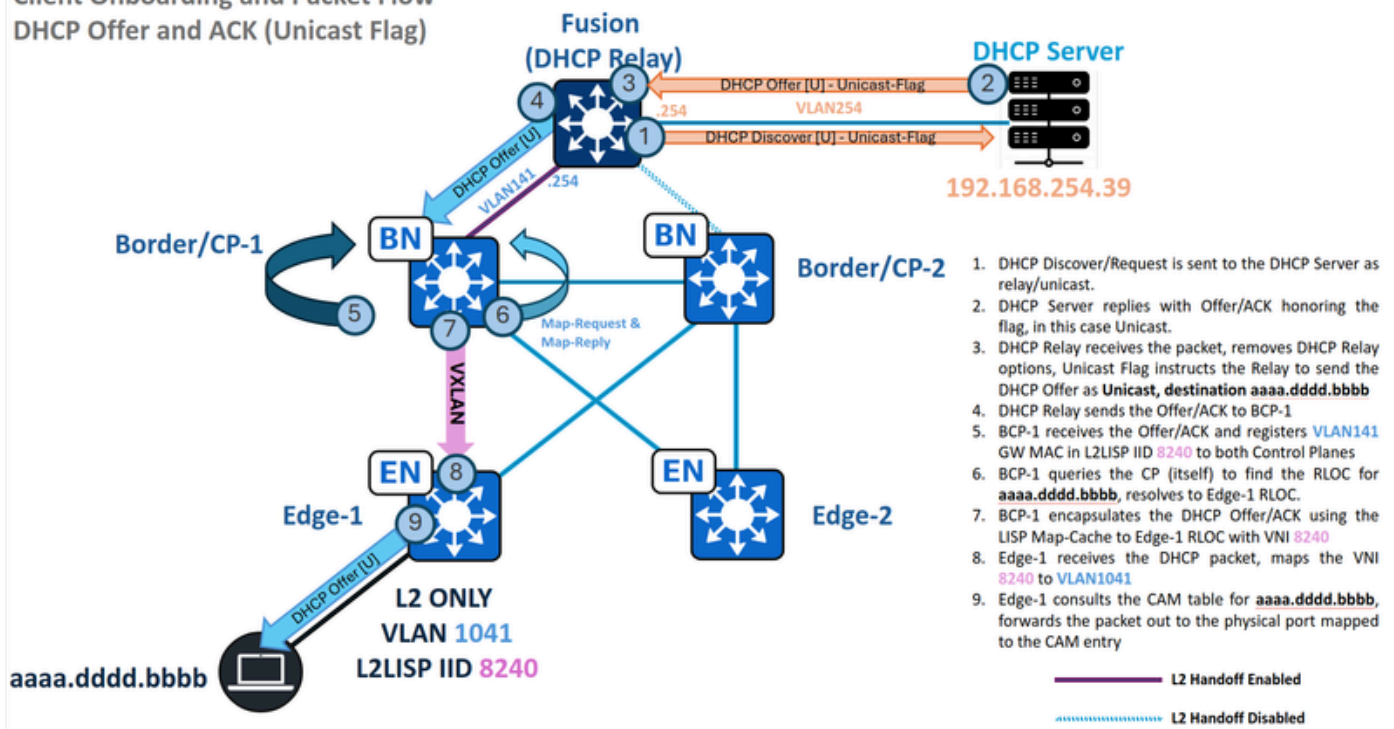
Edge-1#

show ip dhcp snooping binding vlan 1041

MacAddress	IpAddress	Lease(sec)	Type	VLAN	Interface
-----	-----	-----	-----	----	-----
Total number of bindings: 0					

DHCP Offer and ACK - Unicast - L2 Border

## Client Onboarding and Packet Flow DHCP Offer and ACK (Unicast Flag)



Traffic Flow - Unicast DHCP Offer and ACK in L2 Only

Here the scenario is a bit different, the endpoint sets the DHCP Broadcast Flag as unset or "0".

The DHCP Relay does not send the DHCP Offer/ACK as Broadcast, but as a unicast packet instead, with a destination MAC address derived from the client hardware address inside the DHCP payload. This drastically modifies the way the packet is handled by the SD-Access fabric, it uses the L2LISP Map-Cache to forward the traffic, not the Layer 2 Flooding multicast encapsulation method.

## Fabric Border/CP (192.168.0.201) packet capture: Ingress DHCP Offer

```
<#root>
```

```
BorderCP-1#
```

```
show monitor capture cap buffer display-filter "bootp.type==1 and dhcp.hw.mac_addr==aaaa.dddd.bbbb" deta
```

```
Dynamic Host Configuration Protocol (
```

```
Discover
```

```
)
```

```
Message type: Boot Request (1)
```

```
Hardware type: Ethernet (0x01)
```

```
Hardware address length: 6
```

```
Hops: 0
```

```
Transaction ID: 0x00002030
```

```
Seconds elapsed: 0
```

```
Bootp flags: 0x0000, Broadcast flag (Unicast)
```

```
0... .. = Broadcast flag: Unicast
```

```
.000 0000 0000 0000 = Reserved flags: 0x0000
Client IP address: 0.0.0.0
Your (client) IP address: 0.0.0.0
Next server IP address: 0.0.0.0
Relay agent IP address: 0.0.0.0
```

Client MAC address: aa:aa:dd:dd:bb:bb (aa:aa:dd:dd:bb:bb)

In this scenario, L2 Flooding is exclusively used for Discover/Requests, while Offers/ACKs are forwarded via L2LISP Map-Caches, simplifying the overall operation. Adhering to unicast forwarding principles, the L2 Border queries the Control Plane for the destination MAC address (aaaa.dddd.bbbb). Assuming successful "**MAC Learning and Endpoint Registration**" on the Fabric Edge, the Control Plane has this Endpoint ID (EID) registered.

<#root>

BorderCP-1#

show

```
lisp instance-id 8240 ethernet server aaaa.dddd.bbbb
```

LISP Site Registration Information

Site name: site\_uci

Description: map-server configured from Catalyst Center

Allowed configured locators: any

Requested EID-prefix:

EID-prefix:

aaaa.dddd.bbbb/48

instance-id

8240

First registered:	00:36:37
Last registered:	00:36:37
Routing table tag:	0
Origin:	Dynamic, more specific of any-mac
Merge active:	No
Proxy reply:	Yes
Skip Publication:	No
Force Withdraw:	No
TTL:	1d00h
State:	complete
Extranet IID:	Unspecified

Registration errors:

Authentication failures: 0

Allowed locators mismatch: 0

ETR 192.168.0.101:51328

```
, last registered 00:36:37, proxy-reply, map-notify
    TTL 1d00h, no merge, hash-function sha1
    state complete, no security-capability
    nonce 0x1BF33879-0x707E9307
    xTR-ID 0xDEFA4F0B-0xA801409E-0x29F87978-0xB865BF0D
    site-ID unspecified
    Domain-ID 1712573701
    Multihoming-ID unspecified
    sourced by reliable transport
Locator      Local State      Pri/Wgt Scope
192.168.0.101 yes      up          10/10   IPv4 none
```

After the Border's query to the Control Plane (local or remote), the LISP resolution establishes a Map-Cache entry for the endpoint's MAC address.

<#root>

BorderCP-1#

show lisp instance-id 8240 ethernet map-cache aaaa.ddd.ddd

LISP MAC Mapping Cache for LISP 0 EID-table Vlan

141

(IID

8240

), 1 entries

aaa.ddd.ddd/48

, uptime: 4d07h, expires: 16:33:09,

via map-reply

,

complete

, local-to-site

Sources: map-reply

State: complete, last modified: 4d07h, map-source: 192.168.0.206

Idle, Packets out: 46(0 bytes), counters are not accurate (~ 00:13:12 ago)

Encapsulating dynamic-EID traffic

Locator	Uptime	State	Pri/Wgt	Encap-IID
---------	--------	-------	---------	-----------

192.168.0.101				
---------------	--	--	--	--

4d07h	up	10/10	-	
-------	----	-------	---	--

<snip>

With the RLOC resolved, the DHCP Offer is encapsulated in unicast and sent directly to Edge-1 at

192.168.0.101, utilizing VNI 8240.

<#root>

BorderCP-1#

show mac address-table address aaaa.ddd.ddd

Mac Address Table			
Vlan	Mac Address	Type	Ports

141

aaa.ddd.ddd

CP\_LEARN

L2L10

BorderCP-1#

show platform software fed switch active matm macTable vlan 141 mac aaa.ddd.ddd

VLAN	MAC	Type	Seq#	EC_Bi	Flags	machandle	siHandle	riHandle	di
------	-----	------	------	-------	-------	-----------	----------	----------	----

141	aaa.ddd.ddd								
	0x1000001	0	0	64	0x718eb5271228	0x718eb52b4d68	0x718eb52be578	0x0	10

RLOC 192.168.0.101

adj\_id 747 No

BorderCP-1#

show ip route 192.168.0.101

Routing entry for 192.168.0.101/32  
Known via "

isis

", distance 115, metric 20, type level-2  
Redistributing via isis, bgp 65001T  
Advertised by bgp 65001 level-2 route-map FABRIC\_RLOC  
Last update from 192.168.98.3 on TenGigabitEthernet1/0/42, 1w3d ago  
Routing Descriptor Blocks:  
\* 192.168.98.3, from 192.168.0.101, 1w3d ago,

via TenGigabitEthernet1/0/42

Route metric is 20, traffic share count is 1

With the same methodology as in previous sections, capture traffic both ingress from the DHCP Relay and to the RLOC egress interface to observe the VXLAN encapsulation in unicast to the Edge RLOC.

## DHCP Offer and ACK - Unicast - Edge

The Edge receives the unicast DHCP Offer/ACK from the Border, de-encapsulates the traffic and consult its MAC address table to determine the correct egress port. Unlike broadcast Offer/ACKs, the Edge node forwards the packet only to the specific port where the endpoint is connected, rather than flooding it to all ports.

The MAC address table identifies port Te1/0/2 as our client port, which is in FWD state by STP, the packet is forwarded out to the endpoint.

<#root>

Edge-1#

```
show mac address-table interface te1/0/2
```

Mac Address Table			
Vlan	Mac Address	Type	Ports
1041	aaaa.dddd.bbbb	DYNAMIC	Te1/0/2

The DHCP Offer and ACK process remains consistent. Without DHCP Snooping enabled, no entries are created in the DHCP Snooping table. Consequently, the Device-Tracking entry for the DHCP-enabled endpoint is generated by the glean ARP packets. It is also expected that commands like "show platform dhcpsnooping client stats" display no data, as DHCP snooping is disabled.

<#root>

Edge-1#

```
show device-tracking database interface te1/0/2 | be Network
```

Network Layer Address	Link Layer Address	Interface	vlan	prlv1	ag
ARP					



172.16.141.1

aaaa.ddd.bbbb

Te1/0/2

1041

0005 45s REACHABLE 207 s try 0

Edge-1#

show ip dhcp snooping binding vlan 1041

MacAddress	IpAddress	Lease(sec)	Type	VLAN	Interface
-----	-----	-----	-----	----	-----
Total number of bindings: 0					

It is crucial to note that the SD-Access fabric does not influence the use of the Unicast or Broadcast flag, as this is solely an endpoint behavior. While this functionality might be overridden by the DHCP Relay or the DHCP Server itself, both mechanisms are essential for seamless DHCP operation in an L2 Only environment: L2 Flooding with Underlay Multicast for Broadcast Offers/ACKs, and proper Endpoint registration in the Control Plane for Unicast Offer/ACKs.