Implement IPv6 in Software-Defined Access

Contents

Introduction

This document describes how to implement IPv6 in Cisco® Software-Defined Access (SD-Access).

Background Information

IPv4 was released in 1983 and is still in use for the majority of internet traffic. The 32 bits IPv4 addressing allowed over 4 billion unique combinations. However, due to the increase in the number of internet-connected clients, there is a shortage of unique IPv4 addresses. In the 1990s, the exhaustion of IPv4 addressing became inevitable.
In anticipation of this, Internet Engineering Taskforce introduced the IPv6 standard. IPv6 utilises 128 bits and offers 340 undecillion unique IP addresses, which is more than enough to cater to the need of connected devices that grow. As more and more modern end-point devices support dual-stack and or single IPv6 stack, it is crucial for any organization to be ready for the adoption of IPv6. This means that the entire infrastructure must be ready for IPv6. Cisco SD-Access is the evolution from traditional campus designs to the networks that directly implement the intent of an organization. Cisco Software Defined Networks is now ready to onboard dual-stack (IPv6 Devices).

A major challenge for any organization in the adoption of IPV6 is the change management and complexities associated with the migration of legacy IPv4 systems to IPv6. This paper covers all the details about IPv6 feature support on Cisco SDN, strategy, and critical pit spot points, which need to be taken care of when you adopt IPv6 with Cisco Software Defined Networks.

In August 2019, Cisco DNA Centre version 1.3 was first introduced with the support of IPv6. In this release, the Cisco SD-Access campus network supported the host IP address with wired and wireless clients in IPv4, IPv6 or IPv4v6 Dual-stack from the overlay fabric network. The solution is to continuously evolve to bring new features and functionalities that easily onboard the IPv6 for any enterprise.

Cisco SD-Access with IPv6 Architecture

Fabric technology, an integral part of SD-Access, provides wired and wireless campus networks with programmable overlays and easy-to-deploy network virtualization, that permit a physical network to host one or more logical networks to meet the design intent. In addition to network virtualization, fabric technology in the campus network enhances control of communications, that provides software-defined segmentation and policy enforcement based on user identity and group membership. The entire Cisco SDN solution runs on the DNA of the fabric. Hence, it is critical to understand each pillar of the solution with respect to IPv6 support.

â€¢ Underlay - IPv6 functionality for Overlay has a dependency on the underlay as the IPv6 overlay makes use of the IPv4 underlay IP addressing to create LISP control plane and VxLAN data plane tunnels. You can always enable the dual-stack for the underlay routing protocol, just the SD-Access overlay LISP only depends on the IPv4 routing. (This requirement is for the current version of DNA-C (2.3.x) and is removed in later releases where underlay can be dual-stack or single IPv6 stack only).

â€¢ Overlay - When it comes to the overlay, SD-Access supports both IPv6-only wired and wireless endpoints. That IPv6 traffic is encapsulated in IPv4 and VxLAN header within the SD-Access fabric until they reach the fabric border nodes. The fabric border nodes decapsulate the IPv4 and VxLAN header, which pursues the normal IPv6 unicast routing process from then.

â€¢ Control Plane Nodes - The Control Plane node is configured to allow all IPv6 host subnets and the /128 host routes within the subnet ranges to be registered in its mapping database.

â€¢ Border nodes - On the border nodes, IPv6 BGP peering with fusion devices is enabled. The border node decapsulates the IPv4 header from the fabric egress traffic while the ingress IPv6 traffic is encapsulated with the IPv4 header by the border nodes as well.

â€¢ Fabric Edge â€“ All the SVIs which are configured in Fabric Edge have to be IPv6. This configuration is pushed by the DNA Center Controller.

â€¢ Cisco DNA Center â€“ The Cisco DNA Center physical interfaces do not support dual-stack as of the time this document is published. It can only deploy in a single stack with either IPv4 or IPv6 only in the management and or enterprise interfaces of the DNA Center.

â€¢ Clients â€“ Cisco® Software-Defined Access (SD-Access) supports dual-stack (IPv4&IPv6) or single stack either IPv4 or IPv6. However, in the case of an IPv6 single stack is deployed, DNA Center still requires to create a dual-stack pool to support an IPv6-only client. The IPv4 in the dual-stack pool is a dummy address only as the IPv6 the client is expected to disable the IPv4 address.

IPv6 Overlay Architecture in Cisco Software-Defined Access.

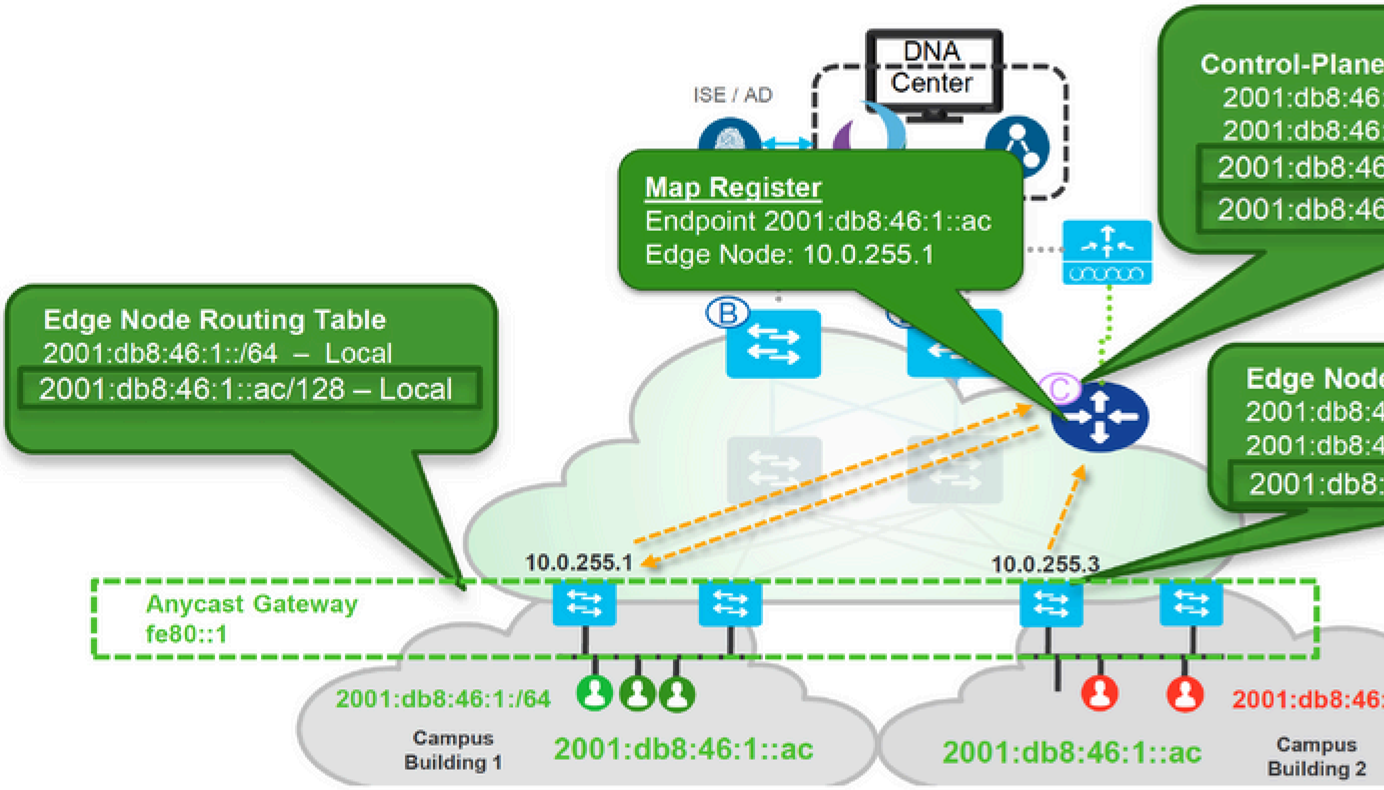

Figure 1.
IPv6 Overlay Architecture in Cisco Software Defined Access

IPv6 Overlay Architecture

Enable IPv6 with Cisco DNA-Center

There are two ways to enable the IPv6 pool in the Cisco DNA Center:
1. Create a new dual-stack IPv4/v6 Pool - greenfield
2. Edit IPv6 on the IPv4 pool that already exists - brownfield migration
The current release (up to 2.3.x) of DNA Center does not support IPv6 only a pool if the user plans to support a single/native IPv6 address-only client, a dummy IPv4 address needs to associate with the IPv6 pool. Please note that from the deployed IPv4 pool that already exists with a site associated with it, and edit the pool with an IPv6 address, DNA Center provides the migration option for the SD-Access Fabric which requires the user to reprovision the fabric for that site. A warning indicator is displayed in the Fabric where the site belongs and indicates that Fabric needs to â€˜reconfigure fabricâ€™. Please see these images for samples.

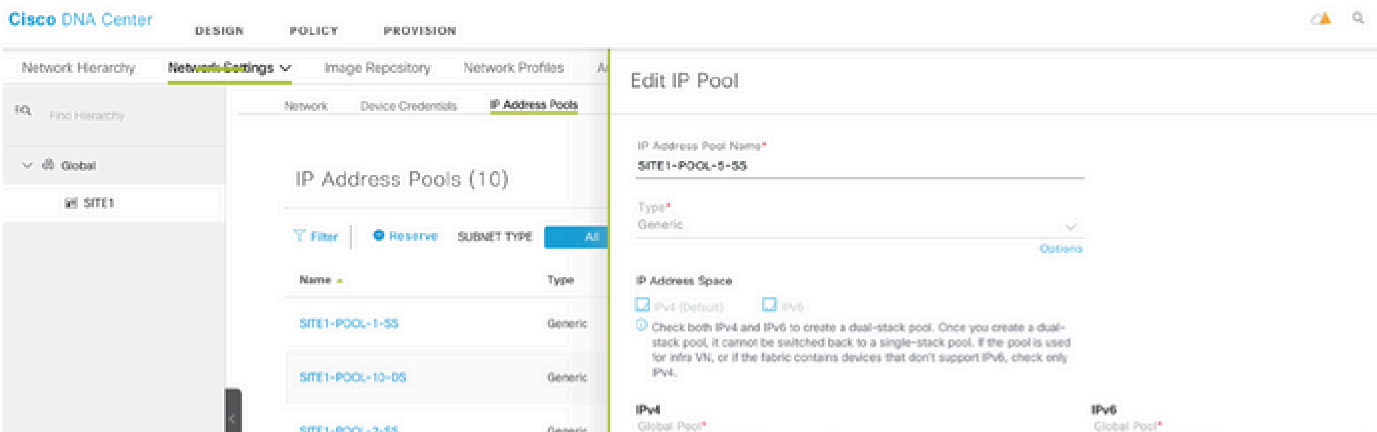

IP Address Pools (10)

Name	Type
SITE1-POOL-1-SS	Generic
SITE1-POOL-10-DS	Generic
SITE1-POOL-2-SS	Generic

Edit IP Pool

IP Address Pool Name*
SITE1-POOL-5-SS

Type*
Generic

IP Address Space
□ IPv4 (Default) □ IPv6
Check both IPv4 and IPv6 to create a dual-stack pool. Once you create a dual-stack pool, it cannot be switched back to a single-stack pool. If the pool is used for infra VN, or if the fabric contains devices that don't support IPv6, check only IPv4.

IPv4
Global Pool*

IPv6
Global Pool*

Step 1. After the IPv6 client configures itself with an IPv6 link-local address, the client sends ICMPv6 Router Solicitation (RS) message to Fabric Edge. The purpose of this message is to gain the global unicast prefix of its connected segment.
Step 2. After the fabric edge receives the RS message, it responds back with an ICMPv6 Router Advertisement (RA) message which contains the global IPv6 unicast prefix and its length inside.
Step 3. Once the client gets the RA message, it combines the IPv6 global unicast prefix with its EUI-64 interface identifier to generate its unique IPv6 global unicast address and set its gateway to the link-local address of the fabric edgeâ€™s SVI which is related to the clientâ€™s segment. Then the client sends out an ICMPv6 Neighbor Solicitation message to perform duplicate address detection (DAD) to make sure the IPv6 address it gets is unique.

> **Note**: All the SLAAC-related messages are encapsulated with the clientâ€™s and the fabric nodeâ€™s SVI IPv6 link-local address.

IPv6 Address Assignment â€“ DHCPv6

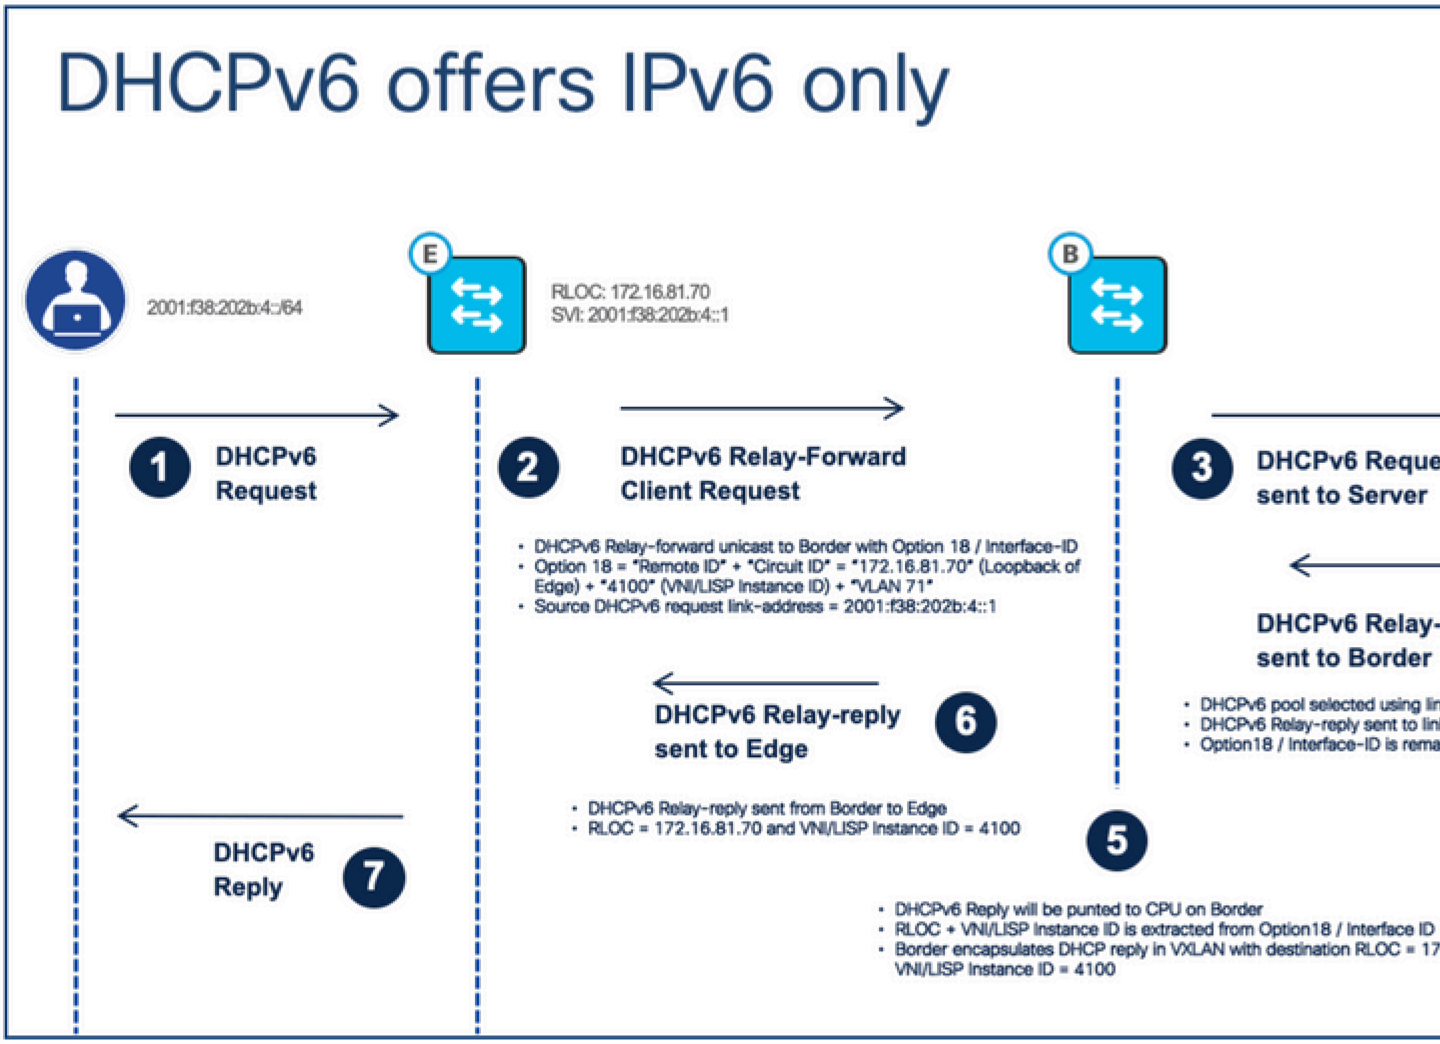

DHCPv6 offers IPv6 only

- DHCPv6 Relay-forward unicast to Border with Option 18 / Interface-ID
- Option 18 = "Remote ID" + "Circuit ID" = "172.16.81.70" (Loopback of Edge) + "4100" (VNI/LISP Instance ID) + "VLAN 71"
- Source DHCPv6 request link-address = 2001:f38:202b:4::1

- DHCPv6 pool selected using lir
- DHCPv6 Relay-reply sent to lin
- Option18 / Interface-ID is rema

- DHCPv6 Relay-reply sent from Border to Edge
- RLOC = 172.16.81.70 and VNI/LISP Instance ID = 4100

- DHCPv6 Reply will be punted to CPU on Border
- RLOC + VNI/LISP Instance ID is extracted from Option18 / Interface ID
- Border encapsulates DHCP reply in VXLAN with destination RLOC = 17
VNI/LISP Instance ID = 4100

IPv6 address assignment â€“ DHCPv6

Call Flow Description:
Step 1. The client sends out the DHCPv6 request to the fabric edge.
Step 2. When the fabric edge receives the DHCPv6 request, itâ€™ll use the DHCPv6 Relay-forward message to unicast the request to the fabric border with DHCPv6 option 18. Compared to DHCP option 82, the DHCPv6 option 18 encodes both â€œCircuit IDâ€� and â€œRemote IDâ€� together. The LISP Instance ID/VNI, IPv4 RLOC and endpoint VLAN are encoding
inside.
Step 3. The fabric border decapsulates the VxLAN header and unicasts the DHCPv6 packet to the DHCPv6 server.
Step 4. The DHCPv6 server receives the relay-forward message, it uses the source link address (DHCPv6 relay agent/clientâ€™s gateway) of the message to select the IPv6 IP pool to assign the IPv6 address. Then sends out the DHCPv6 relay-reply message back to the clientâ€™s gateway address. Option 18 remains unchanged.
Step 5. When the fabric border receives the relay-reply message, it extracts the RLOC and LISP instance/VNI from option 18. Fabric border encapsulates the relay-reply message in VxLAN with a destination which it extracted from option 18.
Step 6. The fabric border sends the DHCPv6 Relay-reply message to the fabric edge to which the client connects.
Step 7. When fabric edge receives the DHCPv6 relay-reply message, it decapsulates the messageâ€™s VxLAN header and forwards the message to the client. Then the client knows its assigned IPv6 address.

: In case the AP broadcasts the Non-Fabric SSIDs and does not broadcast Fabric SSID, please check for the VXLAN tunnel between the Access Point and the Fabric Edge node.

Also, note the AP to WLC communication always happens via Underlay CAPWAP and all WLC to AP communication uses VXLAN CAPWAP via overlay. This means, if you capture packets that go from AP to WLC, you only see CAPWAP while the reverse traffic has a VXLAN tunnel. See this example for communication between AP and WLC.

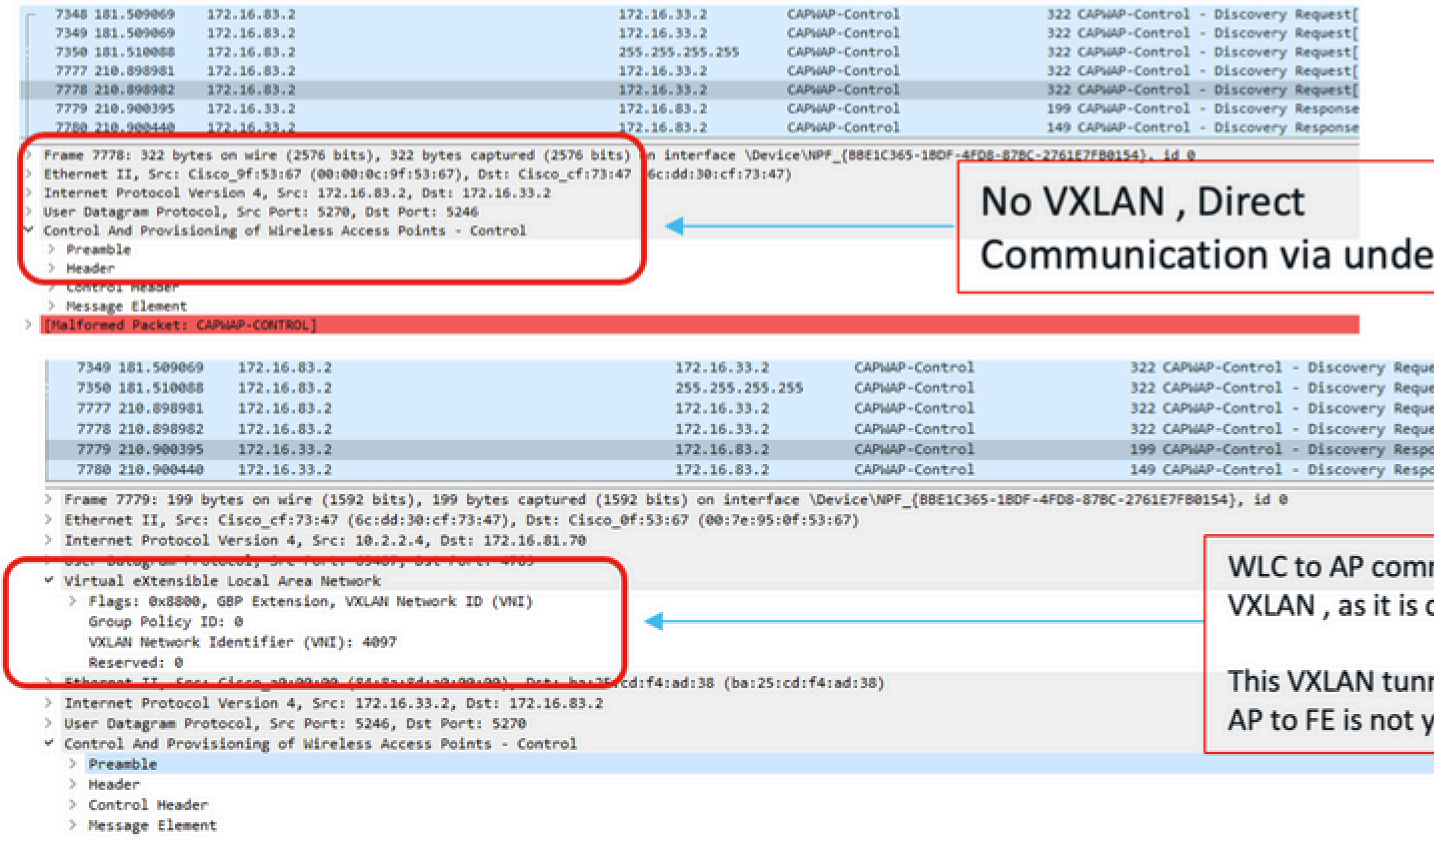

```
7348 181.509069   172.16.83.2          172.16.33.2       CAPWAP-Control   322 CAPWAP-Control - Discovery Request[
7349 181.509069   172.16.83.2          172.16.33.2       CAPWAP-Control   322 CAPWAP-Control - Discovery Request[
7350 181.510088   172.16.83.2          255.255.255.255   CAPWAP-Control   322 CAPWAP-Control - Discovery Request[
7777 210.898981   172.16.83.2          172.16.33.2       CAPWAP-Control   322 CAPWAP-Control - Discovery Request[
7778 210.898982   172.16.83.2          172.16.33.2       CAPWAP-Control   322 CAPWAP-Control - Discovery Request[
7779 210.900395   172.16.33.2          172.16.83.2       CAPWAP-Control   199 CAPWAP-Control - Discovery Response
7780 210.900440   172.16.33.2          172.16.83.2       CAPWAP-Control   149 CAPWAP-Control - Discovery Response
> Frame 7778: 322 bytes on wire (2576 bits), 322 bytes captured (2576 bits) on interface \Device\NPF_{88E1C365-18DF-4FD8-87BC-2761E7FB0154}, id 0
> Ethernet II, Src: Cisco_9f:53:67 (00:00:0c:9f:53:67), Dst: Cisco_cf:73:47 (6c:dd:30:cf:73:47)
> Internet Protocol Version 4, Src: 172.16.83.2, Dst: 172.16.33.2
> User Datagram Protocol, Src Port: 5270, Dst Port: 5246
v Control And Provisioning of Wireless Access Points - Control
  > Preamble
  > Header
  > Control Header
  > Message Element
> [Malformed Packet: CAPWAP-CONTROL]
```

No VXLAN , Direct
Communication via unde

```
7349 181.509069   172.16.83.2          172.16.33.2       CAPWAP-Control   322 CAPWAP-Control - Discovery Reque
7350 181.510088   172.16.83.2          255.255.255.255   CAPWAP-Control   322 CAPWAP-Control - Discovery Reque
7777 210.898981   172.16.83.2          172.16.33.2       CAPWAP-Control   322 CAPWAP-Control - Discovery Reque
7778 210.898982   172.16.83.2          172.16.33.2       CAPWAP-Control   322 CAPWAP-Control - Discovery Reque
7779 210.900395   172.16.33.2          172.16.83.2       CAPWAP-Control   199 CAPWAP-Control - Discovery Respo
7780 210.900440   172.16.33.2          172.16.83.2       CAPWAP-Control   149 CAPWAP-Control - Discovery Respo
> Frame 7779: 199 bytes on wire (1592 bits), 199 bytes captured (1592 bits) on interface \Device\NPF_{8BE1C365-18DF-4FD8-87BC-2761E7FB0154}, id 0
> Ethernet II, Src: Cisco_cf:73:47 (6c:dd:30:cf:73:47), Dst: Cisco_0f:53:67 (00:7e:95:0f:53:67)
> Internet Protocol Version 4, Src: 10.2.2.4, Dst: 172.16.81.70
> User Datagram Protocol, Src Port: 65487, Dst Port: 4789
v Virtual eXtensible Local Area Network
  > Flags: 0x8800, GBP Extension, VXLAN Network ID (VNI)
    Group Policy ID: 0
    VXLAN Network Identifier (VNI): 4097
    Reserved: 0
> Ethernet II, Src: Cisco_a9:00:00 (84:8a:8d:a9:00:00), Dst: ba:25:cd:f4:ad:38 (ba:25:cd:f4:ad:38)
> Internet Protocol Version 4, Src: 172.16.33.2, Dst: 172.16.83.2
> User Datagram Protocol, Src Port: 5246, Dst Port: 5270
v Control And Provisioning of Wireless Access Points - Control
  > Preamble
  > Header
  > Control Header
  > Message Element
```

WLC to AP com
VXLAN , as it is c

This VXLAN tunn
AP to FE is not y

Packet Captures from AP to WLC (CAPWAP Tunnel) vs WLC to AP (VxLAN Tunnel in the Fabric)

Client Onboarding

Dual-Stack/IPv6 Client on-boarding process remains the same but the client uses the IPv6 address assignment methods like SLAAAC/DHCPv6 to get the IPv6 addresses.
1. Client joins the Fabric and enables SSID on the AP.
2. WLC knows the AP RLOC.
3. Client Authenticates and WLC registers the Client L2 details with CP and updates AP.
4. Client initiates the IPv6 Addressing from configured methods â€“ SLAAC/DHCPv6.
5. FE triggers IPv6 client registration to CP HTDB.
AP to. FE and FE to other destinations use the VXLAN and LISP IPv6 encapsulation within IPv4 frames.

Client-Client Communication with IPv6

The image here summarises the IPv6 wireless client communication process with another IPv6 wired client. (this assumes the client is authenticated and got the IPv6 address via configured methods).
1. Client Sends the 802.11 frames to the AP with IPv6 payload.
2. AP removes the 802.11 headers and sends the original IPv6 payload in the IPv4 VXAN tunnel to fabric Edge.
3. Fabric Edge uses the MAP request to identify the destination and sends the frame to the destination RLOC with IPv4 VXLAN.
4. At the destination Switch, the IPv4 VXLAN header is removed and the IPv6 packet is sent to the client.

VXLAN Tunnel Between AP and FE is IPv4
VXLAN
(Data)
C
B

: For any IPv6 communication outside fabric, for example, DHCP/DNS, IPv6 routing must be enabled between the border and Non-fabric infrastructure.

Step 0. The client authenticates and gets the IPv6 address from the configured methods.

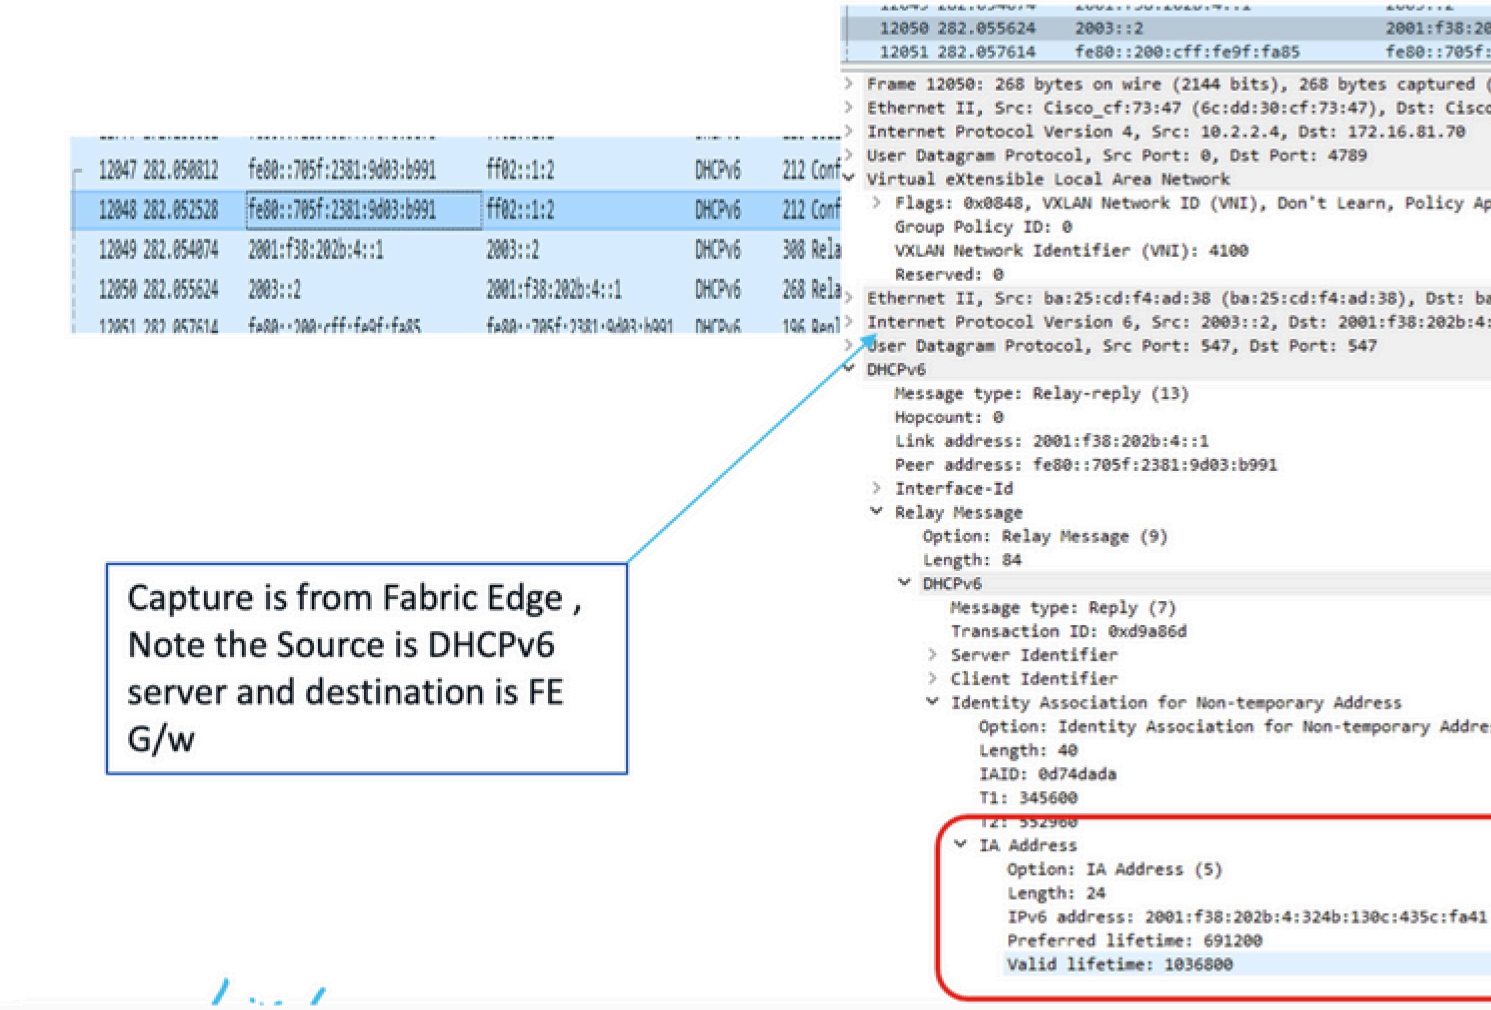

```
12047 282.050812  fe80::705f:2381:9d03:b991  ff02::1:2              DHCPv6  212 Conf
12048 282.052528  fe80::705f:2381:9d03:b991  ff02::1:2              DHCPv6  212 Conf
12049 282.054074  2001:f38:202b:4::1         2003::2                DHCPv6  308 Rela
12050 282.055624  2003::2                    2001:f38:202b:4::1     DHCPv6  268 Rela
```

```
12050 282.055624    2003::2                     2001:f38:20
12051 282.057614    fe80::200:cff:fe9f:fa85     fe80::705f:
> Frame 12050: 268 bytes on wire (2144 bits), 268 bytes captured (
> Ethernet II, Src: Cisco_cf:73:47 (6c:dd:30:cf:73:47), Dst: Cisco
> Internet Protocol Version 4, Src: 10.2.2.4, Dst: 172.16.81.70
> User Datagram Protocol, Src Port: 0, Dst Port: 4789
v Virtual eXtensible Local Area Network
  > Flags: 0x0848, VXLAN Network ID (VNI), Don't Learn, Policy Ap
    Group Policy ID: 0
    VXLAN Network Identifier (VNI): 4100
    Reserved: 0
> Ethernet II, Src: ba:25:cd:f4:ad:38 (ba:25:cd:f4:ad:38), Dst: ba
> Internet Protocol Version 6, Src: 2003::2, Dst: 2001:f38:202b:4:
> User Datagram Protocol, Src Port: 547, Dst Port: 547
v DHCPv6
    Message type: Relay-reply (13)
    Hopcount: 0
    Link address: 2001:f38:202b:4::1
    Peer address: fe80::705f:2381:9d03:b991
  > Interface-Id
  v Relay Message
      Option: Relay Message (9)
      Length: 84
    v DHCPv6
        Message type: Reply (7)
        Transaction ID: 0xd9a86d
      > Server Identifier
      > Client Identifier
      v Identity Association for Non-temporary Address
          Option: Identity Association for Non-temporary Addre
          Length: 40
          IAID: 0d74dada
          T1: 345600
          T2: 552960
        v IA Address
            Option: IA Address (5)
            Length: 24
            IPv6 address: 2001:f38:202b:4:324b:130c:435c:fa41
            Preferred lifetime: 691200
            Valid lifetime: 1036800
```

Capture is from Fabric Edge , Note the Source is DHCPv6 server and destination is FE G/w

Packet Capture from DHCPv6 server to Fabric Edge Node

Step 1. The wireless client sends the 802.11 frames to the Access point with the IPv6 payload.
Step 2. The access point removes the wireless header and sends the packet to the Fabric edge. This uses the VXLAN tunnel header which is IPv4 based as Access Point has the IPv4 address.

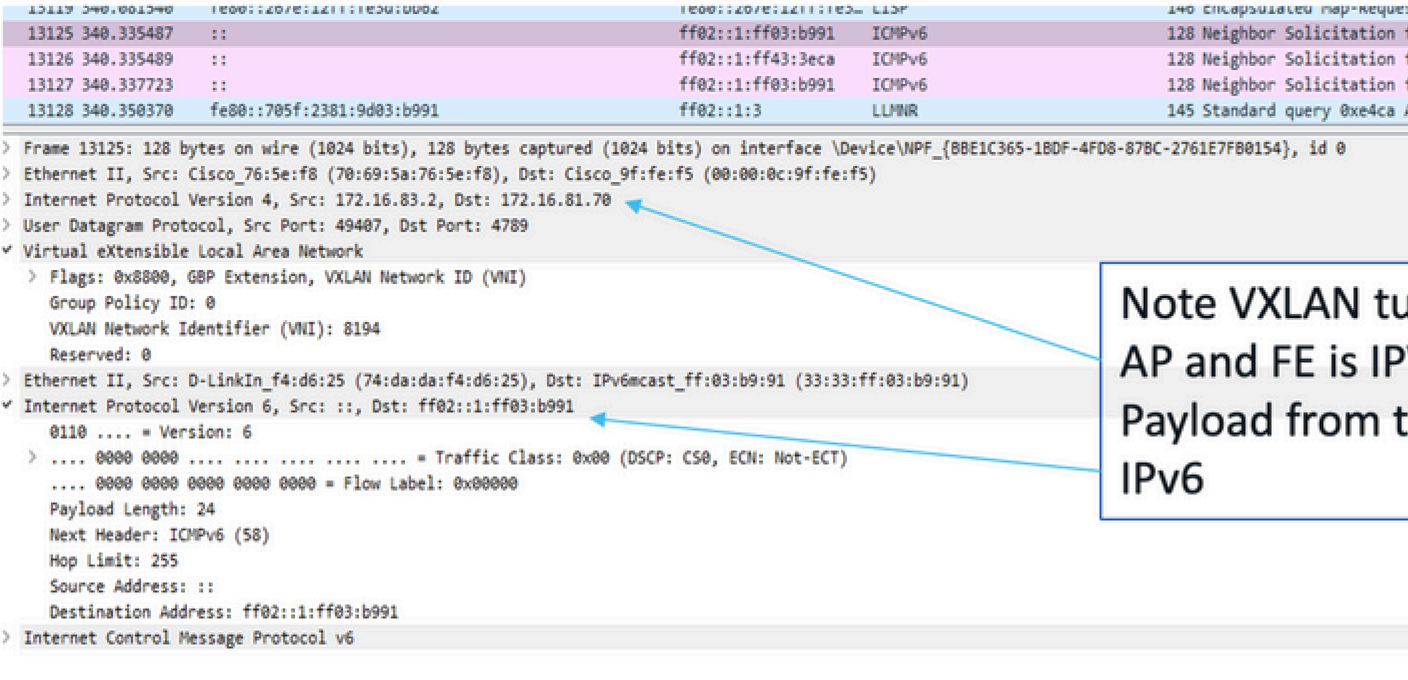

```
13125 340.335487   ::                           ff02::1:ff03:b991   ICMPv6   128 Neighbor Solicitation
13126 340.335489   ::                           ff02::1:ff43:3eca   ICMPv6   128 Neighbor Solicitation
13127 340.337723   ::                           ff02::1:ff03:b991   ICMPv6   128 Neighbor Solicitation
13128 340.350370   fe80::705f:2381:9d03:b991    ff02::1:3           LLMNR    145 Standard query 0xe4ca
> Frame 13125: 128 bytes on wire (1024 bits), 128 bytes captured (1024 bits) on interface \Device\NPF_{BBE1C365-1BDF-4FD8-87BC-2761E7FB0154}, id 0
> Ethernet II, Src: Cisco_76:5e:f8 (70:69:5a:76:5e:f8), Dst: Cisco_9f:fe:f5 (00:00:0c:9f:fe:f5)
> Internet Protocol Version 4, Src: 172.16.83.2, Dst: 172.16.81.70
> User Datagram Protocol, Src Port: 49407, Dst Port: 4789
v Virtual eXtensible Local Area Network
  > Flags: 0x8800, GBP Extension, VXLAN Network ID (VNI)
    Group Policy ID: 0
    VXLAN Network Identifier (VNI): 8194
    Reserved: 0
> Ethernet II, Src: D-LinkIn_f4:d6:25 (74:da:da:f4:d6:25), Dst: IPv6mcast_ff:03:b9:91 (33:33:ff:03:b9:91)
v Internet Protocol Version 6, Src: ::, Dst: ff02::1:ff03:b991
    0110 .... = Version: 6
  > .... 0000 0000 .... .... .... .... .... = Traffic Class: 0x00 (DSCP: CS0, ECN: Not-ECT)
    .... 0000 0000 0000 0000 0000 = Flow Label: 0x00000
    Payload Length: 24
    Next Header: ICMPv6 (58)
    Hop Limit: 255
    Source Address: ::
    Destination Address: ff02::1:ff03:b991
> Internet Control Message Protocol v6
```

Note VXLAN tu
AP and FE is IP
Payload from t
IPv6

Packet Capture for the VxLAN tunnel between FE and AP

: IPv6 Guest Access (web Portal) via Cisco Identity Services is not supported due to limitations on ISE.

Dependency Matrix

It is important to note the dependencies and support for IPv6 from different wireless components which are part of Cisco SD-Access. The table in this image summarises this feature matrix.

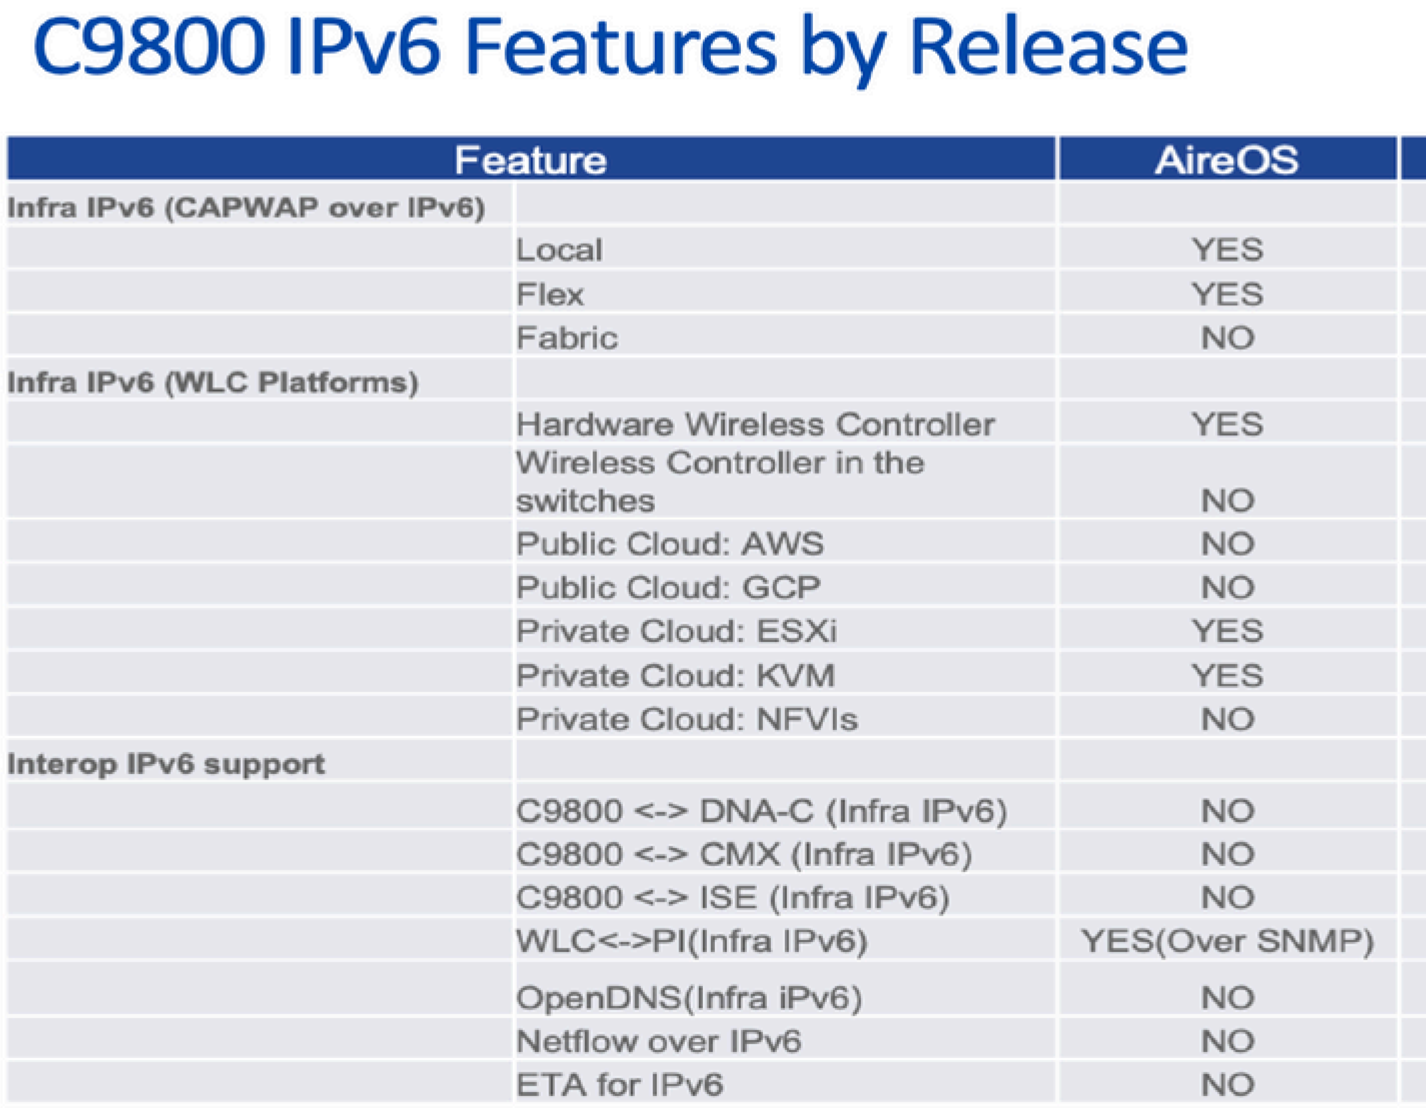

C9800 IPv6 Features by Release

Feature		AireOS
Infra IPv6 (CAPWAP over IPv6)		
	Local	YES
	Flex	YES
	Fabric	NO
Infra IPv6 (WLC Platforms)		
	Hardware Wireless Controller	YES
	Wireless Controller in the switches	NO
	Public Cloud: AWS	NO
	Public Cloud: GCP	NO
	Private Cloud: ESXi	YES
	Private Cloud: KVM	YES
	Private Cloud: NFVIs	NO
Interop IPv6 support		
	C9800 <-> DNA-C (Infra IPv6)	NO
	C9800 <-> CMX (Infra IPv6)	NO
	C9800 <-> ISE (Infra IPv6)	NO
	WLC<->PI(Infra IPv6)	YES(Over SNMP)
	OpenDNS(Infra iPv6)	NO
	Netflow over IPv6	NO
	ETA for IPv6	NO

Cat9800 WLC IPv6 Features by Release

Monitor the Control Plane for IPv6

Once you enable the IPv6, you start to see additional entries about host IPv6 in the MS/MR servers. As a host can have multiple IPv6 IP addresses, your MS/MR lookup table has entries for all the IP addresses. This is combined with the IPv4 table that already exists.
You have to login into the device CLI and issue these commands to check all the entries.

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2003::2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Quick FAQs for IPv6 Design with Cisco SD-Access

Q. Does Cisco Software Defined Network support IPv6 for underlay and overlay networks?
Only overlay is supported with the current release (2.3.x) at the time this paper is written.
Q. Does Cisco SDN support native IPv6 for both wired and wireless clients?
A: Yes. This requires dual-stack pools that are created in the DNA Center while IPv4 is the dummy pool as the clients disable IPv4 DHCP requests and only IPv6 DHCP or SLAAC addresses are offered.
Q. Can I have a native IPv6-only campus network in my Cisco SD-Access Fabric?
A: Not with the current release (up to 2.3.x). It is on the roadmap.
Q. Does Cisco SD-Access support L2 IPv6 hand-off?
A: Not at present. Only L2 IPv4 handoff and or L3 Dual-Stack hand-off are supported
Q. Does Cisco SD-Access support multicast for IPv6?
A: Yes. Only overlay IPv6 with headend replication multicast is supported. Native IPv6 multicast is not yet supported.
Q. Does Cisco SD-Access Fabric Enabled Wireless support guests in dual-stack?
A: Not supported yet in Cisco IOS-XE (Cat9800) WLC. AireOS WLC is supported via a workaround. For details implementation of the workaround, please contact the Cisco Customer Experience team.