



IPv6 Deployment Guide for Cisco Catalyst 9800 Series Wireless Controllers, Cisco IOS XE Amsterdam 17.1

Client and Infrastructure IPv6 Support in Cisco IOS XE Amsterdam 17.1
First Published: March 12, 2020

Table of Contents

<i>Client IPv6 Support in Cisco IOS XE Amsterdam 17.1</i>	4
<i>IPv6 Wireless Client Connectivity Supported in Cisco IOS XE Amsterdam 17.1</i>	4
Solution Components	4
<i>Client IPv6 Support in C9800 Cisco IOS XE Amsterdam 17.1</i>	5
Prerequisites for Wireless IPv6 Client Connectivity	5
SLAAC Address Assignment	6
DHCPv6 Address Assignment	7
Static IPv6 Address Assignment	7
NDP - Neighbor Discovery Protocol	8
Router Solicitation	8
Router Advertisement	8
Neighbor Solicitation	8
Neighbor Advertisement	8
<i>Multicast Handling on C9800 IOS-XE</i>	8
<i>IPv6 Multicast Configuration in IOS XE</i>	9
<i>IPv6 Media Stream</i>	10
<i>IPv6 Client Mobility</i>	12
<i>IPv6 IRCM – Inter Release Controller Mobility</i>	13
Support for Interface Groups	14
<i>First Hop Security for IPv6 Clients</i>	14
RA Management	14
Router Advertisement Guard	15
ND Suppress	16
DHCPv6 Server Guard	16
IPv6 Source Guard	16
IPv6 Access Control Lists	16
<i>Network Resource Efficiency for IPv6 Clients</i>	17
Neighbor Discovery Caching	17
Router Advertisement Throttling	17
IPv6 Router Configuration	18
<i>Configuration for Wireless IPv6 Client Support</i>	18
Configuring Global Controller (Screen Shots from IOS-XE Release)	18

Interface configuration.....	18
Infrastructure IPv6 Support in Cisco IOS XE Amsterdam 17.1.....	20
Configuring IPv6 CAPWAP UDP-Lite Support	29
Enabling IPv6 on Your IOS Infrastructure Device	29
Glossary	30

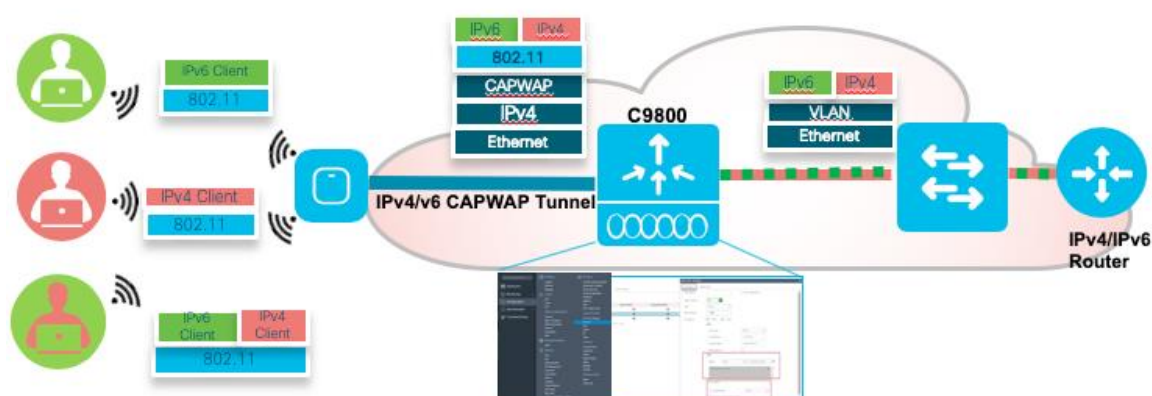
Client IPv6 Support in Cisco IOS XE Amsterdam 17.1

This document provides information about the theory of operation and configuration for Cisco's Unified Wireless LAN solution as it pertains to supporting IPv6 clients.

The [Infrastructure IPv6 Support](#) section in this document provides information about the Infrastructure support for IPv6 protocols in the C9800 controllers in Cisco IOS XE Amsterdam 17.1 and above.

IPv6 Wireless Client Connectivity Supported in Cisco IOS XE Amsterdam 17.1

Diagram below shows a typical IPv6 Wireless Client Configuration Workflow on the C9800 with Cisco IOS XE Amsterdam 17.1.



The IPv6 feature set within the Cisco C9800 software IOS-XE release version 17.1 allows the wireless network to support IPv4, Dual-Stack, and IPv6-only clients on the same wireless network. The overall goal for the addition of IPv6 client support to the Cisco SDA is to maintain feature parity between IPv4 and IPv6 clients including mobility, security, guest access, quality of service, and endpoint visibility.

A client can have multiple IPv6 addresses. For each IPv6 prefix received, the client generates an IPv6 address based on its MAC address, and one or more temporary IPv6 addresses, which are also globally routable. When multiple IPv6 prefixes are present, the client may have more addresses. In C9800 rel 17.1, maximum 8 IPv6 addresses per wireless client are supported. This allows IPv6 clients to have a link-local, SLAAC address, DHCPv6 address, and even addresses in alternative prefixes to be on a single interface.

Every IPv6 enabled interface must contain at least, 1 Loopback and 1 Link-Local address. Optionally, every interface can have multiple Unique-Local and Global IPv6 addresses.

Solution Components

- Wireless controllers supported in the release 17.1 are all forms of the Virtual C9800-CL, HW Appliance C9800-40/80, C9800-L and C9800-CL
- Cisco APs supported in IOS-XE 17.1 - 2700, 3700,1800, 2800, 3800, 4800 series and C9100 series APs.
- Cisco Outdoor APs supported in rel 17.1 – 1540, 1560 and 1570 (IPv4 only) series

Feature	AireOS	16.12	17.1
Infra IPv6 (CAPWAP over IPv6)			
Local	YES	YES	YES
Flex	YES	YES	YES
Fabric	NO	YES	YES
Infra IPv6 (WLC Platforms)			
Hardware Wireless Controller	YES	YES	YES
Wireless Controller in the switches	NO	YES	YES
Public Cloud: AWS	NO	NO	NO
Public Cloud: GCP	NO	NO	NO
Private Cloud: ESXi	YES	YES	YES
Private Cloud: KVM	YES	YES	YES
Private Cloud: NFVIs	NO	YES	YES
Interop IPv6 support			
C9800 <-> DNA-C (Infra IPv6)	NO	TBD	NO
C9800 <-> CMX (Infra IPv6)	NO	TBD	YES
C9800 <-> ISE (Infra IPv6)	NO	TBD	YES
WLC<->PI(Infra IPv6)	YES(Over SNMP)	YES	YES
OpenDNS(Infra IPv6)	NO	YES	YES
Netflow over IPv6	NO	YES	YES
ETA for IPv6	NO	NO	YES

Client IPv6 Support in C9800 Cisco IOS XE Amsterdam 17.1

Prerequisites for Wireless IPv6 Client Connectivity

To enable wireless IPv6 client connectivity, the underlying wired network must support IPv6 routing and an address assignment mechanism such as SLAAC or DHCPv6. The wireless controller must have L2 adjacency to the IPv6 router, and the VLAN must be tagged when entering the controller interfaces.

IPv6 is an important networking feature on wireless controllers when we extend IPv6 support from wired clients to wireless clients. A major IPv6 deployment method is the Stateless Address Autoconfiguration (SLAAC), which allows the clients to generate their own IPv6 addresses and ensures compliance to the network addressing rules. The SLAAC method utilizes the IPv6 Neighbor Discovery Protocol (NDP) to distribute the IPv6 network prefix information to clients and accommodate the detection for potential address conflicts. When address conflict does not exist, the NDP protocol facilitates client IPv6 address learning in the network.

SLAAC and NDP work on a VLAN basis. The IPv6 prefix from the router is specific to the client VLAN and it is distributed in the network only through that VLAN. The technical challenge for wireless clients to acquire IPv6 addresses through SLAAC and NDP is that VLAN is not a native concept in the wireless network between the controller and the wireless clients. Transportation of VLAN specific traffic over the wireless network requires special handling. The wireless controller needs to track the existence and location of wireless clients on a given VLAN in order to deliver the VLAN specific packets to the clients. The problem becomes more complicated when client mobility is present. In addition to tracking the local clients by their AP associations, the wireless controller also needs to track the roamed clients by their foreign controller attachments.

Prior to release 17.1 only IPv4 NDP was supported, with release 17.1 IPv4 and IPv6 NDP is supported.


The Cisco Neighbor Discovery Packet (NDP) is the fundamental tool for RRM and other wireless applications that provides information about the neighbor radio information. NDP is used first by the RF Grouping algorithm, but also by other RRM algorithms like DCA, TPC, FRA, Load balancing, Optimized roaming and Rogue detection.

Access points periodically also send out NDP messages over the air. APs are using the same RF group name to validate NDP messages from each other. When APs on different controllers hear validated neighbor messages at a signal strength above the given threshold, controllers use the neighbor message to form an RF grouping dynamically.

Note: when one controller is IPv4 only and the other controller is IPv6 only, this combination is not supported. In this case,

both controllers will continue to operate as independent RF group leaders, since a controller with IPv4 address cannot communicate with a controller having IPv6 address.

SLAAC Address Assignment




As mentioned above, the most common method for IPv6 client address assignment is Stateless Address Auto Configuration (SLAAC). SLAAC provides simple plug and play connectivity where clients self-assign an address based on the IPv6 prefix. This process is achieved by the IPv6 router sending out periodic Router Advertisement messages which inform the client of the IPv6 prefix in use (the first 64 bits) and of the IPv6 default gateway. From that point, clients can generate the remaining 64 bits of their IPv6 address based on either the MAC address of the adapter or randomly. Duplicate address detection is performed by IPv6 clients to ensure random addresses that are picked do not collide with other clients. The address of the router sending advertisements is used as the default gateway for the client.

The Stateless Address Autoconfiguration (SLAAC) is a newer IPv6 address configuration method other than the conventional DHCP method. DHCP is a stateful mechanism, where the server keeps the individual client addresses to ensure non-conflicting address assignments. SLAAC is stateless in the sense that the router does not track individual addresses. Rather, it is only responsible for advertising the address prefix. It is the client's decision to generate the IPv6 address conforming to the prefix. The address is formed by combining the prefix and a suffix mapped from the client MAC address. To protect the address uniqueness, the Neighbor Discovery Protocol (NDP) is employed to detect potential address conflicts.

The following configuration example from a Cisco-capable IPv6 router has the necessary commands to enable SLAAC addressing and router advertisements:

```
interface Vlan20 description IPv6-SLAAC
ip address 192.168.20.1 255.255.255.0
ipv6 address 2001:DB8:0:20::1/64
ipv6 enable
end
```

DHCPv6 Address Assignment



With DHCPv6, the client looks for the DHCPv6 server once it is connected to the network. The Solicit packet is multicast to the network and the client uses its IPv6 link local address as the source address. Any DHCPv6 server replying to the Solicit unicasts the Advertise packet back to the client using the client link local address as the destination. The client replies with a Request packet that is multicast to all servers to select an IPv6 address. Finally, the chosen server responds by unicasting the Reply packet to the client to confirm the IPv6 address.

The use of DHCPv6 is not required for IPv6 client connectivity if SLAAC is already deployed. There are two modes of operation for DHCPv6 called **Stateless** and **Stateful**.

The DHCPv6 **Stateless** mode is used to provide clients with additional network information not available in the router advertisement. This information can include the DNS domain name, DNS server(s), and other vendor-specific options. The following interface configuration example is for an IPv6 router implementing stateless DHCPv6 with SLAAC enabled, see a configuration example below:

```
interface Vlan20
description IPv6-DHCP-Stateless
ip address 192.168.20.1 255.255.255.0
ipv6 enable
ipv6 address 2001:DB8:0:20::1/64 ipv6 nd other-config-flag
ipv6 dhcp relay destination 2001:DB8:0:20::100
end
```

The DHCPv6 **Stateful** mode operates similar to DHCPv4, that is, it assigns addresses to each client instead of the client generating the address as in SLAAC. The following interface configuration is for an IPv6 router implementing stateful DHCPv6 with SLAAC turned off:

```
interface Vlan20
description IPv6-DHCP-Stateful
ip address 192.168.20.1 255.255.255.0
ipv6 enable
ipv6 address 2001:DB8:0:20::1/64
ipv6 nd prefix 2001:DB8:0:20::/64 no-advertise ipv6 nd managed-config-flag
ipv6 nd other-config-flag
ipv6 dhcp relay destination 2001:DB8:0:20::100
end
```

Static IPv6 Address Assignment

In this method, the user configures the client IPv6 address manually. The network learns the client address when the client announces it through an NA packet. Static addressing is simple, but it becomes difficult to use when the user has many clients to configure. It also has the potential problem of address conflict since the user chosen address is not guaranteed to be collision free.

NDP - Neighbor Discovery Protocol

The NDP protocol consists of several ICMPv6 message types that are designed for hosts in a network to find v6 routers, acquire routing information, and resolve address bindings. When used together with SLAAC.

IPv6 NDP packets have 4 types: Router Solicitation (RS), Router Advertisement (RA), Neighbor Solicitation (NS), Neighbor Advertisement (NA).

Router Solicitation

The RS message is ICMPv6 type 133. Hosts inquire the network for routers using this message after being attached to the network. The host uses its link local address as the source and multicasts it to the all-routers address FF02::2, which will reach all the routers in the link local scope. The RS message prompts routers to generate router advertisements quickly.

Router Advertisement

The RA message is ICMPv6 type 134. Routers advertise their presence with network parameters using this message either periodically or in response to a RS message. It carries the IPv6 network prefix to be used by hosts to construct their addresses. The source address in the packet is the router's link local address. If the RA is solicited by a RS, the destination address is the source address in the invoking RS packet. If it is unsolicited, the destination address is the all-nodes multicast address FF02::1.

Neighbor Solicitation

The NS message is ICMPv6 type 135. Nodes send NS to resolve the link-layer MAC address of a target node while providing their own MAC addresses to the target at the same time. NS is multicast when it is used to resolve an address and unicast when it is to verify the reachability of a neighbor node. The source address in NS is the sender's address except for the duplicate address detection case, where the source address is unspecified since the sender cannot start using its IPv6 address yet. The destination address in NS is the target node address if the sender is verifying the target reachability or the solicited-node multicast address corresponding to the target address.

Neighbor Advertisement

The NA message is ICMPv6 type 136. Nodes send NA in response to NS and also unsolicited NA to propagate node address information in the network. The Solicited flag in the message marks if the NA is solicited by an NS or not. The source address in the message is the sender's address. For solicited NA, the destination address is the source address in the invoking NS or all-nodes multicast address if the solicitation's source address is unspecified. For unsolicited NA, the destination is the all-nodes multicast address.

Multicast Handling on C9800 IOS-XE

Prior to release 16.12 MC2UC (multicast to unicast) was not supported on the C9800 controllers. With the release of the IOS-XE 16.12 the IPv6 MC2UC is supported in Local, Flex, Fabric and ME modes. AP in Flex, ME and Fabric will download all the current IPv6 MC2UC video stream groups from the controller during CAPWAP Join. During CAPWAP Join, APs in Flex Connect mode will download all the current MC2UC video stream groups from the C9800.

IPv6 MC2UC feature will extend the current wireless multicast IPv4 MC2UC while considering below requirements:

1. **Local mode:** Protect the Controller from excessive video client bandwidth usage and over-subscription (that will cause video quality to degrade). To achieve this channel utilization will be used as a metric to determine capacity and perform admission control and traffic QoS treatment. Traffic policies such as the number of maximum numbers of client per stream can also be configured to prevent over usage.
2. **Local Mode:** IPv6 multicast video subscription is determined via an MLD Join request. When the Controller denies a client video stream request, the Controller will notify the client video application that there is no

bandwidth available for the new stream.

3. **Local Mode:** The IPv6 MC2UC heavily depends on the radio resource management on APs and calculation of radio bandwidth which will be computed in the Controller. Controller will not process unnecessary traffic incoming from the multicast source if the Controller denies the request because of bandwidth allocation.
4. **Local Mode:** The network administrators will be able to control the maximum airtime allocation for video as well as for voice and data. The maximum airtime configuration can be made on the per radio band.
5. **Local/Flex/Fabric/ME:** The feature will co-exist with the existing voice architecture and design. The IPv6 MC2UC solution will inter-operate with existing voice QoS.

The multicast traffic for client in any local switching WLAN matching those video stream multicast groups will be converted into unicast frame to send to the client(s) subscribing the traffic through IGMP Join. Clients in the same WLAN not subscribing to the multicast stream through IGMP Join will not receive the traffic. If a particular multi-cast address is not part of the MC2UC configured groups, then traffic destined to such multi-cast addresses will continue to be transmitted as multi-cast traffic (i.e. no unicast conversion).

FlexConnect mode has two sub-modes: local switching mode and central switching mode. In the local switching mode, the data traffic is switched in the AP. Controller does not see any data traffic for that WLAN. So, with respect to multicast traffic, in the local mode, IGMP/MLD snooping is done by the AP itself.

Fabric mode can have fabric and non-fabric WLANs. Non-fabric-WLANs will be centrally switch and expected to work as local mode. For fabric WLANs the data traffic is switched in the AP. Controller does not see any data traffic for that WLAN. So, with respect to multicast traffic, in the fabric WLAN, IGMP/MLD snooping is done by the AP itself.

Similar to IPv4 MC2UC AP will not support CAC in Flex, ME and Fabric modes.

Note: This feature is only supported on Wave-2 APs. Wave-1 APs do not support this feature.

IPv6 Multicast Configuration in IOS XE

By default, multicast forwarding and MLD snooping is not enabled on the C9800 controllers. To enable IPv6 multicast forwarding and IGMP snooping enable both on the controller as shown in the example below by going to Configuration > Services > Multicast.

The screenshot displays the configuration page for IPv6 Multicast. The left sidebar shows navigation options: Dashboard, Monitoring, Configuration, Administration, and Troubleshooting. The main content area is titled 'Configuration > Services > Multicast'. It features several configuration sections:

- Global Wireless Multicast Mode:** Enabled (checkbox checked).
- Wireless mDNS Bidding:** Disabled (checkbox unchecked).
- Wireless Non-IP Multicast:** Disabled (checkbox unchecked).
- Wireless Broadcast:** Disabled (checkbox unchecked).
- AP Capwap Multicast:** Set to 'Unicast' (dropdown menu).
- MLD Snooping:** Enabled (checkbox checked). Below it, 'MLD Query Interval (milliseconds)' is set to 1000.
- IGMP Snooping Querier:** Enabled (checkbox checked).
- IGMP Snooping:** Enabled (checkbox checked). Below it, 'Last Member Querier Interval (milliseconds)' is set to 1000.

On the right side, there is a section for 'IGMP Snooping' with a search bar. It shows two columns: 'Disabled' and 'Enabled'. The 'Enabled' column contains a table with the following data:


Status	VLAN ID	Name
Enabled	1	default
Enabled	70	VI AN0070

Buttons for 'Enable All' and 'Disable All' are located at the bottom of the table.

When Multicast forwarding and IGMP snooping are enabled:

1. C9800 intercepts IGMP reports from IPv6 multicast clients

2. C9800 creates a MGID based on (Source, Group, VLAN) tuple, range 1-4095 (L2) or 4160-8191 (L3)
3. C9800 uses IGMP snooping to determine if report should be forwarded to source router
4. C9800 sends to MGID to WLAN mapping to AP
5. AP keeps track of MGID, WLAN and client association



When Multicast flows from source to client:

1. Controller checks if a MGID exists for that flow
2. If MGID exists:
 - a. With MoM, Controller forwards to all APs
 - b. With MoU, Controller forwards to APs having clients for the multicast group
3. APs forward to each WLAN having clients for this MGID

When Leave message is received from the last client in a MGID:

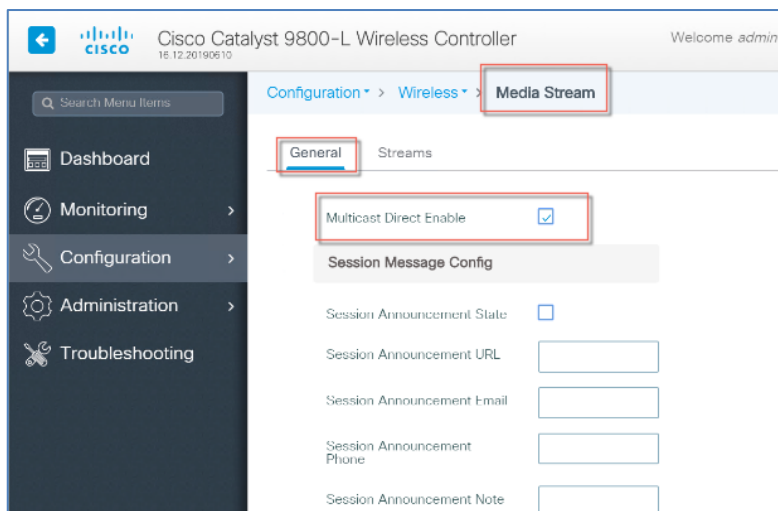
1. Controller removes client from MGID list, removes MGID
2. Controller does not forward leave to wired infrastructure (timeout is used)
3. Controller informs the AP about MGID deletion

IPv6 Media Stream

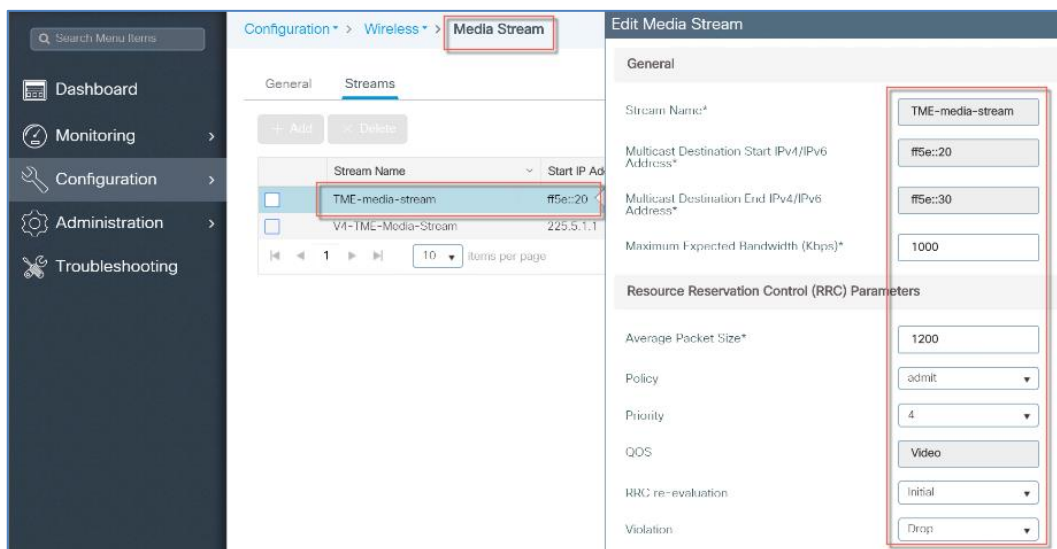
This feature is also called video-stream or media-stream under the Configuration > Wireless.

The Video Stream feature makes the IPv6 multicast stream delivery reliable over the air by converting the **broadcast** frame over the air to a **unicast** frame. Each Video Stream client acknowledges receiving a video IPv6 multicast stream. Below is an example of the Media Stream configuration for the IPv6 Media stream. First configure Media Stream> General > Multicast Direct Enable.

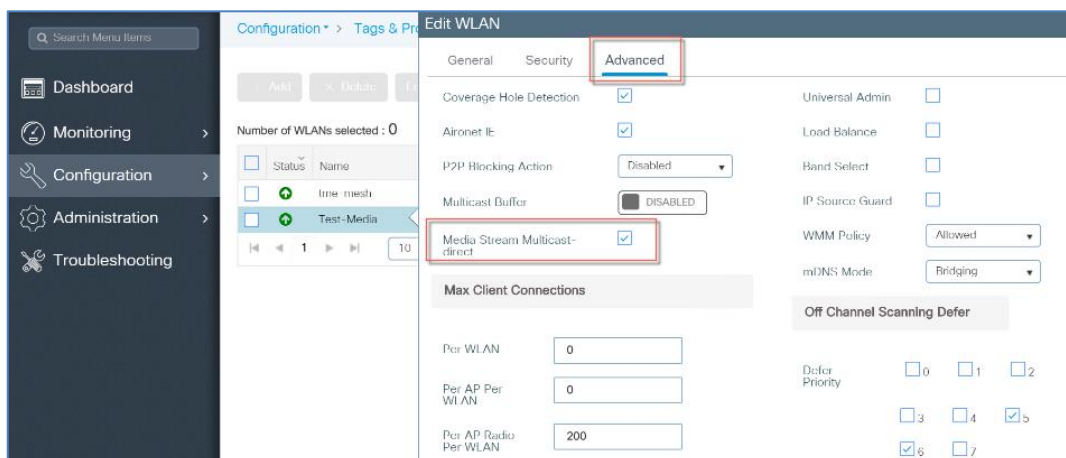
IPv6 Media Stream



And then configure the Media Stream with Multicast Start and End IP addresses.

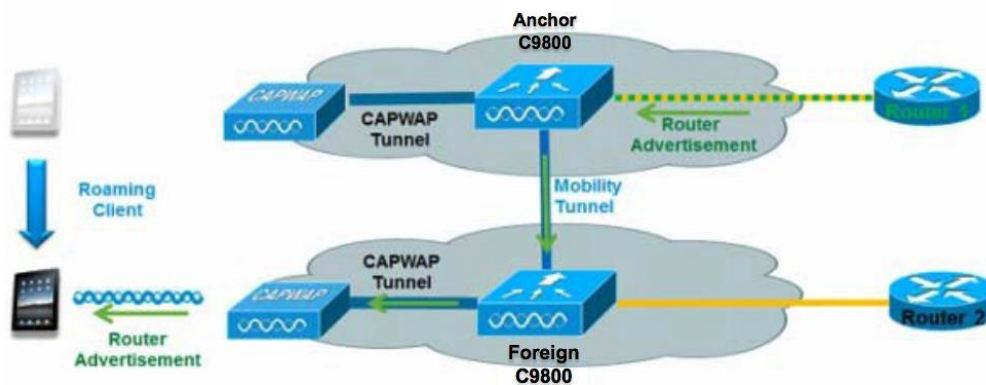


Enable Media stream Multicast-direct under WLAN> Advanced configuration tab.



Optionally, you can also globally enable Unicast Video Redirect and Multicast Direct for 5 GHz or 2.4 GHz radios in the wireless network as shown in the example below.

IPv6 Client Mobility



In order to deal with roaming IPv6 clients across controllers, the ICMPv6 messages such as NS, NA, RA, and RS must be dealt with specially to ensure that a client remains on the same Layer 3 network. The configuration for IPv6 mobility is the same as for IPv4 mobility and requires no separate software on the client side to achieve seamless roaming. The only required configuration is the controllers must be part of the same mobility group/domain.

The process of IPv6 client mobility across controllers is as follows:

1. If both controllers have access to the same VLAN the client was originally on, the roam is simply a Layer 2 roaming event where the client record is copied to the new controller and no traffic is tunneled back to the anchor controller.
2. If the second controller does not have access to the original VLAN the client was on, a Layer 3 roaming event will occur, meaning all traffic from the client must be tunneled via the Mobility Tunnel to the anchor controller.

3. In IOS-XE deployments, we support CAPWAP encrypted tunnel for IPv6 mobility tunnel.
 - a. To ensure that the client retains its original IPv6 address, the Router Advertisements from the original VLAN are sent by the anchor controller to the foreign controller where they are delivered to the client using L2 Unicast from the AP.
 - b. When the roamed client goes to renew its address via DHCPv6 or generate a new address via SLAAC, the Router Solicitation, Neighbor Advertisement, and Neighbor Solicitation packets continue to be tunneled to the original VLAN so that the client receives an IPv6 address that is applicable to that VLAN.

In C9800 IOS-XE release, the RA MGID and RA Mobility tables are created and maintained inside each controller. Both tables are implemented as process internal data structures. Each controller updates the RA MGID table whenever a client joins or leaves to keep track of the current topology of clients. In mobility scenarios, the anchor controller constructs the RA Mobility table to keep an up-to-date roamed client topology so that the anchor controller knows which foreign controllers need to receive RA packets. Entries are created in this table when the anchor controller sends the handoff message for roamed clients. Entries are deleted at the time the clients roam back to the anchor controller or disassociate from the network. On the foreign controllers, the RA MGID tables are used to deliver RAs to the visiting clients in the same way as the local clients.

On client deletion, the controller cleans up the client from the MGID or mobility table, removes the AP or mobility IFID from the multicast group if the client is the last one using this IFID, and requests AP to delete the client from the MGID list. After the cleanup, the controller will not forward RA packets to the client anymore.

AP maintains a list of MGIDs that it receives from the controller on client associations. MGIDs are mapped statically from the client VLAN IDs. AP has an MGID entry in the list as long as it has any client associated on the statically mapped VLAN. All such clients are recorded in the MGID entry. Since multicast over wireless is not robust, the MGID entry helps AP convert RA multicast packet into unicast packet for each wireless client in the entry. The controller informs AP to remove a client from the MGID entry at time of client deletion.

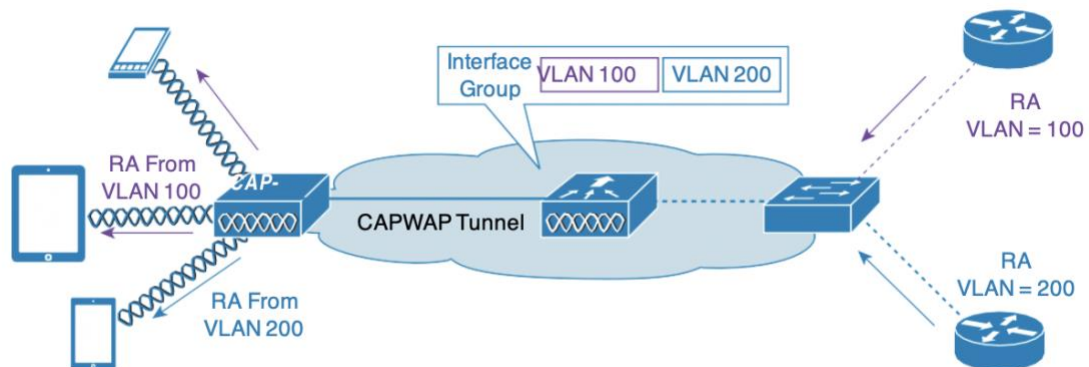
In the current implementations in IOS-XE, the RA MGID table is created and maintained within the controller. The requirement is that every VLAN is assigned a unique RA MGID value.

- Maximum number of VLANs supported – 4096
- MGID range reserved for RA – 8192 to 12287
- Number of RA-MGIDs – 8192

IPv6 IRCM – Inter Release Controller Mobility

AireOS based controllers such as 8540, 5520 and 3504 do not support IPv6 MC2UC feature. When client moves from the C9800 IOS-XE based controllers to the AireOS based controllers, it will be considered as IPv6 Multicast roaming, whereas roaming from AireOS to C9800 IOS-XE controllers, IPv6 MC2UC feature will kick in. When roaming from AireOS to IOS-XE controller, continuous IPv6 multicast stream is not guaranteed if stream belongs to an IPv6 MC2UC group. The behavior is dependent on configuration and available AP bandwidth at the time of roaming.

Support for Interface Groups



The interface groups feature allows an organization to have a single WLAN with multiple VLANs configured on the controller to permit load balancing of wireless clients across these VLANs. This feature is commonly used to keep IPv4 subnet sizes small while enabling a WLAN to scale to thousands of users across multiple VLANs in the group. To support IPv6 clients with interface groups, no additional configuration is required as the system automatically sends the correct router advertisement to the correct clients via L2 wireless unicast. By unicasting the router advertisement, clients on the same WLAN, but a different VLAN, do not receive the incorrect RA.

Note: For the reasons, explained above, it is not recommended to mix IPv4 and IPv6 dual stack clients in the same Interface Group.

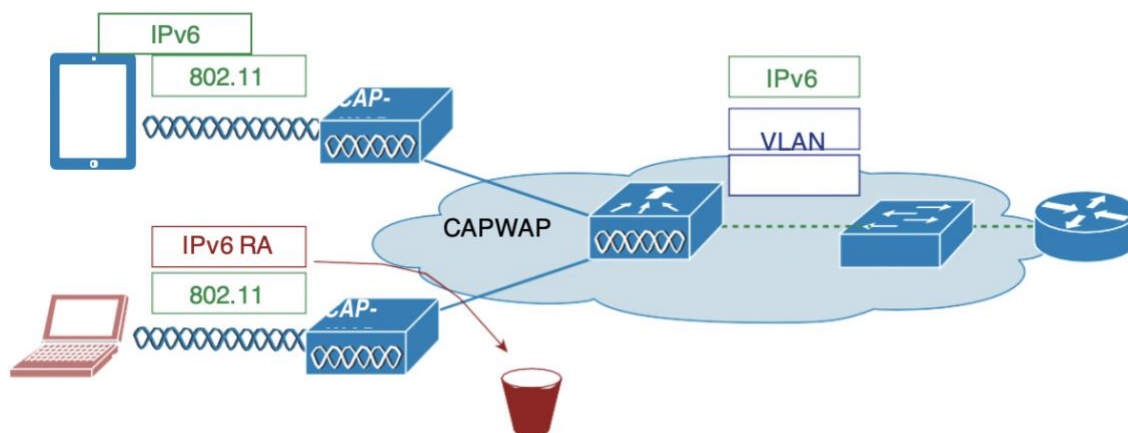
First Hop Security for IPv6 Clients

RA Management

When SLAAC is used, RA delivery to wireless clients is important for the clients to construct IPv6 addresses. Clients need to receive RA periodically for two reasons. First, the IPv6 prefix advertised in RA is valid for a period only. Clients expect new RAs before the prefix validity expires. Second, if a new prefix is configured, clients should receive new RAs in order to be informed on the new prefix. For these reasons, RAs should be delivered to both local and roamed clients continuously regardless of client mobility. In local case, RAs are sent to the clients through AP CAPWAP tunnels. In foreign case, RAs are forwarded by the anchor controller to the foreign controllers through mobility tunnels first, and then they are delivered to wireless clients through AP CAPWAP tunnels on the foreign controllers.

In IOS-XE Controllers the RA Guard needs to be configured. RAs from the Wired side are distributed to all clients on the VLAN (mcast-unicast). RAs from the Wireless clients is always send out/forwarded.

Router Advertisement Guard



The RA Guard feature increases the security of the IPv6 network by dropping router advertisements coming from wireless clients. Without this feature, misconfigured or malicious IPv6 clients could announce themselves as a router for the network, often with a high priority, which could take precedence over legitimate IPv6 routers.

RA guard is a functionality that protects the system from RA packets that are sent from non-trusted, unauthorized or misconfigured devices. This functionality analyzes the RAs and filters them out based on the device that sent them. By default, RA guard is disabled on C9800 but can be enabled through CLI. Once RA guard is enabled, the RA packets received on the wireless interfaces are dropped. It can be considered as a simple mechanism to prevent wireless clients from advertising themselves as routers on a wireless network.

Command Line Interface

Exec Configure Run Command Clear Copy Export

```
sh run int vlan 42
```

Control+X: Clear | Control+M: Switch Mode | Control+Return(J): Execute Command | Control+Y: Copy | Control+Shift+E: Export | Shift+Up Arrow(↑)/Down Arrow(↓): Lookup History

```
Tue Sep 18 2018 02:47:28 GMT-07:00 (Pacific Daylight Time)
-----
#sh wireless int summary
Wireless Interface Summary
Interface Name Interface Type VLAN ID IP Address IP Netmask MAC Address
-----
Vlan42 Management 42 9.8.42.90 255.255.255.0 00a3.8a23.a1eb
fd09:9:8:42::142/64
-----
Tue Sep 18 2018 02:35:24 GMT-07:00 (Pacific Daylight Time)
sh run int vlan 42
Building configuration...
Current configuration : 128 bytes
!
interface Vlan42
 ip address 9.8.42.90 255.255.255.0
 ipv6 address FD09:9:8:42::142/64
 ipv6 enable
 ipv6 nd ra suppress
end
```

```
C9800#sh run int vlan 42
Building configuration...

Current configuration : 128 bytes
!
interface Vlan42
 ip address 9.8.42.90 255.255.255.0
 ipv6 address FD09:9:8:42::142/64
 ipv6 enable
 ipv6 nd ra suppress all
end
```


ND Suppress

Neighbor solicitation to resolve wireless client link-layer address may result in a lot of overhead in the wireless network. In this case, the destination address in NS is the all-node multicast address. Multicasting NS to every wireless client consumes wireless resources heavily. To solve the wireless multicast problem, the controller maintains the wireless client address-to-MAC bindings. The controller can either respond to an NS on behalf of the wireless client or convert the multicast NS into a unicast one to the target client. Both solutions save the wireless resources as the unnecessary delivery to other clients is eliminated.

In IOS-XE controller the NS behavior is as follows:

1. Suppressed (NS multicast for known Wireless clients will be unicasted to the clients)
2. Unknown NS from Wired side is dropped
3. Unknown NS from Wireless is flooded on the Wired VLAN
4. Wireless client sends a DAD NS and if a DAD hit NA is proxied to the originating client

DHCPv6 Server Guard

The DHCPv6 Server guard feature prevents wireless clients from handing out IPv6 addresses to other wireless clients or wired clients upstream. To prevent DHCPv6 addresses from being handed out, all DHCPv6 advertise packets from wireless clients are dropped. This feature operates on the controller, requires no configuration and is enabled automatically.

IPv6 Source Guard

The IPv6 source guard feature prevents a wireless client spoofing an IPv6 address of another client. This feature is analogous to IPv4 source guard. IPv6 source guard is enabled by default. Custom defined pre-auth ACL will be sent from the Controller to the FlexConnect APs.

IPv6 Access Control Lists


In order to restrict access to certain upstream wired resources or block certain applications, IPv6 Access Control lists can be used to identify traffic and permit or deny it. IPv6 Access Lists support the same options as IPv4 Access Lists including source, destination, source port, and destination port (port ranges are also supported).

In IOS-XE release Pre-authentication ACL is supported for LWA and EWA (Local and External Web Auth).

All the custom ACLs must be mapped in the Flex profile. Only the custom ACL definitions will be pushed to AP apart from the generated default ACLs. Custom pre-authentication ACLs are mapped under WLAN profile. Whereas, custom post-authentication ACLs are mapped under default policy profile. All post-authentication ACLs are configured under default Flex profile.

Network Resource Efficiency for IPv6 Clients

Neighbor Discovery Caching



The IPv6 neighbor discovery protocol (NDP) utilizes Neighbor Advertisement (NA) and Neighbor Solicitation (NS) packets in place of ARP to allow IPv6 clients to resolve the MAC address of other clients on the network. The NDP process initially uses multicast addresses to perform address resolution. This process consumes valuable wireless airtime because the multicast addresses are sent.

To increase the efficiency of the NDP process, neighbor discovery caching allows the controller to act as a proxy and responds back to the NS queries that it can support address resolution and duplicate address detection. Neighbor discovery caching is made possible by the underlying neighbor binding table present in the controller. The neighbor binding table keeps track of each IPv6 address and its associated MAC address. When an IPv6 client attempts to resolve another client's link-layer address, the neighbor solicitation packet is intercepted by the controller that responds back with a neighbor advertisement packet to all the clients in the network segment.

NA behavior on the IOS-XE controller is as follows:

- Unicast is treated as Unicast (After Security Validations, i.e. IP Theft and Priority)
- Multicast
- From Wireless is flooded on Wired VLAN (but not sent on Wireless)
- From Wired Un-known NA is currently being forwarded on Wireless
- **Mobility:** If the controllers are in mobility domain and share the same L2 client VLAN the NA from other wireless clients on the other controllers is still treated as unknown NA

Router Advertisement Throttling

Router Advertisement (RA) throttling allows the controller to enforce rate limiting of RAs headed towards the wireless network. By enabling RA throttling, routers that are configured to send RAs frequently (every 3 seconds) can be trimmed back to a minimum frequency that will still maintain IPv6 client connectivity. This allows airtime to be optimized by reducing the number of multicast packets that must be sent. In all cases, if a client sends a Router Solicitation (RS), then an RA will be allowed through the controller and unicast to the requesting client. This is to ensure that new clients or roaming clients are not negatively impacted by RA throttling.

Note: When RA throttling occurs, only the first IPv6 capable router is allowed through. For networks that have multiple IPv6 prefixes being served by different routers, RA throttling must be disabled.

IPv6 Router Configuration

The IPv6 router is usually configured on a switch connected to C9800. The client VLAN should be configured on the router with the IPv6 settings. In addition, the physical link between the router and C9800 should allow the client VLAN traffic to pass through. If the router has other links to reach APs, which is possible since APs join C9800 through an intermediate network, the client VLAN should be blocked on these links to avoid sending RA traffic to APs directly through the intermediate network. By design, RAs should only be sent by C9800 to its APs. Below are example configurations on IPv6 router.

```
ipv6 unicast-routing
interface GigabitEthernet1/0/1
  switchport trunk allowed vlan 13,14,17,37,70,43,44,80-82,99,113,120,122,123,128
  switchport trunk allowed vlan add 70,129,160,161,170
  switchport mode trunk
interface Vlan160
  description "Client vlan - vC9800 Int"
  ip address 160.160.0.1 255.255.0.0
  ipv6 address FE80:20:22:160::1 link-local
  ipv6 address 2200:20:22:160::1/64
  ipv6 enable
```

Configuration for Wireless IPv6 Client Support

Configuring Global Controller (Screen Shots from IOS-XE Release)

Complete these steps:

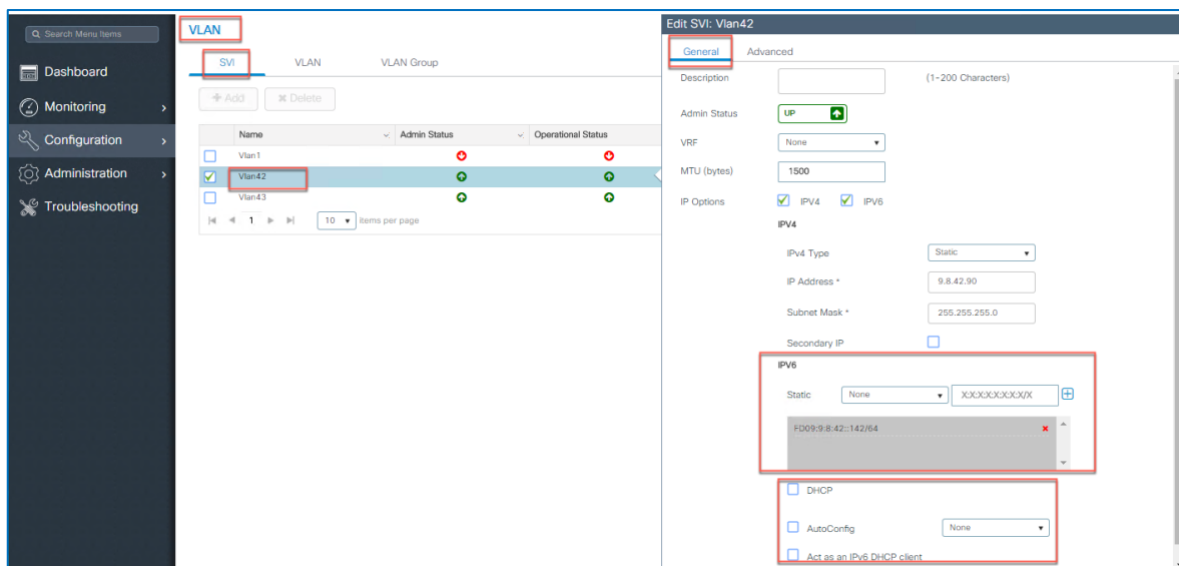
Interface configuration

The link between C9800 and IPv6 router should allow the client VLAN traffic so that C9800 can receive and forward RAs. For this purpose, C9800 has a configuration like in example below.

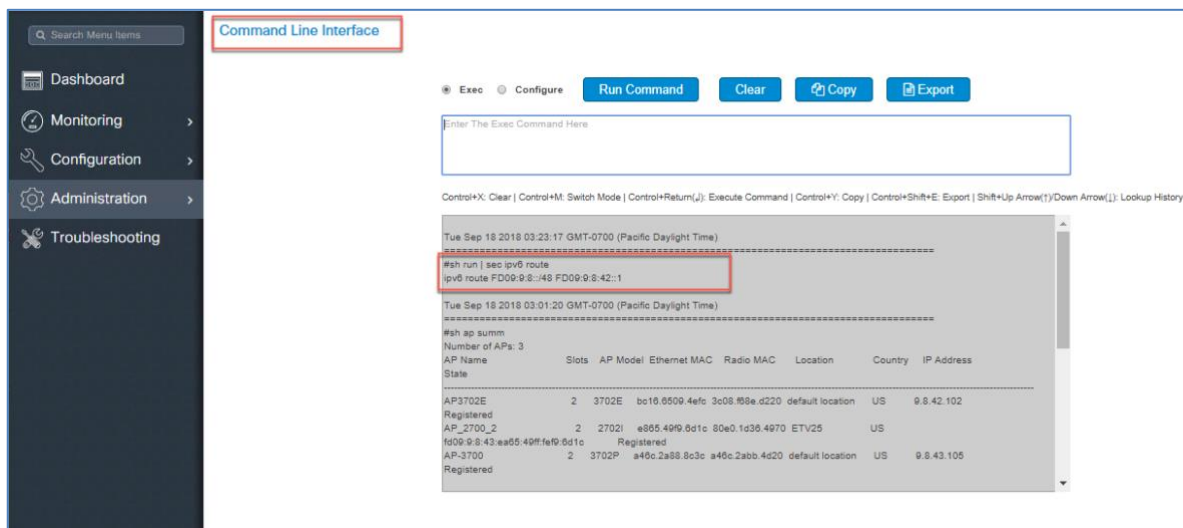
```
interface GigabitEthernet1
  switchport trunk allowed vlan 42,129,160,161
  switchport mode trunk
```

Step 1: Configure via CLI or WebUI one Global Management Interface with IPv6 address:

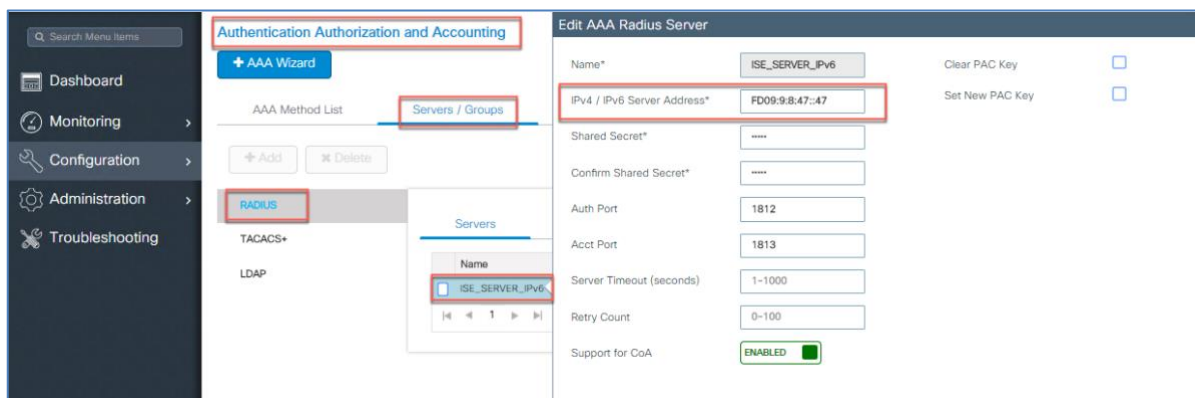
Configure the IPv6 address of the Management VLAN interface



Step 2: Configure IPv6 Static Route on the controller - This is a very important step to ensure all wireless client traffic has only one default route in a pure IPv6 deployments.



Step 3: Configure Authentication and Authorization Servers with IPv6 address



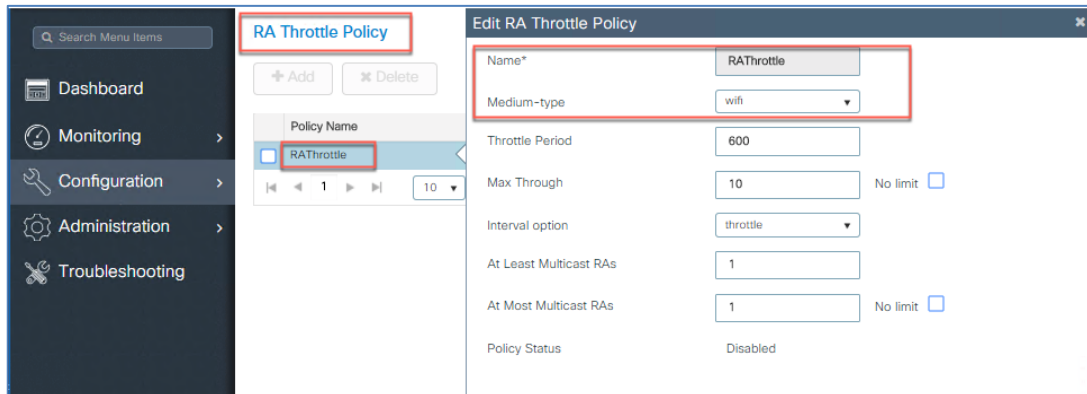
Step 4: Configure RA Throttle Policy via CLI or WebUI. Once the RA Throttle policy is configured apply it on management VLAN as shown in the example below.

To throttle excessive RAs on C9800, these configurations are needed.

```

ipv6 nd ra-throttler policy ndrapol
 throttle-period 600
 max-through 10
 allow at-least 1 at-most 1
vlan configuration 42
 ipv6 nd <ra-throttle> attach-policy <ndrapol>

```



Step 5: Configure IPv6 Multicast on the controller via CLI

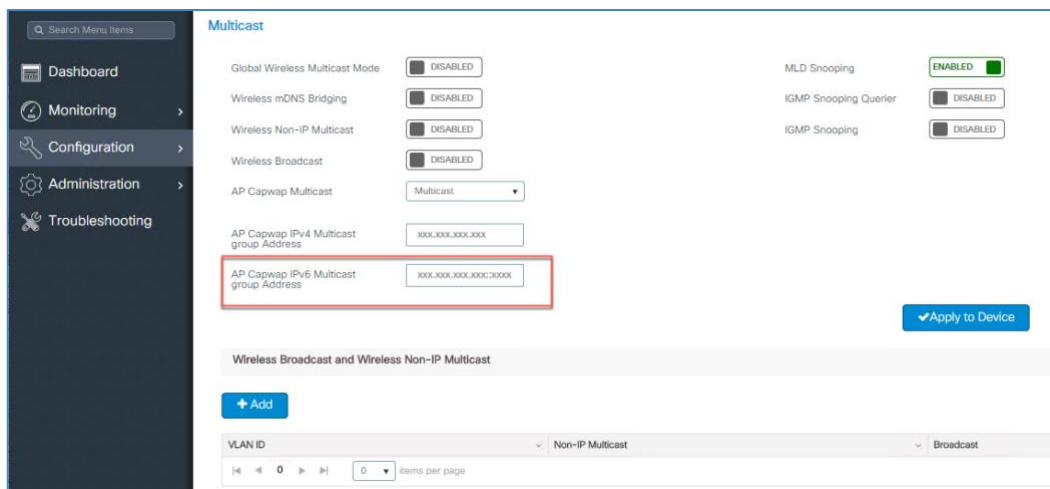
To configure MoM, the following line is needed. To turn off MoM, use 0.0.0.0 as the multicast IP address.

```

wireless multicast ipv6 <MOM IPv6 address>
 wireless multicast <0.0.0.0>

```

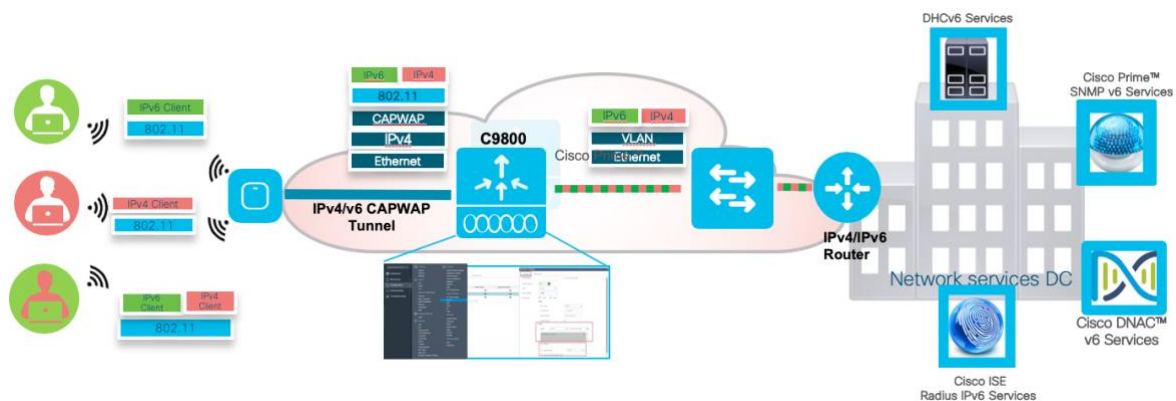
Same can be done from the WebUI



Infrastructure IPv6 Support in Cisco IOS XE Amsterdam 17.1

This section provides a set of instructions to effectively configure native IPv6 features based on C9800 IOS-XE.

The diagram below shows Wireless Infrastructure Configuration Workflow on the C9800 and IOS-XE.

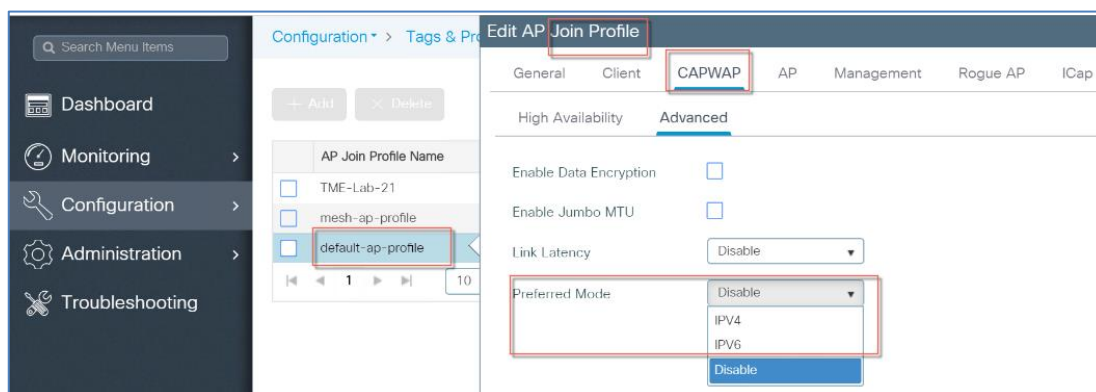


The following are the Infrastructure IPv6 configuration items:

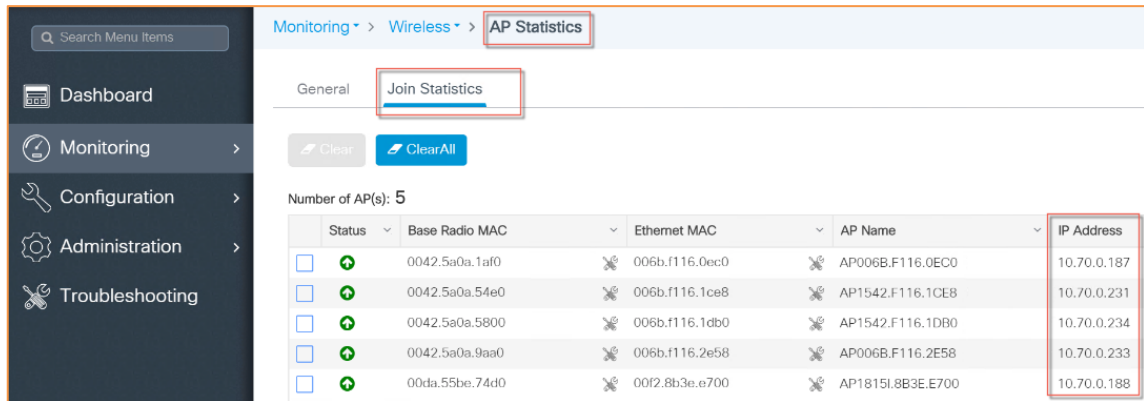
- Support Configuration of Preferred IPv4 or IPv6 Dual Stack Mode APs
- Support of the HA Mode with Primary/Secondary/Tertiary Controllers
- Support of IPv6 Mobility Peers and Guest Networks
- Support of the IPv6 Multicast
- IPv6 support of the TFTP/FTP, NTP, DNS and SNMP servers
- IPv6 Ping and Trace Route Support
- IPv6 CAPWAP UDP-Lite Support

Step 1: Configure Global Preferred AP Joint Profile for IPv6

Note: This step is required for pure IPv6 deployments



To view in what mode the AP has joined the controller, check under Monitoring > AP Statistics > Join Statistics.



Monitoring > Wireless > AP Statistics

General Join Statistics

Clear ClearAll

Number of AP(s): 5

Status	Base Radio MAC	Ethernet MAC	AP Name	IP Address
🟢	0042.5a0a.1af0	006b.f116.0ec0	AP006B.F116.0EC0	10.70.0.187
🟢	0042.5a0a.54e0	006b.f116.1ce8	AP1542.F116.1CE8	10.70.0.231
🟢	0042.5a0a.5800	006b.f116.1db0	AP1542.F116.1DB0	10.70.0.234
🟢	0042.5a0a.9aa0	006b.f116.2e58	AP006B.F116.2E58	10.70.0.233
🟢	00da.55be.74d0	00f2.8b3e.e700	AP1815I.8B3E.E700	10.70.0.188

The same can be verified via the CLI commands

```
C9800>show ap profile default-ap-profile detailed
```

```
eWLC#sh ap prof default-ap-profile det
AP Profile Name       : default-ap-profile
Description           : default ap profile
Stats Timer          : 180
VLAN Tagging         : DISABLED
Link Latency         : DISABLED
Data Encryption      : DISABLED
LED State            : ENABLED
NTP server           : 0.0.0.0
Country Code         : Not Configured
Jumbo MTU            : DISABLED
24ghz Report Interval : 90
5ghz Report Interval  : 90
POE :
  PreStandard 802.3af Switch : DISABLED
  Power Injector State       : DISABLED
  Power Injector Selection   : Unknown
  Injector Switch Mac       : Not Configured
Device Management :
  Telnet                    : DISABLED
  SSH                       : DISABLED
User Management :
  Username                  : Not Configured
TCP MSS :
  Adjust MSS                 : ENABLED
  TCP Adjust MSS             : 1250
CAPWAP Timer :
  Heartbeat Timeout         : 30
  Discovery Timeout         : 10
  Fast Heartbeat Timeout    : 0
  Primary Discovery Timeout : 120
  Primed Join Timeout       : 0
Retransmit Timer :
  Count                     : 5
  Interval                   : 3
Login Credentials :
  Local Credentials         : DISABLED
  Dot1x Credentials        : DISABLED
  Local Username            :
```



```

Ap eap auth info :
  Dot1x EAP Method      : EAP-FAST
  LSC AP AUTH STATE    : CAPWAP DTLS
Syslog :
  Facility Value       : FACILITY_KERN
  Host                 : 255.255.255.255
  Log Level            : SYSLOG_LEVEL_INFORMATION
Backup Controllers :
  Fallback              : ENABLED
  Primary Backup Name  : Not Configured
  Primary Backup IP    :
  Secondary Backup Name : Not Configured
  Secondary Backup IP  :
Hyperlocation :
  Admin State          : DISABLED
  PAK RSSI Threshold Detection: -100
  PAK RSSI Threshold Trigger : 10
  PAK RSSI Threshold Reset : 8
Halo BLE Beacon :
  Interval             :
  TX Power             :
  Enabled              : Unknown
  Apply Global         : Unknown
Group NAS Id         : Not Configured
CDP                  : ENABLED
TFTP Downgrade :
  IP Address           : 0.0.0.0
  Filename             : Not Configured
  Time Limit           : 0
AP packet capture profile : Not Configured
AP trace profile     : Not Configured
Mesh profile name    : default-mesh-profile
Extended Module      : DISABLED
USB Module           : ENABLED
Persistent SSID Broadcast : DISABLED
Preferred Mode       : IPV6
Lawful-Interception : DISABLED
  LI timer           : 60
eWLC#

```

Same is shown with the execution of the CLI command

```
C9800>show ap summary
```

```

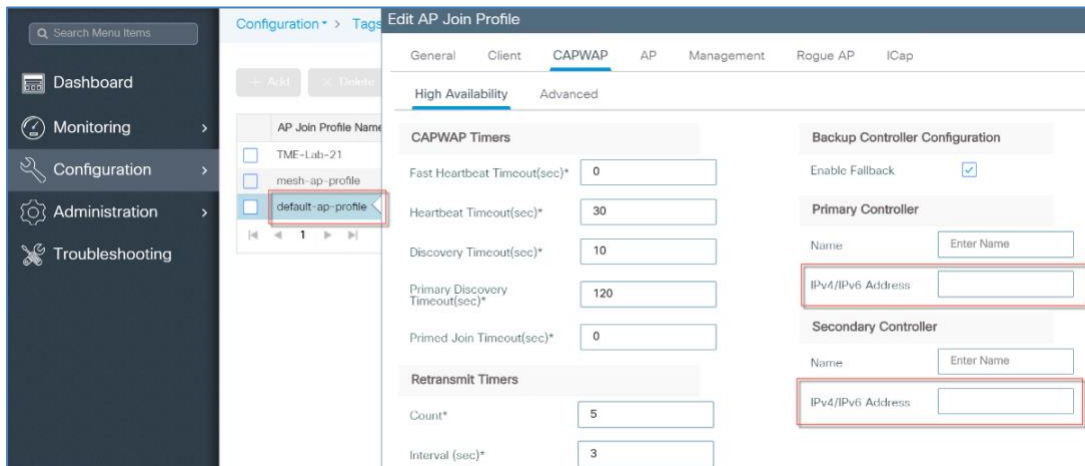
smonaniewlc2#sh ap summary
Number of APs: 7

```

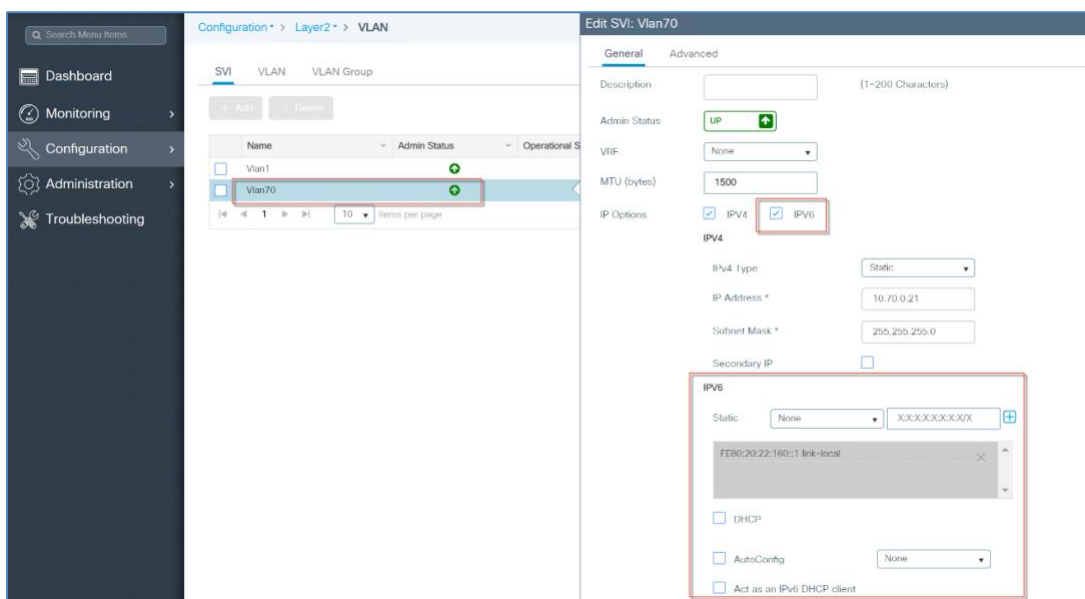
AP Name	Slots	AP Model	Ethernet MAC	Radio MAC	Location	Country	IP Address	State
AP00A6.CA36.F2A	3	3802E	00a6.ca36.f2a	006b.f107.70e0	default location	IN	fd09:9:2:58:6da3:e57e:7a73:d556	Regist
APb4de.3196.caec	2	2702I	b4de.3196.caec	4c77.6d83.6c90	default location	IN	9.2.52.62	Regist
AP4C77.6D67.F260	3	1815T	4c77.6d67.f260	4c77.6dca.b300	default location	US	fd09:9:2:60:2457:dbb2:fa9:22c5	Regist
AP0042.68C5.CBB4	3	3802I	0042.68c5.cbb4	70db.9813.c760	default location	IN	fd09:9:2:49:a187:bb65:e49b:9819	Regist
AP80e0.1d70.d478	2	3702I	80e0.1d70.d478	80e0.1d7b.83c0	default location	US	fd09:9:2:58:3125:ec9f:80b8:317e	Regist
APADf8.49E4.0448	3	1852E	a0f8.49e4.0448	a0f8.49e4.6a40	default location	IN	9.2.52.15	Regist
APf8c2.883a.2344	2	1702I	f8c2.883a.2344	f8c2.884a.b600	default location	US	fd09:9:2:60:a0fc:e27b:fd0:3dca	Regist

Step 2: Configure CAPWAP HA Primary and Secondary Controllers IPv6 addresses

Note: This step is required for pure IPv6 deployments with HA

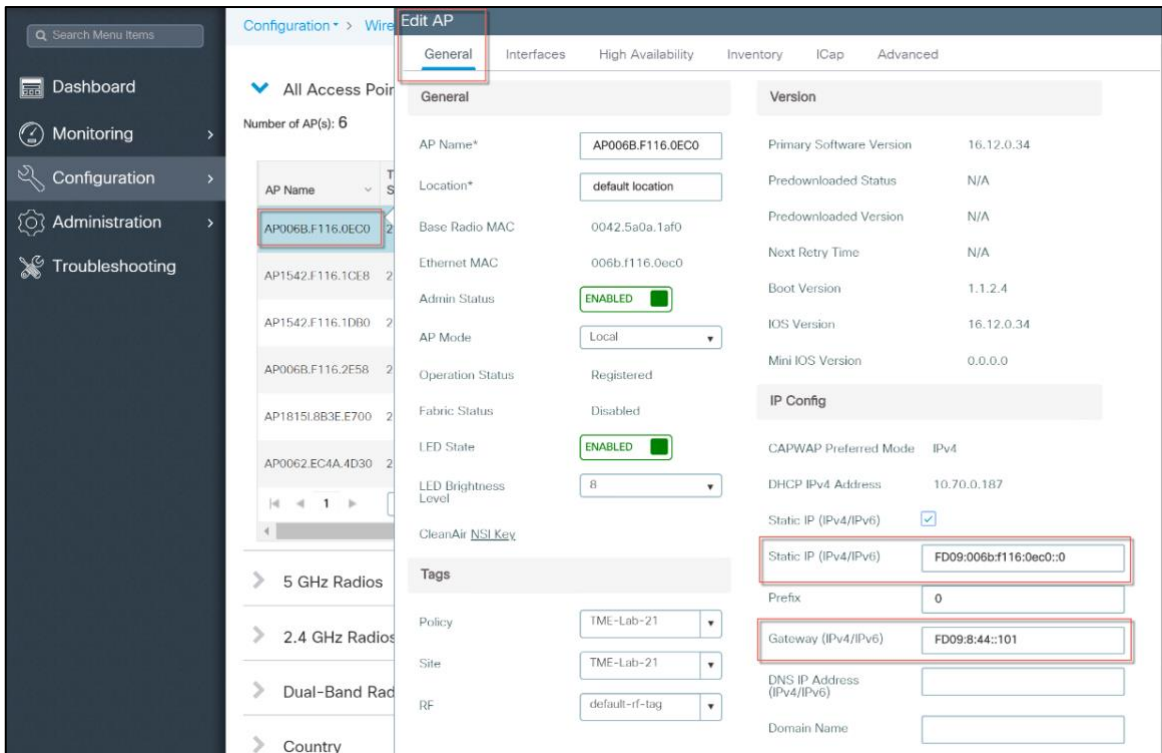


Step 3: Configure IPv6 on the VLAN for the wireless management interface



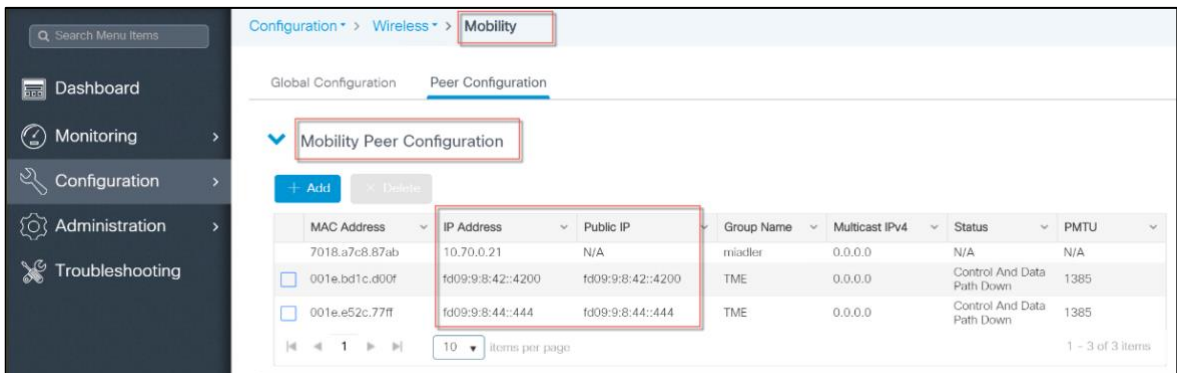
Step 4: Configure Individual AP SLAAC IPv6 address if DHCPv6 is not available or not desired.

Note: This step is required for pure IPv6 deployments with AP static IPv6 address



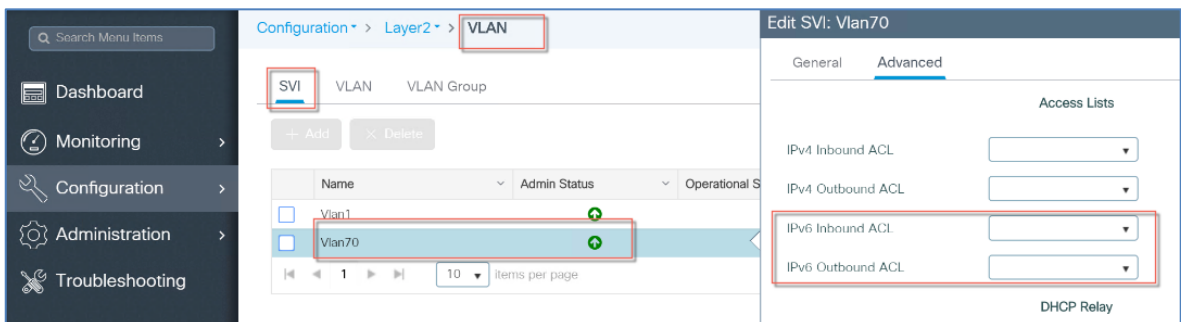
Step 5: Configure Wireless Mobility Peers IPv6 addresses

Note: This step is required for pure IPv6 deployments with Guest access or Mobility



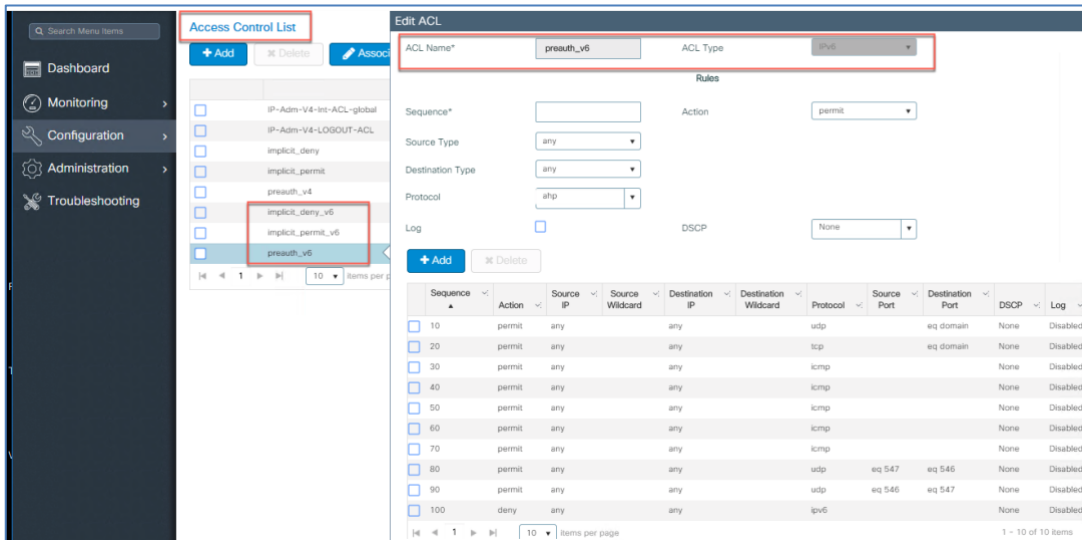
Step 6: Configure IPv6 Inbound and Outbound ACL on the Wireless Management VLAN

Note: This step is required for IPv6 deployments with IPv6 ACLs

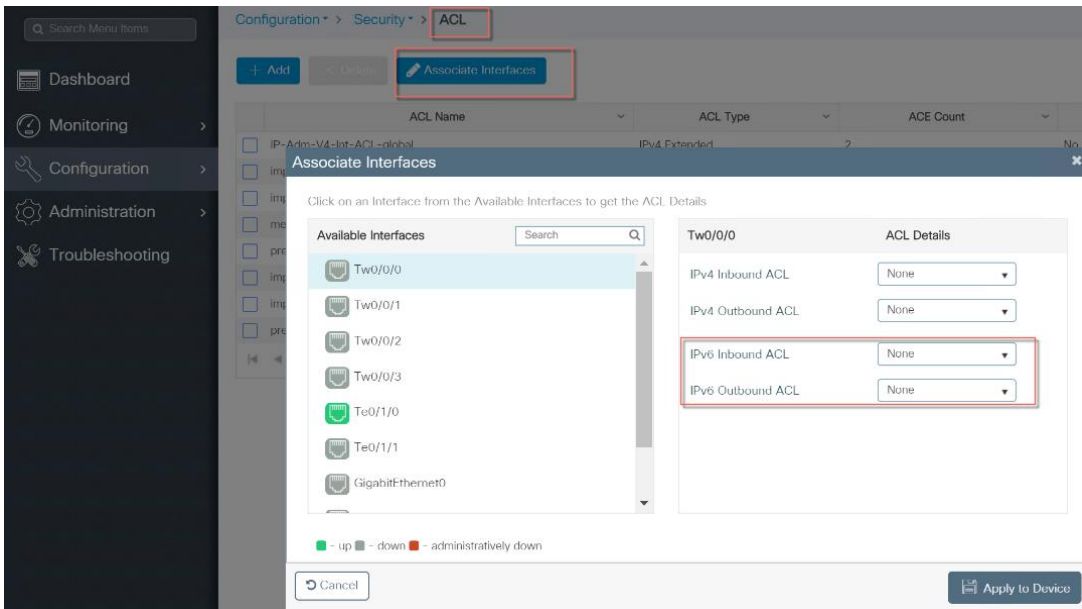


Step 7: Configure IPv6 Implicit-deny, Implicit-permit or pre-auth-v6 ACLs

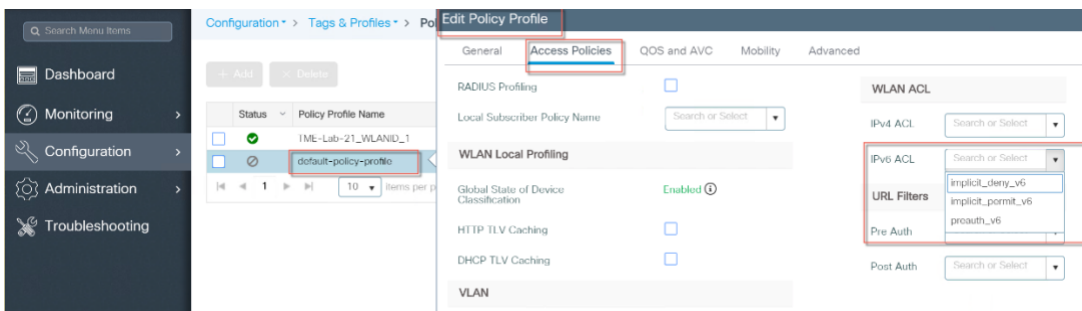
Note: This step is required for IPv6 deployments with Client VLAN ACLs



After ACL have been defined, Associate the desired Interfaces with the configured ACLs

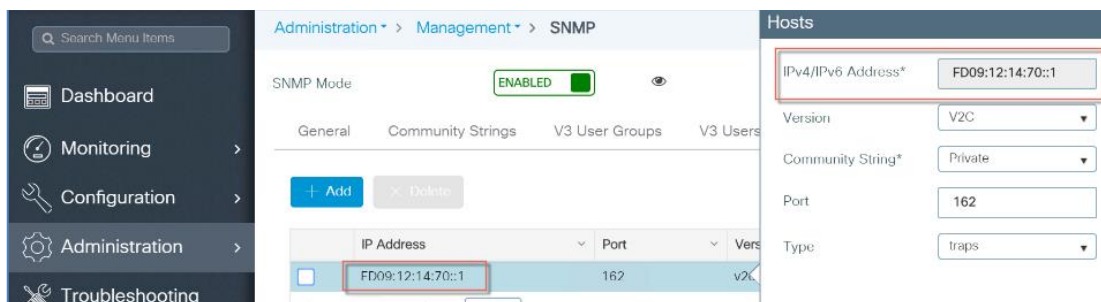


Step 8: Configure Policy Profile with WLAN ACL as shown below

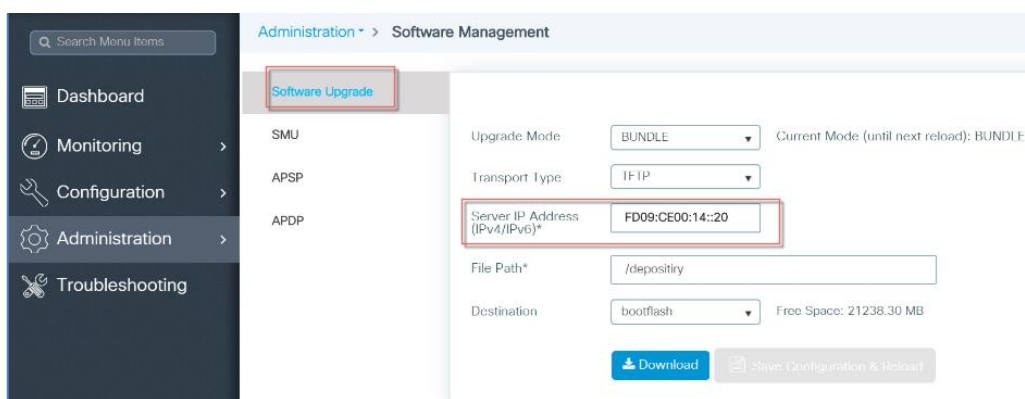


Step 9: Configure Controller SNMP IPv6 address

Note: This step is required for pure IPv6 deployments

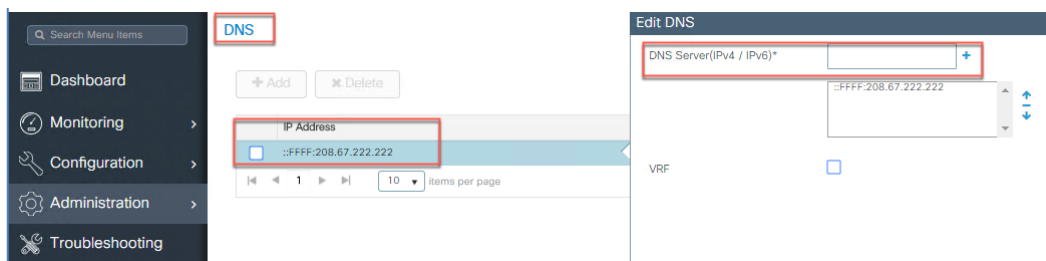


Step 10: Configure Source IPv6 address for Software Upgrade server



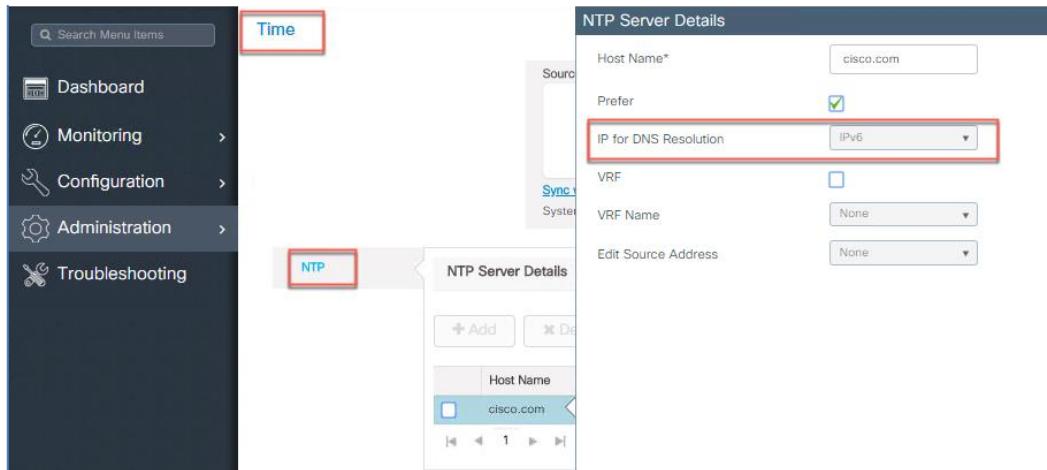
Step 11: Configure IPv6 address for DNS Server

Note: This step is required for pure IPv6 deployments with IPv6 DNS server



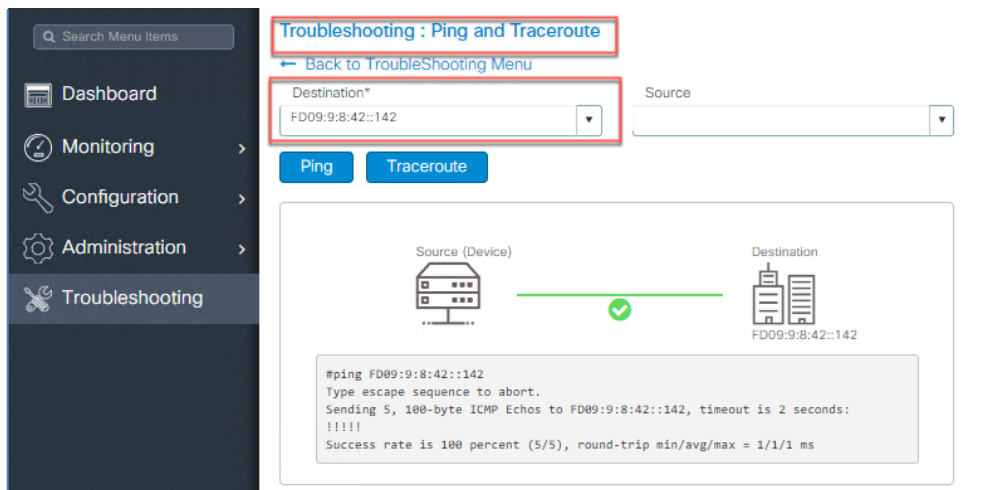
Step 12: Configure IPv6 address for NTP Server

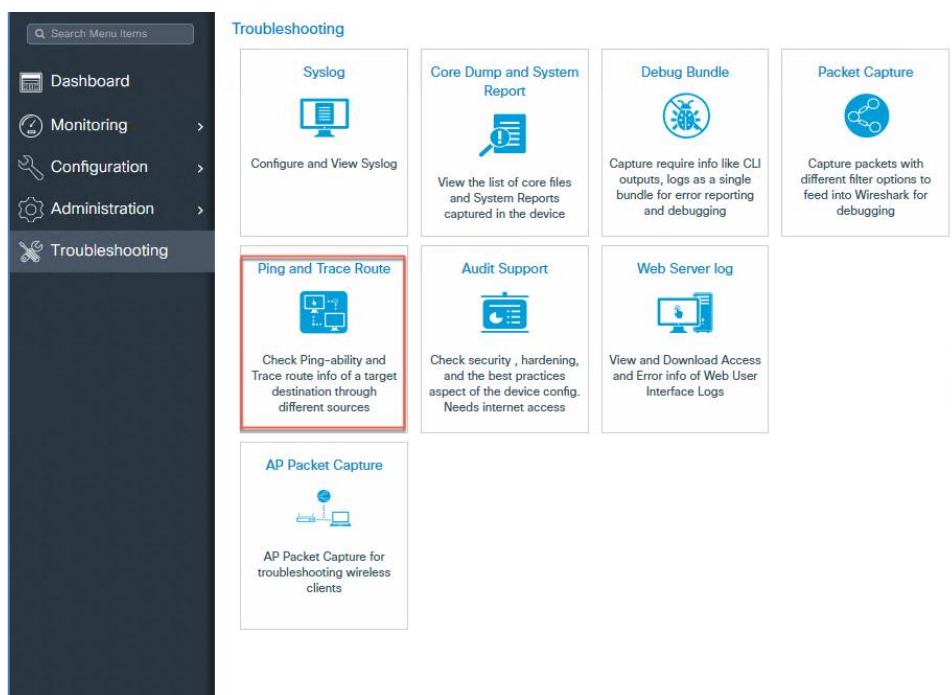
Note: This step is required for pure IPv6 deployments with IPv6 NTP server



Step 13: IPv6 Ping and Trace Route

Note: This command is required for pure IPv6 deployments





Configuring IPv6 CAPWAP UDP-Lite Support

The UDP Lite support feature is an enhancement to the existing IPv6 functionality to support the UDP Lite protocol. This feature is only applicable for IPv6 addresses of the controller and APs. IPv6 mandates complete payload checksum for UDP that results in performance implications. The UDP Lite feature minimizes the performance impact on the controller and AP by restricting the checksum calculation coverage for the UDP Lite header of 8-bytes only.

The use of UDP Lite feature impacts intermediate firewalls to allow UDP Lite protocol (protocol ID of 136) packets. Existing firewalls may not provide the option to open specific ports on UDP Lite protocol. In such cases, the administrator must open up all the ports on UDP Lite.

Note: Mobility IPv6 tunnels do not support the UDP Lite feature.

Below is an example of the UDP-Lite configuration from the CLI

```
C9800(config)# ap profile default-ap-profile
C9800(config-ap-profile)# capwap udplite
```

To verify the CAPWAP UDP Lite status, use the following command:

```
Device# show ap profile name default-ap-profile detailed
CAPWAP UDP-Lite : ENABLED
Lawful-Interception : ENABLED
LI timer : 60
AWIPS : DISABLED
AWIPS Forensic : Unknown
Client RSSI Statistics
Reporting : ENABLED
Reporting Interval : 30 seconds
```

Enabling IPv6 on Your IOS Infrastructure Device

Enabling IPv6 on an individual infrastructure device to which wireless controller will be connected. Refer to the following Cisco documentations for configuring IPv6 on other IOS-XE based devices.

- Cisco Catalyst 9800 Series Wireless Controller Software Configuration Guide, Cisco IOS XE Gibraltar 17.1.x at

<https://www.cisco.com/c/en/us/support/wireless/catalyst-9800-series-wireless-controllers/products-installation-and-configuration-guides-list.html>

- IPv6 Implementation Guide at <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-2mt/ipv6-15-2mt-book.html>
- IPv6 Routing chapter at https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750/software/release/12-2_55_se/configuration/guide/scg3750/swipv6.html

Glossary

- AP - Access Point
- ARP - Address Resolution Protocol
- CAPWAP - Control And Provisioning of Wireless Access Points
- CLI - Command Line Interface
- CPP - Cisco Packet Processor
- CWDB - Common Wireless Database
- DAD - Duplicate Address Detection
- DHCP - Dynamic Host Configuration Protocol
- CWC 9800 - Elastic Wireless LAN Controller
- IOSd - IOS daemon
- IFID – Interface Identifier
- IP - Internet Protocol
- IPC - Inter-Process Communication
- MAC - Medium Access Control
- MC2UC - Multicast to Unicast
- MGID - Multicast Group ID
- MOM - Multicast-Over-Multicast
- MOU - Multicast-Over-Unicast
- NA - Neighbor Advertisement

Glossary

- ND - Neighbor Discovery
- NDP - Neighbor Discovery Protocol
- NGWC - Next Generation Wiring Closet
- NS - Neighbor Solicitation
- ODM - Operational Data Manager
- RA - Router Advertisement
- RS - Router Solicitation
- RP – Redundancy Port
- SISF - Switch Integrated Security Module
- SLAAC - Stateless Auto Address Configuration
- VLAN - Virtual LAN
- WCM - Wireless Controller Module

Legal Information

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies are considered un-Controlled copies and the original on-line version should be referred to for latest version.

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco Trademark

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Cisco Copyright

© 2020 Cisco Systems, Inc. All rights reserved.