

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

First Published: March 12, 2020
Last Revised: February 24, 2025

Table of Contents

Preface *4*

Mesh Network Components *4*

Cisco Catalyst 9800 Series Wireless Controllers **5**

Mesh Access Points **5**

Access Point Roles 5
Network Access 6
Mesh Network Segmentation 6
Cisco Outdoor Mesh Access Points 7
Flexible Antenna Port Configuration 10
Cisco Wireless Controllers 11
Cisco Prime Infrastructure and Cisco DNA Center 11
Architecture 11

Mesh Deployment Modes *15*

Wireless Mesh Network **15**

Wireless Backhaul at 5 and 2.4 GHz **15**

Universal Access 16

Point-to-Multipoint Wireless Bridging **16**

Point-to-Point Wireless Bridging **16**

Mesh Daisy Chaining **17**

Design Considerations *22*

Wireless Mesh Constraints **22**

Wireless Backhaul Data Rate 22

Controller Planning **23**

VM Ware Specifications for IOS-XE 16.10-17.1 24
KVM Specifications for IOS-XE rel 16.10-17.1 24

Site Preparation and Planning *24*

Site Survey **24**

Pre-Survey Checklist **25**

Outdoor Site Survey **25**

Determining a Line of Sight **26**

Weather **26**

Fresnel Zone **26**

Fresnel Zone Size in Wireless Mesh Deployments **27**

Co-Channel Interference **27**

Wireless Mesh Network Coverage Considerations **28**

Background scanning on Wave-2 APs 28
Mesh Convergence 28
Cell Planning and Distance 29
Assumptions for the Cisco Range Calculator 31

Collocating Mesh Access Points 33
DFS and None-DFS Channel Scan 34
Wireless Propagation Characteristics 34
CleanAir and RRM in Mesh 35
CleanAir AP Modes of Operation 35
Pseudo MAC (PMAC) and Merging 36
Event Driven Radio Resource Management and Persistence Device Avoidance 36
CleanAir Access Point Deployment Recommendations 36
CleanAir Advisor 37
Wireless Mesh Mobility Groups 37
Multiple Controllers 37
Increasing Mesh Availability 38
Multiple RAPs 39

Connecting and Monitoring Cisco Mesh Access Points in the Network 39

Adding Mesh Access Points to the Mesh Network 39
MAC Authorization 39

Monitoring Mesh Access Points to the Mesh Network 48

Command Line Configuration for Mesh Access Points 54
Example of the Mesh configuration 55

Air Time Fairness (ATF) in Mesh Deployments 55

Introduction to Air Time Fairness (ATF) 55

Cisco Air Time Fairness (ATF) Use Cases 57
Public Hotspots (Stadium/Airport/Convention Center/Other) 57
Education 57
Enterprise or Hospitality or Retail 57
Time Shared Managed Hotspot 57

ATF Functionality and Capabilities 57

Air Time Fairness in Mesh Deployments IOS-XE 17.1 58

Pre-requisite and Supported Features in IOS-XE 17.1.1 58

ATF on Mesh Feature Overview 58

Mesh ATF Modes of Operation 59

Configuring ATF on Mesh 60

Mesh/Bridge Mode ATF Configuration Sample 65

Preface

This document provides design and deployment guidelines for the implementation of secure enterprise, campus, and metropolitan Wi-Fi networks within Cisco wireless mesh networking solution, a component of the Cisco Catalyst 9800 architecture with IOS-XE release 17.1.

Mesh networking employs Cisco Aironet 1540,1560 and 1570 Series outdoor mesh access points; Cisco Catalyst 9800 wireless controller (C9800), and Cisco DNA Center1.3 to provide scalable, central management, and mobility between indoor and outdoor deployments. Control and Provisioning of Wireless Access Points (CAPWAP) protocol manages the connection of mesh access points to the network.

End-to-end security within the mesh network is- supported by employing Advanced Encryption Standard (AES) encryption between the wireless mesh access points and Wi-Fi Protected Access 2 (WPA2) clients. This document also outlines radio frequency (RF) components to consider when designing an outdoor network.

The features described in this document are for the following outdoor AP products:

- Cisco Aironet 1560 (1562) series outdoor mesh access points
- Cisco Aironet 1540 (1542) Series outdoor mesh access points
- Cisco Aironet 1572 (1572) Series outdoor mesh access points

Note: 1572 series Outdoor AP is not supported in the DNAC.

- Cisco Aironet Wave-1 indoor APs: 1700, 2700 and 3700 series.
- Cisco Aironet Wave-2 indoor APs: 1815i, 1815m, 1830,1850, 2800, 3800 and 4800 series
- Mesh features in Cisco C9800 wireless controller
- Mesh features in Cisco PI 3.7 and DNAC rel 1.4
- Airtime Fairness (ATF) in Mesh Deployments

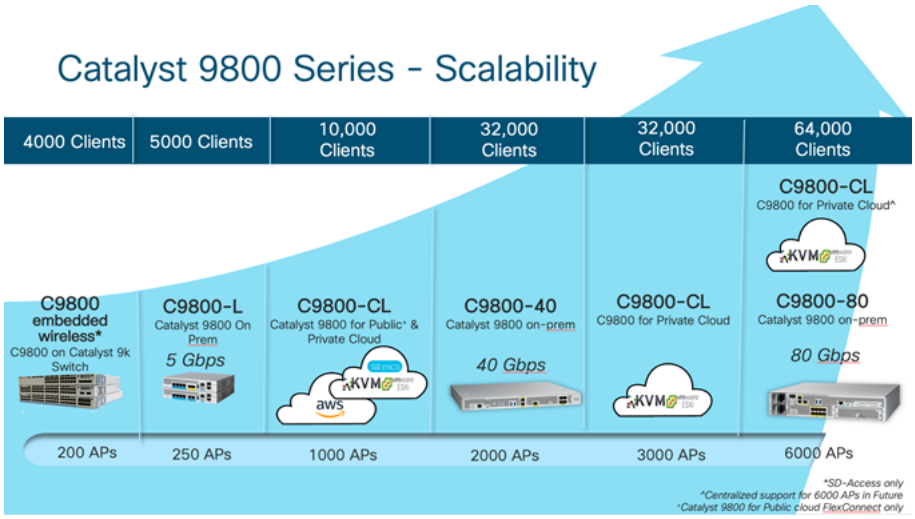
Mesh Network Components

This section describes the mesh network components.

The Cisco wireless mesh network has four core components:

- Cisco C9800 and IOS-XE 17.1
- Cisco Aironet and Catalyst series access points
- Cisco PI and Cisco DNA Center
- Mesh software architecture

Cisco Catalyst 9800 Series Wireless Controllers



Mesh Access Points

Access Point Roles

Mesh networking employs Cisco Aironet outdoor mesh access points and indoor mesh access points along with the Cisco Wireless Controller, and Cisco PI and Cisco DNA Center to provide scalable, central management, and mobility between indoor and outdoor deployments. Control and Provisioning of Wireless Access Points (CAPWAP) protocol manages the connection of mesh access points to the network.

For Mesh mode, Access Point should be configured with Bridge mode. Access points within a mesh network operate in one of the following two ways:

1. Root access point (RAP)

2. Mesh access point (MAP)

Note: All access points are configured and shipped as mesh access points. To use an access point as a root access point, you must reconfigure the mesh access point to a root access point. In all mesh networks, ensure that there is at least one root access point.

While the RAPs have wired connections to their controller, MAPs have wireless connections to their controller via the RAP or another MAP. MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n/ac/ax radio backhaul. MAPs use the Cisco Adaptive Wireless Path Protocol (AWPP) to determine the best path through the other mesh access points to the controller.

End-to-end security within the mesh network is supported by employing Advanced Encryption Standard (AES) encryption between the wireless mesh access points and Wi-Fi Protected Access 2 (WPA2) and WPA3 clients. A mesh access point establishes AWPP link with a parent Mesh AP which is already connected to the Controller before starting CAPWAP discovery.

Note: The RAP or MAP does not generate Bridge Protocol Data Unit (BPDU) itself. However, the RAP or MAP forwards the BPDU to upstream devices if the RAP or MAP received the BPDU from its connected wired or wireless interface across the network.

This figure shows the relationship between RAPs and MAPs in a mesh network

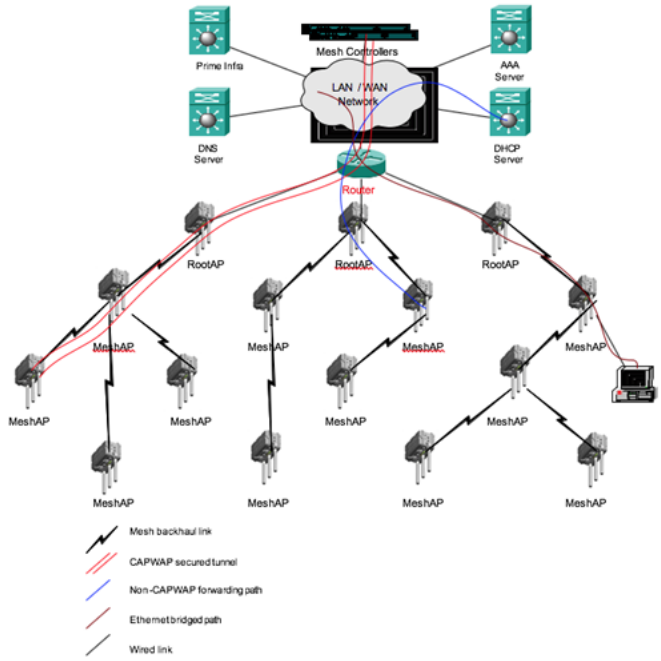


Figure 1 Simple Mesh Network Hierarchy

Network Access

Wireless mesh networks can simultaneously carry two different traffic types. They are as follows:

- Wireless LAN client traffic
- MAP Ethernet port traffic

Wireless LAN client traffic terminates on the controller, and the Ethernet traffic terminates on the Ethernet ports of the mesh access points

Access to the wireless LAN mesh for mesh access points is managed by the following authentication methods:

- MAC authentication—Mesh access points are added to a database that can be referenced to ensure they are provided access to a given controller and mesh network.
- External RADIUS Authentication—Mesh access points can be externally authorized using a RADIUS server such as Cisco ISE that supports the client authentication type of Extensible Authentication Protocol-FAST (EAP-FAST) with certificates and WPA2/PSK on the C9800.

Mesh Network Segmentation

Membership to the wireless LAN mesh network for mesh access points is controlled by the bridge group names (BGNs). Mesh access points can be placed in similar bridge groups to manage membership or provide network segmentation.

Enterprise 11n/ac mesh added to the C9800 controller feature to work with the 802.11n/ac access points. Enterprise 11ac/ax mesh features are compatible with non-802.11ac mesh but adds higher backhaul and client access speeds. The 802.11ac indoor access points are two-radio Wi-Fi infrastructure devices for select indoor deployments. One radio can be used for local (client) access for the access point and the other radio can be configured for wireless backhaul. If Universal Backhaul Access is enabled, the 5-GHz and 2.4–GHz radios in rel 17.1 can be used for local (client) access as well as a backhaul. Enterprise 11ac mesh supports P2P, P2MP, and mesh types of architectures.

You have a choice of ordering indoor access points directly into the bridge mode, so that these access points can be used

directly as mesh access points. If you have these access points in a local mode (non-mesh), then you have to connect these access points to the controller and change the AP mode to the bridge mode (mesh). This scenario can become cumbersome particularly if the volume of the access points being deployed is large and if the access points are already deployed in the local mode for a traditional non-mesh wireless coverage.

Cisco Outdoor Mesh Access Points

The mesh access points, can operate, apart from the mesh mode, in the following modes:

- Local mode—In this mode, the AP can handle clients on its assigned channel or while monitoring all channels on the band over a 180-second period. During this time, the AP listens on each channel for 50 milliseconds for rogue client beacons, noise floor measurements, interference, and IDS events. The AP also scans for CleanAir interference on the channel.

- FlexConnect mode—FlexConnect is a wireless solution for branch office and remote office deployments. The FlexConnect mode enables you to configure and control access points in a branch or remote office from the corporate office through a WAN link without having to deploy a controller in each office. The FlexConnect mode can switch client data traffic locally and perform client authentication locally when the connection to the controller is lost. When connected to the controller, the FlexConnect mode can also tunnel traffic back to the controller.

- Flex+Mesh Mode—In this mode, both the FlexConnect and Bridge mode configuration options are available on the access point.

- Monitor mode—In this mode, the AP radios are in the receive state. The AP scans all the channels every 12 seconds for rogue client beacons, noise floor measurements, interference, IDS events, and CleanAir intruders.

- Rogue Detector mode—In this mode, the AP radio is turned off, and the AP listens only to the wired traffic. The controller passes the APs that are configured as rogue detectors as well as lists of suspected rogue clients and AP MAC addresses. The rogue detector listens for ARP packets and can be connected to all broadcast domains through a trunk link.

- Sniffer mode—In this mode, the AP captures and forwards all packets on a channel to a remote device that decodes the packets with packet analyzer software such as Wireshark.

- Bridge mode—In this mode, the AP is configured to build a wireless mesh network where wired network cabling is not available.

The Mesh Access Point can be changed to a desired mode via the following command:

```
1560-MAP1#capwap ap mode
bridge  Bridge mode
flex-bridge  Flex-Bridge mode
local       Local Mode
1560-MAP1#capwap ap mode local
```

Frequency Bands

Both the 2.4-GHz and 5-GHz frequency bands are supported on the indoor and outdoor access points.

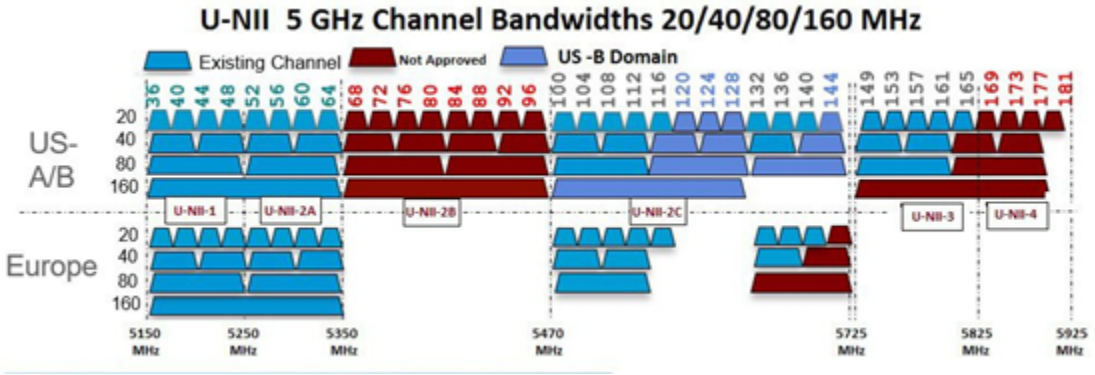


Figure 2 Frequency Bands Supported By 802.11a Radios on WAVE-2 AP15XXs

- FCC United States U-NII-1-This band can now be used indoors and outdoors Maximum power is increased to 30 dBm (1 Watt) assuming antenna is 6 dBi Power should be reduced by 1 dB for every dB antenna gain exceeds 6 dBi

When used outdoors, EIRP power in the upwards direction above 30 degrees is limited to 125 mW (20.9 dBm)

- U-NII-2A and U-NII2C-Must include Dynamic Frequency Selection (DFS) radar detection.

Terminal Doppler Weather Radar (TWDR) bands (channels 120, 124 & 128) are now available with new DFS test requirements

- U-NII-3-Band extended from 5825 MHz to 5850 MHz
- Europe U-NII-1-23 dBm Maximum–Not permitted for outdoor usage
- U-NII-2A-23 dBm Maximum–Not permitted for outdoor usage
- U-NII-2C-30 dBm Maximum
- U-NII-3-Only available in UK at 23 dBm for Indoor usage only

Dynamic Frequency Selection

Previously, devices employing radar operated in frequency sub bands without other competing services. However, controlling regulatory bodies are attempting to open and share these bands with new services like wireless mesh LANs (IEEE 802.11).

To protect existing radar services, the regulatory bodies require that devices wishing to share the newly opened frequency sub band behave in accordance with the Dynamic Frequency Selection (DFS) protocol. DFS dictates that to be compliant, a radio device must be capable of detecting the presence of radar signals. When a radio detects a radar signal, it is required to stop transmitting to for at least 30 minutes to protect that service. The radio then selects a different channel to transmit on but only after monitoring it. If no radar is detected on the projected channel for at least one minute, then the new radio service device may begin transmissions on that channel.

The AP performs a DFS scan on the new DFS channel for 60 seconds. However, if a neighboring AP is already using that new DFS channel, the AP does not perform the DFS scan.

The process for a radio to detect and identify a radar signal is a complicated task that sometimes leads to incorrect detects. Incorrect radar detections can occur due to a large number of factors, including due to uncertainties of the RF environment and the ability of the access point to reliably detect actual on-channel radar.

The 802.11h standard addresses DFS and Transmit Power Control (TPC) as it relates to the 5-GHz band. Use DFS to avoid

interference with radar and TPC to avoid interference with satellite feeder links.

DFS in RAP:

The RAP performs the following steps as a response to radar detection:

1. The RAP sends a message to the controller that the channel is infected with radar. The channel is marked as infected on the RAP and on the controller.
2. The RAP blocks the channel for 30 minutes. This 30-minute period is called the non-occupancy period.
3. The controller sends a TRAP, which indicates that the radar has been detected on the channel. A TRAP remains until the non-occupancy period expires.
4. The RAP has 10 seconds to move away from the channel. This period is called the channel move time, which is defined as the time for the system to clear the channel and is measured from the end of the radar burst to the end of the final transmission on the channel.
5. The RAP enters the quiet mode. In the quiet mode, the RAP stops data transmissions. Beacons are still generated and probe responses are still delivered. The quiet mode exists until the channel move time is over (10 seconds).
6. The controller picks up a new random channel and sends the channel information to the RAP.
7. The RAP receives the new channel information and sends channel change frames (unicast, encrypted) to the MAP, and each MAP sends the same information to its lower children down the sector. Each mesh access point sends the channel change frames once every 100 msecs for a total of five times.
8. The RAP tunes to the new channel and enters into the silent mode. During the silent mode, only the receiver is ON. The RAP keeps scanning the new channel for any radar presence for 60 seconds. This process is called channel availability check (CAC).
9. The MAP tunes to the new channel and enters into the silent mode. During the silent mode, only the receiver is ON. The MAP keeps scanning the new channel for any radar presence for 60 seconds.
10. If radar is not detected, the RAP resumes full functionality on this new channel and the whole sector tunes to this new channel.
11. If Radar interference is detected the Radios will shift to the new channel in the non UNII 2A-C band.

Note: If radar is detected on the RAP or MAP radio. This can trigger a channel change even when whole Mesh tree is connected.

Antennas

Antenna choice is a vital component of any wireless network deployment. There are two broad types of antennas:

- Directional
- Omnidirectional

Each type of antenna has a specific use and is most beneficial in specific types of deployments. Because antennas distribute RF signal in large lobed coverage areas determined by antenna design, successful coverage is heavily reliant on antenna

choice.

An antenna gives a mesh access point three fundamental properties: gain, directivity, and polarization:

- Gain—A measure of the increase in power. Gain is the amount of increase in energy that an antenna adds to an RF signal.
- Directivity—The shape of the transmission pattern. If the gain of the antenna increases, the coverage area decreases. The coverage area or radiation pattern is measured in degrees. These angles are measured in degrees and are called beam-widths.

Note: Beamwidth is defined as a measure of the ability of an antenna to focus radio signal energy toward a particular direction in space. Beamwidth is usually expressed in Note degrees HB (Horizontal Beamwidth); usually, the most important one is expressed in a VB (Vertical Beamwidth) (up and down) radiation pattern. When viewing an antenna plot or pattern, the angle is usually measured at half-power (3 dB) points of the main lobe when referenced to the peak effective radiated power of the main lobe.

Note: An 8-dBi antenna transmits with a horizontal beamwidth of 360 degrees, causing the radio waves to disperse power in all directions. Therefore, radio waves from an 8-dBi antenna do not go nearly as far as those radio waves sent from a 14-dBi patch antenna (or a third-party dish) that has a more narrow beamwidth (less than 360 degrees).

Polarization—The orientation of the electric field of the electromagnetic wave through space. Antennas can be polarized either horizontally or vertically, though other kinds of polarization are available. Both antennas in a link must have the same polarization to avoid an additional unwanted loss of signal. To improve the performance, an antenna can sometimes be rotated to alter polarization, which reduces interference. A vertical polarization is preferable for sending RF waves down concrete canyons, and horizontal polarization is generally more preferable for wide area distribution. Polarization can also be harnessed to optimize for RF bleed-over when reducing RF energy to adjacent structures is important. Most omnidirectional antennas ship with vertical polarization as their default.

Antenna Options

A wide variety of antennas are available to provide flexibility when you deploy the mesh access points over various terrains. Refer to the applicable access point data sheet or ordering guide for a list of supported antennas.

See the Cisco Aironet Antenna and Accessories Reference Guide on Cisco antennas and accessories at https://www.cisco.com/c/en/us/products/collateral/wireless/aironet-antennas-accessories/product_data_sheet09186a008008883b.html

The deployment and design, limitations and capabilities, and basic theories of antennas as well as installation scenarios, regulatory information, and technical specifications are addressed in detail.

https://www.cisco.com/c/en/us/products/collateral/wireless/aironet-antennas-accessories/product_data_sheet09186a008008883b.html

Flexible Antenna Port Configuration

The above HW changes have requirements for SW changes as well. The AP needs to support a flexible antenna port configuration. SW changes are done to let the user configure the antennas to support either in a single band mode or dual band mode. Software configurable Single Band Vs Dual Band mode.

Client Access Certified Antennas (Third-Party Antennas)

You can use third-party antennas with Wave-2 AP15XXs. However, note the following:

- Cisco does not track or maintain information about the quality, performance, or reliability of the noncertified antennas and cables.
- RF connectivity and compliance is the customer's responsibility.
- Compliance is only guaranteed with Cisco antennas or antennas that are of the same design and gain as Cisco

antennas.

- Cisco Technical Assistance Center (TAC) has no training or customer history with regard to non-Cisco antennas and cables.

Cisco Wireless Controllers

The wireless mesh solution with IOS-XE 17.1 is supported on Cisco C9800-CL Virtual and HW Appliances such C9800-80, C9800-40 and C9800-L.

Cisco Prime Infrastructure and Cisco DNA Center

The Cisco DNA Center provides a graphical platform for wireless mesh planning, configuration, and management. Network managers can use the Cisco DNA Center to design, control, and monitor wireless mesh networks from a central location.

With the Cisco DNA Center, network administrators have a solution for RF prediction, policy provisioning, network optimization, troubleshooting, user tracking, security monitoring, and wireless LAN systems management. Graphical interfaces make wireless LAN deployment and operations simple and cost-effective. Detailed trending and analysis reports make the Cisco DNA Center vital to ongoing network operations.

The Cisco DNA Center runs on a server platform with an embedded database, which provides scalability that allows hundreds of controllers and thousands of Cisco mesh access points to be managed. Controllers can be located on the same LAN as the Cisco DNA Center, on separate routed subnets, or across a wide-area connection.

Architecture

Control and Provisioning of Wireless Access Points

Control and provisioning of wireless access points (CAPWAP) is the provisioning and control protocol used by the controller to manage access points (mesh and nonmesh) in the network.

CAPWAP Discovery on a Mesh Network

The process for CAPWAP discovery on a mesh network is as follows:

1. A mesh access point establishes a link before starting CAPWAP discovery, whereas a non-mesh access point starts CAPWAP discovery using a static IP for the mesh access point, if any.
2. The mesh access point initiates CAPWAP discovery using a static IP for the mesh access point on the Layer 3 network or searches the network for its assigned primary, secondary, or tertiary controller. A maximum of 10 attempts are made to connect.
3. **Note**: The mesh access point searches a list of controllers configured on the access point (primed) during setup.
4. If Step 2 fails after 10 attempts, the mesh access point falls back to DHCP and attempts to connect in 10 tries.
5. If both Steps 2 and 3 fail and there is no successful CAPWAP connection to a controller.
6. If there is no discovery after attempting Steps 2, 3, and 4, the mesh access point tries the next link.

Dynamic MTU Detection

If the MTU is changed in the network, the access point detects the new MTU value and forwards that to the controller to adjust to the new MTU. After both the access point and the controller are set at the new MTU, all data within their path are fragmented into the new MTU. The new MTU size is used until it is changed. The default MTU on switches and routers is 1500 bytes.

Adaptive Wireless Path Protocol

The Adaptive Wireless Path Protocol (AWPP) is designed specifically for wireless mesh networking to provide ease of deployment, fast convergence, and minimal resource consumption.

AWPP takes advantage of the CAPWAP WLAN, where client traffic is tunneled to the controller and is therefore hidden from the AWPP process. Also, the advance radio management features in the CAPWAP WLAN solution are available to the wireless mesh network and do not have to be built into AWPP.

AWPP enables a remote access point to dynamically find the best path back to a RAP for each MAP that is part of the RAP's bridge group (BGN). Unlike traditional routing protocols, AWPP takes RF details into account.

To optimize the route, a MAP actively solicits neighbor MAP. During the solicitation, the MAP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor. The path decisions of AWPP are based on the link quality and the number of hops.

AWPP automatically determines the best path back to the CAPWAP controller by calculating the cost of each path in terms of the signal strength and number of hops. After the path is established, AWPP continuously monitors conditions and changes routes to reflect changes in conditions. AWPP also performs a smoothing function to signal condition information to ensure that the ephemeral nature of RF environments does not impact network stability.

Traffic Flow

The traffic flow within the wireless mesh can be divided into three components:

1. Overlay CAPWAP traffic that flows within a standard CAPWAP access point deployment; that is, CAPWAP traffic between the CAPWAP access point and the CAPWAP controller.

2. Wireless mesh data frame flow.

3. AWPP exchanges.

As the CAPWAP model is well known and the AWPP is a proprietary protocol, only the wireless mesh data flow is described. The key to the wireless mesh data flow is the address fields of the 802.11 frames being sent between mesh access points.

An 802.11 data frame can use up to four address fields: receiver, transmitter, destination, and source. The standard frame from a WLAN client to an AP uses only three of these address fields because the transmitter address and the source address are the same. However, in a WLAN bridging network, all four address fields are used because the source of the frame might not be the transmitter of the frame, because the frame might have been generated by a device behind the transmitter.

The figure below Wireless Mesh Frame shows an example of this type of framing. The source address of the frame is MAP:03:70, the destination address of this frame is the controller (the mesh network is operating in Layer 2 mode), the transmitter address is MAP:D5:60, and the receiver address is RAP:03:40.

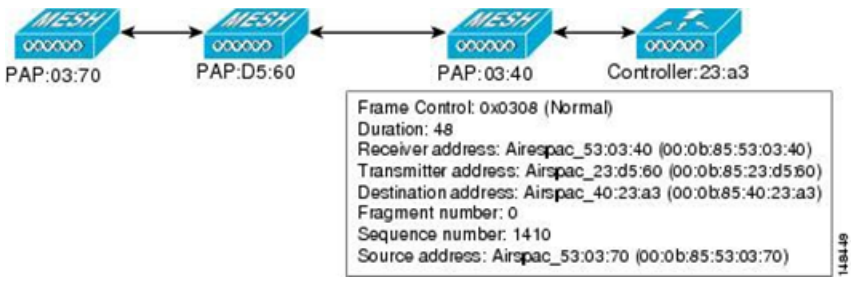


Figure 3 Wireless Mesh Frame

As this frame is sent, the transmitter and receiver addresses change on a hop-by-hop basis. AWPP is used to determine the receiver address at each hop. The transmitter address is known because it is the current mesh access point. The source and destination addresses are the same over the entire path.

If the RAP's controller connection is Layer 3, the destination address for the frame is the default gateway MAC address, because the MAP has already encapsulated the CAPWAP in the IP packet to send it to the controller, and is using the standard IP behavior of using ARP to find the MAC address of the default gateway.

Each mesh access point within the mesh forms an CAPWAP session with a controller. WLAN traffic is encapsulated inside CAPWAP and is mapped to a VLAN interface on the controller. Bridged Ethernet traffic can be passed from each Ethernet interface on the mesh network and does not have to be mapped to an interface on the controller (see Figure 5: Logical Bridge and WLAN Mapping).

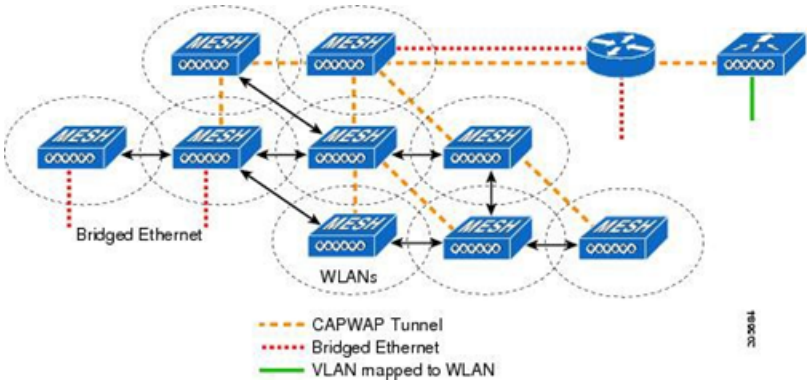


Figure 4 Logical Bridge and WLAN Mapping

Mesh Neighbors, Parents, and Children

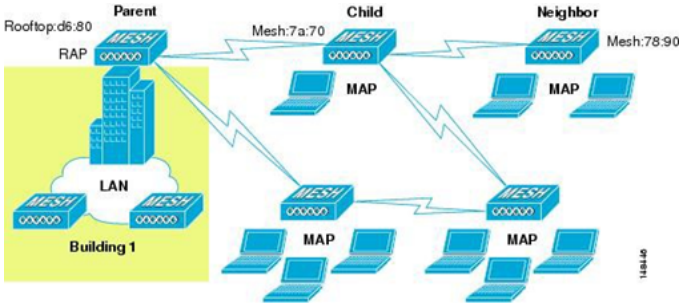


Figure 5 Parent, Child, and Neighbor Access Points

Relationships among mesh access points are as a parent, child, or neighbor (see Figure 6: Parent, Child, and Neighbor Access Points).

- A parent access point offers the best route back to the RAP based on its ease values. A parent can be either the RAP itself or another MAP.
    - Ease is calculated using the SNR and link hop value of each neighbor. Given multiple choices, generally an access point with a higher ease value is selected.
- A child access point selects the parent access point as its best route back to the RAP.
- A neighbor access point is within RF range of another access point but is not selected as its parent or a child because its ease values are lower than that of the parent.
- Mesh networks are half duplex meaning after the 1st hop (RAP to MAP) each additional hop (Map to Map) overall throughput is decreased by 50% per hop. Where Ethernet-bridges client are used in MAPs and heavy traffic is passed, it may result in a high throughput consumption, which may cause the downlink MAPs to disassociate from the network due to throughput starvation.

Criteria to Choose the Best Parent

AWPP follows this process in selecting parents for a RAP or MAP with a radio backhaul:

- A list of channels with neighbors is generated by passive scanning in the scan state, which is a subset of all backhaul channels.
- The channels with neighbors are sought by actively scanning in the seek state and the backhaul channel is changed to the channel with the best neighbor.
- The parent is set to the best neighbor and the parent-child handshake is completed in the seek state.
- Parent maintenance and optimization occurs in the maintain state.

This algorithm is run at startup and whenever a parent is lost and no other potential parent exists, and is usually followed by CAPWAP network and controller discovery. All neighbor protocol frames carry the channel information.

Parent maintenance occurs by the child node sending a directed NEIGHBOR_REQUEST to the parent and the parent responding with a NEIGHBOR_RESPONSE.

Parent optimization and refresh occurs by the child node sending a NEIGHBOR_REQUEST broadcast on the same channel on which its parent resides, and by evaluating all responses from neighboring nodes on the channel.

A parent mesh access point provides the best path back to a RAP. AWPP uses ease to determine the best path. Ease can be considered the opposite of cost, and the preferred path is the path with the higher ease.

Ease Calculation

Ease is calculated using the SNR and hop value of each neighbor, and applying a multiplier based on various SNR thresholds. The purpose of this multiplier is to apply a spreading function to the SNRs that reflects various link qualities.

The figure below shows Parent Path Selection shows the parent path selection where MAP2 prefers the path through MAP1 because the adjusted ease value (436906) though this path is greater then the ease value (262144) of the direct path from MAP2 to RAP.

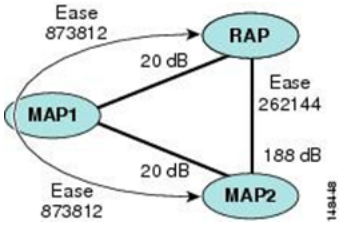


Figure 6 Parent Path Selection

Parent Decision

A parent mesh access point is chosen by using the adjusted ease, which is the ease of each neighbor

divided by the number of hops to the RAP: adjusted ease = min (ease at each hop) Hop count.

SNR Smoothing

One of the challenges in WLAN routing is the ephemeral nature of RF, which must be considered when analyzing an optimal path and deciding when a change in path is required. The SNR on a given RF link can change substantially from moment to moment, and changing route paths based on these fluctuations results in an unstable network, with severely degraded performance. To effectively capture the underlying SNR but remove moment-to-moment fluctuations, a smoothing function is applied that provides an adjusted SNR.

In evaluating potential neighbors against the current parent, the parent is given 20 percent of bonus-ease on top of the parent's calculated ease, to reduce the ping-pong effect between parents. A potential parent must be significantly better for a child to make a switch. Parent switching is transparent to CAPWAP and other higher-layer functions.

Loop Prevention

To ensure that routing loops are not created, AWPP discards any route that contains its own MAC address. That is, routing information apart from hop information contains the MAC address of each hop to the RAP; therefore, a mesh access point can easily detect and discard routes that loop.

Mesh AP Roaming

Mesh Access Points can roam from one parent mesh AP to a new parent mesh AP. Parent Mesh AP and C9800 will use the MESH_ROAM_REQUEST and MESH_ROAM_RESPONSE payloads to handle MAP roaming.

C9800 will support MAP roaming between parent Mesh APs within the same Controller and parent Mesh APs across different Controllers. MAP roaming across parent Mesh APs connected to Aire-OS and C9800 Controllers in the same mobility group will be supported.

Mesh Deployment Modes

Wireless Mesh Network

In a Cisco wireless outdoor mesh network, multiple mesh access points comprise a network that provides secure, scalable outdoor wireless LAN.

The three RAPs are connected to the wired network at each location and are located on the building roof. All the downstream access points operate as MAPs and communicate using wireless links (not shown).

Both MAPs and RAPs can provide WLAN client access; however, the location of RAPs are often not suitable for providing client access. All the three access points in are located on the building roofs and are functioning as RAPs. These RAPs are connected to the network at each location.

Some of the buildings have onsite controllers to terminate CAPWAP sessions from the mesh access points but it is not a mandatory requirement because CAPWAP sessions can be back hauled to a controller over a wide-area network (WAN).

Wireless Backhaul at 5 and 2.4 GHz

In a Cisco wireless backhaul network, traffic can be bridged between MAPs and RAPs. This traffic can be from wired devices that are being bridged by the wireless mesh or CAPWAP traffic from the mesh access points. This traffic is always AES encrypted when it crosses a wireless mesh link such as a wireless backhaul.

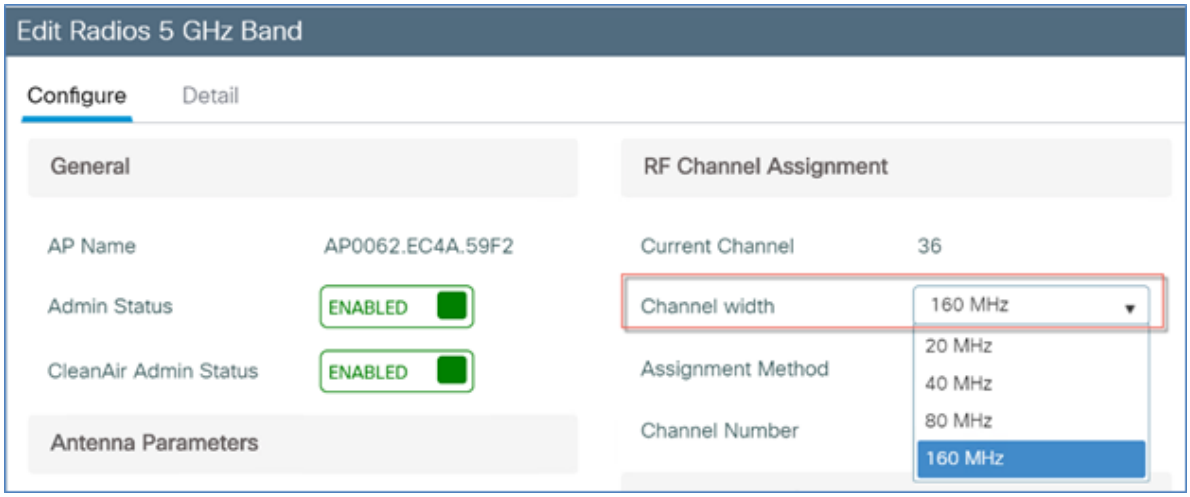
AES encryption is established as part of the mesh access point neighbor relationship with other mesh access points. The encryption keys used between mesh access points are derived during the EAP authentication process.

By default, the backhaul interface for Mesh APs is 802.11a/ac/ax. In certain countries it is not allowed to use Mesh Network with 5 GHz backhaul network or even in the courtiers when 5GHz is permitted customer may prefer to use 2.4 GHz radio frequencies to achieve much larger Mesh or Bridge distances.

When a RAP gets change of the configuration from 5 to 2.4 GHz that selection gets propagated from RAP to all MAPs and they will disconnect from 5GHz network and get reconnected at 2.4 GHz. During this process parent Mesh APs does not send any messages to child MAPs about the change in backhaul slot. MAPs should detect the parent loss and connect to parent APs after the scan in the new backhaul radio band.

Only RAPs are configured with the backhaul frequency of 5 or 2.4GHz.

Note: 160-MHz Channel Width is not supported on Mesh Bridge APs although it is not a restricted configuration on the Controller



Universal Access

You can configure the backhaul on mesh access points to accept client traffic over its 802.11radio. This feature is identified as Backhaul Client Access in the controller When this feature is disabled, backhaul traffic is transmitted only over the 802.11a/ac radio and client association is allowed only over the second radio. Backhaul Client Access is disabled by default. After this feature is enabled, all mesh access points, except slave AP and its child APs in Daisy-chained deployment on 1500 series APs, reboot.

Point-to-Multipoint Wireless Bridging

In the point-to-multipoint bridging scenario, a RAP acting as a root bridge connects multiple MAPs as non-root bridges with their associated wired LANs. By default, this feature is disabled for all MAPs. If Ethernet bridging is used, you must enable it on the controller for the respective MAP and for the RAP.

This figure shows a simple deployment with one RAP and two MAPs, but this configuration is fundamentally a wireless mesh with no WLAN clients. Client access can still be provided with Ethernet bridging enabled, although if bridging between buildings, MAP coverage from a high rooftop might not be suitable for client access.

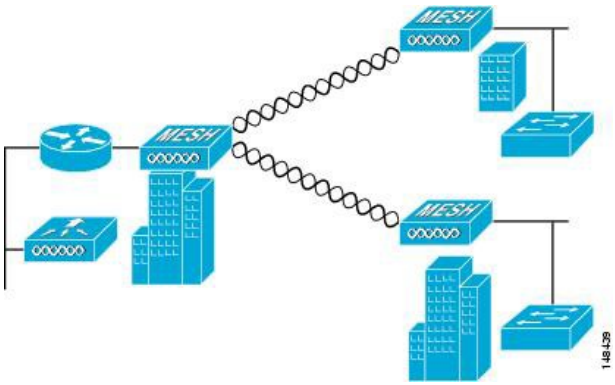


Figure 7 Point-to-Multipoint Bridging Example

Point-to-Point Wireless Bridging

In a point-to-point bridging scenario, a 1500 Series Mesh AP can be used to extend a remote network by using the backhaul radio to bridge two segments of a switched network. This is fundamentally a wireless mesh network with one MAP and no WLAN clients. Just as in point-to-multipoint networks, client access can still be provided with Ethernet bridging enabled, although if bridging between buildings, MAP coverage from a high rooftop might not be suitable for client access.

If you intend to use an Ethernet bridged application, we recommend that you enable the bridging feature on the RAP and on all MAPs in that segment. You must verify that any attached switches to the Ethernet ports of your MAPs are not using VLAN Trunking Protocol (VTP). VTP can reconfigure the trunked VLANs across your mesh and possibly cause a loss in

connection for your RAP to its primary C9800. An incorrect configuration can take down your mesh deployment.

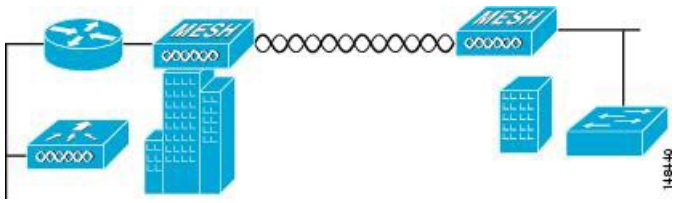


Figure 8 Point-to-Point Bridging Example

For security reasons the Ethernet port on the MAPs is disabled by default. It can be enabled only by configuring Ethernet bridging on the Root and the respective MAPs.

Note: Ethernet bridging has to be enabled for the following two scenarios:

- When you want to use the mesh nodes as bridges.
- When you want to connect Ethernet devices such as a video camera on the MAP using its Ethernet port.

Ensure that you enable Ethernet bridging for every parent mesh AP taking the path from the mesh AP in question to the controller. For example, if you enable Ethernet bridging on MAP2 in Hop 2, then you must also enable Ethernet bridging on MAP1 (parent MAP), and on the RAP connecting to the controller. Optimum distance (in feet) should exist between the root access point (RAP) and the farthest mesh access point (MAP). Range from the RAP bridge to the MAP bridge has to be mentioned in feet.

The following global parameter applies to all mesh access points when they join the controller and all existing mesh access points in the network:

Range: 150 to 132,000 feet

Configuring Mesh Range (CLI)

- To configure the distance between the nodes doing the bridging, enter the **config mesh range** command. APs reboot after you specify the range.
- To view the mesh range, enter the show mesh config command.

Mesh Daisy Chaining

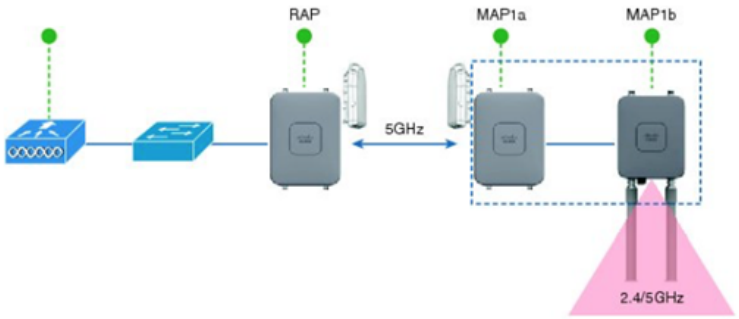
The Cisco Aironet 1540, 1560 and 1572 Series Access Points have the capability to "daisy chain" access points when they function as mesh APs (MAPs). The "daisy chained" MAPs can extend universal access by connecting a local mode or FlexConnect mode Cisco AP1570 to the Ethernet port of a MAP, thus extending the network to provide better client access.

In this case, the daisy chained Mesh AP is called master MAP and the MAP which is connected to master MAP over Ethernet is called the slave MAP or slave RAP since another wireless MAP can connect to the slave RAP if properly configured to prevent loops.

In case of daisy-chaining mode,

- Master MAP should be configured as mesh AP
- Slave MAP should be configured as root AP
- Daisy chaining should be enabled on both master and slave MAP
- Ethernet bridging should be enabled on all the APs in the Bridge mode. Enable Ethernet bridging in Mesh profile and all Bridge mode APs in the sector should be mapped to the same mesh profile.

- VLAN support should be enabled on wired root AP, slave MAP and master MAP along with proper native VLAN configuration

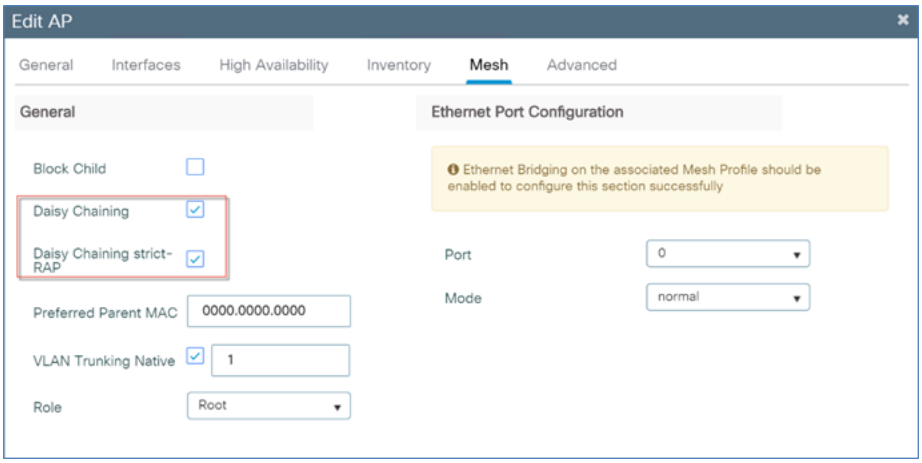


Mesh Daisy-Cain Configuration

This configuration should be applied to both Master MAP and slave RAP.

```
(Cisco Controller) >config ap daisy-chaining <enable/disable> <AP Name>
(Cisco Controller) >config ap strict-wired-uplink <enable/disable>?
enable          Enables Strict Wired Uplink on the Cisco AP.
disable         Disables Strict Wired Uplink on the Cisco AP.

RAP#capwap ap mesh strict-wired-uplink <enable/disable>
  disable  disable strict wired uplink
  enable   enable strict wired uplink
```



Flex+Mesh AP Running Modes

Flex+Mesh Wave-2 APs can be running in connected or standalone mode. Standalone mode in flex connect will undergo some changes to inherit standalone functionality for a mesh network. There is also another mode called 'abandoned' mode discussed below in this section of the guide.

Connected Mode

A Wave-2 Flex+Mesh AP (Root AP or Child Mesh AP) is considered to be in connected mode when it can access and join the C9800 and can exchange periodic keep alive messages with C9800. In this mode, Flex+Mesh AP will be able support locally and centrally switched WLAN's. It shall allow regular client and Child mesh APs to join.

Standalone Mode

A Wave-2Flex+Mesh AP, is considered to be in standalone mode if it loses connection to the controller but it can access the local gateway. In this mode, the Wave-2 Flex+Mesh AP will disable all the centrally switched WLANs, and shall keep the locally switched WLANs up and running. It will also allow the new clients to join on local switched WLANs using local authentication as long as the authentication server is reachable in the local network. Child mesh APs will NOT be allowed

to join in this mode.

Abandoned Mode or Persistent SSID Mode

A Wave-2 Flex+Mesh AP is in abandoned mode when it can no longer access the gateway IP and has no connectivity to the local network. Possible scenarios are:

- AP is still not locked on to any uplink wired or wireless.
- A wireless uplink has been established but has not been authenticated.
- An uplink is established and authenticated, but IP address has the gateway IP has not been configured.
- An uplink is established, authenticated and also IP address and gateway IP has been configured, but the gateway is not reachable for over a minute.

Neither Child Mesh APs nor the clients are allowed to join in this mode. Local as well as centrally switched WLANs will be disabled. AP may still be scanning for an uplink in this mode so no beacons will be transmitted during this time.

Note: For Flex+Mesh Wave-2 APs, in abandoned mode, reboot timer shall be enabled so the AP will have rebooted after 40 minutes, if it does not transition to either standalone mode or connected mode.

Mode/State transitions in the Flex+Mesh Wave-2APs

- Flex+Mesh mode Wave-2AP will always boot up in abandoned mode, in which it would need to scan for the uplink (wired or radio).
- Once a new uplink is selected either during initial stage or during inter gateway roaming scenario, it is expected that the authentication should pass and the CAPWAP connection needs to be formed within 2 minutes, else the selected parent will be blacklisted. This function should be same as a regular Mesh mode Wave-2AP.
- If a Flex+Mesh AP has a valid CAPWAP connection and it loses the CAPWAP connection it will transition to standalone mode, and will stay in standalone mode, as long as the gateway is reachable. A Flex+Mesh AP will keep track of the IP mode (IPV6 or IPV4) used for the last successful CAPWAP connection and with track the reachability of the GW for that IP mode.
- For Flex+Mesh AP in standalone mode, Mesh control will start a timer (20 second) to periodically refresh the ARP entry for GW IP (IPV4 or IPV6) and to also query the GW reachability status from the Path Control Protocol. PCP will maintain the gateway reachability status from that AP either reported by the Root AP via PCP messages or if it is Root AP by doing an ARP lookup for the gateway IP address. If the GW is unreachable for over a minute, the Flex+Mesh AP will blacklist the parent and will transition to abandoned mode and will re-scan for a new uplink.
- To come out of the abandoned mode, AP must connect to the C9800 and transition to the connected mode. Transition from abandoned mode directly to standalone mode is not supported and needs to be considered in future design enhancements.

Design considerations for Flex AP in standalone mode:

- When the Flex AP is in standalone mode, it will stick to the same parent and will NOT try to discover or roam to a better neighbor, even if it is a preferred parent. The reason is that there is no guarantee that the security will pass with the new parent and the roaming will be successful. If the security fails, the perspective parent may get blacklisted unnecessarily. It is best to consider standalone roaming once standalone security is supported for

Mesh APs in future design enhancements.

- BGN timer will be stopped in standalone mode. So, if the child mesh AP is in standalone mode and it joins a parent with a different BGN and goes back into standalone mode after that, BGN timer will be stopped so that the child Mesh AP does not go into re-scan mode after 15 minutes (BGN timer expiry).
- In standalone mode, reboot timer will be stopped so that the AP does not reboot after 40 minutes, in the absence of a CAPWAP connection.
- After moving back to connected mode, from standalone mode, best neighbor selection timer and BGN timer will be restarted, so allow the child mesh AP to roam to the best possible neighbor.

Special standalone mode for Wave-2 Flex RAPs

In this mode the SSID will be broadcasted always (Persistent SSID). In addition, after reboot, when this special Persistent mode is enabled, Flex+Mesh RAP should be able to start broadcasting the SSID even if the gateway is not reachable.

Existing FlexConnect AP mode design

- Locally switched WLANs are stored in config.flex file and Flex-connect AP broadcasts the local WLAN SSIDs as long as it is standalone mode.
- On boot up Flex-connect AP would only start broadcasting the locally switched WLANs if the gateway provisioned.
- If for a Wave-2 Flex connect AP, gateway information is removed at some point, it moves out of the standalone mode and stops broadcasting the locally switched SSIDs and waits for gateway to be provisioned again.
- Once the gateway is provisioned, Flex AP again transitions into the standalone mode and starts broadcasting the locally switched SSIDs again.
- Without a valid gateway, flex-connect AP eventually stops broadcasting SSIDs, since the local network is not reachable so no reason to connect the clients.
- Parts of the existing Flex-connect AP mode design is used to retain WLAN configuration during reboot and to be able to start broadcasting Local SSIDs etc. However, for Flex RAP we have a special standalone mode requirement for NBN deployment as stated below:
- Flex RAP should be able to boot up directly into the standalone mode and start broadcasting SSIDs, even if the gateway is not reachable.
- Flex RAP will continue to be in standalone mode and keep broadcasting SSIDs if the gateway was reachable earlier and becomes unreachable at some point.
- Even if the Flex RAP cannot support any real clients, it still needs to broadcast SSID so that the operator can check if the AP is UP and running.

Ethernet Bridging

For security reasons the Ethernet port on the MAPs is disabled by default. It can be enabled only by configuring Ethernet bridging on the Root and the respective MAPs. This means that traffic from a wired client on a mesh AP gets bridged to the

other clients in the mesh or to the DS and beyond. Typical use of Ethernet ports is to connect cameras for monitoring the APs.

Both tagged and untagged packets are supported on secondary Ethernet interfaces.

In a point-to-point bridging scenario, a 1500 Series Mesh AP can be used to extend a remote network by using the backhaul radio to bridge two segments of a switched network. This is fundamentally a wireless mesh network with one MAP and no WLAN clients. Just as in point-to-multipoint networks, client access can still be provided with Ethernet bridging enabled, although if bridging between buildings, MAP coverage from a high rooftop might not be suitable for client access. To use an Ethernet bridged application, enable the bridging feature on the RAP and on all MAPs in that sector.

Ethernet bridging has to be enabled for the following two scenarios

- When you want to use the mesh nodes as bridges
- When you want to connect Ethernet devices such as a video camera on the MAP using its Ethernet port

Ensure that Ethernet bridging is enabled for every parent mesh AP taking the path from the mesh AP in question to the controller.

Ethernet bridging works without controller knowledge, this means that there's no CAPWAP involved for Ethernet bridging. We only use CAPWAP for configuration purpose.

In a mesh environment with VLAN support for Ethernet bridging, the secondary Ethernet interfaces on MAPs can be assigned a VLAN individually, via "ap exec" commands. All backhaul bridge links, both wired and wireless are trunk links with all VLANs enabled. Non-Ethernet bridged traffic, as well as untagged Ethernet bridged traffic travels along the mesh using the native VLAN of the APs in the mesh. This holds true for all the traffic to/from wireless clients which the APs are servicing.

The VLAN tagged packet will be tunneled through AWPP over Wireless Backhaul links.

VLAN Tagging for MAP Ethernet Clients

The backhaul interfaces of Mesh APs are referred to as primary interfaces and other interfaces are referred to as secondary interfaces. So, the Ethernet interface which is used as backhaul is known as "Primary Ethernet interface" and others are known as "Secondary Ethernet interfaces".

The following are supported:

1. Allow trunk configuration on "Secondary Ethernet interfaces" of mesh APs such that if a VLAN tagged packet comes to one of these interfaces, it gets forwarded based on its tag. "Primary Ethernet Interfaces", which are backhauls, will behave as trunk.

2. Allow VLAN configuration on the "Secondary Ethernet interfaces" of mesh APs such that any untagged packet coming on these interfaces can be tagged and then forwarded.

3. It should be noted that only the "Secondary Ethernet Interfaces" can be configured. The "Primary Ethernet Interfaces" will not be configurable by user as they are required to behave as trunk and are expected to carry data of all the VLANs.

Workgroup Bridge Interoperability with Mesh Infrastructure

A workgroup bridge (WGB) is a small standalone unit that can provide a wireless infrastructure connection for Ethernet-enabled devices. Devices that do not have a wireless client adapter to connect to the wireless network can be connected to the WGB through the Ethernet port. The WGB is associated with the root AP through the wireless interface, which means that wired clients get access to the wireless network.

A WGB is used to connect wired networks over a single wireless segment by informing the mesh access point of all the

clients that the WGB has on its wired segment via IAPP messages. The data packets for WGB clients contain an additional MAC address in the 802.11 header (4 MAC headers, versus the normal 3 MAC data headers). The additional MAC in the header is the address of the WGB itself. This additional MAC address is used to route the packet to and from the clients.

WGB association is supported on all radios of every mesh access point.

In the current architecture, while an autonomous AP functions as a workgroup bridge, only one radio interface is used for controller connectivity, Ethernet interface for wired client connectivity, and other radio interface for wireless client connectivity. dot11radio 1 (5 GHz) can be used to connect to a controller (using the mesh infrastructure) and Ethernet interface for wired clients. dot11radio 0 (2.4 GHz) can be used for wireless client connectivity. Depending on the requirement, dot11radio 1 or dot11radio 0 can be used for client association or controller connectivity.

With two radios, one radio can be used for client access and the other radio can be used for accessing the access points. Having two independent radios performing two independent functions provides you better control and lowers the latency. Also, wireless clients on the second radio for the WGB do not get disassociated by the WGB when an uplink is lost or in a roaming scenario. One radio has to be configured as a Root AP (radio role) and the second radio has to be configured as a WGB (radio role).

Note: If one radio is configured as a WGB, then the second radio cannot be a WGB or a repeater in the IOS-XE release 17.1.

The following features are not supported for use with a WGB:

- Idle timeout
- Web authentication—If a WGB associates to a web-authentication WLAN, the WGB is added to the exclusion list, and all of the WGB-wired clients are deleted (web-authentication WLAN is another name for a guest WLAN).
- For wired clients behind the WGB, MAC filtering, link tests, and idle timeout

Design Considerations

This chapter describes important design considerations and provides an example of a wireless mesh design.

Each outdoor wireless mesh deployment is unique, and each environment has its own challenges with available locations, obstructions, and available network infrastructure. Design requirements driven by expected users, traffic, and availability needs are also major design criteria.

Wireless Mesh Constraints

The following are a few system characteristics to consider when you design and build a wireless mesh network. Some of these characteristics apply to the backhaul network design and others to the CAPWAP controller design:

Wireless Backhaul Data Rate

Backhaul is used to create only the wireless connection between the access points. The backhaul interface is 802.11a/n/ac depending upon the access point. The rate selection is important for effective use of the available RF spectrum. The rate can also affect the throughput of client devices, and throughput is an important metric used by industry publications to evaluate vendor devices.

Dynamic Rate Adaptation (DRA) introduces a process to estimate optimal transmission rate for packet transmissions. It is important to select rates correctly. If the rate is too high, packet transmissions fail resulting in communication failure. If the rate is too low, the available channel bandwidth is not used, resulting in inferior products, and the potential for catastrophic network congestion and collapse.

Data rates also affect the RF coverage and network performance. Lower data rates, for example 6 Mbps, can extend farther from the access point than can higher data rates, for example 1300 Mbps. As a result, the data rate affects cell coverage and consequently the number of access points required. Different data rates are achieved by sending a more redundant signal on the wireless link, allowing data to be easily recovered from noise. The number of symbols sent out for a packet at the 1-Mbps data rate is higher than the number of symbols used for the same packet at 11 Mbps. Therefore,

sending data at the lower bit rates takes more time than sending the equivalent data at a higher bit rate, resulting in reduced throughput.

A lower bit rate might allow a greater distance between MAPs, but there are likely to be gaps in the WLAN client coverage, and the capacity of the backhaul network is reduced. An increased bit rate for the backhaul network either requires more MAPs or results in a reduced SNR between MAPs, limiting mesh reliability and interconnection.

Note: The data rate can be set on the backhaul on a per AP basis. It is not a global command.

- The required minimum Link SNR value is driven by the data rate and the following formula: Minimum SNR + fade margin.
- If we take into account the effect of MRC for calculating Minimum Required Link SNR.
- Link SNR = Minimum SNR - MRC + Fade Margin (9 dB)

With two spatial streams, the MRC gain is halved, that is the MRC gain is reduced by 3 dB. This is because the system has 10 log (3/2 SS) instead of 10 log (3/1 SS). If there were to have been 3 SS with 3 RX, then the MRC gain would have been zero.

- Number of backhaul hops is limited to eight but we recommend three to four hops.
- The number of hops is recommended to be limited to three or four primarily to maintain sufficient backhaul throughput, because each mesh access point uses the same radio for transmission and reception of backhaul traffic, which means that throughput is approximately halved over every hop. For example, the maximum throughput for 24 Mbps is approximately 14 Mbps for the first hop, 9 Mbps for the second hop, and 4 Mbps for the third hop.
- Number of MAPs per RAP.
- There is There is no current software limitation on how many MAPs per RAP you can configure. However, it is suggested that you limit the number to 20 MAPs per RAP.
- Number of controllers
    - The number of controllers per mobility group is limited to 72.
- Number of mesh access points supported per controller.

Please see the Mesh AP 1570, 1560 and 1540 series Data Sheets at the links below for supported Data Rates, Receive Sensitivity, Supported MCS rates and other details at the link below

https://www.cisco.com/c/en/us/products/collateral/wireless/aironet-1570-series/datasheet-c78-732348.html

https://www.cisco.com/c/en/us/products/collateral/wireless/aironet-1560-series/datasheet-c78-737416.html

https://www.cisco.com/c/en/us/products/collateral/wireless/aironet-1540-series/datasheet-c78-738585.html

Controller Planning

- The following items affect the number of controllers required in a mesh network:
- Mesh access points (RAPs and MAPs) in the network.

- The wired network that connects the RAP and controllers can affect the total number of access points supported in the network. If this network allows the controllers to be equally available to all access points without any impact on WLAN performance, the access points can be evenly distributed across all controllers for maximum efficiency. If this is not the case, and controllers are grouped into various clusters or PoPs, the overall number of access points and coverage are reduced.
- Number of mesh access points (RAPs and MAPs) supported per controller.
- For clarity, non-mesh access points are referred to as local access points in this document.

Mesh Access Point Support by Controller Model

Controller Model	Local AP Support (non-mesh)	Maximum Possible Mesh AP Support
C9800 -80 C9800 -40	6000 2000	6000 2000
C9800 -CL C9800 -L	1000 250	1000 250

VM Ware Specifications for IOS-XE 16.10-17.1

- Supported hypervisor: VMware ESXi 6.0 and higher

Model Configuration	Small (16.10)	Medium(16.10)	Large(16.10)*
Maximum Access Points	1,000	3,000	6,000
Maximum Clients Support	6,000	32,000	64,000
Minimum Number of vCPUs	4	6	10
Minimum Memory (GB)	8	16	32
Required Storage (GB)	8	8	8
Virtual NICs (vNIC) – 3rd NIC is for High Availability	2 /(3)	2 /(3)	2 /(3)
vNIC driver	VMXNET3, E1000E, E1000	VMXNET3, E1000E, E1000	VMXNET3, E1000E, E1000
Virtual bridge	Vswitch	Vswitch	Vswitch
vMotion, vNIC teaming, L2 LAG, SRIOV	Planned for 16.11	Planned for 16.11	Planned for 16.11

*Limited scale with Local Mode and Flex Central switching : 3K APs, 32K clients

KVM Specifications for IOS-XE rel 16.10-17.1

- Supported Linux distribution: RHEL 7.1 & 7.2, Ubuntu 14.04, 16.04 LTS

Model Configuration	Small(16.10)	Medium(16.10)	Large(16.10)*
Maximum Access Points	1,000	3,000	6,000
Maximum Clients Support	6,000	32,000	64,000
Minimum Number of vCPUs	4	6	10
Minimum Memory (GB)	8	16	32
Required Storage (GB)	8	8	8
Virtual NICs (vNIC) 3rd NIC is for High Availability	3	3	3
vNIC driver	VIRTIO	VIRTIO	VIRTIO
Virtual bridge	OVS Linux bridge (brctl)	OVS Linux bridge (brctl)	OVS Linux bridge (brctl)
vNIC teaming, L2 LAG, SRIOV	Planned for 16.11	Planned for 16.11	Planned for 16.11

*Limited scale with Local Mode and Flex Central switching : 3K APs, 32K clients

Site Preparation and Planning

Site Survey

We recommend that you perform a radio site survey before installing the equipment. A site survey reveals problems such

as interference, Fresnel zone, or logistics problems. A proper site survey involves temporarily setting up mesh links and taking measurements to determine whether your antenna calculations are accurate. Determine the correct location and antenna before drilling holes, routing cables, and mounting equipment.

Note: When power is not readily available, we recommend you to use an unrestricted power supply (UPS) to temporarily power the mesh link.

Pre-Survey Checklist

Before attempting a site survey, determine the following:

- How long is your wireless link?
- Do you have a clear line of sight?
- What is the minimum acceptable data rate within which the link runs?
- Is this a point-to-point or point-to-multipoint link?
- Do you have the correct antenna?
- Can the access point installation area support the weight of the access point?
- Do you have access to both of the mesh site locations?
- Do you have the proper permits, if required?
- Do you have a partner? Never attempt to survey or work alone on a roof or tower.
- Have you configured the 1500 series before you go onsite? It is always easier to resolve configuration or device problems first.
- Do you have the proper tools and equipment to complete your task?

Note: Cellular phones or handheld two-way radios can be helpful to do surveys.

Outdoor Site Survey

Deploying WLAN systems outdoors requires a different skill set to indoor wireless deployments. Considerations such as weather extremes, lightning, physical security, and local regulations need to be considered.

When determining the suitability of a successful mesh link, define how far the mesh link is expected to transmit and at what radio data rate. Remember that the data rate is not directly included in the wireless routing calculation, and we recommend that the same data rate is used throughout the same mesh.

Design recommendations for mesh links are as follows:

- MAP deployment cannot exceed 35 feet in height above the street.
- MAPs are deployed with antennas pointed down toward the ground.
- Typical 5-GHz RAP-to-MAP distances are 1000 to 4000 feet.
- RAP locations are typically towers or tall buildings.

- Typical 5-GHz MAP-to-MAP distances are 500 to 1000 feet.
- MAP locations are typically short building tops or streetlights.
- Typical 2.4-GHz MAP-to-client distances are 500 to 1000 feet (depends upon the type of access point).
- Clients are typically laptops, Smart Phones, Tablets, and CPEs. Most of the clients operate in the 2.4-GHz band.

Determining a Line of Sight

When you determine the suitability of a successful link, you must define how far the link is expected to transmit and at what radio data rate. Very close links, one kilometer or less, are fairly easy to achieve assuming there is a clear line of sight (LOS)–a path with no obstructions.

Because mesh radio waves have very high frequency in the 5-GHz band, the radio wavelength is small; therefore, the radio waves do not travel as far as radio waves on lower frequencies, given the same amount of power. This higher frequency range makes the mesh ideal for unlicensed use because the radio waves do not travel far unless a high-gain antenna is used to tightly focus the radio waves in a given direction.

This high-gain antenna configuration is recommended only for connecting a RAP to the MAP. To optimize mesh behavior, omnidirectional antennas are used because mesh links are limited to one mile (1.6 km). The curvature of the earth does not impact line-of-sight calculations because the curvature of the earth changes every six miles (9.6 km).

Weather

In addition to free space path loss and line of sight, weather can also degrade a mesh link. Rain, snow, fog, and any high humidity condition can slightly obstruct or affect the line of sight, introducing a small loss (sometimes referred to as rain fade or fade mesh link, the weather should not be a problem; however, if the link is poor to begin with, bad weather can degrade performance or cause loss of link.

Ideally, you need a line of sight; a white-out snow storm does not allow a line of sight. Also, while storms may make the rain or snow itself appear to be the problem, many times it might be additional conditions caused by the adverse weather. For example, perhaps the antenna is on a mast pipe and the storm is blowing the mast pipe or antenna structure and that movement is causing the link to come and go, or there might be a large build-up of ice or snow on the antenna.

Fresnel Zone

A Fresnel zone is an imaginary ellipse around the visual line of sight between the transmitter and receiver. As radio signals travel through free space to their intended target, they could encounter an obstruction in the Fresnel area, degrading the signal. Best performance and range are attained when there is no obstruction of this Fresnel area. Fresnel zone, free space loss, antenna gain, cable loss, data rate, link distance, transmitter power, receiver sensitivity, and other variables play a role in determining how far your mesh link goes. Links can still occur as long as 60 percent to 70 percent of the Fresnel area is unobstructed, as illustrated in Figure 11: Point-to-Point Link Fresnel Zone



Figure 9 Point-to-Point Link Fresnel Zone

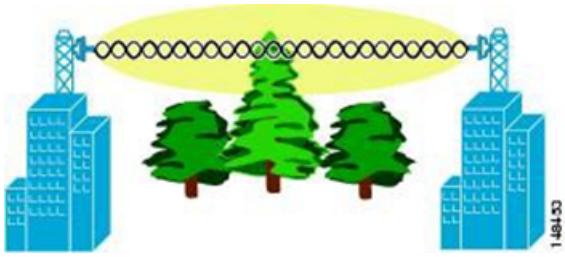


Figure 10 Typical Obstructions in a Fresnel Zone

It is possible to calculate the radius of the Fresnel zone (in feet) at any particular distance along the path using the following equation:

F1 = 72.6 X square root (d/4 x f) where

F1 = the first Fresnel zone radius in feet D = total path length in miles

F = frequency (GHz)

Normally, 60 percent of the first Fresnel zone clearance is recommended, so the above formula for 60 percent Fresnel zone clearance can be expressed as follows:

0.60 F1= 43.3 x square root (d/4 x f)

These calculations are based on a flat terrain.

The figure below shows the removal of an obstruction in the Fresnel zone of the wireless signal.



Figure 11 Removing Obstructions in a Fresnel Zone

Fresnel Zone Size in Wireless Mesh Deployments

To give an approximation of size of the maximum Fresnel zone to be considered, at a possible minimum frequency of 4.9 GHz, the minimum value changes depending on the regulatory domain. The minimum figure quoted is a possible band allocated for public safety in the USA, and a maximum distance of one mile gives a Fresnel zone of clearance requirement of 9.78 ft = 43.3 x SQR(1/(4*4.9)). This clearance is relatively easy to achieve in most situations. In most deployments, distances are expected to be less than one mile, and the frequency greater than 4.9 GHz, making the Fresnel zone smaller. Every mesh deployment should consider the Fresnel zone as part of its design, but in most cases, it is not expected that meeting the Fresnel clearance requirement is an issue.

Co-Channel Interference

In addition to hidden node interference, co-channel interference can also impact performance. Co-channel interference occurs when adjacent radios on the same channel interfere with the performance of the local mesh network. This interference takes the form of collisions or excessive deferrals by CSMA. In both cases, performance of the mesh network is degraded. With appropriate channel management, co-channel interference on the wireless mesh network can be minimized.

Wireless Mesh Network Coverage Considerations

This section provides a summary of items that must be considered for maximum wireless LAN coverage in an urban or suburban area, to adhere to compliance conditions for respective domains.

The following recommendations assume a flat terrain with no obstacles (green field deployment).

Background scanning on Wave-2 APs

The Mesh Background Scanning and Auto parent selection will improve convergence times and parent selection reliability and stability. A MAP should be able to find and connect with a better potential parent across any channels and maintain its uplink with a best parent all the time.

With background scanning disabled, whenever a Mesh AP detects a parent loss, it has to scan all the channels of the regulatory domain to find a new parent and then has to get authenticated through the selected parent to the C9800. This will take a longer time for the mesh AP to connect back to the C9800.

With background scanning enabled, after the detection of parent loss, a mesh AP can avoid scan to find a best parent across other channels. It can directly select a best parent from the neighbor list and establish the AWPP link.

A child MAP maintains its uplink with its parent using AWPP - Neighbor Discovery Request/Response (NDReq/NDResp) messages which are acting as keep-alives. If there are consecutive losses of NDResp messages, a parent is declared to be lost and the child MAP tries to find its new parent. A MAP maintains a list of neighbors of current on-channel, and on losing its current parent, it will try roaming to next best potential neighbor in the same serving channel.

But if there are no other neighbors found in same channel, it has to do scan/seek across all/subset channels to find a parent. Each off-channel list node will have a neighbor list managing all neighbors heard in that channel. Upon each off-channel NDReq broadcasts, the neighbors will be updated with latest SNR values based on their NDResp packets.

Mesh background scanning tries to avoid finding a parent across other channels by scan/seek which are time consuming, but keeps the child MAP updated with all the neighbors across all channels and will help just 'switching' to a neighbor of any channel and use him as its next parent for its uplink.

Mesh Convergence

Mesh convergence is one of the important features of Mesh Access Points in order to re-establish the connection to C9800, whenever a MAP loses its backhaul connection from its current parent which could be a RAP or a successive level MAP.

To improve the Mesh convergence time, each Mesh AP will maintain a subset of channels which will be used for future scan/seek and identify a parent among the subset neighbor list.

Along with this, there are other optimizations done on the Mesh Access Point to reduce the convergence time.

Standard convergence

In this method, to detect a parent loss MAP takes 21 seconds and for seek it takes 3 seconds per channel. Parent/Neighbor keep alive timer will be 3 seconds.

Fast convergence

In this method, parent loss detection is reduced to 7 seconds and for seek 2 seconds per channel. By using the subset of channels, MAP will scan only in a subset of channels which reduces the overall seek time. Parent/Neighbor keep alive time will be 3 seconds.

Very fast convergence

In this method, parent loss detection was reduced to 4 seconds and for seek 2 seconds per channel only on the subset of channels. Parent/Neighbor keep alive time will be 1.5 seconds.

Convergence method	Parent loss adjTimerMN – seconds	Seek per channel adjTimerI1 – seconds	Parent, neighbor keep alive adjTimerMP – seconds

Standard	21	3	3
Fast	7	2	3
Very Fast	4	2	1.5

In case if a MAP gets stranded and fails to find a parent and connect to C9800 , it will reboot after MESH_LWAPP_REBOOT_TIMER (40 minutes) expiry. After this the existing standard convergence will get applied.

Parent loss Detection

Parent loss detection is based on failure to get response for AWPP neighbor request, which evaluates current parent (AE_UPDATED flag) every 21 seconds in the standard method. Each neighbor is sent a unicast request every 3 seconds along with a request to parent. Failure to get response from parent, initiates a roam if there is any best neighbor available on same channel (OR) a full scan for a new parent. Considering a parent loss on 1st second of an update cycle, it takes another 20 seconds to detect parent loss to initiate roam / full channel scan in standard convergence method. For Fast convergence method, timer adjTimerMN was reduced to 7 seconds from 21 seconds. If there is a failure in getting response for unicast AWPP request, after a retry, change of parent logic will be triggered.

For Very Fast convergence method, timer adjTimerMN was reduced to 4 seconds. Also, Parent, neighbor keep alive timer (adjTimerMP) was reduced to 1.5 seconds to keep updated with parent and adjacent neighbors in same channel. This results in increase of number of request messages (1 request per 3 sec to 1 request per 1.5 sec) for parent and 5 best neighbors.

Cell Planning and Distance

For We always recommend that you perform a site survey before taking any real estimations for the area and creating a bill of materials.

Cisco 1500 Series Access Points

The RAP-to-MAP ratio is the starting point. For general planning purposes, the current ratio is 20 MAPs per RAP.

We recommend the following values for cell planning and distance in nonvoice networks:

- RAP-to-MAP ratio—Recommended maximum ratio is 20 MAPs per RAP.
- AP-to-AP distance—A spacing of no more than of 2000 feet (609.6 meters) between each mesh access point is recommended. When you extend the mesh network on the backhaul (no client access), use a cell radius of 1000 feet (304.8 meters).
- Hop count—Three to four hops.
    - One square mile in feet (52802), is nine cells and you can cover one square mile with approximately three or four hops (see Figure 12: Cell Radius of 1000 Feet and Access Point Placement for Nonvoice Mesh Networks, on page 54 and Figure 13: Path Loss Exponent 2.3 to 2.7)
- For 2.4 GHz, the local access cell size radius is 600 feet (182.88 meters). One cell size is around 1.310 x 106, so there are 25 cells per square mile. (See Figure 14: Cell Radius of 600 Feet and Access Point.

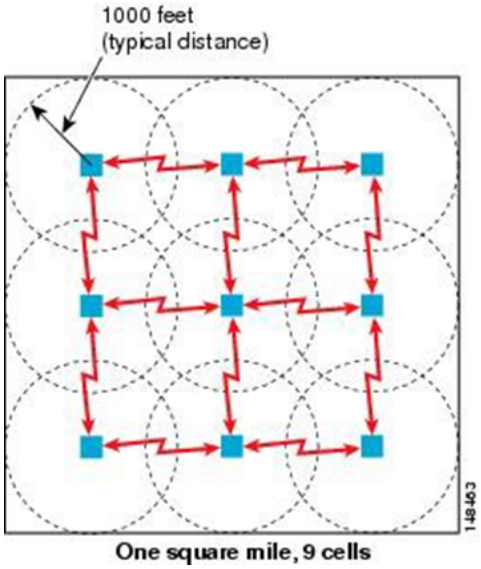


Figure 12 Cell Radius of 1000 Feet and Access Point Placement for Nonvoice Mesh Networks

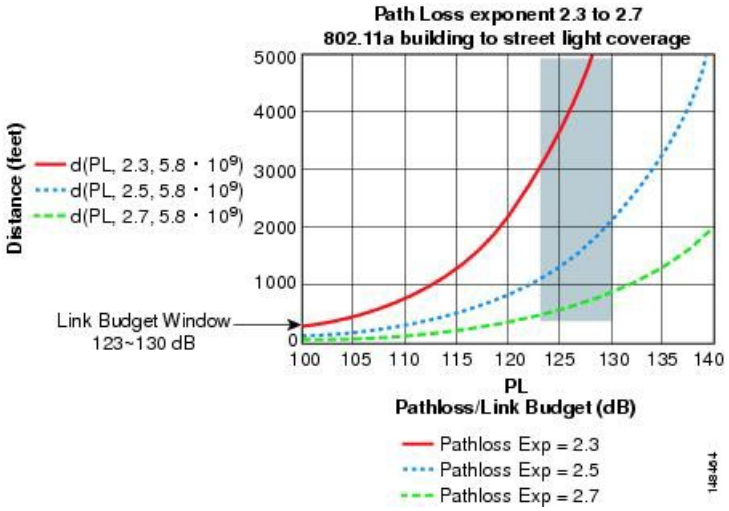


Figure 13 Path Loss Exponent 2.3 to 2.7

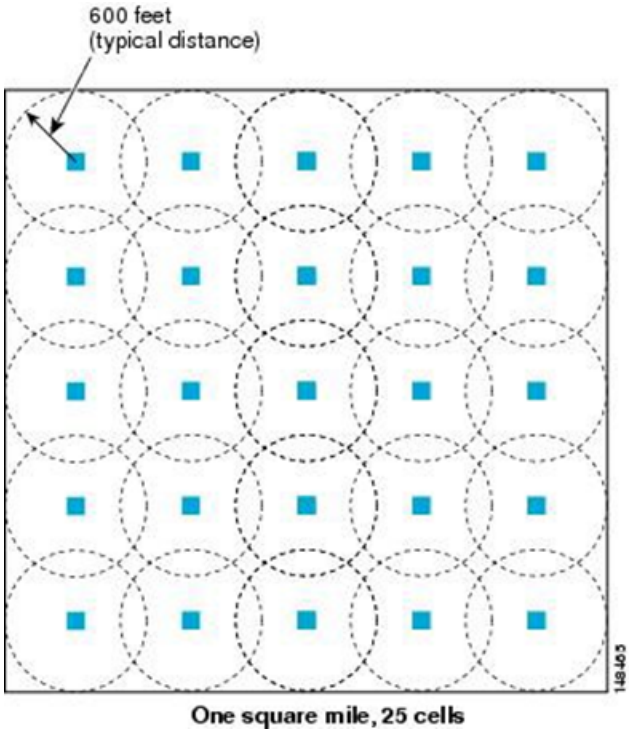


Figure 14 Cell Radius of 600 Feet and Access Point Placement for Nonvoice Mesh Networks

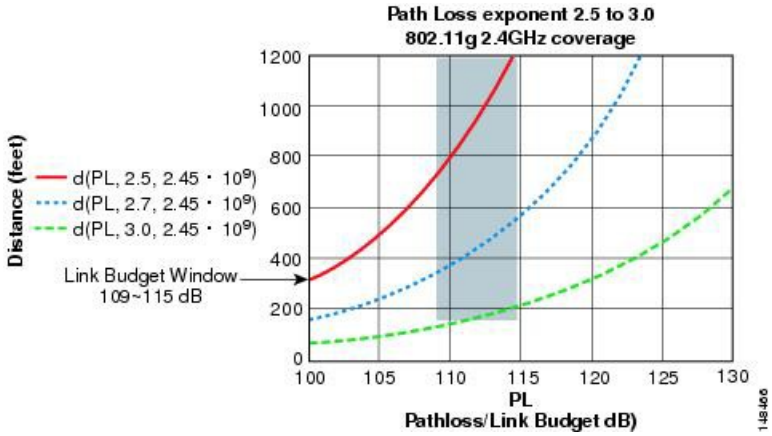


Figure 15 Path Loss Exponent 2.5 to 3.0

For the Cisco 1500 Series Access Points

As seen in the previous section, we recommend a cell radius of 600 feet, and an AP to AP distance of 1200 feet. Normally, an AP to AP distance that is twice the AP to client distance is recommended. That is, if we halve the AP to AP distance, we will get the approximate cell radius.

The WAVE-2 AP15XX series offers comparatively better range and capacity as it has the 802.11n/a/ac functionality. It has advantages of ClientLink (Beamforming) in downstream, better receiver sensitivities because of MRC in upstream, multiple transmitter streams and a few other advantages of 802.11a/n/ac such as channel combining and so on.

Assumptions for the Cisco Range Calculator

- The Range Calculator has been edited to stay within limitations for Tx power and EIRP under the listed regulatory domains. There may be cases where it exceeds the limitations. You must verify that the installation is within the laws of the location in which it is being installed.

- All antenna ports must be used for external antenna models for effective performance. Otherwise, range is significantly compromised.
- The Tx power is the total composite power of both Tx paths.
- Rx sensitivity is the composite sensitivity of all three Rx paths. That is, MRC is included.
- The Range Calculator assumes that ClientLink (Beamforming) is switched on.
- When you use the Range Calculator, available power levels change based upon the regulatory domain, the antenna (or antenna gain) selected, and the data rate selected. You must verify all parameters after making any parameter changes.
- You can select a different antenna than the two that are available by default. If you enter a high gain antenna and choose a power that goes over the EIRP limit, then you get a warning and the range equals 0.
- You can choose only the channels that the access point is certified for.
- You can only select only valid power levels.

The RAPs shown in the figure below are simply a starting point. The goal is to use the RAP location in combination with the RF antenna design to ensure that there is a good RF link to the MAP within the core of the cell, which means that the physical location of the RAPs can be on the edge of the cell, and a directional antenna is used to establish a link into the center of the cell. Therefore, the wired network location of a RAP might play host to the RAP of multiple cells, as shown in the figure below.

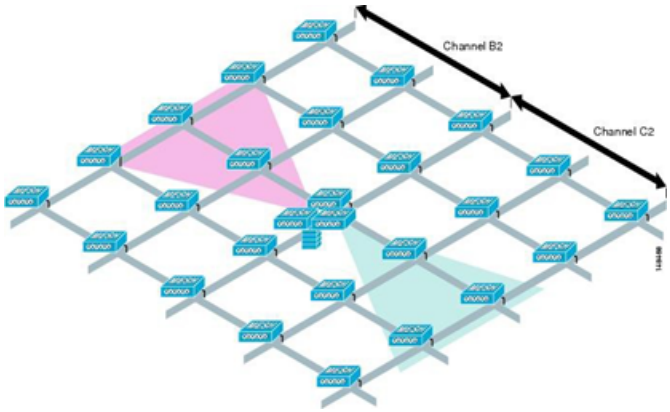


Figure 16 PoP with Multiple RAPs

When the basic cell composition is settled, the cell can be replicated to cover a greater area. When replicating the cells, a decision needs to be made whether to use the same backhaul channel on all cells or to change backhaul channels with each cell. In the example shown in the figure below, various backhaul channels (B2, C2, and D2) per cell have been chosen to reduce the co-channel interference between cells.

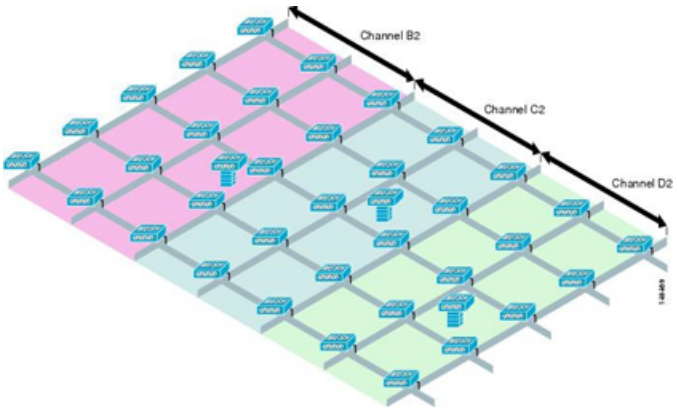


Figure 17 Multiple RAP and MAP Cells

Choosing various channels reduces the co-channel interference at the cell boundaries, at the expense of faster mesh convergence, because MAPs must fall back to seek mode to find neighbors in adjacent cells. In areas of high-traffic density, co-channel interference has the highest impact, which is likely to be around the RAP.

If RAPs are clustered in one location, a different channel strategy is likely to give optimal performance; if RAPs are dispersed among the cells, using the same channel is less likely to degrade performance.

When you lay out multiple cells, use channel planning similar to standard WLAN planning to avoid overlapping channels, as shown in the figure below.

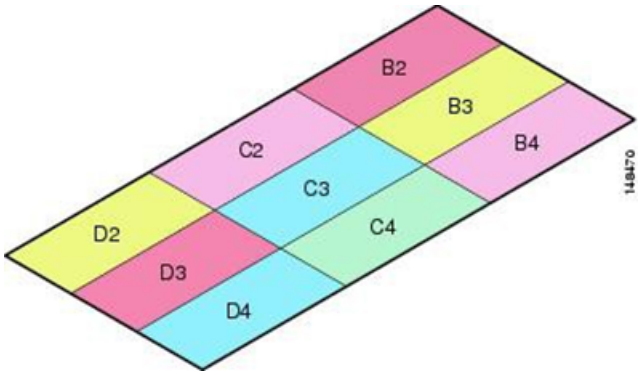


Figure 18 Laying out Various Cells

If possible, the channel planning should also minimize channel overlap in cases where the mesh has expanded to cover the loss of a RAP connection, as shown in the figure below.

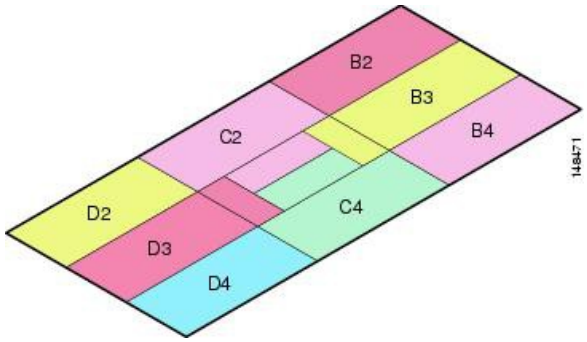


Figure 19 Failover Coverage

Collocating Mesh Access Points

The following recommendations provide guidelines to determine the required antenna separation when you collocate

Wave-2 AP15XX s on the same tower. The recommended minimum separations for antennas, transmit powers, and channel spacing are addressed.

The goal of proper spacing and antenna selection is to provide sufficient isolation by way of antenna radiation pattern, free space path loss, and adjacent or alternate adjacent channel receiver rejection to provide independent operation of the collocated units. The goal is to have negligible throughput degradation due to a CCA hold-off, and negligible receive sensitivity degradation due to a receive noise floor increase.

You must follow antenna proximity requirements, which depend upon the adjacent and alternate adjacent channel usage.

Collocating Wave-2 AP15XX s on Adjacent Channels

If two collocated Wave-2 AP15XX s operate on adjacent channels such as channel 149 (5745 MHz) and channel 152 (5765 MHz), the minimum vertical separation between the two Wave-2 AP15XX s is 40 feet (12.192 meters) (the requirement applies for mesh access points equipped with either 8 dBi omnidirectional or 17 dBi high-gain directional patch antennas).

If two collocated Wave-2 AP15XX s operate on channels 1, 6, or 11 (2412 to 2437 MHz) with a 5.5-dBi omnidirectional antenna, then the minimum vertical separation is 8 feet (2.438 meters).

Collocating Wave-2 AP15XX s on Alternate Adjacent Channels

If two collocated Wave-2 AP15XX s operate on alternate adjacent channels such as channel 149 (5745 MHz) and channel 157 (5785 MHz), the minimum vertical separation between the two Wave-2 AP15XX s is 10 feet (3.048 meters) (the requirements applies for mesh access points equipped with either 8-dBi omnidirectional or 17-dBi high-gain directional patch antennas).

If two collocated Wave-2 AP15XX s operate on alternate adjacent channels 1 and 11 (2412 MHz and 2462 MHz) with a 5.5-dBi omnidirectional antenna, then the minimum vertical separation is 2 feet (0.609 meters).

In summary, a 5-GHz antenna isolation determines mesh access point spacing requirements and antenna proximity must be followed and is dependent upon the adjacent and alternate adjacent channel usage.

DFS and None-DFS Channel Scan

Non–DFS channel scan

- A MAP goes off-channel periodically, transmits NDReq broadcast packets on the selected off-channel, and shall receive NDResp packets from all ‘reachable’ neighbors
- Off-channel scan periodicity will occur every 3 seconds and stay for a maximum of 50 milliseconds per off-channel
- NDReq has to be transmitted every 10 milliseconds to send at least 4 messages within 50 milliseconds dwell time to hear better from each neighbor

Wireless Propagation Characteristics

The 2.4-GHz band provides better propagation characteristics than 5 GHz, but 2.4 GHz is an unlicensed band and has historically been affected with more noise and interference to date than the 5-GHz band. In addition, because there are only three backhaul channels in 2.4 GHz, co-channel interference would result. Therefore, the best method to achieve comparable capacity is by reducing system gain (that is, transmit power, antenna gain, receive sensitivity, and path loss) to create smaller cells. These smaller cells require more access points per square mile (greater access point density).

The table below shows comparison of 2.4-GHz and 5-GHz Bands, provides a comparison of the 2.4-GHz and 5-GHz bands.

Comparison of 2.4-GHz and 5-GHz Bands

2.4-GHz Band Characteristics	5-GHz Band Characteristics

3 channels	22 channels (-A/-B regulatory domain)
More prone to co-channel interference	No co-channel interference
Lower power	Higher power
Lower SNR requirements given lower data rates	Higher SNR requirements given higher data rates
Better propagation characteristics than 5 GHz but more susceptible to noise and interference	Worse propagation characteristics than 2.4 GHz but less susceptible to noise and interference
Unlicensed band. Widely available throughout the world. 2.4 GHz has more penetration capability across the obstacles due to a larger wavelength. In addition, 2.4 GHz has lower date rates which increases the success of the signal to reach the other end.	Not as widely available in the world as 2.4-GHz. Licenses in some countries.

CleanAir and RRM in Mesh

The 1560 series access points contain the CleanAir chipset, allowing full CleanAir support.

CleanAir in mesh can be implemented on the 2.4-GHz radio and provides clients complete 802.11n/a/ac data rates while detecting, locating, classifying, and mitigating radio frequency (RF) interference. This provides a carrier class management and customer experience and ensures that you have control over the spectrum in the deployed location. CleanAir enabled RRM technology on the outdoor platform detects, quantifies, and mitigates Wi-Fi and non-Wi-Fi interference on 2.4-GHz radios. Access points operating in Bridge Mode support CleanAir in 2.4 GHz client access mode.

CleanAir AP Modes of Operation

Bridge (Mesh) Mode AP—CleanAir capable access points offer complete CleanAir functionality in the 2.4 GHz band and CleanAir advisor on the 5 GHz radio. This is across all access points that operate in Bridge mode.

Tight silicon integration with the Wi-Fi radio allows the CleanAir hardware to listen between traffic on the channel that is currently being served with no penalty to throughput of attached clients. That is, line rate detection without interrupting client traffic.

Bridge mode access points support Radio Resource Management (RRM) on the 2.4 GHz band, which helps to mitigate the interference from Wi-Fi interferers. RRM is only available on the 5 GHz band, if a Bridge mode RAP has no child MAPs. Mesh disables RRM off channel activity when first child joins even if RRM is enabled.

A CleanAir Mesh AP only scans one channel of each band continuously. In a normal deployment density, there should be many access points on the same channel, and at least one on each channel, assuming RRM is handling channel selection. In 2.4 GHz, access points have sufficient density to ensure at least three points of classification. An interference source that uses narrow band modulation (operates on or around a single frequency) is only detected by access points that share the frequency space. If the interference is a frequency hopping type (uses multiple frequencies—generally covering the whole band), it is detected by every access point that can hear it operating in the band.

Monitor Mode AP (MMAP)—A CleanAir monitor mode AP is dedicated and does not serve client traffic. The monitor mode ensures that all bands-channels are routinely scanned. The monitor mode is not available for access points in bridge (mesh) mode because in a mesh environment, access points also talk to each other on the backhaul. If a mesh AP (MAP) is in the monitor mode, then it cannot perform mesh operation.

Local Mode AP— When an outdoor access point is operating in local mode, it can perform full CleanAir and RRM on both the 2.4 GHz and 5 GHz channels. It will predominately scan its primary channel, but will periodically go off-channel to scan the rest of the spectrum. Enhanced Local Mode (ELM) wIPS detection is not available on the 1532, 1550, or 1570.

Spectrum Expert Connect Mode (optional) (SE Connect)—An SE Connect AP is configured as a dedicated spectrum sensor that allows connection of the Cisco Spectrum Expert application running on a local host to use the CleanAir AP as a remote spectrum sensor for the local application. This mode allows viewing of the raw spectrum data such as FFT plots and detailed measurements. This mode is intended for remote troubleshooting only.

Pseudo MAC (PMAC) and Merging

PMAC and Merging phenomenon is similar to the one for Generation 2 access points in local mode. A PMAC is calculated as part of the device classification and included in the interference device record (IDR). Each AP generates the PMAC independently. While it is not identical for each report (at a minimum the measured RSSI of the device is likely different at each AP), it is similar. The function of comparing and evaluating PMACs is called merging. The PMAC is not exposed to customer interfaces. Only the results of merging are available in the form of a cluster ID.

The same device can be detected by multiple APs. All the PMACs and IDRs are analyzed on the controller and a report is generated called a device cluster, which shows the APs detecting the device and the device cluster showing the AP which is hearing the device as strongest.

In this merging spatial proximity, RF proximity (RF neighbor relationship) work together. If there are six similar IDRs with 5 APs nearby and another one from an AP that is far away, it is unlikely that it is the same interferer. Therefore, a cluster is formed taking all these into account. MSE and the controller first rely on RF Neighbor lists to establish spatial proximity in a merge.

PMAC Convergence and Merging depends upon the following factors:

- Density of the sensors
- Quality of the observed classification
- RSSI from the interferer to the APs
- RF neighbor list at the APs

So RRM on 2.4 GHz in mesh also plays a key role in deciding the merging aspect. APs should be RF neighbors for any possibility of Merging. RF Neighbor list is consulted and spatial relationships for IDRs are considered for Merging.

Because there is no Monitor Mode in mesh, a single controller merging occurs on the controller. The result of a controller merge is forwarded to the MSE (if present) along with all of the supporting IDRs.

For more than one C9800 (possible in outdoor deployments), merging occurs on the MSE. MSE does more advanced merging and extracts location and historical information for interferers. No Location is performed on controller merged interferers. Location is done on the MSE.

After PMAC signature merging, you can identify which AP can hear the device, and which AP is the center of a cluster. In the figure above, the values are relevant to the band selected. The label R on AP indicates that the AP is a RAP and the line between APs shows the mesh relationship.

Event Driven Radio Resource Management and Persistence Device Avoidance

There are two key mitigation features that are present with CleanAir. Both rely directly on information that can only be gathered by CleanAir. Event Driven Radio Resource Management (EDRRM) and Persistence Device Avoidance (PDA). For mesh networks, they work exactly the same way as for non-mesh networks in the 2.4-GHz band.

Note: EDRRM and PDA are only available in a Greenfield installation and configured off by default.

CleanAir Access Point Deployment Recommendations

CleanAir is a passive technology that does not affect the normal operation of Wi-Fi networks. There is no inherent difference between a CleanAir deployment and a mesh deployment.

Locating a non-Wi-Fi device has a lot of variables to consider. Accuracy increases with power, duty cycle, and the number of channels hearing the device. This is advantageous because higher power, higher duty cycle, and devices that impact multiple channels are considered to be severe with respect to interference to networks.

Note: There is no guarantee of accuracy for location of non- Wi-Fi devices.

There are a lot of variables in the world of consumer electronics and unintentional electrical interference. Any expectation

of accuracy that is derived from current Client or Tag location accuracy models does not apply to non-Wi-Fi location and CleanAir features.

Important notes to consider:

- CleanAir mesh AP supports the assigned channel only.
- Band Coverage is implemented by ensuring that channels are covered.
- The CleanAir mesh AP can hear very well, and the active cell boundary is not the limit.
- For Location solutions, the RSSI cutoff value is –75 dBm.
- A minimum of three quality measurements is required for location resolution.

In most deployments, it is difficult to have a coverage area that does not have at least three APs nearby on the same channel in the 2.4-GHz band. In locations where there is minimal density, while the location resolution is likely not supported, the active user channel is protected.

Deployment considerations are dependent upon planning the network for desired capacity and ensuring that you have the correct components and network paths in place to support CleanAir functions. RF proximity and the importance of RF Neighbor Relations cannot be understated. It is important to keep in mind the PMAC and the merging process. If a network does not have a good RF design, the neighbor relations is affected, which in turn affects CleanAir performance.

The AP Density recommendations for CleanAir remain the same as normal mesh AP deployment.

Location resolution in the Outdoors is to the nearest AP. Devices are located near the AP which is physically closest to the device. It is advisable to assume closest AP resolution.

It is possible to deploy a few 1530 APs (non-CleanAir) with an installation that consists of 1562 and 1572 APs (CleanAir). This deployment can work from a client and coverage standpoint as these access points are fully interoperable with each other. The complete CleanAir functionality depends on all access points being CleanAir enabled. Detection can be affected, and mitigation is not recommended.

A CleanAir AP actively serving clients can only monitor the assigned channel that it is serving. In an area where you have multiple access points serving clients in close proximity, the channels being served by CleanAir access points can drive CleanAir features. Legacy non-CleanAir access points rely on RRM, and mitigate interference issues, but not report the type and severity as CleanAir access points do to the system level.

For more information about mixed systems, see https://www.cisco.com/c/en/us/support/docs/wireless/5500-series-wireless-controllers/112139-cleanair-uwn-guide-00.html

CleanAir Advisor

If CleanAir is enabled on a backhaul radio, CleanAir Advisor is activated. CleanAir Advisor generates Air Quality Index (AQI) and Interferer Detection Reports (IDR) but the reports are only displayed in the controller. No action is taken through event driven RRM (ED-RRM). CleanAir Advisor is only present on the 5-GHz.

Wireless Mesh Mobility Groups

A mobility group allows controllers to peer with each other to support seamless roaming across controller boundaries. APs learn the IP addresses of the other members of the mobility group after the CAPWAP Join process. A controller can be a member of a single mobility group which can contain up to 24 controllers. Mobility is supported across 72 controllers. There can be up to 72 members (C9800s) in the mobility list with up to 24 members in the same mobility group (or domain) participating in client hand-offs. The IP address of a client does not have to be renewed in the same mobility domain. Renewing the IP address is irrelevant in the controller-based architecture when you use this feature

Multiple Controllers

The consideration in distance of the CAPWAP controllers from other CAPWAP controllers in the mobility group, and the

distance of the CAPWAP controllers from the RAP, is similar to the consideration of an CAPWAP WLAN deployment in an enterprise.

There are operational advantages to centralizing CAPWAP controllers, and these advantages need to be traded off against the speed and capacity of the links to the CAPWAP APs and the traffic profile of the WLAN clients using these mesh access points.

If the WLAN client traffic is expected to be focused on particular sites, such as the Internet or a data center, centralizing the controllers at the same sites as these traffic focal points gives the operational advantages without sacrificing traffic efficiency.

If the WLAN client traffic is more peer-to-peer, a distributed controller model might be a better fit. It is likely that a majority of the WLAN traffic are clients in the area, with a smaller amount of traffic going to other locations. Given that many peer-to-peer applications can be sensitive to delay and packet loss, you should ensure that traffic between peers takes the most efficient path.

Given that most deployments see a mix of client-server traffic and peer-to peer traffic, it is likely that a hybrid model of CAPWAP controller placement is used, where points of presence (PoPs) are created with clusters of controllers placed in strategic locations in the network.

The CAPWAP model used in the wireless mesh network is designed for campus networks; that is, it expects a high-speed, low-latency network between the CAPWAP mesh access points and the CAPWAP controller.

Increasing Mesh Availability

In the Cell Planning Distance section, a wireless mesh cell of one square mile was created and then built upon. This wireless mesh cell has similar properties to the cells used to create a cellular phone network because the smaller cells (rather than the defined maximum cell size) can be created to cover the same physical area, providing greater availability or capacity. This process is done by adding a RAP to the cell. Similar to the larger mesh deployment, the decision is whether to use RAP on the same channel, as shown in Figure 20: Two RAPs per Cell with the Same Channel, or to use RAPs placed on different channels, as

shown in Figure 21: Two RAPs per Cell on Different Channels. The addition of RAPs into an area adds capacity and resilience to that area.

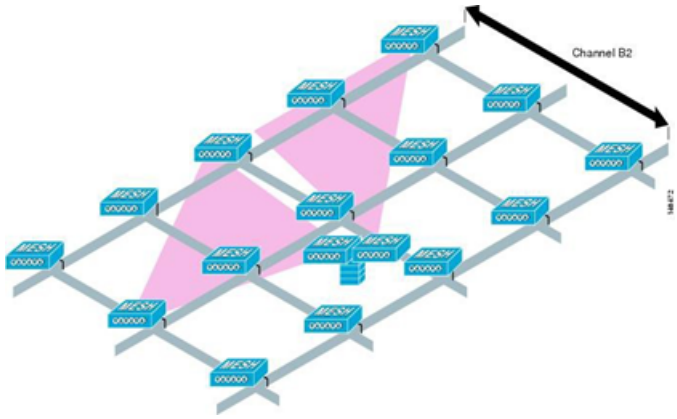


Figure 20 Two RAPs per Cell with the Same Channel

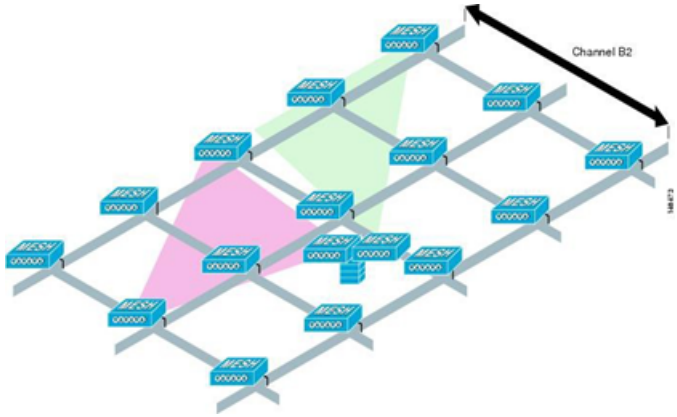


Figure 21 Two RAPs per Cell on Different Channels

Multiple RAPs

If multiple RAPs are to be deployed, the purpose for deploying these RAPs needs to be considered. If the RAPs are being deployed to provide hardware diversity, the additional RAP(s) should be deployed on the same channel as the primary RAP to minimize the convergence time in a scenario where the mesh transfers from one RAP to another. When you plan RAP hardware diversity, consider the 32 MAPs per RAP limitation.

If additional RAPs are deployed to primarily provide additional capacity, then the additional RAPs should be deployed on a different channel than its neighboring RAP to minimize the interference on the backhaul channels.

Adding a second RAP on a different channel also reduces the collision domain through channel planning or through RAP cell splitting. Channel planning allocates different nonoverlapping channels to mesh nodes in the same collision domain to minimize the collision probability. RAP cell splitting is a simple, yet effective, way to reduce the collision domain. Instead of deploying one RAP with omnidirectional antennas in a mesh network, two or more RAPs with directional antennas can be deployed. These RAPs collocate with each other and operate on different frequency channels. This process divides a large collision domain into several smaller ones that operate independently.

If the mesh access point bridging features are being used with multiple RAPs, these RAPs should all be on the same subnet to ensure that a consistent subnet is provided for bridge clients.

If you build your mesh with multiple RAPs on different subnets, MAP convergence times increase if a MAP has to fail over to another RAP on a different subnet. One way to limit this process from happening is to use different BGNs for segments in your network that are separated by subnet boundaries.

Connecting and Monitoring Cisco Mesh Access Points in the Network

This section describes how to connect the Cisco mesh access points to the network and then Monitor the Mesh network.

The wireless mesh terminates on two points on the wired network. The first location is where the RAP attaches to the wired network, and where all bridged traffic connects to the wired network. The second location is where the CAPWAP controller connects to the wired network; this location is where the WLAN client traffic from the mesh network connects to the wired network.

Adding Mesh Access Points to the Mesh Network

MAC Authorization

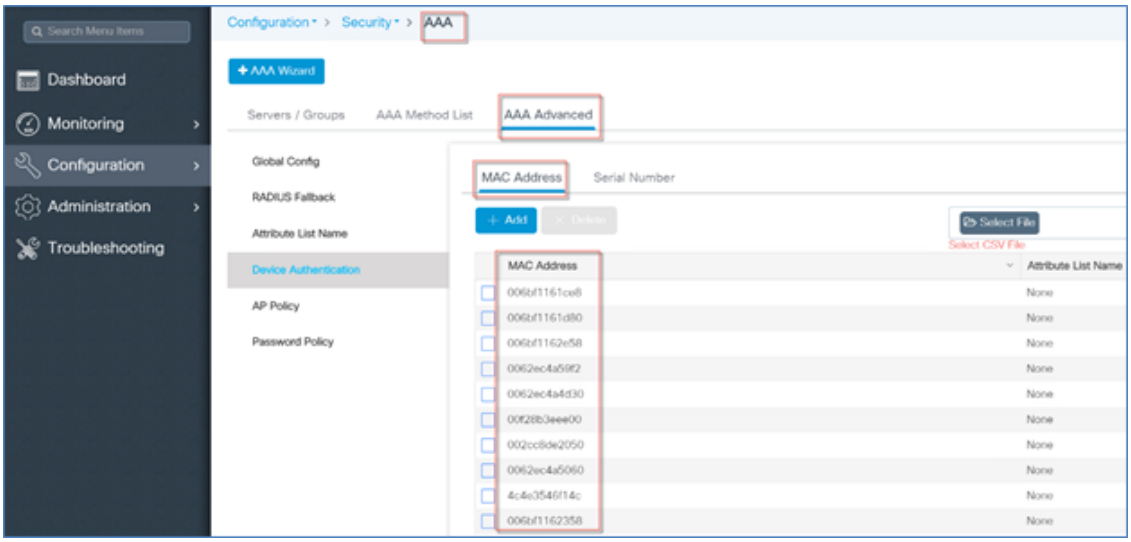
This section assumes that the controller is already active in the network and is operational.

In order to make a Mesh Access Point to join to the Controller, we must enter the MAC address of the access point into the controller. A controller only responds to CAPWAP join requests from mesh access points that appear in its authorization list. MAC filtering for Bridge mode Access Points is enabled by default on the controller, so only the MAC addresses need to be configured. The mac address used here would be the mac address labelled on the backside of the Access Point which is usually the Ethernet mac. MAC Authorization for Mesh APs connecting to C9800 over Ethernet backhaul will happen during

the CAPWAP Join process. MAC Authentication for Mesh APs which are joining to C9800 over Radio backhaul will happen when it tries to establish a secure AWPP link with the parent Mesh AP. For these APs, MAC Authorization will be skipped during CAPWAP Join process.

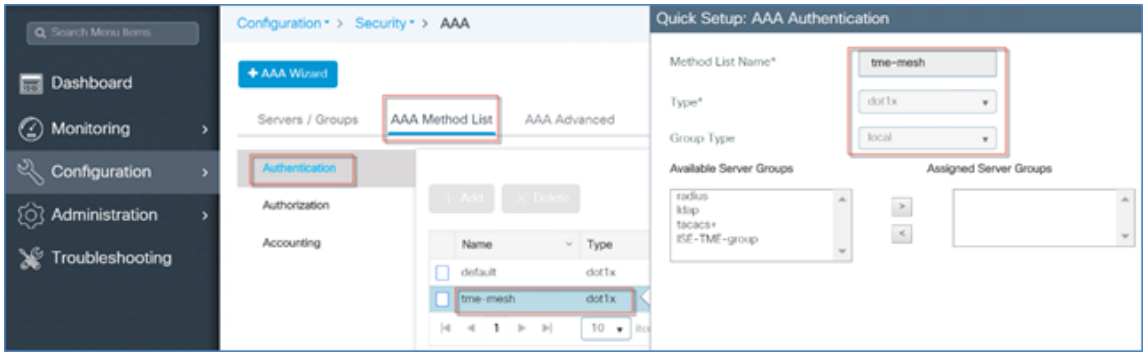
In C9800, MAC Authorization is supported internally, as well as using an external AAA server.

1. Under AAA tab Provision RAP/MAP MAC addresses, by adding the MAC addresses of the Mesh APs to be added to the network. There is also an option to import the Serial numbers or MAC addresses from the CSV file

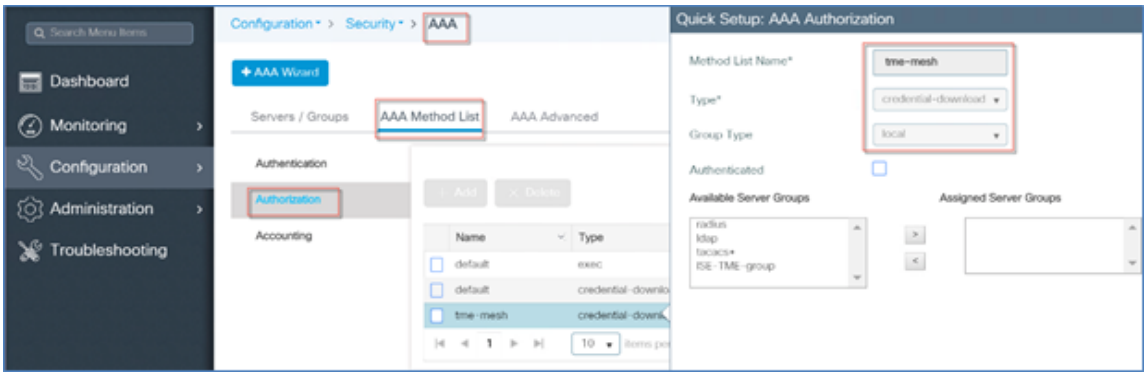


2. Configure AAA Authentication and Authorization Methods if Mesh AP will be using EAP authentication in the AP Join Profile. If PSK or only MAP's MAC address option will be used then AAA configuration step can be skipped.

Authentication Method Configuration

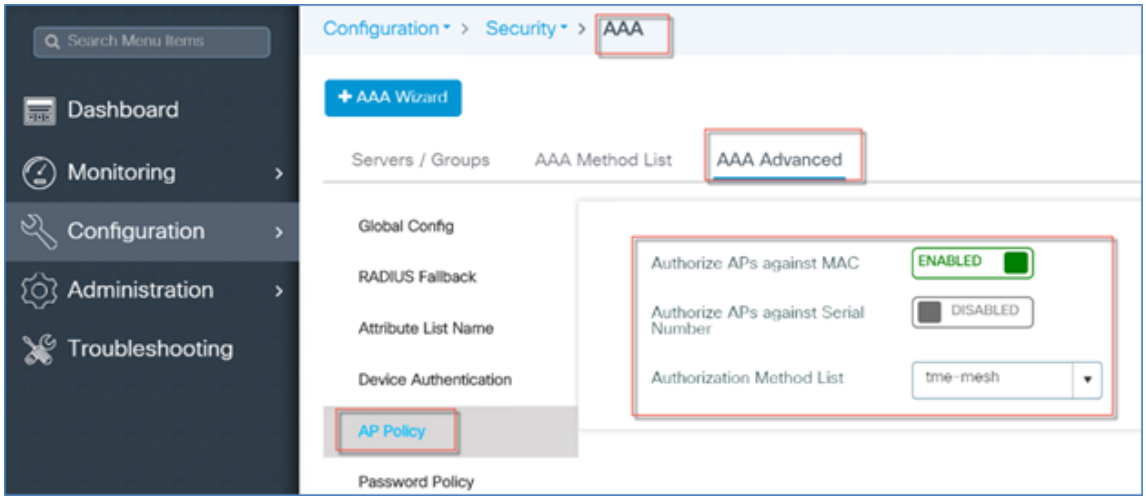


Authorization Method Configuration

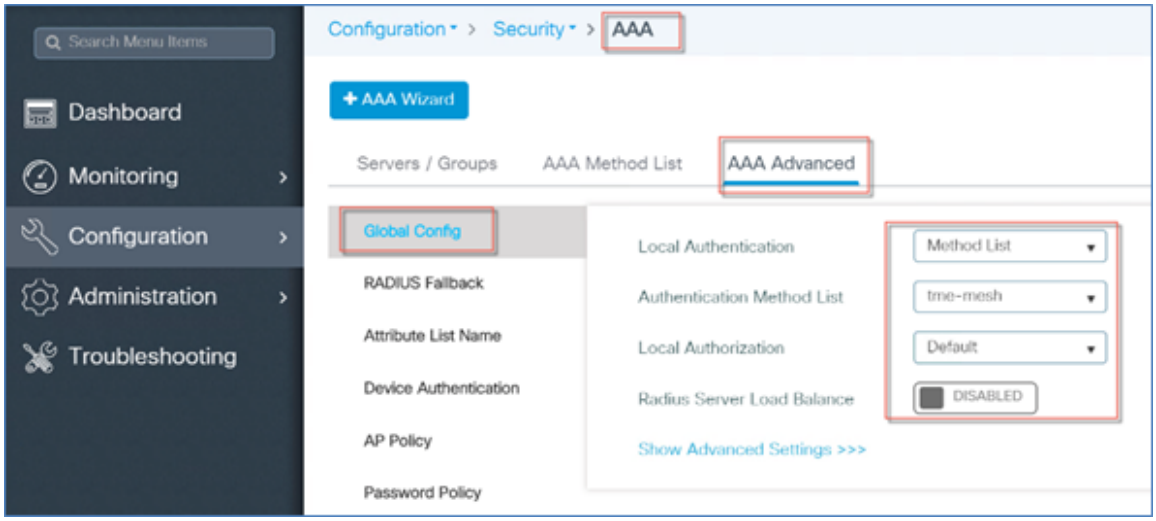


3. Under AAA configuration configure the AP Policy Authorization as shown in the example below, the Mesh AP

authorization is done via MAC addresses

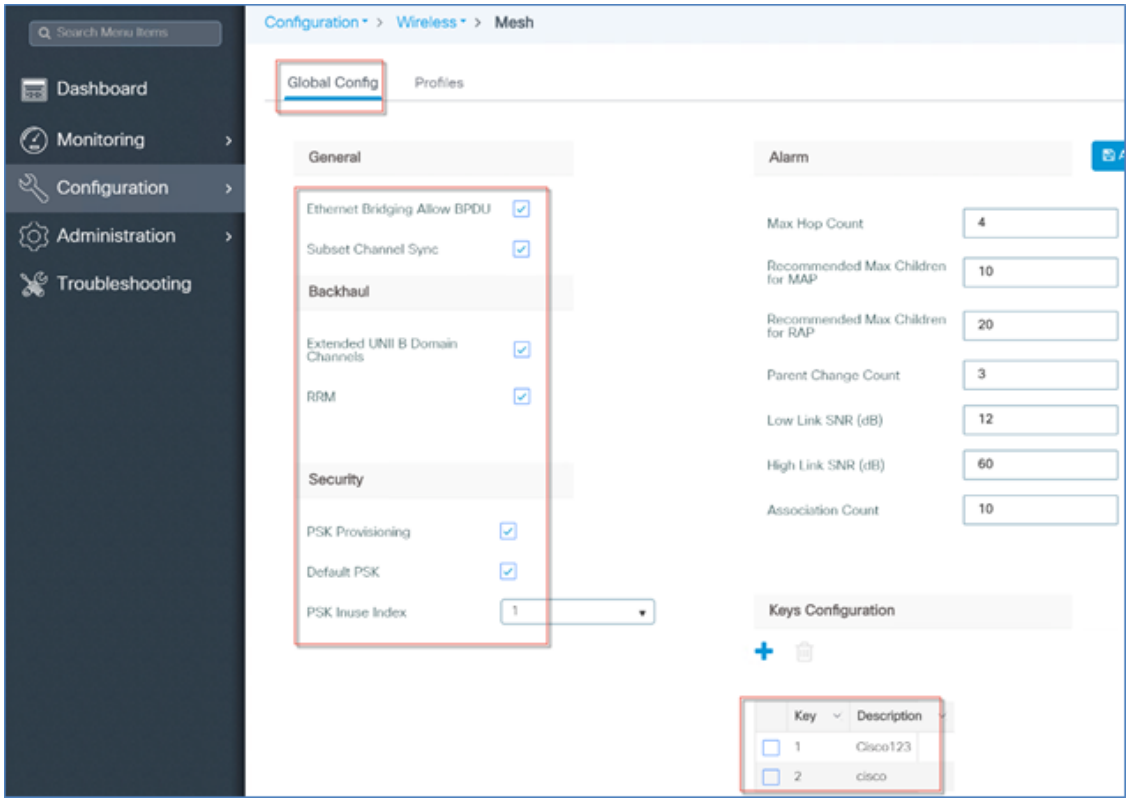


4. And lastly under AAA configuration setup a Global config setup the Local Authentication Method as shown in the example config below



5. Configure Global Mesh Settings and PSK Authentication

C9800 will support PSK based authentication for Mesh Access Points (MAP). Mesh APs with default ‘cisco’ PSK will join to the Wireless Controller, when the security method is configured as PSK. C9800 will also support PSK key configuration and provision-able PSK functionality.



PSK option with default passphrase also presents security risk and hijacking possibility. This issue will be prominently seen in overlapping deployments of two different SPs when the MAPs are used in a moving vehicle (public transportations, ferry, ship and so on.). This way, there is no restriction on MAPs to 'stick' to the SPs mesh network and MAPs can be hijacked / getting used by another SPs network / and cannot serve intended customers of SPs in a deployment.

This feature will help make a controlled mesh deployment and enhance MAPs security beyond default 'cisco' PSK used today. With this feature, MAPs which are configured with a custom PSK, will use this key to do their authentication with their parent Mesh APs and C9800.

C9800 will host a provisioning window upon setting with above command to allow child MAPs to join parent mesh AP with default PSK. After the provisioning window is un-set, child MAPs are not allowed to join with default PSK unchecked and they should only use the provisioned PSK to join to a parent Mesh AP. The provisioning window can be set and un-set at any time by the admin when Default PSK is checked or unchecked. At a time, C9800 will always store the last five provisioned PSK keys. Child Mesh APs will be able to authenticate with any one of these five PSK keys. This will ensure a mesh AP will get authenticated even if it has not received the latest PSK key.

C9800 will push configured PSK and time to all the Mesh APs. Configured PSK keys will not be shared across C9800s in the mobility group. Admin has to configure the same PSK key in all the C9800s in the mobility group. Distribution of new PSK to all the Mesh APs will be done as soon as the key is configured before they get disconnected from current security mode. This is to ensure they have latest PSK available before its next disconnect and able to join back. MAPs receive new PSK in encrypted format via CAPWAP control message, and store in its flash.

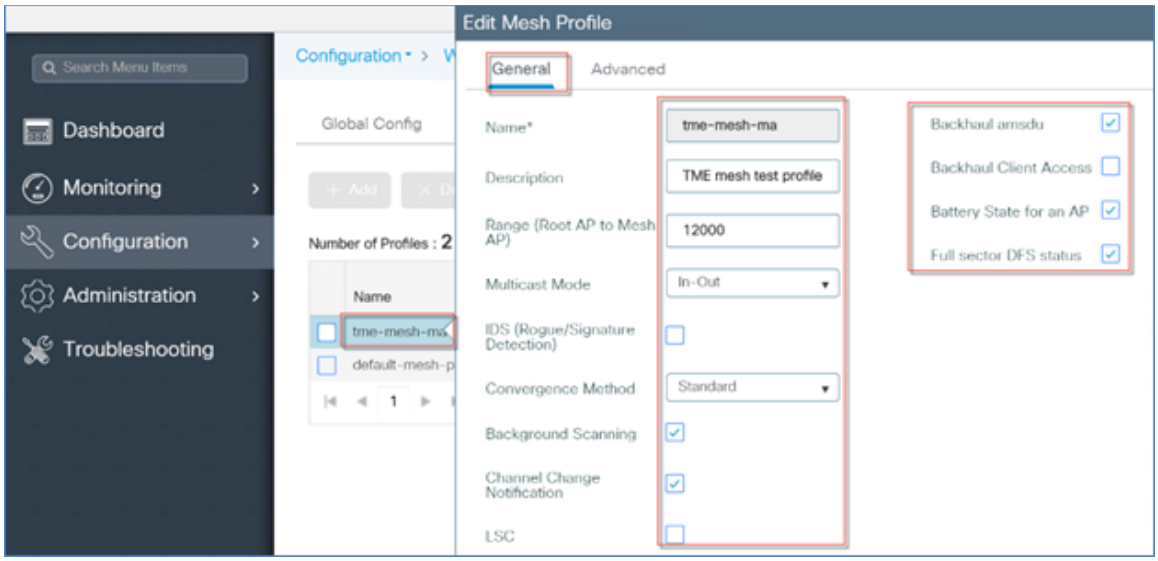
Every time new PSK key is configured in EC9800, only the new PSK key payload is pushed to AP connected. Also, the new PSK key shall be pushed from leaf node MAP towards up root RAP. C9800 expects AP do perform re-authentication with the new PSK pushed.

Configured PSK keys will be saved across reboot in the C9800 as well as on the Mesh AP. A C9800 can have total of 5 PSK keys and a default PSK key. A Mesh AP will delete its provisioned PSK only on factory reset and use default ‘cisco’ PSK for its further authentication. MAP shall never use default ‘cisco’ PSK after receiving first provisioned PSK. This is to avoid a MAP getting hijacked with default PSK.

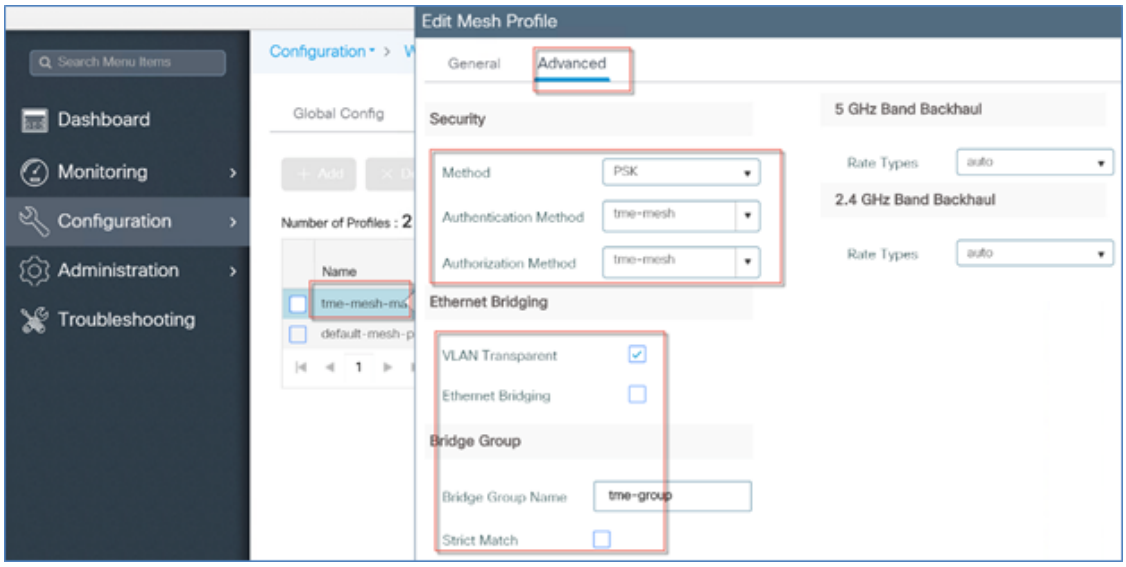
Admin should be able to delete a stale PSK from the list based on the description and it shall be sent to all APs and C9800s post configuration. Admin should be able to delete a provisioned PSK of a MAP from C9800. This is to remove any rogue MAPs joined to C9800 through provisioning window.

The PSK provisioning will be done for all the APs connected to C9800 irrespective of BGN or Mesh Profile.

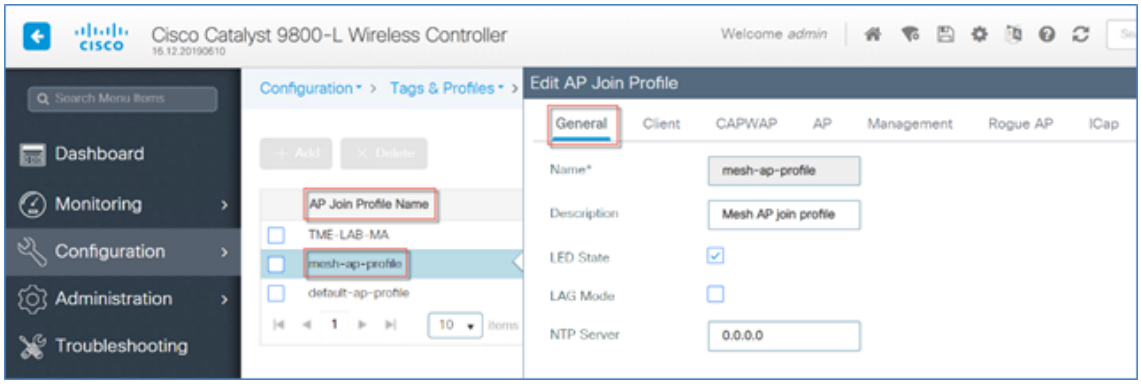
6. Configure Global Mesh Profile under General and Advanced Tabs as shown in the example below



7. Under the Mesh Advanced tab apply the previously configure Authentication and Authorization Methods on the Mesh Profile and also configure the 5 or 2.4GHz backhauls as shown in the example below

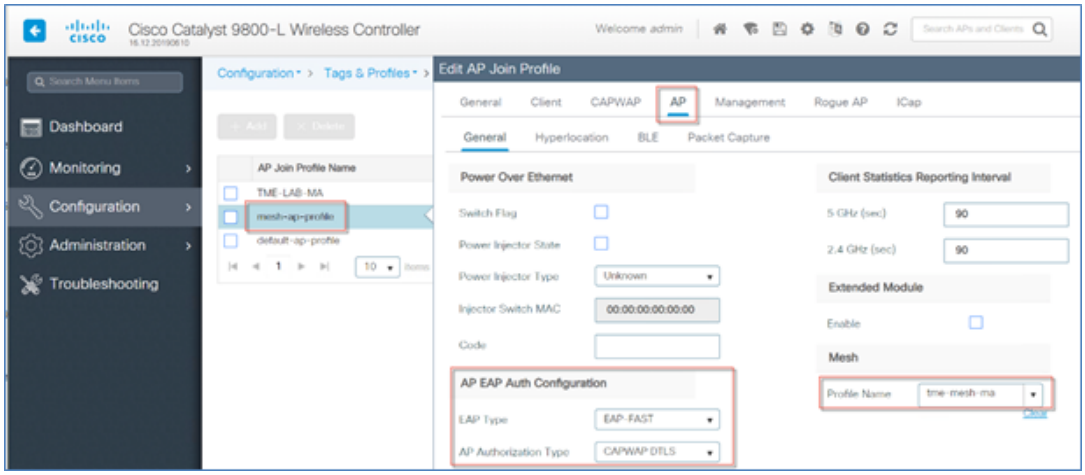


8. Configure AP Join Profile by adding Mesh Profile under the General tab as shown in the example below

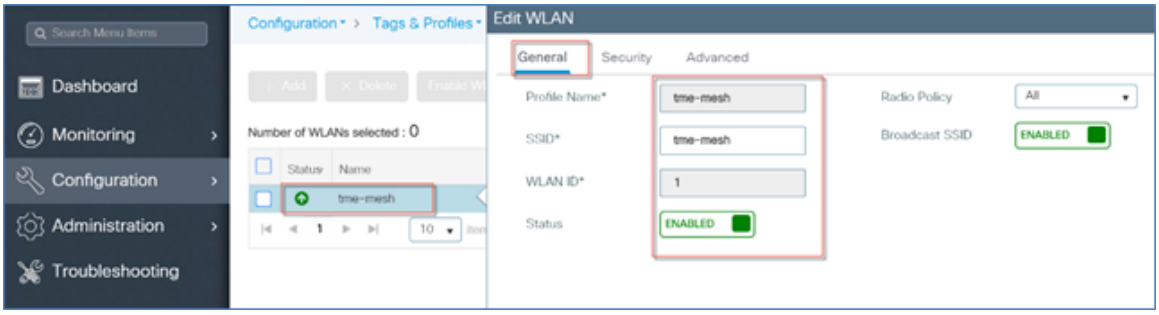


9. In the AP Join Profile apply the Previously configured Mesh AP Profile

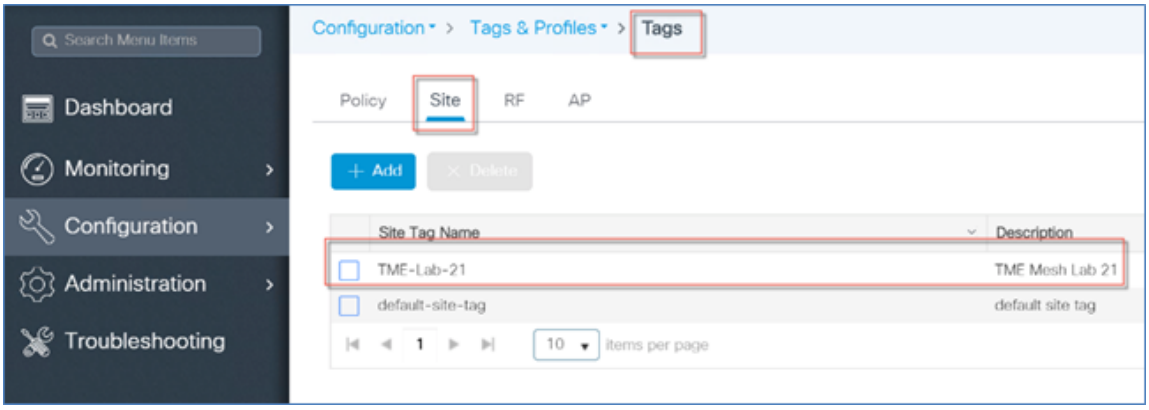
Optionally, configure AP EAP Authentication. EAP-FAST configuration is shown in the example below.



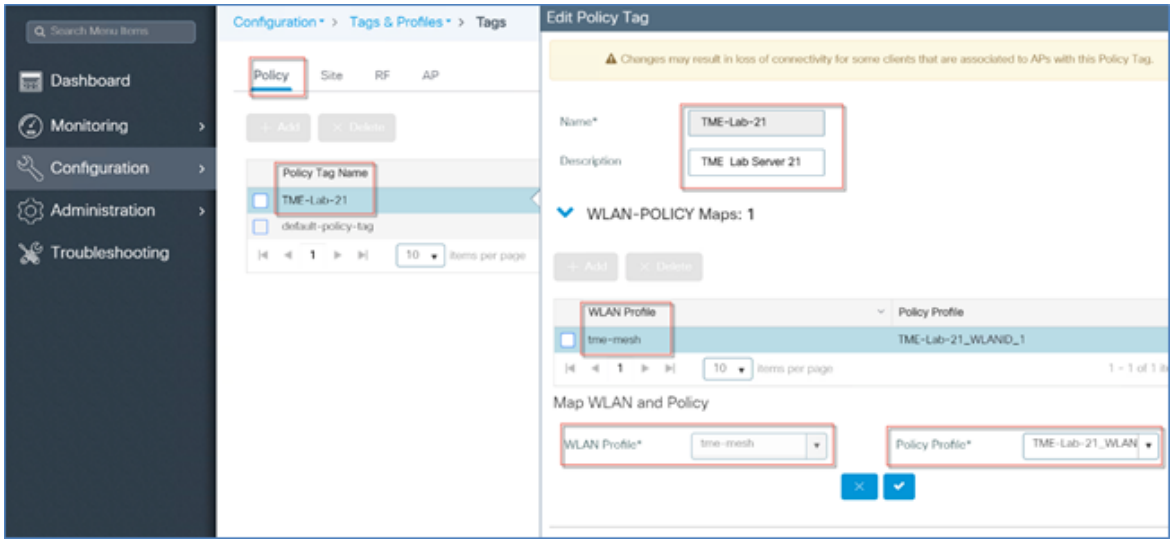
10. Make sure the WLAN is also configured as shown in the example below



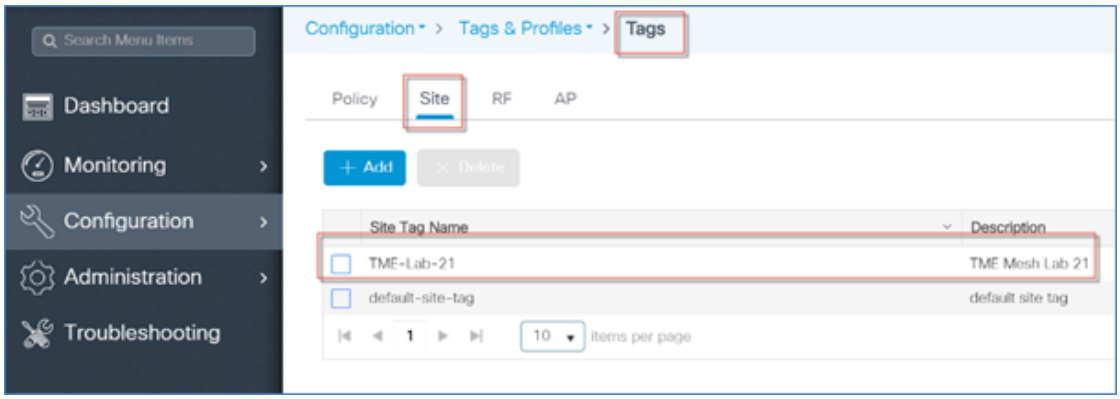
11. Add Policy Tag under Configuration Tags Policy



12. Under Policy Tag also configure WLAN Profile and Policy Profile configured earlier and as shown in the example below



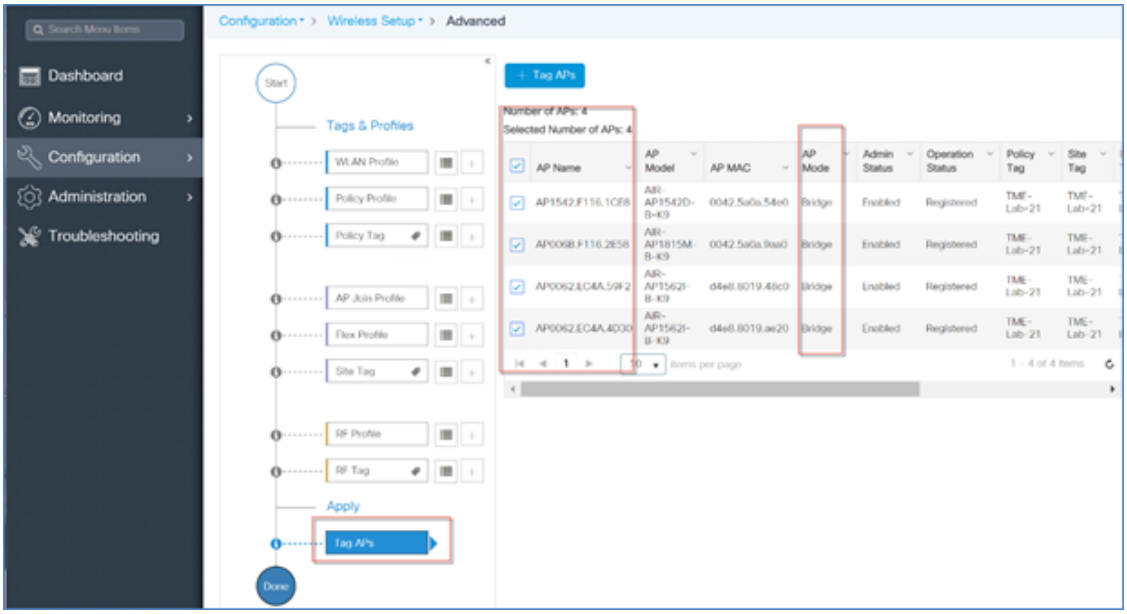
13. Add Mesh Site Tag under the Configure Tag tab

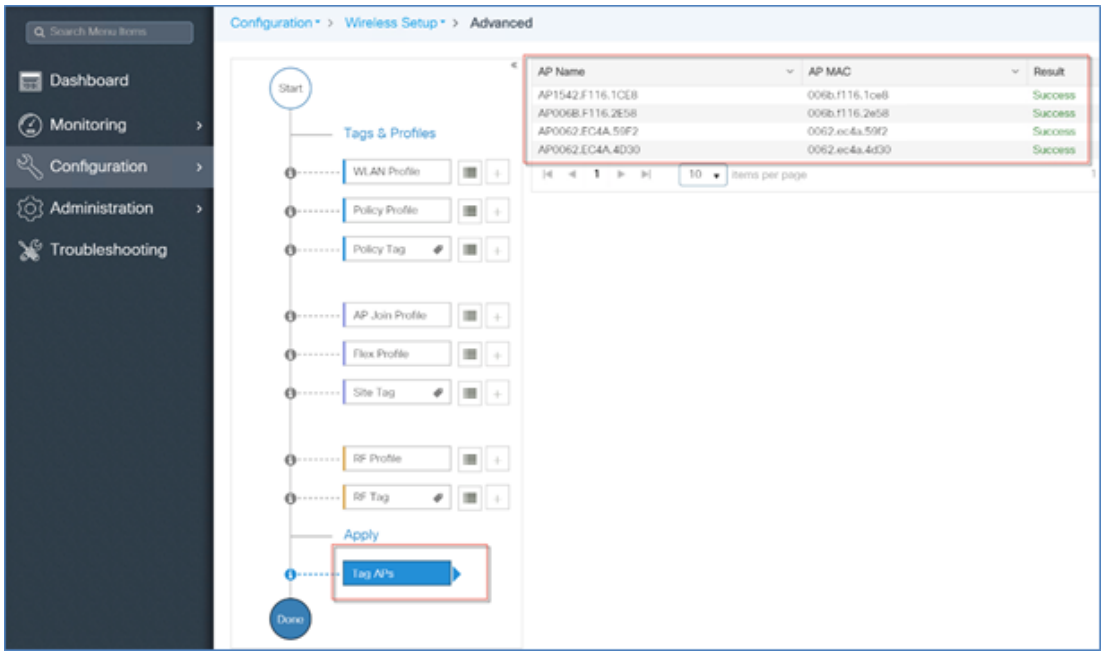


14. After Configuring Site Tag apply to it previously configured Mesh AP Join Profile as shown below to the Mesh Site Tag

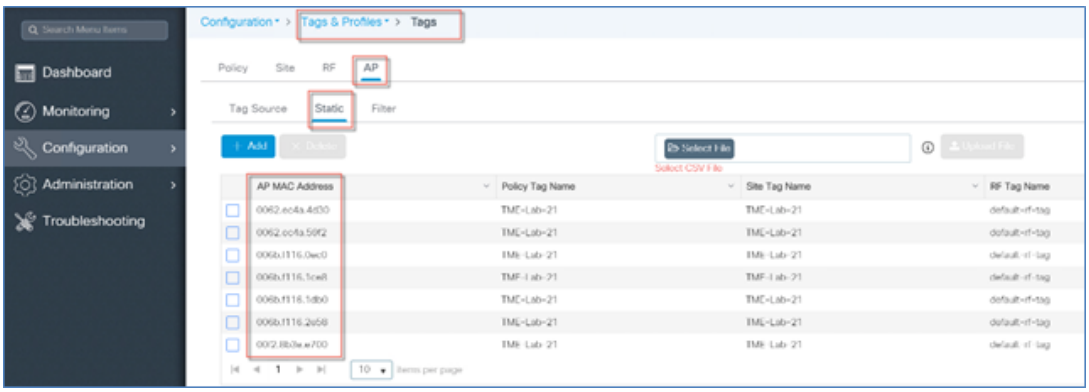


15. Finally Tag Mesh APs from the AP list with the configured Mesh Policies and Profiles as shown in the example below



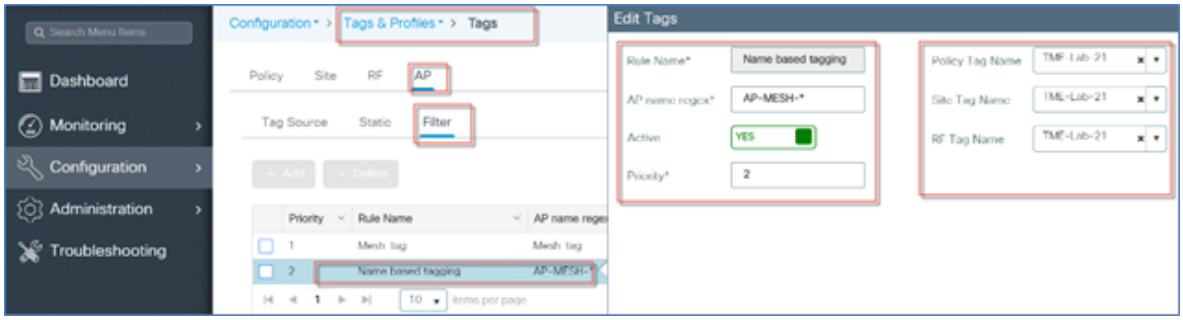


16. As shown in the example above the AP can be tagged statically as one option via the CLI or WebUI interface under Configuration > Tags and Profiles > AP > Static

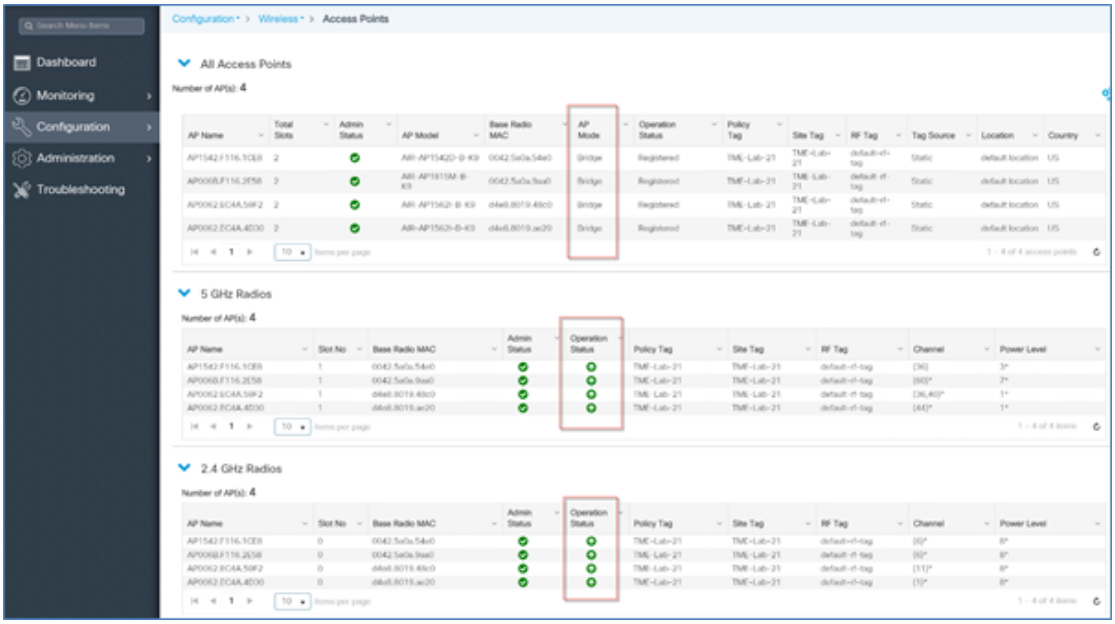


17. If the deployment consists of many Mesh APs then another option is available for tagging Bridge AP with a Tag. This can be accomplished if the AP names include certain characters in their name, for example if the Mesh AP that need to be

tagged include characters “AP-Mesh” then a tagging filter rule can be created as shown below. Then all the AP that include these characters can be tagged at once with a desired priority, Policy Tag, Site Tag and RF Tag. Name of the Mesh AP that need to be tagged have to be renamed during the staging phase to include the above mentioned or any other characters in their name.

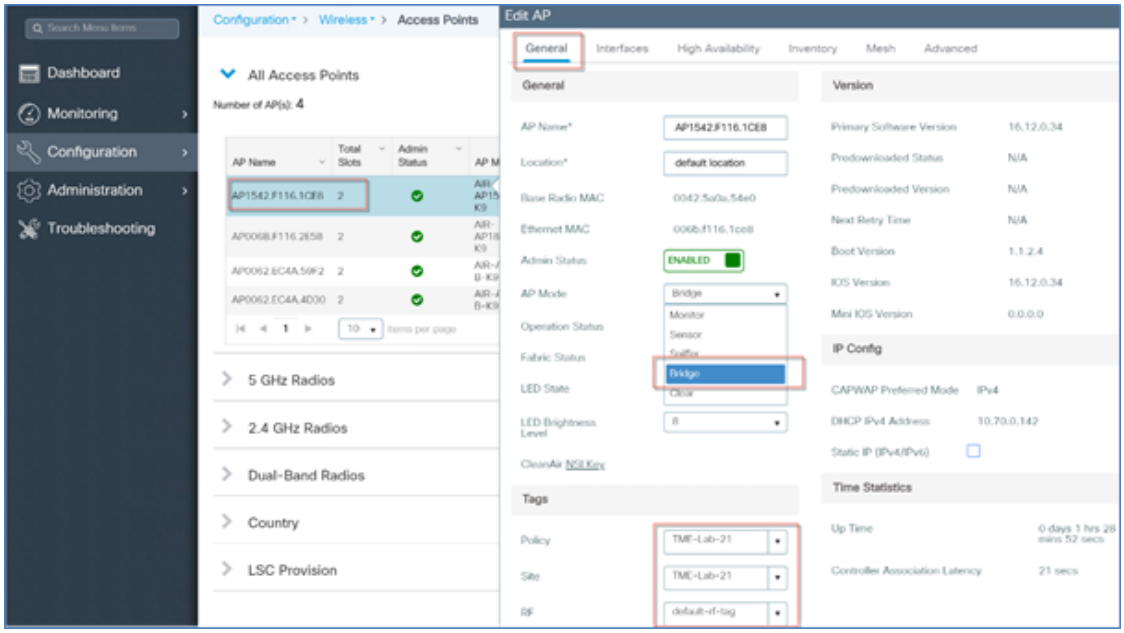


18. Verify the Mesh AP configured with the proper Site Tags and Policy Tags under Wireless > Access Points

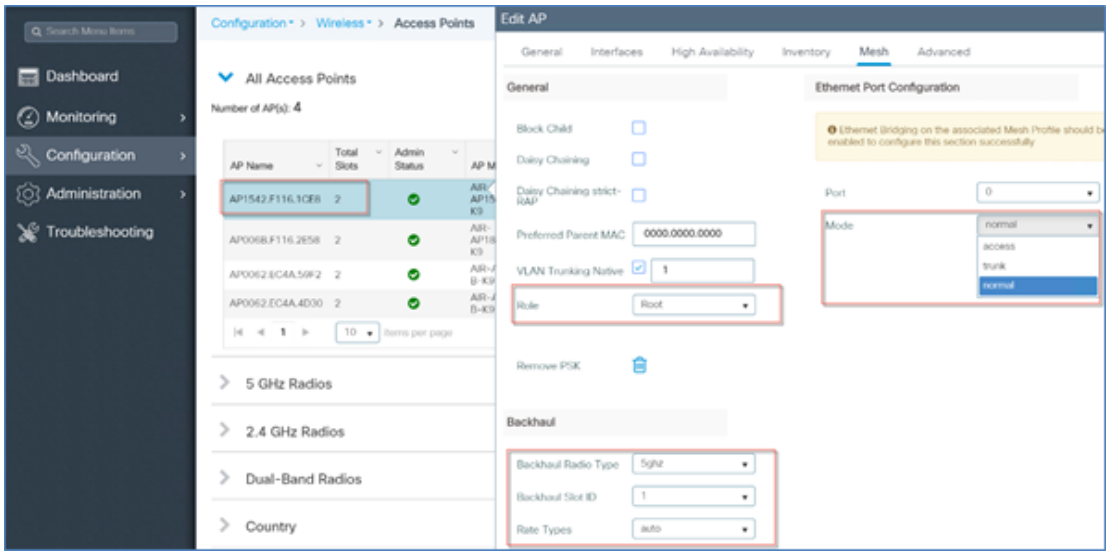


19. From the Configure> Access Point tab choose the earlier provisioned Mesh APs and set the AP Mode of all Mesh APs to “Bridge” mode as shown below.

APs will reboot and come back as Mesh APs and a new “Mesh” Tab will also appear.



20. After the Mesh Mode was configured, select AP from the list of the Mesh APs and Configure Role of Mesh AP as RAP or MAP for each individual Outdoor AP.

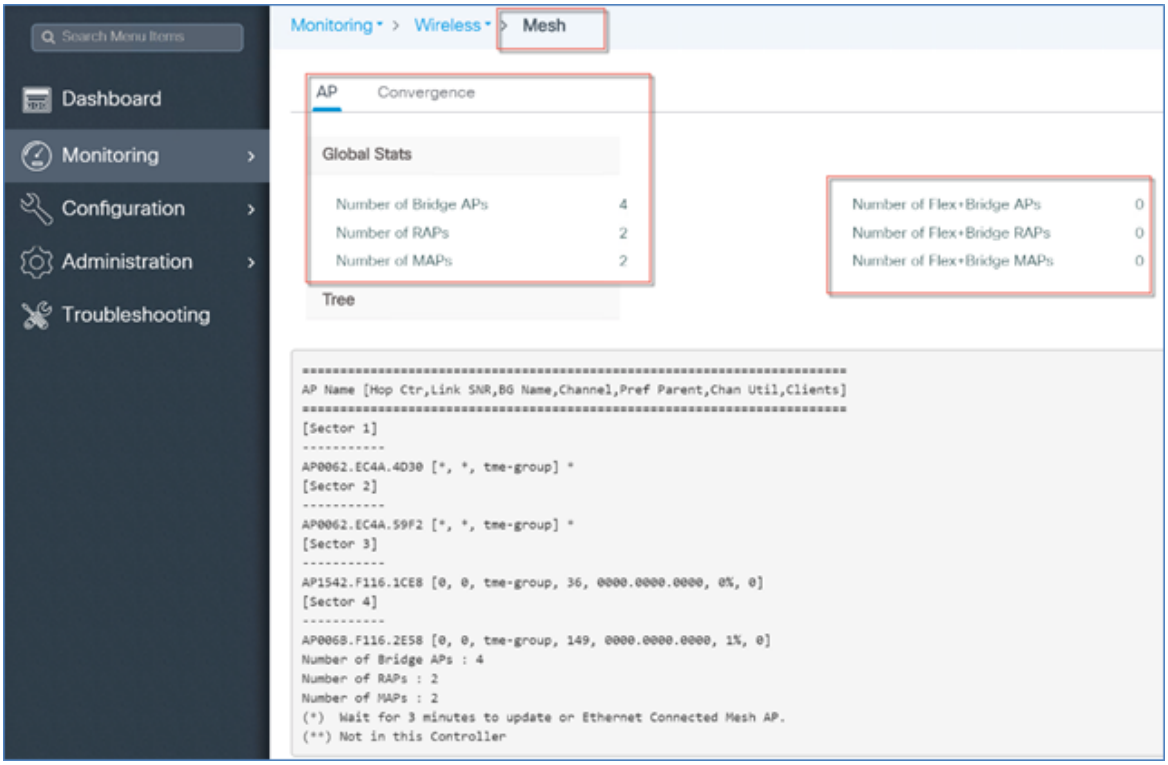


This concludes the Mesh AP configuration steps.

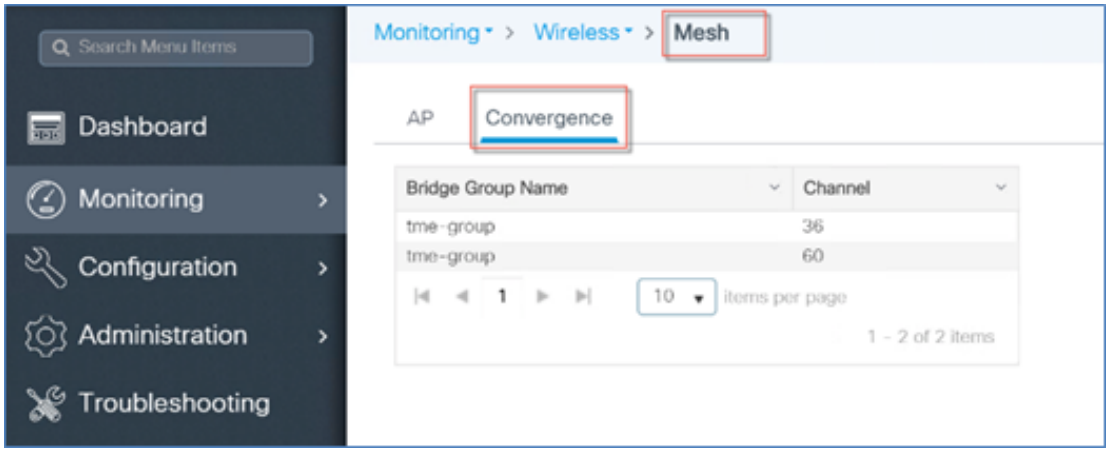
Monitoring Mesh Access Points to the Mesh Network

In IOS-XE 17.1 there are Mesh Network monitoring option available under the Monitoring tab.

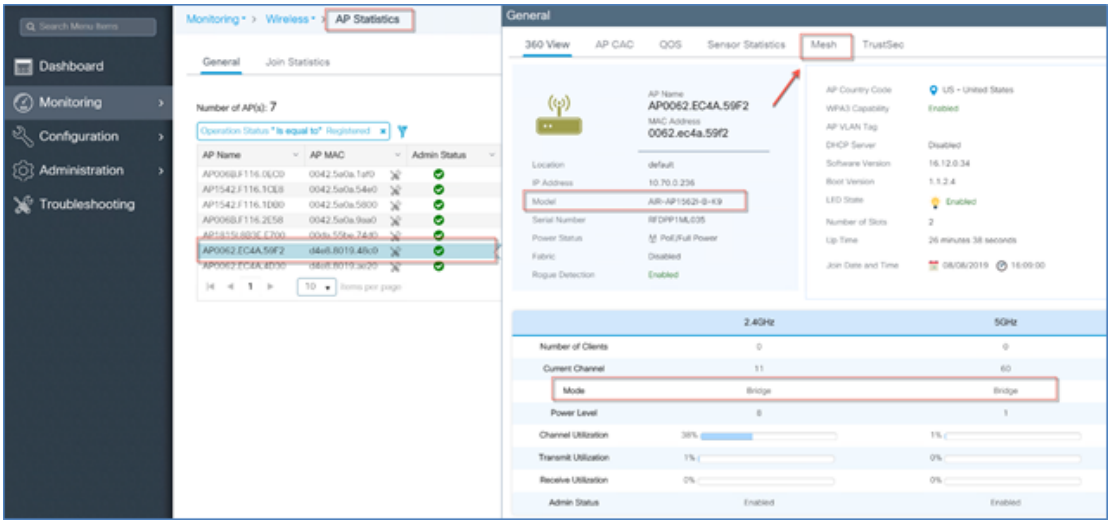
Under the Monitoring Tab choose Mesh >AP and as shown in the example below, global stats and the tree are presented.



The second option under Monitoring Mesh is the Convergence as shown in the example below

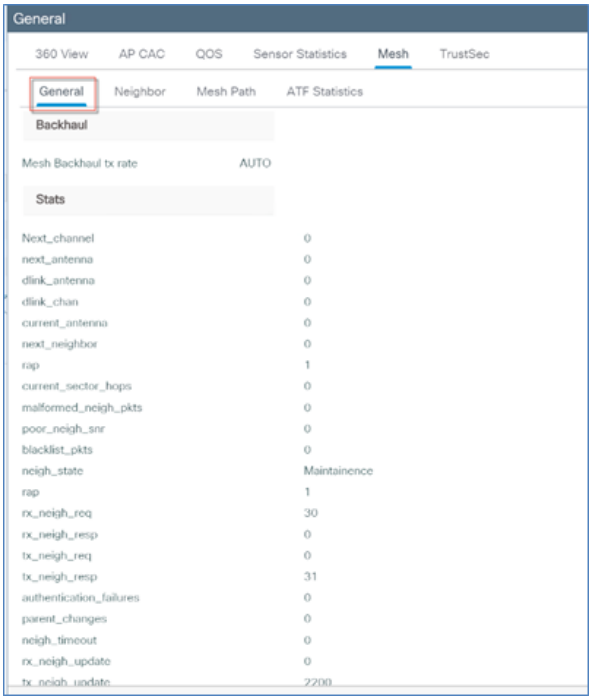


More Mesh details can be found about the Mesh AP by clicking on the Monitoring > Wireless > AP Statistics, and then clicking on the AP name as shown in the example below.



As shown in the screenshot above, there is also a new Mesh tab available if the AP is configured in a bridge mode.

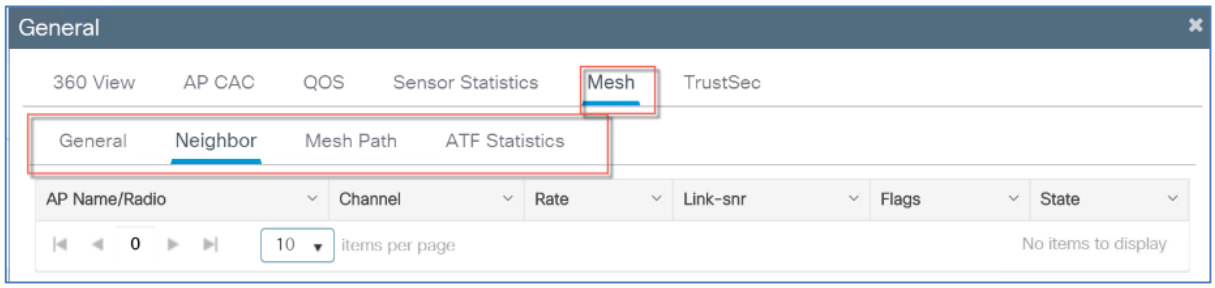
As shown in the screen shot below you can find General Details about backhaul and other Stats.



There is other Monitoring information available under the Mesh Tab

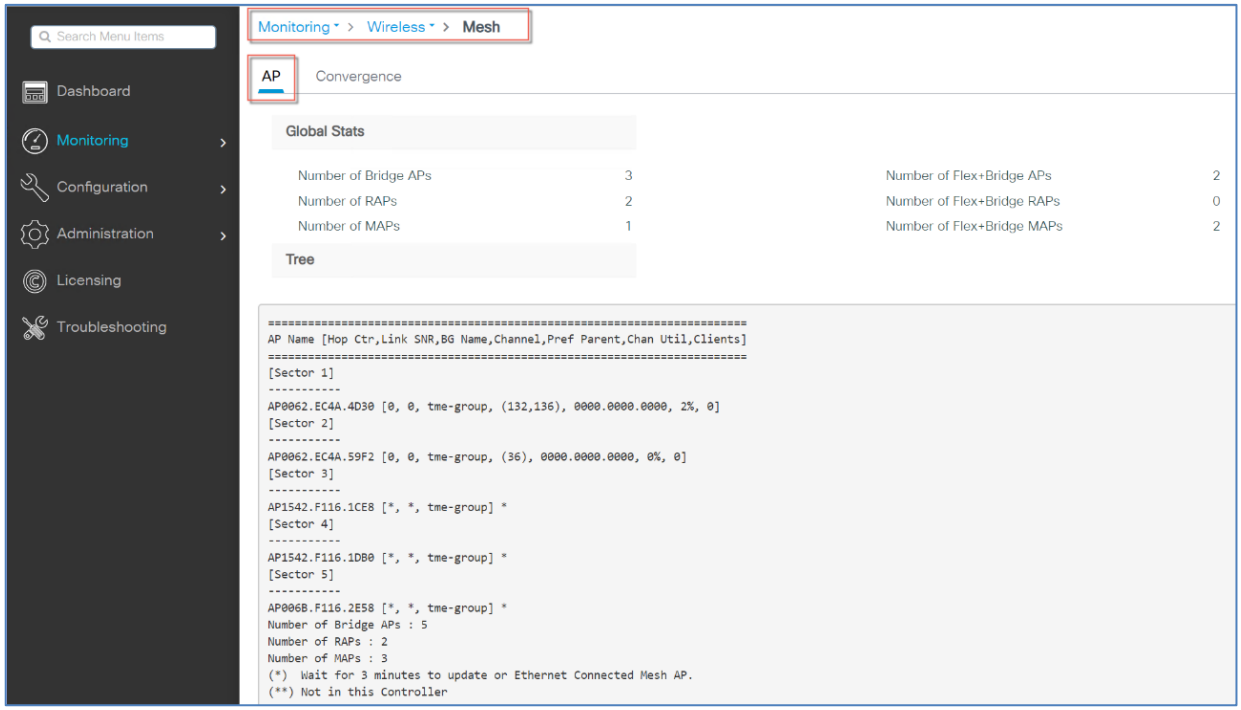


Another Mesh Monitoring/Configuration option is available under Config > Wireless > Access Points and then click on the AP in a Bridge mode as shown in the example below.

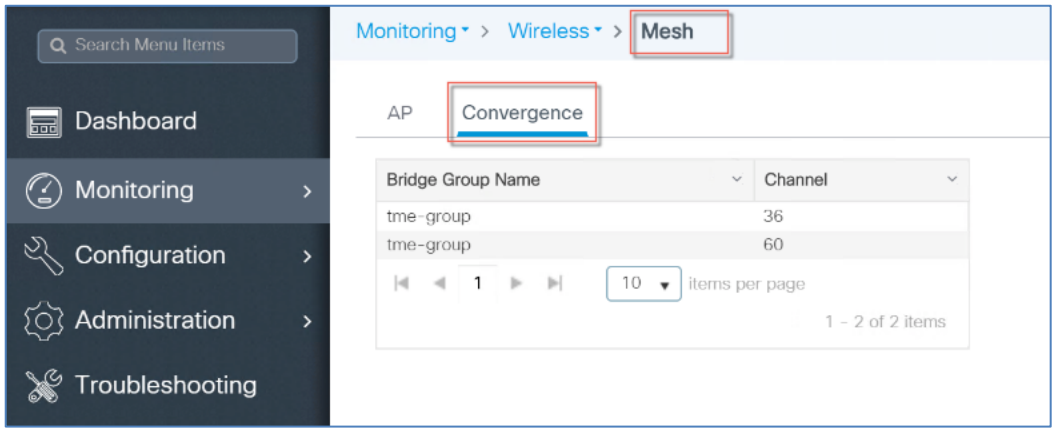


In IOS-XE 17.1 there are Mesh Network monitoring option available under the Monitoring tab.

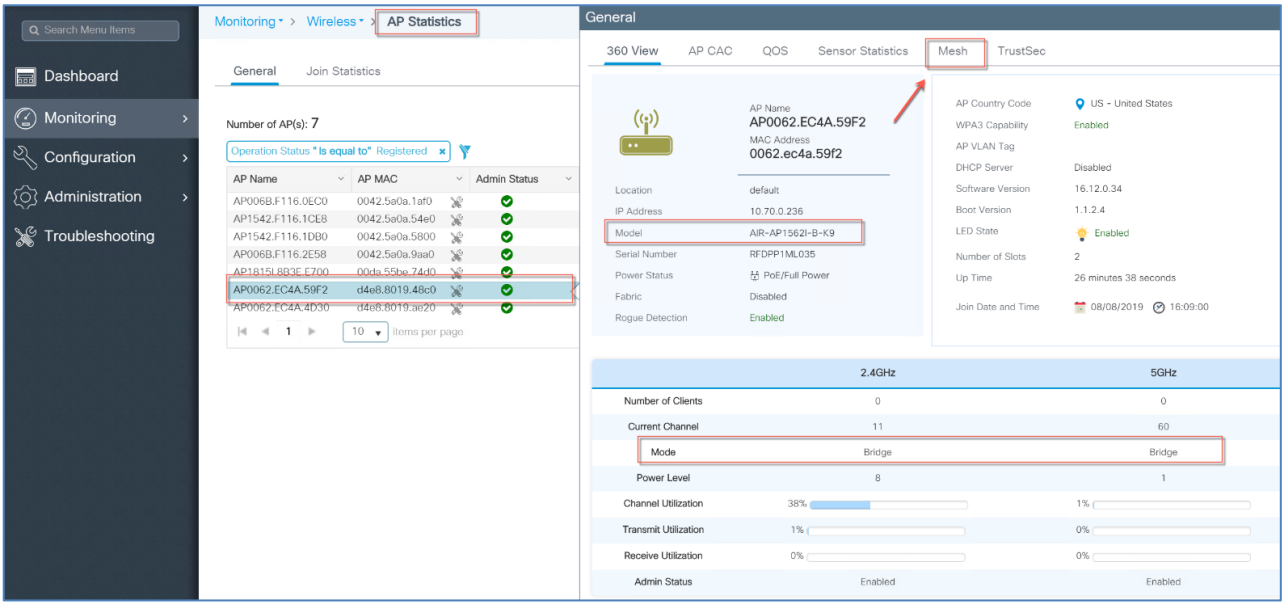
Under the Monitoring Tab choose Mesh >AP and as shown in the example below, global stats and the tree are presented.



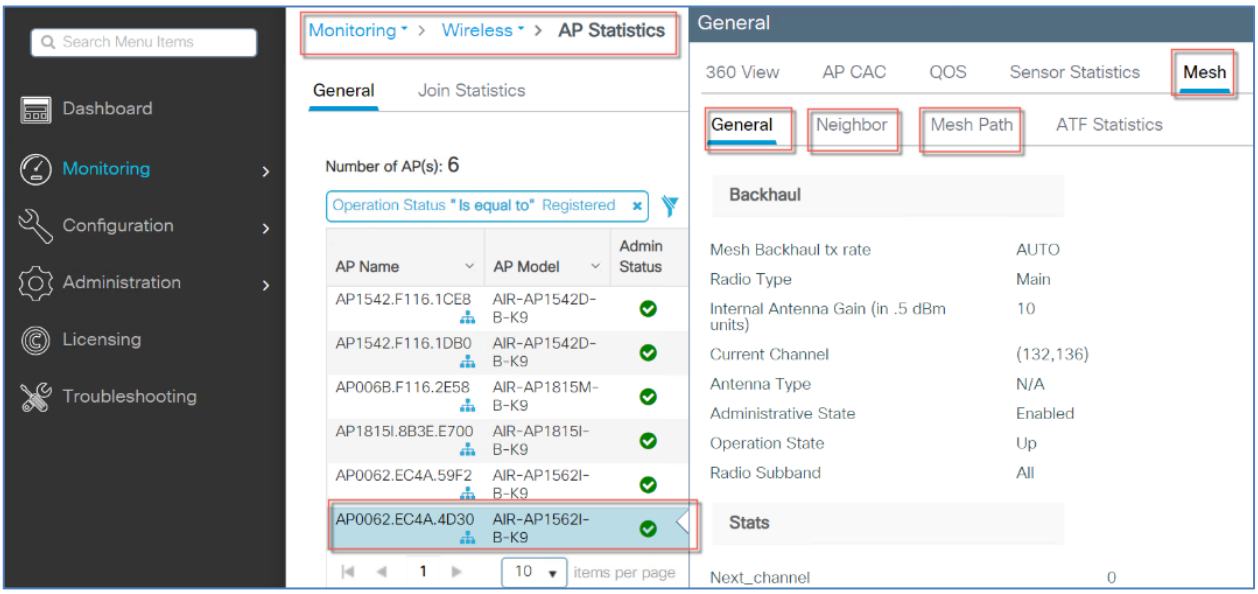
The second option under Monitoring Mesh is the Convergence as shown in the example below

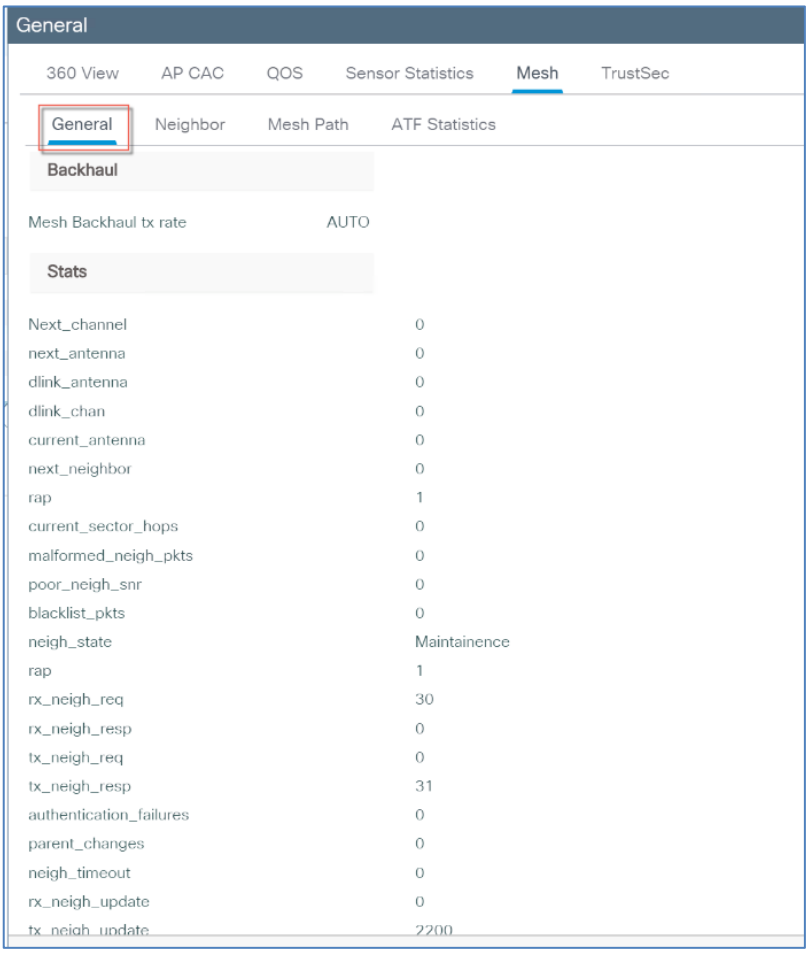


More Mesh details can be found about the Mesh AP by clicking on the Monitoring > Wireless > AP Statistics, and then clicking on the AP name as shown in the example below.

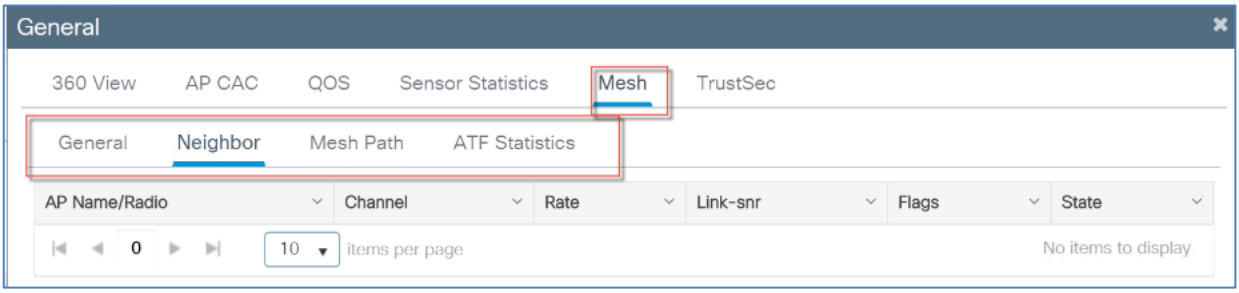


As shown in the screen shot below you can find General Details about Backhaul and other Stats.





There is other Monitoring information available under the Mesh Tab



Command Line Configuration for Mesh Access Points

The following Commands can be used to configure the Mesh APs and Mesh network.

```
wireless profile mesh default-mesh-profile
wire profile flex default-flex-profile
wireless profile policy default-policy-profile
wireless tag site default-site-tag
wireless tag policy default-policy-tag
wlan <wlan name> <no.> <SSID>
```

The Mesh related Show commands

```
sh wire profile mesh detailed default-mesh-profile
sh wire profile flex detailed default-flex-profile
sh wire profile policy detailed default-policy-profile
show wire tag site detailed default-site-tag
sh wireless tag policy detailed default-policy-tag
sh run wlan local-mesh
sh wlan name local-mesh
sh ap profile name default-ap-profile detailed
sh wire mesh ap tree
show ap tag summary
sh run | sec ap <ap mac in aaaa.bbbb.cccc format>
show ap name <ap name> tag info
sh ap dot11 5ghz summary
show wire profile policy summary
show wire profile mesh summary
 show wire tag site summary
show wire tag policy summary
show wire profile mesh summary
```

For more details on the controller configuration please see:

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

https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/16-12/config-guide/b_wl_16_12_cg.html

Cisco Catalyst 9800 Series Wireless Controller Command Reference, Cisco IOS XE Gibraltar 17.1.x

https://www.cisco.com/c/en/us/td/docs/wireless/controller/9800/16-12/cmd-ref/b_wl_16_12_cr.html

Example of the Mesh configuration

```
C9800-MA21#sh run | inc mesh
aaa local authentication tme-mesh authorization default
aaa authentication dot1x tme-mesh local
aaa authorization credential-download tme-mesh local
wireless mesh backhaul bdomain-channels
wireless mesh backhaul rrm
wireless mesh ethernet-bridging allow-bdpu
wireless mesh security psk provisioning
wireless mesh security psk provisioning default_psk
wireless mesh subset-channel-sync
wireless mesh security psk provisioning key 1 0 Cisco123 Cisco123
wireless mesh security psk provisioning key 2 0 cisco cisco
wireless profile mesh tme-mesh-ma
 description "TME mesh test profile"
 method authentication tme-mesh
 method authorization tme-mesh
wireless profile mesh default-mesh-profile
 description "default mesh profile"
 description TME-Lab-21_tme-mesh
 wlan tme-mesh policy TME-Lab-21_WLANID_1
wlan tme-mesh 1 tme-mesh
ap auth-list method-list tme-mesh
 mesh-profile tme-mesh-ma
ap profile mesh-ap-profile
 mesh-profile tme-mesh-ma
```

```
C9800-MA21#sh run | inc mesh
aaa local authentication tme-mesh authorization default
aaa authentication dot1x tme-mesh local
aaa authorization credential-download tme-mesh local
wireless mesh backhaul bdomain-channels
wireless mesh backhaul rrm
wireless mesh ethernet-bridging allow-bdpu
wireless mesh security psk provisioning
wireless mesh security psk provisioning default_psk
wireless mesh subset-channel-sync
wireless mesh security psk provisioning key 1 0 Cisco123 Cisco123
wireless mesh security psk provisioning key 2 0 cisco cisco
wireless profile mesh tme-mesh-ma
 description "TME mesh test profile"
 method authentication tme-mesh
 method authorization tme-mesh
wireless profile mesh default-mesh-profile
 description "default mesh profile"
 description TME-Lab-21_tme-mesh
 wlan tme-mesh policy TME-Lab-21_WLANID_1
wlan tme-mesh 1 tme-mesh
ap auth-list method-list tme-mesh
 mesh-profile tme-mesh-ma
ap profile mesh-ap-profile
 mesh-profile tme-mesh-ma
```

Air Time Fairness (ATF) in Mesh Deployments

This section describes ATF (Air Time Fairness) feature in Mesh Deployments, and provides general guidelines for its deployment.

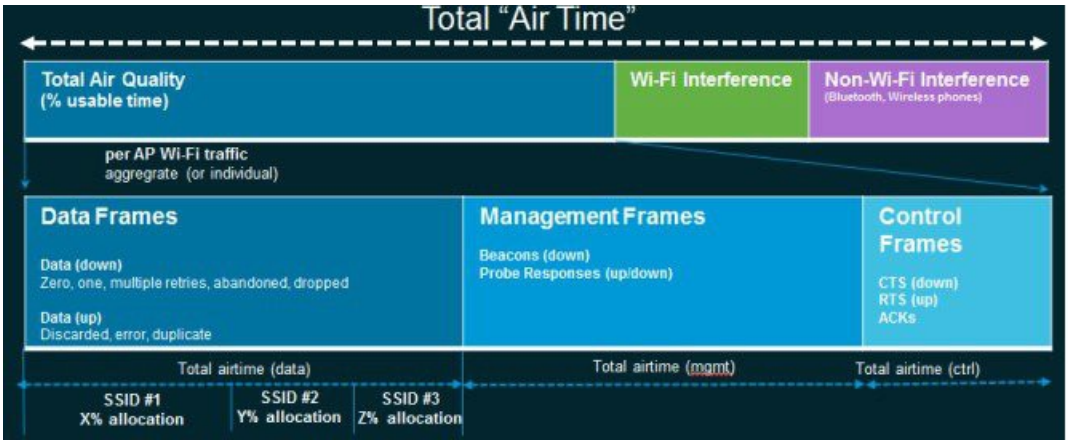
- Provides an general overview of ATF feature, and its deployment within the Cisco SDA Architecture with Cisco Catalyst C9800 controllers.
- Highlight key Service Provider features

Introduction to Air Time Fairness (ATF)

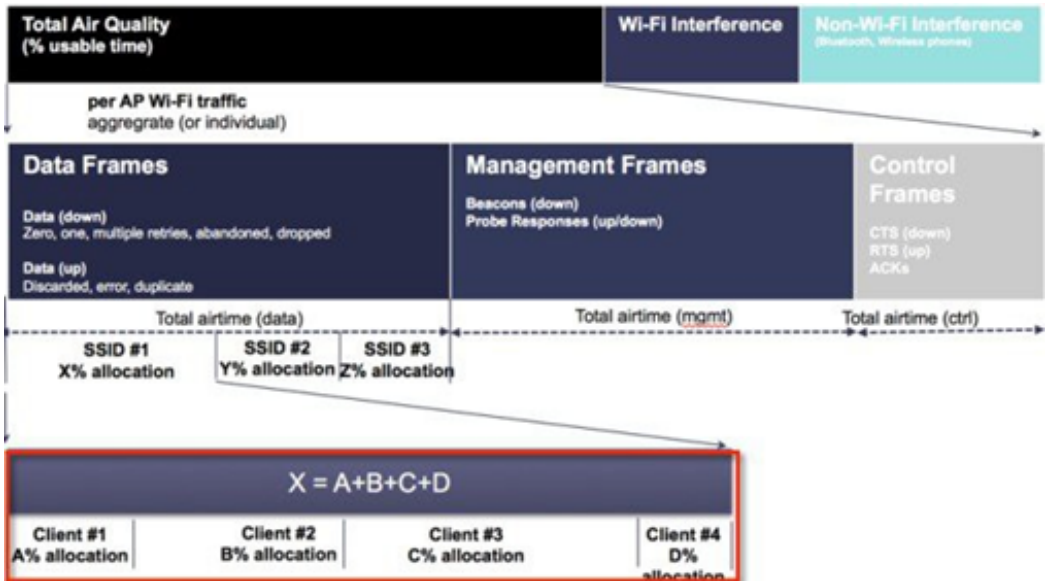
Traditional (wired) implementations of QOS regulate egress bandwidth. With wireless networking, the transmission medium is via radio waves that transmit data at varying rates. Instead of regulating egress bandwidth, it makes more sense to regulate the amount of airtime needed to transmit frames. Air Time Fairness (ATF) is a form of wireless QOS that regulates downlink airtime (as opposed to egress bandwidth). Large scale, high density Wi-Fi deployments are driving this feature. Wireless Network owners are mandating that their applications be allocated some fixed percentage of the total bandwidth of the Wi-Fi network. At the same time, with capital sharing being considered with multiple cellular providers, ATF is needed to ensure fairness of usage across operators.

Before a frame is transmitted, the ATF budget for that SSID is checked to ensure that there is sufficient airtime budget to transmit the frame. Each SSID can be thought of as having a token bucket (1 token = 1 microsecond of airtime). If the token bucket contains enough airtime to transmit the frame, it is transmitted over the air. Otherwise, the frame can either be dropped or deferred. While the concept of dropping a frame is obvious, deferring a frame deserves further explanation. Deferring a frame means that the frame is not admitted into the Access Category Queue (ACQ). Instead, it remains in the Client Priority Queue (CPQ) and may be transmitted at a later time when the corresponding token bucket contains a sufficient number of tokens (unless the CPQ reaches capacity, at which point the frame will be dropped regardless). The majority of the work involved for ATF takes place on the access points. The wireless controller is used simply to configure the feature and display results.

An ATF Policy will be created with weight and client-sharing information. The change of configuration, Database populations will be performed. Client fair share ensures the clients within a SSID/WLAN are treated equally based on their utilization of the radio bandwidth.



ATF within Mesh networks also insures that all backhaul nodes are treaded as equally as possible and there is no bandwidth starvation in a direction of mesh node or clients.



ATF policy has an option to turn on or off client fair sharing among clients associated to a policy. This option can be executed while creating, modifying the ATF policy in the C9800 Controller. Customer can use this option or feature to provide fair sharing of Airtime between clients associated to a WLAN and also configure a weight of each ATF policy.

Cisco Air Time Fairness (ATF) Use Cases

Public Hotspots (Stadium/Airport/Convention Center/Other)

In this instance a public network is sharing a WLAN between two (or more) service providers and the venue. Subscribers to each service provider can be grouped and each group can be allocated a certain percentage of airtime.

Education

In this instance, a university is sharing a WLAN between students, faculty, and guests. The guest network can be further partitioned by service provider. Each WLAN can be assigned a certain percentage of airtime.

Enterprise or Hospitality or Retail

In this instance, the venue is sharing a WLAN between employees and guests. The guest network can be further partitioned by service provider. The guests could be sub-grouped by tier of service type with each WLAN subgroup being assigned a certain percentage of airtime, for example a paid group is entitled to more airtime than the free group.

Time Shared Managed Hotspot

In this instance, the business entity managing the hotspot, such as a service provider or an enterprise, can allocate and subsequently lease airtime to other business entities.

ATF Functionality and Capabilities

- ATF can be globally enabled or disable on the C9800 controllers
- ATF policies are applied only in the downlink direction (AP transmitting frames to client). Only airtime in the downlink direction, that is AP to client, can be controlled accurately by the AP. Although the AP can constrain airtime for packets that it sends to clients, the AP can only measure airtime for packets that it 'hears' from clients because it cannot strictly limit their airtime.
- ATF policies are applied only on wireless data frames; management and control frames get ignored.
- When ATF is configured per-SSID, each SSID is granted airtime according to the configured ATF policy mapped to the Policy Profile.
- Enabling/Disabling of ATF per radio, being associated with the RF Profile of the C9800, provides the flexibility of configuring different ATF policies for different sites and different radios. Maximum of 512 ATF policies can be created.
- If single ATF Policy mapped to the C9800 Controller Policy Profile, a situation can arise where different policies for two radios cannot be configured in same WLAN. To overcome this limitation two ATF policies can be configured in the Policy Profile. One can choose two different polices for different radios or same ATF Policy for both.
- ATF can be configured to either drop or defer frames that exceed their airtime policies. If the frame is deferred, it will be buffered and transmitted at some point in the future when the offending SSID has a sufficient airtime budget. Of course, there is a limit as to how many frames can be buffered. If this limit is crossed, frames will be dropped regardless.

- ATF is supported in IOS-XE release 17.1.1 on the on the 2800, 3800, 4800 and 1560 APs in release 17.1 in Flex and Local mode. In mesh mode only 1560 will support ATF.
- The C9100 series AP will support ATF on the IOS-XE controllers in later releases.

Air Time Fairness in Mesh Deployments IOS-XE 17.1

This section of the document introduces the ATF on Mesh APs and provides guidelines for its deployment.

The purpose of this section is to:

- Provide an overview of ATF on Mesh APs
- Highlight supported Key Features
- Provide details on deploying and managing the ATF on Mesh APs

Pre-requisite and Supported Features in IOS-XE 17.1.1

Mesh ATF is supported on C9800 rel 17.1.1 or higher. Mesh ATF is supported only on the1560 AP.

ATF on Mesh Feature Overview

At the present time, enterprise class, high density stadium and other major Wi-Fi deployments with Cisco Wave-1 and Wave-2 Indoor APs are benefited by "per SSID" based Airtime Fairness and "per Client within a SSID" based Airtime Fairness through IOS-XE 17.1 Release.

In a same way, currently, there is a demand from the Customers with large scale Outdoor wireless mesh deployments to serve their users by providing fairness among the Wi-Fi users across the Outdoor wireless mesh network in utilizing the AP radio Airtime downstream and also provide administrators the key control to enforce SLA (implied on multiple cellular operator through Wi-Fi hotspot) on the Wi-Fi users across the Outdoor wireless mesh network. However, since all Wi-Fi users traffic is bridged between MAPs and RAPs through the wireless backhaul radio and there is no SSID concept on wireless backhaul radio for backhaul nodes to enforce policies through SSID's for each backhaul node, there is no easy solution for Wi-Fi users across the Outdoor wireless mesh network to get treated fairly in terms of utilizing the Wi-Fi airtime through their Outdoor Wireless Mesh Aps. As far as the clients on client access radios are concerned, it's fairly simple to regulate the airtime fairness through SSIDs (w/ or w/o client fair sharing) in a similar way how it is done for Cisco Local and Flex mode APs.

Before the solution overview of supporting ATF on mesh, lets quickly recap ATF - Airtime Fairness (ATF) is basically a concept which provides an ability to regulate/enforce the AP radio airtime in downstream direction for the clients associated through the SSID's. As a result, the Wi-Fi users on wireless network are fairly treated in terms of utilizing the radio WiFi radio airtime. This basically provides the key control either to enforce SLA additionally or simply to avoid certain group or individual from occupying an unfair amount of WiFi airtime on a particular or on a given AP radio. A service level agreement (SLA) is a contract between a service provider (either internal or external) and the end user that defines the level of service expected from the service provider.

SLAs are output-based in that their purpose is specifically to define what the customer will receive.

In general, in the Mesh architecture, the Mesh Aps (Parents, child MAPs) in a Mesh Tree will be accessing the same channel (let's forget about extended sub-backhaul radios for a minute) on backhaul radio for mesh connectivity between Parents and child Maps. Whereas, the Root AP will be connected wired to the controller and MAPs will be connected wireless to the controller. Hence all the CAPWAP, Wi-Fi traffic will be bridged to the controller through the wireless backhaul radio and through RAP. In terms of the physical locations, normally the RAPs will be placed at roof top and the MAPs in multiple hops will be placed some distance apart within each other based on the Mesh network segmentation guidelines. Hence each MAP in a Mesh tree can provide 100% of their own radio airtime downstream to their users though each MAP accessing the same medium. To compare this in non-mesh scenario, where there can be neighboring Local mode APs in the area next to each other in different rooms serving their respective clients on the same channel with each providing 100% radio airtime downstream. Therefore, ATF has no control over enforcing clients in two different neighboring AP's accessing

the same medium. Similarly, it's applicable for MAPs in a Mesh tree.

For Outdoor/Indoor Mesh Aps, Airtime fairness must be supported on client access radios which serve regular clients as same as how we currently support ATF on non-mesh Local mode APs to serve the clients and additionally it must also be supported on backhaul radios which bridge the traffic to/from the clients on client access radios to RAPs (one hop) or through MAPs to RAPs (multiple hops).

Its bit tricky to support ATF on backhaul radio's using the same SSID/Policy/Weight/Client fair sharing model. Since backhaul radios don't have SSIDs and it always bridges traffic through their hidden backhaul nodes. Henceforth, on the backhaul radios either in RAP or MAP, the radio airtime downstream will be fair shared equally based on the number of backhaul nodes. This approach eliminates the problem and provides fairness to users across wireless mesh network in the case where the clients associated to 2nd hop MAP can stall the clients associated to 1st hop MAP where 2nd hop MAP is connected wireless to 1st hop MAP through backhaul radio though the Wi-Fi users in the MAPs are separated by a physical location. In the scenario, when a backhaul radio has an option to serve normal clients through universal client access feature, ATF considers the regular clients into single node and group them into it. It enforces the Airtime by equally fair sharing the radio airtime downstream based on the number of nodes (backhaul nodes + single node for regular clients). We will see more details how this solution is turned into design in the next sections.

Mesh ATF Modes of Operation

To understand the airtime-allocation field present under RF-Profile WebUI page is specific to MESH APs and to understand this we need to be aware of the modes in which ATF on mesh can operate.

Airtime Fairness on mesh can be divided into three parts:

1. ATF on client access radio of a Mesh AP (RAP or MAP).

ATF runs normally for these radios based on profile configurations and Policy Mapping of ATF Profiles.

These Radios are basically treated as regular radios on local/flex mode APs.

2. ATF on backhaul radio of a Mesh AP with universal access (client access) disabled.

For backhaul only radios, ATF divides the available bandwidth equally among backhaul nodes where backhaul node is defined as any radio backhaul link (uplink or downlink).

For example as illustrated in the diagram below: If a RAP (R1) has 3 children MAPs, it has 3 backhaul nodes.

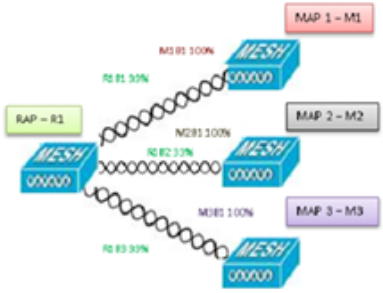
If a MAP (M1) has 2 children MAPs, it has 3 backhaul nodes (2 downlink and 1 uplink).

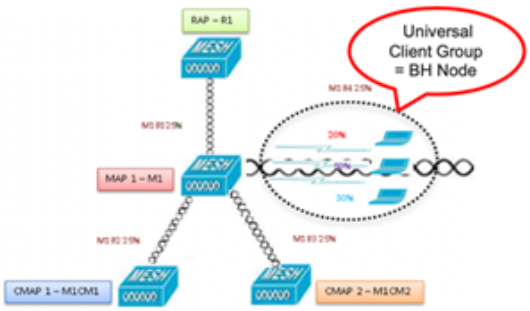
3. ATF on backhaul radio of a Mesh AP with universal access (client access) enabled.

For a backhaul radio with client access enabled (also known as universal access), ATF counts client access as one additional node.

Therefore, in the example as illustrated below, a MAP (M1) with 3 children MAPs and universal access enabled is counted as 4 nodes and airtime is divided equally among these.

Additionally, capability is provided to change the client access allocation on such a radio through CLI or WebUI. This provision is made available via configurable airtime allocation after enabling bridge client access.



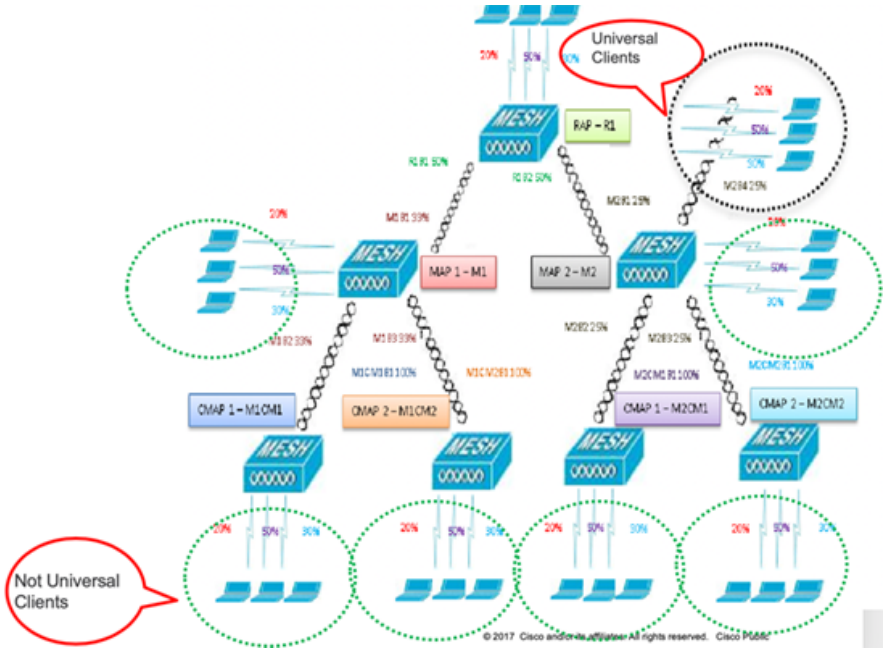


A bigger mesh design with ATF will looks as illustrated below. In this Mesh design ATF is applied on the Backhaul nodes and also on the Access Radios.

Please note again on the Client Access Radios (as shown below in the green dotted circles) ATF runs normally for these radios based on profile configurations and Policy Mapping of ATF Profiles. These Radios are basically treated as regular radios on local/flex mode APs.

For a backhaul radio with Client Access enabled also known as Universal Access, ATF counts client access as one additional node. Additionally, capability is provided to change the Universal client access allocation on such a radio.

For example, if Bridge Client Access is enabled and airtime allocation is set as 30, then 30% of the available airtime will be assigned to the Universal clients and the remaining 70% of the available airtime will be shared between the remaining nodes.



Configuring ATF on Mesh

To configure, ATF on mesh, perform the following steps.

First, ATF has to be enabled Globally. To enable the ATF globally in the Configuration > ATF Global Configuration tab as shown in the example below. In this option the ATF Profile has to be configured for both 5GHz and 2.4 GHz individually.

Enabling/Disabling of the ATF per Radio being associated with the RF Profile of C9800 controller, provides the flexibility of configuring different ATF policies for different RF sites or groups for 5 and 2.4GHz radios. Optimization is effective when the current WLAN reaches the air time limit and the other available WLANs do not use air time to its full extent.

Also enable Bridge Client Access for 5GHz and /or 2.4GHz, as noted above, enabling Globally and configuring ATF profiles enable ATF for Client Access Radios.

Enabling Bridge Client Access enables the airtime allocation for the Universal Client Access.

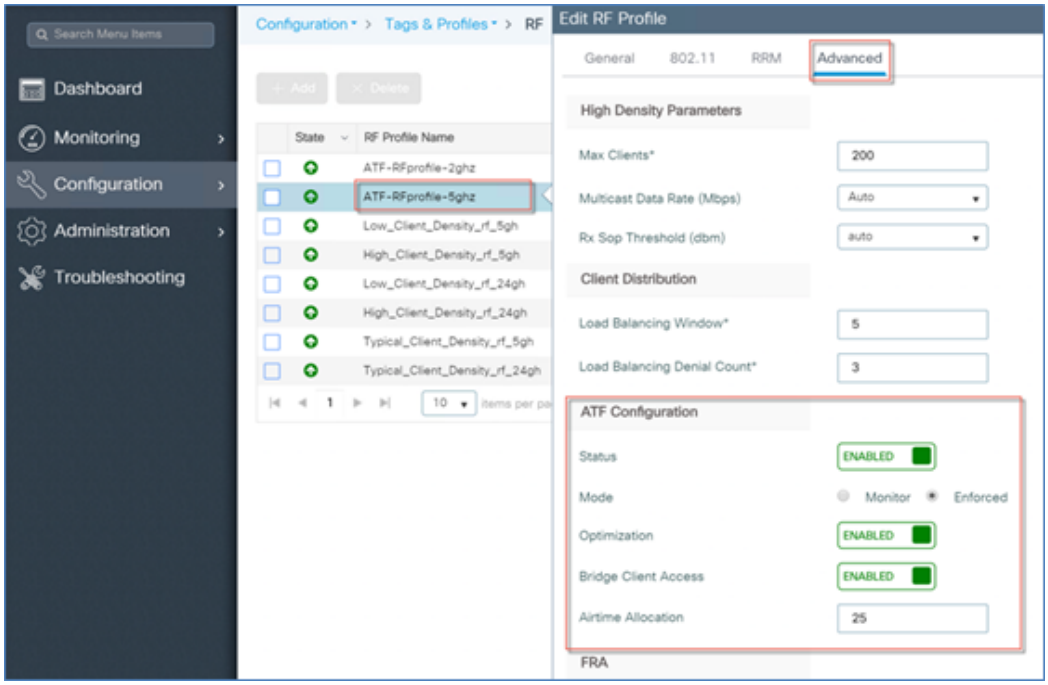


Next, enable the ATF at the Radio level to basically apply the ATF mode at the RF group level based on the RF profile.

From the Configuration > RF Add > Advanced > ATF configuration. Create an ATF RF profile for 5GHz and one for 2.4GHz access radios.

As shown in the example below 5GHz and 2.4GHz radio is configured in the Enforced Mode and Bridge Client Access is also enabled and Airtime Allocation is set to 25.

Note: Bridge client access should be enabled for ATF configuration of APs Mesh/Bridge Mode. This capability is provided for enabling Universal Client access allocation on the Radio in Bridge or Mesh APs



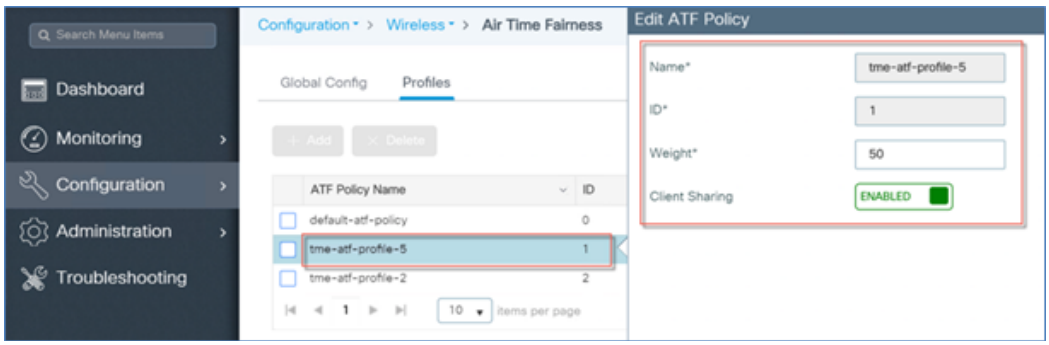
If we attach single ATF Policy with the Controller Policy Profile, we will end up with the situation where we cannot configure different policies for two radios in same WLAN. To overcome this limitation, we will have two ATF policies configured in the Policy Profile. One can choose two different polices for different radios or same ATF Policy as well.

As stated above, ATF policy has an option to turn on or off client fair sharing among clients associated to a policy. This option can be executed while creating, modifying the ATF policy in the C9800 Controller.

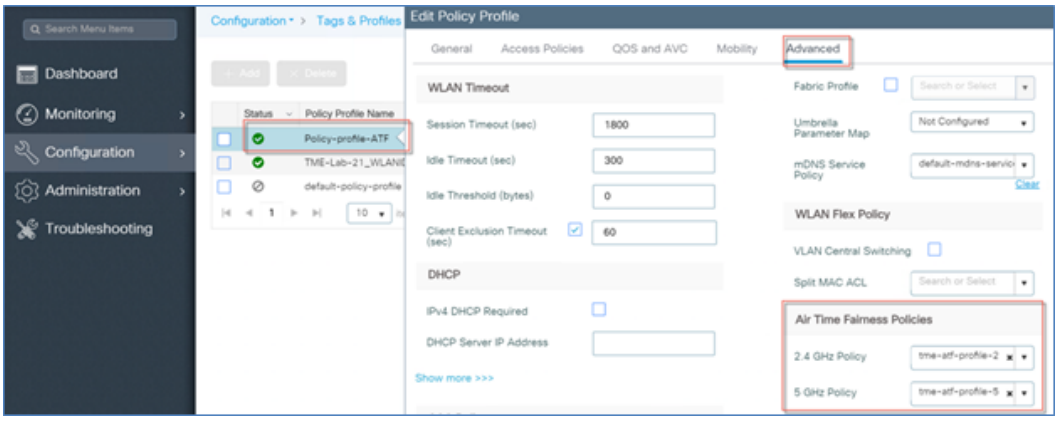
Wireless Network owners are mandating that their applications be allocated some fixed percentage of the total bandwidth of the Wi-Fi network to eliminate complete bandwidth starvation. Customer can use Weight indicator (5-100) to provide weighted sharing of Airtime assigned to each of the ATF policies. Basically, each ATF policy can be configured with its own Air Time weight and then applied to the WLANs.

Additionally, capability is provided to change the ATF profile air time weighted allocation (5-100) on a Client Access Radios

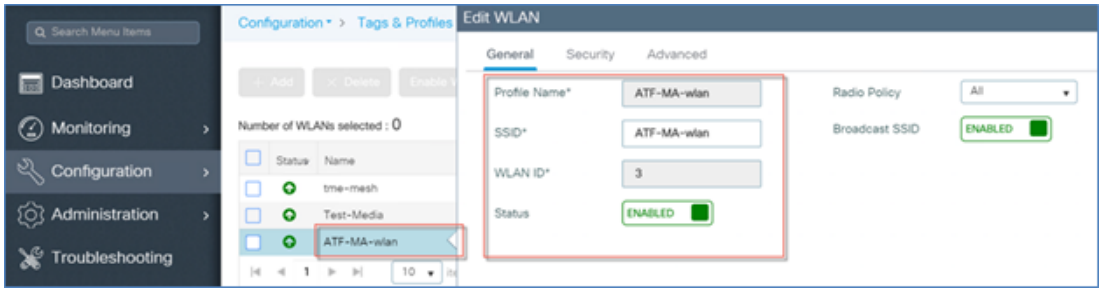
through CLI or WebUI configuration for Client sharing Optimization enabled or disabled. When Optimization is disabled Client Sharing of the Airtime will be strict, however when Optimization is enabled, clients can use airtime allocated to others but not being used, thus airtime allocation will not be strictly enforced or in other terms “shared” or “borrowed”.



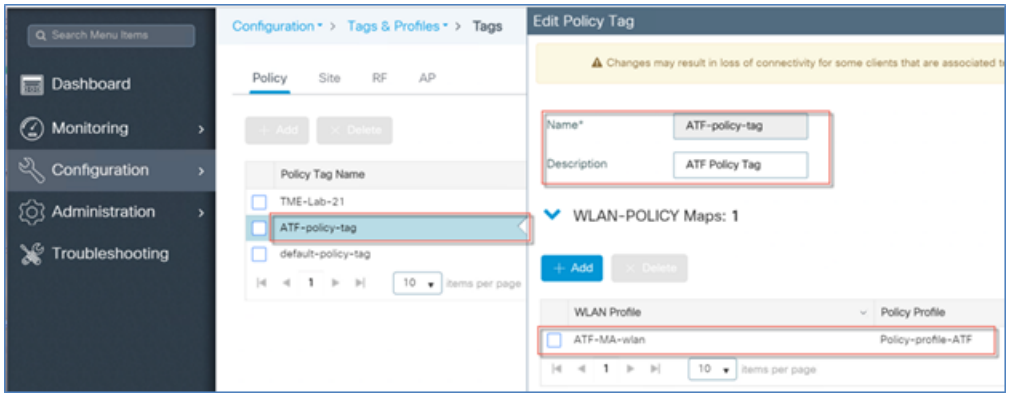
After ATF profile or profiles have been created, create a new Policy Profile or update an existing Policy profile with the ATF profiles configured in the previous step. Go to Policy profile Advanced tab and configure there the ATF profiles for 2.4 and 5GHz policies as shown in the example below. These ATF Policy Profiles will be mapped to the WLANs as shown in the example below.



When the ATF configuration is done, then Create and Enable the WLANs on which ATF will applied by the ATF Policy profile created in the above steps.



Next step is to configure ATF Policy Tag and map it to the WLAN created in the step above

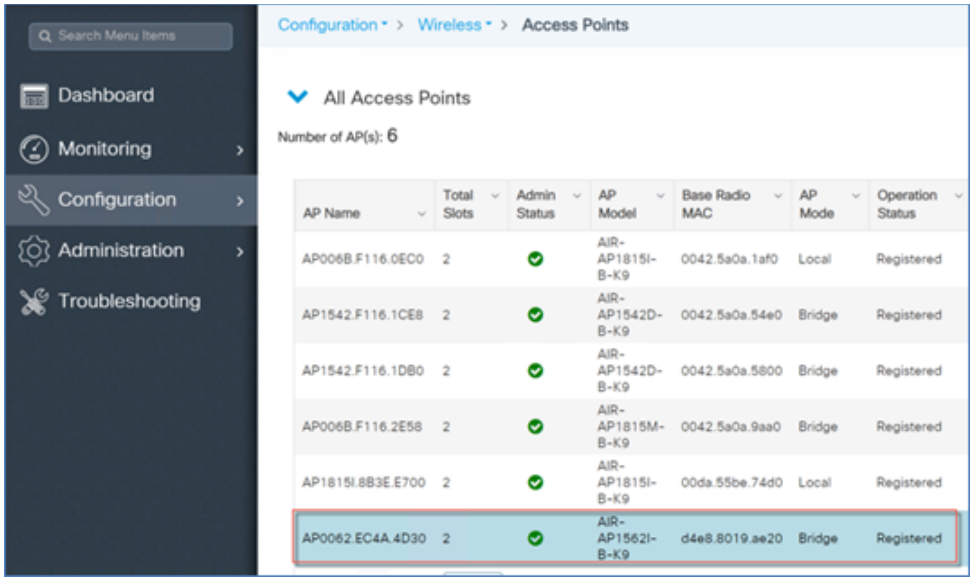


Next, create RF Tag and apply ATF RF Profiles created above to the Mesh and Local/Flex Mode Client Access 5 and 2.4 GHz Radios

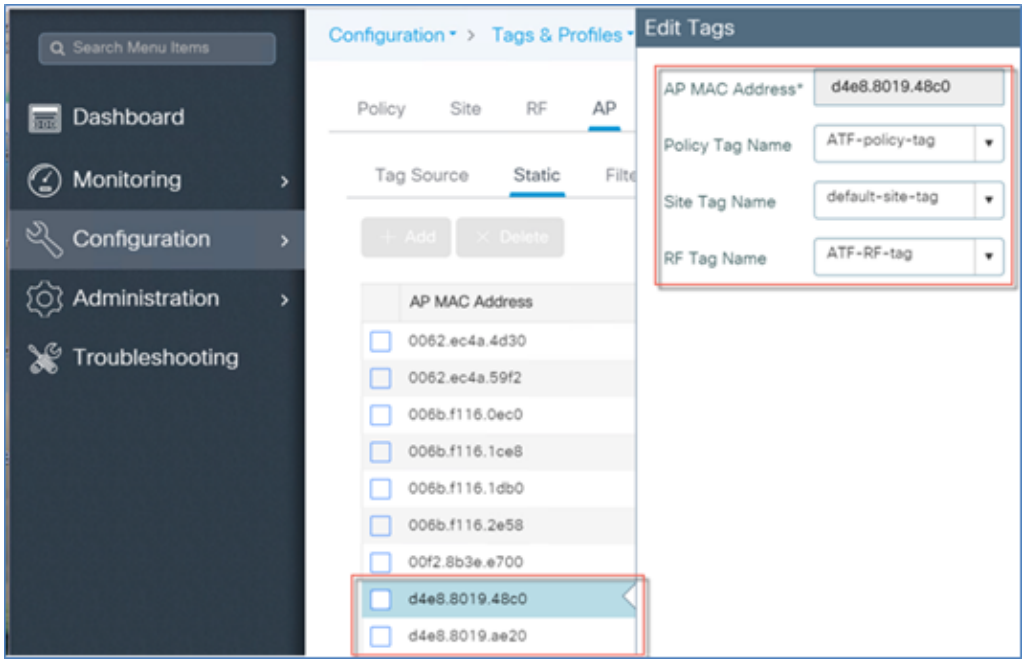


From the list of the available Mesh/Bridge APs select the MAC addresses of the Radios to which the ATF policies should be applied

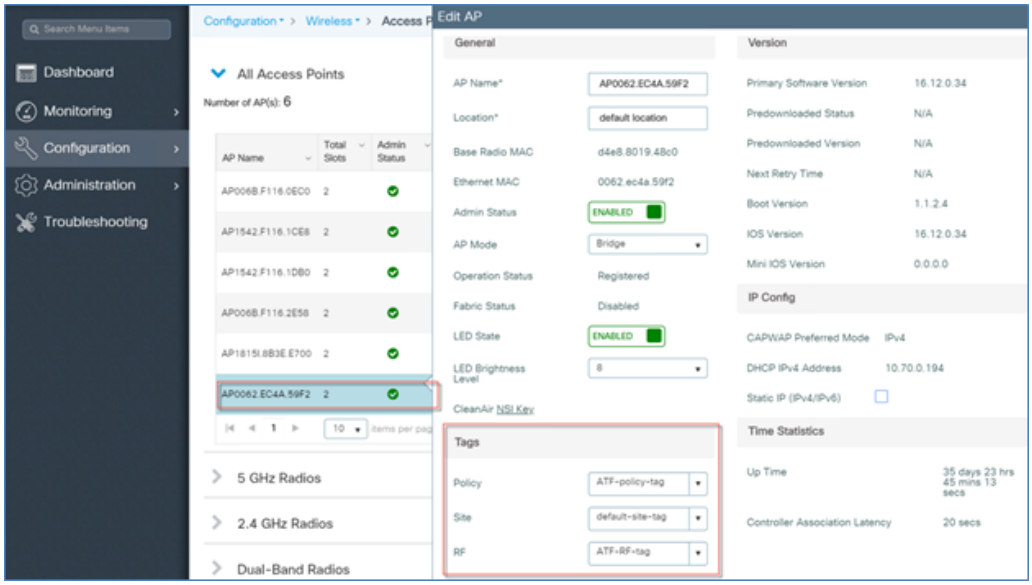
Note: In release 17.1 ATF is only supported on the 1560 series Mesh APs.



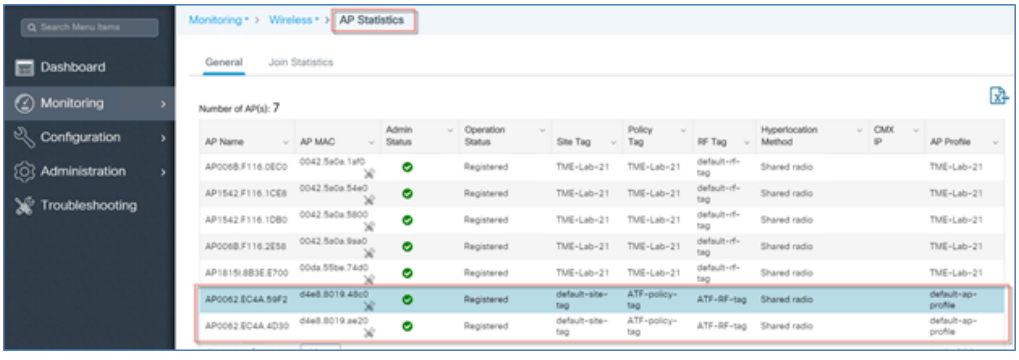
Configure the AP Tag and manually add MAC addresses of the 1560 series APs that the RF Policy Tag and ATF Tag will be mapped to. You can Statically add APs or you can create a Filter to add mac addresses with a * Character, ie. “d4e8.8019.*”



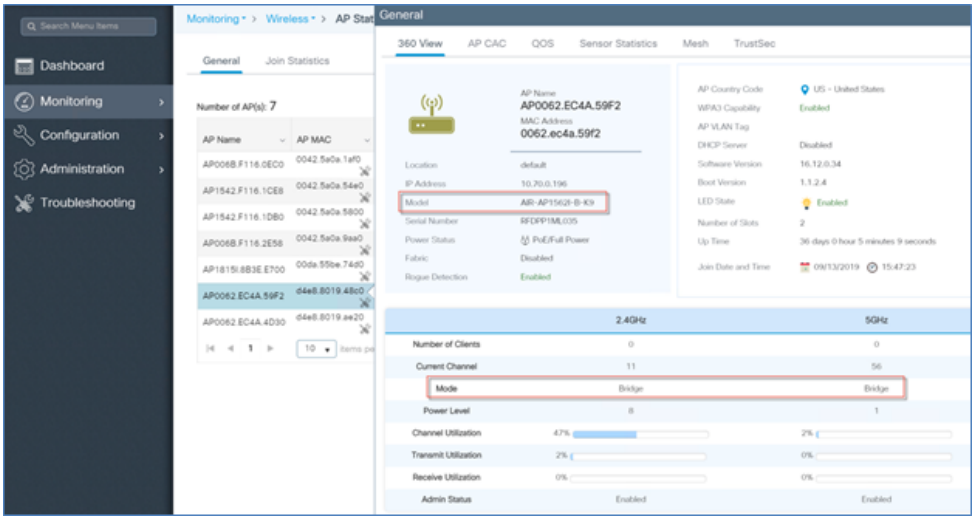
And lastly, apply the Policy and TF Tags to the selected Mesh APs as shown in the example below.



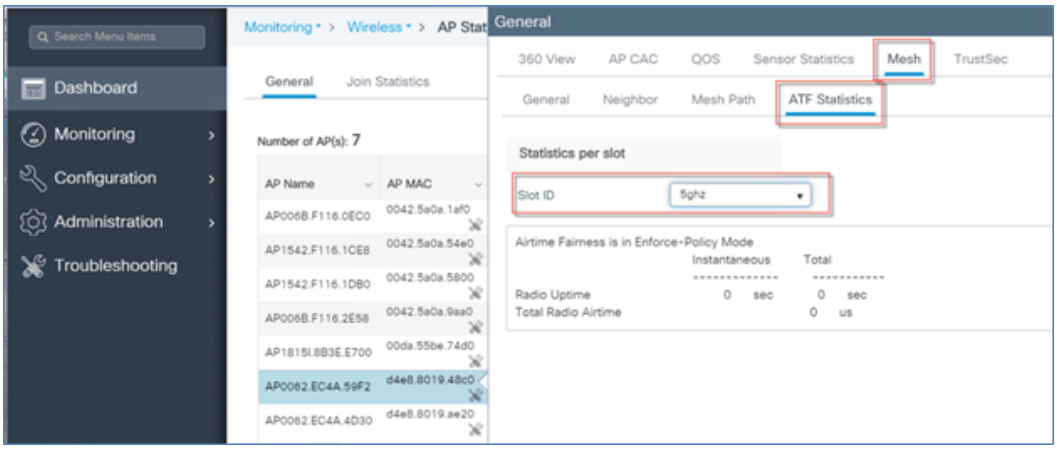
Once everything is mapped and tagged you can see all the mappings in the Monitor mode of the earlier selected AP, as shown below.



By further selecting Monitor > AP Statistics > General and clicking on one of the Mesh APs , additional details will be shown as illustrated below.



When choosing a Mesh tab additional Mesh ATF statistic can be seen, as shown in the example below.



User can run a speed test to verify the ATF by configuring two WLANs with different ATF policies. In the example we have configured two ATF policies, one with weight 90 and other with weight 10.

1. Connect a wireless client to SSID with ATF policy with weight 90 configured and observe the effect of the ATF on the WLAN by running Speedtest from the URL

2. http://www.speedtest.net

3. Connect the same wireless client to SSID with ATF policy configured as 10 and observed the effects of the ATF on that WLAN. You should see Speedtest performance on the download side is much slower. The test results might vary due to the air time availability, interference and so on.

Mesh/Bridge Mode ATF Configuration Sample

The sample of the Run Config file is shown below after the Mesh ATF WebUI configurations as shown in the examples above.

```
wireless mobility group name ma
wireless aaa policy default-aaa-policy
wireless cts-sxp profile default-sxp-profile
wireless management certificate ssc auth-token 0 cisco123
wireless management interface Vlan70
wireless mesh backhaul bdomain-channels
wireless mesh backhaul rrm
wireless mesh ethernet-bridging allow-bdpu
```

```
wireless mesh security psk provisioning
wireless mesh security psk provisioning default_psk
wireless mesh subset-channel-sync
wireless mesh security psk provisioning key 1 0 Cisco123 Cisco123
wireless mesh security psk provisioning key 2 0 cisco cisco
wireless profile airtime-fairness default-atf-policy 0
wireless profile airtime-fairness tme-atf-profile-5 1
client-sharing
weight 50
wireless profile airtime-fairness tme-atf-profile-2 2
client-sharing
weight 50
description "TME mesh test profile"
method authentication tme-mesh
method authorization tme-mesh
security psk
wireless profile mesh default-mesh-profile

description "default mesh profile"
wireless profile policy Policy-profile-ATF
description "ATF Policy Profile"
dot11 24ghz airtime-fairness tme-atf-profile-2
dot11 5ghz airtime-fairness tme-atf-profile-5
no shutdown

wireless tag policy ATF-policy-tag
description "ATF Policy Tag"
wlan ATF-MA-wlan policy Policy-profile-ATF
wireless tag policy default-policy-tag
description "default policy-tag"
wireless tag rf ATF-RF-tag
24ghz-rf-policy ATF-RFprofile-2ghz
5ghz-rf-policy ATF-RFprofile-5ghz
description "ATF RF Tag"

wlan ATF-MA-wlan 3 ATF-MA-wlan
no shutdown
ap dot11 24ghz rf-profile ATF-RFprofile-2ghz
airtime-fairness bridge-client-access airtime-allocation 25
airtime-fairness mode enforce-policy
airtime-fairness optimization
description "ATF profile for 2GHz"
no shutdown

ap dot11 24ghz airtime-fairness bridge-client-access airtime-allocation 30
ap dot11 24ghz airtime-fairness mode enforce-policy
ap dot11 24ghz airtime-fairness optimization
ap dot11 5ghz rf-profile ATF-RFprofile-5ghz
airtime-fairness bridge-client-access airtime-allocation 25
airtime-fairness mode enforce-policy
airtime-fairness optimization
description "ATF RF profile for 5GHz"
no shutdown
ap d4e8.8019.48c0
policy-tag ATF-policy-tag
rf-tag ATF-RF-tag
ap d4e8.8019.ae20
policy-tag ATF-policy-tag
rf-tag ATF-RF-tag
end
```

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.