



Cisco Wireless Mesh Networking

This chapter provides design and deployment guidelines for the deployment of secure enterprise, campus, and metropolitan Wi-Fi networks within the Cisco Wireless Mesh Networking solution, a component of the Cisco Unified Wireless Network solution.



Note

For more detailed information about Cisco Wireless Mesh Networking, including configuration and deployment, refer to the [Cisco Mesh Access Points, Design and Deployment Guide, Release 8.5](#).

Mesh networking employs Cisco Aironet 1500 Series outdoor mesh access points (APs) and indoor mesh APs along with the Cisco Wireless LAN Controller (WLC), and Cisco Prime Infrastructure to provide scalable, central management and mobility between indoor and outdoor deployments. The Control and Provisioning of Wireless Access Points (CAPWAP) protocol manages the connection of the mesh APs 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 chapter are for the following products:

- Cisco Aironet 1570 (1572) series outdoor mesh access points
- Cisco Aironet 1560 (1562) series outdoor mesh access points
- Cisco Aironet 1540 (1542) Series outdoor mesh access points
- Cisco Aironet 1550 (1552) series outdoor mesh access points
- Cisco Aironet 1530 series outdoor mesh access points
- Cisco Aironet 1600, 2600, 3600, 3500, 1700, 2700 and 3700 series indoor mesh access points
- Mesh features in Cisco Wireless LAN Controller
- Mesh features in Cisco Prime Infrastructure

Mesh Access Points

Access Point Roles

The 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, the MAPs have wireless connections to their controller.

MAPs communicate among themselves and back to the RAP using wireless connections over the 802.11a/n 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.


Bridge mode access points support CleanAir in mesh backhaul at 5GHz frequency and provides only the interference device report (IDR) and Air Quality Index (AQI) reports.

**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.

Figure 8-1 shows the relationship between RAPs and MAPs in a mesh network.

Figure 8-1 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 ACS (4.1 and later) or ISE 2.X that supports the client authentication type of Extensible Authentication Protocol-FAST (EAP-FAST) with certificates.

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.

Cisco Indoor Mesh Access Points

Indoor mesh is available on the following access points:

- Cisco Aironet 1600 Series Access Points
- Cisco Aironet 1700 Series Access Points
- Cisco Aironet 2600 Series Access Points
- Cisco Aironet 2700 Series Access Points
- Cisco Aironet 3500 Series Access Points
- Cisco Aironet 3600 Series Access Points
- Cisco Aironet 3700 Series Access Points
- Cisco Aironet 1530 Series Access Points
- Cisco Aironet 1540 Series Access Points
- Cisco Aironet 1550 Series Access Points
- Cisco Aironet 1560 Series Access Points
- Cisco Aironet 1570 Series Access Points
- Cisco Industrial Wireless 3700 Series Access Points

**Note**

The Cisco 1040 Series, 1140 Series, and 1260 Series access points have feature parity with Cisco Wireless Release 8.0. Features introduced in Cisco Wireless Release 8.1 and later are not supported on these access points.

**Note**

For more information about controller software support for access points, see the Cisco Wireless Solutions Software Compatibility Matrix at http://www.cisco.com/en/US/docs/wireless/controller/5500/tech_notes/Wireless_Software_Compatibility_Matrix.html.

Enterprise 11n/ac mesh is an enhancement added to the CUWN feature to work with the 802.11n/ac access points. Enterprise 11ac 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. The backhaul is supported only on the 5-GHz radio. If Universal Backhaul Access is enabled, the 5-GHz radio 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.

The Cisco indoor mesh access points are equipped with the following two simultaneously operating radios:

- From rel 8.2 2.4 GHz radio used for data backhaul and client access if UBA is enable
- 5-GHz radio used for data backhaul and client access if Universal Backhaul Access is enabled

The 5-GHz radio supports the 5.15 GHz, 5.25 GHz, 5.47 GHz, and 5.8 GHz bands.

Cisco Outdoor Mesh Access Points

Cisco outdoor mesh access points comprise of the Cisco Aironet 1500 series access points. The 1500 series includes 1572 11ac outdoor access points, 1552 and 1532 11n outdoor mesh access points, and the 1540 and 1560 11ac wave 2 series.

Cisco 1500 series mesh access points are the core components of the wireless mesh deployment. AP1500s are configured by both the controller (GUI and CLI) and Cisco Prime Infrastructure. Communication between outdoor mesh access points (MAPs and RAPs) is over the 802.11a/n/ac radio backhaul. Client traffic is generally transmitted over the 802.11b/g/n radio (802.11a/n/ac can also be configured to accept client traffic).

The mesh access point can also operate as a relay node for other access points that are not directly connected to a wired network. Intelligent wireless routing is provided by the Adaptive Wireless Path Protocol (AWPP). This Cisco protocol enables each mesh access point to identify its neighbors and intelligently choose the optimal path to the wired network by calculating the cost of each path in terms of the signal strength and the number of hops required to get to a controller.

AP1500s are manufactured in two different configurations:

- Cable—This configuration can be mounted to a cable strand and supports power-over-cable (POC).
- Non-cable—This configuration supports multiple antennas. It can be mounted to a pole or building wall and supports several power options.

Uplinks support includes Gigabit Ethernet (1000BASE-T) and a Small Form-Factor Pluggable(SFP) slot that can be plugged for a fiber or cable modem interface. Both single mode and multimode SFPs up to 1000BASE-BX are supported. The cable modem can be DOCSIS 2.0 or DOCSIS/EuroDOCSIS 3.0 depending upon the type of mesh access point.

AP1500s are available in a hazardous location hardware enclosure. When configured, the AP1500 complies with safety standards for Class I, Division 2, Zone 2 hazardous locations.

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.

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

**Note**

You can configure these modes using both the GUI and CLI. For configuration instructions, see the [Cisco Wireless LAN Controller Configuration Guide, Cisco Wireless Mesh Access Points, Design and Deployment Guide, Release 8.5](#)

**Note**


MAPs can only be configured in Bridge / Flex+Bridge mode regardless of their wired or wireless backhaul. If the MAPs have a wired backhaul, you must change their AP role to RAP before you change the AP Mode.

For complete details and specification of all models of outdoor Mesh AP please visit this link below: <https://www.cisco.com/c/en/us/products/wireless/outdoor-wireless/index.html?stickynav=1>

Frequency Bands

Both the 2.4-GHz and 5-GHz frequency bands are supported on the indoor and outdoor access points. All 1500 series Mesh APs support Channel Bands as indicated below.

Figure 8-2 Frequency Bands Supported by 802.11a/n/ac Radios on Mesh APs



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 subbands 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 subband 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.

Antennas

Overview

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

- Directional
- Omni-directional

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 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 omni-directional antennas ship with vertical polarization by default.

Antenna Options

A wide variety of antennas are available to provide flexibility when you deploy the mesh access points over various terrains. 5 GHz is used as a backhaul and 2.4 GHz is used for client access.

See the [Cisco Aironet Antenna and Accessories Reference Guide](#) on Cisco antennas and accessories.

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.

Client Access Certified Antennas (Third-Party Antennas)

You can use third-party antennas with AP1500s. However, note the following:

- Cisco does not track or maintain information about the quality, performance, or reliability of the non-certified 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 LAN Controllers

The wireless mesh solution is supported on Cisco 2500, 3504, 5500, and 8500 Series Wireless LAN Controllers.

For more information about the Cisco 2500, 3500, 5500, and 8500 Series Wireless LAN Controllers, see <https://www.cisco.com/c/en/us/products/wireless/wireless-lan-controller/index.html>

Cisco Prime Infrastructure

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

With the Prime Infrastructure, 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 Prime Infrastructure vital to ongoing network operations.

The Prime Infrastructure 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 Prime Infrastructure, 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 non-mesh) in the network..




Note

CAPWAP significantly reduces capital expenditures (CapEx) and operational expenses (OpEx), which enables the Cisco wireless mesh networking solution to be a cost-effective and secure deployment option in enterprise, campus, and metropolitan networks.

CAPWAP Discovery on a Mesh Network

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

- Step 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.
 - Step 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.
- 

Note The mesh access point searches a list of controllers configured on the access point (primed) during setup.
- Step 3** If [Step 2](#) fails after 10 attempts, the mesh access point falls back to DHCP and attempts to connect in 10 tries.
 - Step 4** If both [Step 2](#) and [Step 3](#) fail and there is no successful CAPWAP connection to a controller, then the mesh access point falls back to LWAPP.
 - Step 5** If there is no discovery after attempting [Step 2](#), [Step 3](#), and [Step 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.

Air Time Fairness in Mesh Deployments Rel 8.4

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 8.4 release

Mesh ATF is supported on AireOS 8.4 and all other supported APs as indicated in Release Notes. Mesh ATF is supported on 1550/128, 1570 and all other IOS based APs.

Table 8-1

	Access Points						
	1550 (64 MB)	1550 (128 MB)	1570	3700	1530	1540	1560
Features							
Basic Mesh	Yes	Yes	Yes	Yes	Yes	Yes	8.4
Flex+Mesh	Yes	Yes	Yes	Yes	Yes	No	No
Fast Convergence (background scanning)	No	8.3	8.3	Yes	8.3	No	8.4
Wired Clients on RAP	Yes	Yes	Yes	No	Yes	No	No
Wired Clients on MAP	Yes	Yes	Yes	No	Yes	No	8.4
Daisy Chain	7.6	7.6	7.6	No	7.6	No	No
LSC	Yes	Yes	Yes	Yes	Yes	No	No
PSK provisioning: MAP-RAP authentication	8.2	8.2	8.2	8.2	8.2	8.5	8.4
ATF on Mesh	No	8.4	8.4	8.4	No	No	No

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 group 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 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

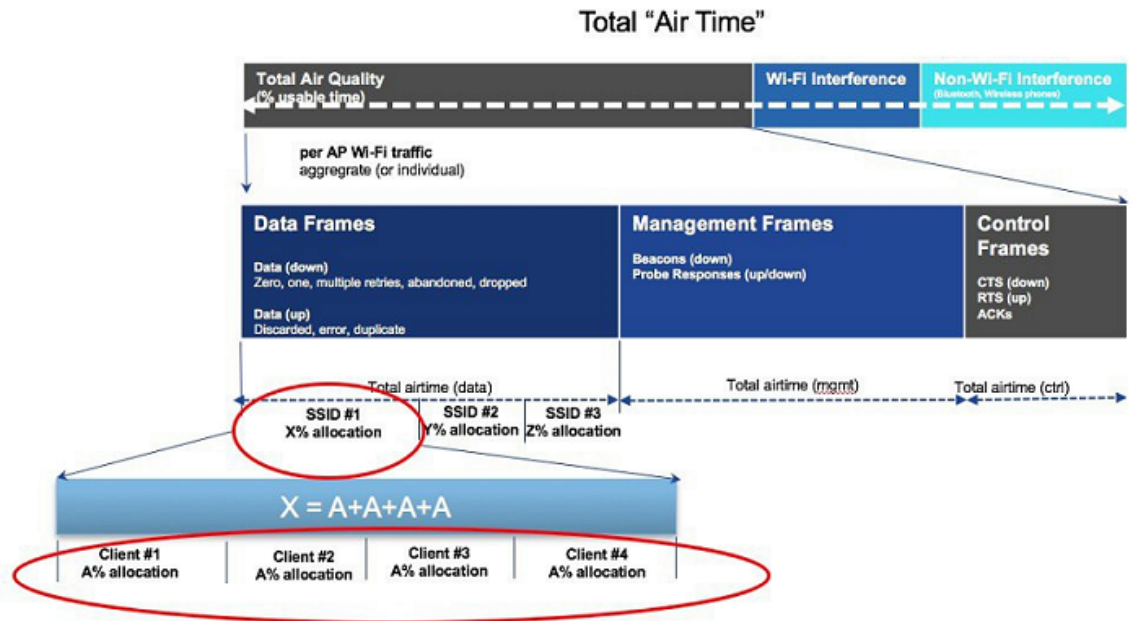
ATF Functionality and Capabilities:

- 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 airtime in the uplink direction, that is client to AP, can be measured, it cannot be strictly controlled. 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 gets ignored
- When ATF is configured per-SSID, each SSID is granted airtime according to the configured policy
- 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 transmit 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 can be globally enabled or disabled
- ATF can be enabled or disabled on an individual access point, AP group or entire network
- Allocation is applied Per SSID and Per Client
- Applies to Downstream only
- Can be configured in WLC GUI/CLI and PI
- Can be applied to all APs on a Network to AP Group or one AP
- Supported on APs in Local mode: **AP1260, 1550-128Mb, 1570, 1700, 2600, 2700, 3500, 3600,3700**



Note

COS based APs or Wave-2 APs do not support ATF in the rel 8.5.



For more details and ATF configuration steps see Mesh Deployment Guide rel 8.5
https://www.cisco.com/c/en/us/td/docs/wireless/technology/mesh/8-5/b_mesh_85.html

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:


- 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.
- Wireless mesh data frame flow.
- 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.

Figure 8-3 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 8-3 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 8-4.)

Figure 8-4 Logical Bridge and WLAN Mapping




Mesh Neighbors, Parents, and Children

Relationships among mesh access points are as a parent, child, or neighbor (see Figure 8-5).

- 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.

Figure 8-5 Parent, Child, and Neighbor Access Points



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.

Figure 8-6 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 8-6 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 Back-Ground Scan rel 8.3

In release 8.3 additional enhancements introduced for faster mesh convergence - Mesh AP background scan feature . There are already two Mesh convergence features implemented in releases 8.0 and 8.1 WLC software releases to reduce convergence time taken for a MAP and to re-converge the mesh network faster:

- Mesh Subset channel based convergence in rel 8.0
- Mesh Clear Channel Notification Convergence in rel 8.1

With both features in place, a 3rd Hop MAP in a mesh tree is able to re-converge and recover its data path in less than 10 seconds.

This new Mesh Background Scanning and Auto parent selection will further 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.



Note

This implementation of BG scanning will be applicable to Marvell-based APs. Specifically, AP1550, AP1560, AP1570 and IW3702.

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. A misscount parameter will indicate the number of times a neighbor did not respond to off-channel scan attempt. Each adjacency neighbor will have its adjusted Ease updated after every BG Scan cycle with latest linkSNR value.

This feature 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. This 'switching' parent procedure need not be a triggered event like parent loss detection, but also on identifying a better parent using 'Auto parent selection algorithm' when the child MAP still has its current parent uplink active. The "Auto Parent selection algorithm" is based on the new "ease" values. For better convergence calculation in rel 8.3 a new "Ease" value introduced for smoother and faster parent or neighbor finding and auto parent connection algorithm. The Ease value is based on SNR, number of Hops, Timers and load values. For the off-channel neighbors the AdjustedEase value will be used and a best neighbor per off-channel shall be identified based on its highest AdjustedEase value. StickyEase shall be applicable only for on-channel parent.

A child MAP will switch between optimal parents based on periodic evaluation of best neighbors across all off-channels. Best next parent is identified with highest adjustedEase value of another off-channel neighbor compared to current on-channel parent's stickyEase.

Table below illustrates the new convergence times based on different convergence configuration options. With the implementation of the latest CCN and Background scan features and fast convergence the First Hop MAP can achieve a 3 to 4 sec convergence.

Table 8-2

Parent Loss Detection/ Keep Alive	Parent Loss Detection/ Keep Alive	Channel Scan/Seek	DHCP/CAPWAP Information	Time per hop (sec)
Standard	21 / 3 sec	Scan/Seek all 2.4 and 5 Ghz Channels	Renew / Restart CAPWAP	48.6*
Fast	7 / 3 sec	Scan/Seek only channels found in same bridge group	Maintain DHCP and CAPWAP	48.6*
Very Fast	4 / 1.5 sec	Scan/Seek only channels found in same bridge group	Maintain DHCP and CAPWAP	15.9*
CCN/BG Scan Fast/VF	4 / 3 sec for 50ms	Scan/Seek only channels found in same bridge group	Maintain DHCP and CAPWAP	8-10 sec

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

DFS channel scan

Per regulatory restrictions, an AP shall not transmit over a DFS channel (when set on radio for off-channel scan) unless the channel is declared to be 'safe to send'. If there are radar signals detected, there should be no transmission and the channel should be avoided for AP's wireless Tx/Rx. One way to make sure that the channel is safe to send is by when the AP does passive scan and it receives any packet from other neighbors who are on DFS on-channel.

- To enable MAPs receive a packet during their off-channel scan over a DFS channel, all other on-channel DFS neighbors shall transmit the AWPP mesh beacons if there is no Tx/Rx in last 50 milliseconds.
- These mesh beacons will help the MAPs which is doing the off-channel on the DFS channel to declare 'safe to send' and do off-channel activities.

Mesh Deployment Modes

In a Cisco wireless outdoor mesh network, multiple mesh APs 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 APs 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 APs 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 APs 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

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 APs. 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 AP neighbor relationship with other mesh APs. The encryption keys used between mesh APs are derived during the EAP authentication process.

Universal Access

You can configure the backhaul on mesh APs to accept client traffic over its 802.11a radio. This feature is identified as Backhaul Client Access in the controller GUI (**Monitor > Wireless**). When this feature is disabled, backhaul traffic is transmitted only over the 802.11a or 802.11a/n radio and client association is allowed only over the 802.11b/g or 802.11b/g/n radio. For more information about the configuration, *see* Configuring Advanced Features.

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.



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. To enable Ethernet bridging using the controller GUI, choose **Wireless > All APs > Details for the AP** page, click the Mesh tab, and then select the **Ethernet Bridging** check box.



Note

The overall throughput of backhaul radio decreases by half for each hop of a mesh tree. When the Ethernet-bridged clients 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.

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

1. When you want to use the mesh nodes as bridges.
2. 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.

To configure range parameters for longer links, choose **Wireless > Mesh**. 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;


Default: 12,000 feet

Point-to-Multipoint Wireless Bridging

In the point-to-multipoint bridging scenario, a RAP acting as a root bridge connects multiple MAPs as nonroot 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.

Figure 8-7 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 8-7 Point-to-Multipoint 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. To enable Ethernet bridging using the controller GUI, choose **Wireless > All APs > Details** from the AP page, click the **Mesh** tab, and then check the **Ethernet Bridging** check box.

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.

To configure range parameters for longer links, choose **Wireless > Mesh**. Optimum distance (in feet) should exist between the root AP (RAP) and the farthest mesh AP (MAP). Range from the RAP bridge to the MAP bridge has to be mentioned in feet.

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

- Range: 150 to 132,000 feet
- Default: 12,000 feet

Wireless Backhaul Data Rate

Backhaul is used to create only the wireless connection between the APs. The backhaul interface by default is 802.11a or 802.11a/n depending upon the AP. 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 AP than can higher data rates, for example 300 Mbps. As a result, the data rate affects cell coverage and consequently the number of APs 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.

ClientLink Technology

Many networks still support a mix of 802.11a/g and 802.11n clients. Because 802.11a/g clients (legacy clients) operate at lower data rates, the older clients can reduce the capacity of the entire network.

Cisco ClientLink technology can help to solve problems related to adoption of 802.11n in mixed-client networks by ensuring that 802.11a/g clients operate at the best possible rates, especially when they are near cell boundaries.

Advanced signal processing has been added to the Wi-Fi chipset. Multiple transmit antennas are used to focus transmissions in the direction of the 802.11a/g client, increasing the downlink signal-to-noise ratio and the data rate over range, thereby reducing coverage holes and enhancing the overall system performance. This technology learns the optimum way to combine the signal received from a client and then uses this information to send packets in an optimum way back to the client. This technique is also

referred to as Multiple-input multiple-output (MIMO) beamforming, transmit beamforming, and it is the only enterprise-class and service provider-class solution in the market that does not require expensive antenna arrays.

The 802.11n systems take advantage of multipath by sending multiple radio signals simultaneously. Each of these signals, called a spatial stream, is sent from its own antenna using its own transmitter. Because there is some space between these antennas, each signal follows a slightly different path to the receiver, a situation called spatial diversity. The receiver has multiple antennas as well, each with its own radio that independently decodes the arriving signals, and each signal is combined with signals from the other receiver radios. This results in multiple data streams receiving at the same time. This enables a higher throughput than previous 802.11a/g systems, but requires an 802.11n capable client to decipher the signal. Therefore, both AP and client need to support this capability. Due to the complexity of issues, in the first generation of mainstream 802.11n chipsets, neither the AP nor client chipsets implemented 802.11n transmit beamforming. Therefore, the 802.11n standard transmit beamforming will be available eventually, but not until the next generation of chipsets take hold in the market. We intend to lead in this area going forward.

We realized that for the current generation of 802.11n APs, while the second transmit path was being well utilized for 802.11n clients (to implement spatial diversity), it was not being fully used for 802.11a/g clients. In other words, for 802.11 a/g clients, some of the capabilities of the extra transmit path was lying idle. In addition, we realized that for many networks, the performance of the installed 802.11 a/g client base would be a limiting factor on the network.

To take advantage of this fallow capacity and greatly enhance overall network capacity by bringing 802.11 a/g clients up to a higher performance level, we created an innovation in transmit beamforming technology, called ClientLink.

ClientLink uses advanced signal processing techniques and multiple transmit paths to optimize the signal received by 802.11a/g clients in the downlink direction without requiring feedback. Because no special feedback is required, Cisco ClientLink works with all existing 802.11a/g clients.

Cisco ClientLink technology effectively enables the AP to optimize the SNR exactly at the position where the client is placed. ClientLink provides a gain of almost 4 dB in the downlink direction. Improved SNR yields many benefits, such as a reduced number of retries and higher data rates. For example, a client at the edge of the cell that might previously have been capable of receiving packets at 12 Mbps could now receive them at 36 Mbps. Typical measurements of downlink performance with ClientLink show as much as 65 percent greater throughput for 802.11a/g clients. By allowing the Wi-Fi system to operate at higher data rates and with fewer retries, ClientLink increases the overall capacity of the system, which means an efficient use of spectrum resources.

ClientLink in the 1552 APs is based on ClientLink capability available in AP3500s. Therefore, the AP has the ability to beamform well to nearby clients and to update beamforming information on 802.11ACKs. Therefore, even if there is no dedicated uplink traffic, the ClientLink works well, which is beneficial to both TCP and UDP traffic streams. There are no RSSI watermarks, which the client has to cross to take advantage of this beamforming with Cisco 802.11n APs.

ClientLink can beamform to 15 clients at a time. Therefore, the host must select the best 15 if the number of legacy clients exceeds 15 per radio. AP1552 has two radios, which means that up to 30 clients can be beamformed in time domain.

For the latest updates please see the link

https://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-5/b_feature_matrix_for_802_11_ac_wave2_access_points.html:

Controller Planning

The following items affect the number of controllers required in a mesh network:

- Mesh APs (RAPs and MAPs) in the network.
- The wired network that connects the RAP and controllers can affect the total number of APs supported in the network. If this network allows the controllers to be equally available to all APs without any impact on WLAN performance, the APs 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 APs and coverage are reduced.
- Number of mesh APs (RAPs and MAPs) supported per controller.

For clarity, non-mesh APs are referred to as local APs in this document.

Table 8-3 Mesh AP Support by Controller Model

Controller Model	Local AP Support (nonmesh) ¹	Maximum Possible Mesh AP Support
5508 ²	500	500
2504 ³	75	75
3504	150	150
WiSM2	500	500
5520	1500	1500
8540	6000	6000

1. Local AP support is the total number of non-mesh APs supported on the controller model.
2. For 5508, controllers, the number of MAPs is equal to (local AP support - number of RAPs).
3. For 2504, controllers, the number of MAPs is equal to (local AP support - number of RAPs)



Note

Mesh is fully supported on all above mentioned Cisco Controllers. The Base License (LIC-CT508-Base) is sufficient for indoor and outdoor APs. The WPlus License (LIC-WPLUS-SW) is merged with the base license. The WPlus License is not required for indoor mesh APs.

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).

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

Cell Planning and Distance

For the Cisco 1520 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 non-voice 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 AP 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.
- For 2.4 GHz, the local access cell size radius is 600 feet (182.88 meters). One cell size is around 1.310×10^6 , so there are 25 cells per square mile.

Collocating Mesh Access Points

The following recommendations provide guidelines to determine the required antenna separation when you collocate AP1500s 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 AP1500s on Adjacent Channels

If two collocated AP1500s operate on adjacent channels such as channel 149 (5745 MHz) and channel 152 (5765 MHz), the minimum vertical separation between the two AP1500s is 40 feet (12.192 meters) (the requirement applies for mesh APs equipped with either 8 dBi omni-directional or 17 dBi high-gain directional patch antennas).

If two collocated AP1500s operate on channels 1, 6, or 11 (2412 to 2437 MHz) with a 5.5-dBi omni-directional antenna, then the minimum vertical separation is 8 feet (2.438 meters).

Collocating AP1500s on Alternate Adjacent Channels

If two collocated AP1500s operate on alternate adjacent channels such as channel 149 (5745 MHz) and channel 157 (5785 MHz), the minimum vertical separation between the two AP1500s is 10 feet (3.048 meters) (the requirements applies for mesh APs equipped with either 8-dBi omni-directional or 17-dBi high-gain directional patch antennas).

If two collocated AP1500s operate on alternate adjacent channels 1 and 11 (2412 MHz and 2462 MHz) with a 5.5-dBi omni-directional antenna, then the minimum vertical separation is 2 feet (0.609 meters).

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

CleanAir

The 1550 series leverages 802.11n technology with integrated radio and internal/external antennas. The 1550 series APs are based on the same chipset as the present CleanAir capable Aironet 3600 and 3700 APs. In other words, the 1550 series APs are capable of doing CleanAir.

With the 7.3.101.0 Release, 2600 series APs can mesh with each other and can also provide CleanAir functionality.

With the 7.2.103.0 Release, 3600 series APs can mesh with each other and can also provide CleanAir functionality.

With the 7.0.116.0 Release, 3500 series APs can mesh with each other and can also provide CleanAir functionality.

CleanAir in mesh (1552, 2600, 2700, 3500, 3600 and 3700) can be implemented on the 2.4-GHz radio and provides clients complete 802.11n 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 11n platform detects, quantifies, and mitigates Wi-Fi and non-Wi-Fi interference on 2.4-GHz radios. AP1552 supports CleanAir in 2.4 GHz client access mode.

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 backhaul radio of APs in bridge mode.

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 (WLCs) 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 APs.

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 APs and the CAPWAP controller.

Increasing Mesh Availability

In the [Cell Planning and 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 8-8](#) or to use RAPs placed on different channels, as shown in [Figure 8-9](#). The addition of RAPs into an area adds capacity and resilience to that area.

Figure 8-8 Two RAPs per Cell with the Same Channel





Figure 8-9 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 non-overlapping 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 omni-directional 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 AP 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.

Indoor Mesh Interoperability with Outdoor Mesh

Complete interoperability of indoor mesh APs with the outdoor ones is supported. It helps to bring coverage from outdoors to indoors. We recommend indoor mesh APs for indoor use only, and these APs should be deployed outdoors only under limited circumstances as described below.



Caution

The indoor APs in a third-party outdoor enclosure can be deployed for limited outdoor deployments, such as a simple short haul extension from an indoor WLAN to a hop in a parking lot. The 1240, 1250, 1260, 1700, 2600, 2700, 3500e, 3600 and 3700 APs in an outdoor enclosure is recommended because of its robust environmental and temperature specifications. Additionally, the indoor APs have connectors to support articulated antennas when the AP is within an outdoor enclosure. Exercise caution with the SNR values as they may not scale and long-term fades may take away the links for these APs when compared to a more optimized outdoor 1500 series AP.


Mobility groups can be shared between outdoor mesh networks and indoor WLAN networks. It is also possible for a single controller to control indoor and outdoor mesh APs simultaneously. The same WLANs are broadcast out of both indoor and outdoor mesh APs.

Connecting the Cisco 1500 Series Mesh APs to the Network

This section describes how to connect the Cisco 1500 Series mesh APs to the 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 (see [Figure 8-10](#)). The WLAN client traffic from CAPWAP is tunneled at Layer 2, and matching WLANs should terminate on the same switch VLAN where the controllers are collocated. The security and network configuration for each of the WLANs on the mesh depend on the security capabilities of the network to which the controller is connected.

Figure 8-10 Mesh Network Traffic Termination



Note

When an HSRP configuration is in operation on a mesh network, we recommend that the In-Out multicast mode be configured. For more details on multicast configuration, refer to the [Cisco Mesh Access Points, Design and Deployment Guide](#).

