CHAPTER 8

Cisco Wireless Mesh Networking

This chapter summarizes the design details for deploying a Cisco Wireless mesh network for outdoor environments. It focuses primarily on design considerations for mesh deployment, but also covers the solution components and interworkings. For further details about the Cisco Wireless Mesh solution, refer to the *Cisco Aironet 1500 Series Wireless Mesh AP Version 5.0 Design Guide* at: http://www.cisco.com/en/US/docs/wireless/technology/mesh/design/guide/MeshAP.html.

Introduction

The Cisco Wireless Mesh solution enables cost-effective and secure deployment of outdoor Wi-Fi networks. Outdoor wireless access takes advantage of the growing popularity of inexpensive Wi-Fi clients, enabling new service opportunities and applications that improve user productivity and responsiveness.

As the demand for outdoor wireless access increases, customers faced with tight budgets and reduced resources must respond with wireless LAN (WLAN) solutions that take full advantage of existing tools, knowledge, and network resources to address ease of deployment and WLAN security issues in a cost-effective way. The Cisco Wireless Mesh solution is an outdoor WLAN solution that excels in the unique attributes of wireless mesh technology, effectively supports current networking requirements, and lays the foundation for the integration of business applications.

Outdoor wireless solutions offer a number of challenges compared to a standard indoor WLAN, particularly in these areas:

- Environment
- Coverage
- Total cost of ownership (TCO)
- Physical device security

The outdoor environment is harsher than the indoor environment and so requires specialized equipment or enclosures to contain and protect indoor equipment that is deployed outdoors.

Outdoor WLAN deployments attempt to cover wider areas than indoor wireless networks, while addressing the challenges of less control over sources of interference, finding a suitable wired connection to connect the wireless mesh network to the wired network, and the availability of power for the mesh network devices.

Outdoor deployments also require specialized radio frequency (RF) skills, may have a lower user density than indoor deployments, and may be deployed in environment that is less regulated than inside a building. These features put pressure on the TCO of the outdoor solutions and require a solution that is easy to deploy and maintain.

The Cisco Wireless Mesh solution has three core components:

- Cisco 1500 Series Mesh AP—Outdoor access point that provides WLAN client access in the mesh and backhaul client connections
- Cisco Wireless LAN controller (WLC)—Provides a central point for AP control functions
- Cisco Wireless Control System (WCS)—Management platform for enhanced scalability, manageability, and visibility of large-scale implementations

Figure 8-1 shows a simple mesh network deployment made up of mesh APs, WLCs, and a WCS. In this example deployment, there are three mesh APs connected to the wired network. These APs are designated as roof-top APs (RAPs); all other APs in the mesh network are known simply as mesh APs (MAPs). All mesh APs, both MAP and RAP, can provide WLAN client access, however in most cases because of the RAPs location it is not well suited for providing client access. In the following example the RAPs are located on the roof of each of the buildings and are connected to the network at each location. Some of the buildings have WLCs located at them to terminate LWAPP sessions from the mesh APs, but it is not necessary for every building to have a WLC. LWAPP sessions can be back hauled across the WAN if needed to other locations where a WLC resides.

Figure 8-1 Mesh Solution Diagram

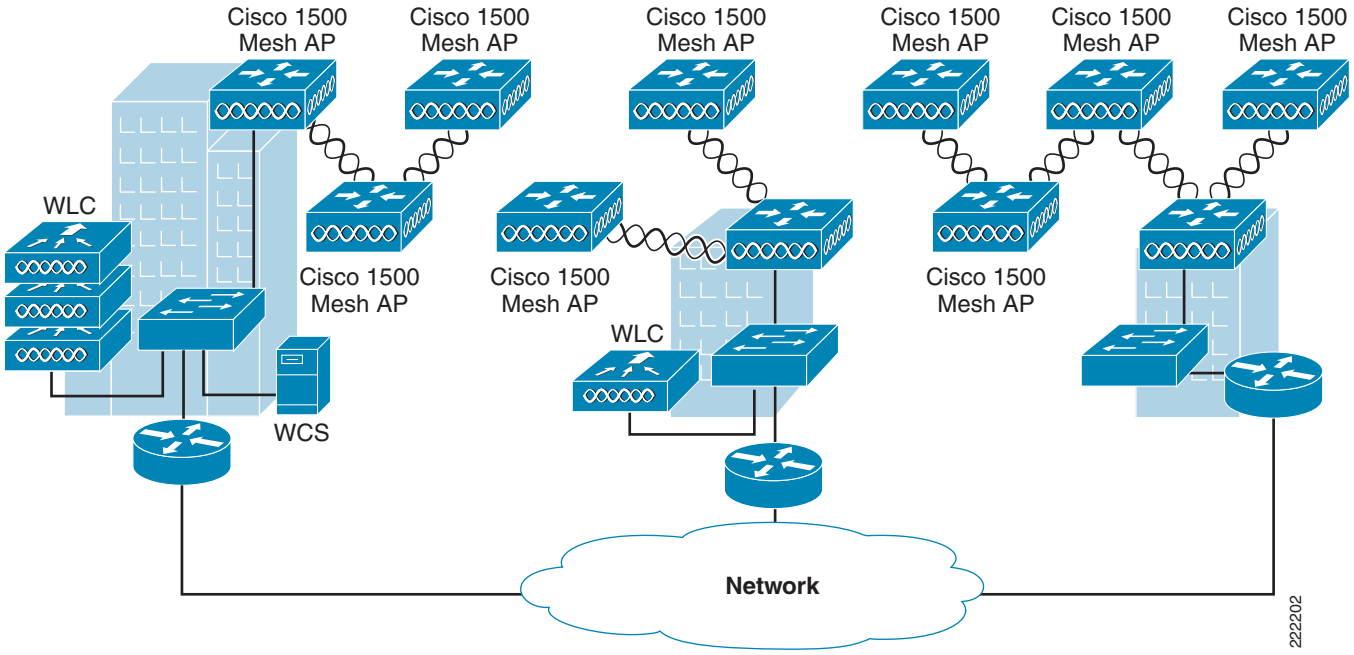

Cisco 1500 Series Mesh AP

The Cisco 1500 Series Mesh AP shown in Figure 8-2 is the core component of the wireless mesh solution and leverages existing and new features and functionality in the Wireless LAN controllers and the WCS.

Figure 8-2 Cisco 1510 and 1520 Wireless mesh APs

There are three types of Cisco 1500 Series Mesh APs:

- The AP1520—An outdoor access point consisting of two simultaneous operating radios:
  - One 2.4 GHz radio that is used for client access.
  - One 5.8/4.9 GHz radio that is used for data backhaul to other 1500 Series Mesh APs.
- The AP1520 also has a modular design and can be configured with the following optional uplink interfaces:
  - Cable Modem DOCSIS 2.0 with Cable Power Supply
  - Fiber Interface with 100BaseBX SFP
  - 1000BaseT Gig Ethernet
- The AP1510—An outdoor access point consisting of two simultaneous operating radios:
  - One 2.4 GHz radio that is used for client access.
  - One 5.8/4.9 GHz radio that is used for data backhaul to other 1500 Series Mesh APs.

  The AP1510 also has an Ethernet port which can be used for connectivity to the WLC or for a connected LAN segment for bridging.
- The AP1505—An outdoor access point consisting of a single 2.4 GHz radio:
  - One 2.4 GHz radio that is used for client access and backhaul.
  - Like the AP1510, the AP1505 also has a wired Ethernet port.

A wide variety of antennas are available to provide flexibility when deploying the 1500 Series Mesh AP over various terrain. The 5.8 GHz frequency radio uses 802.11a technology and is used in the system as the backhaul or relay radio. Wireless LAN client traffic arrives at the AP via the 2.4GHz radio passes either through the AP backhaul radio or is relayed through other 1500 Series Mesh APs until it reaches the WLC Ethernet connection.

The 1500 Series Mesh AP also has a 10/100 Ethernet connection to provide bridging functionality. This Ethernet connection supports power over Ethernet (PoE) through a separate power injection system.

Note The power injector is unique for this product; other Cisco power injection solutions are not suitable for use with the Cisco 1500 Series Mesh AP.

The Cisco 1500 Series Mesh AP uses LWAPP to communicate to a wireless controller and other 1500 Series Mesh APs in the wireless mesh.

The 1500 Series Mesh AP is designed to be mounted upside-down with its antenna oriented vertically, as shown in Figure 8-3.

Figure 8-3 1500 Series Mesh AP Installation

Cisco Wireless LAN Controllers

The wireless mesh solution is supported by the Cisco 4400 Series Wireless LAN Controller (WLC) (shown in Figure 8-4) and the Cisco Wireless Services Module (WiSM) (shown in Figure 8-5). Either platform is recommended for wireless mesh deployments because they can both scale to large numbers of access points and can support both Layer 2 and Layer 3 LWAPP connections.

Figure 8-4 Cisco 4400 Wireless LAN Controller

Figure 8-5 Cisco Wireless Services Module

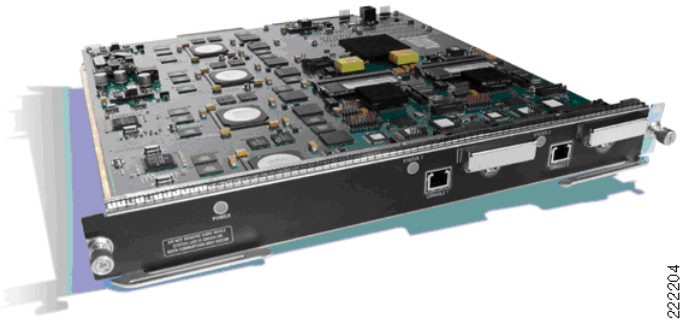

For more information on Cisco Wireless LAN controllers, see: http://www.cisco.com/en/US/prod/collateral/wireless/ps5679/ps6548/prod_brochure0900aecd8036884a_ns621_Networking_Solutions_Brochure.html.

Wireless Control System (WCS)

The Cisco Wireless Control System (WCS) is the platform for wireless mesh planning, configuration, and management. It provides the tools to allow network managers to design, control, and monitor wireless mesh networks from a central location.

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

Wireless Mesh Operation

In a wireless mesh deployment, there are multiple 1500 Mesh APs deployed as part of the same network. Mesh APs form parent, child, and neighbor relationships with each other to form the mesh and establish a LWAPP tunnel back to their specified primary WLC. Parent, child, and neighbor relationships are discussed further in Mesh Neighbors, Parents, and Children, page 8-10.

MAPs use the Adaptive Wireless Path Protocol (AWPP) to determine the best path through other 1500 Mesh APs to their WLC. The wireless links between the MAPs and RAP(s) form a wireless mesh that is used to carry traffic from WLAN clients (through LWAPP tunnels) to the WLC and also to carry bridge traffic between devices connected to the MAP Ethernet ports.

A wireless mesh can simultaneously carry two different traffic types:

- WLAN client traffic through LWAPP tunnels
- MAP bridge traffic

WLAN client traffic terminates on the WLC, but the bridge traffic terminates on the Ethernet ports of the MAPs of the wireless mesh.

MAP membership in the wireless mesh can be controlled in a variety of ways. The default mesh AP authentication is EAP, but Pre Shared Key (PSK) authentication can also be configured. Bridge Group Name (BGN) is used in addition to authentication to control mesh membership or to segment a wireless mesh.

Bridge Authentication

When a mesh AP is turned on and connected to the network via a wired Ethernet connection, it joins a WLC using the following steps:

1. When the AP booted, it optionally obtained an IP address via DHCP if a static IP has not been previously configured.
2. The mesh AP s out a LWAPP discovery request.
3. If a WLC receives the request, it responds with a discovery response.
4. At this point the mesh AP issues a LWAPP join request.
5. The WLC issues an LWAPP join response and proceeds with EAP authentication.
6. Depending on the mesh AP's current image version, it may download a new image and re-boot.
7. After the reboot, the mesh AP requests to join the WLC again and re-authenticate.

Note PSK may be used in place of EAP if configured on the WLC.

If there is no wired connection for the mesh AP to use to connect to a WLC, it uses the following procedure to join the controller.:

1. After boot, the mesh AP forms a 802.11 association and issues a LWAPP discovery request via its 802.11a connection.
2. When a mesh AP with a connection to the WLC is discovered, it uses DHCP to obtain an IP address if one has not been statically configured.
3. At this point the mesh AP issues a LWAPP join request.
4. The WLC issues an LWAPP join response and proceeds with EAP authentication.
5. Depending on the mesh AP's current image version, it may download a new image and re-boot.
6. After the reboot, the mesh AP rediscovers its parent and requests to join the controller again and re-authenticate.

Wireless Mesh Encryption

As discussed above, the wireless mesh bridges traffic between the MAPs and the RAPs. This traffic can be from wired devices being bridged by the wireless mesh or LWAPP traffic from the mesh APs. This traffic is always AES encrypted when it crosses a wireless mesh link (see Figure 8-6).

The AES encryption is established as part of the mesh AP neighbor relationships establishment with other mesh APs. The encryption keys used between mesh APs are derived during the EAP authentication process.

Figure 8-6 Mesh Encryption

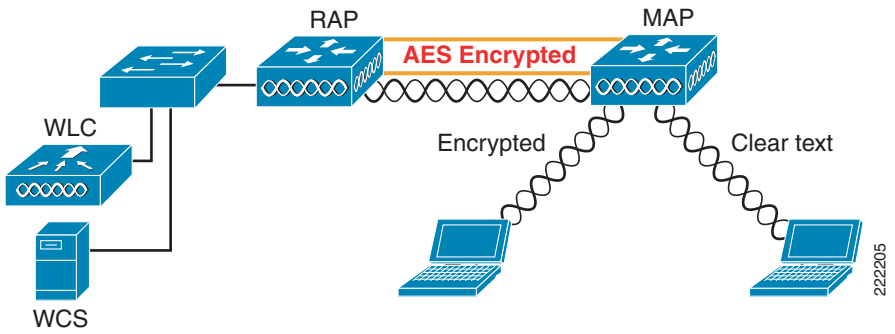

AWPP Wireless Mesh Routing

The core of the Cisco Wireless Mesh network is the Cisco Adaptive Wireless Path Protocol (AWPP).

This protocol is designed specifically for wireless mesh networking in that its path decisions are based on both link quality and the number of Mesh AP hops. AWPP is also designed to provide ease of deployment, fast convergence, and minimal resource consumption.

For more information on AWPP, refer to:
http://www.cisco.com/en/US/prod/collateral/wireless/ps5679/ps6548/prod_brochure0900aecd8036884a_ns621_Networking_Solutions_Brochure.html.

Example Simple Mesh Deployment

The key network components of a simple mesh deployment design shown in Figure 8-7 are the following:

- WCS—Key component in the management, operation, and optimization of the mesh network.
- Wireless LAN Controller— Provides real-time communication between lightweight access points and other wireless LAN controllers to deliver centralized security policies, wireless intrusion prevention system (IPS) capabilities, RF management, quality of service (QoS), and mobility.
- Router between the network and the mesh—Provides a Layer 3 boundary where security and policy enforcement can be applied. The router also provides Layer 2 isolation of the RAP. This is necessary because the RAP bridges traffic from its local Ethernet port to the mesh, so this traffic must be limited to that necessary to support the solution so that resources are not consumed by the unnecessary flooding of traffic.
- RAP— The wired network connected Mesh AP that provides the “path” home for the wireless mesh APs.
- A number of MAPs.

Note The RAP wireless connection is to the center of the MAP mesh, which is an optimal configuration that minimizes the average number of hops in the mesh. A RAP connection to the edge of a mesh would result in an increase of hops.

Figure 8-7 Simple Mesh Deployment

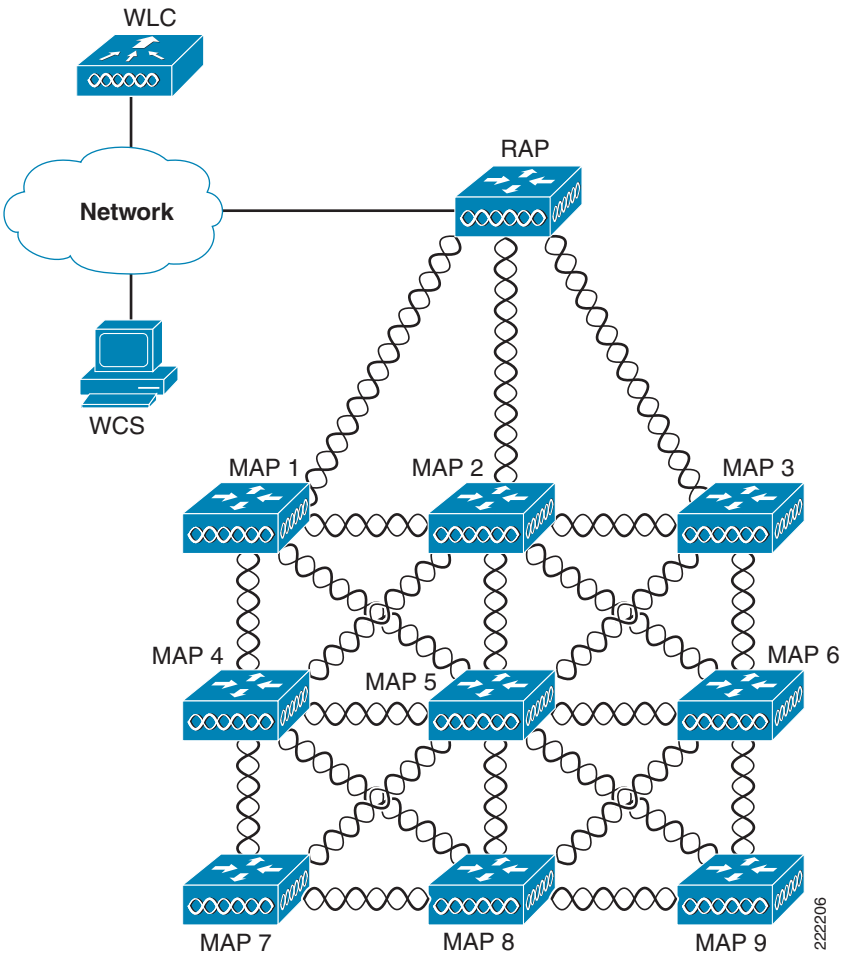

Figure 8-8 shows one possible logical view of the physical configuration, with MAP5 as the path home for all other MAPs.

Figure 8-8 Logical Mesh View

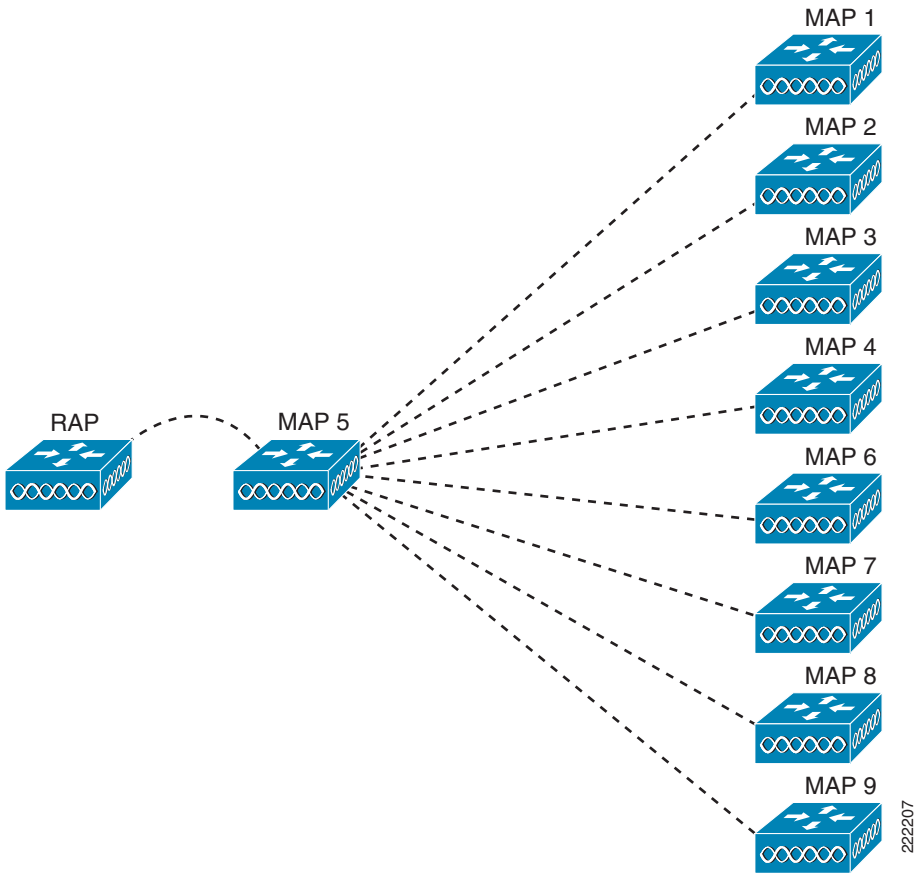

Figure 8-9 shows an alternate logical view, in which the signal-to-noise ratio (SNR) on the indirect paths to MAP5 is small enough for the other MAPs to consider taking an extra hop along a greater NSR link to get to MAP5.

Figure 8-9 Unequal Mesh Paths

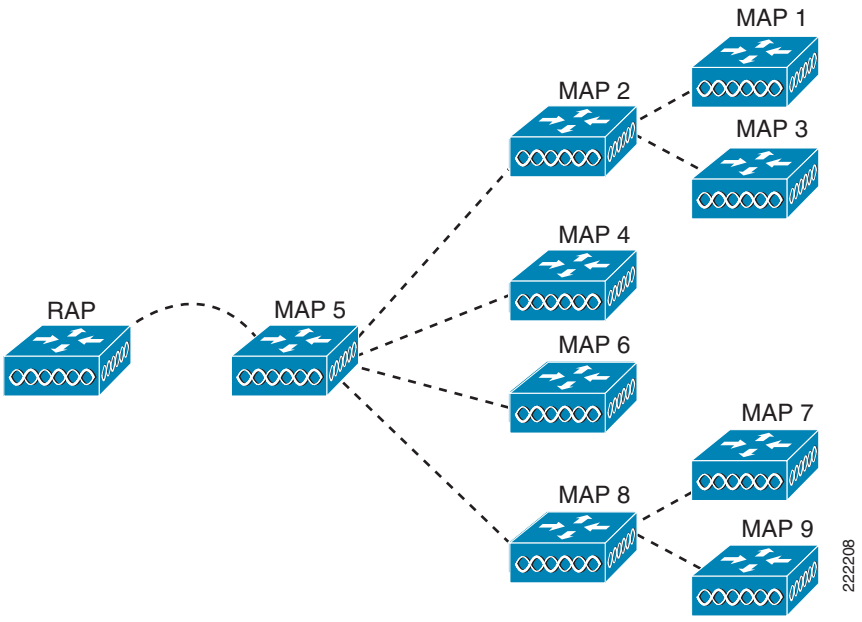

In both the cases above, MAP5 is the path home for all traffic. Ideally, the coverage from the RAP should be such that other MAPs, such as MAP2 for example, have a path back to the RAP, and traffic could be routed via MAP 2 in case of a loss of signal to MAP 5.

Mesh Neighbors, Parents, and Children

There are three relationships mesh APs can have with another:

- A *neighbor* within a mesh is an AP that is within RF range but has not been selected as a parent or a child because its "ease" values are lower than another neighboring AP (refer to Ease Calculation, page 8-14).
- A *parent* AP is one that is selected as the best route back to the RAP based on the best ease values. A parent can be either the RAP itself or another MAP.
- A *child* AP is one that has selected the parent AP as the best route back to the RAP. The example in Figure 8-10 illustrates a small mesh. In this example, MAP:7a:70 parent is RAP:d680 and the MAP:7a:70 child is MAP:78:90. Map:7a:70 also has a neighbor relation with MAP:d8:90.

Note A mesh AP can be both a parent other mesh APs and a child of another mesh AP, however a RAP is the only mesh AP that is not a child of any AP.

Figure 8-10 Parent, Child, and Neighbor

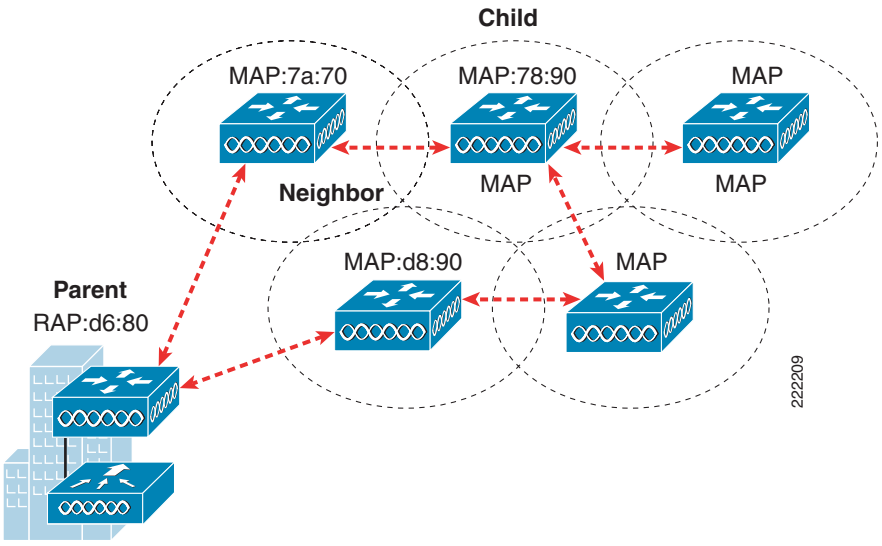

The goal of AWPP is to find the best backhaul link path for a MAP through the mesh back to a RAP. To do this, the mesh AP actively solicits for neighbor APs. During the solicitation, the mesh AP learns all of the available neighbors back to a RAP, determines which neighbor offers the best path, and then synchronizes with that neighbor.

Figure 8-11 shows the state diagram for a mesh AP when it is trying to establish a connection.

A mesh AP must first decide whether it is a RAP. A mesh AP becomes a RAP if it can communicate with an WLC through its Ethernet interface. If the mesh AP is a RAP, it can go straight to the maintain state. In the maintain state, the mesh AP has established an LWAPP connection to the controller so it does not need to seek other mesh APs, but simply responds to solicitations. If the mesh AP is not a RAP, it starts a scan process where the mesh AP scans all available channels and solicits information from other mesh APs.

Figure 8-11 Mesh AP State Diagram

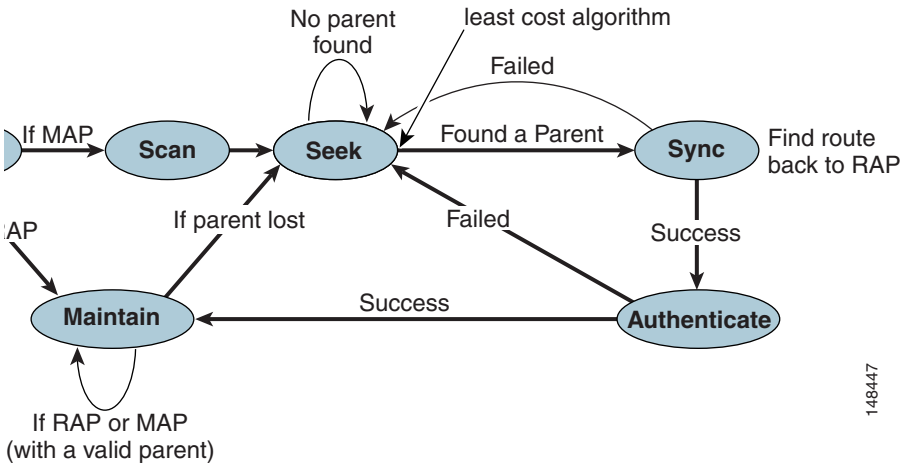

This behavior has two main implications:

- The RAP does not change channels, and therefore the channel used to build the mesh from a RAP is defined in the RAP configuration. By default, the RAP uses channel 161.

- The mesh is built from the RAP out, because initially only the RAP can respond to solicitations.

If the mesh AP is not a RAP, it follows the state diagram above in the following modes:

- Scan—The AP scans all the backhaul channels using mesh beaconing. This mechanism is similar to the 802.11 beaconing mechanisms used by wireless access networks. The frame used for beaconing is called NEIGHBOR_UPDATE. Essentially, NEIGHBOR_UPDATE frames are advertised by the network so that new nodes can scan and quickly discover neighbors. Each RAP and MAP broadcast NEIGHBOR_UPDATE frames after being connected to the network (via a WLAN controller). Any neighbor updates with SNRs lower than 10 dB are discarded. This process is called passive scanning.
- Seek—Solicits other members of the mesh. Mesh APs issuing successful responses to these solicitations become neighbors.
- Sync—The mesh AP learns the path information from each of its neighbors and the neighbor with the greatest ease becomes the parent of the soliciting mesh AP. If the neighbors report multiple RAPs, the RAP with the greatest ease is chosen.
- Authenticate—The mesh AP authenticates to the WLC through a connection established via its parent AP. This is a standard certificate-based LWAPP AP authentication.
- Maintain—The mesh AP responds to other mesh AP solicitations and regularly issues solicitations to determine any changes in the mesh. It is only after entering the maintain state that the mesh AP is visible to the WLC and WCS. Note that in the maintain state, the solicitations occur only on the channel defined by the mesh RAP, whereas a mesh AP in seek mode solicits on all channels, only stopping when it has found a parent AP.

Background Scanning in Mesh Networks

Background scanning allows Cisco 1500 series APs to actively and continuously monitor neighboring channels for optimum paths and parents. Because the access point is searching on neighboring channels as well as the current channel, the list of possible alternate optimal paths and parents is greater.

Identifying this information prior to the loss of a parent results in a faster switch over and the best link possible for the access point. Additionally, access points might switch to a new channel if a link on that channel is found to have a better cost metric (fewer hops, stronger SNR) than its current channel.

Background scanning on other channels and collecting of data from neighbors on those channels is done on the backhaul between two access points:

- For 1510 access points, the backhaul (primary) operates on the 802.11a link.
- For 1505 access points, the backhaul operates on the 802.11b/g link.

Background scanning is enabled on a global basis on the controller using the command-line interface:

config mesh background-scanning {enable | disable}

Enter this command to verify background scanning is enabled.

show mesh background-scanning

It is enabled by default.

Note Latency might increase for voice calls when they are switched to a new channel.

If channels requiring Dynamic Frequency Selection (DFS) are used, locating neighbors in other channels might take longer.

A few operating scenarios are provided below to better understand the operation of background scanning. In Figure 8-12, when the mesh access point, MAP1, initially comes up it is aware of both root access points, RAP1 and RAP2, as possible parents. RAP2 is chosen as the parent for MAP1 because the route through RAP2 provides a better cost metric.

Once the link is established, background scanning continuously monitors all channels in search of a better path and parent. RAP2 continues to act as parent for MAP1 and communicates on channel 2 until either the link goes down or a better path is located on another channel.

Figure 8-12 Mesh Access Point, MAP1, Selects Parent

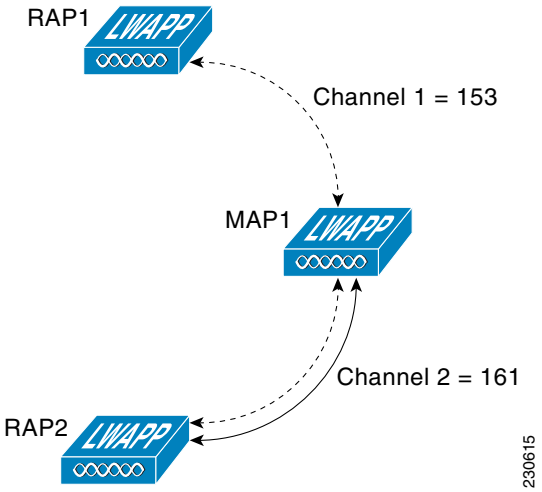

In Figure 8-13, the link between MAP1 and RAP2 is lost. Data from ongoing background scanning identifies RAP1 and channel 1 as the next best parent and communication path for MAP1, so that link is established immediately without the need for additional scanning after the link to RAP2 goes down.

Figure 8-13 Background Scanning Identifies New Parent

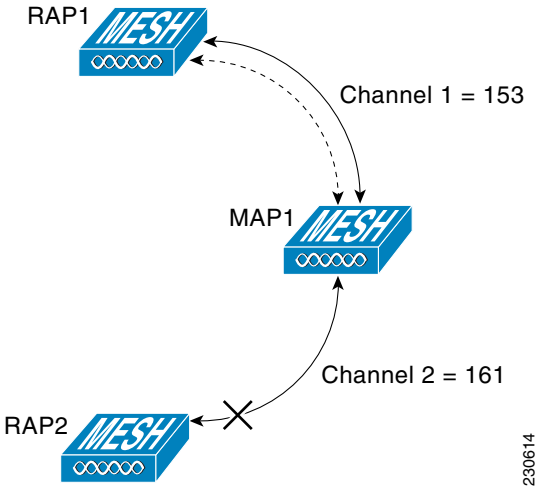

Enter this command to enable or disable background scanning on the controller:

config mesh background-scanning {enable | disable}

Enter this command to verify background scanning is enabled:

show mesh background-scanning

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.

A parent AP 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

In Figure 8-14, MAP2 prefers the path through MAP1 because the adjusted ease (436906) though this path is greater then the ease value (262144) of the direct path from MAP2 to RAP.

Figure 8-14 Parent Path Selection

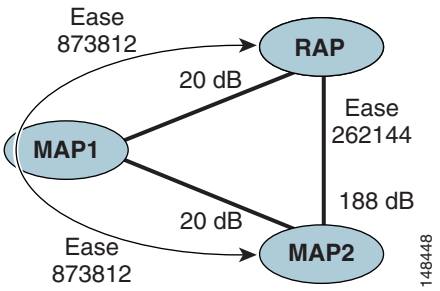

SNR Smoothing

One of the challenges in WLAN routing is the ephemeral nature of RF. This must be considered when determining 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; 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% of “bonus-ease” on top of the parent’s calculated ease to reduce flapping between parents. The assignment of a bonus-ease mandates that a potential parent must offer a significantly better route to the RAP in order for a child to make a switch. Parent switching is transparent to LWAPP and other higher-layer functions.

Loop Prevention

To ensure that routing loops are not created, AWP 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. This enables a 1500 Series mesh AP to easily detect and discard routes that loop.

Choosing the Best Mesh Parent

The OPS algorithm is implemented in the Seek state of the AWPP state machine. The basic steps of the parent selection algorithm in the AWPP (for both a RAP and MAP with radio backhaul) is as follows:

- 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 present are actively scanned in the Seek state and the backhaul channel is changed to the channel with the best neighbor ease.
- The parent is set to the best neighbor and the parent-child handshake is completed in 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, usually followed by an LWAPP network and controller discovery. All neighbor protocol frames carry the channel information. Both parent maintenance and optimization techniques remain unchanged, as described in the following:

- 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 as that of its parent and evaluating all responses from neighboring nodes on this channel. In most practical mesh networks, only a single channel backhaul is designed.
- A parent MAP is the MAP that has 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.

Routing Around an Interface

This feature is optional and is user configurable via Controller CLI only. If this feature is enabled, it transmits packets on secondary backhaul (b/g radio) when there is transient interference on the primary backhaul (A radio).

There are two modes of operation for Routing Around an Interface (RAI):

- **Config mesh secondary-backhaul enable**—This enables RAI globally on all APs. In order for RAI to work properly, the user has to configure the same “b/g” channel on all APs beyond the first HOP to the one that is being used on the first HOP “b/g” radio. If RRM (auto-rf) is enabled, then it changes the channels on APs and RAI will not work.
- **Config mesh secondary-backhaul enable force-same-secondary-channel**—This forces the whole subtree rooted at one hop MAPs to have the same secondary channel. Ignore RRM or manually assigned for MAPs at two hops and deeper.

Design Details

Each outdoor wireless mesh deployment is unique, with its own challenges regarding locations, obstructions, and network infrastructure availability. Such challenges must typically be addressed in addition to design requirements that are based on users, traffic, and availability. This section discusses important design considerations and provides an example of a wireless mesh design.

Wireless Mesh Design Constraints

When designing and building a wireless mesh network with the 1500 Mesh AP, there are a number of system characteristics to consider. Some of these apply to the backhaul network design and others to the WLC design:

- Recommended backhaul is 18 Mbps—18 Mbps is chosen as the optimal backhaul rate because it aligns with the maximum WLAN coverage distances of the MAP; that is, the distance between MAPs using 18 Mbps backhaul should allow for seamless WLAN client coverage between the MAPs. A lower bit rate may allow a greater distance between 1500 Mesh APs, 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 1500 Mesh APs or results in a reduced SNR between mesh APs, limiting mesh reliability and interconnection. The wireless mesh backhaul bit rate, like the mesh channel, is set by the RAP.
- Number of backhaul hops should be limited to three to four—The number of hops is recommended to be limited to three to four primarily to maintain sufficient backhaul throughput, because each mesh AP uses the same radio for transmission and reception of backhaul traffic. This means that throughput is approximately halved over every hop. For example, the maximum throughput for an 18 Mbps is approximately 10 Mbps for the first hop, 5 Mbps for the second hop, and 2.5 Mbps for the third hop.
- Number of MAPs per RAP—There is no current software limitation of how many MAPs per RAP you can configure. However, it is suggested that you limit this to 20 MAPs per RAP to avoid bottle necks in your mesh.
- Number of APs per controller—The number of APs per controller is determined by the controller capacity.
- Number of controllers—The number of controllers per mobility group is limited to 24.

Client WLAN

The mesh AP client WLAN delivers all the WLAN features of a standard 802.11bg LWAPP deployment with the full range of security and radio management features.

The goals of the client WLAN must be considered in the overall mesh deployment:

- What bit rates are required? Higher bit rates reduce coverage and are limited by the mesh backhaul.
- What throughput is required? What are the application throughput requirements and how many simultaneous clients are expected on a Cisco 1500 Mesh AP?
- What coverage is required? Is the coverage between different 1500 Mesh APs required to be contiguous or is the mesh deployment a collection of separate active zones?
- What security mechanisms are required? Is the WLAN intended for public consumption or private? What security is needed for client access?

Bridging Backhaul Packets

Bridging services are treated a little differently from regular controller-based services. There is no outer DSCP value in bridging packets because they are not LWAPP encapsulated. Therefore, the DSCP value in the IP header as it was received by the AP is used to index into the table as described in the path from AP to AP (backhaul).

Bridged frames received from a station on a LAN connected to a MAP are not modified in any way. There is no override value for an 802.1p classification. Therefore, in bridging mode the LAN traffic classification must be properly secured.

Frames are transmitted to the MAP LAN precisely as they are received upon the ingress to the wireless mesh bridge.

The 1500 does not tag modify DSCP:

- On the ingress port, the 1510 sees a DSCP marking and encapsulates the IP packet and applies the corresponding 802.1p priority.
- On the egress port, the 1510 decapsulates the IP packet and places it on the wire with an unmodified DSCP marking.

For this prioritization to be effective, Ethernet devices, such as IP video cameras, must have the capability to mark the DSCP of packets.

Client Access on Backhaul Connections

It is possible to allow client access on the 5.8 and 4.9 GHz backhaul connection while simultaneously transmitting backhaul traffic. This feature is particularly beneficial for deployments that need to support both 2.4 and 5GHz clients. This optional feature is turned off by default and can be enabled via the CLI command interface with the following command.

(Cisco Controller) >config mesh client-access *enable/disable*

In the GUI you can enable this feature in the mesh feature section as shown in Figure 8-15.

Figure 8-15 Client Access

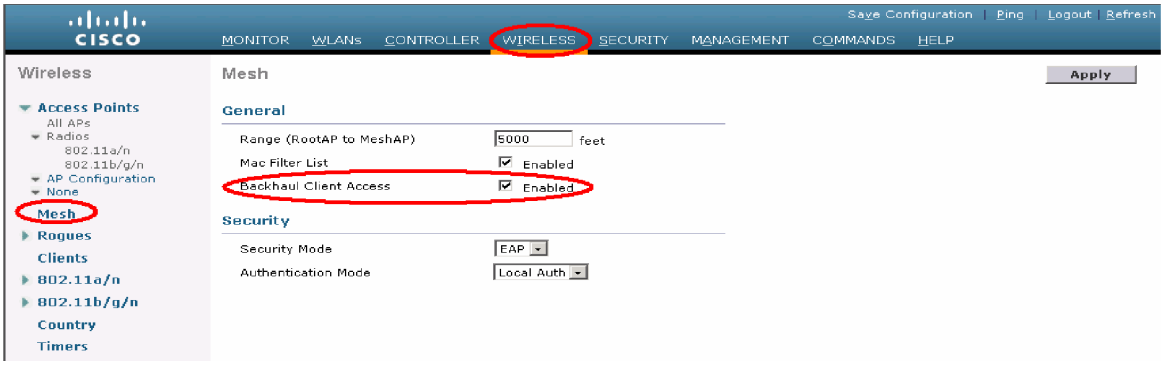

Increasing Mesh Availability

A wireless mesh cell has similar properties to the cells used to create a cellular phone network. The technology may define the maximum size of the cell; smaller cells can be created to cover the same physical area, providing greater availability or capacity. This is done by adding RAPs to the cell. Just as in the larger mesh deployment, the decision is whether to use RAPs on the same channel, as shown in Figure 8-16, or use different channels, as shown in Figure 8-17. The addition of RAPs into an area adds capacity and resilience to that area.

Figure 8-16 Two RAPs per Cell with the Same Channel

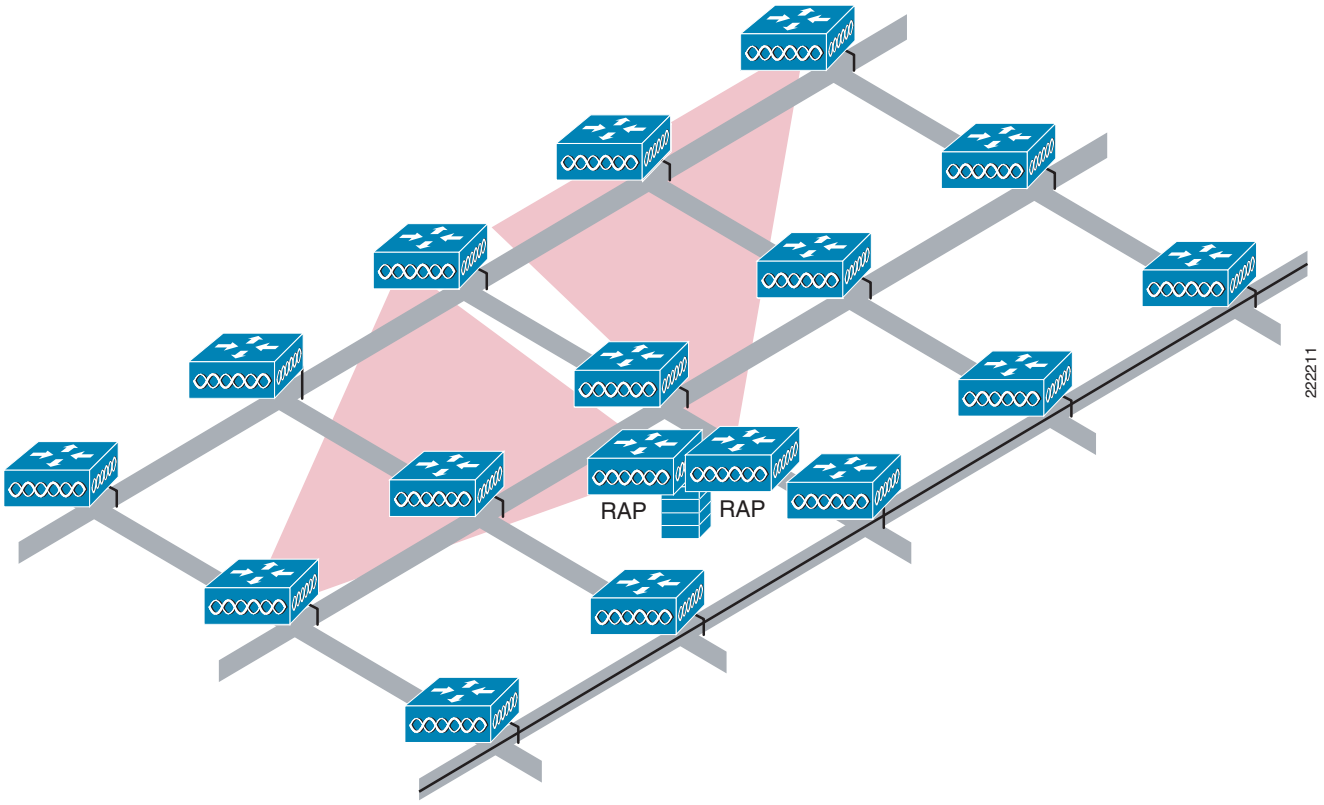

Figure 8-17 Two RAPs per Cell on Different Channels

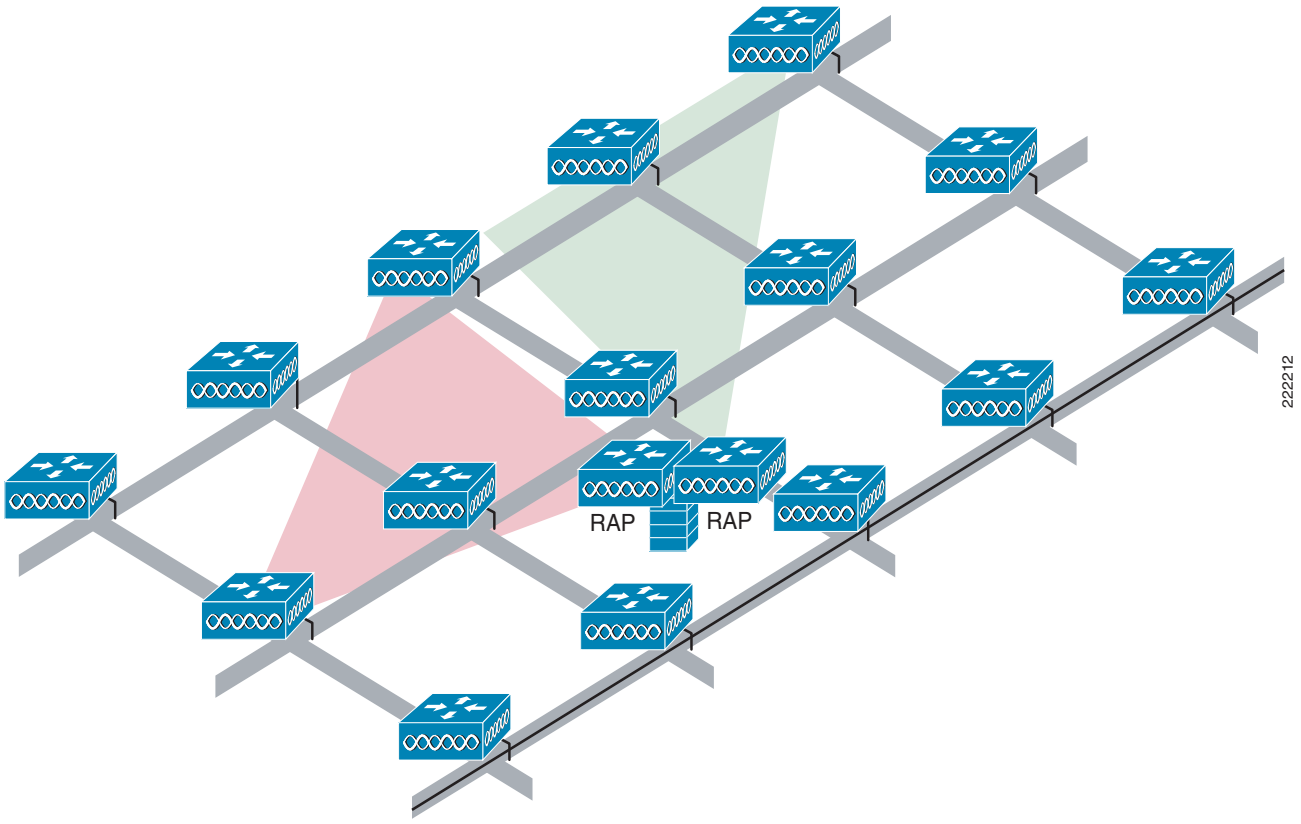

Multiple RAPs

Before deploying multiple RAPs, the purpose for deploying these RAPs needs to be considered. If additional RAPs are being considered to provide hardware diversity, they should be deployed on the same channel as the primary RAP. The reason for this is to minimize the convergence time in a scenario where the mesh transfers from one RAP to another. When planning RAP hardware diversity, the 32 MAPs per RAP limitation should be remembered.

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

When adding a second RAP on a different channel, channel planning or RAP cell splitting can be used to reduce the extent of potential collision domains. Channel planning allocates different non-overlapping channels to RAPs 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 are collocated but operate on different frequency channels, thus dividing a large collision domain into several smaller ones that operate independently.

If the Wireless Mesh 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 can increase in the event of a fail over as the MAP has to failover to another RAP on a different subnet and DHCP for an appropriate IP address. One way to limit this from happening is to use different BGN for segments in your network that are separated by subnet boundaries. In segmenting in this manner, MAPs do not associate with a RAP on a different subnet and you avoid slow convergence issues and the expense of the higher availability offered by the additional RAPs.

Multiple Controllers

There are operational advantages to centralizing WLCs and these advantages need to be traded off against the speed and capacity of the links to the LWAPP APs and the traffic profile of the WLAN clients using these 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 (see Figure 8-18)

Figure 8-18 Centralized Controllers

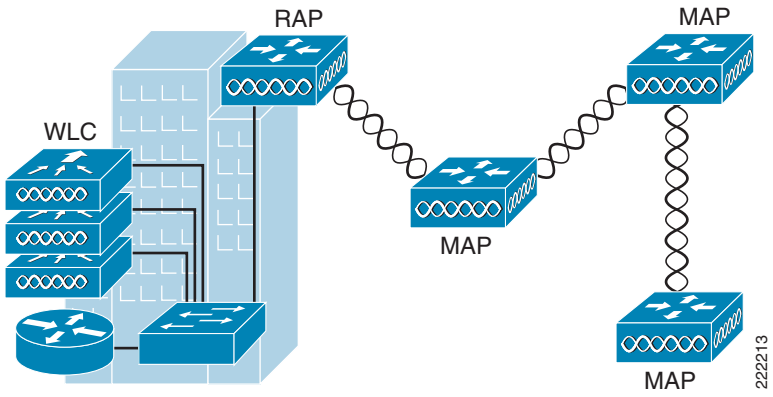

If the WLAN client traffic is predominantly peer-to-peer, a distributed controller model as shown in Figure 8-19 may be a better fit. In such cases, it may be that a majority of the WLAN traffic are clients in the area, with a smaller amount of traffic going to other locations. Given that many peer-to-peer applications can be sensitive to delay and packet loss, it is best to ensure that traffic between peers takes the most efficient path.

Figure 8-19 Distributed Controllers

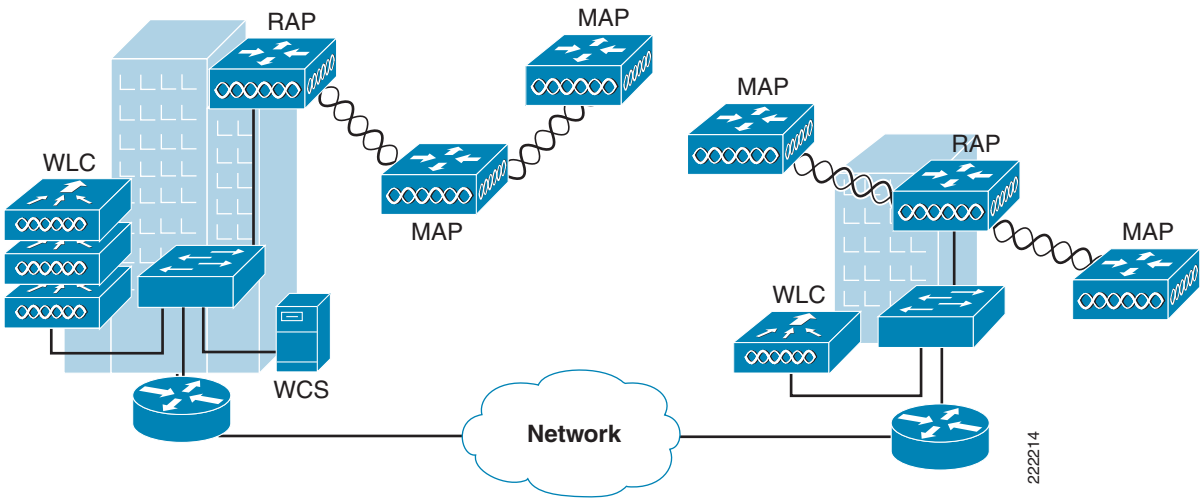

Given that most deployments see a mix of client server traffic and peer-to peer traffic, it is likely that a hybrid model of WLC placement is used, where points of presence (PoPs) are created with clusters of controllers placed in strategic locations in the network.

In all cases, remember that the LWAPP model used in the wireless mesh network is designed for a high-speed, low-latency network between the LWAPP APs (RAPs and MAPs in a Wireless Mesh) and the Wireless LAN Controller.

Multiple Wireless Mesh Mobility Groups

A wireless mesh WLAN coverage is not limited by the maximum number of controllers allowed in a mobility group. The WLANs that are part of a mobility group can be replicated in another mobility group and a WLAN client is able to roam between these mobility groups.

When roaming between mobility groups, the roaming may occur at Layer 2 or Layer 3, depending on the network topology behind the wireless mesh networks. When a Layer 3 roam occurs between mobility groups, mobility tunneling does not occur. Because of this the client must request a new DHCP address and will experience a session interruption.

Design Example

This section describes a design example of WLAN coverage in an urban/suburban area. It is important to understand cell size limitations and channel spacing for proper coverage. The following example explains how to make these preparations for a mesh deployment.

MAP Density and Distance

In cell planning there are two distances that should be considered, one in the typical backhaul radius as the other distance is the typical 2.4 client access radius. If you are simply building out a mesh to back haul data, then the 1000 foot radius is your limit. However if you are trying to provide complete client coverage, adhere to the 600 foot radius limit.

If you are designing your mesh deployment around the backhaul and do not intend to provide seamless WLAN coverage, you can have a typical cell size radius of 1000 feet. One square mile is 27,878,400 square feet; in this case, the approximate number of MAPs to cover a square mile given some cell overlap is nine and you can cover one square mile with approximately three or four hops (see Figure 8-20).

Figure 8-20 1000 Foot Distance Example

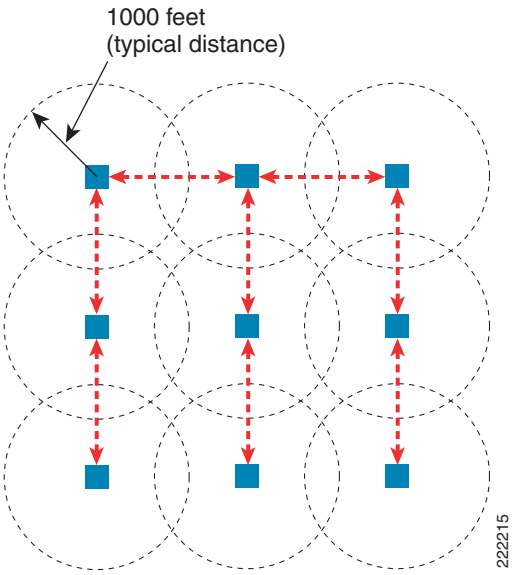

For deployments that provide seamless WLAN coverage it is suggested to have an approximate cell radius of 600 feet. One cell size comes out to be 1,130,973 square feet, so the approximate number of cells given some cell overlap is 25 cells per square mile (see Figure 8-21).

Figure 8-21 600 Foot Distance Example

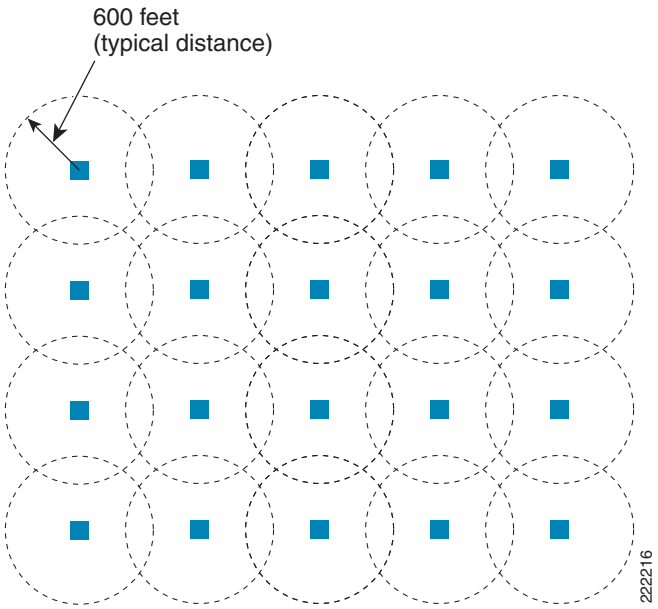

When finding a location for a RAP the goal is to use the RAP location in combination with RF antenna design to ensure that there is a good RF link to the MAPs within the core of the cell.

This means that the physical location of a RAP may be on the edge of the cell and a directional antenna is used to establish a link into the center of the cell.

Figure 8-22 Schematic of the Wireless Mesh Layout

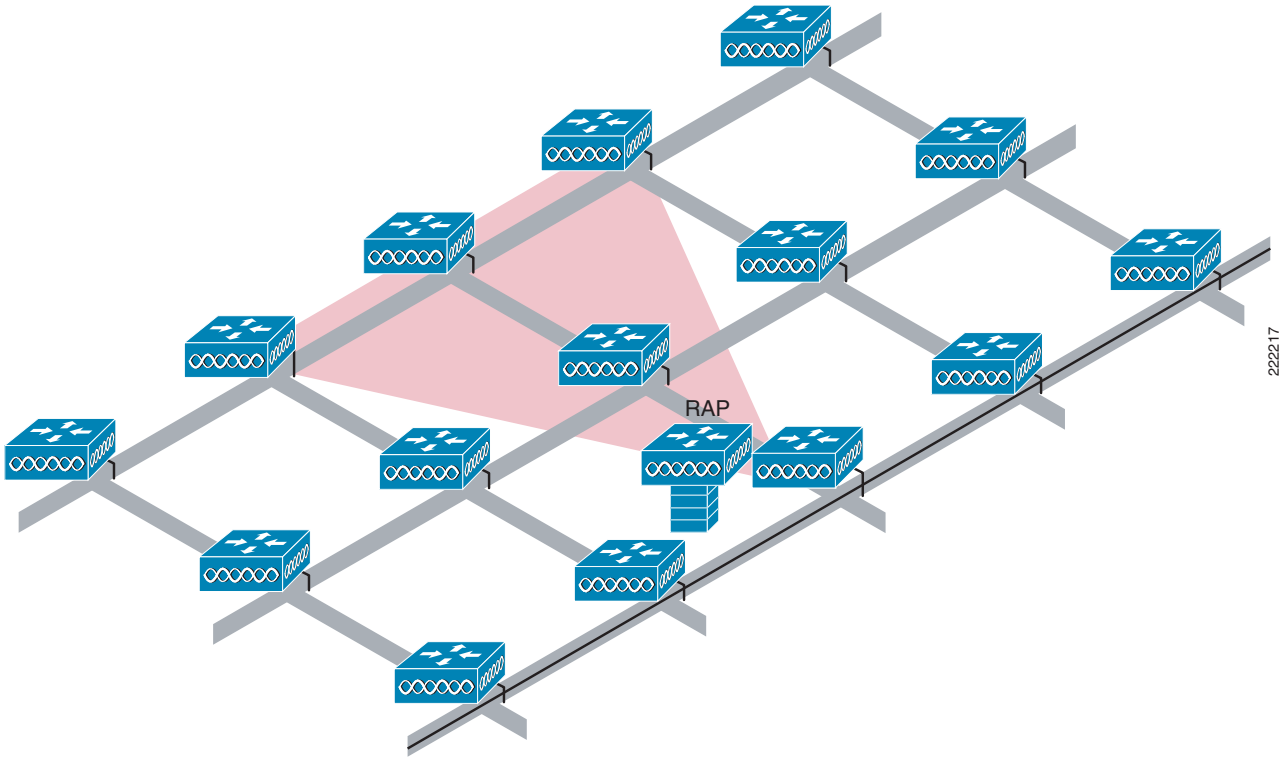

When laying out multiple cells, use channel planning similar to standard WLAN planning to avoid overlapping channels. As shown in Figure 8-23, both B2 and D3 share the same channel but do not overlap. This is the same for the other cells in the diagram that share the same channel. The cells sharing the same channels are:

- B2 and D3
- B3 and D2
- B4 and C2
- C3 and D4

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 Figure 8-24.

Figure 8-23 Laying out Various Cells

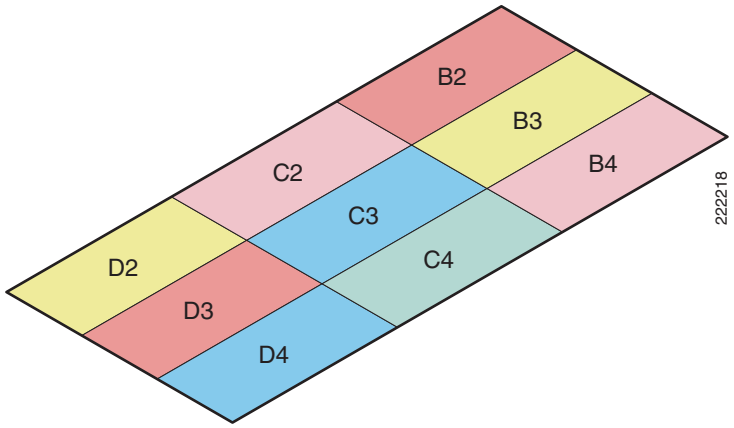

Figure 8-24 Failover Coverage

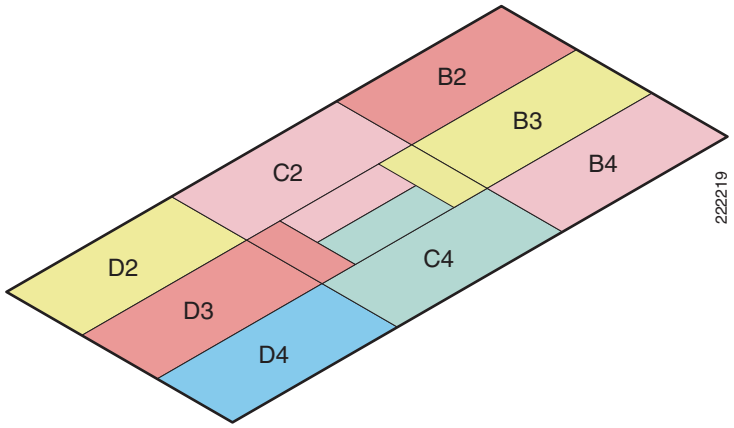

Connecting the Cisco 1500 Mesh AP to your Network

The wireless mesh has two locations where both bridged or WLAN client traffic terminate on the wired network. The first location is where the RAP attaches to the wired network, If bridging is enabled, this is where all bridged traffic connects to the wired network. The second location is where the WLC connects to the wired network; this is where WLAN client traffic from the mesh network connects to the wired network. This is shown schematically in Figure 8-25. The WLAN client traffic from the mesh is tunneled to the Wireless LAN Controller and then terminated on a VLAN to which the WLAN is assigned. 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 8-25 Mesh Network Traffic Termination

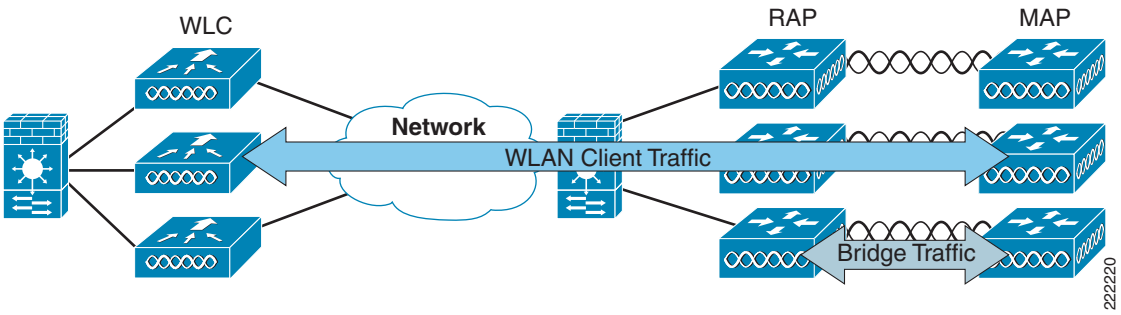

The connection to an Outdoor AP, unlike an indoor AP, may need to be firewalled from the wired network because the MAP that may be used for a bridging application and have limited security on the wired MAP ports.

Physical Placement of Mesh APs

When choosing a location for your MAPs, keep in mind issues like building height obstructions, light pole locations, and power options. In most environments there are light poles, but not all of them are equipped with an electric eye, which is a common feature used on light poles to automatically turn them on at night and off during the day. Street light power taps can be inserted between the light pole circuit and the electric eye to tap power from the street light. If a light pole does not have an electric eye, another method for powering the AP is required. Make note of what types of light poles you have and options for tapping power. When placing the roof top MAP, a directional antenna may be of use to direct coverage to a specific MAP or group of MAPs designated as the first hops into the mesh. If you plan to use omni directional antennas for the RAP, make sure to mount it towards the edge of the building so the radio coverage is not blocked. Figure 8-26 shows coverage concerns between the RAP and MAPs in the mesh.

Figure 8-26 AP Placement

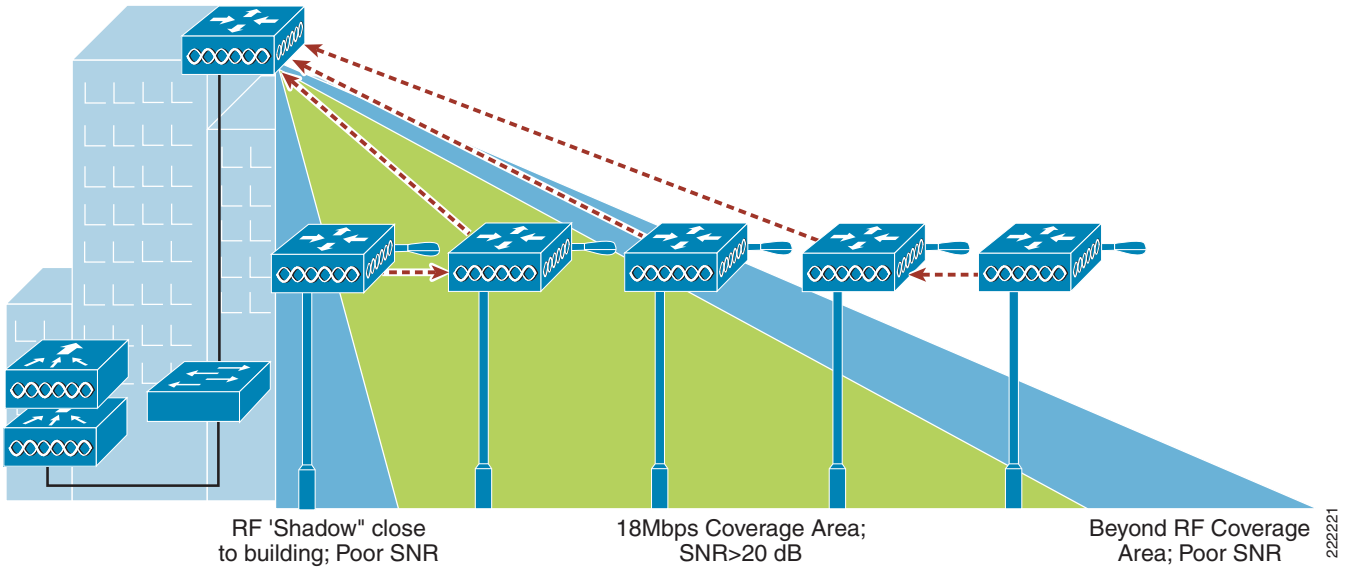

AP 1500 Alternate Deployment Options

The Cisco 1500 Series Mesh AP solution supports alternate deployment modes, including the following:

- WLAN backhaul
- Point-to-multipoint wireless bridging
- Point-to-point wireless bridging

These deployment methods can be useful for connecting LAN segments in a metropolitan environment or can be used to supplement backup connectivity for LAN segments. Client WLAN support can coexist in a bridged network configuration. Any of the following alternate deployment methods can simultaneously support mesh WLAN client traffic.

Wireless Backhaul

Cisco 1500 Mesh APs can provide a simple wireless backhaul solution where the 1500 Mesh AP is used to provide 802.11b/g services to WLAN and wired clients. This configuration is basically a wireless mesh with one MAP. Figure 8-27 shows an example of this deployment type.

Figure 8-27 Wireless Backhaul

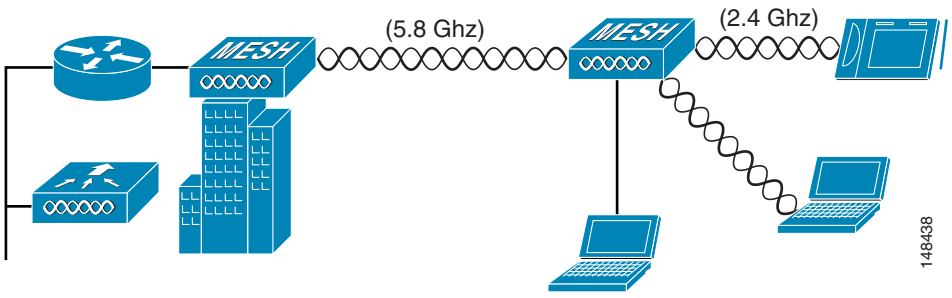

Point-to-Multipoint Wireless Bridging

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 turned off for all MAPs.

If Ethernet bridging is used, you must enable it on the controller for each MAP. Figure 8-28 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 may not be suitable for client access.

Figure 8-28 Point-to-Multipoint Wireless Bridging

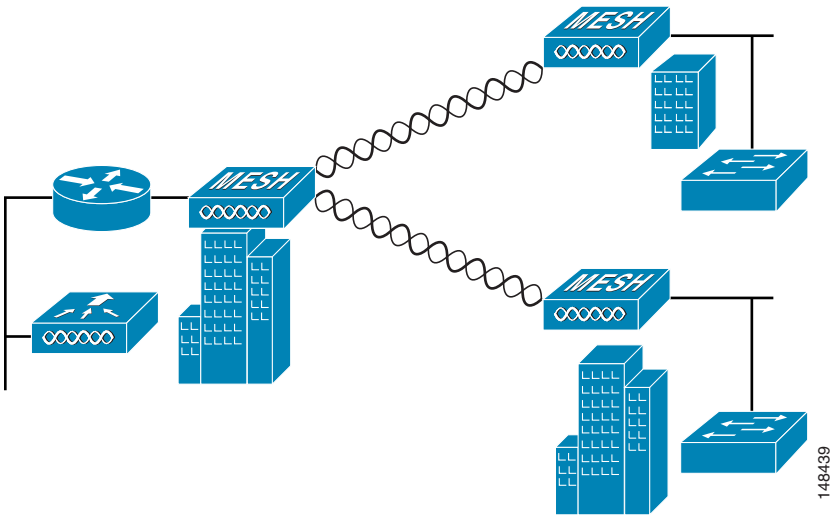

10.6.3 Point-to-Point Wireless Bridging

In a point-to-point bridging scenario, a 1500 mesh AP can be used to extend a Layer 2 network by using the backhaul radio to bridge two segments of a switched network, as shown in Figure 8-29. This is fundamentally a wireless mesh 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 may not be suitable for client access.

Figure 8-29 Point-to-Point Wireless Bridging