



Site Preparation and Planning

This chapter describes the site preparation and planning for your mesh network and contains the following sections:

- [Site Survey, on page 1](#)
- [Wireless Mesh Network Coverage Considerations, on page 8](#)
- [Indoor Mesh Interoperability with Outdoor Mesh, on page 31](#)

Site Survey

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



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

Pre-Survey Checklist

Before attempting a site survey, determine the following:

- How long is your wireless link?
- Do you have a clear line of sight?
- What is the minimum acceptable data rate within which the link runs?
- Is this a point-to-point or point-to-multipoint link?
- Do you have the correct antenna?
- Can the access point installation area support the weight of the access point?
- Do you have access to both of the mesh site locations?
- Do you have the proper permits, if required?

- Do you have a partner? Never attempt to survey or work alone on a roof or tower.
- Have you configured the 1500 series before you go onsite? It is always easier to resolve configuration or device problems first.
- Do you have the proper tools and equipment to complete your task?



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

Outdoor Site Survey

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

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

Design recommendations for mesh links are as follows:

- MAP deployment cannot exceed 35 feet in height above the street.
- MAPs are deployed with antennas pointed down toward the ground.
- Typical 5-GHz RAP-to-MAP distances are 1000 to 4000 feet.
- RAP locations are typically towers or tall buildings.
- Typical 5-GHz MAP-to-MAP distances are 500 to 1000 feet.
- MAP locations are typically short building tops or streetlights.
- Typical 2.4-GHz MAP-to-client distances are 500 to 1000 feet (depends upon the type of access point).
- Clients are typically laptops, Smart Phones, Tablets, and CPEs. Most of the clients operate in the 2.4-GHz band.
- In release 8.2 and above 2.4GHz radios can be used for backhaul and slightly longer distances can be achieved; however at the same time lower throughput is expected.

Determining a Line of Sight

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

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

This high-gain antenna configuration is recommended only for connecting a RAP to the MAP. To optimize mesh behavior, omnidirectional antennas are used because mesh links are limited to one mile (1.6 km). The

curvature of the earth does not impact line-of-sight calculations because the curvature of the earth changes every six miles (9.6 km).

Weather


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

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

Fresnel Zone

A Fresnel zone is an imaginary ellipse around the visual line of sight between the transmitter and receiver. As radio signals travel through free space to their intended target, they could encounter an obstruction in the Fresnel area, degrading the signal. Best performance and range are attained when there is no obstruction of this Fresnel area. Fresnel zone, free space loss, antenna gain, cable loss, data rate, link distance, transmitter power, receiver sensitivity, and other variables play a role in determining how far your mesh link goes. Links can still occur as long as 60 percent to 70 percent of the Fresnel area is unobstructed, as illustrated in [Figure 1: Point-to-Point Link Fresnel Zone](#), on page 3.

Figure 1: Point-to-Point Link Fresnel Zone



[Figure 2: Typical Obstructions in a Fresnel Zone](#), on page 3 illustrates an obstructed Fresnel zone.

Figure 2: Typical Obstructions in a Fresnel Zone



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

$$F1 = 72.6 \times \text{square root} (d/4 \times f)$$

where

F1 = the first Fresnel zone radius in feet

D = total path length in miles

F = frequency (GHz)


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

$$0.60 F1 = 43.3 \times \text{square root} (d/4 \times f)$$

These calculations are based on a flat terrain.

[Figure 3: Removing Obstructions in a Fresnel Zone](#), on page 4 shows the removal of an obstruction in the Fresnel zone of the wireless signal.

Figure 3: Removing Obstructions in a Fresnel Zone



Fresnel Zone Size in Wireless Mesh Deployments

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

Hidden Nodes Interference

The mesh backhaul uses the same 802.11a channel for all nodes in that mesh, which can introduce hidden nodes into the WLAN backhaul environment.

Figure 4: Hidden Nodes




Figure 4: Hidden Nodes, on page 5 shows the following three MAPs:

- MAP X
- MAP Y
- MAP Z

If MAP X is the route back to the RAP for MAP Y and Z, both MAP X and MAP Z might be sending traffic to MAP Y at the same time. MAP Y can see traffic from both MAP X and Z, but MAP X and Z cannot see each other because of the RF environment, which means that the carrier sense multi-access (CSMA) mechanism does not stop MAP X and Z from transmitting during the same time window; if either of these frames is destined for a MAP, it is corrupted by the collision between frames and requires retransmission.

Although all WLANs at some time can expect some hidden node collisions, the fixed nature of the MAP makes hidden node collisions a persistent feature of the mesh WLAN backhaul under some traffic conditions such as heavy loads and large packet streams.

Both the hidden node problem and the exposed node problem are inherent to wireless mesh networks because mesh access points share the same backhaul channel. Because these two problems can affect the overall network performance, the Cisco mesh solution seeks to mitigate these two problems as much as possible. For example, the AP1500s have at least two radios: one for backhaul access on a 5-GHz channel and the other for 2.4-GHz client access. In addition, the radio resource management (RRM) feature, which operates on the 2.4-GHz radio, enables cell breathing and automatic channel change, which can effectively decrease the collision domains in a mesh network.

There is an additional solution that can help to further mitigate these two problems. To reduce collisions and to improve stability under high load conditions, the 802.11 MAC uses an exponential backoff algorithm, where contending nodes back off exponentially and retransmit packets whenever a perceived collision occurs. Theoretically, the more retries a node has, the smaller the collision probability will be. In practice, when there are only two contending stations and they are not hidden stations, the collision probability becomes negligible after just three retries. The collision probability increases when there are more contending stations. Therefore, when there are many contending stations in the same collision domain, a higher retry limit and a larger maximum contention window are necessary. Further, collision probability does not decrease exponentially when there are hidden nodes in the network. In this case, an RTS/CTS exchange can be used to mitigate the hidden node problem.

Preferred Parent Selection

You can configure a preferred parent for a MAP. This feature gives more control to you and enables you to enforce a linear topology in a mesh environment. You can skip AWPP and force a parent to go to a preferred parent.

Preferred Parent Selection Criteria

The child AP selects the preferred parent based on the following criteria:

- The preferred parent is the best parent.
- The preferred parent has a link SNR of at least 20 dB (other parents, however good, are ignored).
- The preferred parent has a link SNR in the range of 12 dB and 20 dB, but no other parent is significantly better (that is, the SNR is more than 20 percent better). For an SNR lower than 12 dB, the configuration is ignored.
- The preferred parent is not in the blocked list.
- The preferred parent is not in silent mode because of dynamic frequency selection (DFS).
- The preferred parent is in the same bridge group name (BGN). If the configured preferred parent is not in the same BGN and no other parent is available, the child joins the parent AP using the default BGN.

Configuring a Preferred Parent

To configure a preferred parent, enter the following command:

```
(Cisco Controller) > config mesh parent preferred AP_name MAC
```

where:

- *AP_name* is the name of the child AP that you have to specify.
- *MAC* is the MAC address of the preferred parent that you have to specify.



Note When you configure a preferred parent, ensure that you specify the MAC address of the actual mesh neighbor for the desired parent. This MAC address is the base radio MAC address that has the letter f as the final character. For example, if the base radio MAC address is 00:24:13:0f:92:00, then you must specify 00:24:13:0f:92:0f as the preferred parent. This is the actual MAC address that is used for mesh neighbor relationships.

The following example shows how to configure the preferred parent for the MAP1SB access point, where 00:24:13:0f:92:00 is the preferred parent's MAC address:

```
(Cisco Controller) > config mesh parent preferred MAP1SB 00:24:13:0f:92:0f
```

To configure a preferred parent using the controller GUI, follow these steps:

1. Choose **Wireless > Access Points > AP_NAME > Mesh**.

2. Enter the MAC address of the preferred parent in the **Preferred Parent** text box.



Note To clear the Preferred Parent value, enter **none** in the Preferred Parent Text box.

3. Click **Apply**.



Note When the preferred parent is entered, no other mesh configurations can be made at the same time. You must apply the changes and wait for 90 seconds before other mesh changes can be made.

Related Commands

The following commands are related to preferred parent selection:

- To clear a configured parent, enter the following command:

```
(Cisco Controller) > config mesh parent preferred AP_name none
```

- To get information about the AP that is configured as the preferred parent of a child AP, enter the following command:

```
(Cisco Controller) > show ap config general AP_name
```

The following example shows how to get the configuration information for the MAP1SB access point, where 00:24:13:0f:92:00 is the MAC address of the preferred parent:

```
(Cisco Controller) > show ap config general MAP1
```

```
Cisco AP Identifier..... 9
Cisco AP Name..... MAP1
Country code..... US - United States
Regulatory Domain allowed by Country..... 802.11bg:-A 802.11a:-A
AP Country code..... US - United States
AP Regulatory Domain..... 802.11bg:-A 802.11a:-A
Switch Port Number ..... 1
MAC Address..... 12:12:12:12:12:12
IP Address Configuration..... DHCP
IP Address..... 209.165.200.225
IP NetMask..... 255.255.255.224
CAPWAP Path MTU..... 1485
Domain.....
Name Server.....
Telnet State..... Disabled
Ssh State..... Disabled
Cisco AP Location..... default location
Cisco AP Group Name..... default-group
Primary Cisco Switch Name..... 4404
Primary Cisco Switch IP Address..... 209.165.200.230
Secondary Cisco Switch Name.....
Secondary Cisco Switch IP Address..... Not Configured
Tertiary Cisco Switch Name..... 4404
```

```

Tertiary Cisco Switch IP Address..... 3.3.3.3
Administrative State ..... ADMIN_ENABLED
Operation State ..... REGISTERED
Mirroring Mode ..... Disabled
AP Mode ..... Local
Public Safety ..... Global: Disabled, Local: Disabled
AP subMode ..... WIPS
Remote AP Debug ..... Disabled
S/W Version ..... 5.1.0.0
Boot Version ..... 12.4.10.0
Mini IOS Version ..... 0.0.0.0
Stats Reporting Period ..... 180
LED State..... Enabled
PoE Pre-Standard Switch..... Enabled
PoE Power Injector MAC Addr..... Disabled
Power Type/Mode..... PoE/Low Power (degraded mode)
Number Of Slots..... 2
AP Model..... AIR-LAP1252AG-A-K9
IOS Version..... 12.4(10:0)
Reset Button..... Enabled
AP Serial Number..... serial_number
AP Certificate Type..... Manufacture Installed
Management Frame Protection Validation..... Enabled (Global MFP Disabled)
AP User Mode..... CUSTOMIZED
AP username..... maria
AP Dot1x User Mode..... Not Configured
AP Dot1x username..... Not Configured
Cisco AP system logging host..... 255.255.255.255
AP Up Time..... 4 days, 06 h 17 m 22 s
AP LWAPP Up Time..... 4 days, 06 h 15 m 00 s
Join Date and Time..... Mon Mar 3 06:19:47 2008

Ethernet Port Duplex..... Auto
Ethernet Port Speed..... Auto
AP Link Latency..... Enabled
  Current Delay..... 0 ms
  Maximum Delay..... 240 ms
  Minimum Delay..... 0 ms
  Last updated (based on AP Up Time)..... 4 days, 06 h 17 m 20 s
Rogue Detection..... Enabled
AP TCP MSS Adjust..... Disabled
Mesh preferred parent..... 00:24:13:0f:92:00

```

Co-Channel Interference

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

Wireless Mesh Network Coverage Considerations

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

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

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 1500 Series Access Points

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

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

- RAP-to-MAP ratio—Recommended maximum ratio is 20 MAPs per RAP.
- AP-to-AP distance—A spacing of no more than of 2000 feet (609.6 meters) between each mesh access point is recommended. When you extend the mesh network on the backhaul (no client access), use a cell radius of 1000 feet (304.8 meters).
- Hop count—Three to four hops.
 - One square mile in feet (52802), is nine cells and you can cover one square mile with approximately three or four hops (see [Figure 5: Cell Radius of 1000 Feet and Access Point Placement for Nonvoice Mesh Networks](#), on page 9 and [Figure 6: Path Loss Exponent 2.3 to 2.7](#), on page 10.)
- For 2.4 GHz, the local access cell size radius is 600 feet (182.88 meters). One cell size is around 1.310 x 106, so there are 25 cells per square mile. (See [Figure 7: Cell Radius of 600 Feet and Access Point Placement for Nonvoice Mesh Networks](#), on page 10 and [Figure 8: Path Loss Exponent 2.5 to 3.0](#), on page 11.)

Figure 5: Cell Radius of 1000 Feet and Access Point Placement for Nonvoice Mesh Networks




Figure 6: Path Loss Exponent 2.3 to 2.7




Figure 7: Cell Radius of 600 Feet and Access Point Placement for Nonvoice Mesh Networks





Figure 8: Path Loss Exponent 2.5 to 3.0



For the Cisco 1550 Series Access Points

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

The AP1500 series offers comparatively better range and capacity as it has the 802.11n functionality. It has advantages of ClientLink (Beamforming) in downstream, better receiver sensitivities because of MRC in upstream, multiple transmitter streams and a few other advantages of 802.11n such as channel combining and so on. The 1552 access points can provide comparatively larger and higher capacity cells.



Note Link budgets are different for different country domains. The discussion in this section takes into account the most widely distributed and large country domains: -A and -E.

Comparison of Link Budgets of AP1572 Series and AP1552 Series in 2.4- and 5-GHz Bands (-A Domain)

See [Table 1: Link Budget Comparison for the 2.4-GHz band in -A/-B Domain, on page 11.](#)

Table 1: Link Budget Comparison for the 2.4-GHz band in -A/-B Domain

Parameter	Cisco 1552 (-A domain)	Cisco 1532 (-A Domain)	Cisco 1562 (-A Domain)	Cisco 1572 (-B Domain)
Frequency Band	2412 – 2462 MHz	2412 – 2462 MHz	2412–2462 MHz	2412 – 2462 MHz
Air Interface	802.11b/g/n	802.11a/b/g/n/acW2	802.11a/b/g/n/acW2	802.11 a/b/g/n
Channel Bandwidth	20 MHz	20 MHz	—	20 MHz
No. of Tx Spatial Streams	2	3SS for 1562I, 2SS for 1562E/D models	3SS for 1562I, 2SS for 1562E/D models	3SS

Parameter	Cisco 1552 (-A domain)	Cisco 1532 (-A Domain)	Cisco 1562 (-A Domain)	Cisco 1572 (-B Domain)
PHY Data Rates	Up to 144 Mbps ¹	Up to 216 Mbps with 3SS 144 Mbps for 2SS	Up to 216 Mbps with 3SS 144 Mbps for 2SS	Up to 216 Mbps
Tx Power Conducted	28 dBm, Composite ²	29 dBm for 1562I 27 dBm for 1562E/D	29 dBm for 1562I 27 dBm for 1562E/D	30 dBm
Rx Sensitivity	-94 dBm at 6 Mbps -79 dBm at 54 Mbps -73 dBm at 150 Mbps	- 92 dBm at 6 Mbps - 76 dBm at 54 Mbps - 71 dBm at 216 Mbps	- 92 dBm at 6 Mbps - 76 dBm at 54 Mbps - 71 dBm at 216 Mbps	-93 dBm at 6 Mbps -81 dBm at 54 Mbps -76 dBm at 216 Mbps
No. of Receive Channels	3	3 or 2	3 or 2	4
Rx Diversity	MRC	MRC	MRC	MRC
Antenna Cable loss	0.5 dB, with external antenna	0.5 dB external antenna	0.5 dB external antenna	0.5 dB external antenna

¹ 40-MHz channel bonding in 2.4 GHz is not applicable. Therefore, the maximum data rate is 144 Mbps.

² Composite power is the power when we have two Tx streams enabled in AP1552.

For the 5-GHz band, see [Table 2: Link Budget Comparison for the 5-GHz band in -A/-B Domain, on page 12](#).

Table 2: Link Budget Comparison for the 5-GHz band in -A/-B Domain

Parameter	Cisco 1552 (-A Domain)	Cisco 1532 (-A Domain)	Cisco 1562 (-A/B Domain)	Cisco 1572 (-B Domain)
Frequency Band	5745 – 5825 MHz	5.180 – 5.240 GHz 5.260 – 5.320 GHz 5.500 – 5.560 GHz 5.680 – 5.720 GHz 5.745 – 5.825 GHz	5.180 – 5.240 GHz 5.260 – 5.320 GHz 5.500 – 5.560 GHz 5.680 – 5.720 GHz 5.745 – 5.825 GHz	5.180 – 5.240 GHz 5.260 – 5.320 GHz 5.500 – 5.560 GHz 5.680 – 5.720 GHz 5.745 – 5.825 GHz
Air Interface	802.11a/n	802.11a/b/g/n/acW2	802.11a/b/g/n/acW2	802.11a/n/ac
Channel Bandwidth	20 MHz, 40 MHz	20 MHz, 40 MHz, 80 MHz	20 MHz, 40 MHz, 80 MHz	20 MHz, 40 MHz, 80 MHz
No. of Tx Spatial Streams	2	2	3 or 2	3
PHY Data Rates	Up to 300 Mbps	Up to 300 Mbps	1.300 / 867 Mbps	Up to 1.3 Gbps

Parameter	Cisco 1552 (-A Domain)	Cisco 1532 (-A Domain)	Cisco 1562 (-A/B Domain)	Cisco 1572 (-B Domain)
Tx Power Conducted	28 dBm, Composite	27 dBm	29 or 27 dBm	30 dBm
Rx Sensitivity	-92 dBm at 6 Mbps -76 dBm at 54 Mbps -72 dBm at 300 Mbps	-94 dBm at 6 Mbps -80 dBm at 54 Mbps -65 dBm at 1300 Mbps	-94 dBm at 6 Mbps -80 dBm at 54 Mbps -65 dBm at 1300 Mbps	-92 dBm at 6 Mbps -80 dBm at 54 Mbps -60 dBm at 1300 Mbps

The 20-MHz channel bonding to form a 40-MHz channel is available in 5 GHz. Therefore, we can go up to a data rate of 300 Mbps.

As discussed in the previous section, Path Loss Exponents (PLE) and Link Budget windows work together. For a full clear path, PLE is 2.0. For AP to AP, there is comparatively more clearance than AP to client. For AP to AP, PLE can be taken as 2.3 because it can be assumed that the height of both APs is about 10 meters, which means a good line of sight (but without Fresnel zone clearance).

For AP to client, PLE should be greater than or equal to 2.5 because the client is only 1 meter high. Therefore, there will be less Fresnel zone clearance. This applies to both the 2.4-GHz and 5-GHz bands.

Let us consider AP to AP link budget in 5 GHz for -A domain because 5 GHz is used as a backhaul for mesh. We can take a legacy data rate of 9 Mbps to estimate the range.



Note This is the lowest data rate for outdoor 802.11n APs, which carries the Cisco's ClientLink (Beamforming for Legacy clients) advantage. It provides a gain of up to 4 dB in the downlink direction.


Assumptions for the Cisco Range Calculator

- The Range Calculator has been edited to stay within limitations for Tx power and EIRP under the listed regulatory domains. There may be cases where it exceeds the limitations. You must verify that the installation is within the laws of the location in which it is being installed.
- All antenna ports must be used for external antenna models for effective performance. Otherwise, range is significantly compromised.
- The Tx power is the total composite power of both Tx paths.
- Rx sensitivity is the composite sensitivity of all three Rx paths. That is, MRC is included.
- The Range Calculator assumes that ClientLink (Beamforming) is switched on.
- When you use the Range Calculator, available power levels change based upon the regulatory domain, the antenna (or antenna gain) selected, and the data rate selected. You must verify all parameters after making any parameter changes.
- You can select a different antenna than the two that are available by default. If you enter a high gain antenna and choose a power that goes over the EIRP limit, then you get a warning and the range equals 0.

- You can choose only the channels that the access point is certified for.
- You can only select only valid power levels.


The RAPs shown in [Figure 9: PoP with Multiple RAPs, on page 14](#) are simply a starting point. The goal is to use the RAP location in combination with the RF antenna design to ensure that there is a good RF link to the MAP within the core of the cell, which means that the physical location of the RAPs can be on the edge of the cell, and a directional antenna is used to establish a link into the center of the cell. Therefore, the wired network location of a RAP might play host to the RAP of multiple cells, as shown in [Figure 9: PoP with Multiple RAPs, on page 14](#).

Figure 9: PoP with Multiple RAPs



When the basic cell composition is settled, the cell can be replicated to cover a greater area. When replicating the cells, a decision needs to be made whether to use the same backhaul channel on all cells or to change backhaul channels with each cell. In the example shown in [Figure 10: Multiple RAP and MAP Cells, on page 14](#), various backhaul channels (B2, C2, and D2) per cell have been chosen to reduce the co-channel interference between cells.


Figure 10: Multiple RAP and MAP Cells



Choosing various channels reduces the co-channel interference at the cell boundaries, at the expense of faster mesh convergence, because MAPs must fall back to seek mode to find neighbors in adjacent cells. In areas of high-traffic density, co-channel interference has the highest impact, which is likely to be around the RAP. If RAPs are clustered in one location, a different channel strategy is likely to give optimal performance; if RAPs are dispersed among the cells, using the same channel is less likely to degrade performance.


When you lay out multiple cells, use channel planning similar to standard WLAN planning to avoid overlapping channels, as shown in [Figure 11: Laying out Various Cells](#), on page 15.

Figure 11: Laying out Various Cells



If possible, the channel planning should also minimize channel overlap in cases where the mesh has expanded to cover the loss of a RAP connection, as shown in [Figure 12: Failover Coverage](#), on page 15.

Figure 12: Failover Coverage



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 access points equipped with either 8 dBi omnidirectional 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 omnidirectional 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 access points equipped with either 8-dBi omnidirectional 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 omnidirectional antenna, then the minimum vertical separation is 2 feet (0.609 meters).

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

Special Considerations for Indoor Mesh Networks

Note these considerations for indoor mesh networks:

- For outdoors, voice is supported on a best-effort basis on a mesh infrastructure.
- Quality of Service (QoS) is supported on the local 2.4-GHz client access radio and on the 5-GHz.
- Cisco also supports static Call Admission Control (CAC) in CCXv4 clients, which provides CAC between the access point and the client.
- RAP-to-MAP ratio—The recommended ratio is 3 to 4 MAPs per RAP.
- AP-to-AP distance:
 - For 11n and 11ac mesh APs, a spacing of no more than 250 feet between each mesh AP with a cell radius of 125 feet is recommended.
- Hop count—For data, the maximum is 4 hops. No more than 2 hops is recommended for voice.
- RF considerations for client access on voice networks:
 - Coverage hole of 2 to 10 percent
 - Cell coverage overlap of 15 to 20 percent
 - Voice needs RSSI and SNR values that are at least 15 dB higher than data requirements
 - RSSI of -67 dBm for all data rates should be the goal for 11b/g/n and 11a/n
 - SNR should be 25 dB for the data rate used by client to connect to the AP
 - Packet error rate (PER) should be configured for a value of one percent or less

- Channel with the lowest utilization (CU) must be used
Check the CU when no traffic is running
- Radio resource manager (RRM) can be used to implement the recommended RSSI, PER, SNR, CU, cell coverage, and coverage hole settings on the 802.11b/g/n/ac radio.

Figure 13: Cell Radius of 100 Feet (30.4 meters) and Access Point Placement for Voice Mesh Networks





Figure 14: Cell Radius of 125 Feet (38 meters) and Access Point Placement for Indoor 11n Mesh Networks



Note Although you can use directional antenna and have an AP-to-AP distance longer than 250 feet (76.2 meters), for seamless roaming, we recommend that you have an AP-to-AP distance no more than 250 feet.

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, AP1570, AP 1560 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.

	Parent Loss Detection/ Keep Alive Timers	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	20.5*
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.

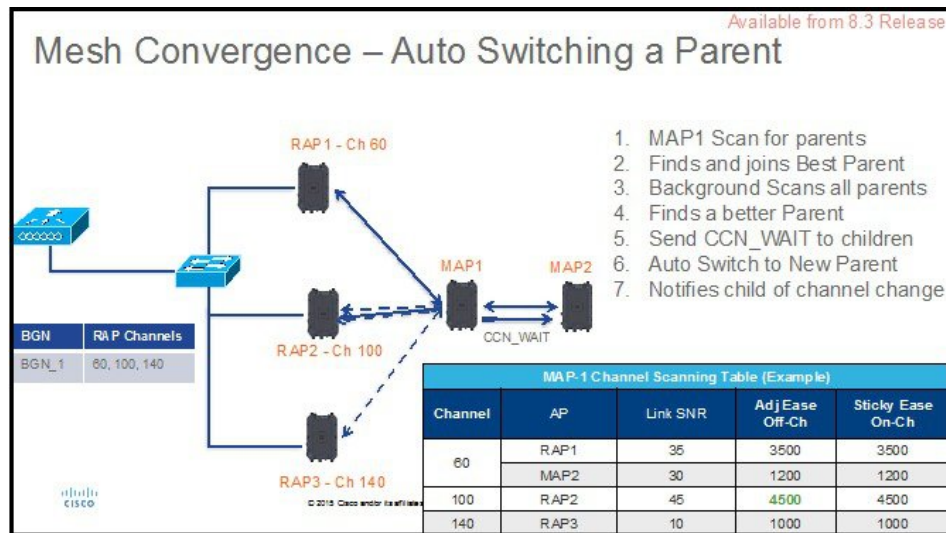
- To enable MAPs receive a packet during their off-channel scan over a DFS channel, all other on-channel DFS neighbors shall transmit the AWPP mesh beacons if there is no Tx/Rx in last 50 milliseconds
- These mesh beacons will help the MAPs which is doing the off-channel on the DFS channel to declare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



The illustration above shows a typical Off-Channel convergence process during the "standard" or "fast/very fast" configuration.



Note The timers in the table are for the illustration purpose only



The illustration below shows Mesh convergence and Parent Auto Switching when the "ease" value dictates switching to a better parent is prompted by new "ease value" even if the original parent is still available.

Configuring Mesh Convergence

Configuration steps are very simple to invoke new Background Scan feature.

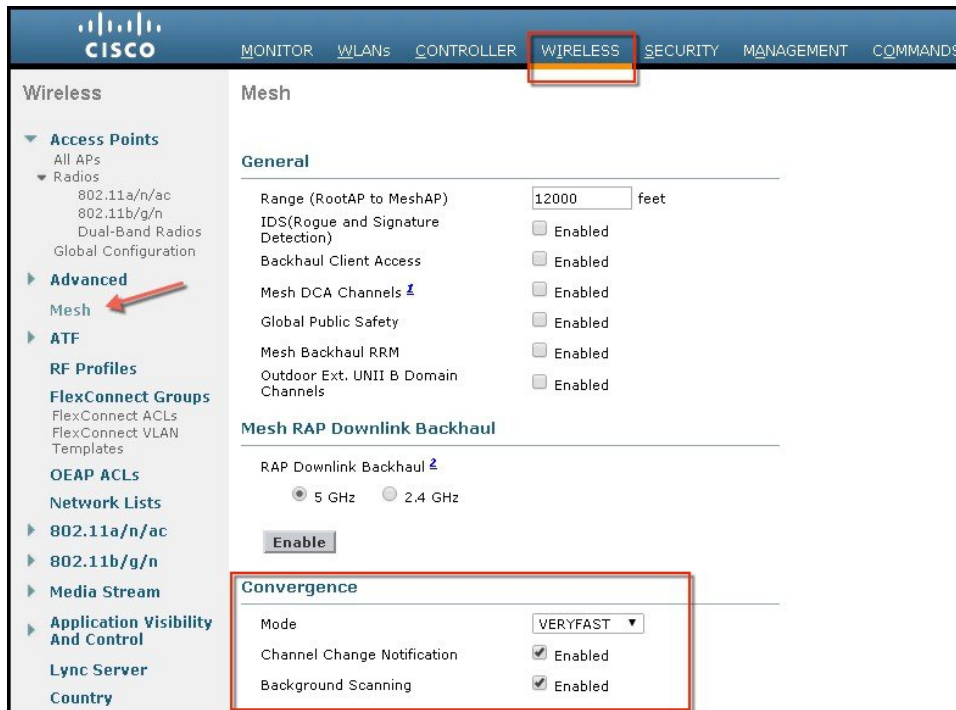
Perform the following steps to configure controller using GUI:

SUMMARY STEPS

1. On the controller choose **Wireless > Mesh** tabs and then under the "Convergence" section of the Mesh configuration choose Mode and then enable CCN and Background Scan.
2. Note, that under Mode there are three options to select the convergence mode, as indicated above the convergence times will change drastically depending on the mode selected.

DETAILED STEPS

Step 1 On the controller choose **Wireless > Mesh** tabs and then under the "Convergence" section of the Mesh configuration choose Mode and then enable CCN and Background Scan.



Background scan configure from CLI using the following commands:

```
(Cisco Controller) >config mesh background-scanning ?
enable          Enable background scanning on Mesh
disable         Disable background scanning on Mesh

(Cisco Controller) >config mesh background-scanning enable
```

Configure CCN using the following CLI command:

```
(Cisco Controller) >config mesh ccn ?
enable          Enables channel change notification
disable        Disables channel change notification

(Cisco Controller) >config mesh ccn enable
```

Step 2 Note, that under Mode there are three options to select the convergence mode, as indicated above the convergence times will change drastically depending on the mode selected.

Convergence

Mode	VERYFAST ▾	
Channel Change Notification	STANDARD	
Background Scanning	FAST	
	<input checked="" type="checkbox"/> VERYFAST	
	<input type="checkbox"/> Enabled	

From the CLI the same convergence configured using the following commands:

```
(Cisco Controller) >config mesh convergence ?
fast          Set fast convergence method
noise-tolerant fast Set noise-tolerant fast convergence method to handle unstable RF environment
standard      Set standard convergence method
very-fast     Set very fast convergence method

(Cisco Controller) >config mesh convergence very-fast all
```

Note In Standard mode CCN and BG scan options do not apply

Managing Mesh Features

There are several commands introduced to debug and troubleshoot convergence issues

Debug mesh convergence enable—debug trace

```
AP1572-7a7f.09c0#debug mesh ?
adjacency      MESH Adjacency debug
channel        Mesh Channel debug
convergence    Mesh convergence debug
error          Mesh error debug
ethernet       Mesh Ethernet debug
event          Mesh event debug
forwarding     Mesh forwarding debug
link           MESH Link debug
mperf         MESH BW test tool
node           Mesh node debug
port-control   Mesh port control debug
reliable       Mesh Reliable Delivery debug
security       MESH Security debug
trace          trace address
```

Debug mesh bgscan enable/disable

```
Cyprus_MAP1#debug mesh ?
adjacency      MESH Adjacency debug
bgscan         Mesh bgscan debug
channel        Mesh Channel debug
convergence    Mesh convergence debug
error          Mesh error debug
ethernet       Mesh Ethernet debug
event          Mesh event debug
forwarding     Mesh forwarding debug
link           MESH Link debug
mperf         MESH BW test tool
node           Mesh node debug
port-control   Mesh port control debug
reliable       Mesh Reliable Delivery debug
security       MESH Security debug
trace          trace address
```

Show mesh convergence—for state and counters

```

AP1572-7a7f.09c0#sh mesh ?
 adjacency      MESH Adjacency
 backhaul       MESH backhaul
 channel        MESH channel
 config         MESH config parameter
 convergence    MESH convergence info
 dfs           MESH dfs information
 ethernet       show mesh ethernet bridging
 forwarding     MESH Forwarding
 inventory      platform inventory
 linktest       MESH linktest stats
 lsc           MESH lsc details
 module         MESH module detail
 mperf         MESH BW tool
 security       MESH Security show
 simulation     MESH simulated configuration
 status        MESH status

```

Show mesh bgscan

```

Cyprus_MAP1#sh mesh ?
 adjacency      MESH Adjacency
 backhaul       MESH backhaul
 bgscan        MESH Background scanning info
 channel        MESH channel
 config         MESH config parameter
 convergence    MESH convergence info
 dfs           MESH dfs information
 ethernet       show mesh ethernet bridging
 forwarding     MESH Forwarding
 inventory      platform inventory
 linktest       MESH linktest stats
 lsc           MESH lsc details
 module         MESH module detail
 mperf         MESH BW tool
 security       MESH Security show
 simulation     MESH simulated configuration
 status        MESH status

```

```

Cyprus_MAP1#sh mesh bgscan
show MESH BG Scan

Background Scanning: Enabled

off channel Neighbors
-----
Channel:149 MissCnt:0
Mac:1c6a.7a7f.11ef MissCnt:0 NDRspCnt:72972 HopCnt:1 AdjustedEase:15448576
Flags: UPDATED NEIGH BEACON OCNEIGH

Channel:153 MissCnt:0
Mac:1c6a.7a7f.107f MissCnt:0 NDRspCnt:2579 HopCnt:1 AdjustedEase:17048576 StickyEase:21848576
Flags: UPDATED NEIGH PARENT BEACON
Mac:5835.d9aa.e46f MissCnt:0 NDRspCnt:0 HopCnt:0 AdjustedEase:0
Flags: BEACON
Mac:18e7.28aa.e87f MissCnt:0 NDRspCnt:0 HopCnt:0 AdjustedEase:0
Flags: UPDATED CHILD BEACON

Aligned Offchannel neighbors
-----
Channel:149 (POTENTIAL OFFCHANNEL)
Mac:1c6a.7a7f.11ef Ease:15448576

Channel:153 (ON-CHANNEL)
Mac:1c6a.7a7f.107f Ease:17048576

offChannel Requests Statistics
-----
Mac:18e7.28aa.e87f NDReqCnt:64 ch:149 last NDReq rx at: 10:54:21 UTC Mar 28 2016

Cyprus_MAP1#

```

Wireless Propagation Characteristics

[Table 3: Comparison of 2.4-GHz and 5-GHz Bands, on page 24](#) provides a comparison of the 2.4-GHz and 5-GHz bands.

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

Table 3: Comparison of 2.4-GHz and 5-GHz Bands

2.4-GHz Band Characteristics	5-GHz Band Characteristics
3 channels	22 channels (-A/-B regulatory domain)
More prone to co-channel interference	No co-channel interference
Lower power	Higher power
Lower SNR requirements given lower data rates	Higher SNR requirements given higher data rates
Better propagation characteristics than 5 GHz but more susceptible to noise and interference	Worse propagation characteristics than 2.4 GHz but less susceptible to noise and interference
Unlicensed band. Widely available throughout the world.	Not as widely available in the world as 2.4-GHz. Licenses in some countries.

2.4 GHz has more penetration capability across the obstacles due to a larger wavelength. In addition, 2.4 GHz has lower data rates which increases the success of the signal to reach the other end.

CleanAir

The 1550/1560/1570 series access points contain the CleanAir chipset, allowing full CleanAir support.

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

CleanAir AP Modes of Operation

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

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

Bridge mode access points support Radio Resource Management (RRM) on the 2.4 GHz band, which helps to mitigate the interference from WiFi interferers. RRM is only available on the 5 GHz band, if a Bridge mode RAP has no child MAPs.

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

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

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

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

Pseudo MAC (PMAC) and Merging

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

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

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

PMAC Convergence and Merging depends upon the following factors:

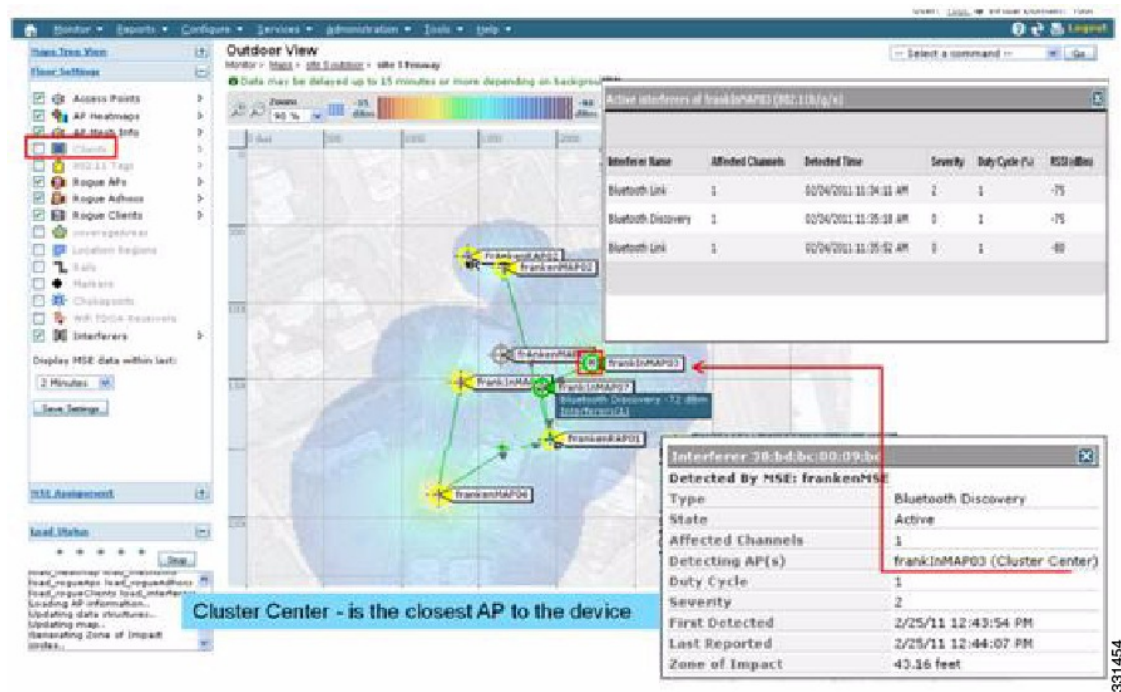
- Density of the sensors
- Quality of the observed classification
- RSSI from the interferer to the APs
- RF neighbor list at the APs

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

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

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

Figure 15: Pseudo MAC Merging in Outdoors



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

Event Driven Radio Resource Management and Persistence Device Avoidance

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



Note

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

CleanAir Access Point Deployment Recommendations

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

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



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

There are a lot of variables in the world of consumer electronics and unintentional electrical interference. Any expectation of accuracy that is derived from current Client or Tag location accuracy models does not apply to non-Wi-Fi location and CleanAir features.

Important notes to consider:

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

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

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

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

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

It is possible to deploy a few 1530 APs (non-CleanAir) with an installation that consists of 1552 APs and 1572 APs (CleanAir). This deployment can work from a client and coverage standpoint as these access points are fully interoperable with each other. The complete CleanAir functionality depends on all access points being CleanAir enabled. Detection can be affected, and mitigation is not recommended.

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

For more information about mixed systems, see http://www.cisco.com/en/US/products/ps10315/products_tech_note09186a0080b4bdc1.shtml

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 the 1552 access points in Bridge mode. In all other AP modes, the 5-GHz backhaul radio of the 1552 access points operates in CleanAir mode.

Enabling CleanAir

To enable CleanAir functionality in the system, you first need to enable CleanAir on the controller through **Wireless > 802.11a/b > CleanAir**. Although CleanAir is disabled by default, CleanAir is enabled by default on the AP interface.

After you enable CleanAir, it takes 15 minutes to propagate air quality information because the default reporting interval is 15 minutes. However, you can see the results instantly at the CleanAir detail level on the radio by going to **Monitor > Access Points > 802.11a/n or 802.11b/n**.

Licensing

A CleanAir system requires a CleanAir AP and a controller that is running release 7.0 or later releases. Adding the Cisco Prime Infrastructure allows the displays to be enhanced and additional information to be correlated within the system. Adding the MSE further enhances the available features and provides the history and location of specific interference devices. There is no additional license requirement for the CleanAir feature because the CleanAir AP is the license. Adding the Prime Infrastructure can be done with a basic license. Adding the MSE to the system requires a Prime Infrastructure Plus license and a context-aware license selection for the MSE.

For purposes of interference location with the MSE or CMX, each interference device counts as a location target in Context-Aware. One hundred Permanent Interferer licenses are embedded in the MSE. Interferer Licenses open as CleanAir APs are detected, in stages of five licenses per CleanAir AP. This process is applicable to AP1552/1562/1572. An Interference device is the same as a client or a tag from a license quantity standpoint. Only a small percentage of the available licenses are used because there should be far less interference devices than clients or tags to track. Users do have control over what types of interference devices to detect and located from the controller configuration menus.

Cisco context-aware licenses can be managed and limited by the class of target (client, tag, interference), which gives users complete control over how licenses are used.



Note Each interference device requires one context-aware service (CAS) license.

If you have too many Bluetooth devices, it is advisable to switch off the tracing of these devices because they might take up too many CAS licenses.

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 access points.

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

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

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

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

Increasing Mesh Availability

In the Cell Planning Distance section, a wireless mesh cell of one square mile was created and then built upon. This wireless mesh cell has similar properties to the cells used to create a cellular phone network because the smaller cells (rather than the defined maximum cell size) can be created to cover the same physical area, providing greater availability or capacity. This process is done by adding a RAP to the cell. Similar to the larger mesh deployment, the decision is whether to use RAP on the same channel, as shown in [Figure 16: Two RAPs per Cell with the Same Channel, on page 30](#), or to use RAPs placed on different channels, as shown in [Figure 17: Two RAPs per Cell on Different Channels, on page 30](#). The addition of RAPs into an area adds capacity and resilience to that area.

Figure 16: Two RAPs per Cell with the Same Channel

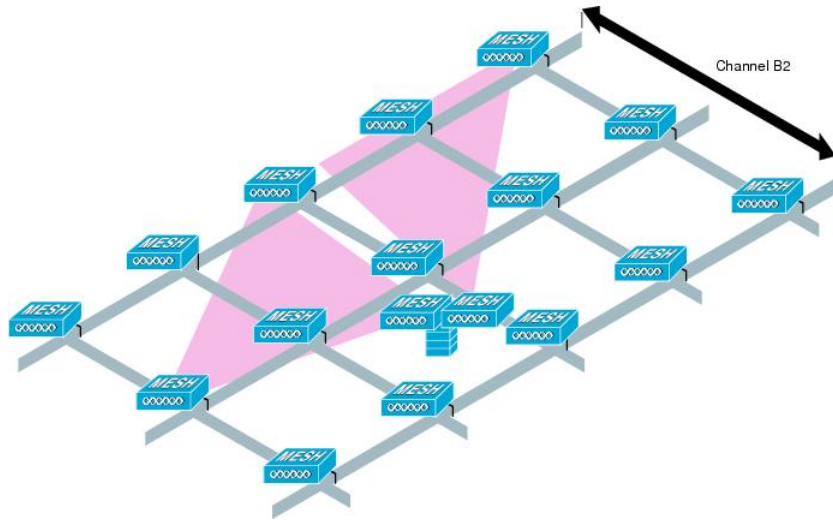



Figure 17: Two RAPs per Cell on Different Channels



Multiple RAPs

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

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

Adding a second RAP on a different channel also reduces the collision domain through channel planning or through RAP cell splitting. Channel planning allocates different nonoverlapping channels to mesh nodes in the same collision domain to minimize the collision probability. RAP cell splitting is a simple, yet effective, way to reduce the collision domain. Instead of deploying one RAP with omnidirectional antennas in a mesh

network, two or more RAPs with directional antennas can be deployed. These RAPs collocate with each other and operate on different frequency channels. This process divides a large collision domain into several smaller ones that operate independently.

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

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

Indoor Mesh Interoperability with Outdoor Mesh

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



Caution

The indoor access points 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 1700, 1800, 2600, 2700, 2800, 3500e/i, 3600, 3700 and 3800 series access points in an outdoor enclosure is recommended because of its robust environmental and temperature specifications. Additionally, the indoor access points 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 access point.

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 access points simultaneously. The same WLANs are broadcast out of both indoor and outdoor mesh access points.

