CHAPTER 20

Post-Deployment Radio Frequency Tuning

September 4, 2014

This chapter discusses post deployment RF tuning that should be done regularly on the deployment and includes using RRM for channel planning, CleanAir to mitigate RF interference, and a regular post-site survey assessment to ensure that optimum RF health is maintained.

Once a CMX solution is deployed, it is highly recommended that administrators turn on both Radio Resource Management and CleanAir to plan, optimize, and mitigate interference in the network.

Radio Resource Management

The Radio Resource Management (RRM) software embedded in the Cisco Wireless LAN Controller acts as a built-in RF engineer to consistently provide real-time RF management of your wireless network. RRM enables Cisco WLCs to continually monitor their associated lightweight access points for the following information:

- Traffic load—The total bandwidth used for transmitting and receiving traffic, which enables wireless LAN managers to track and plan network growth ahead of client demand.
- Interference—The amount of traffic coming from other 802.11 sources.
- Noise—The amount of non-802.11 traffic that is interfering with the currently assigned channel.
- Coverage—The received signal strength (RSSI) and signal-to-noise ratio (SNR) for all connected clients.
- Other—The number of nearby access points.

Using this information, RRM can periodically reconfigure the 802.11 RF network for best efficiency. To do this, RRM performs these functions:

- Radio resource monitoring
- Transmit power control
- Dynamic channel assignment
- Coverage hole detection and correction

RRM automatically detects and configures new Cisco WLCs and lightweight access points as they are added to the network. It then automatically adjusts associated and nearby lightweight access points to optimize coverage and capacity.

Lightweight access points can simultaneously scan all valid 802.11a/b/g/n/ac channels for the country of operation as well as for channels available in other locations. The access points go "off-channel" for a period not greater than 60 ms to monitor these channels for noise and interference. Packets collected during this time are analyzed to detect rogue access points, rogue clients, ad-hoc clients, and interfering access points.

Note In the presence of voice traffic (in the last 100 ms), the access points defer off-channel measurements.

Each access point spends only 0.2 percent of its time off-channel. This activity is distributed across all access points so that adjacent access points are not scanning at the same time, which could adversely affect wireless LAN performance.

Note When there are numerous rogue access points in the network, the chance of detecting rogues on channels 157 or 161 by a FlexConnect or local mode access point is small. In such cases, the monitor mode AP can be used for rogue detection.

Transmit Power Control

The Cisco WLC dynamically controls access point transmit power based on real-time wireless LAN conditions. You can choose between two versions of transmit power control: TPCv1 and TPCv2. With TPCv1, typically power can be kept low to gain extra capacity and reduce interference. With TPCv2, transmit power is dynamically adjusted with the goal of minimum interference. TPCv2 is suitable for dense networks. In this mode, there could be higher roaming delays and coverage hole incidents.

The Transmit Power Control (TPC) algorithm both increases and decreases an access point's power in response to changes in the RF environment. In most instances, TPC seeks to lower an access point's power to reduce interference, but in the case of a sudden change in the RF coverage—for example, if an access point fails or becomes disabled—TPC can also increase power on surrounding access points. This feature is different from coverage hole detection, which is primarily concerned with clients. TPC provides enough RF power to achieve desired coverage levels while avoiding channel interference between access points.

Table 20-1 shows the power level mapping.

Table 20-1 Power Level Table

Power Level	2.4 GHz	5 GHz
1	23 dBm (200 mW) CCK Only	20 dBm (100 mW)
2	20 dBm (100 mW)	17 dBm (50 mW)
3	17 dBm (50 mW)	14 dBm (25 mW)
4	14 dBm (25 mW)	11 dBm (12.5 mW)
5	11 dBm (12.5 mW)	8 dBm (6.25 mW)
6	8 dBm (6.25 mW)	5 dBm (3.13 mW)
7	5 dBm (3.13 mW)	2 dBm (1.56 mW)
8	2 dBm (1.56 mW)	-1 dBm (0.78 mW)
	-1 dBm (0.78 mW)	

Overriding the TPC Algorithm with Minimum and Maximum Transmit Power Settings

The TPC algorithm balances RF power in many diverse RF environments. However it is possible that automatic power control will not be able to resolve some scenarios in which an adequate RF design was not possible to implement due to architectural restrictions or site restrictions—for example, when all access points must be mounted in a central hallway, placing the access points close together, but requiring coverage out to the edge of the building.

In these scenarios, you can configure maximum and minimum transmit power limits to override TPC recommendations. The maximum and minimum TPC power settings apply to all access points through RF profiles in a RF network.

To set the Maximum Power Level Assignment and Minimum Power Level Assignment, enter the maximum and minimum transmit power used by RRM in the text boxes in the Tx Power Control page. The range for these parameters is -10 to 30 dBm. The minimum value cannot be greater than the maximum value; the maximum value cannot be less than the minimum value.

If you configure a maximum transmit power, RRM does not allow any access point attached to the controller to exceed this transmit power level (whether the power is set by RRM TPC or by coverage hole detection). For example, if you configure a maximum transmit power of 11 dBm, then no access point would transmit above 11 dBm, unless the access point is configured manually.

Dynamic Channel Assignment

Two adjacent access points on the same channel can cause either signal contention or signal collision. In a collision, data is not received by the access point. This functionality can become a problem, for example, when someone reading e-mail in a cafe affects the performance of the access point in a neighboring business. Even though these are completely separate networks, someone sending traffic to the cafe on channel 1 can disrupt communication in an enterprise using the same channel. Controllers can dynamically allocate access point channel assignments to avoid conflict and to increase capacity and performance. Channels are “reused” to avoid wasting scarce RF resources. In other words, channel 1 is allocated to a different access point far from the cafe, which is more effective than not using channel 1 altogether.

The controller’s Dynamic Channel Assignment (DCA) capabilities are also useful in minimizing adjacent channel interference between access points. For example, two overlapping channels in the 802.11b/g band, such as 1 and 2, cannot both simultaneously use 11/54 Mbps. By effectively reassigning channels, the controllerkeeps adjacent channels separated.

Note We recommend that you use only non-overlapping channels (1, 6, 11, and so on).

The controller examines a variety of real-time RF characteristics to efficiently handle channel assignments as follows:

- Access point received energy—The received signal strength measured between each access point and its nearby neighboring access points. Channels are optimized for the highest network capacity.
- Noise—Noise can limit signal quality at the client and access point. An increase in noise reduces the effective cell size and degrades user experience. By optimizing channels to avoid noise sources, the controller can optimize coverage while maintaining system capacity. If a channel is unusable due to excessive noise, that channel can be avoided.

- 802.11 Interference—Interference is any 802.11 traffic that is not part of your wireless LAN, including rogue access points and neighboring wireless networks. Lightweight access points constantly scan all channels looking for sources of interference. If the amount of 802.11 interference exceeds a predefined configurable threshold (the default is 10 percent), the access point sends an alert to the controller. Using the RRM algorithms, the controller may then dynamically rearrange channel assignments to increase system performance in the presence of the interference. Such an adjustment could result in adjacent lightweight access points being on the same channel, but this setup is preferable to having the access points remain on a channel that is unusable due to an interfering foreign access point. In addition, if other wireless networks are present, the controller shifts the usage of channels to complement the other networks. For example, if one network is on channel 6, an adjacent wireless LAN is assigned to channel 1 or 11. This arrangement increases the capacity of the network by limiting the sharing of frequencies. If a channel has virtually no capacity remaining, the controller may choose to avoid this channel. In very dense deployments in which all nonoverlapping channels are occupied, the controller does its best, but you must consider RF density when setting expectations.
- Load and utilization—When utilization monitoring is enabled, capacity calculations can consider that some access points are deployed in ways that carry more traffic than other access points (for example, a lobby versus an engineering area). The controller can then assign channels to improve the access point with the worst performance reported. The load is taken into account when changing the channel structure to minimize the impact on clients currently in the wireless LAN. This metric keeps track of every access point's transmitted and received packet counts to determine how busy the access points are. New clients avoid an overloaded access point and associate to a new access point. This parameter is disabled by default.

The controller combines this RF characteristic information with RRM algorithms to make system-wide decisions. Conflicting demands are resolved using soft-decision metrics that guarantee the best choice for minimizing network interference. The end result is optimal channel configuration in a three-dimensional space, where access points on the floor above and below play a major factor in an overall wireless LAN configuration.

Note Radios using 40-MHz channels in the 2.4-GHz band or are not supported by DCA and cannot be configured.

The RRM startup mode is invoked in the following conditions:

- In a single-controller environment, the RRM startup mode is invoked after the controller is rebooted.
- In a multiple-controller environment, the RRM startup mode is invoked after an RF Group leader is elected.

You can trigger RRM startup mode from CLI.

RRM startup mode runs for 100 minutes (10 iterations at 10-minute intervals). The duration of the RRM startup mode is independent of the DCA interval, sensitivity, and network size. The startup mode consists of 10 DCA runs with high sensitivity (making channel changes easy and sensitive to the environment) to converge to a steady state channel plan. After the startup mode is finished, DCA continues to run at the specified interval and sensitivity.

Coverage Hole Detection and Correction

The RRM coverage hole detection algorithm can detect areas of radio coverage in a wireless LAN that are below the level needed for robust radio performance. This feature can alert you to the need for an additional (or relocated) lightweight access point.

If clients on a lightweight access point are detected at threshold levels (RSSI, failed client count, percentage of failed packets, and number of failed packets) lower than those specified in the RRM configuration, the access point sends a “coverage hole” alert to the controller. The alert indicates the existence of an area where clients are continually experiencing poor signal coverage without having a viable access point to which to roam. The controller discriminates between coverage holes that can and cannot be corrected. For coverage holes that can be corrected, the controller mitigates the coverage hole by increasing the transmit power level for that specific access point. The controller does not mitigate coverage holes caused by clients that are unable to increase their transmit power or are statically set to a power level because increasing their downstream transmit power might increase interference in the network.

Benefits of RRM

RRM produces a network with optimal capacity, performance, and reliability. It frees you from having to continually monitor the network for noise and interference problems, which can be transient and difficult to troubleshoot. RRM ensures that clients enjoy a seamless, trouble-free connection throughout the Cisco unified wireless network.

RRM uses separate monitoring and control for each deployed network: 802.11an/ac and 802.11bgn. The RRM algorithms run separately for each radio type (802.11an/ac and 802.11b/g). RRM uses both measurements and algorithms. RRM measurements can be adjusted using monitor intervals, but they cannot be disabled. RRM algorithms are enabled automatically but can be disabled by statically configuring channel and power assignment. The RRM algorithms run at a specified updated interval, which is 600 seconds by default.

For more detailed discussion and configuration of RRM, refer to:
http://www.cisco.com/c/en/us/td/docs/wireless/controller/7-6/configuration-guide/b_cg76/b_cg76_chapter_01111110.html.

CleanAir

Cisco CleanAir is a spectrum intelligence solution designed to proactively manage the challenges of a shared wireless spectrum. It allows you to see all of the users of the shared spectrum (both native devices and foreign interferers). It also enables you or your network to act upon this information. For example, you could manually remove the interfering device or the system could automatically change the channel away from the interference. CleanAir provides spectrum management and RF visibility.

A Cisco CleanAir system consists of CleanAir-enabled access points, Cisco Wireless LAN Controllers, and Cisco Prime Infrastructure. These access points collect information about all devices that operate in the industrial, scientific, and medical (ISM) bands, identify and evaluate the information as a potential interference source, and forward it to the Cisco WLC. The Cisco WLC controls the access points, collects spectrum data, and forwards information to Cisco Prime Infrastructure or a Cisco mobility services engine (MSE) upon request.

For every device operating in the unlicensed band, Cisco CleanAir tells you what it is, where it is, how it is impacting your wireless network, and what actions you or your network should take. It simplifies RF so that you do not have to be an RF expert.

Wireless LAN systems operate in unlicensed 2.4- and 5-GHz ISM bands. Many devices, such as microwave ovens, cordless phones, and Bluetooth devices also operate in these bands and can negatively affect Wi-Fi operations.

Some of the most advanced WLAN services, such as voice over wireless and IEEE 802.11n radio communications, could be significantly impaired by the interference caused by other legal users of the ISM bands. The integration of Cisco CleanAir functionality into the Cisco Unified Wireless Network addresses this problem of radio frequency (RF) interference.

CleanAir is supported on mesh AP backhaul at a 5-GHz radio of mesh. You can enable CleanAir on backhaul radios and can provide report interference details and air quality.

Role of the Cisco Wireless LAN Controller in a Cisco CleanAir System

The Cisco WLC performs the following tasks in a Cisco CleanAir system:

- Configures Cisco CleanAir capabilities on the access point.
- Provides interfaces (GUI, CLI, and SNMP) for configuring Cisco CleanAir features and retrieving data
- Displays spectrum data.
- Collects and processes air quality reports from the access point and stores them in the air quality database. The Air Quality Report (AQR) contains information about the total interference from all identified sources represented by the Air Quality Index (AQI) and summary for the most severe interference categories. The CleanAir system can also include unclassified interference information under per interference type reports, which enables you to take action in cases where the interference due to unclassified interfering devices is more.
- Collects and processes interference device reports (IDRs) from the access point and stores them in the interference device database.
- Forwards spectrum data to Prime Infrastructure and the MSE.

Interference Types that Cisco CleanAir Can Detect

Cisco CleanAir can detect interference, report on the location and severity of the interference, and recommend different mitigation strategies. Two such mitigation strategies are persistent device avoidance and spectrum event-driven RRM.

Wi-Fi chip-based RF management systems share these characteristics:

- Any RF energy that cannot be identified as a Wi-Fi signal is reported as noise.
- Noise measurements that are used to assign a channel plan tend to be averaged over a period of time to avoid instability or rapid changes that can be disruptive to certain client devices.
- Averaging measurements reduces the resolution of the measurement. As such, a signal that disrupts clients might not look like it needs to be mitigated after averaging.
- All RF management systems available today are reactive in nature.

Cisco CleanAir is different and can positively identify not only the source of the noise but also its location and potential impact to a WLAN. Having this information allows you to consider the noise within the context of the network and make intelligent and, where possible, proactive decisions. For CleanAir, two types of interference events are common:

- Persistent interference
- Spontaneous interference

Persistent interference events are created by devices that are stationary in nature and have intermittent but largely repeatable patterns of interference. For example, consider the case of a microwave oven located in a break room. Such a device might be active for only 1 or 2 minutes at a time. When operating, however, it can be disruptive to the performance of the wireless network and associated clients. Using Cisco CleanAir, you can positively identify the device as a microwave oven rather than indiscriminate noise. You can also determine exactly which part of the band is affected by the device and because you can locate it, you can understand which access points are most severely affected. You can then use this information to direct RRM in selecting a channel plan that avoids this source of interference for the access points within its range. Because this interference is not active for a large portion of the day, existing RF management applications might attempt to again change the channels of the affected access points. Persistent device avoidance is unique, however, in that it remains in effect as long as the source of interference is periodically detected to refresh the persistent status. The Cisco CleanAir system knows that the microwave oven exists and includes it in all future planning. If you move either the microwave oven or the surrounding access points, the algorithm updates RRM automatically.

Note Spectrum event-driven RRM can be triggered only by Cisco CleanAir-enabled access points in local mode.

Spontaneous interference is interference that appears suddenly on a network, perhaps jamming a channel or a range of channels completely. The Cisco CleanAir spectrum event-driven RRM feature allows you to set a threshold for air quality (AQ) that, if exceeded, triggers an immediate channel change for the affected access point. Most RF management systems can avoid interference, but this information takes time to propagate through the system. Cisco CleanAir relies on AQ measurements to continuously evaluate the spectrum and can trigger a move within 30 seconds. For example, if an access point detects interference from a video camera, it can recover by changing channels within 30 seconds of the camera becoming active. Cisco CleanAir also identifies and locates the source of interference so that more permanent mitigation of the device can be performed at a later time.

In the case of Bluetooth devices, Cisco CleanAir-enabled access points can detect and report interferences only if the devices are actively transmitting. Bluetooth devices have extensive power save modes. For example, interference can be detected when data or voice is being streamed between the connected devices.

Persistent Devices

Some interference devices such as outdoor bridges and microwave ovens only transmit when needed. These devices can cause significant interference to the local WLAN due to short duration and periodic operation remain largely undetected by normal RF management metrics. With CleanAir the RRM DCA algorithm can detect, measure, register, and remember the impact and adjust the DCA algorithm. This minimizes the use of channels affected by the persistent devices in the channel plan local to the interference source. Cisco CleanAir detects and stores the persistent device information in the Cisco WLC and this information is used to mitigate interfering channels.

Persistent Devices Detection

CleanAir-capable Monitor Mode access point collects information about persistent devices on all configured channels and stores the information in the Cisco WLC. Local/Bridge mode AP detects interference devices on the serving channels only.

Persistent Devices Propagation

Persistent device information that is detected by local or monitor mode access points is propagated to the neighboring access points connected to the same Cisco WLC to provide better chance of handling and avoiding persistent devices. Persistent device detected by the CleanAir-enabled access point is propagated to neighboring non-CleanAir access points, thus enhancing channel selection quality.

Detecting Interferers by an Access Point

When a CleanAir-enabled access point detects interference devices, detections of the same device from multiple sensors are merged together to create clusters. Each cluster is given a unique ID. Some devices conserve power by limiting the transmit time until actually needed, which results in the spectrum sensor to temporarily stop detecting the device. This device is then correctly marked as down. A down device is correctly removed from the spectrum database. In cases when all the interferer detections for a specific devices are reported, the cluster ID is kept alive for an extended period of time to prevent possible device detection bouncing. If the same device is detected again, it is merged with the original cluster ID and the device detection history is preserved.

For example, some Bluetooth headsets operate on battery power. These devices employ methods to reduce power consumption, such as turning off the transmitter when not actually needed. Such devices can appear to come and go from the classification. To manage these devices, CleanAir keeps the cluster IDs longer and they are remerged into a single record upon detection. This process smoothens the user records and accurately represents the device history.

For more detailed discussion CleanAir and configuration of the same, refer to: http://www.cisco.com/c/en/us/td/docs/wireless/controller/7-6/configuration-guide/b_cg76/b_cg76_chapter_010000011.html.

Post-Deployment RF Tuning

In addition to having both RRM and CleanAir turned on, it is recommended that regular post-deployment RF Site Survey is performed. A post-deployment RF Site Survey ensures that spectrum is being effectively managed by the CMX solution and there are no surprises. Just like the Pre-Deployment RF Site Survey, a Post-Deployment Site Survey should look at the following:

Location Assessment

- Assess building type and materials used—Has the building under gone renovation with newer material used or newer obstacles erected?
- Areas where full coverage and full performance is needed—Have coverage zones changed? Zones that were thinly populated with APs before because of no coverage requirement may now be required to have full coverage.
- Areas where RF free zones exist—Are there new RF free zones that did not exist before?
- Obtain location maps for RF surveys later—Have any maps been updated with newer construction or layouts.
- High Density—Has this location become a high density location that was not anticipated? In this case a reevaluation of location should be performed.

Business Needs of WLAN

- Critical Services—Are the services still the same? Is the network able to guarantee the same performance as before? Features like AVC/QOS should be explored.
- Long term Network View—Planning for future needs is always better than planning for present. With explosion of smartphones and mobile devices in markets, more and more devices are constantly on the Wi-Fi network than ever before. Internet of Everything is enabling machine-to-machine and machine-to-network communication via Wi-Fi. It is not just enough to plan for smartphones and laptops alone, but also for potential IoT devices that may use Wi-Fi.
- Guarding against potential RF explosion—Have number of devices gone up way more than anticipated? In this case better (and more) APs may be required leading to a high density design later.
- 802.11n/ac readiness and expectations—Is the customer looking at upgrading their wireless infrastructure? A new network should be re planned with all new objectives in mind, which may lead to a different RF plan than before.

Constraints on Deployment

- DFS and radar avoidance requirement—Are there any new radar and DFS requirement that did not exist before.
- Aesthetic design requirements—Aesthetics requirements may change over time. Administrators should always keep a watch on whether aesthetic requirements over time have changed AP location and antenna direction.
- Are there newer channels available—The FCC makes newer spectrum available time to time for WLAN to use. Can these enabled if the software supports it?

Existing 802.11 Surveys

It is a good idea to capture existing 802.11 state of the location after CMX is deployed and at regular intervals. Because Wi-Fi and other technologies are hugely prevalent today, it is a good investment of time to have a pre-study of existing RF done. Use tools like Metageek site survey to record existing RF at the location. In a big location it is advisable to move through the entire location and take a RF base reading.

- Plan for persistent non-movable interferers—Are there newer persistent interferers in the location that were not there before (microwaves, video cameras, etc.)
- Capture current state of Wi-Fi network—How many SSIDs exist already? What is their signal strength? Are there potential Wi-Fi networks around the location that are strong enough to interferer with a CMX deployment? What can be done to mitigate such potential sources of interference? Can power be increased on the APs?
- Antenna Evaluation—Evaluate use external antennas instead of internal antennas to ensure good coverage based on above parameters. Will use of external antenna be an aesthetic constraint? Can directional antennas be used in corner of the building to direct more coverage into building rather than outside of it? In a large environment this may be useful.—