



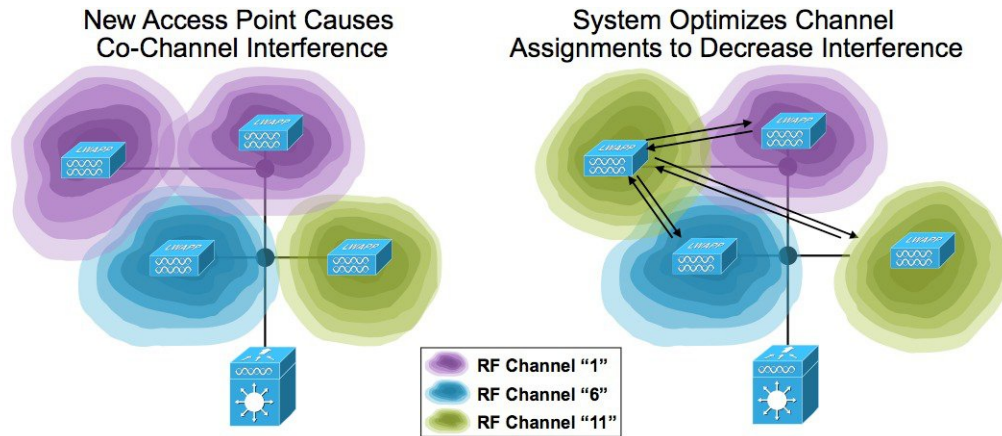
Dynamic Channel Assignment (DCA)

- [What does Dynamic Channel Assignment do?, on page 1](#)
- [The Dynamic Channel Assignment \(DCA\) Algorithm, on page 2](#)
- [DCA in a Nutshell, on page 3](#)
- [DCA Modes of Operation, on page 5](#)
- [DCA 20, 40, 80, and 160 MHz support , on page 7](#)
- [Dynamic Bandwidth Selection–DBS , on page 10](#)
- [Device Aware RRM , on page 12](#)

What does Dynamic Channel Assignment do?

- Dynamically manages channel assignments for an RF group.
- Evaluates the assignments on a per AP per radio basis
- Makes decisions using an RSSI based cost metric function which evaluates performance based on interference for each available channel
- Dynamically adjusts the channel plan to maintain performance of individual radios
- Actively manages 20/40/80/160 MHz bandwidth OBSS's
- Can dynamically determine best bandwidth for each AP (DBS v.8.1)

Figure 1: When a new AP is added, its radio conflicts with an existing AP's radio causing contention. DCA adjusts the channel plan for the best solution for the new AP



DCA's job is to monitor the available channels for the RF group and track the changing conditions. Optimizing the RF separation between AP's (minimizing co-channel interference) by selecting channels that are physically diverse which maximizes RF Efficiency. DCA monitors all available channels and develops the Cost Metric (CM) that will be used to evaluate various channel plan options. The CM is an RSSI value comprised of interference, noise, a constant (user sensitivity threshold), and load (if enabled). The Cost Metric equates to a weighted SNIR (Signal to Noise Interference Ratio). See RRM Data Collection Activities above for a complete discussion.

Competitive Note - our competitors radio management systems also must monitor off channel in order to develop information used for decisions. Cisco's RRM implementation has consistently tested as the least disruptive. Conducting throughput testing can validate this; Cisco AP's maintain fluid information flows. Competitor's products typically show distinct drops in throughput when subjected to the same test suites. Aruba by default requires a 110 ms dwell off channel. Off Channel scans are used for many things, implementation of wIDS/wIPS typically requires extensive off channel scanning, not just on DCA channels but typically on Country Channels which is a much larger list to visit. Turning off RRM, disables these off channel scans - but it also eliminates wIDS and rouge detection as well.

DCA uses all of these measurements and sums them up into an RSSI based Cost Metric that will be used in the equation. The cost function is a single numeric value expressed as RSSI that represents the overall goodness of a given channel option.

Changing the channel of an AP is potentially disruptive. Care must be taken in the evaluation of apparent improvements. This is where next generation DCA excels. Determining if an AP's performance can be improved without negatively impacting neighbors in the neighborhood is a multi-step process.

The Dynamic Channel Assignment (DCA) Algorithm

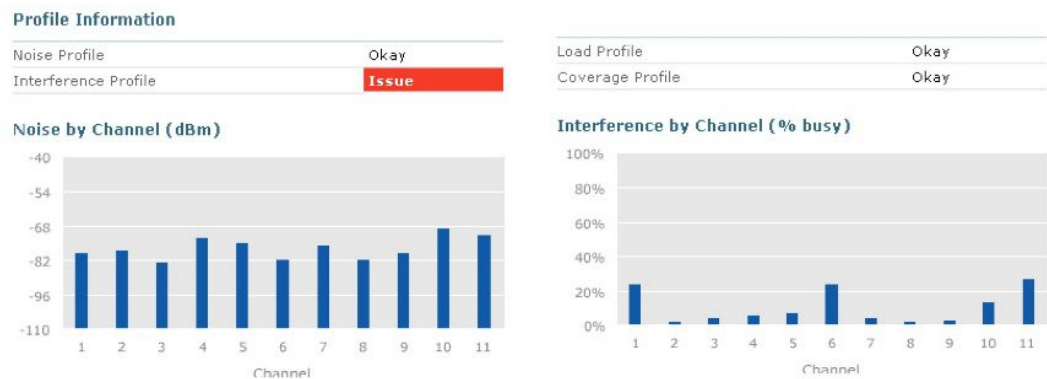
The Group Leader maintains the neighbor lists for all AP's in the RF Group, and organizes these neighbors into RF Neighborhoods. The following metrics are also tracked for each AP in the RF Group.

- **Same Channel Contention**—other AP's/clients on the same channel - also known as Co-Channel interference or CCI
- **Foreign Channel - Rogue**—Other non RF Group AP's operating on or overlapping with the AP's served channel

- **Noise**—Non-Wi-Fi sources of interference such as Bluetooth, analog video, or cordless phones - see CleanAir for useful information on using CleanAir to detect noise sources
- **Channel Load**—through the use of industry standard QBSS measurements - these metrics are gathered from the Phy layer - very similar to CAC load measurements.
- **DCA Sensitivity**—A sensitivity threshold selectable by the user that applies hysteresis to the evaluation on channel changes

The impact of each of these factors is combined to form a single RSSI based metric known as the Cost Metric (CM). The CM then represents complex SNIR of a specific channel and is used to evaluate the throughput potential of one channel over another. The goal is to be able to select the best channel - for a given AP/Radio while minimizing interference . Using the CM, the Group Leader is able to evaluate every AP and every channel for maximum efficiency. Of course conditions change in RF, so these statistics are dynamically collected and monitored 24 hours 7 days per week.

Figure 2: View of Interference and Noise from the Radio Page on a Controller



Using the CM for the currently served local channels on the AP's, the RF group leader develops a list stack ranked worst to best. This becomes the CPCI list (Channel Plan Change Initiator) which indicates which AP's are suffering the worst performance in the RF Group. For simplicity - lets take a quick look at a single AP and what DCA does - then we'll apply that concept to the more complicated job of an entire RF group with channel bonding and multiple AP capabilities.

DCA in a Nutshell


A DCA run starts with selecting a CPCI - by default, DCA will always pick the AP with the worst CM to start with, and alternate for successive iterations between a random AP and then the next worst on the remaining list. DCA takes the CPCI, along with all of it's 1st hop and 2nd hop neighbors as a group to see if a channel plan can be calculated that provides a better selection for the current CPCI.

A first hop neighbor is any AP our CPCI knows about through direct observation (neighbor relation), a second hop neighbor is an AP that is in our neighborhood and we know about because our first hop friends know them. In the evaluation, channels for the CPCI and all first hop neighbors may be changed to achieve a solution. Channels for second hop neighbors - while evaluated for impact, cannot be changed. This allows isolation of local groups of AP's and prevents the possibility of a change impacting AP's across the entire RF group.

Once the calculations are complete the result is often several possible channel plans which will improve the CPCI. Each channel plan, which yields improvement, is subjected to another gating feature known as the

NCCF (normalized cumulative cost function). This non-RSSI based function evaluates the resulting channel plans for overall CPCI group goodness, in other words the CPCI must see an improved CM, but only if it's neighbors, as a group, either improve or stay the same for the channel plan to be recommended.

Figure 3: CPCI with First and Second hop RF Neighbors



Once the calculation is complete, the CPCI and its first hop neighbors are removed from the CPCI list, and the next iteration begins with a random selection out of the remaining AP's on the list. The DCA process will alternate between worst and random selections until the entire CM list is empty. In this way - all AP's are evaluated in the context of every other AP that can hear them. DCA completes when the CM list is empty, NCCF is completed and channel changes are processed.

DCA Sensitivity Threshold

Wi-Fi is a bursty medium, meaning that things can look really bad for a short period of time, but over all be pretty good. Since changing the channel of an AP is potentially disruptive care is taken to ensure that if a change is made - it is for a non trivial performance improvement and not a knee jerk response to a short term trend. A user selectable sensitivity threshold is provided that allows dampening of the channel change algorithm. The default value is medium (10 dB), and essentially says that in order for a channel change to be made, the new channel must have a CM of 10 dB better in order for it to be recommended. The low sensitivity value is 20 dB and the medium value is 10-15 dB depending on band. NCCF will process this threshold since it has final say on a recommended channel plan. Any channel plans not meeting that criteria will not be processed at the AP.

Table 1: DCA Sensitivity Thresholds by Band

Band	Low	Medium	High
2.4 GHz	20 dB	10 dB	5 dB

Band	Low	Medium	High
5 GHz	20 dB	15 dB	5 dB

The evaluation is simple. NCCF asks, is the Delta between current and proposed channel cost metrics equal to, greater than or less than DCA sensitivity threshold value? If equal or greater than, then the channel change is recommended. This serves to dampen temporary or short term gains and thrashing of channels in response to loads which can have a bad effect on client connectivity.

DCA Modes of Operation

Scheduled DCA

DCA operates by default every 10 minutes (600 seconds) in steady state once it has been initialized unless some other interval is defined and DCA is running in Scheduled mode. Scheduled DCA allows customers to plan around potential disruptions associated with channel changes, however it should be noted that the DCA algorithm will only run at this selected time and may not be evaluating the users environment at peak loads. The same environment when loaded with clients could be significantly different. To increase the effectiveness it is recommended that customers select the highest sensitivity level which will maximize the changes made during off peak hours. It's also a good idea to periodically re-evaluate the environment for its tolerance to channel changes. As clients are refreshed this will improve and most modern clients do just fine managing a channel change.



Note Whenever an AP's channel is changed clients will be briefly disconnected. Depending on client roaming behavior, clients may either reconnect to the same AP (on its new channel), or roam to a nearby AP. The clients ability to roam properly will determine it's effectiveness during a channel change.

Start-up Mode



Note When AP's boot up for the first time (new out of the box), they transmit on the first non-overlapping channel in the band(s) they support (channel 1 for 11b/g/n and channel 36 for 11a/n/ac). When AP's power cycle, they use their previous channel settings (stored in the AP's memory). Dynamic Channel Assignment adjustments will subsequently occur as needed.

Any time that a controller in the RF Group enters or departs the RF group (a reboot for instance) Start-up mode is assumed. This means that if the controller was the RF Group Leader and it returns as the RF Group leader then DCA will run startup mode - regardless of the user settings- every 10 minutes for the next 100 minutes. Now, obviously this is something that should be considered before rebooting a controller, however it's not as bad as it may seem. If the network was previously at steady state, then the AP's channel assignments should already be optimized. If the controller is a new addition, and you've added AP's then DCA will need to run to optimize the new channel assignments required. Plan accordingly.

Start up mode is aggressive and ignores NCCF and the user sensitivity threshold. It will produce a channel plan that maximizes the RF Distance between AP's without regard to the dampening mechanisms designed to slow the rate of change in a live network.

Since version 7.3 of code, there is a command line argument for initializing DCA startup mode. It is present on all controllers in an RF Group - but will only affect the DCA mode of the controller whom is the RF Group Leader. Running the command **config 802.11a/b channel global restart** from the command line of the Group leader will re-initialize RRM's DCA and provide an optimal answer based on measured values over the air.

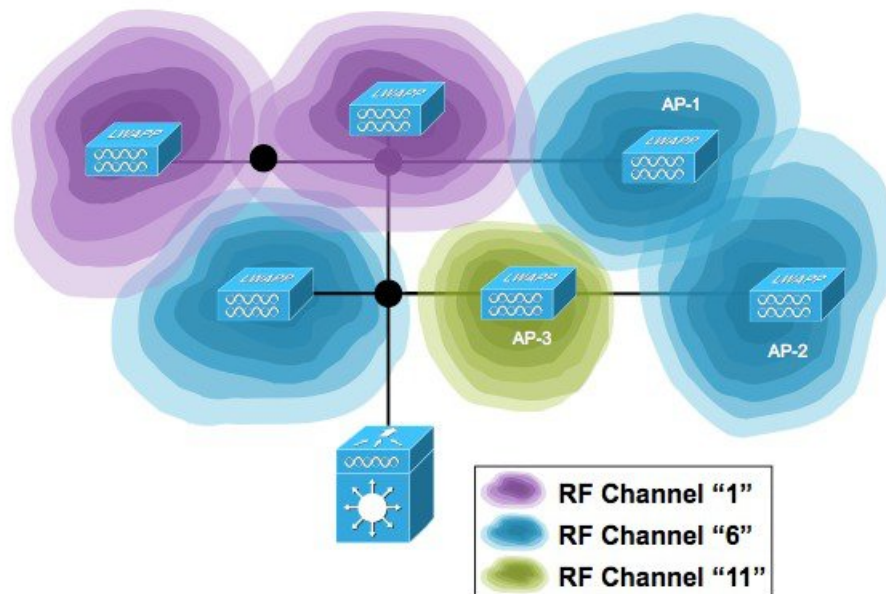
Steady State Mode

DCA runs by default every 10 minutes. If the user schedules DCA with an Anchor time and interval - DCA runs on the scheduled intervals. Cisco recommends a minimum of 2 intervals per day - even though it is possible to run only 1. See Scheduled DCA above for additional considerations.

Over time, and especially with changes in the network architecture the user sensitivity threshold (dampening) can lead to sub optimal channel assignments. Most network architectures change over time, and DCA's rules assume a steady state network. If AP's have been added or removed, or channel bandwidths have been changed network wide, it's very possible that you could have AP's that could see a 9 dB improvement in the cost metric, but because the hysteresis is 10 dB (default) a change is not made.

When making changes to the architecture it is a best practice to restart the DCA algorithm by placing it into Startup Mode which suspends all user settings (the sensitivity threshold) and the NCCF functions and permits an aggressive channel search for a good baseline on the new architecture.

Figure 4: DCA operational example



Using Figure 12 above, lets suppose that AP-1 is on channel 6 and has the worst CM for the group at -60 dBm (Remember, less is more. The lower the CM the lower the noise floor and the better the throughput).

- DCA Evaluates Channels 1 and 11 for AP-1's location and determines that the CM could be -80 dBm on channel 11
- This represents a potential $\Delta(\text{CM}) = 20$ dB if we change channel 6 to channel 11 for AP-1
- DCA would change the channel if sensitivity set to High or Medium or Low (5,15,20) are all = to or < 20).

- If the CM for Channel 11 were -75, then the delta would be 15 dBm and a change would only be made if the sensitivity threshold were High or Medium (5 or 15 dBm) but not low as 15 dB does not meet the 20 dB hysteresis.
- Additionally, if the new channel plan results in neighbor changes and the neighbors CM will be driven lower – NCCF will NOT Recommend the channel plan for implementation

Without diving heavily into the math, NCCF provides a normalization of the CM data for the CPCI and its first hop neighbors and prevents making a channel change if the CPCI would negatively impact its neighbors. Think of NCCF as an overall "goodness" rating of the change for the group. This breaks down like this.

NCCF is applied as such to each radio being affected by the recommended change (CPCI and its 1st and 2nd hop neighbors)

- +1 - if CM improves by +5 dBm or better
- 0 - If CM \pm 4 dBm
- -1 - If CM worsens by 5 dBm or better


If NCCF evaluates the recommended change as being beneficial for the CPCI and its neighbors then the change is implemented.

DCA 20, 40, 80, and 160 MHz support

Keeping in mind that everything that is evaluated by RRM is based on actual over the air observations. How then does RRM handle coexistence and the challenges of 20/40/80/160 MHz OBSS channel selections. What if we are deploying a mixture of 802.11a/n/ac (or perhaps we have 802.11a radios as neighbors) how does RRM's DCA address this? Things have become complicated for sure, but the goal of DCA is always to create a channel plan that favors constructive coexistence. Constructive coexistence doesn't mean we can eliminate the other radios in the air, they are usually there and have a legal right to be, but rather make a decision that reinforces a complementary plan and supports everyone's contention needs and provides equal - shared - access to the medium.

DCA, the OBSS and Constructive Coexistence

Figure 5: OBSS Channel Architecture



The OBSS or Overlapping BSS became a reality with the introduction of 802.11n and continues with 802.11ac. Both of these protocols allow for dynamically linking multiple 20 MHz channels together to form a wider

channel in which more data can be transmitted simultaneously. Channel positions within the bonded channel are important, as not all channels behave the same.

Table 2: OBSS Bonded Channel Segment Names and Function

Abbreviation	Proper Name	Function and Notes
P20	Primary Channel	All management and signaling frames, HT and VHT headers are on the P20 only
S20	Secondary 20	added to the primary for additional capacity to form a 40 MHz channel - may be +/- of the primary channel position
S40	Secondary 40	Added to an P20 and S20 to make an 80 MHz channel. Bonded channels must in the same band (Unii 1,2,2e,3)
S80	Secondary 80	Added to an P20 and S20 to make an 80 MHz channel. Bonded channels must in the same band (Unii 1,2,2e,3)

For the purposes of this discussion we will focus on 5 GHz. It is legal to have an 802.11n BSS use a 40 MHz channel in 2.4 GHz, however Cisco does not support this. There are simply not enough channels in 2.4 GHz spectrum for this to be effective. 802.11ac - ONLY operates in 5 GHz spectrum.

802.11a clients do not understand 802.11n HT headers, and both 802.11a and 802.11n don't understand 802.11ac's VHT header. In order to maintain backward compatibility and satisfy all three protocols requirements - all 3 share the primary channel architecture and definition as a common signaling channel using the 802.11a protocol. Both 802.11n and 802.11ac add an additional headers (HT and VHT) to the standard 802.11a frame format used to advise 802.11n and 802.11ac clients on specifics such as channels and selected bandwidth as well as supported data rates for each protocol. All management (broadcast) traffic will use the 802.11a protocol on the primary channel. To an 802.11a device - it's all 802.11a.

Wi-Fi is contention based. Each station listens to the channel to determine when it is quiet (listen before talk or LBT). However, not all 20 MHz segments are treated equally in within a bonded channel. Secondary channels have less contention to ensure that when the primary channel is clear, the secondary(s) have a higher probability of also being clear. For this reason it is important to understand the impact this can have in a design where multiple protocols are being supported (at a minimum today you will have 802.11n and 802.11ac AP's present either as infrastructure or rogue neighbors).

In the table below, CCA thresholds example, the RSSI values are the thresholds at which the receiver must listen to determine if the channel is busy or idle. CCA assessment is done by segment, and the first not clear segment suspends checking the rest of the channel segments and reports not clear to the host. Energy at or above the threshold indicates a carrier busy or not clear - and no TX will happen. Any energy falling below the threshold, represents a distant station and we consider the channel idle and we can clear the next segment or transmit if all are completed.

Note that all three protocols share the same value for the primary channel - this makes them equal with regards to contention -they will all get fair access to the medium. You can also see that the values for the Secondary

20, and all other secondary's are more generous (in that the threshold is higher representing less contention - and with a higher probability of winning contention than a station that is listening at a lower value).

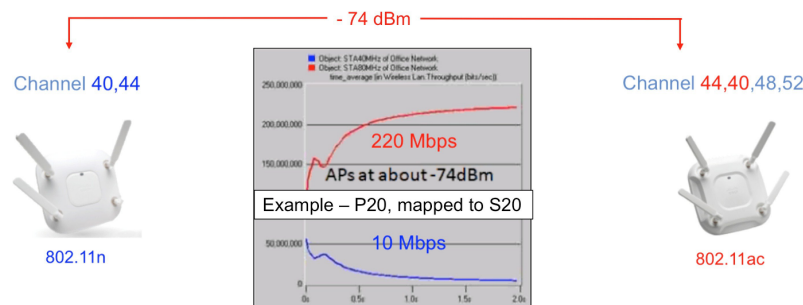
Table 3: CCA Threshold Examples

CCA Threshold Example				
Protocol	P20	S20	S40	S80
802.11a	-82	—	—	—
802.11n	-82	-62	—	—
802.11ac	-82	-72	-76/-79	-76/-79

DCA's job is to provide a channel plan accounting for the variables, as they exist, in the air around each individual AP. Critical to this is the overall number of available channels, and that changes based on both the regulatory of the equipment and the channel width selected. An 80 MHz channel is 4x20 MHz channels so depending on your regulatory; you can chew through channels pretty quickly and leave yourself without enough spectrum to build an efficient network. We also have to make these decisions in a way that promotes and supports a constructive coexistence between different specifications or someone will go wanting.

For instance, referencing the table above for CCA thresholds, If I place an 802.11n 40 MHz P20 channel on an 802.11ac (or 802.11n for that matter) S20 channel, I am forcing the 802.11n AP to compete for airtime against a stacked deck - since the 802.11n AP will need to wait until the channel is quiet at -82 dBm to win contention - while the 802.11ac AP only has to clear the same channel down to -72 dBm. This sets up a very unfair match in which the 802.11ac AP can starve the 802.11n AP for access - simply because every time they both need the channel - the 802.11ac AP will likely win. This assumes that the two AP's are close enough to hear one another at the affected range say -74 dBm (there will be plenty of these close enough in a moderately dense network).

Figure 6: Destructive Coexistence Example



The graphic below shows two RF Coverage plots made using average device (client) power of 10 dBm. The AP listening at -82 dBm (CCA for a P20), is in contention with every station within the -82 dBm plot area. The coverage area for -76 dBm (CCA for an S20 channel) is much smaller - and represents a lot less stations to compete with.

DCA's algorithms are looking for 3 possible solutions to work out compromises, each for both our AP's and neighbors or rogues. In order of preference, if there are no free channels available DCA

- Primary channels aligned = P20 to P20 = BEST
- Primary Channel aligned Secondary 40 or 80 = P20 to S40/S80 = OK

- Primary Channel aligned with Secondary 20 = P20 to S20 = Better than nothing

After that, DCA runs as normal - seeking to resolve the channel plan with the given mix of radios. Assignments with someone's 20 MHz channel as a secondary channel are given a higher cost metric to lessen the likelihood of their selection as a valid assignment for any radio in the domain.

In RRM, you may select either 20/40/80 MHz channels from the DCA dialogue, however if the radio is an 802.11a Radio, it can only support a 20 MHz channel - and that is all it will receive. Likewise for 802.11n radios, if you select 80 MHz - they will be assigned a 40 MHz channel.

Is there any benefit to running the 802.11n or 802.11ac protocol even if you choose to not support 40 or 80 MHz channels? Certainly, Higher Data Rates, better multipath immunity, and Client Link are three examples of big benefits that can be enjoyed by legacy as well as 802.11n/ac clients. There is all upside and no downside to implementing 802.11n or ac regardless of the clients operating on the infrastructure - that's pretty rare in networking.

Dynamic Bandwidth Selection–DBS

The DBS feature was introduced in version 8.0 of the code and represents a flexible and intelligent way to allow RRM to assign bandwidth to AP's that have clients associated that can benefit from the additional bandwidth. This approach is dynamic, and since it is based on analysis of what the client capabilities are as well as what they are doing allows RRM to Right Size the network channels.

As previously discussed, the advantage to having a wider channel is obvious - more data with each transmission. However, this only holds true if we can balance this against contention needs and spectrum availability. Moving more data with every transmission is not better if I have to wait 3 times as long to send a single packet - the result could be worse than sending what I have more frequently, in smaller bits. Not all applications actually benefit from bonded channels; Voice for instance relies on small packets that are time sensitive (jitter). Video however benefits greatly - but still has a sensitivity to Jitter in some cases (real time video). Neither are the channels within the bonded channel equal in function. The Primary channel is the only one that will be transmitting signaling information where the other bonded channels will simply send payload associated with a packet defined on the signaling channel. Secondaries are less loaded than the primary's as a rule.

The Best Practice for most organizations today is to use no more than 40 MHz in enterprise deployments. However it really comes down to how many channels you have and how close your AP's are to one another. For this reason DBS relies on the tremendous amount of information available within RRM to dynamically adjust the channel width in conjunction with it's other duties.

DBS will evaluate:


- Associated client capabilities and types
- RF Neighbor Channel Widths
- OBSS channel Overlap ratios
- Channel Utilization
- Non Wi-Fi Noise
- Wi-Fi interference

In practice, the main objection to this feature has been - "but I want an 80 MHz channel, and it wont give it to me here.....". You can still override this feature and set a manual bandwidth on the AP, however be warned that RRM didn't think it was a good idea, it is usually pretty right on these things.

Flex DFS - Flexible Dynamic Frequency Selection

With the inclusion of DBS, another challenge that is observed in the modern OBSS world is resolved as well. If the channel definition is 80 Mhz, comprised of 4x 20 MHz segments and we are using UNII 2 channels (DFS) then if a radar is detected on any of the 4 20 MHz segments forces abandonment of the entire channel by the AP and the users. Without DBS and Flex DFS this equates to an 80 MHz chunk of spectrum which is marked as unusable for 30 minutes. With DBS and Flex DFS - we simply mark the affected 20 MHz channel - and reconfigure the AP accordingly to use either the remaining 40 MHz channel or the 20 MHz channel, either way - the AP and clients no longer have to switch gears - the AP does not have to find space that is less optimal for it's position- and you only loose 20 MHz - not 80 MHz of spectrum.

Figure 9: Example of Flex DFS channel options



This seems like a simple thing - and it makes sense. However if I have told the system to only assign 80 MHz channels - this is what it will look to do. With DBS and Flex DFS we give the system the ability to do what makes the best sense while maintaining compliance.

Device Aware RRM

CleanAir shares information with RRM that normal Wi-Fi radios do not have access too at the physical layer. Non-Wi-Fi radio interference (known as noise to Wi-Fi) is actionable information for RRM in some instances. For instance, a Microwave oven, most offices have at least one - and it represents a significant source of noise for Wi-Fi. There are two CleanAir features that interact with RRM in different ways; we will discuss those here.

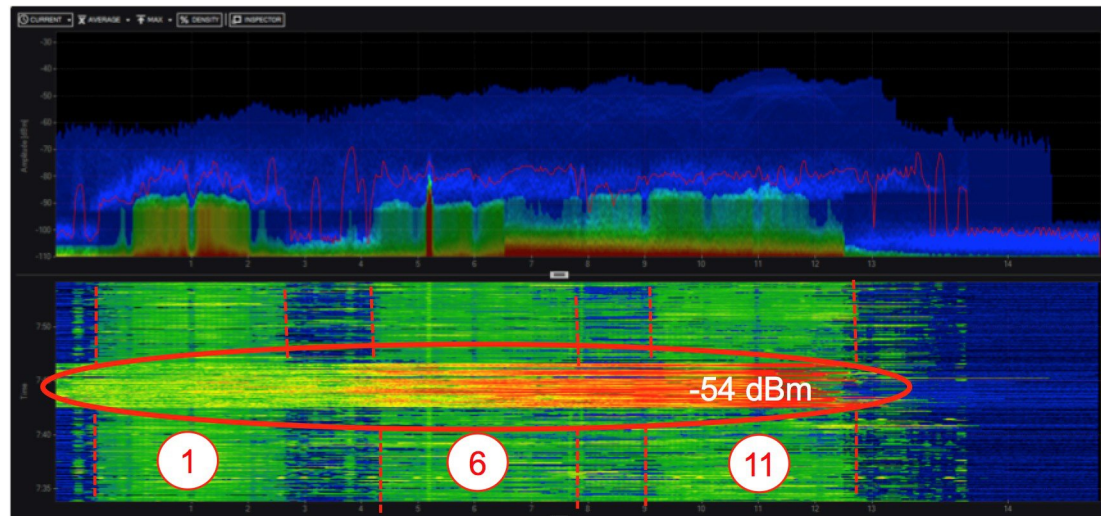
Persistent Device Avoidance

Persistent device avoidance identifies sources of Wi-Fi interference, which are frequently present within installations and some which are not. If present, these devices represent a factor, which, while perhaps not constant, will negatively impact any channel that they interfere with and as a result, should be avoided. RRM's normal data collection and action cycle will be aware of the interference and will avoid it. However, once the source goes quiet, the channel that was avoided will likely look good to RRM again and in that case RRM will likely re-assign the radio to the previously bad channel. Microwave Ovens, Outdoor Ethernet bridges are two classes of devices that qualify as persistent, since once detected, it is likely that these devices will continue to be a random problem and are not likely to move. For these types of devices we can tell RRM of the detection

and Bias the affected channel so that RRM "remembers" that there is a high potential for client impacting interference for the Detecting AP on the detected channel.

Lets use a Microwave oven as an example. Most workplaces have at least one, and some have many. While in operation an MWO will impact the 2.4 GHz band with high duty cycle noise. MWO's operate anywhere from 700-1200 watts for consumer units, and can range higher for commercial grade units. MWO's are shielded to avoid harmful radiation leakage, but the concern here is for the humans, not the Wi-Fi and operating at a fraction of a watt, there is enough energy left over to seriously impact communications. MWO's operate anywhere within the 2.4 GHz spectrum, generally at the higher end (channel 11) but frequently impacting channel's 11,6 or even the entire band.

Figure 10: Microwave Oven impact - Channelizer Pro



MWO's do not run continuously, generally first thing in the morning - on and off for a couple of hours around lunch - then again for the afternoon popcorn. Persistent Device Avoidance allows us to Mark and AP and it's detection channel so that RRM knows the device exists. PDA registers the interference, and then starts a countdown timer which refreshes with each new detection. If at the end of 7 days, no more detections were processed, the bias is removed and the PDA detection is reset.

Biasing an affected AP/Channel does not guarantee that RRM will not use that channel for that AP, but it decreases the likelihood by increasing the cost metric. The end result is up to DCA as even with the cost metric bias, this could still be the best channel available.

You can view an AP's PDA status on the controller under Wireless>802.11b/g/n (or 802.11a/n/ac)>details, at the bottom of the Details page is the current PDA devices being tracked with their last detection date.

Figure 11: Persistent Device Table from MMAP

CleanAir Parameters				
Operational Status	Up			
Persistent Devices				
Class Type	Channel	DC(%)	RSSI(dBm)	Last Seen Time
Microwave Oven	9	23	-44	Fri Jan 29 13:12:19 2016
	10	15	-57	Fri Jan 29 18:58:55 2016
	11	12	-65	Fri Jan 29 18:58:55 2016
WiMax Fixed	2	0	-60	Thu Jan 28 17:17:00 2016
	3	0	-60	Thu Jan 28 17:17:00 2016
	4	0	-60	Thu Jan 28 17:17:00 2016
	5	0	-60	Thu Jan 28 17:17:00 2016
	6	0	-60	Thu Jan 28 17:17:00 2016

CleanAir PDA devices include:

- Microwave Oven
- WiMax Fixed
- WiMax Mobile
- Motorola Canopy

PDA is based on an actual device classification - so we know that this device exists, and we know which AP's could hear it at a level that was impacting. This allows RRM to work around these devices to come up with an alternate channel plan that works around the affected channels for the areas where there is an issue. PDA only affects the AP that detected the device.

A secondary feature to PDA, which was added, is called Persistent Device Propagation or PDP. This feature was designed to share CleanAir information with non-CleanAir AP's through RRM. This feature (disabled by default) if enabled shares the PDA report with neighbors of the detecting AP and applies the same bias for the same channel to neighbors of the detecting AP. This is a secondary function, which happens completely outside of CleanAir. Once detection is logged on a CleanAir AP - RRM will propagate the same bias, which is applied to the detecting AP with all neighbors that are above -70 dBm to the detecting AP.

802.11b > CleanAir

CleanAir Parameters

CleanAir	<input checked="" type="checkbox"/> Enabled
Report Interferers ¹	<input checked="" type="checkbox"/> Enabled
Persistent Device Propagation	<input type="checkbox"/> Enabled



Note This feature should be used with great caution - as some installations can have a lot of neighbor AP's that can be heard at or above -70 dBm and you could potentially exclude a channel from an entire RF neighborhood - potentially.

This feature was created as a stopgap for customers use while implementing CleanAir AP's, it should not be used as part of a plan to mix some CleanAir AP's in with existing non-CleanAir AP's unless you are deeply familiar with CleanAir behaviors and understand the risks.

Channel Change traps related to PDA will have "Device Aware" as the reason code.

ED-RRM

ED-RRM is not directly related to RRM, but will cause channel changes if invoked. ED-RRM stands for Event Driven-RRM and is intended to quickly resolve catastrophic interference events. Because Wi-Fi is Listen Before Talk (LBT) If there is energy on the channel above the CCA threshold - all stations will hold off using the channel until it has cleared. Certain non-Wi-Fi devices are classified as continuous, meaning 100% or near 100% duty cycle, in short they never turn off. An analogue video camera is an example of such a device. If this device is present, neither the AP or it's clients that hear it will ever attempt to transmit, since the energy is always present. This would be corrected by normal RRM DCA activities, however correction could take up to 10 minutes (DCA interval) or more if DCA timing has been changed.

CleanAir at the AP allows us to recognize such a device, and positively classify it as such a device (can not be confused with normal Wi-Fi Oversaturation). This is a distinct advantage, since we know for certain if this device exists, it will not yield the channel or get better on it's own unless disabled. We can however detect this very quickly at the AP interface, and allow the AP to make a temporary channel change to quickly avoid this energy and restore service. Following that change a normal DCA cycle will find a better permanent home for the AP that avoids the now unusable channel in that location.

ED-RRM is based entirely on the Air Quality metric on the AP. Air Quality or AQ for short is entirely comprised of CleanAir classified non-Wi-Fi interference metrics, so can not be driven by unclassified or normal Wi-Fi related noise. Simply relying on noise for this would be very bad since Wi-Fi noise can have very high short duration peaks followed by relative calm - this is quite normal. However relying on the AQ metric avoids all of this since we know for certain that it is a problem that is not just going to go away.

In version 8.0 a new component was included in ED-RRM functionality. Rogue Contribution, which allows ED-RRM to trigger based on identified Rogue Channel Utilization, which is completely separate from CleanAir metrics. Rogue Duty Cycle comes from normal off channel RRM metrics, and allows us to invoke a channel change based on neighboring rogue interference as well. Because this comes from RRM metrics and not CleanAir, the timing - assuming normal 180 second off channel intervals - would be within 3 minutes or 180 seconds worst case. It is configured separately from CleanAir ED-RRM and is disabled by default. This allows the AP to become reactive to Wi-Fi interference that is not coming from our own network and is measured at each individual AP. Other than the source trigger, Rogue Contribution in ED-RRM follows the same rules as CleanAir contribution.

The AP calculates AQ on a 15 second rolling window, and any two consecutive AP level AQ threshold violations will trigger ED-RRM is configured (disabled by default). It also has the following protections:

- Once triggered, the AP is desensitized for ED-RRM for 60 seconds on the new channel – to prevent immediate flapping
- Once a channel has been identified with an ED-RRM trigger event – that channel is locked out for 60 minutes.

Using 2.4 GHz as an example, let's say that we trigger an ED-RRM channel change on Channel 1 and switch to channel 6. Let's assume that the interference covers the entire 2.4 GHz band, and we trigger again on channel 6 after a 60 second rest and move to channel 11. In our scenario channel 11 is also affected and so also triggers an ED-RRM alert in 60 seconds. At this point - there are no other channels to move too, since both channel 1 and 6 are now in a 60 minute lock out. The AP would continue to sit on channel 11 until such time that either the 60 minute timers are cleared - or the interference is disabled/corrected. This prevents flapping or a runaway condition.

Configuring ED-RRM is done through the **Wireless>802.11a/b>DCA configuration dialogue**.

Figure 12: ED-RRM config Dialogue - WLC GUI

Event Driven RRM

EDRRM	<input checked="" type="checkbox"/> Enabled
Sensitivity Threshold	Low
Rogue Contribution	<input checked="" type="checkbox"/> Enabled
Rogue Duty-Cycle	80

Configuration consists of enabling ED-RRM (disabled by default) and selecting the AQ threshold level:

Low sensitivity = AQ at 35%

Medium sensitivity = AQ at 50%

High sensitivity = AQ at 60%

Custom = custom - but be very careful here

Remember that AQ is a scale which shows the collective impact of all CleanAir classified Interferers, a good AQ is 100% and a very bad one is 0%.

To enable and use Rogue Contribution, ED-RRM must be enabled first, then enable Rogue Contribution, Rogue Duty cycle is just that - the default is 80 which means if Rogue devices are using 80% of the channels capacity, you should leave and find a better channel.

While neither of these triggers and responses are driven by DCA, they will be honored by DCA and channel changes to re-balance the surrounding AP's will likely happen after a trigger event. Channel Change traps resulting from ED-RRM triggers will include "Major AQ event" for the reason code.