



Radio Resource Management White Paper

First Published: 2016-02-18

Last Modified: 2016-06-14

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2016 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1	Radio Resource Management	1
	Introduction	1
	A Brief History of RRM in Cisco	2

CHAPTER 2	Radio Resource Management Concepts	7
	Pre-requisites and Assumptions	7
	Key Terms	8

CHAPTER 3	RRM Data Collection Activities	11
	RRM Data Collection Activities	11

CHAPTER 4	RF Grouping	15
	How RF Groups are formed	16
	Neighbor Discovery Protocol–NDP	16
	NDP and DFS	19
	What do we use NDP for?	20
	RF Group Leader Election	21
	RF Grouping Automatic mode	23
	Static RF Grouping	24
	RF Group Scalability	24
	RF Group Backward Compatibility	26
	WSSI and WSM, WSM2 Modules and RRM	26
	Troubleshooting RF Grouping	26
	RRM Data Collection	26
	RF Grouping Trouble	28
	Summary of the Reason Codes	31

CHAPTER 5	Flexible Radio Assignment (FRA) and Redundant Radios	33
	FRA–Coverage Overlap Factor–COF	34
	FRA–Radio Role Assignment	38
	FRA–Client Network Preference	40
	FRA–Steady State Operations	42
	FRA and Dual 5 GHz Operations	42
	802.11v BSS Transition–Default Enabled	43
	802.11k–Default Enabled	44
	FRA–Troubleshooting	45
	FRA–CLI and Debug Commands	46
	FRA–Best Practices	49

CHAPTER 6	Dynamic Channel Assignment (DCA)	51
	What does Dynamic Channel Assignment do?	51
	The Dynamic Channel Assignment (DCA) Algorithm	52
	DCA in a Nutshell	53
	DCA Sensitivity Threshold	54
	DCA Modes of Operation	55
	DCA 20, 40, 80, and 160 MHz support	57
	DCA, the OBSS and Constructive Coexistence	57
	Dynamic Bandwidth Selection–DBS	60
	Flex DFS - Flexible Dynamic Frequency Selection	62
	Device Aware RRM	62
	Persistent Device Avoidance	62
	ED-RRM	65

CHAPTER 7	Transmit Power Control (TPC) Algorithm	67
	What does TPC do?	67
	TPCv1	68
	Calculating Tx_Ideal–Ideal Power	69
	Evaluating a TPCv1 Change Recommendation	69
	Implementing a Recommended Power Change	69
	TPCv2	70

TPC Min/Max 74

CHAPTER 8 Coverage Hole Detection and Mitigation Algorithm 77

Coverage Hole Detection (CHD) 77

Coverage Hole Mitigation 79

Optimized Roaming 79

CHAPTER 9 RF Profiles 81

TPC 82

DCA 83

Coverage Hole Detection 83

Profile Threshold for Traps 83



CHAPTER 1

Radio Resource Management

- [Introduction](#) , on page 1
- [A Brief History of RRM in Cisco](#) , on page 2

Introduction

Wireless connectivity is truly ubiquitous and Wi-Fi is one of the fastest growing wireless technologies of all time, everything has a Wi-Fi chipset and client installed and IoT is just getting warmed up. An analysts report back in 2011 said that the number of wireless stations vs wired stations on the internet will likely flip with wireless users exceeding wired nodes by 2015. We beat that estimate in 2014, and the future is clear, there will be more. Once upon a time we counted seats to evaluate capacity, however most users have more than one active device operating at all times. As I'm sitting here writing this, I count 3 devices between the laptop I am writing this on, my smartphone, and the tablet that is in hibernation mode at the moment. I'm not turning any of these off - and you probably aren't either. All of this persistent connectivity requires bandwidth and resources, wireless spectrum is becoming even more precious than in previous times and the pressure on available spectrum doesn't look to be easing anytime soon. What has not changed significantly is the spectrum with which we have to work. All of this together makes managing what you have the primary mission of any wireless administrator or network operator.

Most of the pressure to date has been in the 2.4 GHz spectrum, however this will be spreading to a 5 GHz band near you soon (if it hasn't already). If this is your first foray into that deep and mystical world of the RF physical layer, fear not, the rules on this have changed pretty regularly so you're not behind but rather just in time! Most of the work will be done for you, but like any manager it's a good idea to to understand the goal of your team and their individual strengths. To that end, we'll discuss what RRM is, what it does, and how it does this. We will also discuss how to characterize your operating environment so you can ask RRM good questions.

Why is this important? Spectrum is the physical layer. Unlike wired networks - our spectrum is free to propagate in all directions. This means that if two cells overlap one another on the same channel, that they are sharing the spectrum normally reserved for each. Not only are users of each cell sharing the single channel of available spectrum, it's doubled the management traffic. The result is higher consumption of air time and less throughput. This is commonly known as co-channel interference. Assuming that all wireless devices are operating on your network and not on a neighbors, there is only two things that can be manipulated to adjust any given cell in response to co-channel interference:

- **Channel Plan:** adjusting the channel plan to facilitate the maximum separation of one AP from another
- **Power Levels:** power levels increase or decrease the size of the effective cell

Both of these are separate arguments but work together to produce an effective solution.

Cisco's Radio Resource Management (abbreviated RRM) allows Cisco's Unified WLAN Architecture to continuously analyze the existing RF environs, automatically adjusting each AP's power and channel configurations to help mitigate such things as co-channel interference and signal coverage problems. RRM reduces the need to perform exhaustive site surveys, increases system capacity and provides automated self-healing functionality to compensate for RF dead zones and AP failures.

This paper details the functionality and operation of RRM and provides an in-depth discussion of the algorithms behind the features

A Brief History of RRM in Cisco

RRM was introduced originally as a feature on AireSpace AP's and controllers, and became part of the Cisco CUWN with the acquisition of AireSpace in 2005.

In 2005 if you had 150 AP's in a network, that was a large Wi-Fi network. Today we routinely see RF installations with 3000-5000 and more AP's installed in campus deployments, stadium environments, conference centers, metro deployments, and hospitals. Much has changed in this short history - and as the questions have changed - so too have the answers that RRM must deliver. Since 2007, every release of CUWN (Cisco Unified Wireless Network) code has included several features related to RRM, as well as features designed to increase spectral efficiency and enhance RRM's effectiveness.

As AP spacing continues to decline, installations have migrated from simply providing Coverage models to demanding dense capacities for thousands of devices as the only edge technology. The investment in RRM as a core technology has kept pace. Smartphones and tablets with no wired connection have gone from being an accessory to being the main computing platform for users, and with this some growing pains as both the design methodologies and the network as a whole have had to adapt to different design goals, technologies and strategies.

Figure 1: Visual Timeline of Wi-Fi



Today, the wireless office is not just a cool idea - it is being implemented around the world as the only network connectivity possible between the diverse range of devices we require to do business and provide core services. Yes, Wi-Fi is mission critical.

RRM has kept pace as the technology has changed. We've gone from legacy single radio interfaces to 80 MHz 4 spatial stream 802.11ac in the last 5 years. Many of these changes have required new radios to take advantage of the advances in efficiency. We not only need to upgrade the core network but the clients as these changes impact our environments. When 802.11n entered the scene in 2003 we began to discuss the concept of an

OBSS (Overlapping Base Service Set) and instead of modulating a single half duplex radio stream, we began modulating simultaneous spatial streams as well as linking existing 20 MHz channels together to increase the channel width and spatial efficiency.

Table 1: Current Wi-Fi Protocols and Capabilities Compared

Protocol	Date	Characteristics	Spatial Steams	20 MHz Channels
802.11	1997	1,2 Mbps, infra Red, spread and DSSS, 802.11FH 2.4 GHz	1	1
802.11b	1999	1,2,5.5,11 Mbps, DSSS 2.4 GHz	1	1
802.11a	1999	6,9,12,18,36,48,54 Mbps - OFDM - 5 Ghz	1	1
802.11g	2003	6,9,12,18,36,48,54 Mbps - OFDM - 5 Ghz	1	1
802.11n	2005	MCS 1-15-23 1-3 SS, OFDM, 20,40 MHz, 2.4 and 5 GHz	1-3	1-2
802.11ac	2012/15	1-8 SS MCS 1-9, OFDM, 20-40-80-160 MHz, 5 GHz	1-8	1-8

The client market was slow to embrace 802.11n as most people where just starting to rely on smartphones and functionality was slowly increasing - the majority of smartphone clients where strictly 2.4 GHz capable. As time went on, we saw more functionality and subsequent adoption. BYOD and the concept of everyone bringing their favorite platform to work created demand and as hardware technology improved, the market started changing to dualband smart devices. At the peak of this revolution 802.11AC makes it's debut, and we are off to the races. The good news is that the market is catching up and we largely have consensus for the devices that people rely on at least supporting 5 GHz (not perfectly, but then it never is).

The point is that even if you update your network to the latest and greatest standard, the client market and what you have to support on your network remain somewhat variable and define how efficient you can use essentially the same spectrum. Backwards compatibility has always been a part of networking technologies, with wireless we are limited by airtime - and the efficiency we can gain in that finite airtime is affected by the technology in use as well as the number of clients you are supporting.

With 802.11n - we got an important boost, however it was barely keeping ahead of demand in most extreme cases and still falling behind in the worst examples. Not all devices supported more than a single spatial stream, or even bonded channels. The ability to use 40 MHz bonded channels was a waste of channel space unless your user base where all using laptops only.

Welcome to the 802.11ac evolution. Every client must support up to an 80 MHz channel in order to pass WFA certification, so that levels the playing field a bit. Spatial streams capabilities vary but tend to be matched to the size and power source of the device being implemented. Each spatial stream requires an additional radio

and corresponding power requirements still limit but have improved what is possible. Battery efficiency/capacity plays a role in the design decisions with smaller entry level devices still supporting only 1 SS. However, all devices can benefit from the expanded channel widths and along with what's being called wave 2 implementations - we can now simultaneously address individual single and dual spatial stream devices from the same BSS radio to achieve Multi User - Multiple Input Multiple Output (MU-MIMO). Multi User - Multiple Input Multiple Output radios allow us to service multiple single spatial stream clients in the same time block by using spatial stream diversity. Add to this that most clients are now releasing with 802.11ac radios - the time has never been better to start taking control of airtime efficiency and seeing big gains that were simply not possible only a couple of years ago.

In most environments, we are seeing overwhelming client support for 802.11n as a minimum. There are still pockets of legacy clients out there, but most of these are limited to application specific devices such as scanners, printers or devices that are purpose built for a specific industry tasks (retail, logistics). BYOD has enabled users to stay up with the latest technology trends by placing them in a continuous update cycle. If your implementation still relies on legacy clients - it is in your best interest to update these devices as soon as is feasible. While the new technologies allow for backward compatibility, they require more airtime and contribute heavily to what we now consider spectrum waste. It may still work, but you will never see most of the benefits gained in the current specifications while supporting the older less efficient radios and designs. For most users, this means 802.11n and the news is pretty good there. A mixed 802.11n and 802.11ac deployment has a tremendous amount of capacity and if designed properly will continue to service client needs over a wide range of demands.

Obviously, not every one has the same use case in mind – and RRM is designed to be flexible in its implementation to fit multiple use cases today without an exhaustive user understanding of the underlying RF challenges. RRM can be applied intelligently to multiple use models through the use of RF Profiles. Many new features can be found under the heading of HDX (High Density Experience) features, however all of these features actually support allowing RRM to do its job better, and under a wider range of conditions. We will touch on some of these features in this document, as they apply to managing user architectures, however full documentation for these features should be referenced in their deployment guides located here: [HDX High Density Experience deployment guide](#). Also refer to this document [Air Time Fairness \(ATF\) deployment guide](#) which covers additional protections which can be implemented for multiple roles ensuring airtime fairness for multiple deployment roles.

Most issues with RRM result from either too many (yes, too many) or not enough AP's/channels serving applications at a specific site. For the last few years trouble reports with RF are generally related to over saturation of the 2.4 GHz band. This should not be a surprise, increased density is mitigated by channel isolation and with only 3 channels that results in a much quicker need to reuse those channels which results in higher co-channel interference. The 2.4 GHz is largely considered a junk band for Wi-Fi users now as many devices that do not use Wi-Fi as well as many IOT devices take advantage of this band for ease of implementation as well as favorable propagation to power characteristics. These devices generally do not have the same requirements as a data or voice client, so it works out ok for them. More of these devices are coming and this will continue to make 2.4 GHz less favorable for most infrastructure users.

There is a finite limit to the number of radios that can operate in close proximity, and with many new devices entering the market, exceeding an RF designs capacity is becoming much more common. While this is sometimes initially blamed on RRM, RRM can only manage the resources that it has to work with. Architecture and radio placement need to be considered as part of the overall design. It is likely not good enough to just assume that the site survey that was conducted even 5 years ago meets the needs of today's user base. The good news is that once deployment density and design decisions have been adjusted to accommodate the increased demands on our networks today, RRM manages the result quite well. Poor planning can lead to unintended results with RRM. Improved diagnostics and instrumentation have made this information clearer, easier to understand, and more available at all levels of the organization.

This document seeks to provide you, the architect or technician with the details of how and why RRM makes its decisions. Knowing this will lead to better design decisions and quicker issue resolution. Continuing focus on Cisco's RF Excellence will continue to bring value to the users experience. A proven track record of changes and continuous development to stay ahead of the curve means that RRM is well established to continue to manage RF for our continuously growing needs.



CHAPTER 2

Radio Resource Management Concepts

RRM is a collection of algorithms which together provide a comprehensive management solution- the key algorithmic groups to be discussed here are:

- RF Grouping–The algorithm responsible for determining the RF Group Leader and members
- FRA–Flexible Radio Assignment – a new algorithm charged with identifying redundant radios resources and re-assigning the resource to a better role – strictly applicable to 2800, 3800 series AP’s with XOR radio for slot 0.
- DCA–A Global Algorithm, runs on the RF Group leader
- TPC–A Global Algorithm, runs on the RF Group Leader
- CHDM–A Local Algorithm, runs on each individual controller

In addition to RRM, there are several features which manage specific traffic types or client types which can greatly increase the spectral efficiency and assist RRM in providing a better experience for users. These will be discussed in context with the algorithms.

RRM is organized under the following Hierarchy:

RF Group Name ⇒ **RF Group leader(s)** ⇒ **RF Neighborhood(s)**

For any RF Group Name, multiple RF group Leaders may exist (a minimum of 2, one for 2.4 GHz and one for 5 GHz will always be present). An RF Group Leader will manage multiple RF Neighborhoods.

- [Pre-requisites and Assumptions, on page 7](#)
- [Key Terms , on page 8](#)

Pre-requisites and Assumptions

It is assumed that readers have a detailed knowledge of the following:

- Knowledge of and experience with common WLAN/RF design considerations (knowledge comparable to that of CWNA certification)
- Unified wireless access methodologies and hardware

Key Terms

Readers should fully understand the following terms used throughout this document with regard to Cisco's RRM algorithms:

1. **Signal:** refers to RF emanating from AP's belonging to the same RF group or our AP's.
2. **Interference:** Wi-Fi signals that do not belong to our network (rogues).
3. **Noise:** any signal that cannot be demodulated as an 802.11 signal. This can either be from a non-802.11 source (such as a microwave or Bluetooth device) or from an 802.11 source whose signal is below sensitivity threshold of the receiver or has been corrupted due to collision or interference.
4. **dBm:** an absolute, logarithmic mathematical representation of the strength of an RF signal. dBm is directly correlated to milliwatts, but is commonly used to easily represent output power in the very low values common in wireless networking.
5. **RSSI, or Received Signal Strength Indicator:** an absolute, numeric measurement of the strength of the signal in a channel.
6. **Noise floor:** the ambient RF Noise level (an absolute value expressed in dBm) below which received signals are unintelligible.
7. **SNR:** the ratio of signal strength to noise floor. This value is a relative value and as such is measured in decibels (dB).
8. **RF Group:** The logical container that an instance of RRM is configured through. All devices belonging to a single RF Network will be configured as a member of a particular RF group.
9. **RF Group leader:** The device where the algorithms for the RF group will be run. The RF group leader is either automatically selected through an election process or may be manually assigned through configuration. Two are required – one for each Spectrum band 2.4 and 5 GHz. And more may be present given the equipment and scale being employed.
10. **RF Neighborhood:** A group of AP's that belonging to the same RF group which can hear each other at ≥ -80 dBm. This is a physical grouping based on RF proximity.
11. **TPC:** Transmit Power Control is the RRM algorithm that monitors and manages transmit power level for all AP's in the RF group. There are two versions – each with their strengths – this document will cover both with recommendations.
12. **DCA:** Dynamic Channel Assignment is the RRM algorithm responsible for selecting the operating channel for all AP's in the RF group.
13. **CHDM:** Coverage Hole Detection and Mitigation—consists of the Coverage Hole Detection algorithm and the Coverage Hole Mitigation algorithm – CHD and CHM. This also has intersection with the HDX feature of Optimized Roaming as it relies on the measurements obtained from CHD.
14. **CM:** Cost Metric—an RSSI based metric which combines AP load, Co-channel interference, Adjacent channel interference and non wi-fi sourced interference into a goodness metric used by DCA to evaluate effective channel throughput potential.
15. **COF:** Coverage Overlap Factor – output of FRA algorithm, represents percentage of cell covered to -67 dBm by other 2.4 GHz radios.

16. **FRA:** Flexible radio Assignment – RRM algorithm that manages flexible radios and determines coverage redundancy and best roles for Flexible Radios to play based on coverage requirements and density.

**Note**

- RRM (and RF Grouping) is a separate function from inter-controller mobility (and Mobility Grouping). Confusion can arise through the default use of a common ASCII string assigned to both group names (RF Group, Mobility Group) during the initial controller configuration wizard. This is done for a simplified setup process and can be changed later.
- It is normal for multiple logical RF Group Leaders to exist. An AP on a given controller will help join their controller with another controller **only** if an AP or AP's from each controller can hear one another. In large Campus environments it is quite normal for multiple RF Neighborhoods to exist, spanning small clusters of buildings.

How Does RRM do and what it does?

The high level view of RRM is quite simple. It is a framework of services used to gather relevant over the air information and store it for analysis. Each AP spends time listening within its environment and collecting a variety of utilization statistics. The information collected drives many algorithms (wIDS and rogue detection are examples outside of RRM's algorithms). Each AP will gather information regarding Neighbors (Neighbor Discovery Protocol) channel conditions - Load, Interference, Noise. This information is collected by the RF Group Leader for the entire RF Group and used to determine the structure of the RF Domain first and break down the domain into RF Neighborhoods. An RF Neighborhood is a group of AP's that can hear one another, and as such must have channel and power solutions calculated together.

So the RF Group Leader is the designated controller that will run RRM Algorithm's on information that it collects from Member controllers. It does this by first identifying groups of AP's that are physically close enough to one another and organizing these into groups of RF Neighborhoods. The RF Group Leader is also the repository for the current RRM configurations (for channel and power) that will be used to configure the Algorithms for the RF Group.



CHAPTER 3

RRM Data Collection Activities

- [RRM Data Collection Activities](#) , on page 11

RRM Data Collection Activities

The RRM processes collect data to use in the organization of RRM as well as for processing channel and power selections for the connected APs. A base understanding of where RRM gets its information and how is essential to understand the algorithms. For now, how and where to configure monitoring tasks, and what that translates to in an operational environment will be covered. For each RRM algorithm discussed, what data is used and how it is used will be covered in the discussion.

The channel list monitored is configured under **Wireless > 802.11a/b > RRM > General-Noise/Interference/Rogue/CleanAir Monitoring Channels**

Figure 2: RRM General Configuration Dialogue

The screenshot displays the Cisco RRM General Configuration Dialogue for 802.11a. The interface includes a navigation menu on the left with categories like Access Points, Radios, Advanced, Mesh, ATF, RF Profiles, FlexConnect Groups, OEAP ACLs, Network Lists, and 802.11a/n/ac. The main content area is titled '802.11a > RRM > General' and contains several sections:

- Profile Threshold For Traps:** A table with four rows: Interference (0 to 100%) set to 10, Clients (1 to 200) set to 12, Noise (-127 to 0 dBm) set to -70, and Utilization (0 to 100%) set to 80.
- Noise/Interference/Rogue/CleanAir¹ Monitoring Channels:** A section with a 'Channel List' field and a 'Country Channels' dropdown menu.
- Monitor Intervals (60 to 3600 secs):** A table with three rows: Channel Scan Interval set to 180, Neighbor Packet Frequency set to 180, and Neighbor Timeout Factor (5 to 60) set to 5.
- Factory Default:** A section with a 'Set to Factory Default' button and a note: 'Set all Auto RF 802.11a parameters to Factory Default.'
- Foot Notes:** A note stating: '1. CleanAir monitoring is done on these channels only when the AP is in monitor mode.'

The choices for monitoring are:

- All Channels**— RRM channel scanning occurs on all channels supported by the selected radio, which includes channels not allowed in the country of operation. (Passive only). **Country Channels**—RRM channel scanning occurs only on the data channels in the country of operation. This is the default
- DCA Channels**— RRM channel scanning occurs only on the channel set defined by the DCA algorithm, which by default includes all of the non-overlapping channels allowed in the country of operation. However, you can modify the channel set to be used by DCA if desired.

Two types of off-Channel events are defined:

- Passive Dwell**—used to detect Rogues, and collect noise and interference metrics. The dwell time is 50 ms.
- Neighbor Discovery Protocol Tx**—used to send the NDP message from all channels defined under the monitor set.

The Channel Scan Frequency (**Wireless > 802.11a/b > RRM > General**) is 180 seconds (default value). This means all channel dwells must be completed within 180 seconds. So depending on the number of channels defined by the selection in the Monitor list, the interval between dwells will increase or decrease. For instance:

- Channel List = DCA, slot=0 (2.4 GHz) – DCA defines channels 1,6,11 for a total of 3 channels. So $180 \text{ (seconds)} / 3 \text{ (channels)} = 60$, the AP will go off channel every 60 seconds to listen.
- Channel List = Country, slot=1 (5 GHz) in the –A regulatory domain (US) with UNii 2e enabled - gives us 22 channels defined – so $180 \text{ (seconds)} / 22 \text{ (channels)} = 8.18$, the AP will go off channel every 8 seconds or so to listen for 50 ms.

Neighbor Packet Frequency is also defined on the same page, the default value for 7.6 and below is 60 seconds. This means that the radio must go off channel and send a single NDP packet for every channel defined by the channel monitoring list within 60 seconds. Using the same example from above where Channel List = Country and slot=1 (5 GHz) this translates to $60 \text{ (seconds)} / 21 \text{ (channels)} = 3 \text{ seconds}$, so for every three second the radio is sending an NDP packet on a channel other than the one it is currently serving.

In 8.0 and above Neighbor Packet Frequency is set to 180 seconds. The same as off channel frequency. And it is quite likely that we will drop the packet frequency all together in a future release. The default value is now 180 seconds - doing the same math as above this puts $180/21=8.57$ so every 9 seconds - in conjunction with the off channel passive - we will send 1 NDP packet per channel. See Neighbor Discovery Protocol - NDP Page 16 for more on NDP function.

Both the Channel Scan Interval, and the Neighbor Packet Frequency should be left at the default values. The Monitoring Channels list is set by default to use country channels; this is best for wIPS configurations. However if wIPS is not a primary concern, you can select DCA channels and reduce off channel activity to just the channels which you are using.



CHAPTER 4

RF Grouping

RRM RF Grouping is a central function for RRM. RF Grouping forms the basis for two management domains within the RF Network - the administrative and the physical.

- Administrative domain—For RRM to work properly it must know which APs and controllers are under our administrative control. The RF Group name is an ascii string that all controllers and APs within the group will share.
- Physical RF Domain—In order for RRM to calculate channel plans and power settings it is essential that RRM be aware of the RF Location of our APs and their relation to one another. Neighbor messaging uses the RF Group Name in a special broadcast message that allows the APs in the RF group to identify one another and to measure their RF Proximity. This information is then used to form RF Neighborhoods (A group of AP's that belong to the same RF Group that can physically hear one another's neighbor messages above -80 dBm) within the RF Group.

Each RF Group must have at least one RF Group Leader per band. The RF Group Leader is the physical device responsible for:

- Configuration
- Running the active algorithms
- Collection and storage of RF Group Data and metrics

There will be a minimum of two RF Group Leaders, one for each band 802.11b and 802.11a (2.4 and 5 GHz) respectively. While RF Group Leaders for different bands can coexist on the same physical WLC, they often do not. It's also not uncommon for there to be more than one group leader per band in larger systems that have geographic diversity.

Two modes of RF grouping algorithm exist in the system today. RF Group Leaders can be selected automatically (legacy mode) or assigned statically. Both methods of assignment were overhauled with the addition of static RF Grouping in version 7.0 of the CUWN code.

- [How RF Groups are formed](#) , on page 16
- [Neighbor Discovery Protocol—NDP](#) , on page 16
- [RF Group Leader Election](#) , on page 21
- [RF Group Scalability](#) , on page 24
- [RF Group Backward Compatibility](#) , on page 26
- [WSSI and WSM, WSM2 Modules and RRM](#) , on page 26
- [Troubleshooting RF Grouping](#) , on page 26

How RF Groups are formed

When the WLC initializes as new, it creates a unique Group ID using the IP address of the WLC and a Priority Code. The Priority Code is assigned based on the controller model and MAX license count (hardware limit) to create a hierarchical model and ensure that the controller with the most processing capacity is assigned the job of GL (Group Leader). The Group ID and the RF Group Name will be used together in messages to other WLC's and AP's to identify them. Devices having the same RF Group Name will interoperate as members of the same RF Group.

The current controller hierarchy is as such:

8500 > 7500 > vWLC(large) > 5520 > 5760 > WiSM2 > 5508 > vWLC(small) > 3850 > 2500



Note See [Table 2: Ports required for RRM operation](#) below along with RF group scalability numbers below.

When comparing Group IDs for leader election, the priority code is primary criteria and IP address is secondary. For instance, if there are 3 other controllers, none of which has the same or higher priority code than myself - I become the Group Leader. If all 3 have the same priority code as myself, then the one with the highest IP address wins and assumes the GL role.

For two WLCs to form an RF Group there is an infrastructure as well as OTA (Over The Air) component:

- WLCs must be reachable to one another on the distribution network
- They must each also have at least one AP that can hear the other's NDP messages above -80 dBm

The distribution network communicates over unicast UDP:

Table 2: Ports required for RRM operation

	Source Port	Destination Port
RRM Manger 11b(11a)	12134(12135)	12124(12125)
RRM Client 11b(11a)	12124(12125)	12134(12135)

The OTA component relies on two functions NDP - Neighbor Discovery Protocol and collection of off channel metrics. Think of NDP as the Off Channel TX cycle, and monitoring of off channel metrics as the off channel RX cycle. Both NDP and monitoring are critical to the topic of RF Grouping and RRM in general, so we'll discuss them here before going any deeper.

Neighbor Discovery Protocol–NDP

One of the most unique things about Cisco's RRM implementation is that it uses Over The Air (OTA) messages and runs centralized even in large deployments. This gives us the advantage of being able to monitor and manage all APs and their RF experience from a single point in the network. Not only manage - but understand how every AP relates to any other AP in the RF Group/Neighborhood. This is unique in the industry as most other implementations run AP to AP at the edge in a distributed fashion with only configuration elements being managed centrally.

Neighbor Discovery Protocol or NDP, is sent from every AP/Radio/Channel every 60 seconds or less. The NDP packet is a special broadcast message that APs all listen for and it allows us to understand how every radio on every channel hears every other radio. It also gives us the actual RF path loss between APs.

Neighbor messages are sent to a special Multicast address of **01:0B:85:00:00:00**, and are done so:

- At the Highest Power allowed for the Channel/Band
- The Lowest data rate supported in the band

For 802.11b this means that the message is sent at power level 1 (always the highest power for a particular radio) at 1 Mbps, and for 5 GHz radio's 6 Mbps. This function is hard coded into the radio firmware, there is no user control. NDP power and modulation is not changed by user configured data rates or power levels.

For 802.11b this means that the message is sent at power level 1 (always the highest power for a particular radio) at 1 Mbps, and for 5 GHz radio's 6 Mbps. This function is hard coded into the radio firmware, there is no user control. NDP power and modulation is not changed by user configured data rates or power levels.

An NDP message contains the following information:

Table 3: Contents of NDP Packet

Field Name	Description
Radio Identifier	Slot ID for the sending radio
Group ID	IP Address and Priority code of senders WLC
Hash	RF Group name converted to a hash for authentication
IP Address	The IP address of the sending AP's RRM Group Leader
Encrypted ?	Are we using Encrypted NDP ?
Version	Version of NDP
APs Channel	The operating channel of the sending radio
Encryption Key Length	Key Length
Encryption Key Name	Key Name
Message Channel	The channel the NDP was sent on
Message Power	The power (in dBm) the message was sent at
Antenna	Antenna pattern of the sending radio

When an AP hears an NDP message, it:

- Validates that the message is from a member of its RF Group (hash); if not it is dropped
- If valid forwards the message along with the received channel and RSSI to the controller

The forwarded message is added to the neighbor database, which in turn is forwarded to the RF group leader periodically. For each AP, each radio can store up to 34 neighbors ordered by RSSI high to low.

Post processing of this information develops 2 distinct measurements:

- RX Neighbors: How I hear other APs
- TX Neighbors: How other APs hear me

Neighbor entries on the controller are pruned at different default intervals depending on your code release. Before 8.0 the pruning interval is fixed at 60 Minutes. In 8.0, it is hard coded for 15 minutes. In 8.1 and above there is a new function–Neighbor Timeout Factor. On the GUI wireless=>802.11a/802.11b=>RRM=>General.

The screenshot shows the Cisco Prime GUI for configuring 802.11a RRM General settings. The navigation tabs at the top are MONITOR, WLANs, CONTROLLER, WIRELESS, SECURITY, and MANAGEMENT. The breadcrumb trail is 802.11a > RRM > General. The 'Profile Threshold For Traps' section includes: Interference (0 to 100%) set to 10, Clients (1 to 200) set to 12, Noise (-127 to 0 dBm) set to -70, and Utilization (0 to 100%) set to 80. The 'Noise/Interference/Rogue/CleanAir¹ Monitoring Channels' section has a 'Channel List' field and a 'Country Channels' dropdown. The 'Monitor Intervals (60 to 3600 secs)' section, highlighted with a red box, contains: Channel Scan Interval set to 180, Neighbor Packet Frequency set to 180, and Neighbor Timeout Factor (5 to 60) set to 5. Below this is the 'Factory Default' section with a 'Set to Factory Default' button. A 'Foot Notes' section at the bottom states: '1. CleanAir monitoring is done on these channels only when the AP is in monitor mode.'

The Neighbor Timeout Factor allows the user to directly configure how long to maintain the neighbor list on the controller.

Neighbor Timeout Factor is the multiplier that will be applied to Neighbor Packet Frequency (seconds) - $NTF \times NPF =$ Pruning interval in seconds. In the example above, the Neighbor Packet Frequency is 180 seconds; the Neighbor Timeout Factor is 5, so $5 \times 180 = 900$ seconds/15 Minutes (Default Configuration).



Note 15 Minutes is the default pruning interval in WLC version 8.0 and above. This is fine in most installations. Under sustained heavy loads (high client density, Channel Utilization consistently 60% or higher), or in environments that enjoy a lot of neighbor interference this may be inadequate to ensure stable neighbor resolution. Channel Changes will increase if NDP is changing between DCA cycles (default 10 minutes). If an AP's neighbors are changing then DCA will supply a different answer to the new equation. Changing the Neighbor Timeout Factor to 20 restores 3600 second (1 hour) the legacy interval.

The Neighbor Pruning interval affects the stability of everything that displays or uses neighbor information. Cisco Prime maps and show auto RF outputs use the neighbor list. NDP is a packet sent over the air interface,

and if off channel transmission is suspended it simply waits 180 (default - should never be longer) seconds to get back to the channel and retry the off channel Tx . High Channel Utilization and poor SNR in a busier network will lower the opportunities NDP has to transmit. A Voice SSID will defer off channel activity any time there is has been traffic in the voice QOS que within the last 100 ms. Retaining the historical information over a period of time compensates for short duration inconsistencies.

What does this mean in terms of monitoring performance? If a radio is admin disabled or disconnected, that radio will exist in the neighbor list until the NDP pruning interval kicks in. You will still see it in all lists and reports - until it times out. If a new neighbor is discovered the list is flushed and refreshed in its entirety to capture what the new neighbor can contribute.

What does all this mean? In short, leave the default values alone unless you have a reason to change them, and then only longer - not shorter than 15 minutes. When to set the interval to longer? If you upgrade from 7.x to 8.x and see a definite increase in channel changes after reaching steady state operations (no new bandwidth changes, DBS added, DCA start up isn't running) then increase the Neighbor Timeout Factor in increments of 5 (fifteen minutes for each increase of 5) until you see "Your" normal operations.

You can observe neighbor messages over the air using a packet capture tool and filtering on the multicast address **01:0B:85:00:00:00**.

Figure 3: Sample Packet Capture of NDP Messaging

18	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0.000000
24	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:01:00.005975
29	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:01:59.910124
34	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:02:59.915850
40	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:03:59.922653
46	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:04:59.930237
51	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:05:59.935790
56	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:06:59.946686
62	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:07:59.950317
68	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:08:59.955871
74	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:09:59.964819
80	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:10:59.971166
96	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:13:59.990219
101	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:14:59.994158
115	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:17:59.911287
120	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:18:59.919573
125	Airspace:52:A0:A0	01:0B:85:00:00:00	802:11 Data	SNAP	0:19:59.925931



Caution

Unless you use the AP sniffer mode to capture the packets the RSSI values you see in your capture tool will likely be different from what is recorded in the neighbor lists - AND - the neighbor list will quite likely have more entries than you can hear simply because the APs radio sensitivity and position are generally favorable (on the ceiling) to a mobile tool's.

NDP and DFS

NDP is transmitted on all regulatory channels selected under monitor channels list. However DFS channels represent a special case as in order to transmit on a DFS channel a station must either be a Primary, or in the case of the client - associated directly to a legal Primary. In order to become a Primary, an AP must monitor the channel for 60 seconds to verify that no Radar is present before transmitting on that channel. A client hearing a beacon on a DFS channel can infer that the channel is owned by a Primary and transmit to that Primary. In order for us to transmit NDP on a channel in the DFS bands that we are not the Primary of, we need to first hear either a Beacon or a directed Probe from a client in order to mark that channel as clear, then we can follow up with a transmitted NDP packet within 5 seconds. If there are no other APs, and there are no clients on other DFS channels, we will never send an NDP on any DFS channel except the one on which we are the Primary.

What do we use NDP for?

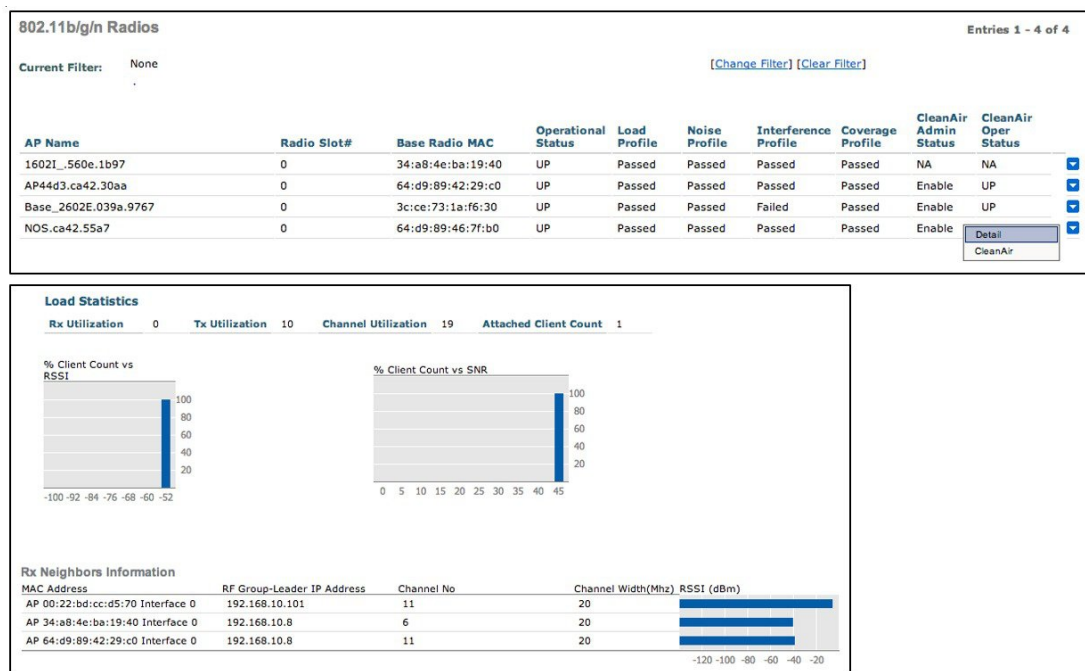
NDP forms the foundation for our understanding of the RF Propagation domain and inherent path losses encountered within the deployment. NDP is very important to RRM, and as such it should go without saying then that if NDP is broken, RRM is broken. NDP is used first by the RF Grouping algorithm, but also by:

- TPC (Transmit Power Control) - third neighbor opinion of our NDP or the basis for calculation as in TPCv2
- FRA (Flexible Radio Assignment) NDP messages from all AP's are the basis for Coverage Overlap Factor
- Rogue Detection—any AP that is either not sending NDP, or sends an unintelligible NDP is considered a rogue
- CleanAir Merging and PMAC functions—CleanAir uses neighbor relations to understand if interference reports are coming from AP's that are close enough to all hear the same interference device
- CMX—for AP RF Distance and pathloss measurements

All of these things require a detailed understanding of where the APs are in relation to each other in RF. And, that's what NDP does.

You can see neighbor relations in several places within the system, on the WLC select **Monitor=>Access Points=>802.11a/b=>details=>RX Neighbors Information**

Figure 4: Examples of where to see Neighbor Relations per AP



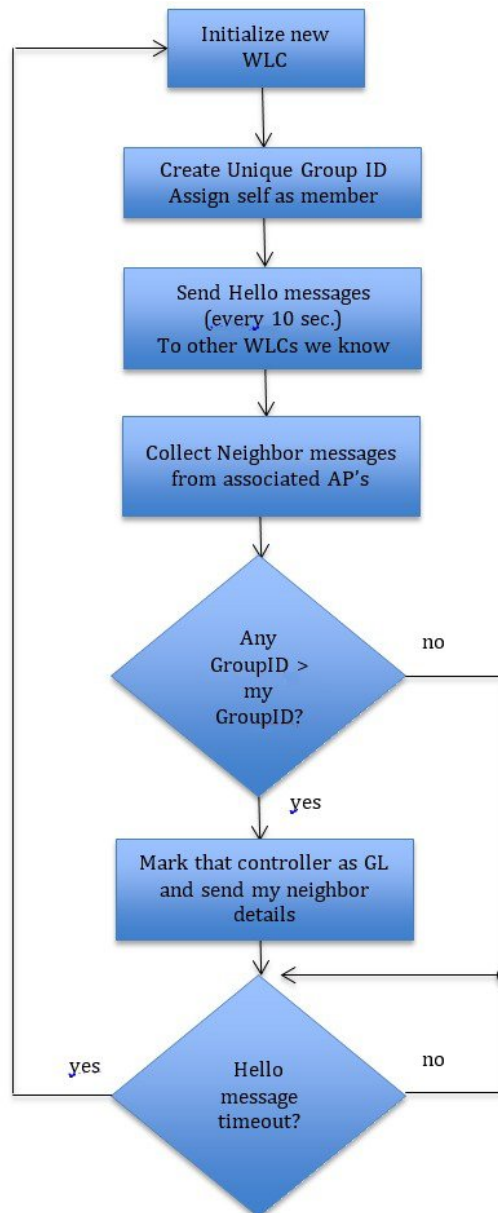
Or from the command line:

```
(Cisco Controller) show ap auto-rf 802.11a/b {AP_Name}
(Cisco Controller) show ap auto-rf 802.11-abgn {AP_Name} (for Flexible Radios)
```

RF Group Leader Election

Now that we've discussed the components, let's have a look at what happens when a brand new controller is initialized and an RF group is formed. We'll cover automatic Grouping first, and then identify how this differs with Static Grouping assignment last. See the flow chart below for RRM state machine initialization:

Figure 5: RF Grouping Process Flow Chart



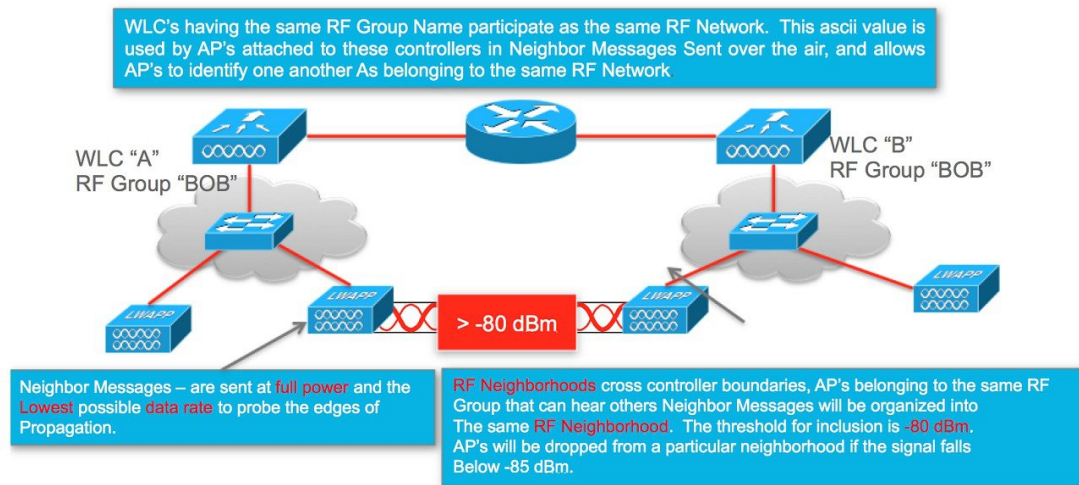
When a WLC is initialized for the first time the only WLC that it's aware of is itself. The WLC generates the GroupID and initially assumes the role of Group Leader taking the RF Group name entered during initial

startup configuration and passing this to any connected AP's for use in their neighbor string. The new leader will have itself as a member. The WLC initializes the hello timers and begins sending over the wire to other WLCs that it knows about. The Hello message is a unicast that is sent to all WLCs stored in the RF Group History. If Auto Grouping, having just been initialized, this list is empty. If Static configuration, then the list is or will be populated by manual assignment.

For Auto Grouping, the received OTA NDP message contains the sender's WLC Group ID and RF Group Hash as well as the IP address of the senders RF Group leader. The new WLC compares all received Group IDs, and any one having a larger value than our own then becomes our Group Leader. RF Grouping completes and the election process ends. Every 10 seconds we'll receive a hello message from our Group Leader that serves as a heartbeat for the RF group. If the Hello messages stop coming - we'll assume that the RF Group has changed - and the election process begins again. By this time we'll normally have a list of WLCs to send Hello packets.

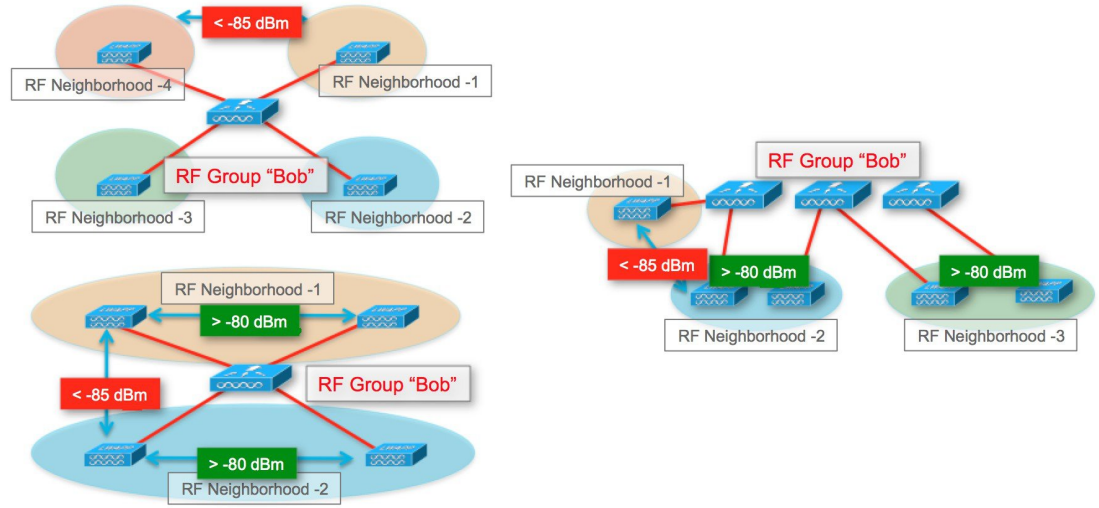
Once the Group Leader is established, neighbor lists from all members will be sent to the GL and APs in the group will be formed into RF Neighborhoods or groups of APs that are close enough to require RF Power and channel be calculated together. For another AP to belong in our neighborhood we'll need to see that APs neighbor message at -80 dBm or above. Once an AP is added to a neighborhood, as long as we see the neighbor message at or above -85 dBm it remains part of the neighborhood. Any neighbor message below -85 dBm is dropped. The neighbor list purges every 60 minutes up through version 8.0 code. In 8.1 the neighbor retention time was adjusted to match 3x the scan interval (so at default 180 seconds, the neighbor list will be purged every 15 minutes). Any AP that remains consistently below -85 dBm will be purged from the list and the neighborhood. In this way, we identify groups of APs that are in the same geographic location.

Figure 6: RF Group and Neighborhood example



RF Neighborhoods can span multiple controllers, or a single controller can be managing multiple neighborhoods, some examples are presented here.

Figure 7: Examples of how RF Neighborhoods are organized



RF Grouping Automatic mode

The default mode of RF grouping is the legacy method of forming RF Groups. You can view the current status of the RF grouping algorithm, learn the identity of the Group Leader and members, and on the RF Group leader WLC see a count of current WLC's and AP's contained in the group on the WLC:

Wireless=>802.11a/b=>RRM=>RF Grouping =>group mode

Figure 8: RF Grouping Configuration Dialogue

802.11a > RRM > RF Grouping

RF Grouping Algorithm

Group Mode	auto	<input type="button" value="Restart"/>
Group Role	Auto-Leader	
Group Update Interval	600 secs	
Group Leader	Cisco_69:9a:64 (192.168.10.8) (::)	
Group Name	test2	
Protocol Version(MIN)	100(30)	
Packet Header Version	2	
Maximum/Current number of Member	20/1	
Maximum/Current number of AP	500/5	
Last Group Update	140 secs ago	

RF Group Members

*If the member has not joined the group, the reason of failure will be shown in brackets

Controller Name	IP Address(Ipv4/Ipv6)
Cisco_69:9a:64	192.168.10.8

Static RF Grouping

In version 7.0 a static method of selecting an RF group leader was introduced. This allows a more deterministic outcome to the grouping process. The Group ID is not needed here (Priority Code and IP address of the WLC) but the Priority Code will be compared to members; this prevents a lower capacity WLC from becoming the group leader of a higher capacity WLC.



Note You cannot assign a 2504 to be the group leader and have a 5508 added as a member.

Static grouping allows the user to designate a particular WLC as the Static leader, and manually add the members to be managed. Members must be in auto mode, and running a compatible version of RRM. Once the Static leader is assigned, members are assigned to it and a special join message is sent to prospective members that overrides the automatic function and provides the member with a new Group leader assignment.

Under **Wireless=>802.11a/b=>RRM=>RF Grouping**

Figure 9: Example of Static and Automatic RF Grouping Configurations

802.11a > RRM > RF Grouping Apply

RF Grouping Algorithm

Group Mode: leader leader Restart

Group Role: Static-Leader

Group Leader: Cisco_69-9a-64 (192.168.10.8) (::)

Group Name: test2

Protocol Version(MIN): 100(30)

Packet Header Version: 2

Maximum/Current number of Member: 20/2

Maximum/Current number of AP: 500/65

Static GL

RF Group Members

Controller Name	IP Address(Ipv4/Ipv6)
Cisco_did:fb:e4	192.168.10.20
Cisco_69-9a-64	192.168.10.8

802.11a > RRM > RF Grouping Apply

RF Grouping Algorithm

Group Mode: auto auto Restart

Group Role: Static-Member

Group Leader: Cisco_69-9a-64 (192.168.10.8) (::)

Group Name: test2

Protocol Version(MIN): 100(30)

Packet Header Version: 2

Static Member

RF Group Members

Controller Name	IP Address(Ipv4/Ipv6)
Cisco_69-9a-64	192.168.10.8
Cisco_did:fb:e4	192.168.10.20

Changing the group mode to leader, and hitting apply opens the member assignment dialogue. You then assign members and when complete select restart to re-initialize group leader elections for the new assignments. In order for a member to be added, the prospective member must be in Auto grouping mode - else it assumes it is its own leader. The new Group Leader controller is automatically added as the first member. Additional members can be added manually at any time. Member controllers should stabilize within 10 minutes or so once the RF Group is restarted.

There are no rules on spectrums, meaning leaving 5 GHz in Auto, and 2.4 GHz as Static is just fine. Or do both static, but on different controllers, your choice. The sky is the limit as both interfaces are different RF Group instances. However, and this is always good advice, Cisco best practice is keep it simple.

RF Group Scalability

The maximum size for an RF Group is dependent on the model of the controller and the number of APs physically connected. The maximum sizes for RF groups can be calculated using the following rules. An RF Group can contain up to 20 WLCs, and have the noted Maximum APs.

Table 4: WLC RF Grouping Hierarchy and Scalability

Group Leader WLC	Maximum APs	Maximum AP per RF Group
2500	75	500
WLCM2	50	500
3850	50	500
vWLC (small)	200	1000
5508	500	1000
5520	1500	3000
WiSM2	1000	2000
5760	1000	2000
vWLC (large)	2000	2000
7500	6000	6000
8500 /8540	6000	6000

What happens if I exceed the RF group size? A popular question, relax, the world does not come to an end, please read on.

If you exceed the maximum allowed number of APs for a given RF Group, the group simply splits and creates a new RF Group Leader using the same RF Group Name on the controller that the AP joined to create the condition. This sounds a lot worse than it is, and in practice most folks are generally not even aware of it until they look for the RF Group Leaders and notice that there is more than one per band.

What's the downside of having two or more RF Groups? There are now more RF group leaders that have to be addressed when you want to make configuration changes (additional GLs for both 802.11a and 802.11b assuming dual radio APs). This adds some complexity, but is easily managed with controller templates and configuration audit tools. Two AP's belonging to two different RF groups will not see one another as neighbors as they have different hashes of the same RF Group name. For this reason, some planning of which AP's go to which controllers is important. It is best to plan for AP's that are co-located to be on the same controller or under the same RF Group Leader.

The RF Group Leader stores the global RRM parameters for the RF Group and if a new Group Leader is created, that new WLC's RRM configurations will govern the global group settings. If you've not taken advantage of config audit features under Monitor=>RRM in NCS or Prime Infrastructure, it is possible that you have different configurations on the new GL (the worst case scenario). This could be quite disruptive if the configurations are seriously out of synch. However if the configurations are matching, DCA and TPC will mitigate the boundary quite seamlessly.

When planning your network keep these things in mind:

1. Groups of APs that are close enough to hear one another as neighbors (above -80 dBm) should reside in the same RF Group.
2. If you have multiple controllers, geographically group your AP's on like RF Groups of controllers – depending on your configuration static assignment of GL's and members may be the best approach.

3. Two otherwise diverse groups of APs only require a single AP in common to join together and form a neighborhood.
4. If you have two groups of APs that are joined together by only a few APs, you can force a split by creating a second RF group. This will change the RF group advertised in NDP messages and separate the two groups.

RF Group Backward Compatibility

In version 7.2 RF Profiles were introduced. This represented a major change to how RRM operated. RF Profiles assigned to AP groups could be configured differently from the global RF Group. Versions from 7.2 and forward are not compatible in an RF Group with older versions. About the same time Converged Access was introduced, and feature parity (RF Profiles) was not achieved immediately. Check the [Cisco Wireless Solutions Software Compatibility Matrix](#) Inter Release Controller Mobility table to ensure compatibility for mixed release integrations. Pay attention to the notes. From version 7.5 on, there are feature differences, however all can be successfully included in a single RF Grouping.

WSSI and WSM, WSM2 Modules and RRM

One of the great additions to make if you own a 3 series AP (3600, 3700, 3800) and can install a module is the Wireless Security Module which contains radios strictly dedicated to monitoring. There are two models of this module now, but both operate with respect to RRM in the same way - they off load the off channel functions of the serving radios to the module. This allows the serving radios to remain dedicated to the channel they are serving and increases the dwell time on each channel based on the role of the dwell (i.e. off channel, location, WIPS, CleanAir). This offloading is a benefit in almost every situation in that it brings a higher resolution to the data that is being collected with longer and more frequent dwells driving the collection. The module relies on its own internal antenna's for collection and the antenna pattern is matched with that of an internal antenna AP model.

One caveat to this approach however is external highly directional antennas used in High Density designs (most omni patch antenna's are just fine and this does not apply to them). The data that is being collected relies on the over the air results matching what the AP and serving interfaces actually see. In a High Density solution using the Stadium antennas, this will differ significantly. For this reason, achieving a good channel plan for the antennas used in the design requires shutting down the module and collecting over the air metrics using the AP's native interfaces and antenna to develop a good channel solution. Once this has been done, freezing DCA will allow the module to continue driving benefit without negatively impacting the channel and power solution.

Troubleshooting RF Grouping

RRM Data Collection

Data Collection at the AP level can be viewed using debugs.

debug capwap rm measurements—the output should be self explanatory. This is useful to compare the intervals of different intervals at the AP.


```

AP44d3.ca42.30aa#deb capwap rm measurements
CAPWAP RM Measurements display debugging is on
AP44d3.ca42.30aa#

*Jan 14 11:36:57.403: CAPWAP_RM: Timer expiry
*Jan 14 11:36:57.403: CAPWAP_RM: Interference onchannel timer expired, slot 1, band 0
*Jan 14 11:36:57.403: CAPWAP_RM: Starting rx activity timer slot 1 band 0
*Jan 14 11:36:57.419: CAPWAP_RM: RRM measurement completed. Request 2003, slot 1 status
TUNED
*Jan 14 11:36:57.483: CAPWAP_RM: RRM measurement completed. Request 2003, slot 1 status
SUCCESS
*Jan 14 11:36:57.483: CAPWAP_RM: noise measurement channel 48 noise 93
*Jan 14 11:37:06.355: CAPWAP_RM: Timer expiry
*Jan 14 11:37:06.355: CAPWAP_RM: Interference onchannel timer expired, slot 1, band 0
*Jan 14 11:37:06.355: CAPWAP_RM: Starting rx activity timer slot 1 band 0
*Jan 14 11:37:06.423: CAPWAP_RM: RRM measurement completed. Request 2004, slot 1 status
TUNED
*Jan 14 11:37:06.487: CAPWAP_RM: RRM measurement completed. Request 2004, slot 1 status
SUCCESS
*Jan 14 11:37:06.487: CAPWAP_RM: noise measurement channel 52 noise 92
*Jan 14 11:37:08.711: CAPWAP_RM: Timer expiry
*Jan 14 11:37:08.711: CAPWAP_RM: Neighbor interval timer expired, slot 0, band 0
*Jan 14 11:37:08.711: CAPWAP_RM: Scheduling neighbor request on ch index:
*Jan 14 11:37:08.711: CAPWAP_RM: Sending neighbor packet #2 on channel 11 with power 1
slot 0
*Jan 14 11:37:08.823: CAPWAP_RM: Request id: 4011, slot: 0, status 1

```

For a granular look at the neighbor activity at the AP specifically: Debug capwap rm neighbors.

```

*Jan 14 17:29:36.683: LWAPP NEIGHBOR: NDP Rx: From 64d9.8946.7fb0 RSSI
[raw:norm:avg]=[-37:-39:-38] Channel [Srv:Tx]=[1 :6 ] TxPower [Srv:Tx]=[4 :22 ]

```

This debug is about the NDP received from a neighbor.

NDP RX from x.x.x.x RSSI (raw:norm:avg)=(n:n:n) Channel (Srv:Tx) SRV = the channel the sending AP is serving clients on, TX= the channel the message was sent on. TxPower (Srv:Tx) Srv= the power in dBm that the AP is currently serving clients at Tx = the power in dBm that the NDP message was sent at.

```

*Jan 14 17:29:37.007: LWAPP NEIGHBOR: NDP Tx: Channel [Srv:Tx]=[64 :64 ] TxPower [Srv:Tx]=[2
:17 ]

```

NDP TX-this sends a NDP message, channel (Srv:Tx) Srv - the channel we are serving clients on, Tx - the channel we sent the NDP message on. TxPower (Srv:Tx) Srv - power in dBm we are serving clients at, Tx - the power in dBm that we sent the message at.

```

*Jan 14 17:29:40.007: LWAPP NEIGHBOR: skipping chan 100; not clear for DFS
*Jan 14 17:29:43.007: LWAPP NEIGHBOR: skipping chan 104; not clear for DFS
*Jan 14 17:29:46.007: LWAPP NEIGHBOR: skipping chan 108; not clear for DFS

```

Channels not clear for transmit for DFS:

```

*Jan 14 17:29:48.299: LWAPP NEIGHBOR: Updating existing neighbor 34a8.4eba.194f(1), rssi
-51 on channel: 48 with encryption: 0
*Jan 14 17:29:48.299: LWAPP NEIGHBOR: Neighbor update 34a8.4eba.194f(avg -45), new rssi
-45, channel 48

```

An update of a change in a neighbor's information being sent to the controller and ultimately the RF Group Leader.

Neighbor messaging issues are pretty easy to spot, if NDP is broken, then APs that are next to one another will not have a relationship.

RF Grouping Trouble

Often the reason for trouble with RF groups is simply compatibility. Since version 7.0 of code and the introduction of Static Grouping, there have been many changes to RRM and how it behaves. Backward compatibility has been preserved where it could be, however, changes in the RRM header were required to implement some of these changes and the header version number is checked on grouping.

RRM Header version 30.0 was used through version 7.0, version 30.1 was introduced with release 7.2 and RF Profiles. 7.3 added more structure to RF Profiles and also saw the introduction of Converged Access Architecture, the header version changed to 30.2. This is the last change required for the foreseeable future.

Table 5: Excerpt of IRCM RRM compatibility matrix

CUWN Service	4.2x	5.0x	5.1x	6.0x	7.0x.x	7.2.x.x	7.3.x.x	CA10.1	7.4.x.x
Radio Resource Management (RRM)	X	-	-	X	X	-1	-2	-3	-2



- Note**
1. In the 7.2.x.x release, RF Groups and Profiles were introduced. RRM for 7.2.x.x and later releases is not compatible with RRM for any previous release.
 2. In the 7.3.x.x release changes were made to RF Profiles, not backwardly compatible with 7.2.
 3. CA 10.1 release will form RF groups with 7.3.101.0 - however there is NO support for RF Profiles.

RF Grouping functions can be observed on the controller using the "*sh advanced 802.11a/b group*" command.

```
(controller) > show advanced 802.11b group
Radio RF Grouping
 802.11b Group Mode..... STATIC
 802.11b Group Update Interval..... 600 seconds
 802.11b Group Leader..... GRP_Leader (1.2.3.4)
   802.11b Group Member..... GRP_Member (1.2.3.4)
   802.11b Group Member..... GRP_Member (1.2.3.5)
 802.11b Last Run..... 594 seconds ago
```

You can view the status on the **WLC GUI at Wireless=>802.11a/b=>RRM=>RF Grouping:**

Figure 10: RF Grouping information on the WLC GUI

802.11b > RRM > RF Grouping

RF Grouping Algorithm

Group Mode: auto

Group Role: Auto-Leader

Group Update Interval: 600 secs

Group Leader: Cisco_69:9a:64 (192.168.10.8)

Last Group Update: 526 secs ago

RF Group Members

*If the member has not joined the group, the reason of failure will be shown in brackets

Controller Name	IP Address
Cisco_69:9a:64	192.168.10.8
Cisco_dd:f8:e4	192.168.10.20
Cisco_dc:bb:24	192.168.10.30

For Automatic RF Grouping, if a WLC that you feel certain should be in an RF Group somehow will just not join, it is either because:

- The RF Group size is above capacity
- The RF Group Name assigned to the WLC is different
- There is no network path for Hello Messages

For Static RF Grouping, if an assigned member will not join the statically assigned group leader - the most common reason is version compatibility, RF Group Name and Controller Hierarchy are high on the list to evaluate.

Useful Debugs from the WLC console

- debug airwave-director error—displays all errors for RRM and RF Grouping
- debug airwave-director group—shows RF Grouping activities in a steady state network, this equates to a split calculation ensuring that the RF Group still meets the criteria on size and neighbor relations.

You can force a re-grouping to occur by selecting the reset button on the **Wireless=>802.11a/b=>RRM=>RF Grouping menu**

Watch the RF group form

```
*emWeb: Jan 16 18:46:49.717: Airewave Director: Group 802.11bg attempting to remove entry
C0.A8.0A.14.00.4B, IP Addr 192.168.10.20
*emWeb: Jan 16 18:46:49.717: Airewave Director: removing entry C0.A8.0A.14.00.4B from
802.11bg group
*emWeb: Jan 16 18:46:49.719: Airewave Director: Group 802.11bg attempting to remove entry
C0.A8.0A.1E.00.32, IP Addr 192.168.10.30
```

*emWeb: Jan 16 18:46:49.719: Airewave Director: removing entry C0.A8.0A.1E.00.32 from 802.11bg group

Deleting the current members

*RRM-MGR-2_4: Jan 16 18:46:49.746: Airewave Director: adding entry C0.A8.0A.08.01.F4 (500) to 802.11bg group

Current group Leader-adding itself as a member

*RRM-MGR-2_4: Jan 16 18:49:03.614: Airewave Director: Group received Join Request from 802.11bg group C0.A8.0A.14.00.4B(63131), IP addr 192.168.10.20

RF Group Leader receives a Join Request

*RRM-MGR-2_4: Jan 16 18:49:03.614: Airewave Director: Deny join request from IP addr 192.168.10.20 to 802.11bg group C0.A8.0A.14.00.4B(63131) with reason Non matching group ID

Join Denied, non matching group ID

*RRM-MGR-2_4: Jan 16 18:51:07.651: Airewave Director: Group received Join Request from 802.11bg group C0.A8.0A.14.00.4B(63131), IP addr 192.168.10.20

Second Join Request received

*RRM-MGR-2_4: Jan 16 18:51:07.651: Airewave Director: Member in join request from source IP addr 192.168.10.20 to 802.11bg group, member IP 192.168.10.20

our Id 500 srcType 75

*RRM-MGR-2_4: Jan 16 18:51:07.651: Airewave Director: adding entry C0.A8.0A.14.00.4B (75) to 802.11bg group

The request is honored and we add the WLC to the group

*RRM-MGR-2_4: Jan 16 18:56:59.958: Airewave Director: Group received Join Request from 802.11bg group C0.A8.0A.1E.00.32(63131), IP addr 192.168.10.30

The second WLC sends it's join request

*RRM-MGR-2_4: Jan 16 18:56:59.958: Airewave Director: Member in join request from source IP addr 192.168.10.30 to 802.11bg group, member IP 192.168.10.30

our Id 500 srcType 50

*RRM-MGR-2_4: Jan 16 18:56:59.958: Airewave Director: adding entry C0.A8.0A.1E.00.32 (50) to 802.11bg group

And it is added to the group—complete

*RRM-MGR-2_4-GRP: Jan 16 18:57:20.909: Airewave Director: prep to join 802.11bg group C0.A8.0A.65.03.E8(63126) due to rssi -8

*RRM-MGR-2_4: Jan 16 18:57:36.839: Airewave Director: Group 802.11bg attempting to join group IP Address 192.168.10.101, ctrl count 3

Now our group leader attempts to join another WLC whose Group ID is higher than ours - with a controller count of 3 (himself and the two new additions)

*RRM-MGR-2_4: Jan 16 18:57:36.857: Airewave Director: Group received join failure from 802.11bg C0.A8.0A.65.03.E8(63126) (192.168.10.101) for reason

Not a configured static member

*RRM-MGR-2_4: Jan 16 18:57:36.857: Airewave Director: Group validated join failure from 802.11bg C0.A8.0A.65.03.E8(63126) for reason Not a configured static member

But we are denied access - 192.168.10.101 is configured as a static Group leader, and we are not configured as members under that group.

Summary of the Reason Codes

1. **Invalid IP:** This suggests that the controller IP is invalid or doesn't match against the controller system name.
2. **Group Size exceeded:** When the operational limits of a leader controller has reached either because of AP numbers or number of member controllers additions, the leader rejects addition of more controllers and display this reason for rejection.
3. **Invalid Group order:** If the grouping order is not in the way they have been formulated for reasons such as memory corruption or if the data-structures have been corrupted while transmission or a unknown controller type is attempting to join –Then this error msg is displayed.
4. **Source Not Included:** No valid source identification.
5. **Weak Signal Strength:** (Not applicable to static RF grouping) nearest neighbor is not close enough
6. **Join Pending:** When a member controller is waiting to complete and exit one RRM state to another, when it can join as a member.
7. **Not a Manager:** An unlikely scenario, When a RF group member is wrongly being acknowledged as a RF leader.
8. **RRM Assigning:** in progress
9. **Grouping disabled:** When RF grouping is switched “OFF” at the configured member
10. **Invalid Protocol Version:** If the RF member controller image is of an incompatible version or if there's a version mismatch.
11. **Country code mismatch:** Configured country mismatch
12. **Invalid hierarchy:** if lower priority controller is trying to add higher priority controller.
13. **Already a static leader:** If trying to add a member who's already been manually configured to be a static leader.
14. **Already Static Member:** When trying to add a member who's already been accepted a static member of another RF leader.
15. **Non-Static Member:**
16. **Not Intended:**
17. **Member Deletion Error:** If error is specifically known to occur due to improper memory allocation of de-allocation.
18. **RF-domain mismatch:** If the RF domain of the configured member and the RF leader is different.
19. **Split for invalid-state request:** An error state if there's a member split because of an RRM state transition that was not expected.
20. **Transitioning to static from auto:** While moving from auto to static state.

21. **Split due to user action:** When there's a user triggered transition because of reset while modifying country code, sys-name change or other
22. **Switch Size Exceeded:**



CHAPTER 5

Flexible Radio Assignment (FRA) and Redundant Radios

What does FRA do?

Cisco's Flexible Radio Assignment (FRA) is a new feature introduced in Release 8.2.110.0 to take advantage of the hardware released in the Cisco AP 3800/2800. FRA is a new core algorithm added to RRM to analyze the NDP measurements and manage the hardware used to determine the Role the new Flexible Radio (2.4 GHz, 5 GHz, or Monitor) will play in your network. There are many new and exciting functions integrated into RRM to cover, and we will cover them all here.

Before we start, the Dual-Band radio is a new product concept creating a flexible AP hardware architecture. Traditional legacy dual band APs have always had 2 radio slots, 1 per band and were organized by the band they were serving - slot0= 802.11b,g,n and slot1=802.11a,n,ac. The Dual-Band (XOR) radio offers the ability to serve either the 2.4 or 5 GHz bands or passively monitor both bands on the same AP. The AP models offered are designed to support dual 5 GHz band operations with the I model supporting a dedicated Macro/Micro architecture and the E and P models supporting Macro/Macro. This will be discussed in detail later - for now understand that the model offerings are designed specifically for different use cases which will be covered.

Cisco's FRA does 3 things:

- Calculates and maintains a measurement of redundancy for 2.4 GHz radios and represents this as a new measurement metric called COF (Coverage Overlap Factor)
- Manages radio “role” assignment or re-assignment for interfaces that are marked as redundant interfaces
- For Macro/Micro implementations (2800/3800 I models) FRA Manages Client load balancing across dual 5 GHz interfaces (Macro/Micro Transition)

Before jumping into the detail, a few name changes need to be covered. FRA is tightly integrated into existing RRM and will have to run in mixed environments with legacy AP's. There are some new naming conventions and behaviors to cover conceptually in order to have a clear understanding of the new operations. The AP MODE is an existing concept that required some modification. The Mode selection sets the entire AP (slot 0 and slot1) into one of several operating modes including:

- Local Mode
- Monitor Mode
- Flex Connect Mode
- Sniffer Mode

- Spectrum Connect Mode

Under the new FRA architecture - these "MODES" are still present and function as they always have. Changing an AP's mode - assumes that you will propagate that change to the entire AP, both radio slots 0/1. This changes with the 2800/3800 AP's. With the addition of the Flexible (XOR) radio in the slot0 position - we now have the ability to operate a single radio interface in many of the previous modes. There is no longer a need to place the whole AP into a MODE. When we apply this concept at the single radio level, we will refer to this as a ROLE. There are 3 Roles that can be assigned at the product release.

- Client Serving
 - Either 2.4 GHz⁽¹⁾ or 5 GHz⁽²⁾
- Monitor–Monitor Mode⁽³⁾



Note Client Serving is the same as Local MODE - in that the radio will support clients in the selected band. The change to Client Serving as a Role is to support post FCS when the Flex Connect MODE will be supported as a role and must be selected for the entire remote AP this will replace using local and remote interface terminology and consolidate under the role of client serving. What's important to remember is:

- A "MODE"–is assigned to a whole AP (slot0/1)
 - A "ROLE"–is assigned to a single radio interface (slot0)
-

- [FRA–Coverage Overlap Factor–COF , on page 34](#)
- [FRA–Radio Role Assignment, on page 38](#)
- [FRA–Client Network Preference, on page 40](#)
- [FRA–Steady State Operations, on page 42](#)
- [FRA–Troubleshooting , on page 45](#)


FRA–Coverage Overlap Factor–COF

Lets discuss the operational theory and a little more detail about how FRA does what it does. First and foremost, FRA evaluates 2.4 GHz coverage **only** and determines if there is overlapping coverage creating interference. The AP 3800/2800's initialize by default the same as a Legacy AP with 1x 2.4 GHz interface and 1 x 5 GHz interface. Once FRA completes it's analysis, any radio judged to be redundant - can then be assigned a different and more beneficial role to support the networks goals.


**Note**

- Lets discuss the operational theory and a little more detail about how FRA does what it does. First and foremost, FRA evaluates 2.4 GHz coverage **only** and determines if there is overlapping coverage creating interference. The AP 3800/2800's initialize by default the same as a Legacy AP with 1x 2.4 GHz interface and 1 x 5 GHz interface. Once FRA completes it's analysis, any radio judged to be redundant, it can then be assigned a different and more beneficial role to support the networks goals.
- Customers using AP groups to manage SSID assignments, may have AP's mixed in the same airspace supporting different AP Groups, FRA should be evaluated closely under these scenarios - with the AP's in Manual mode to evaluate the accuracy of COF in a discontinuous grouping. This is the correct way to evaluate complete coverage however - since if an SSID is added to only one AP, then there are no other AP's that can cover for that one AP with the additional or single SSID.

FRA uses the well-established Neighbor Discovery Protocol output from RRM to locate each radio in RF distance and evaluate overlapping coverage by cell. First, using the NDP measurements from the AP's, FRA plots the x and y coordinates relative to every other AP contained in the solution set (AP Group, physical neighbors). The circumference of each cell is calculated based on the present TX power level of each AP. This produces a logical matrix of the AP's coverage intersections.



Based on this understanding, FRA uses a multipoint analysis, to determine the percentage of overlapping coverage for each evaluated AP.




The output of this calculation is the COF (Coverage Overlap Factor %). Coverage Overlap Factor is the percentage of the analyzed cell that is covered at -67 dBm or higher by other - existing radios in service. In the process of calculating this coverage -FRA keeps track of radios that are Coverage Contributors to other radios COF - and will not allow those radios to be marked redundant as long as a radio they are a contributor for is Marked redundant. While FRA is aware of every radio that sends NDP above -80 dBm (system minimum value) it only uses the 4 loudest neighbors to perform the calculation against.



Note In 8.2MR1/8.3 only radios in 2.4 GHz are evaluated for COF. Once a radio is marked redundant, it is available to be assigned another role. Once a radio has been transitioned into another role - COF for that radio will eventually become not available - as it is no longer operating in the 2.4 GHz role where the evaluation is calculated.

RRM and FRA have no knowledge of a map (like Prime Infrastructure has) so the concept of walls, outside coverage, coverage area borders and the like are non-consequential to FRA. We do have a real idea about RF distances as NDP is sent at a known power level - and received packets are reported at actual RSSI so any obstacles show up as RF distance. Here is what's important to understand, COF is telling you the coverage for a theoretical 360 degree cell.

If the AP is in the middle of a large room - then the coverage expectations and what COF reports are likely in alignment. However - if your AP's are distributed down a long hallway with outside walls on the 8th floor, and you don't particularly care if you are covering the outside you can still use COF to evaluate the candidates, but adjust the expectation to the actual coverage area required. See below.



Using the example above, 4 AP's implemented in a staggered fashion in a hallway or similar architectural feature. None of these AP's would likely see 100% COF since:

- AP1 and AP3 are only covering roughly 60% of AP2's calculated cell coverage
- AP2 and AP4 are also only covering about 60% of AP 3's calculated cell

The lines in this drawing represent the outside walls that we know about, and within those lines each AP's required coverage is clear- to us. Each AP is really contributing about 30 % of it's calculated coverage to the actual required coverage area. It is very possible to use the COF as guidance combined with local knowledge to make a good decision here. Very Likely that AP3 should be converted to another role - as AP2 and AP4 cover almost 95% of the required coverage area. AP1, and AP4 extend beyond the center - but with more AP's in this particular hall, we would easily identify more redundant AP's. But this requires YOU to evaluate it as guidance.

Another example would be an enterprise location design where AP's are placed against outside walls. Again - you likely don't care about the coverage beyond those walls. If you do care about the coverage beyond the walls - leave it all alone or consider outdoor AP's. However if you don't and COF does not show 90-100% covered - no worries, If you know you only need 50% of the cell seeing a COF of 65% should send you out to take a quick couple of measurements to be certain before you assign a new role manually.

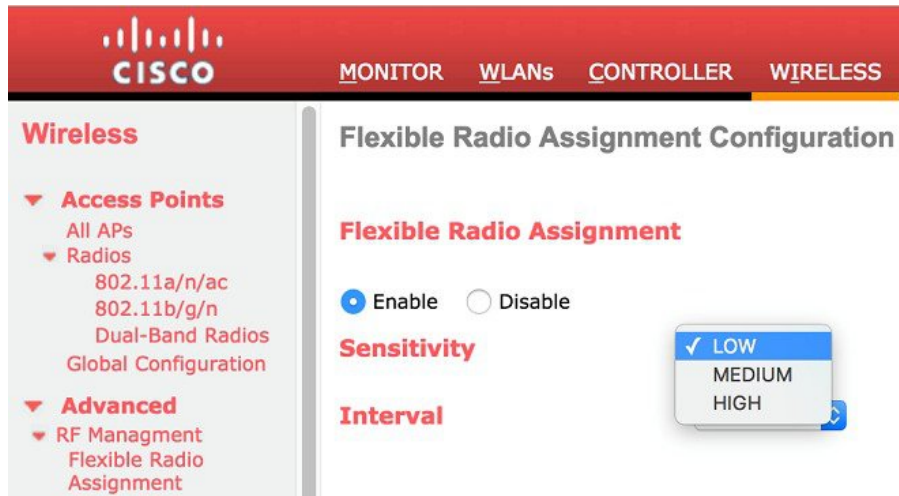
A Radio is marked as redundant when the calculated coverage meets or exceeds the FRA sensitivity threshold:

- Low–100% COF
- Medium–95% COF
- High–90% COF



Note Before enabling the algorithm, it is highly recommended that you place all your Flexible AP's in FRA Manual Mode. This will prevent any radios from being re-assigned while you are testing or becoming familiar with the commands.

You will need to enable the FRA algorithm. Unless FRA is already enabled - there will be no COF generated until it is - it is disabled by default. On the GUI - Wireless=>Advanced=>Flexible Radio Assignment, CLI > *config advanced fra enable/disable*



- Enable and disable the FRA algorithm – it is disabled by default.
- Adjust the FRA Sensitivity Threshold, this sets the percentage of COF required to consider a radio as redundant. Supported values are Low (100%-default) Medium (95%) and High (90%)
- Set the FRA run interval (Valid Choices 1-24 hours must meet or exceed DCA interval)

The FRA run interval is default 1 hour. Selections are available up to 24 hours. However, FRA depends on DCA heavily and the FRA interval cannot be shorter than the DCA interval. See below under role assignment.

FRA–Radio Role Assignment

Once a Radio is marked redundant what happens next is determined by the radio configuration. There are two operational states that the flexible radio can be assigned through - the radio is either FRA-auto (default value) or Manual. In Auto, the radio will be assigned directly by FRA/DCA based on the COF calculated with no user intervention. In Manual, FRA will still produce a COF for the radio - however the user will have complete control over which radios to assign manually.

Radio Role Assignment

Auto Manual

Client Serving Monitor

Band

RF Channel Assignment

Current Channel

Channel Width

Assignment Method Global Custom

Tx Power Level Assignment

Current Tx Power Level

Assignment Method Global Custom



Note

- FRA - Auto or Manual applies only to the assignability of the radio to FRA. RRM, channel and power are separate and users should not confuse the two. In either FRA auto or Manual, the user can still assign channel and power globally (RRM auto) or as a custom setting (no RRM - Manual/custom assignment).
- 2802/3802 models will enforce channel separation of 100 MHz, if attempted less in Dual 5 GHz mode an error message will be displayed.

With the radios in "FRA Auto", FRA then looks to DCA to decide "what" to do with the now redundant radio(s). DCA's priorities are simple:

1. Try to assign the redundant radio in 5 GHz and increase capacity
2. If DCA determines that there is already maximum 5 GHz coverage, it will say so and a Monitor role will be assigned instead

The results of this calculation are shown on the GUI under Monitor=>Access Points=>Dual Band Radios as "suggested mode".

AP Name	Radio Slot#	Base Radio MAC	Coverage Overlap Factor	Suggested Mode	Operational Status	Load Profile	Radio Role
AP00F2.8B26.89D0	0	00:f2:8b:26:ed:d0	84 %	2.4GHz	UP	NA	2.4 GHz * (Local)
3800_e_BC80	0	58:ac:78:df:7e:e0	46 %	2.4GHz	UP	NA	2.4 GHz * (Local)
3802_i_D9A0	0	58:ac:78:df:9c:20	100 %	5GHz	UP	NA	5 GHz * (Local)



Note Suggested mode will report (Static) indicating that the radios are in manual mode if FRA Auto is not selected.

```
Cisco Controller) >show advanced fra
```

```
FRA State..... Enabled
FRA Sensitivity..... high (90)
FAR Interval..... 1 Hour(s)
  Last Run..... 304 seconds ago
  Last Run Time..... 0 seconds
```

AP Name	MAC Address	Slot	Current Band	COF %
Suggested Mode				
-----	-----	-----	-----	-----
3802_i.D9A0 2.4GHz	58:ac:78:df:9c:20	0	2.4GHz	84
AP00F2.8B26.89D0 2.4GHz	00:f2:8b:26:ed:d0	0	2.4GHz	76
3800_e_BC80 2.4GHz	58:ac:78:df:7e:e0	0	2.4GHz	46

COF : Coverage Overlap Factor

Suggested Role's for the Flexible Radio include:

- 2.4 GHz
- 5 GHz
- 5 GHz/Monitor
- (Static)–manual assignment selected at the radio config level



Note The Cisco AP 2802/3802 e and p models do not have a second set of antenna's built into them. For these AP's you will need the DART connector and a second antenna. Neither model will be marked redundant without the insertion of the DART connector it is not available to the Role Selection algorithms.

FRA–Client Network Preference

That takes care of "suggesting" a role. However when that transition is made is governed by another NEW feature in 8.2 MR1 and 8.3 software releases - Client Network Preference. Think of this like a personality adjustment for your network services. It allows an end user to characterize how it wants to prioritize RRM's behaviors. FRA has an entry here.

RRM Algorithm	Connectivity	Throughput	Default
DCA	Higher Bias on current Channel Low Sensitivity for DCA (override in DCA) Prefer current channel when Voice/Video	Less bias on current channel High DCA Sensitivity (override in DCA) All 11n capable clients	No Change from Current DCA behavior
DTPC	DTPC should not reduce power when weak signal clients detected (<-79)	DTPC will not increase power when number of clients is over the configured max threshold	No Change from current TPC behavior
DBS	Prefer Current bandwidth Prefer current Primary channel Lower bias on client type Voice/Video clients increased contribution	Favor Higher Channel Widths Higher bias for high capacity (11n,11ac) client types present	No Change from current DBS behavior
FlexDFS/ED–RRM	Priority on maintaining primary channel Voice/Video client increased contribution	Priority on higher bandwidth channel width	No Change from current flexDFS/ED-RRM behavior
FRA	No Band/Role switching when more than 3 clients exist	Band-switching, regardless of clients	Defaults to connectivity profile

This command is only at the CLI level for now, however it allows customization of processes that different users want to behave differently. Within the context of FRA, what happens to clients connected to the 2.4 GHz interface if it changes to 5 GHz - and do you care? There are 3 options including the default which essentially leaves things as they are today for the customer - except in the case of FRA which by default uses the connectivity setting.

```
(Cisco Controller) >config/show advanced 802.11a/b client-network-preference ?
```

- Connectivity–Preference on Connectivity
- Default–No Client Network Preference Applied
- Throughput–Preference on Throughput

In this case then, the radio would wait to switch roles until the condition (less than 3 clients associated) is met and process the switch on the next DCA run. You can bypass this by selecting the "Throughput" option for Client Network Preference.

A radio will not return to 2.4 GHz unless reverted (through CLI) or manually configured from the GUI/CLI. Automatic reversal only happens in the case of coverage hole alerts on any of the original contributing coverage radios. That means - once a radio is placed in another role - if any of the radios that are contributing to the

coverage to allow an alternate role experience coverage hole alerts, the flexible radio is reverted back to 2.4 GHz immediately.



Note 3802i AP's have two sets of antenna's, a dual band set, and a micro cell antenna for 5 GHz. The 3802i is specifically engineered to support dual 5 GHz as a Macro/Micro cell implementation or operate as a monitor interface. More on this later.

3802e and p models have a single set of external antenna adapters. So there are no built in extra antenna's to support the flexible radio roles. You cannot use the same antenna for two radios operating in the same band simultaneously. For dual 5 GHz operations the Dart connector is needed to support a second antenna that is dedicated to the Flexible (slot0) radio.



Note You do not need a Dart connector to treat the AP as a standard dual band AP - This is only needed support Dual 5GHz and Monitor Roles for the Flexible radio.

These AP's are designed to support Macro/Macro or two full serving cells. Without the Dart connector- no role re-assignment is possible and while it will show COF, it cannot be assigned a different role without the supporting hardware.

FRA—Steady State Operations

FRA is designed in the initial release to evaluate and remediate over population in 2.4 GHz spectrum. In almost all cases - if you are designing a network to optimize 5 GHz, then you will have redundant 2.4 GHz interfaces. It manages the selection, transition, and assignment of redundant radios.

Once it has run, and assuming the network design and number of AP's has not changed, there won't be much for FRA to do but monitor. Keep in mind that while you will not be seeing COF for any AP that's not in the 2.4GHz role, the part about role selection and DCA stays active. What that means - is that if you add AP's or you change the bandwidth you are serving, FRA may choose a different role for your Flexible Radio if 5 GHz is assigned and it becomes challenging to keep that many interfaces in the air peacefully. You may see a 5 GHz interface switch to monitor mode. One sure way to prevent this is lock down the radios by changing them to Manual. That will prevent FRA from taking any more action on their behalf. Keep in mind that if you do this, you could also create issues when making changes that alter the balance of the spectrum - such as Bandwidth Changes which should always be treated as a major upgrade (because - it is).

COF will be "not available" after some time for any Flexible interface not in the 2.4 GHz role in this revision of code. If at any time you wish to reset FRA and re-run the calculations from scratch simply revert the radios:

```
(Cisco Controller) >config advanced fra revert Auto-only/All Static/Auto
```

You can select radios in Auto, or all and revert them to Auto or Static. Reverting to Auto, automatically assigns 2.4 GHz client serving role.

FRA and Dual 5 GHz Operations

Managing a dual 5 GHz cells is one of the most most important things that FRA does. There are two modes of operation for Dual 5 GHz AP's:

- Macro/Micro—A large cell with a smaller cell internal – creating double the capacity within the boundaries of a single cell.
- Macro/Macro—Dual- independent 5 GHz cells, doubling the coverage of a single traditional Dual Band AP.

Macro/Micro is a use case applicable to the Cisco AP 3800/2800 I Models only. The antennas on this AP are designed to support a Cell within a Cell deployment. In order for this to be effective, much design is committed to isolation of the two cells starting at the silicon level, and ending with separation in antenna polarity and frequency separation.

FRA and DCA will enforce many configuration requirements when operating as Dual 5 GHz Macro/Micro:

- Minimum 100 MHz Channel Separation (Frequency Diversity)
- Micro Cell Power is restricted to minimum
- Same SSID's on each cell

Having a Macro/Micro cell architecture is attractive because it solves the problem of having one large cell - in which a very diverse client experience exists. Clients closer to the AP are using higher data rates and enjoy a higher SNR than will clients at the edge of the cell. Macro Micro allows isolation of different clients within the cell that can serve them best and preserves airtime by increasing overall efficiency.

FRA also monitors the two cells and will optimize client connections between the two cells - keeping like performing clients with like performing clients maximize throughput. There are presently three methods by which we move or transition clients between Macro and Micro cells.

- 802.11v BSS Transition
- 802.11k
- Probe Supression

Transition thresholds are configurable by the user - but the defaults are highly recommended and worked well in our testing and Beta.

```
(Cisco Controller) >show advanced client-steering
Client Steering Configuration Information
1 Macro to micro transition threshold..... -55 dBm
2 micro to Macro transition threshold..... -65 dBm
3 micro-Macro transition minimum client count.... 3
4 micro-Macro transition client balancing win.... 3
<snip>
```

Adjustments can be made to these values at the CLI using the commands:

```
(Cisco Controller) >config advanced client-steering transition-threshold ?

balancing-window Configures micro-Macro client load balancing window
macro-to-micro Configures Macro to micro transition RSSI
micro-to-macro Configures micro to Macro transition RSSI
min-client-count Configures micro-Macro minimum client count for transition
802.11v BSS Transition - enabled by default
```

802.11v BSS Transition—Default Enabled

An 802.11v capable client is known to the AP on association. If such a client connects initially to the wrong cell (Macro is most likely, but Micro is possible) on connection the AP will send an 802.11v BSS transition

message with only the BSSID of the target cell as a neighbor. This will be followed by a mandatory De-auth (40 seconds). In testing and beta this method worked for all 802.11v capable clients very well.

802.11k–Default Enabled

Similar to 802.11v, using a site report message with only the target cell's BSSID listed will leave the client little choice. The issue with this method is that ultimately the client needs to request the site report, and in most implementations today - that generally

Probe Suppression–Default Disabled

Non-associated clients usually send probe requests on all channels before using the information in the responses to initiate an association sequence. When both Micro and Macro cell radios receive probe request from non-associated client within a short duration and if the client is within the coverage area of the Micro or Macro cell, the AP can suppress probe responses to the client from the the radio we do not want them to join. Since this effectively "hides" the cell from the client, the client is likely to initiate association sequence with the cell that responded. If a client is "sticky" and tries to associate to the other cell even after not receiving a probe response on that channel, depending on the operational mode, the AP can also suppress a single AUTH response on that cell to force the client to try associating to the desired cell.

Similarly associated clients usually send probe requests on all channels before using the information in the responses to move to a radio with higher transmission power. When both radios on the AP receive probe request from a client associated to either cell, in order to prevent the client from prematurely roaming to the other cell due to its higher transmission power, the AP continues to suppress the other cell's probe responses to the client as long as the client is within the coverage area of the target cell.

Like in the other two methods of client steering, we rely on the same client RSSI thresholds to target a coverage cell. Unlike the previous two, this method is disabled by default.

Other than the Micro to Macro and Macro to Micro client thresholds, there is not much that needs or should be configured for 802.11v or k methods. Probe Suppression has some options we will cover here.

- **Operational Mode**–This parameter governs the mode of operation for micro-Macro probe suppression feature on the AP. It supports the following options:
 - **Disable**–Disables micro-macro probe suppression feature on the AP
 - **Probe Only**–Suppresses only probes on Macro cell subject to constraints
 - **Probe And Auth**–Suppresses probes on Macro cell along with a single Auth subject to constraints
- **Probe Validity Window**–This parameter specifies the time interval for which the probe on the micro cell is considered valid in reference to the probe received on the Macro cell
- **Suppress RSSI Threshold**–The client RSSI threshold within which probe responses from the Macro cell are suppressed
- **Enable RSSI Threshold**–The client RSSI threshold beyond which probe responses from the Macro cell are no longer suppressed
- **Probe Aggregate Window**–Clients can sometimes send multiple probe requests in a short span of time. This parameter specifies time interval during which probes from a client are not considered distinct
- **Transition Aggressiveness**–This parameter specifies number of dis-contiguous probe aggregate windows for which probes on the Macro cell are suppressed

```
(Cisco Controller) >config advanced client-steering probe-suppression ?
```

```
disable      Disables micro-Macro probe suppression
enable       Enables micro-Macro probe suppression
hysteresis   Configures hysteresis
transition-aggressiveness Configures probe cycles to be suppressed
```

All of the Macro/Micro transition thresholds can be seen using the *show advanced client-steering* command at the cli:

```
(Cisco Controller) >show advanced client-steering
```

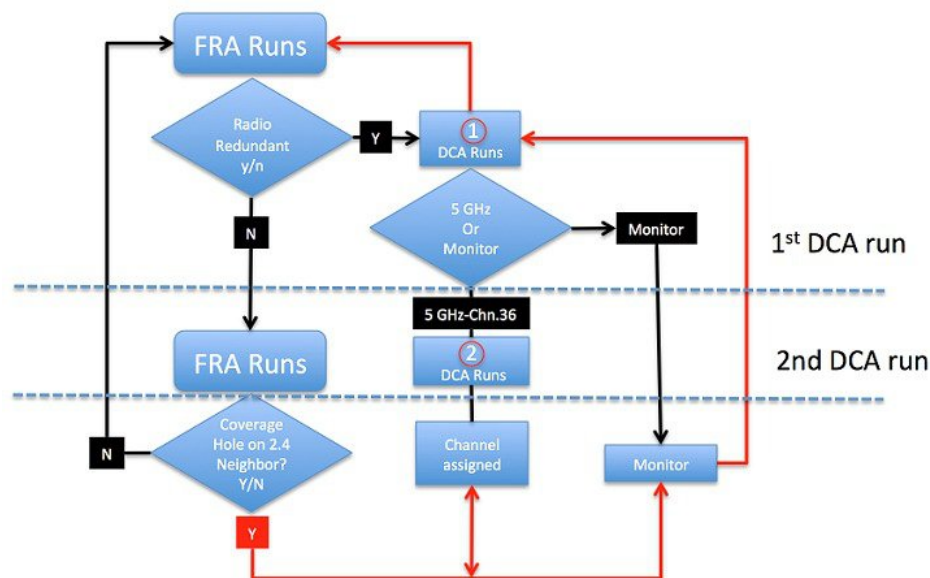
```
Client Steering Configuration Information
```

```
Macro to micro transition threshold..... -55 dBm
micro to Macro transition threshold..... -65 dBm
micro-Macro transition minimum client count... 3
micro-Macro transition client balancing win... 3
Probe suppression mode..... probe-and-auth
Probe suppression validity window..... 100 s
Probe suppression aggregate window..... 200 ms
Probe suppression transition aggressiveness... 3
Probe suppression hysteresis..... -6 dBm
```

FRA–Troubleshooting

FRA is very dependent on DCA, as DCA is used to both evaluate the role selection, and assign channels and band to the radio.

This means that DCA must run at least as often as FRA, and while FRA will make recommendations, it is up to DCA to implement and manage them. This can cause a lot of confusion, unless you know your DCA and FRA configurations, and how they are expected to behave. See the example below, we will walk through the scenarios.



1. FRA Runs and establishes that a Radio is Redundant
2. The Radio now Marked Redundant awaits Role assignment from DCA

3. When DCA runs, we assign it to either 5GHz or monitor based on DCA's assessment of the channel plan/widths/interference and so on
 - a. Because the radio was reverted (running in 2.4 GHz mode, with no 5 GHz channel assigned previously, default channel for 5 GHz is 36, radio is assigned 5 GHz chan 36)
4. Next DCA run, 5 GHz radio has neighbors, and a channel is assigned
5. Any time a neighbor of any of these radios shows a coverage hole at 2.4 GHz being reported FRA will revert any radios that had roles switched back to 2.4 GHz and begin again

If DCA is set for an interval of 12 Hours, FRA must also be set for 12 hours (minimum), the second DCA cycle could be 24 hours

FRA-CLI and Debug Commands

Show Advanced fra shows you the current FRA state (enable/disable), sensitivity, current run interval, last run time, and status and state of all Flexible radios

```
(Cisco Controller) >show advanced fra
```

```
FRA State..... Enabled
FRA Sensitivity..... high (90)
FAR Interval..... 1 Hour(s)
  Last Run..... 620 seconds ago
  Last Run Time..... 0 seconds
```

AP Name Mode	MAC Address	Slot	Current Band	COF %	Suggested
---	---	---	---	---	---
3800_e_BC80	58:ac:78:df:7e:e0	0	2.4GHz	46	2.4GHz
2802_i.89D0	00:f2:8b:26:ed:d0	0	5GHz	84	5G/Monitor
3802_i.D9A0	58:ac:78:df:9c:20	0	5GHz	100	5G/Monitor

COF : Coverage Overlap Factor

Show Advanced 802.11a/b summary will show you all radios in a given band, Flexible radios show up on slot 0 when assigned an 802.11a role. You can also show just the flexible radios with show advanced 802.11-abgn summary

```
(Cisco Controller) >show advanced 802.11-abgn sum
```

```
Member RRM Information
AP Name                MAC Address           Slot Admin   Oper   Channel  TxPower
-----
3800_e_BC80            58:ac:78:df:7e:e0    0  ENABLED  UP     1*      *8/8 (
2 dBm)
2802_i.89D0            00:f2:8b:26:ed:d0    0  ENABLED  UP     (56,52)* *7/7 (
2 dBm)
3802_i.D9A0            58:ac:78:df:9c:20    0  ENABLED  UP     (36,40)* *7/7 (
2 dBm)
```

Config advanced fra revert Can be used to quick flip all or just the flexible radios in auto back to default state of 2.4 GHz client serving and mark them Auto FRA or Manual (very useful for recalculating COF's after adding or moving AP's)

```
(Cisco Controller) >config advanced fra revert all auto
```

```
Reverting XOR radios can cause disruption for associated clients. Are you sure you want to
continue? (y/n)y
```

```
(Cisco Controller) >show advanced fra
```

```
<snip>
```

AP Name	MAC Address	Slot	Current Band	COF %	Suggested Mode
3800_e_BC80	58:ac:78:df:7e:e0	0	2.4GHz	46	2.4GHz
2802_i.89D0	00:f2:8b:26:ed:d0	0	2.4GHz	84	5G/Monitor
3802_i.D9A0	58:ac:78:df:9c:20	0	2.4GHz	100	5G/Monitor

Debug `airewave-director fra enable` will invoke the FRA debug, we'll take this in sections focusing on what's important. There is a lot in this command that's only important to developers so for brevity we'll skip some of the redundant lines.



Warning There is a lot of output in a debug command- debug FRA should never be run on a fully loaded production system against all AP's. You can use the `debug mac addr <mac1,mac2,mac3>` to isolate the debug to just the devices you are concerned with.

The first part of the debug lists all the radios that will be included in the solution set as well as if it is an RRF (Redundant Radio Flexible or a 2800/3800) candidate. Note that the AP's displayed - should be the AP's with 2.4 GHz interfaces. The last column - is extremely useful - NDM (Network Density Metric) this is the number of AP's that each of the indexed AP's have as neighbors they can hear above -67 dBm. The higher that number the denser (AP's closer together) the network is.

```
(Cisco Controller) >debug airewave-director fra enable
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.014: Executing FRA on 802.11bg
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.014: -----
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.014:
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.014: RDCI List :
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.014: -----
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.014: Index          AP                      RRF Candidate
      NDM
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.014: -----
-----
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.014: 0
6                               58:AC:78:DF:9C:20(0)    YES
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.014: 1
6                               00:F2:8B:26:ED:D0(0)    YES
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.014: 2
6                               F0:7F:06:5E:FD:C0(0)    NO
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.014: 3
6                               E4:C7:22:AA:38:20(0)    NO
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.015: 4
6                               E4:C7:22:AA:34:40(0)    NO
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.015: 5
6                               08:CC:68:B4:20:00(0)    NO
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.015: 6
5                               58:AC:78:DF:7E:E0(0)    NO
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.015: 7
4                               64:D9:89:46:7F:B0(0)    NO
```

The next step starts the calculation for each RRF Candidate AP in the above list, in this example we will have 2 AP's that are marled Yes. I have 1 3800i, one 2800i, and a 3800 e with NO Dart connector so not currently flexible, the remaining AP's on the list are 3500/3600/3700.

For AP Index 0 (58:AC:78:DF:9C:20) the AP mac address is displayed- and immediately we get the neighborhood list. This is the list of Neighbors to the target AP that will be used in AP index 0's calculation. This is where you see the Neighbors that matter. Only the loudest 4 neighbors will be used- all must be in the same AP group. If an AP is not on this list that you think honestly should be - due to proximity- it is not known as one of the 4 loudest neighbors to the target AP. This could be because it is in a different AP group or that Neighbor information is not yet available (newly powered on).

In the output, you see the first AP on the list selected index 0, and its loudest 4 neighbors displayed along with their neighbor RSSI and their TX power (because of a cosmetic bug, current TX power is incorrect and always displays as max power in release 8.2 MR1, but the algorithm is using current TX power). This establishes the AP's that can be considered for coverage of AP 58:AC:78:DF:9C:20 coverage area.

```
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.015: 58:ac:78:df:9c:20
-----
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.015: 58:ac:78:df:9c:20 Redundancy Iteration 0 , RDCI
0
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.015: 58:ac:78:df:9c:20
-----
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.029: 58:ac:78:df:9c:20 Neighborhood list for AP 0
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.029: 58:ac:78:df:9c:20
-----
Jun 07 16:22:05.029: 58:ac:78:df:9c:20 RDCI Index AP RSSI Tx
Power (dBm)
Jun 07 16:22:05.029: 58:ac:78:df:9c:20 -----
-----
Jun 07 16:22:05.029: 58:ac:78:df:9c:20 0 58:AC:78:DF:9C:20(0) 0
23
Jun 07 16:22:05.029: 58:ac:78:df:9c:20 2 F0:7F:06:5E:FD:C0(0) -14
22
Jun 07 16:22:05.029: 58:ac:78:df:9c:20 7 64:D9:89:46:7F:B0(0) -30
22
Jun 07 16:22:05.029: 58:ac:78:df:9c:20 5 08:CC:68:B4:20:00(0) -31
23
Jun 07 16:22:05.029: 58:ac:78:df:9c:20 4 E4:C7:22:AA:34:40(0) -43
20
Jun 07 16:22:05.029: 58:ac:78:df:9c:20 3 E4:C7:22:AA:38:20(0) -48
20
Jun 07 16:22:05.029: 58:ac:78:df:9c:20 6 58:AC:78:DF:7E:E0(0) -66
23
Jun 07 16:22:05.029: 58:ac:78:df:9c:20
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.029: 58:ac:78:df:9c:20 Neighborhood Path Loss Matrix
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.030: 58:ac:78:df:9c:20 -----
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.030: 58:ac:78:df:9c:20 0 37 52 53 54
66 71 89
<snip> Calculating every AP to every AP's path loss
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.030: 58:ac:78:df:9c:20 Neighborhood Distance Matrix
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.030: 58:ac:78:df:9c:20 -----
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.030: 58:ac:78:df:9c:20 0 0 4 5 5
14 17
<snip>Calculating RF distances between AP's
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.030: 58:ac:78:df:9c:20 Unable to triangulate
neighborhood, assigning single axis coordinates <normal behavior- perfectly fine we are
operating in 2D>
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.030: 58:ac:78:df:9c:20 Computed Coordinates
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.030: 58:ac:78:df:9c:20 -----
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.030: 58:ac:78:df:9c:20 X Y R
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.031: 58:ac:78:df:9c:20 ----
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.031: 58:ac:78:df:9c:20 0 0 18
```

```

<snip> many lines of X, Y positions - and calculated radius of each cell
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.031: 58:ac:78:df:9c:20 Neighborhood Multi-Point Check
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.031: 58:ac:78:df:9c:20 Using 4 neighbors for overlap
computation
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.031: 58:ac:78:df:9c:20 Point Coordinates Coverage
Score
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.031: 58:ac:78:df:9c:20 -----
-----
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.031: 58:ac:78:df:9c:20 0 , 0 2
<snip> Many more lines - deleted for brevity. This is the output and score from the
multipoint overlap calculation.
<comment> followed by the RDCI Overlap Factor for this radio and it's conclusion
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.031: 58:ac:78:df:9c:20 RDCI Overlap Factor : 92 %
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.031: 58:ac:78:df:9c:20 RDCI is redundant
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.032: 00:f2:8b:26:ed:d0 Redundancy Iteration 1 , RDCI
1(00:F2:8B:26:ED:D0(0))
<snip> Redundancy iterations continue for each radio marked as Yes for RRF candidate
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.034: 58:ac:78:df:7e:e0 Redundancy Iteration 2 , RDCI
6(58:AC:78:DF:7E:E0(0))
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.037: Redundancy Identification Complete
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.037: -----
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.037:
<snip> The last entry below - is the output you see in show advanced fra
*RRM-MGR-2_4-GRP: Jun 07 16:22:05.037: RDCI List :
*RRM-MGR-2_4-GRP: Index AP Overlap % Redundant overlap
Contributor
*RRM-MGR-2_4-GRP: -----
-----
*RRM-MGR-2_4-GRP:0 58:AC:78:DF:9C:20(0) 92 YES NO
*RRM-MGR-2_4-GRP:1 00:F2:8B:26:ED:D0(0) 69 NO YES
*RRM-MGR-2_4-GRP:6 58:AC:78:DF:7E:E0(0) 53 N/A NO

```

The debug concludes with the conclusions. Note that AP index 0 is marked as redundant, and is NOT a overlap contributor. The current threshold for FRA is High (90%) and 92% is above 90 %. Had sensitivity been set to medium (85%) it would not have been marked redundant.

Also notice that AP Index 1, is marked as a coverage contributor. This means regardless of the percentage of COF, this AP's 2.4 GHz interface is being counted on to provide coverage for another Radio that has been Marked as Redundant already (AP index 0).

Remember, FRA only identifies a redundant interface and once identified DCA must run to assign a new Role. Things that will prevent a radio from being assigned a new role:

- Client-network-preference–by default for FRA only we use the connectivity profile. This enforces that 3 or more clients associated to the radio – it will not transition. This is evaluated at DCA run time.
- DCA not running–in which case you will not get a suggested role in the show output since it can not be evaluated.

FRA–Best Practices

FRA is new. It has the power to automatically assign 2.4 GHz radios to other roles completely automatic. It is highly recommended that you familiarize yourself with the operations. Operating in Manual Band Selection at the AP is perfectly safe, and will prevent FRA algorithm from acting. This allows you to view the COF, and validate and become familiar with the findings. FRA is conservative by nature.

To Run in Manual Mode:

1. config advanced fra revert all static

2. config advanced fra enable

You can view the COF results and validate without changing bands. Then select the radios you want to change manually later or just turn on auto and let the radios shift.



CHAPTER 6

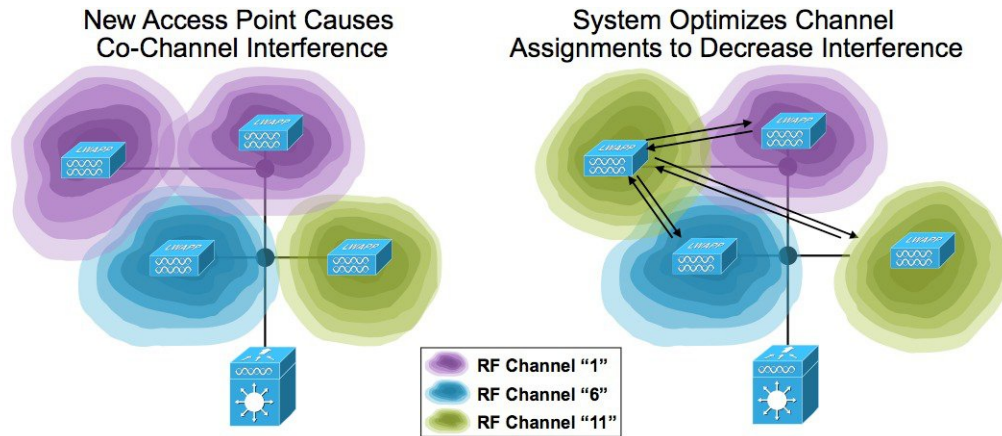
Dynamic Channel Assignment (DCA)

- [What does Dynamic Channel Assignment do?, on page 51](#)
- [The Dynamic Channel Assignment \(DCA\) Algorithm, on page 52](#)
- [DCA in a Nutshell, on page 53](#)
- [DCA Modes of Operation, on page 55](#)
- [DCA 20, 40, 80, and 160 MHz support , on page 57](#)
- [Dynamic Bandwidth Selection–DBS , on page 60](#)
- [Device Aware RRM , on page 62](#)

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 11: When a new AP is added, it's 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

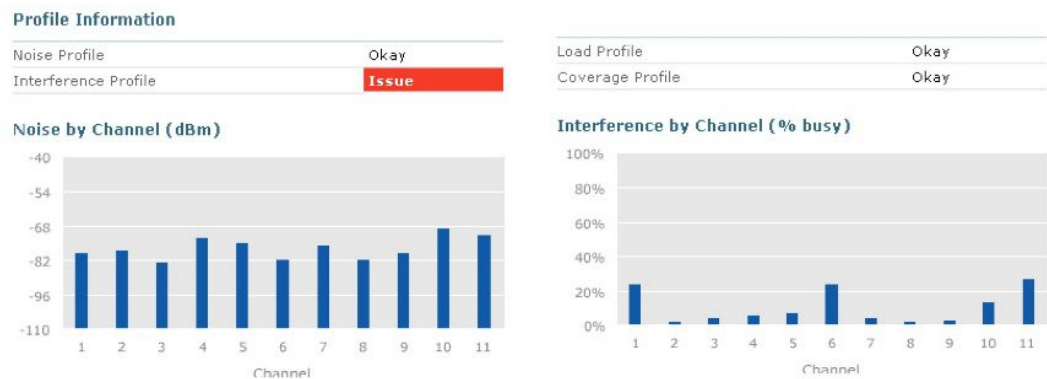
The Group Leader maintains the neighbor lists for all AP's in the RF Group, and organizes these neighbors into RF Neighborhoods. The following metrics are also tracked for each AP in the RF Group.

- **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 12: 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 13: 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 6: 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.

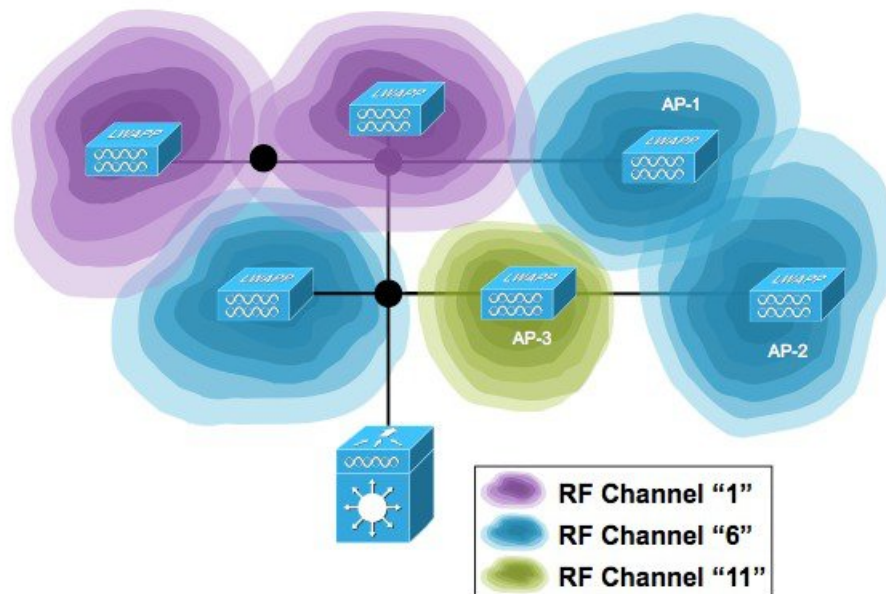
Steady State Mode

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 14: 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 15: 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 7: 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).

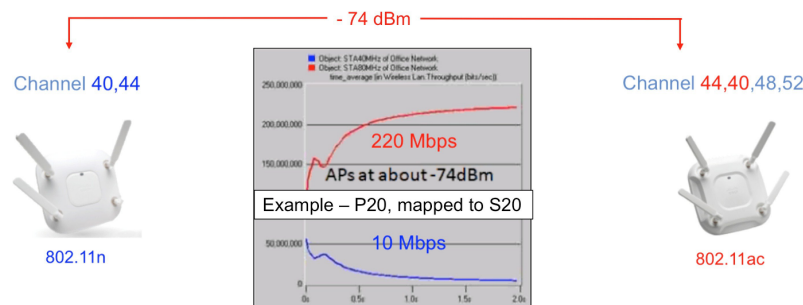
Table 8: CCA Threshold Examples

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 16: 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

Figure 17: Debug Channel Output for DBS

```

>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>> groupDCACheck
08:cc:68:b4:20:60 : Computing channel assignment for AP 08:CC:68:B4:20:60(1)
: LRAD is profile member. Fetching profile data...

08:cc:68:b4:20:60 RRM Chan Assignment Mode: 2 Lrad Capability: 2 DCA Channel Width: 2

08:cc:68:b4:20:60 : Not using chan = 40 on AP 08:CC:68:B4:20:60(1) because of secondary20
08:cc:68:b4:20:60 : Not using chan = 52 on AP 08:CC:68:B4:20:60(1) because of secondary20
08:cc:68:b4:20:60 : Not using chan = 136 on AP 08:CC:68:B4:20:60(1) because of secondary20
08:cc:68:b4:20:60 : Not using chan = 157 on AP 08:CC:68:B4:20:60(1) because of secondary20
08:cc:68:b4:20:60 DBS bs 0 #ac/n/a/vo/vi 2/0/0/0/0 p 80/40/20/vo/vi 0.00/6.00/6.00/0.00/
08:cc:68:b4:20:60 80Mhz Recommended Channel Set: 112 108 100 104 with Best Metric:-71.04
08:cc:68:b4:20:60 Alternative 80Mhz channel set: 0 0 0 0 with Best Metric: -71.04
08:cc:68:b4:20:00

```

In the graphic above (output from debug airwave-director channel enable) note the DBS bs line ac/n/a/vo/vi = 802.11ac/802.11n/802.11a/voice/video 2 /0/0/0 /0 = 2 associated 802.11ac clients, no 802.11n, no 802.11a, no voice and no video.

Following this count we have the bias score –the bias is added to the cost metric for a particular bandwidth, more bias = less likely to choose.

$P 80/40/20/vo/vi = 80 \text{ MHz}/40 \text{ MHz}/20 \text{ MHz}/\text{voice}/\text{video } 0/6/6/0/0 = \text{no bias}$ – against 80 MHz, bias against 40 and 20 MHz, no Bias for voice or video – this is RRM for – recommending an 80 MHz channel – because the only clients are 802.11ac capable. Does this mean I will get an 80 MHz channel–NO. However the likelihood is increased and we will have to weigh it against the other factors within the environment.

Looking at the whole network, a small one to be sure – the same debug and it’s recommendations for each radio look like extracted output as shown.

Figure 18: DBS Conclusions from Channel Debug–Excerpt

- *RRM-MGR-5_0-GRP: Oct 06 17:53:12.424: 64:d9:89:46:7f:b0 DBS bs 0 #ac/n/a/vo/vi 0/0/0/0/0 p 80/40/20/vo/vi 0.00/0.00/6.00/0.00/0.00 0/0/1536/0/0
- *RRM-MGR-5_0-GRP: Oct 06 17:53:12.424: 64:d9:89:46:7f:b0 Mix Mode Recommended Channel Set: 132 0 0 0 with Best Metric:-80.91
- *RRM-MGR-5_0-GRP: Oct 06 17:53:12.426: 64:d9:89:43:4d:50 DBS bs 1 #ac/n/a/vo/vi 0/1/0/0/0 p 80/40/20/vo/vi 0.00/0.00/6.00/0.00/0.00 0/0/1536/0/0
- *RRM-MGR-5_0-GRP: Oct 06 17:53:12.426: 64:d9:89:43:4d:50 Mix Mode Recommended Channel Set: 56 0 0 0 with Best Metric:-80.85
- *RRM-MGR-5_0-GRP: Oct 06 17:53:12.430: f4:0f:1b:b2:8d:80 DBS bs 0 #ac/n/a/vo/vi 0/0/0/0/0 p 80/40/20/vo/vi 0.00/6.00/6.00/0.00/0.00 0/1536/1536/0/0
- *RRM-MGR-5_0-GRP: Oct 06 17:53:12.430: f4:0f:1b:b2:8d:80 80Mhz Recommended Channel Set: 149 153 157 161 with Best Metric:-71.02
- *RRM-MGR-5_0-GRP: Oct 06 17:53:12.434: 08:cc:68:b4:20:60 DBS bs 0 #ac/n/a/vo/vi 2/0/0/0/0 p 80/40/20/vo/vi 0.00/6.00/6.00/0.00/0.00 0/1536/1536/0/0
- *RRM-MGR-5_0-GRP: Oct 06 17:53:12.434: 08:cc:68:b4:20:60 80Mhz Recommended Channel Set: 112 108 100 104 with Best Metric:-71.04
- *RRM-MGR-5_0-GRP: Oct 06 17:53:12.437: 08:cc:68:b4:20:00 DBS bs 0 #ac/n/a/vo/vi 1/1/0/0/0 p 80/40/20/vo/vi 0.00/6.00/6.00/0.00/0.00 0/1536/1536/0/0
- *RRM-MGR-5_0-GRP: Oct 06 17:53:12.437: 08:cc:68:b4:20:00 80Mhz Recommended Channel Set: 36 40 44 48 with Best Metric:-70.30

Other AP's in the configuration example above either have NO 802.11ac clients - or are split between a single 802.11ac and an 802.11n client. Bandwidth is set accordingly for the channels and AP's that are in use.


Arguably - this is a simple configuration and things get more complex at scale - however the logic which is being used is good logic. It matches best practice recommendations that are based on - how many of what type of client are you supporting? If you set 80 MHz channels for everything, when most of your clients are still 802.11n then you are wasting a lot of bandwidth that 802.11n clients can not use. In fact - it is optional for 802.11n clients to support a bonded channel and most smartphones do not, this is something more commonly supported on laptops and upper end tablets only.

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 19: 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 20: 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 21: 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 22: 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.



CHAPTER 7

Transmit Power Control (TPC) Algorithm

Choosing an operating channel for an AP with the best SNR (Signal To Noise ratio) is important. But since one of the major sources of interference in our network is our own clients and AP's - Transmit Power Control is equally as important. DCA and TPC work hand in hand to manage the RF in our environment. Power largely determines our cell boundaries, there are other variables (see [High Client Density Design Guide](#)), but power is one of the primary determining factors. The goal is to maximize the RF coverage in the environment - without causing co-channel interference. It's a balancing act of sorts. Since we cannot control the clients TX power (not all clients will support DTPC, an optional portion of the CCX specification) we only have our AP's to work with. Maximizing the AP's coverage and minimizing it's interference potential then is the job of TPC.

- [What does TPC do?](#) , on page 67
- [TPCv1](#) , on page 68
- [TPCv2](#) , on page 70
- [TPC Min/Max](#) , on page 74

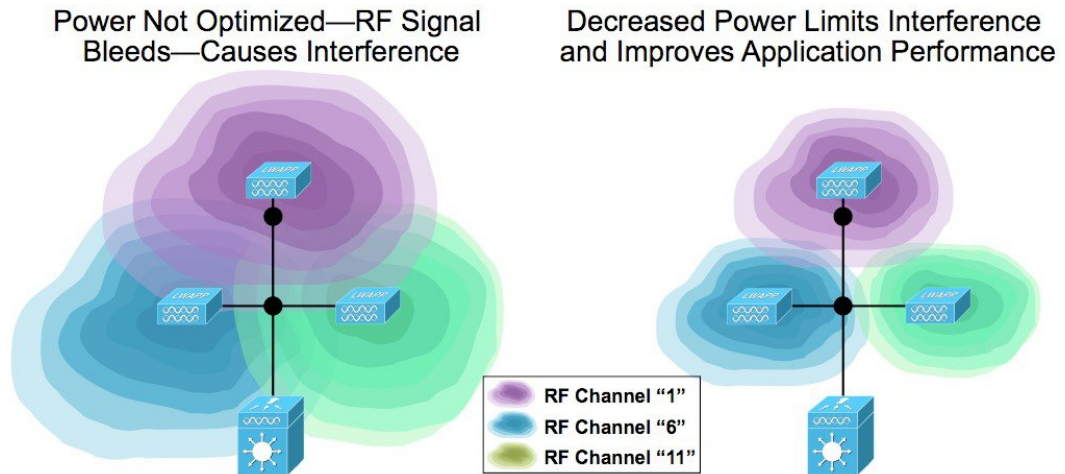
What does TPC do?

TPC uses the TX neighbor and RF Neighbor lists generated by the NDP process. RSSI organized lists built on how we hear other AP's (RX Neighbor) and how other AP's hear us (TX Neighbor), to form a picture of how every AP is heard by every other AP within the RF Neighborhood and RF Group. Based on this information TPC sets the transmit power of each AP to maximize the coverage and minimize co-channel interference. TPC will adjust the Tx power up or down to meet the required coverage level indicated by the TPC Threshold.

Like DCA, TPC runs on the RF Group leader and is a global algorithm that can be sub configured in RF profiles for groups of AP's in an AP group.

There are two versions of this algorithm since version 7.2 known as TPCv1 (or just TPC), and TPCv2. The purpose of these two algorithms is essentially the same - the calculations and how they are implemented differ greatly. We will discuss each below and give their strength's and potential caveats.

Figure 23: TPC Maximizes Coverage, Minimizes CCI



TPC algorithms run on the RF Group Leader along with DCA. They are configured and run separately from DCA. Three modes may be configured:

- Automatic—Runs every 10 minutes
- On Demand—invoke a power level change once, then freeze till the next demand request
- Fixed—allows a user selected power level to be applied to all AP's in the RF group

TPCv1

TPCv1 is known as coverage optimal mode - and is the default method for power control in RRM. The algorithm runs on the RF Group Leader, and calculates Tx power on a per AP per radio basis for every member of the RF group. Control over TPCv1 relies on a user tunable setting - RRM Power Threshold. This combined with the information gathered from the third neighbor in an AP's neighbor list, is used to make decisions on an AP's transmit power.

The RF Power threshold is used to control the cell boundaries of the AP's and hence the coverage behavior of the system. The default value of -65 dBm was in place until rel. 4.1.85.0 (4.1 MR1), which changed the default value to -70 dBm. Valid entries for this are -50 to -80 dBm. The RF Power Threshold is set on the controller and should be the same for every controller in an RF Group. Good results are generally observed with values ranging between -68 dBm and -75 dBm. However certain scenarios will require higher or lower settings. The RF Power threshold determines what an AP does in response to how other AP's hear him. In a situation like a high ceiling - the AP's might hear each other fine - but down on the floor the clients are having issues.

The Third loudest TX neighbor is used because this is the number of non-overlapping channels available for 2.4 GHz. As an aside, tests have been done with the 8th neighbor and more -while this would seem to make more sense in 5 Ghz, there is no real advantage realized. Also, very few regulatory domains agree on the number of channels they permit - so 3 is the number.

Calculating Tx_Ideal-Ideal Power

The TPCv1 algorithm runs as a two stage process—first determining what the ideal Tx power for a radio would be (Tx_Ideal).

$$\text{Tx_ideal} = \text{Tx_max} + (\text{TPCv1_Threshold} - \text{RSSI_3rd})$$

1. Tx_Max—the maximum supported power for a given radio
2. TPCv1_Threshold—User selectable RRM power threshold - default -70 dBm version 4.2 and forward -65 dBm before
3. RSSI_Third—The Third loudest AP in the AP TX Neighbor list



Note This is the TX—not the RX neighbor—see above

If Tx_Ideal is higher than Tx_current, then a power increase is recommended.

If Tx_Ideal is lower than Tx_current, then a power decrease is recommended.

Evaluating a TPCv1 Change Recommendation

The second part of the process involves evaluating the recommended results of the first part and deciding to implement it or not. Since changing the Tx power of an AP also changes the cell boundaries - it can be disruptive to clients. To ensure that a change is necessary a hysteresis is applied.

- For a TX power Increase—Hysteresis = 3 dB
- For a Tx Power decrease—Hysteresis = 6 dB

So determining if a Tx change is recommended looks like this:

$$\text{Tx_Curr} - \text{Tx_Ideal} = N$$

- If N is \leq Hysteresis - No Change is recommended
- If N is $>$ Hysteresis a change is recommended

Implementing a Recommended Power Change

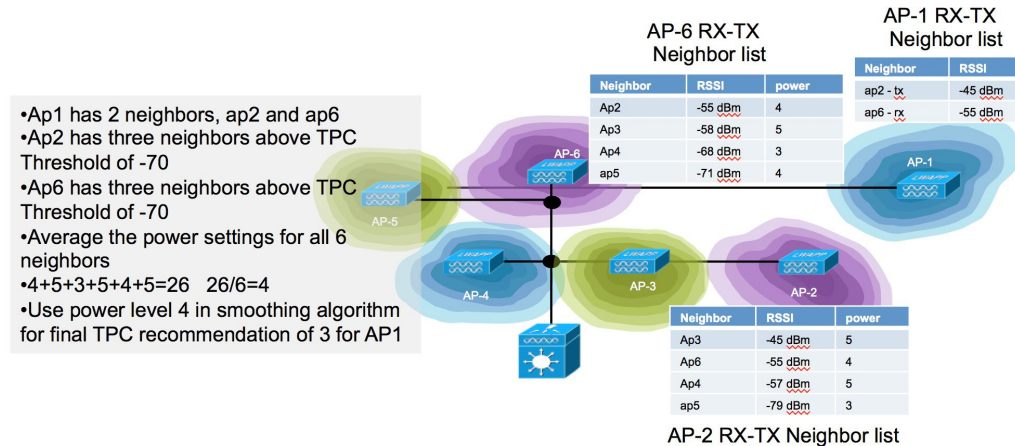
If a power change is recommended it will be implemented by the following rules:

- Decreasing Power—power is decreased 1 level (3 dB) at per TPCv1 interval (600 seconds Default) and the effect would be a gradual reduction and allow for settling in the environment
- Increasing Power—the power level is set to increase 1 level(3 dB) per iteration of TPCv1 until Tx_Ideal or hysteresis is reached.

Before a power change is recommended- it is sanity checked for validity in the RF neighborhood. As an example, we'll use a use case where the AP being set does not have a 3rd neighbor (classically resulted in a power level 1 assignment before smoothing was introduced in version 6.1 of the RRM algorithms). Before applying the new power level to the AP, a check of the AP's neighbors is made to see what they're operating

at. In the case where there is no 3rd neighbor we would look at the two existing neighbors and ensure that the recommended power level is in alignment with they're neighbors and themselves.

Figure 24: TPCv1 Smoothing Algorithm function



The recommendation is matched against the power levels being used for our neighbors neighbor list of AP's, and an average of averages is developed for any neighbor who is on the list at or above the TPCv1 threshold (-70 dBm by default).

Once the recommendation is validated it is passed to the AP for implementation by TPCv1. Let's say the power needs to be decreased by 12 dB to match Tx_curent to Tx_ideal. In all cases the power will be decreased by 3 dB (one step) during each TPCv1 cycle. Then on the next run of TPCv1, the entire process is repeated and if power still needs to be changed- will apply another single 3 dB step. The process is the same for a TX increase to be implemented, 1 power level per iteration will be applied and evaluated.

Like DCA, TPCv1 runs on the RF Group leader, changes in DCA - will affect changes in TPCv1 and vice versa - so it's good that they work together to balance the infrastructure. TPCv1 runs at the same interval as DCA, by default every ten minutes. TPCv1 has no knowledge of channel and assumes that a neighbor entry could be on the same channel at any time (reasonable since DCA runs independently).

On Demand mode schedules a TPCv1 run on the next regularly scheduled boundary, that is if running in default mode, TPCv1 will run on the next 10 minute interval - however the calculation will be held then and TPCv1 will not run or update until you select On Demand again.




Note When all APs boot up for the first time (new out of the box), they transmit at their maximum power levels (power level 1). When APs are power cycled or rebooted, they use their last configured power settings. Transmit Power Control adjustments will subsequently occur as needed.

TPCv2

Cosmetically TPCv2 is almost invisible to the user, in fact it shares all the same configuration parameters on the GUI that are involved in TPCv1. It is however very different under the hood. First, it does not use the 3rd neighbor method or anything else really other than the neighbor lists from TPCv1. The primary difference is that TPCv1 is based on a received energy measurement at the neighboring AP. TPCv2 calculates a cell

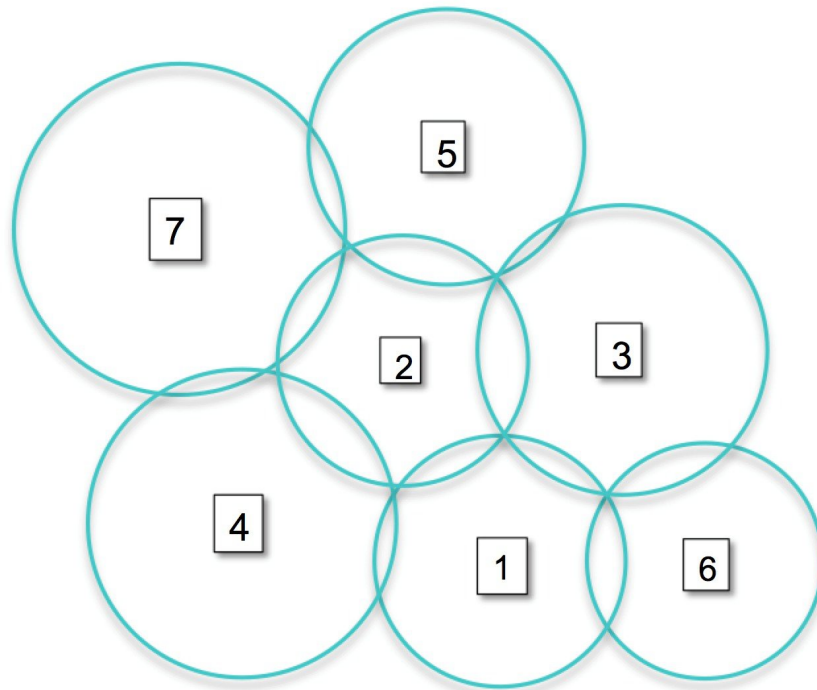
boundary between two AP's based on the measured RF distance between them, and optimizes the coverage based on that calculation.

Figure 25: TPCv2 Two AP example



The problem to be solved is similar to filling a box with balls. Larger balls fill the box quicker, but leave larger amounts of open space between the balls that still could contain more balls. So - add some smaller balls, and fill in those spaces as well. This is essentially the same problem to be solved in getting maximum cellular coverage - larger balls will fill the space quicker but also leave larger open areas, so the equation seeks to optimize the size of all the balls in order to maximize the coverage provided. By increasing and decreasing the cell size, and minding overlap the solution arranges variable size cells for optimum coverage.

Figure 26: TPCv2 Multi-ap Result Example



TPCv2 runs to completion, and power changes are handled continuously. If a cycle concludes that a particular AP's power needs to be raised or lowered the algorithm will run continuously adjusting the power and re-checking the results until Ideal_TX is reached. TPCv2 runs considering all AP's neighbors as equals regardless of channel (default mode), and in channel mode which only considers AP's that are on the "SAME" channel as the AP being solved. Channel mode is enabled at the command line of the RF group leader only. Channel mode is a good choice for High Density Deployments - as this will minimize co-channel interference, while Maximizing coverage and signal between 2 AP's on the same channel. Adjacent channel interference will be addressed by DCA.

```
(Cisco Controller) >config advanced 802.11a tpcv2-per-chan enable/disable (disabled by default).
```

Using Channel mode increases the power significantly.

TPCv2 also adds a radio Utility feature. As the algorithm runs, it models different combinations of power to reach a solution. As it does this it keeps track of the utility of a given radio. TPCv2 will run 10 iterations every TPC Interval (600 seconds) if a particular radio is only need for 3 of the 10 solutions – that radio is marked as 30% utility. You can see the results of this in the show advanced 802.11a/b sum command.

```
(Cisco Controller) >show advanced 802.11b sum
Leader RRM Information
-----
AP_1 : [b2:8d:80] Ch 6* TxPower 1dBm (Level 8)* CHDM 0dBm AP Util 0% dBm
[22/19/16/13/10/7/4/1]
    RxNbrs:: total 5[ 3:-20][ 6:-27][ 5:-30][ 4:-41][ 2:-49]
    TxNbrs:: total 5[ 3:-23][ 6:-23][ 5:-28][ 4:-44][ 2:-46]
AP_2 : [43:4d:50] Ch 1* TxPower 7dBm (Level 6)* CHDM 0dBm AP Util 30% dBm
[22/19/16/13/10/7/4/.]
    RxNbrs:: total 5[ 6:-46][ 1:-46][ 4:-47][ 3:-52][ 5:-59]
```

```

TxNbrs:: total 5[ 4:-39][ 6:-46][ 1:-49][ 3:-50][ 5:-51]
AP_3 : [ba:19:40] Ch 11 TxPower 7dBm (Level 6) CHDM 0dBm AP Util 100% dBm
[22/19/16/13/10/7/./.]
RxNbrs:: total 5[ 1:-23][ 6:-26][ 5:-30][ 2:-50][ 4:-52]
TxNbrs:: total 5[ 1:-20][ 6:-21][ 5:-26][ 4:-49][ 2:-52]
AP_4 : [b4:20:60] Ch 6* TxPower 20dBm (Level 2)* CHDM 0dBm AP Util 70% dBm
[23/20/17/14/11/8/5/2]
RxNbrs:: total 5[ 2:-39][ 1:-44][ 3:-49][ 6:-52][ 5:-59]
TxNbrs:: total 5[ 1:-41][ 2:-47][ 3:-52][ 6:-53][ 5:-56]
AP_5 : [b4:20: 0] Ch 1 TxPower 20dBm (Level 2) CHDM 0dBm AP Util 100% dBm
[23/20/17/14/11/8/5/2]
RxNbrs:: total 5[ 3:-26][ 1:-28][ 6:-34][ 2:-51][ 4:-56]
TxNbrs:: total 5[ 1:-30][ 3:-30][ 6:-35][ 2:-59][ 4:-59]
AP_6 : [cc:d4:20] Ch 6* TxPower -1dBm (Level 8)* CHDM 0dBm AP Util 0% dBm
[20/17/14/11/8/5/2/-1]
RxNbrs:: total 5[ 3:-21][ 1:-23][ 5:-35][ 2:-46][ 4:-53]
TxNbrs:: total 5[ 3:-26][ 1:-27][ 5:-34][ 2:-46][ 4:-52]

```

Other data in this show command includes the last 3 octets of the AP BSSID, and RX neighbor and TX neighbor counts and values for each AP. For instance - AP_1 has a total of 5 RxNbrs - his closest RX neighbor is AP3 at -21 dBm. This is a very useful output and can be used to select AP's for shutting off 2.4 Ghz radios in overly dense situations - such as this example above. From above - we could get good coverage with 3 of the 6 AP's.



Note It is highly advisable to look at the AP's position and coverage on the map with a good understanding of the desired coverage requirements before shutting off radios. This data should be used only as a guide.

The other show command you will want to know is:

```
show advanced 802.11a/b txpower
```

```
(Cisco Controller) >show advanced 802.11b txpower
```

```

Leader Automatic Transmit Power Assignment
Transmit Power Assignment Mode..... AUTO
Transmit Power Update Interval..... 600 seconds
Transmit Power Threshold..... -70 dBm
Transmit Power Neighbor Count..... 3 APs
Min Transmit Power..... -10 dBm
Max Transmit Power..... 30 dBm
Update Contribution
  Noise..... Enable
  Interference..... Enable
  Load..... Disable
  Device Aware..... Disable
Transmit Power Assignment Leader..... Cisco_69:9a:64 (192.168.10.8) (::)
Last Run..... 539 seconds ago
Last Run Time..... 0 seconds
TPC Mode..... Version 2 Per-Channel NO
TPCv2 Target RSSI..... -67 dBm
TPCv2 VoWLAN Guide RSSI..... -67.0 dBm
TPCv2 SOP..... -85.0 dBm
TPCv2 Default Client Ant Gain..... 0.0 dBi
TPCv2 Path Loss Decay Factor..... 3.6
TPCv2 Search Intensity..... 10 Iterations

TPCv2 Plan Quality Index..... Overall -0.5 Coverage 33.2 CCI 3.0 Ratio 1.0
TPCv2 Target Plan..... To be reached in 7 TPC runs

```

AP Name	Channel	TxPower	Allowed Power Levels

```

upstairs_3602e          *1          *7/7 ( 4 dBm) [22/19/16/13/10/7/4/4]
AP_2702E               *6          *8/8 ( 1 dBm) [22/19/16/13/10/7/4/1]
downac                 1           2/8 (20 dBm) [23/20/17/14/11/8/5/2]
upac                   *6          *2/8 (20 dBm) [23/20/17/14/11/8/5/2]
NOS_3600               1           6/7 ( 7 dBm) [22/19/16/13/10/7/4/4]
AP_3502                *6          *8/8 (-1 dBm) [20/17/14/11/8/5/2/-1]
1602I_.560e.1b97      11          6/6 ( 7 dBm) [22/19/16/13/10/7/7/7]

```

This command shows you the configurations for both TPCv1, TPCv2, identifies which is in use (you can only use one) as well as where the TPC threshold and the TPCv2 Target RSSI are set for the global configuration. TPCv1 Threshold and TPCv2 Target RSSI can both be set differently within an RF Profile, and you will not see that in this command. The show command also lists all the AP's WITH their allowed powers as well as the current power level. Keep in mind – the allowed powers are for the current channel assignment, see this example where that makes a difference in 5 Ghz

AP Name	Channel	TxPower	Allowed Power Levels
upstairs_3602e	*(36,40)	*2/5 (11 dBm)	[14/11/8/5/2/2/2/2]
AP_2702E	*(64,60,52,56)	*1/6 (17 dBm)	[17/14/11/8/5/2/2/2]
downac	*(149,153,157,161)	*1/8 (23 dBm)	[23/20/17/14/11/8/5/2]
upac	*(149,153,157,161)	*4/8 (14 dBm)	[23/20/17/14/11/8/5/2]
NOS_3600	*36	5/5 (2 dBm)	[14/11/8/5/2/2/2/2]
AP_3502	*(36,40)	*6/7 (2 dBm)	[17/14/11/8/5/2/-1/0]
1602I_.560e.1b97	*(100,104)	*1/4 (17 dBm)	[17/14/11/8/8/8/8/8]

TPC Min/Max

Regardless of the TPC method you choose you will have the option to limit either the maximum or minimum power settings allowed. TPC Min/Max is a setting that unlike TPC, runs on every controller. It is designed as a safety to prevent going too low or too high in power. The effect is that no matter what TPC sends to the radio, if it is above the Max or below the minimum, the TPC Max or Minimum Value overrides the global assignment.

Why would you want to do this you may ask? Well, we don't always get the AP where we would like it in an installation. The classic use case for which this feature was created is all the AP's being mounted in a central hallway with the intended coverage being on either side of the hall. Because the AP's can see one another quite well in the hallway - power will be reduced to meet the criteria between the AP's and may not be loud enough to reach the edges of the rooms on either side of the hall. Moving the AP's into the rooms - and staggering them down the hall would be better solution - but also may not be possible for a number of reasons. In this case you could use TPC min to ensure that the power levels required to reach the users was indeed honored. Now, this will increase co-channel interference in the hallway itself - however very few of the users will be in the hallway.

Another good example of where to use this would be a lecture hall or classroom that is configured for High Density. People absorb RF, and when the room is full of people - the amount of RF energy you see at the floor could be attenuated by 5 or as much as 10 dB in extremely dense cases. When the room is empty, the power levels of the AP's will drop because the propagation has improved and when it is full, you will need more power (5-10 dB more). If you let TPC manage this without any guidelines, it will eventually apply enough power, but that could be 30 minutes into the class that lasts 1 hour. Setting TPC minimum at the required full class level will ensure that they have plenty of signal at the beginning of class. Yes, the AP's will be louder for all other times, however the unused AP's will only be sending beacons during those times - so not a concern really as the interference will be minimal.

TPC Min and Max settings are entered in dBm NOT Power level index. For this you will want to know the allowed powers for the AP model you are configuring. Power level index is a scale 1-8 from (1)Max to (8) Minimum power for the AP. Not all AP's support 8 Power Levels. The max power an AP can transmit differs

by band, and in 5 GHz will be lowest in the UNii1 band (channels 36-48), higher in Unii2 and Unii2e (52-64, 100-140) and highest in Unii3 (149-165). The advantage to entering this in dBm is that all AP's regardless of channel will exhibit the same power assuming the selection is supported in all 3 ranges, else the AP will be set to the power level it supports in the band closest to the dBm value entered. An AP's allowed powers list is just that – the power levels that the AP can support. To see the allowed powers for the AP's on your network, from the controller CLI - show advanced 802.11a/b summary. Both the channel and the allowed powers for the AP are displayed. From the AP CLI show controller d0/d1 (d0=2.4 and d1=5 GHz).

An AP 2702e -A this list looks like this:

UNii1–15,12,9,6,3 dBm (5 levels 1-5 supported)

UNii2/e–17,14,11,8,5,2 dBm (6 levels 1-6 supported)

UNii3–17,14,11,8,5,2 dBm (6 levels 1-6 supported)

And a 3702e looks like this:

UNii1–15,12,9,6,3 dBm (5 levels 1-5 supported)

UNii2/e–17,14,11,8,5,2 dBm (6 levels 1-6 supported)

UNii3–23,20,17,14,11,8,5,2 dBm (8 levels 1-8 supported)

Power level 1 always relates to the max power that can be made at 6 Mbps - Non-BF (Beam Formed).

Entry of TCP Min/Max values can be done at the GUI - **Wireless> 802.11a > RRM > Tx Power Control(TPC)**

Figure 27: TPC configuration dialogue WLC GUI

802.11a > RRM > Tx Power Control(TPC) Apply

TPC Version

Interference Optimal Mode (TPCv2)
 Coverage Optimal Mode (TPCv1)

Tx Power Level Assignment Algorithm

Power Level Assignment Method	<input checked="" type="radio"/> Automatic <input type="radio"/> On Demand <input type="radio"/> Fixed	Every 600 sec: <input type="button" value="Invoke Power Update Once"/> <input type="text" value="1"/>
Maximum Power Level Assignment (-10 to 30 dBm)	<input type="text" value="30"/>	
Minimum Power Level Assignment (-10 to 30 dBm)	<input type="text" value="-10"/>	
Power Assignment Leader	Cisco_69:9a:64 (192.168.10.8)	
Last Power Level Assignment	478 secs ago	
Power Threshold (-80 to -50 dBm)		<input type="text" value="-67"/>

It is also supported within RF profiles under RRM.

Figure 28: RF Profile - RRM config dialogue - WLC

RF Profile > Edit 'hiCub'

General	802.11	RRM	High Density	Client Distribution
TPC		Coverage Hole Detection		
Maximum Power Level Assignment (-10 to 30 dBm)	30	Data RSSI(-90 to -60 dBm)	-80	
Minimum Power Level Assignment (-10 to 30 dBm)	-10	Voice RSSI(-90 to -60 dBm)	-80	
Power Threshold v1(-80 to -50 dBm)	-70	Coverage Exception(1 to 200 Clients)	3	
Power Threshold v2(-80 to -50 dBm)	-67	Coverage Level(0 to 100 %)	25	

TPC Min/Max Values apply regardless of TPC version in use. Note in the examples above - the default values are - Max =30 dBm and Min = -10 dBm which is effectively off - as AP's do not support these power levels.

Again, this is a per controller/RF profile setting and does not apply to all AP's in the RF Group but only to ones local to the controller where the setting is made.



CHAPTER 8

Coverage Hole Detection and Mitigation Algorithm

The coverage hole detection and Mitigation algorithm is responsible for four things.

1. Coverage Hole Detection
2. Validation of the Coverage Hole
3. Mitigation if Prudent

The first order of business is to detect coverage holes, and second to mitigate them (if possible and wise) by increasing power/coverage. CHDM runs independent of RRM and the RF Group leader. In order to facilitate making decisions at a local level, it runs on every controller. Each individual controller performs coverage hole detection monitoring all associated AP's and thus monitoring every attached client and their received signal levels. Mitigation involves increasing the power on an AP, or group of AP's to improve coverage levels to a certain area where client signals fall below a customer selectable threshold.

The coverage hole algorithm was designed initially as a way for admins to evaluate coverage requirements as the network grows and changes. By monitoring coverage hole alerts an administrator can effectively track and identify areas of the network that might require additional AP's or the re-assignment of existing inventory. Given the dynamic nature of network growth, this is a good thing. Coverage hole correction was envisioned as a way to address short term lapses in coverage by temporarily increasing coverage where needed by extending the reach of existing assets. It had an added benefit in that it could be relied upon, in a properly designed and sufficiently dense network, of providing fault tolerance in the event of an AP failure.

- [Coverage Hole Detection \(CHD\)](#) , on page 77
- [Coverage Hole Mitigation](#) , on page 79
- [Optimized Roaming](#) , on page 79

Coverage Hole Detection (CHD)

Coverage hole detection is based on a 5 second (CHD measurement period) histogram of each Clients Received RSSI values maintained by the AP. Values between -90 dBm and -60 dBm are collected in a histogram in 1 dB increments. A client falling below the configured RSSI thresholds for 5 seconds is marked as a pre-coverage hole event. Pre coverage holes are immediately reported to the WLC and tracked upstream by Prime. At Prime an administrator can review pre coverage hole alarms, and with a location appliance can locate the pre coverage hole on the map.

No Mitigation action is performed on a pre-coverage hole. Pre coverage holes are tracked at the WLC in a 90 second cumulative histogram. A pre coverage hole becomes a coverage hole when it continues to operate below threshold for the entire 90 seconds.

Coverage Hole Detection is based on upstream RSSI metrics observed by the AP. Configurable values are:

- Data RSSI (-60 to -90 dBm) Default -80
- Voice RSSI (-60 to -90 dBm) Default -75
- Min Failed Client Count per AP (1-75) Default 3
- Coverage Exception Level per AP (1-100%) Default 25%

The RSSI value sets the minimum receive threshold for both voice and data separately. Minimum failed client count per AP determines the minimum number of clients that must be in a coverage hole before mitigation can be considered. Coverage Exception level sets a percentage of the overall clients that must be in a coverage hole in order for mitigation to be considered. Both conditions Min failed Clients and Coverage Exception level must be satisfied for a coverage hole to be considered for mitigation.

(Failed Client Count > or = 3) AND (% failed Clients > or = 25%) = Mitigation

It's important when a coverage hole is detected to validate it as best we can and ensure that it's not a false positive. False positives can come from a client that just simply has poor roaming logic and is refusing to move to a better AP option, known as a sticky client.

Additional granularity exists for configuring the thresholds of an individual client as well. What is being tracked at the AP is the overall number of packets that fall below the RSSI thresholds established by the algorithm. Two values are passed from the WLC to the AP for evaluating failed packets against the threshold.

- Num_Failed_Packets—the number of packets received in a 5 second CHD measurement period that where below the RSSI threshold for the associated voice or data client type
- %_failed_Packets—the percentage of total packets received during the CHD measurement period that where below threshold

Both of these conditions must be true in order for a client to be considered in pre coverage hole alarm state. These values are configurable from the CLI only, and the default values should be used unless there is a directed reason to change them.

Notice first that both Voice and data clients are tracked separately based on the WMM UP (user priority). Each received packet from a client is evaluated. This is done through additional configuration values available through the CLI.

At the AP, the 5 second results are collected in a cumulative 90 second histogram, and once every 90 seconds this information is sent in an IAPP message to the WLC and the histogram is re-set. If the client remains in a pre-alarm condition for 90 seconds - it is then considered a coverage hole at the WLC. The WLC will next determine if this coverage hole can and should be mitigated by first determining if the client has a roaming option that it just isn't using. It checks this by determining the location of the client and evaluating its RSSI at other AP's that can hear it. If Other AP's can hear the client above the threshold - the report is marked false and the alarm is re-set.

Coverage Hole Mitigation

Coverage hole mitigation is a fairly simple process once the decision to mitigate is made. If a coverage hole exists AND it meets the criteria (minimum number of clients AND minimum percentage) for mitigation, the AP will increase power by one step. CHDM will then continue to run, and if additional mitigation is called for will re-qualify and power will again be increased by 1 step. This prevents wild and unstable swings in power. Coverage hole mitigation, while operating independent of RRM's DCA and TPC, can have a profound effect on surrounding AP's and the balance of the RF in an environment. Part of the decision to mitigate is to evaluate if mitigation could be successful. Increasing the power of a given AP independently of the RF Group metrics stands a pretty good chance of negatively impacting surrounding AP's. So mitigation is applied very judiciously. The combination of the new detection metrics and the power limits included in mitigation make this a very stable algorithm.

Optimized Roaming

Optimized Roaming was introduced in version 8.0 of the code and without going into detail on the feature (see [HDX High Density Experience deployment guide](#)), borrows the Data RSSI threshold setting from CHDM to set the optimized roaming threshold at which a client will be gracefully dis-associated from the current AP radio. Enabling Optimized Roaming, disables Data RSSI Coverage Hole Detection.



CHAPTER 9

RF Profiles

Throughout this document RF profiles has been mentioned in association with the various algorithms and their functions. Lets take a few moments to run through RF Profiles and the rules for their use as it is important. First getting back to the hierarchy of control implemented in the system, there is the Global Level which encompasses functions that affect every AP attached to either the WLC or the RF group. Below that for RRM, is the RF Profile which inherits properties from the Global choices, but can limit or change the behavior against a group of Access points contained in an AP group. In order for an option to be available in or for an RF profile, in many cases it will need to be enabled first at the global level. We will discuss RF profiles in the context of RRM and its functions here.

Figure 29: RF Profile - RRM configuration Dialogue - WLC

General | **802.11** | **RRM** | **High Density** | **Client Distribution**

TPC

Maximum Power Level Assignment (-10 to 30 dBm) 30
 Minimum Power Level Assignment (-10 to 30 dBm) -10
 Power Threshold v1(-80 to -50 dBm) -70
 Power Threshold v2(-80 to -50 dBm) -67

Coverage Hole Detection

Data RSSI(-90 to -60 dBm) -80
 Voice RSSI(-90 to -60 dBm) -80
 Coverage Exception(1 to 200 Clients) 3
 Coverage Level(0 to 100 %) 25

DCA

Avoid Foreign AP interference Enabled
 Channel Width 20 MHz 40 MHz 80 MHz Best

Profile Threshold For Traps

Interference (0 to 100%) 10
 Clients (1 to 200) 12
 Noise (-127 to 0 dBm) -70
 Utilization (0 to 100 %) 80

High-Speed Roam

HSR mode Enabled
 Neighbor Timeout Factor 5

DCA Channel List

DCA Channels
 36, 40, 44, 48, 52, 56, 60, 64, 100, 104, 108, 112, 116, 132, 136, 140, 144, 149, 153, 157, 161

Select	Channel
<input checked="" type="checkbox"/>	36
<input checked="" type="checkbox"/>	40
<input checked="" type="checkbox"/>	44
<input checked="" type="checkbox"/>	48
<input checked="" type="checkbox"/>	52
<input type="checkbox"/>	56
<input type="checkbox"/>	60
<input type="checkbox"/>	64
<input type="checkbox"/>	100
<input type="checkbox"/>	104
<input type="checkbox"/>	108
<input type="checkbox"/>	112
<input type="checkbox"/>	116
<input type="checkbox"/>	132
<input type="checkbox"/>	136
<input type="checkbox"/>	140
<input type="checkbox"/>	144
<input type="checkbox"/>	149
<input type="checkbox"/>	153
<input type="checkbox"/>	157
<input type="checkbox"/>	161

Extended UNII-2 channels Enabled

- [TPC](#) , on page 82
- [DCA](#), on page 83
- [Coverage Hole Detection](#) , on page 83
- [Profile Threshold for Traps](#) , on page 83

TPC

You can assign a separate Minimum and Maximum TPC power level at the RF Profile level. This will affect only the AP's within the AP group the profile is assigned to. This approach makes it easy to raise or lower the power for a whole group of AP's at once simply by increasing or decreasing the value for either min or max TPC entries. While TPC itself does a fine job of adjusting and maintaining power levels at a correct level for normal installations - this can be very useful for tuning a new high density implementation where SNR will change as more users enter the venue.

You can also assign a different power Threshold's to be used for either TPCv1 or TPCv2. TPC version selection is only available at the Global RF group level and once decided is the same for the entire RF group. You cannot run a different TPC version from the global setting through an RF Profile, you can adjust the target

RSSI or Threshold to match the environments requirements. This is useful if your installation has some areas (a warehouse for instance) where the ceiling height is markedly higher than the rest. Increasing the TPC Power Threshold for a higher ceiling environment will allow achieving the desired coverage at the floor level.

DCA

You can select to enable or disable avoid foreign AP interference contribution to the DCA algorithm. This is particularly useful for areas where there are a high number of rogue interference sources.

You can also change the bandwidth selection that DCA will assign specifically for a group of AP's only. For instance, if you're preference is for 40 MHz channels in most of your installation, but you wind up with a use case in a specific area where you only want 20 MHz (high density deployment for instance) or 80 MHz (classroom with large files stored on a server) that you must support.

You can also modify the DCA channel list. Channels can be customized, however the channel must first be enabled at the global DCA algorithm at the RF group leader WLC as well as at the local WLC that will run the RF Profile (if different from the GL) in order to make it available in an RF profile. This has many practical uses:

1. Managing multi country deployments
2. Assigning groups of channels based on use case
3. Eliminating locally problematic channels (un avoidable interference for instance)
4. Allowing or disallowing UNii2e channels on a case by case basis

Coverage Hole Detection

Defining a coverage hole is very architecture dependent as it is intended to alert when a client is in trouble while in an intended coverage area. Coverage areas differ greatly by architecture. You can customize these values for differing architectures through an RF profile.

If using the Optimized Roaming feature, you can customize the threshold to match the density of the installation. Optimized roaming must be enabled at the global level for the threshold and feature to be applied through an RF Profile.

Profile Threshold for Traps

Trap thresholds for the metrics that RRM monitors may be customized as well through an RF profile. Note that Trap thresholds affect only generation of a trap and have no affect on RRM's operation.

