



Wi-Fi in Healthcare

ekahau
WIRELESS DESIGN



Understanding the bigger picture

Before we begin, we need to level set. This is a technical whitepaper...sort of. This is also a philosophical position on wireless in healthcare. Along with a great deal of useful technical information, our larger goal is to offer important insights. Greater insight yields superior results. If you're reading anything labeled a white paper you obviously care enough to improve your skills and increase the value you bring to projects. To accomplish this, we first need to take you on a journey.

Through this journey, we're asking something from you, the reader. Why? While perhaps unconventional to ask something from the reader, our market focus in this document is on *healthcare*. Without a proper journey visiting the critical ways in which wireless is used in healthcare settings you're just deploying technology. Context is key. Context is gained from real use cases and how minor or severe technology plays a role. Context is also gained from outcomes when everything goes smoothly and when things do not. Context is gained from those affected by these outcomes. In other words, context creates a unique lens to view a common subject—wireless—in a new way.

Technology exists to provide some intended value. Technology inherently has no *business* value itself. Yes, while it may be exciting and enjoyable to many technologists, technology is also incredibly expensive to procure and even more expensive to maintain in proper working conditions. Put another way, technology's role is to serve. In the case of our vertical focus in *healthcare* wireless, technology largely has the intended value to improve patient outcomes. We must never lose context of this goal.

Our journey begins with a brief stint on the proverbial 'soap box'. As IT people—of which I am—we often lose understanding of the business aspects of what we do and get inwardly absorbed in the technology itself. Technologists enjoy the complex and innovative aspects of what technology itself can offer. This isn't inherently bad. However, this pursuit too commonly yields leveraging some intricate and albeit brand new "feature" that results in instability and risk to business operations or more bluntly stated, the core value proposition for which it is designed to serve. My career has informed me that this is the norm; most technologists lack proper context for the service they offer.

By taking this journey with us hopefully you will gain insight and new wisdom to approach wireless and other aspects of your job differently. Moreover, *appreciating* this journey and committing to tracking where it is evolving is what separates technologists from true professionals focused on creating value from their craft.

Importance of wireless' role in healthcare

Caregivers are inherently mobile. A 2006 study published by MEDSURG Nursing observed that out of 146 nurses, each of them walked an average of nearly 5 miles per 12-hour shift. Nurses are constantly moving between patient rooms, fetching items, tending to patients, documenting vital signs and results, administering drugs, tending to medical devices serving patients, moving patients between departments, collaborating and communicating with medical teams and ...you get the idea. New capabilities are now allowing patients to engage directly with their clinical care teams via technology to improve communication, responsiveness and overall quality of care. Oh, and there is this thing people talk about called IoT. Wireless plays a role with that, too, right? The point is that wireless plays an ever increasingly significant role in so many aspects of clinical workflow, hospital operations and the patient experience. This will never change in our lifetimes. Assuming you're a wireless professional yourself, congratulations. You picked a good career field.

Is wireless appropriate for so many aspects of healthcare? Anything inherently mobile benefits from using wireless rather than a fixed piece of furniture that only serves sedentary users. A 150 lb. printer that sits on a desk is well served using Ethernet, but many of the devices we use are inherently mobile. The challenge for wireless is *reliability*. With even relatively modern advances in wireless technology (starting with IEEE 802.11n), I argue speed is no longer a major challenge.

Healthcare has one of the widest diversity of Wi-Fi end-user devices than any vertical.

Types of wireless in healthcare

While likely not a big surprise, not all types of wireless in hospitals is Wi-Fi. Other types of wireless technologies are quite popular in healthcare and include: DECT, legacy paging, WMTS (used for patient telemetry in FCC regulatory domain), cellular, RFID (including all various sub-types), Bluetooth, Zigbee and even more. These are the most common and even more are on the emerging horizon. What's important about comprehending this list is managing the *wireless spectrum* upon which these technologies must compete or rather align with each other. You may only consider yourself a Wi-Fi professional. However, to have reliable Wi-Fi you must have clean or suitably clean radio frequency spectrum. That means some of these other wireless technologies may require more intense investigation to ensure neither interferes with each other.

For the purposes of this paper, we will focus on Wi-Fi. Without question, Wi-Fi is the most leveraged wireless technology in healthcare. It's in practically everything we buy now. The Wi-Fi Alliance (Alliance, 2018) in their 2018 Predictions state that 3 billion new devices will ship resulting in an install base of 9.5 billion Wi-Fi enabled devices. That makes total shipments to date in excess of 20 billion. As for healthcare, by many measures, Wi-Fi is the primary access layer method. Use of Wi-Fi is also used for patient care data and staff communications, alerts and alarms. More on this later.

It is paramount to understand that healthcare has one of the widest diversity of Wi-Fi end-user devices than any vertical I've ever seen. The list continues to grow. It is not uncommon to see well in excess of 50 different Wi-Fi device *types* used within a single hospital. There may be thousands of a single one of these types deployed. The modern hospital usually has both large quantities and high device diversity. While device density is important, it is diversity that has profound effects on design decisions and, in turn, daily operations.

Healthcare wireless use contrasted to other industry verticals

From a pure business perspective, wireless is often critical to daily operations in other industry verticals. Critical implies that substantial financial losses may occur. For manufacturing businesses this may mean that factory lines may stop. For shipping operations, this may mean that inventory cannot be located, labels cannot be printed and packages will not be sent on time. Retail business may not be able to perform point-of-sale transactions for purchases to occur. Hospitality businesses may lose the ability to communicate among staff and guests may lose Internet access, which may impact conferences held in these facilities as well as important business meetings for business travelers. Educational institutions may have to change lesson plans or potentially cancel class depending on that day's instruction. Make no mistake, wireless reliability can sometimes incur heavy financial losses and impact to customer experiences even impacting future revenues. Wireless has gone from a nice-to-have to a necessity. Wi-Fi can be strongly argued that it is now on par as being considered as a utility.

In healthcare, it is often said that patient lives are at stake. Yes, sometimes this is embellished. However, depending on what devices are used, how they are leveraged by clinical staff and how application and device manufacturers depend on wireless in their core function, this can literally be true. We will discuss how mobile devices—namely smartphones—are being used whereby the assertion that patient lives can be affected is indeed real. There are many other examples that involve care team communications and clinical alerts and alarms. From an IT perspective, the lens you need to view the impact of poor reliability and the service you offer is quite different compared to other industry verticals. It's very, very serious.

IT personnel are frequently quite disconnected from the core business for which they work.

Establishing the right mindset

Again, referring to context, the required mindset of IT staff needs to be properly established. It's important to note that the typical healthcare IT person is shielded from clinical activities and the implications of how exactly the services they deliver are leveraged. With larger healthcare organizations it is even more prevalent because IT personnel are often not even located on hospital campuses. Rather, they reside in standard office style buildings. Some likely feel no different than if they, for example, worked for a bank.

Unfortunately, it is my experience that hospital IT personnel are frequently quite disconnected from the core business for which they work. Most may indeed might as well be working for a bank. Working environment and lack of interactions with their internal end customers cannot be minimized as to how it impacts mindset and the context of their work. While IT work is largely relegated to behind-the-scenes type of activities, it is critical to grounding in realizing where the service is being consumed and whom is affected. IT personnel are, in a way, extensions of the clinical teams enabling them to be efficient and effective at improving patient outcomes.



I've often made the following statement to prove a point. *If the mission of healthcare doesn't give significant purpose to your everyday work, healthcare may not be the right industry for you.*

Our healthcare *context* journey is now complete and hopefully this introduction has provided some valuable perspective that I hope motivates your work...today.

Larger trends

Wi-Fi as well as other forms of wireless have continually improved. As 802.11ax is being ratified, we are also now experiencing 5G availability hitting the cellular world. Our expectations are very predictable—we expect more. With those expectations we gain more comfort in the ability to use wireless for more and more applications. In most of the devices we purchase now, a wired Ethernet jack isn't even an option. As we continue to deploy more and more devices and applications wirelessly we have now raised our dependency. So, here we are. We are more dependent than ever on wireless and we should not expect it to change in the foreseeable future except for getting worse. Well, worse or better is a matter of perspective.

High Density

Density...what does this mean? We hear this term in the WLAN world a lot these days. This term refers to more client devices within a given area. Because we're deploying more and more wireless client devices, WLAN vendors have been working on ways to deploy APs in a higher density to allow multiple channels to be operating in the same geographic space. Because 5 GHz is supported in *most* WLAN client devices nowadays, we now have the option to legitimately deploy more APs at 5 GHz than we realistically need to have sufficient coverage. However, density is not without its caveats.

1. Less channel separation
2. Smaller coverage area per AP¹ (cells); more device roaming
3. 2.4 GHz becomes even messier
4. You spend more (Yay for AP vendors!)

¹ Technically cell size doesn't have to be reduced, but it is a common method using this design

5. More APs = more switchports (Yay for switch vendors!)

Less channel separation

Unless you're dealing with fancy, specialized directional antennas to shape RF coverage areas per AP, you're just like everyone else deploying APs with built-in antennas. Both approaches are valid design practices, but co-channel overlap is more greatly introduced using APs with built-in antennas that have omnidirectional coverage for high density deployments. Omnidirectional APs are literally designed to hear in all directions. Hospitals have large, expansive floors that require a great deal of APs. Even at 5 GHz, there are only so many channels available. It's fairly common to see only UNII-1 and UNII-3 channels (36-48, 149-161) deployed in facilities because there are less side effects and incompatibilities related to those eight 20 MHz channels largely due to DFS or lack of certification on newer channel allocations.




Figure 1 - FCC Report & Order 14-30, June 2014 indicating 5 GHz UNII channel allocations

Due to DFS rules, Wi-Fi client device roaming algorithms behave differently. Devices take longer to discover DFS channels and thereby exhibit poor roaming. Many hospitals avoid them unless deploying redundant APs for additional density. If DFS is used, channels in UNII-2 (or UNII-2A as per FCC 14-30 Order) are typically the only DFS channels allowed. Too many legacy devices do not support UNII-2e (or UNII-2C as per FCC 14-30).

Using high density methods for deployment with a limited channel plan, it's not hard for one AP to hear another AP on the same channel with fewer channels to go around. You end up with APs and their clients contending for channel availability time to transmit information between the two cells. You're effectively self-inducing CCI.

Smaller cells result in more roaming

Roaming is still one of the largest reliability and performance issues plaguing all WLANs today. It's not going to significantly change any time soon. Therefore, we must consider it a fact of life. If we're significantly reducing transmit power from APs because they are deployed closer, the client device will perceive it's at the cell edge when the APs signal drops to a certain (client proprietary) threshold. Assuming we're dealing with a well-designed client device, it should roam sooner. While that's a big assumption and a topic for a longer discussion, we can generally assume more APs equal more roaming if the infrastructure is deployed properly and devices are working correctly.

Adding more APs with the same channels within the same coverage area creates problems with greater airtime utilization. Using more APs deployed with a finite number of channels results in less channel separation. When two APs on the same channel are close enough to each other to negatively affect performance, we start to receive interference from adjacent APs on the same channel—CCI. Keep in mind that APs aren't the only transmitters. Client devices equate for roughly half of the traffic and as they are dispersed across the coverage zone of each AP, they also induce a form of interference to adjacent cells operating on the same channel. This causes negative artifacts such as increased retransmissions and,

in turn, creates greater airtime utilization.

When airtime utilization is high, we can expect less transmission opportunities because something else is using the same RF channel. When APs sharing the same RF channel are in close enough proximity to *hear* the transmissions beyond a certain threshold, you end up negatively affecting the performance in both cells. Again, remember that client devices at the edge of an APs cell can create the biggest performance challenge.

Digging a little deeper, every 802.11 frame has what is known as a header and preamble. A frame header (PLCP header) is a small amount of data before an 802.11 frame is sent. The part that is interesting to CCI is how a PLCP header describes what type of transmission will occur and how long. Due to backwards compatibility requirements inherent in Wi-Fi, these headers are sent at 6 Mbps (BPSK) for 5 GHz devices which can be heard at quite a distance. As little as 3.5 dB of SNR may only be needed to decode this header which will can result in all devices in this reception range to refrain from transmitting. Receivers being able to decode the header alone may be enough to cause CCI. ² Per clause 18.3.2.1, we can quote, “In addition, the CCA mechanism is augmented by predicting the duration of the packet from the contents of the RATE and LENGTH fields, even if the data rate is not supported by the STA.” The term “CCA mechanism” is the clear channel assessment state machine that determines wireless medium availability. STA refers to either an AP or 802.11 end-user device.

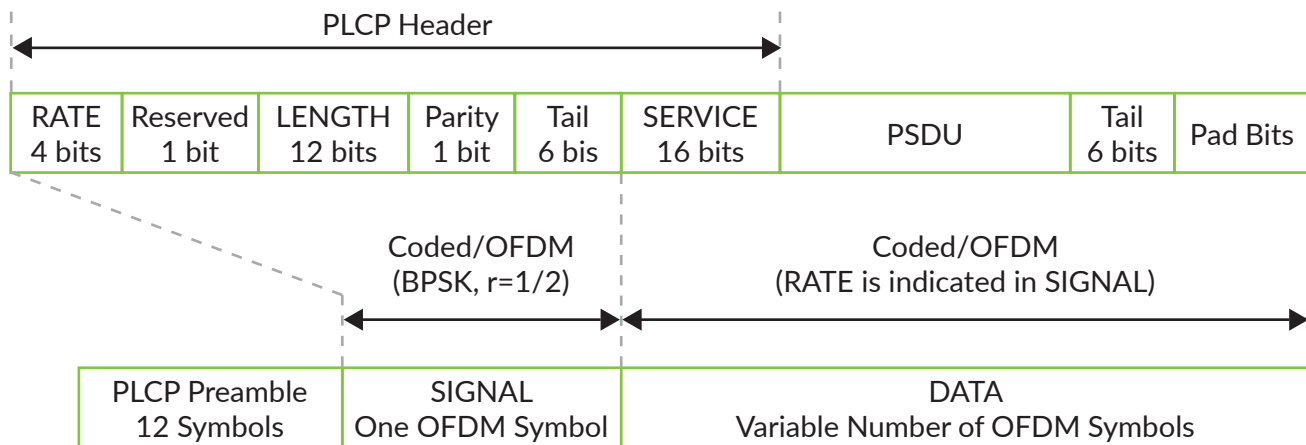


Figure 18-1-PPDU frame format

Figure 2 – OFDM frame format for 802.11a showing PLCP Header with RATE and LENGTH fields. Additional elements are added to subsequent 802.11 amendments, but the core concept still remains.

Alternatively, without more APs per given area, as you add more client devices to an AP other performance issues arise. This is largely due to protocol issues with Wi-Fi and results in more retransmissions, lower PHY rates being used, and more management and control traffic. Wi-Fi isn't very efficient as more clients show up to the party. Hence, as we will explore later, this is why 802.11ax aims to address this challenge.

Frustratingly, most people I encounter think you shrink cell size by simply adding more APs and reducing transmit power of the AP. Yes, the downlink coverage from that AP shrinks, but the AP can still hear just as far. This is synonymous with you lowering your voice. Magically your ears do not get less sensitive, right? The same goes for APs except under special circumstances where vendor proprietary features are available to tweak receive thresholds and are actually deployed correctly. I very rarely see those features leveraged in healthcare.

² Refer to IEEE 802.11-2012 clause 18.3.10.3 CCA requirements



High density results in 2.4 GHz hell

We all know that 2.4 GHz only has three non-overlapping channels in most regulatory domains. Three channels just isn't enough for healthcare deployments where we have many APs in close proximity to each other. What's more is that, all things being equal, 2.4 GHz travels further than 5 GHz. Really, 2.4 GHz attenuates less through our common environments, which results in more distance and therefore has a more usable range.

Stuck with only three 20 MHz channels, signals that travel further and APs deployed closer, we have a problem. You'd end up with channel 1 being at or near full coverage everywhere along with channel 6 and 11 alike. This is the poster child for CCI (co-channel interference). Therefore, we must get tricky and plan for the 2.4 GHz radio on some APs to be effectively turned off. These radios can be in listen only mode for other purposes than serving client devices, but they must not cause CCI in order for 2.4 GHz to be usable. This, in turn, can be a little tricky to figure out which ones should be disabled. It takes very careful attention to RF design to have a high density 5 GHz design while still maintaining solid and usable 2.4 GHz coverage.

Take out your wallet

There are many ways to create high density designs. Regardless of differences in approach you are buying more APs, more AP licenses, more money in annual maintenance spend, and more switches to power them. Some high-density approaches might also add specialized antennas, mounting systems, more backend infrastructure, etc. Even with the simple high-density approach you might also need more electrical power to your closets; modern day APs do not just sip electrical power. Modern APs sip power more like a frat house on a wet college campus. By the way, when adding more electrical power, this creates more heat, which results in more air conditioning load. More density will certainly result in more horizontal cable drops to more APs. Yet, with more switches you may also be adding more fiber strands for more switches back to

I've seen customers with over 80% of their Wi-Fi traffic being used for wireless guest access.

the network backbone. We're also talking about the expensive switchports and not the ones you likely have sitting idle. Get the point? High density AP deployments gets expensive—fast.

Internet of (Healthcare) Things – IoT

IoT is certainly a buzzword we've all become accustomed to. If you ask four people for a definition of IoT you'll likely get four different answers. For the purposes of this document, we'll assume we're talking about more WLAN clients. Let's also assume the term "things" as used in IoT is a term to address the stuff that we don't already have a convenient name for like laptop, workstation on wheels, infusion pump, VoIP phone, smartphone, etc. We know what that stuff is. I would argue that most people have already accounted for those. And because that stuff is largely known, we're not that scared of it.

What if your facilities team said to you that they need to hook up every light bulb to your WLAN? What about every clock? Thermostat? Door lock? Let's assume IoT is that stuff. This is pretty scary because we're talking about swarms, or rather, herds of net new devices that you know nothing about. A proverbial million questions ensue from each one of these new device types. I argue that when we're talking about IoT we're referencing the potential for all of this stuff to emerge as a going concern for our WLAN designs.

For perspective, I have dabbled with electronics a bit and realize just how easy (and cheap!) it is to Wi-Fi enable "things" nowadays. I've made a few smart home devices using WLAN radios with microcontrollers purchased for less than \$3 (check out the ESP8266 or ESP32 microcontroller). Therefore, when I think about healthcare environments and how hard it is to manage what we already have, I immediately stood up and took notice and was quite concerned. That was about five or six years ago.

Realistically speaking, I've not seen IoT manifest and cause any single change to fundamental design practices. Sure, we may end up with it some day, but it will likely be on your next lifecycle refresh of hardware and you'll know even more about what devices you'll need to support by then.

Hospitals are ISPs

I've seen customers with over 80% of their Wi-Fi traffic being used for wireless guest access. This was largely attributed to the fact that they had a very high Internet circuit and no bandwidth cap for guest users. It was a free for all. It's also important to note that employees are frequent users of guest networks for their personal devices. Sometimes there are legitimate business uses for employees to use guest Wi-Fi networks that can be attributed to patient and visitor traffic. While 80% is a large number, it is very common to see 50% of Wi-Fi traffic of many hospital networks attributable to guest users. Make no mistake, if you add high Internet capacity to a network it will get consumed.

An important point to draw out of this is that when guest users do not have any type of restrictions on the bandwidth they consume, this generates Wi-Fi channel contention for clinical devices and applications. Four to five megabits per user is sufficient for real-time video as well as viewing videos through various online video companies. I've not seen reasonable justification for higher throughput limits than that. If the hospital network is designed with 802.11n or higher infrastructure with even reasonable best practices, this should leave sufficient bandwidth available for clinical use. The caveats are really important.

Cellular is critical

At least some of us can remember life before DAS (distributed antenna systems) and indoor cellular. A lot has changed since those days where indoor cellular was a nice-to-have. Physicians often work outside of hospital campuses and only come to hospitals for their rounds with patients. Their cellular phone is a critical piece of their job. For them to pick up a VoWiFi phone that isn't their normal phone number or lacks useful apps on them for their job is a hinderance. The business reasons for providing indoor cellular to enable physician phones is greater than ever. They want to use their smartphone.

Beyond hospital employees and key constituents of the care teams, patients and their loved ones play a huge role in driving the need for indoor cellular. Many healthcare delivery organizations feel at times they are now in in the hospitality business. The hospitality business, like hotels, have needed to cater to guests' needs to ensure their locations are attractive and have the amenities business travelers have come to expect. With healthcare growing more and more competitive, organizations also feel the pressure to make their facilities more attractive.

Unfortunately, it's an expensive proposition to offer cellular services. Hospitals must work with each cellular carrier to deploy infrastructure to transmit their signals indoors. What's more, the hospital IT teams have virtually no control over the service, yet are accosted by physicians, employees and important VIPs when they are not experiencing the level of service they need. Even if guest Wi-Fi is available for use, they usually have to intentionally connect their device(s) to it. Most users don't really care or want to care if they are on Wi-Fi or cellular; they just want their device and applications to work.

As a side note, there are developments happening with the 3.5 GHz innovation band (CBRS) that may offer hospitals the ability to offer an LTE network in their own facilities. It is feasible that the future holds the possibility for hospitals to allow mobile operators to roam on their LTE network. The point is that while Wi-Fi is critical and will not go away, there are even more aspects of wireless that need to be investigated,


Smartphones using Wi-Fi often reveal the smallest weaknesses in Wi-Fi design. In most cases, hospital Wi-Fi networks need to be redesigned/augmented to properly support them.

Wi-Fi calling

Some carriers support Wi-Fi based calling whereby the user can make and receive calls as if they are operating as normal on their cellular network. It is worth mentioning that for users that have this feature available to them, Wi-Fi may relieve some pressure of indoor cellular availability.

Smartphones undergoing mass adoption

Smartphones deserve a special mention. While smartphones have been married to the cellular providers for where you typically purchase them, as IT professionals we should think of them as a mobile computing



platform. Think beyond the phone. These mobile computing platforms run more and more clinical applications as the days progress. If you take a step back for a moment and view the big picture with these types of devices, you have a device that fits in your pocket with incredible computing power, and incredible developer ecosystem and has incredible acoustics for performing audio and video phone calls. If you can provide the same services over Wi-Fi, doesn't that change the possibilities and economics of how we may leverage them?

That is exactly what is going on. There are more Wi-Fi only smartphone phones being deployed in hospitals than ever. Nobody—including CIOs—in their right mind views deploying a traditional VoWiFi handset as a strategic move. Ask a user if they would rather use a legacy VoIP handset and traditional computers or a smartphone for certain tasks. By an overwhelming majority, you will get a smartphone as the answer.

The challenge has been Wi-Fi reliability for smartphones. There certainly are other challenges, but most people would put up with them if they could reliably use these devices for their core job. It has taken a special partnership with Apple and Cisco to close the gap. Android is a more fragmented market, but by and large, at the time of this paper's publishing, Android is further behind. As a notable exception, some enterprise handset makers have delivered enterprise-focused, customized and tuned Android handsets. Many customers have been successful with them, but it takes serious engineering to make them work. I know this because this has been the foundation of more than 50% of our customer engagements over the past several years. Regardless, if it's Android-based or Apple, make no mistake these devices are a challenge to operate with high reliability over Wi-Fi. And, everyone wants them.

Let's paint a picture. Back to the journey we started this article with, consider that you are working as a nurse in a Pediatric Intensive Care Unit administering care to critically ill children. If you have children of your own, hopefully this example hits home. As a nurse, you have been provided a handheld communications device. This is your key to communicating with fellow caregivers, physicians, laboratory technicians as well as key clinical applications that produce time sensitive data that affect the care and urgency of this care to your patients.

In a single device you can eliminate a plethora of devices you are already likely carrying and do more with it. Healthcare specific clinical communications applications are widely available that offer so much value

by integrating with clinical applications, nurse call systems, physiological monitoring, real-time clinician databases and more.

While a handheld device can offer amazing value, the impact can be severe when they do not work reliably. It's important to note just how important it is to attain high reliability and bluntly put—most hospital networks are not properly designed to support smartphones. Traditional VoWiFi devices are a cake walk compared to smartphones. Smartphones using Wi-Fi often reveal the smallest weaknesses in Wi-Fi design. In most cases, hospital Wi-Fi networks need to be redesigned/augmented to properly support them.

Navigating guests through complex hospital

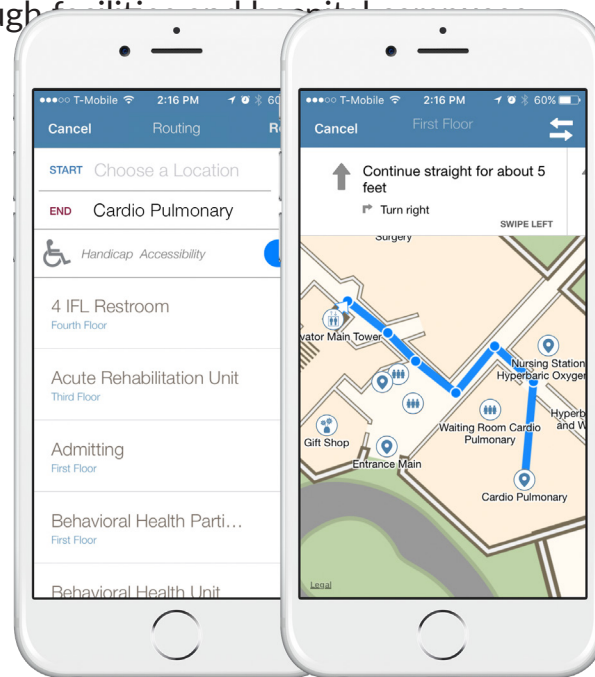


Figure 3 - Image courtesy of Phunware, Inc.

Better known as wayfinding, hospitals are large and ever-changing places that are hard to navigate. Wayfinding is a newer technology that commonly uses Bluetooth Low Energy (BLE) that our smartphones can detect to help users navigate locations using an app. Indoor navigation is beneficial to patient/visitor experiences with healthcare organizations. It can also be helpful for healthcare delivery organizations as well because a department awaiting your arrival can be alerted when you arrive and automatically trigger certain actions. Knowing a patient is already in the waiting room may lead to the possibility of patients being seen earlier if schedule openings occur, automating the check-in process or more

Wi-Fi protocol issues

Different wireless technologies take many different design approaches at the lowest of levels. Some use a one-way broadcast method with no acknowledgement. Some use different frequencies for transmissions than where they receive data. Others use hopping patterns spread out over a large amount of spectrum. Some are designed for peer to peer connections or even point-to-point links. Methods using very low frequencies cover more distance but are subject to channel congestion and low throughput as more devices operate in the same spectrum. In contrast, very high frequencies have shorter ranges and can more easily attain much higher data rates. Some use the same exact spectrum to transmit and receive, thereby



Protection mechanisms always slow down a network.

necessitating methods to contend for transmission time. There are simple methods like in the earliest days of Wi-Fi and there are more elegant ones used with TDD LTE.

Protection mechanisms

The Wi-Fi Alliance touts the benefit of backwards compatibility. A 10+ year Wi-Fi device will likely work with the most modern AP produced today. The other side of this is that backwards compatibility is both a blessing and a curse. In order to support older devices with newer APs using newer 802.11 standards, protocol details must be invoked known as protection mechanisms. Depending on which amendment we are referring to there are different protection mechanisms that are used. Regardless of the finer details, protection mechanisms always slow down a network.

If your network requires higher performance and you are required to support legacy devices, you have a major challenge on your hands. Perhaps you can put the legacy devices at 2.4 GHz only and then use 5 GHz for your modern devices. Otherwise, you will experience some impact to performance. How much depends on the quantity of these devices, the frequency of transmissions and how your network is configured.

Listen before talk

Just as humans do, before we speak we do not want to speak over others. Granted, when my kids were young they didn't follow this rule, but as adults we understand this is the norm. Wi-Fi operates in the same way. As humans, we are effectively using one channel for all verbal communication. A single Wi-Fi AP is also serving wireless devices using one channel (per radio). Therefore, all devices connected to this AP share that same channel for all communications. If another device is sending a lot of data or perhaps time-sensitive QoS data, other devices will wait until its turn arises. Wi-Fi follows a medium contention algorithm that all

devices—including APs—must comply with. This algorithm is referred to as CSMA/CA (carrier sense multiple access with collision avoidance).

When a Wi-Fi device wants to transmit, it will first ensure that there isn't a larger than expected threshold of energy using that channel regardless of if it's an 802.11 transmission—referred to as physical carrier sense. This allows for other unlicensed systems to be sharing the channel or more advanced Wi-Fi modulation to be occurring without legacy devices stepping over the signal. Second, there is a virtual carrier sense mechanism that is based on a countdown system to determine who is allowed to talk next. In Figure 4, an example CSMA/CA state machine is provided. The beginning of the figure includes a potential for physical carrier sense and how virtual carrier sense activities follow from medium activity.

You can think about virtual carrier sense as a random numbering system used to establish a common playing field. Without getting too in-depth in the mechanics, when the contention window starts, a random back-off timer is determined on a per-device basis. The back-off timer is gated by the QoS level of the frame to be transmitted. A lower number is best for gaining medium access to transmit. The contention window is composed of slot times (slots) that each device keeps track of and counts down from the number it started with referred to as a NAV timer. As a slot time passes, the NAV timer is decremented until it reaches zero. When that happens, the device attempts to transmit.

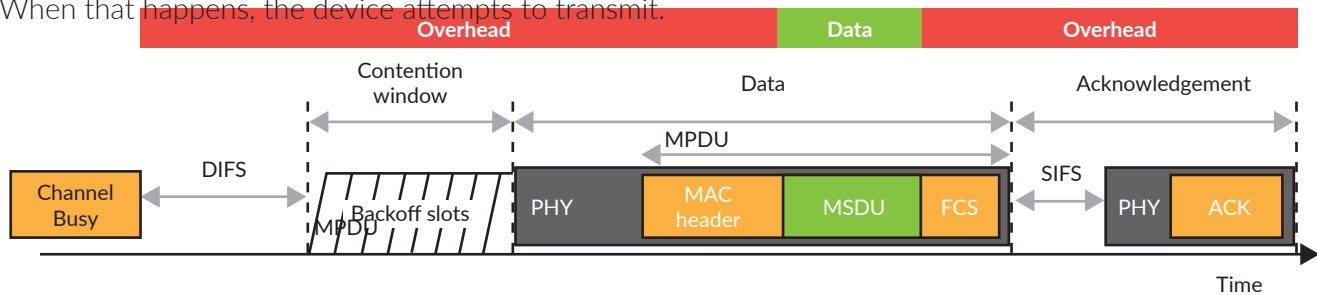


Figure 4 – Example state machine for a Wi-Fi transmission

It's important to note that 802.11 is a very polite protocol when it comes to transmitting. The rules of the game are intended for all parties to operate using the same rules. If the rules are abused, we have unpredictable and unmanageable circumstances. I know of at least two device manufacturers who do not strictly adhere to these rules and it's unfortunate. Other wireless technologies that use the same wireless spectrum are often not quite as polite or friendly.

For a good read, the following article is dated but still highly relevant to today's WLANs. Arbitration schemes have further built on this, but the core tenants are the same.

https://www.cwnp.com/uploads/802-11_arbitration.pdf

Inefficient use of spectrum

I have a close friend who designs radios for a living. He's worked on spacecraft radio systems and cellular systems for much of his career. When I exposed the details of Wi-Fi to him many years ago he was shocked at how brutally inefficient Wi-Fi used wireless spectrum. "20+ MHz of spectrum to send up to 11 Mbps?!", he said. This was an interesting perspective to me because he's used to dealing with often times less than a single megahertz of frequency bandwidth for certain applications.

As technology innovation has increased and brought price points down to consumer levels, we now have incredible radio advances giving us the speeds and capabilities we now enjoy with the more modern Wi-Fi

technologies. In the section on 802.11ax, you will learn how inefficient our current methods still are and how we need to learn from more advanced wireless technologies, like LTE, to build a better Wi-Fi network. All I can say is that Wi-Fi is getting there. There is a new approach to modulation that Wi-Fi will soon be employing, something called OFDMA (orthogonal frequency-division multiple access).

Roaming is Wi-Fi's achilles heal

While improvements have been made, Wi-Fi devices roaming from AP to AP is still one of the most broken parts of the Wi-Fi protocol. We hear of the relatively newer 802.11r protocol (circa 2008) a lot these days. IEEE 802.11r deals with roaming as quickly as possible while still using higher security schemes. I argue that from a roaming perspective, it makes high security methods (IEEE 802.1X) relatively on par with roaming performance associated to personal levels of security (pre-shared keys).

Often alongside 802.11r you hear of 802.11k. This amendment introduced several features, which most notably involves neighbor lists. While not a perfect solution it is better than trying every channel when your signal conditions are bad (i.e. when a device needs to roam). A neighbor list, when working correctly, will provide the channel and address and other relevant information for a client to optimize its search for

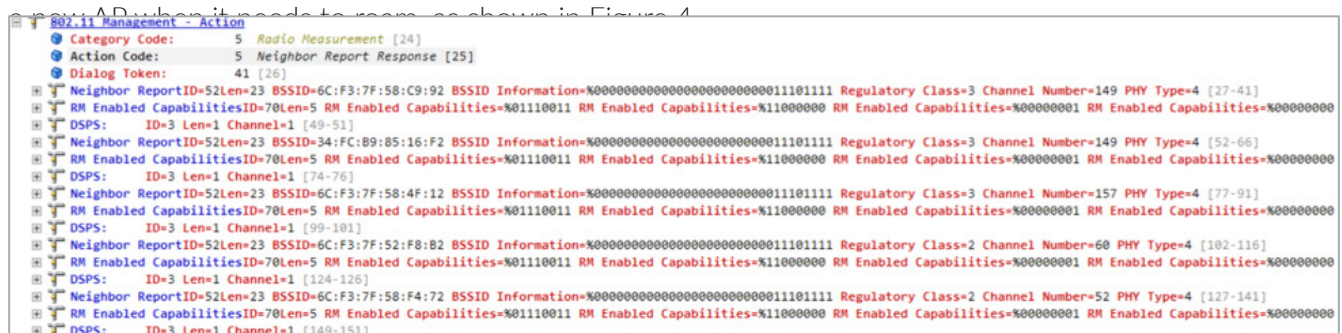


Figure 5 - Example Neighbor List Report

I know of many Wi-Fi experts who don't believe this fundamentally solves the roaming problem—a camp of which I fall into. We often have dinner debates over this topic while completely annoying the tables next to us.

You may also hear of another protocol called 802.11v. There are multiple features included in 802.11v, but we will focus on the roaming related parts. This 802.11 amendment is optional whereby when employed can force or provide a suggestion to a client device to roam. The forceful method should never be used in healthcare. This method effectively disconnects a device from the network and it then needs to find a new AP to associate to. The more polite or suggestive method can send a new roaming target to an 802.11v compatible device and the device can then determine if it's the right course of action. Some manufacturers tout that they support 802.11v but I have never found this feature to solve a problem that couldn't be solved using more fundamental means. Only devices that support 802.11v will benefit, unless you use the harsh method.

It's important to note that a Wi-Fi client device makes the decision when and where to roam. From a protocol design perspective, other wireless technologies/protocols do not leave as much decision power to the client device. Cellular networks, for example, move phones from base station to base station (i.e. AP in a Wi-Fi world) as determined by the infrastructure.



QoS is overrated

IEEE 802.11e introduced the capability of allowing for favoring certain traffic over others. If you enjoyed the section title *Listen before talk* you have a taste for how 802.11 transmitters contend for transmission time. Wireless QoS is not the same as wired QoS whereby you have much more predictability over engineering traffic queues. In a wireless world, QoS is fundamentally based on who gets to talk first or quicker. If we contrast a wireless environment to a wired world we're fundamentally operating over a shared medium much like the old Ethernet hub days. The more modern Ethernet switch allows for simultaneous traffic to be occurring in unison, but a hub fundamentally shares a common bus that all devices on it must share time on.

Because Wi-Fi operates more on the hub model, the IEEE 802.11e amendment and some subsequent amendments, further tweaked the NAV timer and contention windows. For example, under Wi-Fi QoS mechanisms, an 802.11 voice frame would receive a shorter back-off timer than a normal data frame. The sending device must be programmed to detect and take advantage of this capability. If a device is not configured to do this, then it will send traffic using normal methods, even if you enable QoS on your infrastructure.

When Wi-Fi networks were much slower than what we have today, QoS with modern 802.11 deployment was much more important. It's always something you should use whenever designing a network but misconfigured or non-enabled QoS usually is not the source of serious network performance issues. One important point to note is that QoS operates quite differently on wireless than it does with wired networks.

Policy and segmentation

Other non-Wi-Fi wireless protocols assumed from the initial design of their protocol that all devices aren't equal. Some devices must be identified upon connecting to the network that they must have different bandwidth constraints or security restrictions placed on them. For Wi-Fi, all devices are assumed equal. Allowing for higher levels of QoS on a particular SSID provides no real favoritism unless the devices actually participate in the Wi-Fi QoS protocol schemes on a frame-by-frame basis.

From a security perspective, we may want devices to be placed on certain network segments or have a certain ACL or firewall policy applied to them.

For example, a temperature sensor may only need to communicate with one server to send and receive data or commands. This would be a prime use case to apply a restrictive policy on that device, in order to protect it from malicious activity.

Segmentation has been possible using standard RADIUS attributes for years. The problem is that it requires RADIUS and 802.1X authentication. This was available from the relatively early days of Wi-Fi, but it was rarely used, especially in healthcare, except for a few organizations. All RADIUS servers and infrastructure devices needed to be configured perfectly for this to work well. When operating on foreign Wi-Fi networks - all bets were off. As we transition from traditional RADIUS servers to more policy-based access servers, with built-in RADIUS servers, we now have a bigger tool chest to identify, provision, authenticate and categorize devices and thereby apply policy.

Optional features

Wi-Fi has many amendments that sometimes offer a great deal of benefit that often goes unrealized. The IEEE 802.11 protocol is like other protocols that have a lot of optional components vendors can implement as they see fit. It's important to take notice of this when seeking to implement a particular feature while making a general assumption that you can actually implement it.

Wi-Fi challenges not getting easier in short run

For the next five years, I don't expect these fundamental Wi-Fi challenges to change in healthcare networks. We will still see new features and manufacturers will continue to differentiate themselves, but, for example, IEEE 802.11ac wave 2 didn't give us anything we're realistically benefiting from when compared to the initial release. While 802.11ax will change some of how Wi-Fi works, we will still have legacy devices for many years to come. Based on the historical record, we realistically didn't benefit from many Wi-Fi amendments until about five years after ratification. Examples of this include 802.11r, 802.11k, newer 5 GHz channels and more.

Case in point, roaming mechanisms are not significantly changing. WPA 3 is in the process of being introduced and affects from KRACK still being resolved. Patches that addressed the KRACK attack have introduced instability in many hospital and enterprise networks. I generally argue that the newer 5 GHz spectrum in UNII 2e still is not ready for full use into a normal channel plan. Legacy devices will still weigh down healthcare Wi-Fi networks and reduce overall performance, until those devices undergo lifecycle upgrades to newer Wi-Fi technologies.

It's important to put into perspective all of the commentary in this section regarding Wi-Fi protocol issues. Much of this wording could be misconstrued as negative. That is not the intent and to all of the issues we have there is so much benefit that we frankly need to celebrate what we have.

Mind you, no protocol is perfect. All are built on certain assumptions. If you're dealing with an unlicensed RF spectrum, this presents significant additional challenges but you aren't spending billions for dedicated spectrum. Some protocols appear much closer to ideal in comparison. However, if you've ever heard of the "CHEAP, FAST or GOOD? Pick two" philosophy, we can pick one of these immediately because Wi-Fi is first and foremost as mass market, consumer technology so CHEAP must rule to roost. If we choose GOOD, a premium is placed on quality and sacrificing total reliability, performance or speed of realizing the capability in a product. If we want it FAST, then we're would be willing to sacrifice some quality but make mass production and market roll-out the focus. Wi-Fi resides in the CHEAP and FAST world, but we have



been steadily improving in quality whereby some could argue GOOD (enough) at least.

802.11ax primer and expectations for healthcare

What can we likely expect with 802.11ax? The stated goal of 802.11ax is to achieve higher throughput for users in higher client density scenarios. Today, the more client devices connected to an AP, the less efficient the network gets. In fact, the goal of the IEEE 802.11ax task group is to achieve a 4x goal of throughput in this scenario. This will largely be achieved by taking a fairly drastic step in how the radio schedules transmissions at a very deep, fundamental level.

You can look at 802.11n and 802.11ac as a means of increasing over-the-air data rates. You can think of this as the “connection speed” that’s reported by your wireless client. IEEE 802.11a/g topped out at 54 Mbps and 802.11n/ac has increased this speed drastically. What 802.11n/ac did NOT do was increase efficiency when many devices are connected to an access point. Addressing this type of efficiency is a major goal for 802.11ax.

Perspective is important. It seems that IEEE 802.11ac just arrived only a few short years ago and we’re still not widely deployed with it. A vast number of client devices and hospitals are still using 802.11n. I always have argued that the speed problem largely went away from Wi-Fi when 802.11n was introduced. Finally, we could achieve actual throughput (a.k.a. goodput) with a common 802.11n device of 100 Mbps. While 802.11ac added onto this we don’t usually experience as many of these benefits in the healthcare world. This is due to backwards compatibility and not using anything more than 40 MHz channels, along with client devices being limited to one or two spatial streams.

If you look deeper, 802.11ac only introduced a few rates above what 802.11n offered in the 20 and 40 MHz channel width usage. The issue is that these rates (VHT MCS rates 8 and 9) required a much higher modulation called 256-QAM (quadrature amplitude modulation). The requirement to use 256-QAM was at a signal-to-noise ratio (SNR) in excess of 30 dB. Only locations in very close proximity to a ceiling mounted access point will provide this. In healthcare it is largely uncommon to see these rates in use.

In order to understand 802.11ax, we need to first understand some of the fundamental communication




Figure 6 - Example OFDM transmission. Each colours represents a different subcarrier.

mechanics of how OFDM transmissions take place. For reference, 802.11a/n/ac all use OFDM. Let's start with understanding sub-carriers in an OFDM transmission.

Under normal and existing conditions, within a given time segment, a single Wi-Fi OFDM transmission comes from a single source to a single destination. Think about a traditional Wi-Fi transmission like a warehouse distribution center with packages flowing down conveyor belts. Each package is from a single sender to a single recipient. A wireless signal from a single source to a single destination uses all of these OFDM subcarriers to send the message. It's like using multiple beams of information sent in parallel to transfer information faster.

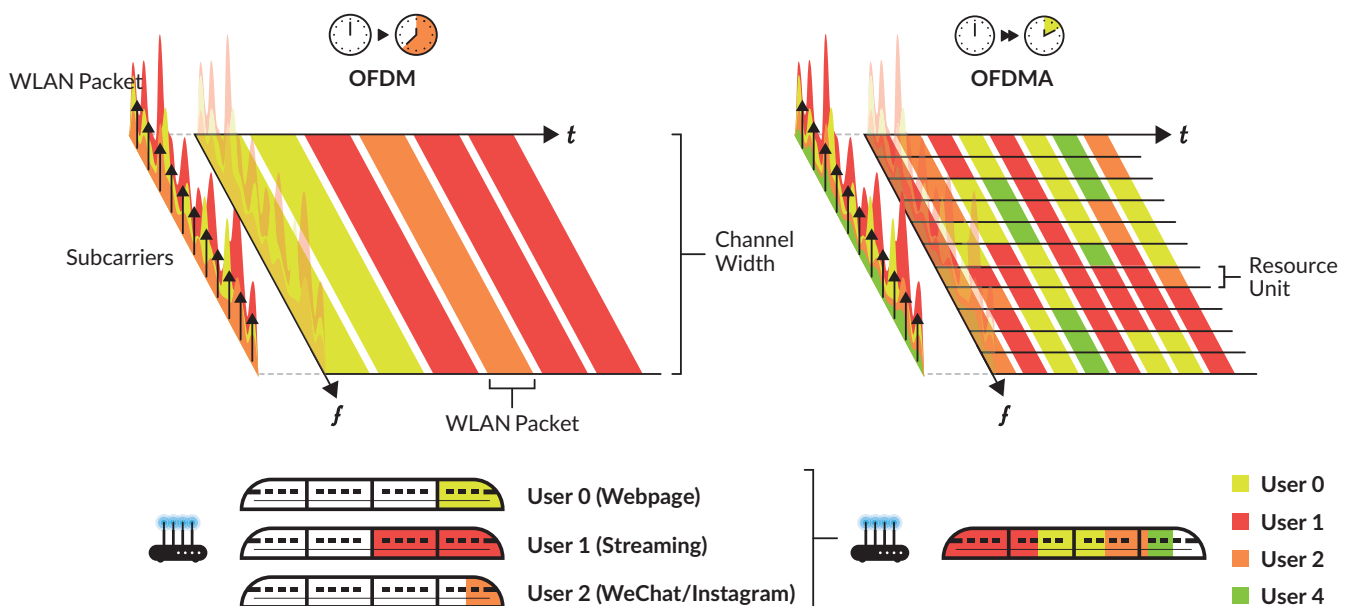


Figure 7 - 802.11ac vs. 802.11ax

Referring to the figure below, picture the OFDM transmission oriented along the ‘subcarriers’ axis. Each vertical arrow indicates a separate OFDM sub-carrier.

If we use our distribution warehouse with packages moving down conveyor belts, the reality is that the width of the belts is large with Wi-Fi. It would be like having large belts to carry packages and only a single package on it at a given time. Today, that’s exactly what happens. Really small packages use those belts and we don’t have a way to address this lack of efficiency. What if you could align multiple packages strategically to consume the full width of the belt and make those other slots available for other packages? In those now free slots, you could do the same thing by again combining other packages all for different destinations. This method increases the overall efficiency and throughput of the system.

There is a new term called Resource Units (RU) you will become more familiar with. RUs are the individual packages stacked up side by side on a conveyor belt in an OFDMA (802.11ax) environment. What’s interesting with 802.11ax is the potential for variable transmit power for individual RUs. This fundamentally changes many aspects of today’s design restrictions of data rates, range and client variance.

802.11ax will use 2.4 GHz and 5 GHz bands. While 802.11ac is only a 5 GHz standard, 2.4 GHz really never went away. We just kept deploying 802.11g/n (2.4GHz) capabilities there even with 802.11ac access points. Yes, 2.4 GHz frequently gets a bad rap, but it is available real estate that we could use more efficiently. In addition, improved power saving mechanisms will also be incorporated. This will greatly benefit IoT types of devices.

BSS Coloring is another feature that will be introduced and has the potential to help with co-channel interference (CCI). Some vendors have proprietary features to modify values that affect whether a device will transmit over another signal it hears from a neighboring Wi-Fi network or AP on the same ESS (multi-AP deployments). It is a complex feature and has sharp, knife edges if you don’t know how to use it. Figuratively speaking, I’ve definitely seen many engineers with cuts on their fingers. I expect BSS Coloring to have a fair amount of complexity and recommend

You should start purchasing 802.11ax chipsets in your client devices ASAP.

There are several other enhancements coming with 802.11ax, but these largely do not relate to normal healthcare scenarios, so we will not cover them here. Most experts I speak to have a general lack of excitement about the addition of a higher modulation—1024 QAM. This isn’t very exciting for nearly all real-world conditions based on what we know today.

Oh, and another important point. It’s a big one. We’re not going to see a lot of these benefits until a significant number of 802.11ax end-user devices are running in an environment. It’s too early to tell where that threshold mix is at this point. So, don’t feel an overwhelming urge to upgrade your infrastructure just yet. However, you should start purchasing 802.11ax chipsets in your



client devices ASAP. Because of Wi-Fi's core tenant of backwards compatibility, you should be safe all things being equal.

How does design differ in healthcare?

This is where we relate context to more technical details for healthcare wireless design. With our journey as a backdrop, we now can look at high level design aspects and make important decisions about how to address key requirements, use cases and the type of device diversity we simply don't see in other industry verticals.

More importantly, you have likely read Best Practice documents or design guides from infrastructure manufacturers. Those documents often have a great deal of useful information, but they do not always put into context the caveats of using or not using certain features. Remember that vendors unique features are promoted in these documents but it doesn't mean they are necessary for your design. This section aims to give certain design aspects relevance to healthcare and your organization's requirements.

Requirements

While requirements come in many forms, the most important set of requirements must be completely based on business aspects. For example, clinicians may have a requirement to leverage a mobile communication solution that allows care teams to communicate by audio, texting and perhaps video. How this is delivered must be further dissected, but it is critical to formulate a subsequent set of requirements that map or align to the business goals and requirements based on a wide range of criteria. Other criteria will range from operational, usability, financial, regulatory, security or any other aspects important to the organization.

The further the flow continues from the initial business requirement, the more the details should get specific and technical.

Let's follow this example a bit further. By interviewing the users and understanding more about their current environment, existing assets and other complementary initiatives, you will learn a great deal more about the solution. Other requirements, such as device consolidation, may arise because they may already be carrying multiple devices. Perhaps there are other initiatives underway that you learn from this process that also require a handheld device to perform barcode functionality. This process will likely lead you to a few potential options whereby cost considerations start to reveal tough decisions. It's important to understand that support and maintenance costs are not free and the fully loaded cost of a new solution should be fully realized in the initial purchase decision. My career has informed me that the ongoing maintenance and support of a solution is usually higher than the acquisition and deployment cost. This is known as the impact of investment into a new solution.

By further continuing with our example, healthcare organizations must adhere to a series of compliance regulations. Some of these are regulatory compliance and others may be internal compliance requirements. By dissecting the entire requirement down to the most fundamental components, this will lead you to requirements or constraints that the final solution must embody. HIPAA is a regulatory requirement for healthcare. For example, one of the criteria to comply with is that patient data must be encrypted at rest. This requirement means that text messages, clinical images or voicemail messages must be encrypted on the device. While this is just one relatively simple example, this leads us to a structured way of design using a wide breadth of categories or



Design for 'time'

If you can picture a time graph whereby you could visualize wireless transmissions occurring and how long each of them takes, this would help paint a picture of how busy a network is. The type of these transmissions and whether there were subsequent retransmissions due to interference, low signal or other means would also reveal important details. You may also see a single set of transmissions consuming a great deal of the available time slots. By looking further, you may

notice excessive low PHY/MCS rates being used. In a healthcare network, this would be cause for concern.

It's important to note that healthcare facilities often have thousands of people traverse them on a daily basis and thousands of medical or computing devices. Often you cannot predict when high quantities of devices may show up in a small area at a given time. Occurrences where high device counts appear in small areas would absolutely wreak havoc on a user experience. To further exacerbate the situation, if client devices were needing to connect at lower signal levels and therefore were using low data rates with even moderate retransmissions, network performance would come to a screeching halt. Healthcare designs must have high signal quality. We will define more what that means, but high signal quality generally means to be in alignment with voice-grade design principles.

Time is a precious commodity in wireless networks. Should a device need to transmit in a Wi-Fi network, it must sense if the medium is busy, contend for a transmission slot and deliver its payload ideally as fast as possible so no more time is consumed than necessary.

A higher quality signal condition is the most fundamental requisite for allowing devices to send their data and receive an acknowledgement. Since this greater spectral efficiency also allows for more available transmission slots, that means that more devices can be supported within the same channel.

Designing for time may cause a wireless engineer to think a little differently about how to diagnose problems, lifecycle upgrade criteria and guide purchase decisions.

General guidelines

The following guidelines are recommended for standard healthcare deployments:

- Use no 2.4 GHz data (PHY) rate lower than 11 Mbps. If no 802.11b devices are still in operation, eliminate all DSSS modulation (1, 2, 5.5 and 11 Mbps) and move to OFDM exclusively. This means not even listing them as 'supported' or 'available' rates. Strictly configure them out of your rate plans on all SSIDs. If your network doesn't work well without them, you likely have areas of poor signal quality. If your equipment vendor allows you to vary rates by SSID, you gain nothing by turning them off on everything else, but leave them on even a single SSID due to backwards compatibility (a.k.a. protection mechanisms). If you must have 802.11b devices still in your network, attempt to put a mandatory timeline on end of life support for the users or departments of those devices. It almost seems surreal, but we still encounter them in production networks.
- Use no 5 GHz PHY rate lower than 12 Mbps but leave on MCS 0. MCS 0 is a special condition because while some devices may have no problem with the lower legacy 802.11a rates disabled, we have seen some devices have difficulty with MCS 0 disabled. The actual data rate MCS 0 uses depends on your configured channel widths and guard intervals, but it can be as low as 6.5 Mbps. While this is less than the recommendation provided of 12 Mbps you will find that MCS 0 operates under different conditions. The lower legacy rates often enabled and even set to mandatory (a.k.a. Basic Rate) by default on many WLAN infrastructure devices and are

It is important to note that disabling lower data rates should only be done if your network is designed for SNR conditions that do not require them.

usually where devices operate in poor signal conditions by default. The amount of traffic I have observed using MCS 0 on even a moderately well-designed WLAN is quite infrequent and not worthy of getting fixated on the actual rate it is operating at. That being said, if you are analyzing a network with MCS 0 as the lowest rate, and you see a high percentage of traffic using it, this means that something is wrong and warrants further diagnosis. It is important to note that disabling lower data rates should only be done if your network is designed for SNR conditions that do not require them. Disabling them if you have even small pockets where SNR is low will result in client disassociations.

- By contrast, do not purchase 802.11ac or higher APs and disable the higher MCS data rates unless there is a bug or serious challenge encountered that is not yet resolved. If this is the case, the goal should be to resolve the issue ASAP and re-enable these high rates. This may sound silly, but there are plenty of examples we have seen. For example, one facility had the ability to disable all high data rates on the guest SSID only. Their goal was to eliminate high bandwidth utilization for guest clients. They certainly achieved that goal. As guest devices used the network, they were constrained to use only low PHY rates. The consequence was a massive increase in channel utilization, creating high airtime utilization, which was starving organizational devices. If rate limiting is required, use other means of accomplishing this goal.
- Too many networks are still using (and reliant on) lower data rates used with 802.11a/g. While it is fortunately rare, some are still reliant on 802.11b. Poor RF design and AP layout may be a reason why these rates are being used. If you haven't had a proper RF design (a.k.a. RF survey) in several years you should strongly consider one.

Common challenges

There are certain root causes to WLAN network woes that always seem to recur. This section lists some of those reasons and other items for consideration when dealing with performance or reliability challenges.

Delusions of quality assurance

Consider yourself a device or infrastructure manufacturer. You conceive of a product, develop it, determine a launch date, maybe even miss the date at least once and continue to work in earnest to launch a product that meets minimum viability requirements. Once it gets out the door you work feverishly to resolve at least the known issues and maybe even in parallel, start to add features to the software. You test these products on what you're able to get and configure them how you think a customer may configure them. Perhaps the manufacturer even has equipment from other



companies to test their product against.

Imagine for a moment you are in charge of quality assurance (Q/A) for that product. Let's assume it's a medical device that uses Wi-Fi. How would you test it? What test gear would you need? Let's say you bought Wi-Fi infrastructure from company XYZ, but which model(s) do you buy? Do you buy gear from the top 3 infrastructure vendors? Assuming you only chose company XYZ, perhaps you order some of each model? Stop and think for a moment about the ramifications that may arise in only selecting one model. Conversely, consider the issue of buying them all and needing to determine a rational method to test all models and have available staff time to test them all.

Assuming you follow this far, which software/firmware should you put on the equipment? Do you keep up with all new software/firmware upgrades? Now, what configuration makes sense for the infrastructure? There is a saying that if you've seen one hospital, you've seen one. It seems that no two are alike and that is often even true for two hospitals owned and managed by the same IT team.

Wow, think about the possible permutations. Even if the equipment was free, time certainly is not and it is also finite. I've never seen an organization able to perform all testing it deemed important due to the many constraints organizations have.

The point of this is to open your eyes to reality. It's frankly unrealistic for device manufacturers to test against every possible permutation, discover a bug, submit it to engineering, issue a new release and retest. While manufacturers would like to do it all they would likely have to charge too



much money to their customers for the product to cover the massive Q/A organization required to do all of this work.

Employing unnecessary complexity

As a designer of hospital networks, it is highly advisable to minimize complexity from designs and keep it simple. As a technologist, this is where you have to keep yourself in check. If you want to employ what seems like an amazing feature, you must consider how refined this feature is and what could possibly go wrong. If you cannot answer that, then stop. If you can, and it's reasonable, then test for it thoroughly before introducing it into a hospital environment.

Quit flipping nerd knobs

Develop a master configuration template by a thoughtful and thorough testing procedure using real devices and applications in a test facility. If testing is successful, deploy it to one more facility to refine and address what you may have missed. Once this is successful, deploy the same master configuration template into another production environment and monitor for performance issues. If a configuration change is warranted, repeat the process from the beginning. Do not start changing configuration settings in production environments without testing. This is the same methodology that mature software companies follow in releasing products for general use. Network designers and operations teams need to take a page from their book.

Quit adding APs willy-nilly

If there are performance problems, find the root cause. Use tools and be trained on how to use them. RF is unfortunately not taught very widely and it's a shame. However, ignorance will not be an excuse your manager or higher executives will appreciate. Quantify the problem and seek help if necessary. You will learn so much from the process. When performance issues occur and even poor signal strength on a single device occurs, do not jump to adding APs as a first resort.

Quit upgrading firmware when there are problems

The same problem resolution process applies when it comes to firmware. Find the root cause of a problem before doing anything. If your TAC or support team is telling you to upgrade, be very

suspect and make them convince you before potentially introducing more issues than you started with. The more you dig the more you will learn.

No hot dog propagation patterns

While this is a tongue and cheek description of an RF propagation pattern, it is a funny way to think of what APs placed in hallways result in. When APs are placed in long hallways the signal travels quite far down the length of the hallways. If a user walks into a patient room at a distance, the signal drops so quickly that the device is forced to roam and likely the user experiences some degree of poor connectivity or worse. Simply put, keep APs out of hallways to prevent long

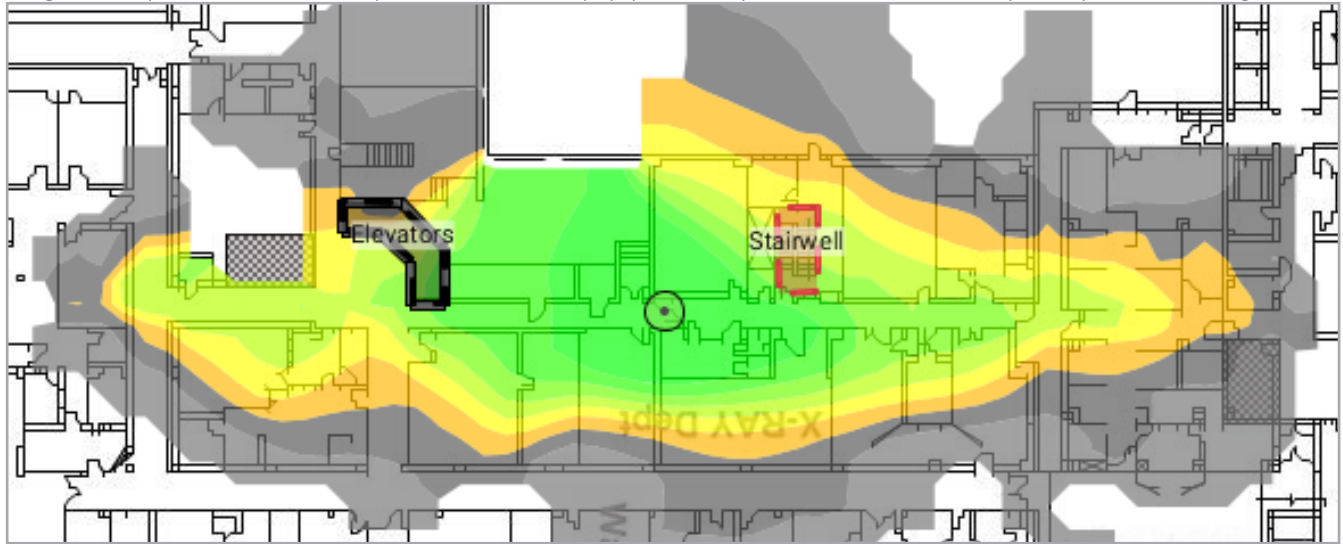


Figure 8 - Hallway AP placement showing elongated shape and sharp decline in coverage as entering rooms away from AP

Be skeptical of automated RF management

Automated RF management features aren't always your friend. These features usually only get you 50% of the way there. Sometimes they cause more harm than good.

Dynamic channelization usually works quite well, but that is mostly useful after initial deployment. Consider freezing channel changes after the system develops a good channel plan. There are not many good reasons for a channel change after an initial channel plan. A single channel change can have a ripple effect. If the cause is interference, find the interference and eliminate it. Remember that some forms of interference move channels, so you gain very little.

As for transmit power changes, this is the most troublesome. Using a transmit power greater than your devices can use is pointless. In fact, it will likely cause your client devices to roam later and operate with poor performance on the uplink (traffic returning to the AP) resulting in retransmissions and a poor user experience. Using exceptionally low power levels usually doesn't solve any major problem in real world conditions, either. Transmit power rarely affects how client devices talk back (DTPC or 802.11h power constraint is an optional 802.11 feature). In addition,



the AP still hears traffic the same, regardless of what its transmit power is set to. Balance your power and design your AP placements according to client transmit power limits.

Different SSIDs share the same radio

Separating traffic on multiple SSIDs doesn't magically solve QoS or performance challenges. All SSIDs share the same physical radios, using the same RF channel, all contending for RF medium time. Valid reasons for using different SSID settings are usually due to different security schemes or a specific configuration setting that renders some devices to perform better. If you find that a device supports the same security scheme as one of your existing SSIDs, and there isn't some special circumstance warranting a quirky configuration setting that will cause harm to the new device, you should use it. If a device manufacturer asks for their own SSID and they do not have a seriously valid reason for it, you need to tell them to pound sand. It happens all the time. Device manufacturers ask for their own SSIDs, but when you really peel back the reasons, most fail to have a valid reason that warrants one.

If you're still using different SSIDs for different VLANs or ACLs you need to employ 802.1X with standard RADIUS or a policy-based access technology (kind of a RADIUS server on steroids). Both can dynamically assign the correct VLAN and/or security policy ideally based on device authentication credentials. There is nothing new about this and it's been available since the early days of WLANs. If you cannot use some type of 802.1X then you may be forced to use far less elegant means, such as MAC filtering, which should be a last resort.

Too many SSIDs

If you're using more than 5 SSIDs, you should be looking at other ways to accomplish your goals. If you understand the concept behind the previous point, this should make more sense.

Be mindful of staff turn-over, operations and maintenance activities

Don't leave a complex environment for the next person. If your design makes sense to you, but you have a hard time explaining it and others are unable to reproduce it, you're likely creating a serious problem for a future engineer and business continuity. Operations teams usually have no time to decipher complex design schemes when an outage occurs and likely will make a change that completely breaks your design. Remember, when outages occur business operations must be restored. These are stressful times for anyone in the position of recovering systems from outages.

If you do not document your work, you are creating a massive liability for your organization. Engineers produce a lot of documents. Barring the risk of sounding insulting, if you are unable to produce good documentation or cannot work with someone in your team to help you, perhaps you should consider a different line of work.

Be very aware of AP outages

No healthcare or high-density network can be completely immune for a performance hit from a single AP outage. If your design needed an AP in a specific location, the likelihood is that if it's offline you will suffer at least minimal, but potentially even major performance issues. Monitoring for AP reboots and outages is critical, which may include PoE switchports that power them. Sometimes APs may fail, but do not show as non-operational. Some infrastructure tools have the ability to run utilization reports. APs with zero clients on them over the course of a full day is likely suspect and should be investigated. Sometimes this is the source of a serious bug or system issue that needs to be addressed.

Call admission control isn't necessary

This may be a topic of argumentation, but from my experience, this feature has caused more bugs and performance issues over the years than it has ever helped. When we were stuck with legacy 802.11 technologies like 802.11b, there was at least an argument why a 300 kbps phone call might need to have an AP reserve channel time. As we moved to 802.11a/g, the argument was much weaker, but one could conceive of a high utilization situation whereby this feature could allow a call to proceed. The argument was usually weak. The reality was that when 20 or so calls occurred under the same AP, there are likely other issues causing performance issues. If you're using 802.11n or higher don't waste your time, because you're just asking for unneeded complexity.

Wireless Security

We will only touch on security as it is a large topic worthy of an entire discussion unto itself. The one point worth making is to state that it is possible to have 100% 802.1X hospitals. This means no WPA2-Personal pre-shared keys(PSKs). SSIDs using a single PSK are nearly impossible to change as device counts and diversity increase. They create management/maintenance

headaches over time.

I used to work for a large healthcare organization and we did not have a single corporate SSID that was using anything less than WPA2 Enterprise. We were not necessarily special, but the point is that it's possible. It takes caring about security and protecting corporate liabilities. We cared and our team was well versed in 802.1X/EAP and knew what to look for when inspecting and testing devices prior to deployment. There is a learning curve, but it is not large. Assuming you use unique credentials per device or device category, your maintenance issues and ability to categorize devices is amazing. This paves the way to deploy authorization rules using new RADIUS server systems protecting where devices are allowed to traverse on your network, including down to protocols, throughput rates, QoS and other controls. While it may seem onerous up front, you're effectively paying it forward in saved time—while attaining high levels of security.

Device consistency and management

If you manage infrastructure you likely go to great lengths to manage it and ensure software and configuration consistency. Do you do the same thing for the other end of your wireless connection? We're talking about the devices connecting to your infrastructure.

Consider the last few tickets or troubleshooting events you can remember. How many of them had a significant component that involved the client device? If we understand how important code levels are to infrastructure, it should be considered equally as important to have well tested and prescribed code levels for every single client type on your network. Consider Figure 9, that uses the same make and model laptop, with the same OS and WLAN hardware. Do you perceive a problem with this behavior? When you get a trouble ticket, are you looking how this may play a



Figure 9 - Same make and model laptop with different WLAN driver versions

The same issue goes for configuration of the driver. Your goal should be to determine the correct version and deploy it using identical configuration settings. In the case of laptops, you have management mechanisms to push updates to them, to update drivers and configuration settings. Smartphones have MDM platforms to aid in at least configuration. Other platforms may have similar capabilities if you investigate.

Again, if you value managing your infrastructure, managing client devices is equally as important.

Wireless is easy?

Many of you have likely heard how wireless is easy. Most of these comments are from people who use it at home and never deal with the complexities of large design, regulations, security and users demanding a great experience. I've even found that those same people who say this ask me for

advice in how to deal with coverage gaps or resolve pesky device disconnects. “Wait, I thought you said wireless was easy?,” I say with a smirk.

The reality is that wireless is far from easy. If you’ve been doing wireless for any length of time, you likely realize how elusive high reliability and performance is to achieve for all devices and areas of your network. As soon as you deploy a fully tested ecosystem of wireless devices, infrastructure hardware, firmware and configuration settings with a full implementation of WPA2 Enterprise security, you’re back in your safe lab environment or non-clinical test facility to handle the next new device rollout, firmware upgrade to WPA3 that’s soon to emerge, security patch, feature enhancement or critical operational bug that wasn’t caught in testing and is causing operational impact. That sentence was a bit of a run-on on purpose. There are so many variables at play.

Sound engineering and testing is like painting that Golden Gate Bridge with a new coat of paint. You started at one end a year ago and just finished at the other end. You look back at the beginning and notice the paint is already fading, cracking and peeling whereby the salty air is eroding the metal bridge structure it exists to protect. If you have many hospital buildings or campuses, it is like having many Golden Gate Bridges to maintain all at once...perhaps even while you’re building a new one or two.

Good tools necessary

I’ve never seen a pair of RF goggles to see RF, but until the day it comes to pass we must rely on tools to provide visibility to the unseen. Tools are absolutely necessary to gain visibility to what is happening at the microsecond level of activity on your WLAN. Without them you are effectively blind to what is really happening.

Just as a carpenter isn’t going to show up to a job site and ‘will’ the nails to sink themselves or expect to cut wood to exact lengths without a measuring device, it is equally unrealistic for you to

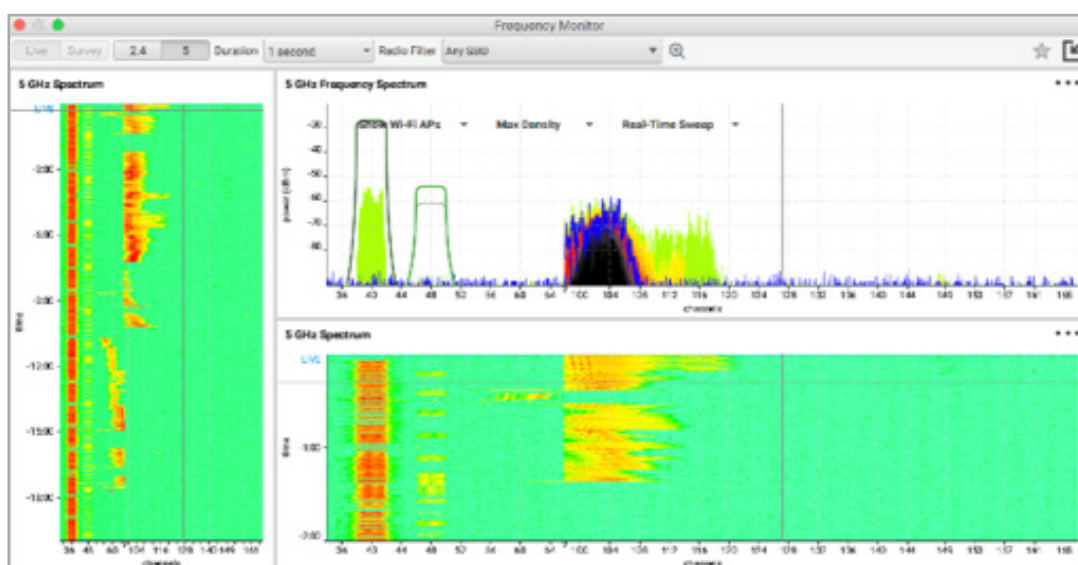


Figure 10 - Ekahau Spectrum Analyzer

Don't ever, ever assume the power from a wireless client device can match your AP's power.

properly test and troubleshoot a WLAN without the right tool. The following tools are laid out in this order on purpose. You always should start with layer 1 and work your way up the OSI stack.

Spectrum Analysis

When running wire for wired networks, you do this right—once—and generally never look back. It's an exercise you perform and properly test so you can generally remove cabling from your list of concerns when troubleshooting issues.

The wiring used in Ethernet networks is your physical layer. With wireless, your radio frequency channels are literally synonymous with your Ethernet wiring. RF is the physical layer for Wi-Fi networks. What's starkly different is that you'll never check off the physical layer from your wireless troubleshooting exercises—ever. We operate using unlicensed (a.k.a. shared) frequency spectrum. In a world where we aren't exactly deploying less wireless (sarcasm intended) we can expect more sources of RF interference.

The only way to gain visibility to spectrum issues is with a spectrum analyzer. There are several ways to accomplish this. Because many Wi-Fi infrastructure APs have built-in spectrum analysis nowadays, this has lessened the need to carry a spectrum analyzer with you in troubleshooting exercises. We always have one when we are performing troubleshooting because you may need to be mobile and perform analysis with the tool beyond the capability that's built into infrastructure APs. Many of the spectrum analysis features built-in APs have their limitations




Figure 11 - Ekahau Spectrum Analyzer example FFT illustrating high level of data fidelity

RF Validation

When you are performing a troubleshooting or analysis exercise you should move to this step once foreign RF interference is ruled out. The next critical component to wireless connection

health is symmetrical RF link analysis. This class of tool is often known as an RF survey tool. To properly perform an RF survey you require an 'ear' that performs equally across all RF channels and bands you are analyzing. We refer to this as a calibrated adapter. Almost all adapters are not. If you plan on purchasing RF survey software you must have a calibrated adapter or you are wasting your time.

When performing an RF validation survey you should perform a series of walks through the area of interest and plot the measurements on a scaled map. These class of tools allow you to import a floor plan of your facility and scale it properly. What is produced from this process is commonly referred to as a heatmap. Wi-Fi heatmaps show you empirical field measurements of RF signals at exactly the location you recorded them. This is powerful data.

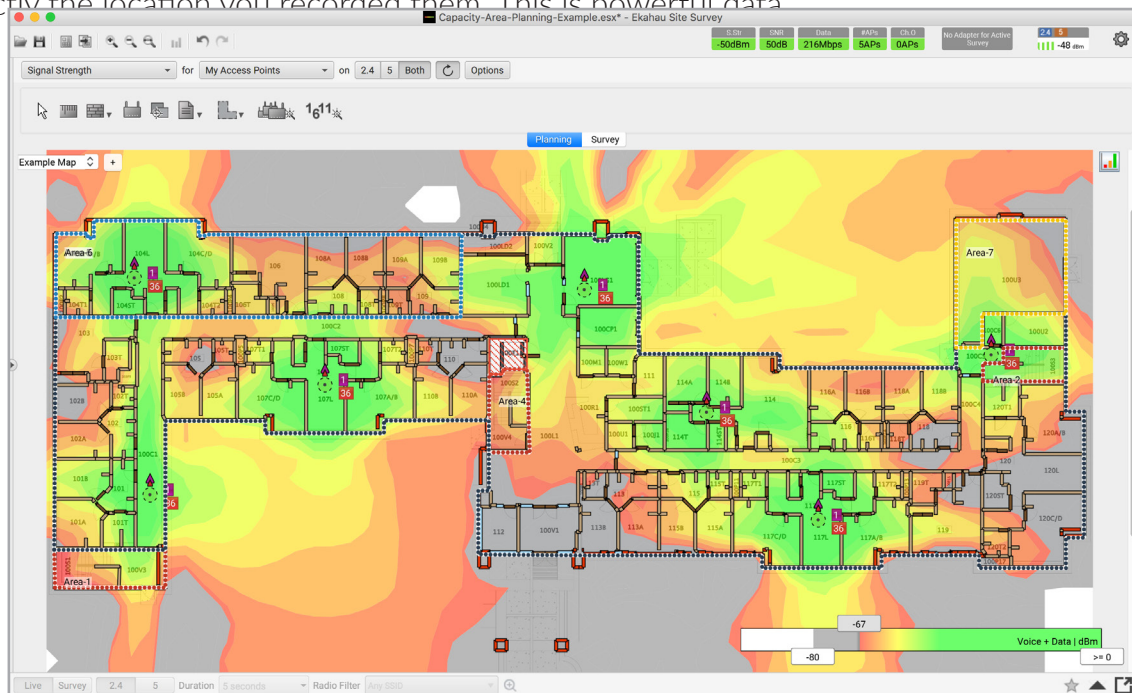


Figure 12 - Ekahau Site Survey software to design and validate RF coverage

Don't be misled

Once a Wi-Fi heatmap is generated be careful about interpreting the data. There are three major mistakes even professionals make. These include:

- Not filtering out rogue or foreign APs
- Assuming downlink power equals uplink power
- Matching receive power of survey adapter to your devices

Ensuring you're looking at only your infrastructure is critical to gaining a proper picture. Also, don't ever, ever assume the power from a wireless client device can match your AP's power. You must determine the upper limits separately. Furthermore, receive power varies sometimes quite wildly from device type to device type, including what frequency channel is used. Placement




Figure 13 - Multiple Wi-Fi radios monitoring dedicated channels for deeper analysis

and orientation of your survey adapter, compared to how devices are placed and oriented in real operating conditions, should also be considered and factored into RF link validation.

802.11 Frame Analysis

The first category of Wi-Fi analysis is gaining a big picture of what is going on with actual 802.11 communication. Some tools provide the ability to sweep many RF channels to detect all Wi-Fi communication happening within a given area. These tools provide great visibility to the number of channels used, number of APs, what clients are connected, what type of security is being used, 802.11 retransmissions occurring and many other important pieces of data to provide visibility to the unseen.

These tools usually allow you to lock into a particular channel to focus on a given client device or access point for richer analysis. Packets captured can be saved and analyzed at a later date or sent off to a peer or consultant. Sometimes these tools do not have the level of raw 802.11 frame analysis found in pure frame/packet capture utilities but are often much more friendly to the novice or casual user to provide insights. If you're a network administrator, this type of tool is a must for getting the big picture.

802.11 Multi-channel Frame Analysis

There is a subtle difference with separating capture from analysis. Sometimes you simply need to capture data. If we're honest with ourselves, 802.11 frame analysis is not for the faint at heart and the skillset is often a scarcity in the vast majority of customer environments. However, the data is absolute gold when analyzing difficult performance issues to determine the root cause.

Generally speaking, if an issue is reproducible, this is one of the first tools that goes out of our tool bag to gain visibility to what is really going on. Simply put, multi-channel frame analysis is the only way to gain true visibility to the unseen when troubleshooting device performance in a production network.


Possessing and knowing how to properly use tools is frankly the difference of calling yourself a professional versus a practitioner. Professionals go deeper. Don't expect that having tools gets you there, you must practice with them and learn how basic RF principles and the 802.11 protocol works for these tools to have any chance of helping you. Without you, the aspiring professional, investing your time with these tools is nothing more than expensive eye candy.

About author



Shawn Jackman, Clinical Mobility
<http://clinical.mobi>

Shawn Jackman is the founder and CEO of Clinical Mobility, Inc., specializing on wireless exclusively in healthcare. Previously, Shawn led wireless standards for a large healthcare organization and built one of healthcare's largest Wi-Fi deployments including the industry's first device certification lab. Co-author of two wireless industry study guides (CWDP & CWSP), co-led the industry's first vendor-neutral wireless design certification, co-chair of AAMI Wireless Strategy Task Force, active participant on customer advisory boards and Certified Wireless Network Expert #54.



About Ekahau

Wi-Fi is like electricity in today's business world. Work stops when the Wi-Fi is down. That's why Ekahau is used by the world's biggest brands (New York Yankees, Amazon, Google, Microsoft, Stanford University, and more) and for the most important events, including Super Bowls and product launches. [Ekahau Connect](#) revolutionizes Wi-Fi planning, analysis, and troubleshooting; [Sidekick](#) sets the industry standard for enterprise Wi-Fi measurement; [Survey for iPhone & iPad](#) changes the game for Wi-Fi site surveys; and [Analyzer for iPhone & iPad](#) redefines the future of troubleshooting through faster, easier, and more accurate network health validation. Whether a corporate office, hotel, hospital, or university – if the Wi-Fi works well, it was likely built with Ekahau.



Find out more

www.ekahau.com

ekahau
WIRELESS DESIGN