



Designing Good Wi-Fi: It's Not Just Plug-and-Play

Introduction

For over 20 years, the Wi-Fi industry has been marketing Wi-Fi on how “easy” it is to deploy Wi-Fi. Most IT professionals with a wired networking background have been sold on the concept that Wi-Fi is a “black box” (though physically, typically an off-white box) that you can simply plug into your network, anywhere you please, and you magically have wireless connectivity for all your users.

While Wi-Fi has been designed to be an ultra-robust protocol, the last 20 years have seen advertised half-duplex speed improvements in the technology from 1-2 Mbps to now >> 1 Gbps. Furthermore, Wi-Fi was originally offered as a best effort tool of convenience, but its use case has morphed into needing Wi-Fi for performance-critical infrastructure to support real-time applications, such as video streaming and voice calling applications, as well as being the primary source of simple Internet connectivity for an ever-increasing array of wireless client devices.

These improvements in potential Wi-Fi performance were achieved over 20 years by increasing “complexity,” which at its core ultimately means using ever-increasingly complex mathematics to bend the laws of physics. Increasing complexity, however, comes at the price of increased sensitivity and fragility; the more performance we want to squeeze out of the Wi-Fi, the more sensitive that performance becomes to factors in its environment. This is most notable in terms of interference, which come from several sources:

- **External non-Wi-Fi Interference:** These are the interference effects from unrelated radio networks operating on the unlicensed band, such as microwave ovens, baby monitors, Bluetooth, cordless phones, etc.
- **External Wi-Fi Interference:** The interference effects of neighboring or co-located (but separate) Wi-Fi networks operating on the same unlicensed band.

- **Self Interference:** The interference effects of your own neighboring APs on the same network.

When somebody declares that “the Wi-Fi sucks” at a particular location, it almost always comes down to a matter of interference from one or more of these sources.

Therefore, to achieve good Wi-Fi performance, the Wi-Fi network must be designed to minimize self-interference and to accommodate the effects of external networks. The quality of the IT infrastructure can only ever be as good as its design. Engineers do not build highways by just paving roads and building bridges wherever they feel like it, and architects don’t build buildings by just randomly throwing together whatever materials that happen to be lying around. How, therefore, can we possibly expect our digital wireless infrastructure network (ie Wi-Fi) to just work by random chance?

Wi-Fi Design as a Formal Process

Design, in any engineering venture, is not a random activity. Rather, design is a well-defined process by which requirements and constraints are gathered, the best options available to meet those requirements and constraints are evaluated and selected, the system is deployed, and the deployed system is measured to ensure it meets the original requirements and is adapted to changes in requirements over time.

For the design of Wi-Fi networks, there are numerous resources available. The Wireless LAN Association (WLA) has recently published a design lifecycle process called the “Standardized WLAN Enterprise Engineering Process,” or “SWEEP.” SWEEP is intended to standardize the design process for Wi-Fi networks and to encourage best practices across the industry, independent of vendor choices or specific applications.

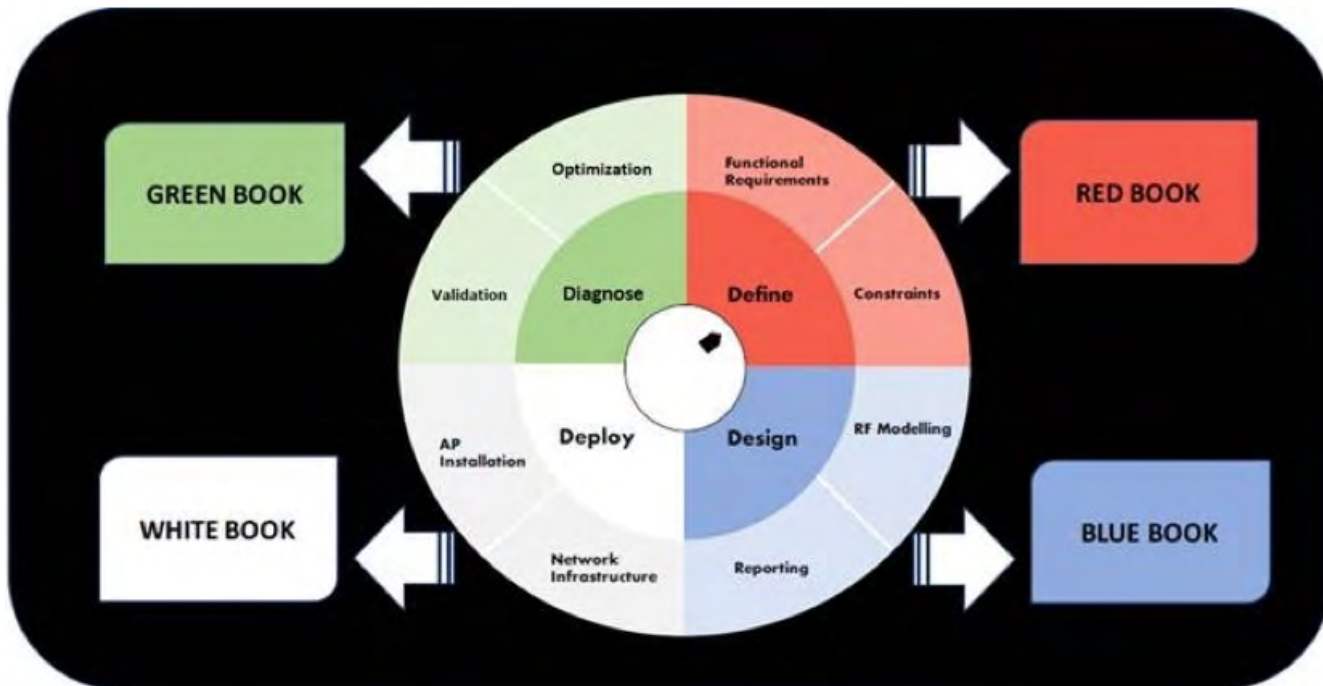


Figure 1: SWEEP design lifecycle process for Wi-Fi networks.²

An overview of SWEEP is shown in Figure 1.

Step 1: Define

The define step requires understanding the “what” of the system, which can be understood in two general categories:

- **Functional Requirements:** What does the system have to do
- **Constraints:** What does the system have to work around

Strictly speaking, requirements are both independent of each other and independent of the design choices you select in later steps. Functional requirements should always be phrased as actions (verbs). Some typical functional requirements for a Wi-Fi network may be as follows:

- Connect smartphones and tablets to the Internet
- Provide at least 5 Mbps downstream / 5 Mbps upstream bandwidth to each client device

- Cover all pre-defined areas with Wi-Fi signal measured at -67 dBm or stronger

These requirements are independent of each other and must be satisfied by whatever design solution you develop (e.g. what vendor you select for your APs).

While every facility is different and thus every Wi-Fi deployment must be uniquely “tuned” to its location, there are general requirements applicable to virtually all types of Wi-Fi networks, as is illustrated below.

¹ <https://wlanassociation.org/blog/,2/19/2018>

² WLA – Transparent Book v1.0. <https://wlanassociation.org/blog/,2/19/2018>




Figure 2: Concept of Virtual Local Area Networks (VLANs). One set of hardware is used to create multiple parallel virtual networks, each with their own unique set of usage requirements.³

By contrast, constraints are highly coupled, usually to one or more requirements and often to each other, and will limit your design options. In extreme cases, the constraints will drive the design and may compromise the ability to satisfy one or more requirements. Common constraints encountered in Wi-Fi deployments are as follows:

- Available budget
- Aesthetics
- Inability to run Ethernet cabling to specific locations
- Building materials
- Pre-selection of a particular AP vendor (e.g. all of our deployments utilize the “Acme AP”)

Functional Requirement 1: Usage

One of the key requirements in defining a Wi-Fi system is to identify the types of devices that will be using the network, the types of applications to be run on those devices, and the level of security that needs to be provided. Usage is naturally unique to each environment, though tends to cluster around the needs of particular

vertical markets. As an example, a hotel will have a different set of users and devices on a nightly basis, but the types of devices brought in (e.g. smartphones, laptops) and the need to separate the network for guests from the network for hotel staff (e.g. maintenance, housekeeping, front desk, etc.) are common across different hotel properties.

When more than one type of network is needed in a facility, which is increasingly common even in small deployments, each wireless service is provided with its own SSID and placed on its own virtual local area network (VLAN), which allows the division of one physical Wi-Fi system into multiple parallel virtual systems that are normally isolated from each other, as illustrated in Figure 2.

Generically, most enterprise Wi-Fi environments will have some or all of the following types of network usage requirements.

- **Visitor / Guest Network:** This network is for people visiting the facility on a temporary basis, generally with smartphones, tablets, and /or laptops, and typically require access to the Internet but not access to any other resources on the internal network. In addition to isolating

³ <https://www.imperialnetsolutions.com/portfolio/>

users from internal network resources, users should also be isolated from each other (i.e. not see each other on the local network). Encryption of the wireless signal is generally not used, so as to facilitate ease of access to the network.

- **Staff Network:** This network is for the staff of the facility, and is intended for devices necessary for operations, which can include tablets, laptops, PCs, and/or dedicated network appliances. This network is typically encrypted to prevent unauthorized access, utilizing either pre-shared key security or access control via a RADIUS or Active Directory Server.
- **Voice Network:** For facilities using Voice over IP and/or Voice over Wi-Fi devices (VoIP or VoWiFi), a dedicated SSID is established with network traffic prioritization. Wireless traffic is usually encrypted with a pre-shared key to facilitate roaming between access points.
- **Security Network:** This network utilizes both appliances (e.g. IP cameras, access control locks, etc.) which may be connected wired or wirelessly, along with monitoring from fixed and/or portable security stations. Wireless traffic is usually encrypted with a pre-shared key to prevent unauthorized access.
- **Network Appliances:** As IoT becomes more commonplace, enterprises start using appliances for improving operational efficiencies. This network generally contains appliances for lighting and temperature control (e.g. NEST thermostats), along with multimedia and other “smart home” appliances. These appliances may not individually be passing a lot of data, but there could potentially be quite a lot of them on the network, depending on the size of the facility.

Functional Requirement 2: Coverage

There is a requirement to provide coverage in all of the locations of a facility where Wi-Fi devices and applications are to be used. This requirement

may seem straightforward, but is influenced heavily by two types of constraints:

- **Physical Constraints:** These are the physical attributes of the facility, including the layout and the (lack of) availability of cabling paths to connect access points. If the facility is new, wiring can often be run inexpensively before the walls are fully erected and sealed. In an existing facility, running new wiring can often be extremely challenging and/or expensive. Even more importantly, however, are the building materials that make up the walls, ceilings, and floors, because radio signals pass much more easily through certain types of materials than others. As an example, open air cubicles, uncoated glass windows, and even drywall are fairly transparent to RF, whereas brick or concrete will highly attenuate Wi-Fi transmissions. There are some materials, like wire-mesh stucco, low-e glass commonly used for outer windows of LEED-certified buildings, and reinforced concrete that serve to block virtually all RF penetration.
- **Logical Constraints:** These are the constraints imposed by owners and operators of the facility. The largest of these is generally budget, though aesthetics, especially in high end properties, can be a major factor. There may also be other radio systems that have to be worked around, such as a legacy or neighboring Wi-Fi or unrelated radio systems using the same frequencies.

Functional Requirement 3: Capacity

In addition to the types of devices using the network, the number of simultaneous devices and the average and peak traffic loads expected by those devices is a critical factor to designing a Wi-Fi network. In high capacity areas, such as classrooms, conference centers, arenas, or stadiums, one AP may be more than sufficient to provide coverage, but multiple APs could be needed to handle the high number and traffic volume of devices.

Accordingly, the area may need to be covered in smaller “cells” in order to handle the quantity of devices and the traffic volume they generate.

Within this capacity requirement is a projection of how capacity needs are going to grow over time. Large enterprise Wi-Fi deployments have a typical expected lifespan of 3-5 years, and smaller enterprise deployments have an expected lifespan of 5-7 years. Meanwhile, the requirements for both bandwidth consumption and quantity of devices tends to more closely follow Moore’s Law (i.e. doubling every 18 months). Accordingly, the network installed today may be perfectly adequate and have margin for today’s need, but will be inadequate in perhaps only 2-3 years. Hence, when selecting components and designing for capacity, it is tomorrow’s loosely defined, requirements, not today’s, which drive the design.

Functional Requirement 4: Control

Consideration needs to be made during the requirements definition phase as to how client devices will access the network, and how the network itself will be monitored and managed.

- **Authentication / Association:** A client device is required to establish and maintain a connection to the network.
- **Access Control:** Generally, there is a need to control who (i.e. what devices) are allowed to access the network. When there are multiple VLANs in place, these user types must be distinguished and placed on the appropriate virtual local area network.
- **Security / Encryption:** Many applications require the use of encryption across the wireless link to prevent unauthorized access to data.
- **Roaming:** When devices move from one area of the facility to another, they will eventually lose contact with the original access point. Hence, there are typically requirements for client devices to move around the facility

and roam between access points without a noticeable impact on application performance.

- **Monitoring and Maintenance:** Most Wi-Fi systems need to be monitored to ensure that all components are online and functioning properly. In case of failure events, some systems may need automatic processes enabled to attempt to resolve the issue or fail into a reduced, but still functional, state.

Functional Requirement 5: Support & Integration

This category of requirements covers how the APs will be supported and integrated into the rest of the network. Common considerations are as follows:

- Provide power to access points
- Provide data connectivity to access points
- Provide infrastructure to facilitate communication between the access points and the central router, central controller, and/or the Internet
- Provide enclosures to protect access points from physical tampering and/or environmental effects (e.g. harsh temperature or weather conditions)

Step 2: Design

Once the requirements and constraints are quantified, an appropriate design solution can be determined.

When assembling a Wi-Fi design, there are only a few fundamental “knobs” that can be controlled, making up the design parameters.

Design Parameter 1: Access Point / Antenna Type

There are multiple access point (AP) vendors providing products with multiple types of capabilities. Many vendors focus their product lines to meet particular types of challenging environments (e.g. stadiums / arenas, K-12 education, etc.), and are optimized and priced accordingly. The most critical, and often most challenging, design decision is the decision of vendor and product line.

While not comprehensive, the following list characterizes the major differences between most product vendors and models.

- **Wi-Fi generation:** The IEEE has extended the original 1997 definition of Wi-Fi numerous times, sometimes with small changes, and other times with entire new product generations. At the time of this writing, 802.11ac is the current product generation, with “wave 2” products beginning to supplant “wave 1” products introduced in 2014. Most vendors also provide models that are of the previous generation (e.g. 802.11n) at lower price points. The most modern APs on the market today will likely be 2-3 generations behind in 4-5 years. Newer generations allow for more complex modulation and coding schemes, which enable faster data rates.
- **Frequency band(s):** 802.11n allowed the use of both the 2.4 GHz frequency band and the 5 GHz frequency band. 802.11ac works entirely on the 5 GHz frequency band, though most access points of this generation will incorporate dual band capability, with 802.11n support on the 2.4 GHz band. Some single band access points are manufactured for particular applications, such as point-to-(multi) point wireless links.
- **MIMO:** 802.11n introduced MIMO (multi-in / multi-out) to allow for parallel data transmissions between multiple transmit and receive antennas. More radios allow for higher throughput, at the expense of power

consumption and space. 802.11ac wave 2 introduces multi-user MIMO (MU-MIMO), allowing the AP to talk to simultaneously communicate with multiple clients in the same environment, for client devices that support the technology. MIMO and MU-MIMO allow for denser utilization of the channel, thus increasing total potential channel capacity and throughput per client device.

- **External vs. Internal Antennas:** From an aesthetic standpoint, most people don’t want to see antennas, so built-in antennas are quite popular. External antennas, however, offer additional flexibility for mounting, along with the ability to add 3rd party antennas with particular profiles for custom applications.

Design Parameter 2: Locations

This design parameter covers the number and layout of the access points within the facility. While seemingly straightforward, there are many factors to consider. First and foremost, where can APs be placed such that cabling can reach it to provide power and data backhaul.

Aesthetic constraints may also limit where APs can be placed within a facility. Placing APs near, or on, metal objects such as pipes, ductwork, I-beams, etc., can change its coverage profile, as those materials serve to act as antennas and distort the coverage profile. Furthermore, APs with internal antennas are generally designed for a particular orientation (e.g. a horizontal ceiling mount), and thus mounting them differently (e.g. vertically on the wall) will change the coverage profile. Finally, APs that are in line of sight of each other (e.g. APs mounted down a long hallway) are more likely to interfere with each other than if there is a physical structure between the APs that will serve to partially attenuate the signals between the APs (e.g. APs mounted every few rooms on alternate sides of the hallway).




Figure 3: 20 MHz channels on the 2.4 GHz frequency band. ⁴




Figure 4: 40 MHz channels on the 2.4 GHz frequency band. ⁵

Design Parameter 3: Channel

Each access point broadcasts a signal on a particular channel, which is specified as a particular center frequency and channel width. On the 2.4 GHz band (802.11b/g/n) in North America, there are 11 channels of 20 MHz size allowed by the FCC. (Channels 12-14 are allowed in some other countries, such as Japan). However, the center frequencies of channels 1-13 are only 5 MHz apart, leading to only three non-overlapping channels, as shown in Figure 3.

The 802.11n spec allows for the optional use of 40 MHz channels on the 2.4 GHz band, by bonding two neighboring channels together. However, given that the entire usable band is only 72 MHz wide, there are no two 40 MHz channel sizes that are independent, as shown in Figure 4. This makes the use of 40 MHz channels impractical in multi-AP deployments, though it is still unfortunately fairly common to see in practice as many vendors allow this channel width in their default settings.

The 5 GHz band is much larger (over 555 MHz, semi-contiguous), and thus makes selecting independent channels and using larger channel widths via bonding neighboring channels much simpler. 802.11a allowed the use of 20 MHz channels. 802.11n allows the use of 40 MHz channels, and 802.11ac allows the use of up to 80 MHz or 160 MHz channels. This is shown in Figure 5. Note that over 2/3 of the frequency space, however, is also used by legacy military, radar, and weather systems, leading to FCC requirements to detect and move off those channels if such external systems are detected. As a result, much of the UNII-2 and UNII-2e bands are not supported by some consumer devices, leading to only two 80 MHz channels and zero 160 MHz channels.

Design Parameter 4: Transmit Power

Transmit power of a radio is proportional to its effective range – the higher the transmit power,

⁴ Coleman, D. and Westcott, D. CWNA Certified Wireless Network Administrator Official Study Guide: Exam CWNA-106. 4th edition. John Wiley & Sons, Inc., Indianapolis, IN. ISBN 978-1-118-89370-8. Copyright 2014.

⁵ Coleman, D. and Westcott, D. CWNA Certified Wireless Network Administrator Official Study Guide: Exam CWNA-106. 4th edition. John Wiley & Sons, Inc., Indianapolis, IN. ISBN 978-1-118-89370-8. Copyright 2014.

802.11ac Channel Allocation (N America)




Figure 5: Channels on the 5 GHz frequency band. ⁶

the more distance that a signal can travel, and/or the more physical materials that it can penetrate, and still be resolved at the receiver. Additionally, a stronger signal at a given distance generally results in a higher signal to noise ratio, allowing for more complex modulation and coding schemes and thus faster data speeds.

In early Wi-Fi deployments, which were primarily driven by the functional requirement for coverage, it was common to turn up the power on the AP transmitter as high as is allowed by FCC and IEEE regulations. This approach was sufficient when most clients had reasonably strong transmitters themselves, such as laptops.

With the emergence of smartphones, tablets, and network appliances, however, there is often a transmit power mismatch that leads to a range mismatch. Most smartphone, tablet, and appliances use relatively weak transmitters in order to preserve both space and battery life. As a result, the situation develops where the client device can receive the relatively strong transmissions of the access point, but the access point cannot receive the relatively weak transmissions of the client device in response. Accordingly, though non-intuitively, the effective coverage area is driven more by the client devices, and the AP power levels should be set to better match the limitations of the clients.

Finally, as compared to 5 GHz, 2.4 GHz has less free space path loss and attenuation through standard building materials, giving it a larger effective range at a given transmit power level. When using a dual band access point, one generally wants to have the coverage area equivalent for both bands. This generally leads to a 4-6 dB difference in power levels on the 2.4 GHz band as compared to the 5 GHz band. In high density environments, it is not unusual to install a denser deployment of APs and then disable the 2.4 GHz band on some of them.

Design Parameter 5: Network Management System

Wi-Fi network activities need to be controlled, coordinated, and monitored. Some access point vendors use access point controllers with relatively “thin client” APs, so that the intelligence of the network is coordinated by a central appliance. Other vendors use standalone APs (i.e. “thick client” APs) where the APs coordinate directly amongst themselves, using a network management system (NMS) to collect usage statistics and log data. There are three types of AP controller architectures that are commonly implemented:

⁶ Adapted from <https://twimags.com/networkcomputing/news/2013/10/graphic-80211-acChannels-all.png>

- **Central Architecture:** In this architecture, all of the intelligence of the network is at the AP controller appliance on the network, and all traffic from the access point is tunneled to the AP controller before being routed to the appropriate destination. As Wi-Fi speeds have increased, the AP controller can become the bottleneck for performance, so this approach is generally no longer used.
- **Distributed Architecture:** In this architecture, all of the intelligence of the network is at the APs themselves, and an AP controller may not even be installed on the network, or if it is, only serves to collect usage statistics and coordinate AP configuration and firmware upgrades. This approach can prove problematic in more complex environments, due to the difficulties in coordinating operational functions across APs, such as client device roaming.
- **Split Architecture:** In this architecture, the intelligence of the network is split between the AP controller and the individual APs. The implementation of this varies by vendor, though typically all data handling functions would be handled by the AP, while management and control functions are handled by the AP controller.

It is also common for wireless networks to be monitored and managed remotely from a remote location, such as a centralized network operations center (NOC). Many vendors have also introduced “cloud controllers,” which are AP controllers that are located on a hosted server on the Internet, managing multiple individual network locations, each consisting of multiple APs.

Design Parameter 6: Wired Network

Fundamentally, an access point is a device that allows one or more wireless client devices to connect to a wired network. The wired network supporting the wireless access points is, in and of itself, a complex system that requires many components, such as cabling, switches, routers, and

modems. The application, coverage, capacity, and control functional requirements drive the need to provide a low-voltage cabling and switch infrastructure that meets these requirements and does not itself become a bottleneck in communications.

Step 3: Deploy

When requirements are being gathered and design solutions are being evaluated, it is very typical to do one or more types of site surveys.

Depending on the project size and scope, funding may not be available on a project to do all of these steps, though performing these surveys are highly recommended, as they can prove invaluable in validating and tuning the network design before and/or during installation. There are generally three types of site surveys that can be performed with specialized software, such as Ekahau. These surveys all generally require accurate floor plans (to scale), and most of them require site access to take measurements.

- **Predictive modeling:** This involves building a mathematical model of the facility, using a software package such as Ekahau Site Survey Pro shown in Figure 6. The floor plans are loaded in, and the building materials of the walls are specified so as to account for their attenuation and reflectivity characteristics. If CAD drawings are available, Ekahau Site Survey Pro can automatically draw in the walls, though if not, walls can also be placed manually.

Once the walls and their characteristics are defined, access points are placed to see how the signal will propagate and self-interfere.



Figure 6: Ekahau Site Survey Pro.⁷

Ekahau Site Survey Pro has a library of the antenna patterns for most common AP vendors and models. APs can be moved around in the model, and their channel and transmit power settings varied, to evaluate how those changes impact the coverage and interference characteristics. These models are fairly straightforward and inexpensive to generate and do not require a site visit, making them fairly easy to perform.

The down-side of predictive modeling is that it is a model based on simplified assumptions – if the floor plans are inaccurate or outdated, or the actual building materials behave differently than what was modeled, then the resulting design will be incorrect. Accordingly, predictive modeling is generally considered a “first step” in a Wi-Fi design to provide a coarse initial estimate as to the quantity, locations, and channel and transmit power settings of the APs for the design.

Pre-deployment or “AP on a Stick” Survey:

This is an on-site survey. If there is no existing Wi-Fi, an AP is temporarily positioned in the environment, and measurements of the room and surrounding rooms are taken to measure actual signal coverage in the environment. The results can be used to refine the predictive model by measuring actual attenuation characteristics of the walls.

This type of survey can also be used to find third party Wi-Fi and non-Wi-Fi devices in the area, so that their presence can be accommodated in the design. Tools like Ekahau Site Survey are specifically designed for these types of measurements, though keep in mind it is only a snapshot in time. Thus, the survey may not be accurate if construction changes are made to the facility or if new neighboring Wi-Fi or other RF systems are installed after the survey is done.

⁷ <https://www.ekahau.com/products/ekahau-site-survey/overview/>




Figure 7: Ekahau Site Survey for passive and active site surveys.⁸

- Post-deployment Survey:** This survey is performed after the network is installed and operational. This is typically performed immediately after network installation to validate that requirements are being satisfied, as well as later on as a diagnostic tool in case of future performance issues. Mechanically, this works very similarly to a pre-deployment site survey, in that the surveyor is walking around the facility with a tool like Ekahau Site Survey and marks their position on the floor plan, so as to build up a complete picture of Wi-Fi performance throughout the facility. Again, keep in mind that this measurement is a snapshot in time, and thus may not reflect performance at a future time.

Step 4: Diagnose

Most of us are not just in the business of designing and deploying networks, but also in maintaining those networks. MSPs make money by having customers pay support contracts, but these are only profitable if the network is working and customers aren't actually pestering you for support issues.

A good quality design can help avoid many support problems. Nonetheless, even with a good quality design, networks occasionally need maintenance and tuning, especially as new client devices get added over time, or

changes are made to the existing Wi-Fi network (e.g. new coverage areas, AP model upgrades, etc.).

One of the key challenges in Wi-Fi is that there is a perpetual stream of new client devices being added to the network, but the Wi-Fi network itself is essentially static once it is deployed. Thus, the expectations for the Wi-Fi network (i.e. the functional requirements) are changing over time while the network itself (i.e. the design parameters) are fixed.

To accommodate this, some level of margin / overcapacity needs to be designed in, as this excess capacity will, in most cases, get consumed over time.

When there are issues on the network, the site survey tools such as Ekahau Site Survey can be used to re-perform an onsite survey. If prior surveys exist, these measurements can be invaluable in determining whether anything has changed in the environment.

Step 5: Repeat

Ultimately, it may come time for a network upgrade, which is driven by the fact that the functional requirements for the network (i.e. usage, capacity, even coverage) have changed to such an extent that the existing design as



deployed is inadequate. Many enterprise networks are designed with a 3 year to 5 year lifecycle, though this is often 7 years or longer in SMB environments.

As the Wi-Fi technology has undoubtedly changed in the interim, the process of gathering requirements and evaluating design parameters needs to be repeated. Existing APs, and existing locations for those APs, may no longer be appropriate for the new solution.

Conclusion

The design of a high-performance Wi-Fi network is a complex engineering task subject to ever-increasing demands on its requirements and constraints. As such, design processes and measurement tools are necessary to identify and validate best practices in Wi-Fi design, and to troubleshoot problems in existing deployments, and thus maximize the performance of the network. Using these tools, high quality Wi-Fi designs can be generated and deployed to maximize the customer's expectations for performance.

Author: Jason D. Hintersteiner, CWNE #171, CWNT

About Ekahau

Ekahau is the global leader in solutions for enterprise wireless network design, optimization and troubleshooting. More than 15,000 customers, including 30% of Fortune 500 companies, run their networks with Ekahau's Wi-Fi planning and measurement solutions.

Our software and hardware solutions design and manage superior wireless networks by minimizing network deployment time and ensuring sufficient wireless coverage across all industries, project sizes, building infrastructures and levels of complexity. We are recognized for delivering the easiest-to-use, most reliable solutions for Wi-Fi planning, site surveys, troubleshooting and optimization. Whether a corporate office, hotel, hospital or university – if the Wi-Fi works well, it has likely been built using Ekahau's Wi-Fi Design solutions.

Learn more about Ekahau's solutions to design, optimize and troubleshoot Wi-Fi networks at www.ekahau.com or contact us at **1-866-435-2428**.



Find out more

www.ekahau.com

ekahau
WIRELESS DESIGN