



Wi-Fi in Hospitality



Introduction

Hospitality is an extremely challenging environment for designing and deploying Wi-Fi networks. **The client device population of the guests literally changes on a nightly basis, so there is a constant barrage of new and unknown devices that are using the network.** It is also quite common for guests to have multiple client devices (i.e., laptops, tablets, and smartphones) connected simultaneously, and to use their own devices for entertainment, such as streaming media from Netflix or Amazon Prime. Guests expect the same level of experience that they are used to at home, and the quality of the Wi-Fi (or lack thereof) is often considered to be more important than basics like clean sheets and towels. Thus, the demands from guests on the network are extremely high.

Add into this that the guest access may not necessarily be the only traffic on the local area network (LAN). Many hotels are using the Wi-Fi network for internal operations, such as directing housekeeping and maintenance staff, security and access control, VoIP systems for the in-room phones, IPTV, minibar monitoring, as well as for creating “smart guest rooms” with IoT devices to control temperature and lighting when the rooms are unoccupied. In New York City, there is also an ordinance that all staff members have to wear an emergency panic button that will indicate the approximate location of the staff

member in the hotel when pressed, and most of these systems work over Wi-Fi.

This white paper will examine the key technologies and design considerations when deploying a wired and wireless network in hospitality environments.

Understanding Your Requirements and Constraints

Before undertaking a wireless deployment, it is critical to understand the key requirements (i.e., what the design needs to do) and constraints (i.e., what the design needs to work around) for the hospitality project. The requirements can generally be categorized as follows:

- 1. Usage:** It is important to capture how the internal LAN is going to be used, and which applications and devices are going to be wireless vs. wired. In all hospitality environments, the network will at least be used for guest access. Guests will generally have relatively modern client devices, though devices like e-readers and low-end tablets and smartphones will generally be 2.4 GHz only, so expect to need to support dual-band 2.4 GHz and 5 GHz operation. Guest access is typically on an unencrypted SSID, and access may be controlled by means of a *captive portal*.

There is usually a distinction between front-of-house, where guests and customers are located and back-of-house, which is purely for operations and facilities.

Client isolation is critical for security, along with strong VLAN isolation from any other network applications. The trend has been for most properties to provide Wi-Fi to guests for free, but some higher-end venues still charge for it or allow for a hybrid model where some level of service is free but a higher tier service needs to be paid for. When doing paid deployments, the captive portal usually needs to be integrated with the hotel's *Property Management System* (PMS), and not all captive portal hardware vendors support this.

Many hospitality environments may also incorporate several other network applications, not all of them wireless, that need to be considered in scoping out the project. Some example applications include security and surveillance, access control, IPTV, VoIP, minibar monitoring, point-of-sale (for restaurants and shops), panic button badges, and in-room IoT devices. If there are multiple applications, each application will likely need to have a dedicated VLAN and, if wireless, a dedicated SSID as well. Note that these additional wireless network applications generally must have WPA2 security to prevent access by hotel guests, but at least some of these wireless appliances likely will not support WPA2 Enterprise (i.e., authentication via a RADIUS server), so at least some of these networks will need to rely upon WPA2 Personal security with a dedicated passphrase. Many of these wireless appliances, especially IoT appliances, are likely to be 2.4 GHz only.

2. **Coverage:** There is usually a distinction between front-of-house, where guests and customers are located and *back-of-house*, which is purely for operations and facilities. Depending on the application, the available wireless SSIDs are often different in these areas; for instance, one generally does not want to provide guest access in the back-of-house areas, nor does one want to provide point-of-sale access on the guest room floors.

Typical deployments will generally allow for one AP to cover a few guestrooms, depending on location and building materials, so the client density per AP is generally not high on the guestroom floors. That said, many hospitality venues also may incorporate an attached banquet facility or conference center, which will require a high capacity design for those specific areas. Thus, in hospitality, there may be a mix of *design for coverage vs. design for capacity* within the same project.

The layout of the property is also extremely important. There may or may not be Ethernet wiring into individual guest rooms. A hotel may be a single mid-rise or high-rise building, or it may consist of multiple buildings that may or may not have fiber or Ethernet infrastructure connecting them. If designing for multiple buildings, then a separate design for the wireless backhaul infrastructure needs to be performed. Motels will generally have outside hallways, meaning that outdoor APs may be required to provide outside-in coverage. Some facilities may require outdoor coverage in particular areas, such as pools, outdoor dining or banquet areas, and even parking facilities.

Peak usage times are usually evenings and weekends, when the guest rooms are occupied.

Expect to need to deal with many different players relying on the same router, switch, and Wi-Fi infrastructure equipment.

- 3. Capacity:** Capacity is generally moderate to high in terms of bandwidth utilization by guests, as users are often streaming movies or television on their devices. Peak usage times are usually evenings and weekends, when the guest rooms are occupied. For hotels with conference centers, peak usage will also be event-based. By comparison, most applications for back-of-house operations usually do not take up significant bandwidth.
- 4. Monitoring and Control:** Most hospitality facilities are generally not running the network themselves, but rather outsource this to a separate *Wireless Internet Service Provider (WISP)*. There may be multiple providers who need to interact, as usually the company managing the guest network is different than the company managing the corporate PCs and property management system, which may be different from the VoIP supplier, which may be different than the IPTV supplier, etc. Expect to need to deal with many different players relying on the same router, switch, and Wi-Fi infrastructure equipment.
- 5. Integration and Infrastructure:** The wireless network is only as good as the wired infrastructure that supports it. If each room has a wired Ethernet port, either for wired guest access or for other applications like VoIP phones and IPTV, it is important to scope out the correct quantity of PoE and non-PoE switches for each intermediate distribution frame, and to ensure that the backhaul (either wired or wireless) to the main distribution frame is not a bottleneck in the flow of traffic. The bandwidth coming into the property

will generally also need to scale based on the number of guest rooms, typically with a fairly low oversubscription ratio of around 10:1 or less.

By comparison, constraints are what the design needs to work around. Constraints will often drive the design solution and may prevent some of the desired requirements from being fulfilled. While constraints are unique to each project, the typical constraints encountered in hospitality deployments are as follows:

- **Budget:** The desired functionality of the network may exceed the funds available to implement it. That said, budget can often help focus the designer on what requirements are critical vs. nice to have. Budget may constrain the choice of AP hardware and vendor, though often this helps from overdesigning the network. In many instances, the requirements can be satisfied by using \$100 access points, so selecting \$1,000 access points with numerous bells and whistles that will never actually get used can prove to be a bad design decision.
- **Aesthetics:** In general, people don't like seeing antennas, which tends to lead to using APs with internal antennas that are less noticeable. Many higher-end hospitality environments will take it a step further, where so much energy and money has been invested in the look and feel of the décor that they do not want to see the access points at all, irrespective of how that affects the actual functionality of the APs. Most APs with plastic enclosures can be safely painted (with non-metallic paint) and have their LEDs disabled, to make them extremely hard to notice. Some venues may require physically hiding the APs in access panels or above drop ceiling tiles.
- **Mounting Restrictions:** In a hospitality environment, this generally means whether or not there is Ethernet wiring available in the individual guest rooms / suites. It is almost always better to put wired APs into particular guest rooms, but hallway deployments are sometimes unavoidable, especially in motels and in older hotels that don't have the wiring infrastructure and where installing such wiring would be prohibitively expensive.



- **External Interference:** If the hotel is reasonably isolated from other buildings, there is unlikely to be any external interference from other sources of Wi-Fi. Most hospitality locations, however, are specifically placed so as to be convenient to access for their intended guests. Hotels in urban or dense suburban environments are therefore likely to have significant exposure to Wi-Fi interference from neighboring buildings. Even when using auto-channel and auto-power algorithms, the levels of interference may prove problematic in certain parts of the building.

It is almost always better to put wired APs into particular guest rooms, but hallway deployments are sometimes unavoidable

Design: AP Selection

Generally, aesthetic constraints will dictate the selection of an indoor AP with internal antennas. If there is Ethernet wiring available in the guest rooms, a wall-plate AP is often a very good option for guest rooms. Wall-plate APs are designed to fit over an existing Ethernet wall-jack, and generally have a built-in switch with 3-4 Ethernet ports, often with one port providing IEEE 802.3af PoE pass-through. The built-in switch allows the connection of wired devices in the room, such as an IPTV or a VoIP guest phone.


For most hospitality applications, APs rated for SMB applications are generally sufficient. A large enterprise access point is likely to be overkill, incurring additional up-front and ongoing licensing costs for features that really are never going to be

used in practice. Unfortunately, most WISPs tend to pick the AP vendor and models that they are most comfortable with, not necessarily the vendor and model that is most suitable to the application. This is why it is important to fully understand your requirements and constraints before selecting an AP model to design around.

For public areas like the lobby, hotel restaurant, banquet and meeting rooms, and conference centers, it may be appropriate to select APs more suited for high-capacity usage. While most WISPs generally will avoid mixing and matching AP vendors at a property, it can be done successfully and may be both advisable and appropriate when the coverage requirements for the guest room floors are very different from the capacity requirements for the public areas.


APs Deployed in the Hallways

Wi-Fi Coverage – 2.4 & 5 GHz



APs Deployed in Rooms

Wi-Fi Coverage – 2.4 & 5 GHz



Channel Overlap - 2.4 & 5 GHz



Channel Overlap - 2.4 & 5 GHz



Figure 1: Hallway vs. in-room deployment for a typical hotel guest room floor.

AP Locations

The job of an AP is to communicate with client devices. Since client devices tend to have much weaker transmitters than APs, the placement of the AP needs to be done so as to facilitate the client's ability to talk to the AP, which is usually counterintuitive to most installers. As a general rule, the APs should be placed as close as possible to the client devices with as few walls or other obstructions as possible between the AP and the client devices. Conversely, to minimize self-interference and to maximize channel re-use, neighboring APs in a deployment should be placed as far apart from each other with as many walls or other structures as possible to minimize how well neighboring APs can see each other. A hallway deployment violates both of these conditions, so it is almost always advisable to deploy APs in the guest rooms vs. the hallways whenever it is possible to do so. Even if it is only possible to get the AP just inside the guest room (i.e., drill a hole above the door and mount the AP on the wall or ceiling just above the hallway door), that alone can be a significant improvement in overall Wi-Fi performance by preventing the

APs from being in line of sight of each other. While hallways are clearly easier from an installation and maintenance perspective, the amount of self-interference in a hallway deployment can be significant, especially on the 2.4 GHz band, which will lower the overall capacity of the network under peak usage conditions and lead to complaints about the poor quality of the Wi-Fi. This is illustrated with an Ekahau predictive design model in Figure 1.

Channel Selection

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 2.

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 3. 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.




Figure 2: 20 MHz channels on the 2.4 GHz frequency band.¹




Figure 3: 40 MHz channels on the 2.4 GHz frequency band.²

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 4. 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.

advisable to use a static and alternating channel pattern on both the 2.4 GHz and 5 GHz bands for all APs on the guest room floors to minimize the chances of self-interference. Only 20 MHz channel widths should be used on the 2.4 GHz band. On the 5 GHz band, some deployments may be able to use 40 MHz or even 80 MHz channels, depending on the density of the rooms and the AP placement. As an example, a typical mid-rise or high-rise hotel should generally use 20 MHz or 40 MHz channels. A hotel consisting of multiple separated buildings with only a few guest suites per building (i.e., only one or two APs per building) may be able to safely utilize 40 MHz or even 80 MHz channels on the 5 GHz band.

When deploying in hospitality environments, the goal is to avoid self-interference between neighboring APs. Even in environments where the APs are located in the guest rooms, it is




Figure 4: Channels on the 5 GHz frequency band. 2.4 GHz frequency band.³

¹ 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.
² 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.
³ Adapted from <https://twings.com/networkcomputing/news/2013/10/graphic-80211-acChannels-all.png>



Most smartphone, tablet, and appliances use relatively weak transmitters in order to preserve both space within the device and battery life.

Transmit Power

Transmit power of a radio is proportional to its effective range – the higher the transmit power, 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, which are commonplace in hospitality environments, 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 within the device 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.



Control and Monitoring: Captive Portal

A *captive portal* is a Layer 3 appliance used to authenticate users prior to allowing them full access to the Internet. Captive portals are typically deployed in network environments that don't utilize any Layer 2 authentication or encryption (i.e., open networks). Captive portals are very common for regulating guest network access in hospitality deployments, and some captive portal appliances are specifically designed for the hospitality vertical. Nonetheless, captive portals can and do get deployed in any vertical where guest access to the Internet is required on an isolated VLAN. Thus, captive portals are fairly ubiquitous for guest networks across many different verticals, such as apartment buildings and student housing that offer Wi-Fi as an amenity, hospitals and healthcare facilities that provide Wi-Fi access for patients and visitors, government offices, schools, etc.

Since the captive portal is a Layer 3 application, it is usually deployed as a dedicated appliance, either as part of the site's gateway or as a bridged Ethernet appliance in-line with the gateway.

As captive portals generally only intercept HTTP traffic, any other traffic, such as HTTPS or POP3/IMAP from an email application will simply get blocked without explanation. This is especially problematic for devices that do not have web browsers built into them

Higher end captive portal appliances were generally cloud managed long before the term “cloud management” was even coined, though lower end appliances operate as standalone devices. There are captive portals that are purely cloud-based, though still require an on-site router or wireless controller that can support the service.

When a user connects to the network and first attempts to reach the Internet, the captive portal will intercept outgoing HTTP traffic and redirect the user to a *splash page*. The user generally has to enter some type of credential and/or accept *Terms of Use* to be allowed through the captive portal to the Internet. At a minimum, the splash page is usually branded with the logo of the property and/or the network operator, and may include images relevant to the site.

Typical login methods for captive portals are as follows:

- **Terms and Conditions (T&C):** In virtually all captive portal networks, the user must acknowledge agreement to a set of terms and conditions to use the network. On most free guest networks, this is the only requirement. To login, the user simply checks a box indicating agreement or clicks directly on the login button.
- **Username and Password:** This login method is common for MDU networks like apartment buildings and student housing, to ensure the user is authorized to use the network. This method is also what is used for paid access, where paid access is based on an hourly, daily, weekly, or monthly subscription. This way, the network owner can track who has paid for

Thus, most captive portals violate the (usually unstated but still important) requirement to make access easy for users.

access, and disable accounts automatically once the subscription expires. The captive portal can also be used to restrict upstream and downstream bandwidth based on the user’s paid subscription level, as well as limit the number of different client devices that can use the same login account.

- **Code:** This method is common in hospitality environments as well as conference centers, to ensure that only guests of the facility are allowed to use the network. The guest is given a code upon check-in to access the network, and enters this code on the splash page. Most hospitality environments are setup to rotate the code on a weekly or monthly basis, though some environments will have a fixed code for the life of the network, which ultimately proves insecure.
- **Name and Room Number:** This method is common in hospitality environments, and generally the captive portal appliance communicates with the hotel’s *property management system (PMS)*, so the user can only log into the network while they are checked into the hotel. This setup is required if the hotel is charging for the Wi-Fi service and having the charge included on the hotel bill, though even hotels that provide free access often use this method to ensure usage only while a user is checked in as a hotel guest.
- **Social Media / Email:** Many network operators use the splash page to collect social media or email addresses from users to populate their marketing lists. This method is common in retail outlets like shops and restaurants.
- **Advertisements:** In an attempt to *monetize* their networks, some network operators use the splash page to show either image or even video advertisements from 3rd parties. By definition, the users are a captive audience.

In most scenarios, once a client device logs in, it will be *MAC Authenticated* for a certain amount of time, so the client device will bypass the splash page and automatically log in. In hospitality environments, this typically expires after 24 hours. In MDU environments, this assignment is usually permanent.

As captive portals generally only intercept HTTP traffic, any other traffic, such as HTTPS or POP3/



IMAP from an email application will simply get blocked without explanation. This is especially problematic for devices that do not have web browsers built into them (e.g., gaming consoles, VoIP handsets, IoT devices, etc.). Such devices usually need to be manually MAC authenticated to be granted access to the network.

Most captive portal appliances are either quite expensive and/or very difficult to program and manage. Usually, both are true. Accordingly, in recent years, many AP vendors have added captive portal capabilities to their controllers because of market demand. Because APs are Layer 2 devices and captive portals are fundamentally a Layer 3 mechanism, the AP is fundamentally the wrong place for this functionality. (Strictly speaking, it is an OSI Layer violation). The captive portals offered by most AP vendors therefore tend to be very limited in their capabilities, and even that basic level of functionality requires convoluted liberties to be taken with the control architecture. Those AP vendors that also bundle Layer 3 solutions as part of their architecture have a better chance of doing this properly.

To make matters worse, most captive portal splash pages are not designed and implemented very well. The login process often requires clicking on multiple screens to get access, in addition to having to provide email or social media contact information or, in some cases, having to watch an advertisement. Thus, most captive portals violate the (usually unstated but still important)

requirement to make access easy for users. Captive portals are still widely used, however. They are necessary if the property is charging for the service, either via credit card or via integration with the hotel's property management system (PMS). While most hotel chains have stopped charging for Wi-Fi access, some higher-end hotel chains still charge for the service. Of these, most hotels have migrated to a hybrid model. In one common iteration, members of the hotel's frequent guest program can enter their membership number and get the service for free, whereas non-members have to pay. In another hybrid model, a fixed bandwidth constraint is applied for the free service, but guests can optionally buy higher speeds. In reality, hotels don't make money off of their Wi-Fi, despite years of trying to figure out ways to *monetize* their Wi-Fi infrastructure. Very few guests actually elect to pay for service, so the hybrid paid plans tend to get used as marketing tools. For example, hotels that specialize in hosting conferences will use captive portals to issue free login vouchers for conference attendees.

Even facilities that do not charge for service still utilize captive portals, based on the "requirement" to provide the T&C legalese warnings that every user agrees to but virtually no user ever actually reads. Such text is unique to each property and/or service provider, though generally include statements such as the user agrees that the service is provided as-is, the user won't download copyrighted material, the user won't hack anybody from this location, the user shall willingly give up his or her first-born child to a particular Satanic



cult for ritualistic sacrifice, etc. It is not clear that stating such legalese and forcing users to acknowledge it is actually required to legally protect the network operator in case of an incident, though most network operators don't want the legal battles and expense to find out through litigation.

Network operators do need to have a mechanism to be compliant with the Communications and Law Enforcement Act (CALEA), which allows a Federal law enforcement agency to obtain a warrant to track the online connection activities of a specific subscriber. Most captive portal appliances will enable compliance by allowing a specific user to be tagged and have all of their metadata (e.g., IP addresses of external networks they are communicating with) to be logged to a syslog server. CALEA warrants are not very common, especially in the hospitality vertical, but there can be hefty fines levied against the network operator and property if a warrant is issued but the network equipment is not in place to comply.

Control and Monitoring: Network Management

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:

- **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.

Hospitality environments can be made to work with any of these types of control architectures. Many hospitality properties that are severely budget-constrained may still operate in standalone mode, where all of the APs are configured individually and there is no central monitoring of the network. The choice comes down to the availability / uptime requirements for the network and the preferences of the WISP doing the installation and maintenance of the property.

Wired Network Infrastructure

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.

Deployment Considerations

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 5. 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. *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.
- 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,

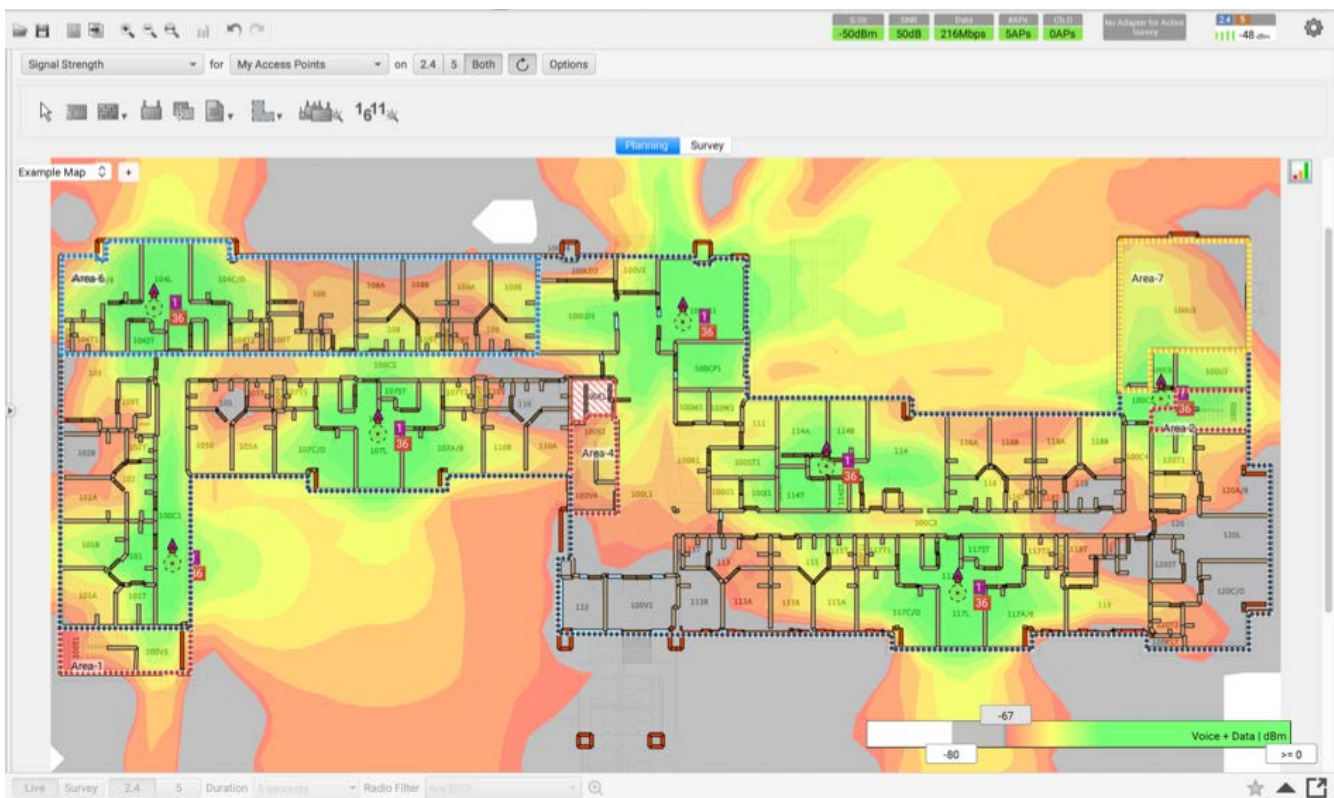


Figure 5: Ekahau Site Survey Pro. ⁴

⁴ <https://www.ekahau.com/products/ekahau-site-survey/overview/>

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.


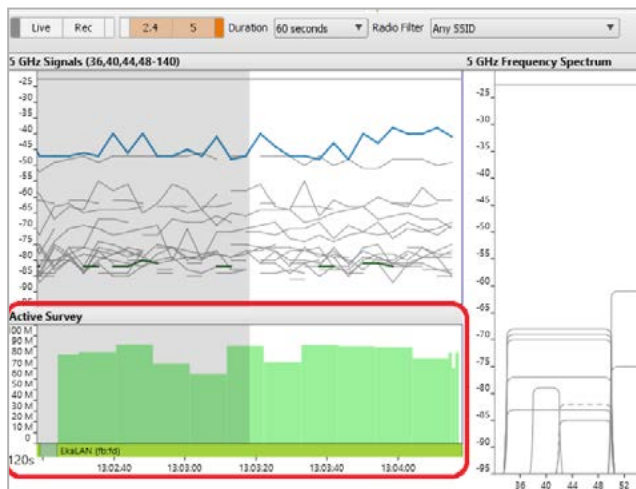


Figure 6: Ekahau Site Survey for passive and active site surveys.⁵

Conclusion

The design of a high-performance Wi-Fi network for hospitality applications 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. It is critical to understand the requirements and constraints of the hospitality project, as not all projects are created equal, and what may be appropriate in some environments may be impractical or too cost-prohibitive in others.

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.

⁵ <https://www.ekahau.com/products/ekahau-site-survey/overview/>



Find out more

www.ekahau.com

ekahau
WIRELESS DESIGN