Capacity Planning and High Density

September 4, 2014

This chapter discusses planning a network while keeping capacity and application requirements in mind. Today's WLAN needs are heavily dependent on mobile devices and applications. Capacity planning involves looking at application needs and designing a network around them, while High Density networks may be required where far too many clients are expected to connect in a location.

Typically the most frequent question that a network designer faces is, "How many Access Points do I need?" While RF planning tools like Ekahau may attempt to answer that question theoretically, and one can assume that the recommendation of the tools are within range of 15-20% error rate, it is still a very useful exercise to plan for capacity planning before rather than after deployment.

Access Point Density

The RF network can support these applications based on the expected accuracy (delta between estimated and actual location) and currency (time between location estimates) required. The accuracy is a direct function of the AP density (and AP height) while the currency is a function of the AP density and client type. The types of applications generally supported given an approximate AP density are shown in Table 18-1.

Table 18-1 CMX Venue and Density Required

Application	Venue types	Density (sq. ft/AP)	Accuracy	Currency
Presence	Mall, airport, etc	10+K	9-18m	N/A
Proximity	Retail, etc	<2.5K	5-8m	~30s
Asset-tracking	Enterprise, mall, etc.	5K	7-12m	>> 1min
Mobility-tracking	Mall, airport, etc.	<2.5K	5-8m	~30s

Consider an area of 500 feet x 500 feet floor of a mall space = 250,000 sq. ft.

From above, for a mobility tracking a density of 2.5K square feet per AP can be assumed for location accuracy of 5-8m and location currency of ~30 seconds.

So total number of APs required, roughly, would be = 250,000/2500 = 100 Access Points.

Note that if the only requirement were to track presence of a client in a CMX, but not engage the client in any application activity, then the AP requirements would change to:

Presence requirements = 10k / AP with accuracy of 9-18m

Mall area = 250,000 qft

Number of APs required = 250,000/10000 = 25

It is recommended that for CMX deployment, the WLAN network be designed with mobility tracking in mind with data and voice as an additional requirement. The CMX solution not only enables presence analytics, but also provides ways of providing guest access and engaging the user. Hence it is highly recommended that designers plan for voice, data, and location instead of presence. While 2,500 sq. ft./ AP is generally considered a good recommendation that covers a wide range of use cases, it is by no means a fit for every situation. Different RF considerations, type of materials, and applications provided need to be considered before planning for a deployment.

High Density Deployment

While a detailed High Density Deployment discussion is beyond the scope of this validated design guide, here are a few guidelines when preparing for a high density deployment. A high density CMX deployment is a network that has may have a higher number of Wi-Fi clients than an a traditional cell may have. High-density is defined as any environment with a large concentration of users, such as a classroom, lecture hall, or auditorium where the users are connected wirelessly, sharing applications, and using other network services individually.

High-density WLAN design refers to any environment where client devices will be positioned in densities greater than coverage expectations of a normal enterprise deployment, in this case a traditional, carpeted office. For reference, a typical office environment has indoor propagation characteristics for signal attenuation. User density is the critical factor in the design. Aggregate available bandwidth is delivered per radio cell and the number of users and their connection characteristics (such as speed, duty cycle, radio type, band, signal, and SNR) occupying that cell determines the overall bandwidth available per user.

A typical office environment, as shown in Figure 18-1, may have APs deployed for 2,500 to 5,000 square feet with a signal of -67 decibels in millowatts (dBm) coverage and a maximum of 20 to 30 users per cell. That is a density of one user every 120 square foot (sq. ft.) and yields a minimum signal of -67 dBm at the edge of the cell.

Figure 18-1 Typical Office WLAN

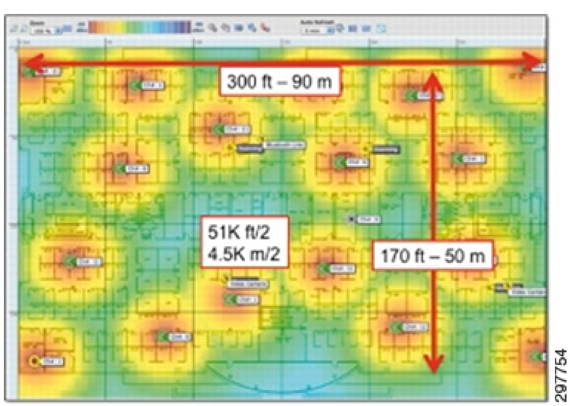

In planning and deploying such a WLAN, an AP is typically placed in an area expected to have a higher user density, such as in a conference room, while common areas are left with less coverage. In this way, pre-planning for high-density areas is anticipated. Conference rooms are often placed in clusters, so it is best to design for the maximum capacity of the area. For example, if maximum occupancy for the three rooms is 32, user density would be one user per 28 square feet, as shown in Figure 18-2.

Figure 18-2 Calculate User Density

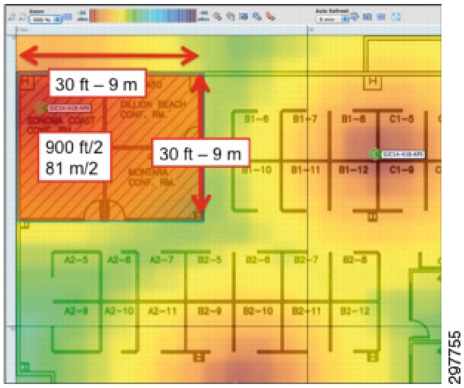

In a high-density environment such as a lecture hall or auditorium, the densities of users in the occupied space increase dramatically. User seating is typically clustered very close together to achieve high occupancy. The overall dimensions of the space are really only useful for getting an idea of the free space path loss of the AP signal. User densities are not evenly distributed over the entire space as aisle ways, stages, and podiums represent a percentage of space which is relatively unoccupied. The RF dynamics of the AP are very different from those experienced at the user level. The APs are exposed with an excellent view of the room and the user devices will be packed closely together with attenuating bodies surrounding them.

The single biggest sources of interference in the room are the client devices themselves. For each user sitting in the auditorium who can rest their hand comfortably on the back of the seat in front of them, the distance is approximately three feet, with an average seat width of 24 inches. This yields what is defined as a high-density environment, with less than 1 square meter per device deployed, assuming one or more devices connected per seat.

Figure 18-3 Seating and Interference

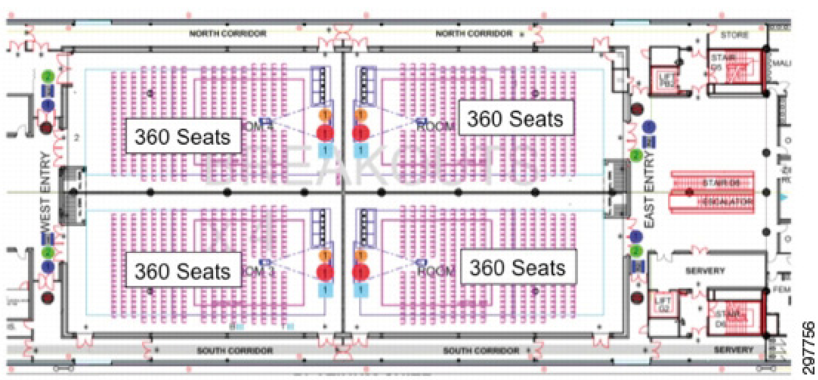

What is ultimately going to affect the client devices more than any other factor is the degradation of signal-to-noise ratio (SNR) through both co-channel and adjacent channel interference driven by co-located devices. Proper system engineering can minimize the impact by maximizing proper spatial reuse, but it cannot be eliminated in highly dense environments entirely. Operating margins become more critical as space is condensed and a bad radio or behavior in the mix can have a large impact within a cell. Client behavior under these conditions will vary widely and trends based on environment and event type have also been reported. There is not much that can be done about the particular client mix or behavior. The design goal is to engineer the network side as robustly as possible and to control and understand all variables.

Within environments that qualify as high-density, there are also submodels built by use case. For example, in a high-density environment such as a public venue or stadium, capacity is planned based on what percentage of users are likely to be active on the network at any one time. In higher education there is a different model, where casual WLAN activity is one use case while activity when a professor is lecturing may increase dramatically, up to 100 percent.

The WLAN design process can begin in many ways, but generally it begins with an expressed desire to provide connections to a specific area where a number of users will participate in a focused activity. To evaluate what is possible, it is first necessary to understand what is required as well as what is possible. There is generally a primary application that is driving the need for connectivity. Understanding the throughput requirements for this application and for other activities that will take place on the network will provide the designer with a per-user bandwidth goal. Multiplying this number by the number of expected connections yields the aggregate bandwidth that will be required.

The required per connection bandwidth is used to drive subsequent design decisions.

Establish and Validate a Per-Connection Bandwidth Requirement

How much bandwidth does each user require on average? In Table 18-2, the nominal throughput requirements for several popular applications and use cases in a higher education setting are shown.

Table 18-2 Bandwidth Requirements per Application

Application by Use Case	Nominal Throughput
Web—Casual	500 kilobits per second (Kbps)
Web—Instructional	1 Megabit per second (Mbps)
Audio—Casual	100 Kbps
Audio—Instructional	1 Mbps
On-demand or Streaming Video—Casual	1 Mbps
On-demand or Streaming Video—Instructional	2-4 Mbps
Printing	1 Mbps
File Sharing—Casual	1 Mbps
File Sharing—Instructional	2-8 Mbps
Online Testing	2-4 Mbps
Device Backups	10-50 Mbps

In all cases, it is highly advisable to test the target application and validate its actual bandwidth requirements. Software designers are often required to pick just one average number to represent the application's requirements when there are actually many modes and deployment decisions that can make up a more accurate number. It is also important to validate applications on a representative sample of the devices that are to be supported in the WLAN. Additionally, not all browsers and operating systems enjoy the same efficiencies, and an application that runs fine in 100 kilobits per second (Kbps) on a Windows laptop with Microsoft Internet Explorer or Firefox, may require more bandwidth when being viewed on a smart phone or tablet with an embedded browser and operating system.

Once the required bandwidth throughput per connection and application is known, this number can be used to determine the aggregate bandwidth required in the WLAN coverage area. To arrive at this number, multiply the minimum acceptable bandwidth by the number of connections expected in the WLAN coverage area. This yields the target bandwidth needed for the need series of steps.

Calculate the Aggregate Throughput Required for the Coverage Area

In a WLAN, a channel's speed is affected by multiple factors including protocols, environmental conditions, and operating band of the adapter. Before calculating aggregate throughput, several things must be considered.

In the aggregate throughput calculation, the connections instead of the seats were used as the basis for calculation. The number of connections in a cell is what determines the total throughput that will be realized per connection instead of the number of seats. Most users today carry both a primary computing device (such as a smartphone, tablet computer, or laptop) as well as a second device (such as a smartphone). Each connection operating in the high-density WLAN consumes air time and network resources and will therefore be part of the aggregate bandwidth calculation. An increase in numbers of device connections is one of the primary reasons older WLAN designs are reaching oversubscription today.

Users and applications also tend to be bursty (a measure of the unevenness or variations in the traffic flow) in nature and often access layer networks are designed with a 20:1 oversubscription to account for these variances. Application and end user anticipated usage patterns must be determined and also accounted for. Some applications, such as streaming multicast video, drive this oversubscription ratio down while others may drive this factor even higher to determine an acceptable SLA for each cell's designed capacity.

For 802.11 wireless networks or any radio network in general, air is the medium of propagation. While there have been many advances in efficiency, it is not possible to logically limit the physical broadcast and collision domain of an RF signal or separate its spectrum footprint from other radios operating in the same spectrum. For that reason, Wi-Fi uses a band plan that breaks up the available spectrums into a group of non-overlapping channels. A channel represents a cell. Using the analogy of Ethernet, a cell represents a single contiguous collision domain.

How many users can access an AP comfortably? Hundreds. But the question should not be how many users can successfully associate to an AP but how many users can be packed into a room and still obtain per-user bandwidth throughput that is acceptable. The question revolves back to around what is the expectation of a CMX deployment.

802.11 and Scalability—How Much Bandwidth Will a Cell Provide?

To scale 802.11 networks to reliably deliver consistent bandwidth to a large number of users in close proximity, it is important to examine certain WLAN fundamentals under reasonably ideal conditions. Once the rules are understood, the ways to manipulate them to maximum advantage are presented.

In real WLANs, the actual application throughput is what matters to the end user and this differs from the signaling speed. Data rates represent the rate at which data packets are carried over the medium. Packets contain a certain amount of overhead that is required to address and control the packets. The application throughput is carried as payload data within that overhead. Table 18-3 shows average application throughput by protocol under good RF conditions. Account for and manage all energy within the operating spectrum to ensure all of it is available for use

Table 18-3 Average Application Throughput by Protocol

Protocol	Throughput (Mbps)
802.11b	7.2
802.11b/g mix	13
802.11g	25
802.11a	25
802.11n—HT20 one spatial stream (1ss) Modulation Coding Scheme 7 (MCS7)	25
802.11n—HT20 2ss Modulation Coding Scheme 15 (MCS15)	70
802.11n—HT40 2ss Modulation Coding Scheme 15 (MCS15)	160
802.11ac—HT80, 1ss Modulation Coding Scheme 8, Short Guard Interval (SGI)	390
802.11ac—HTC80, 2ss Modulation Coding Scheme 8, Short Guard Interval (SGI)	867
802.11ac—HTC80, 3ss Modulation Coding Scheme 8, Short Guard Interval (SGI)	1.3 Gbps

The discussion until now has centered on a use case where every client in the room will be competing for bandwidth simultaneously. This is the case when the users in the room simultaneously access a resource on queue. However there are many instances where the design requirement is to offer access to resources or the Internet for casual use at an event or within a venue such as a sports arena. Planning and sizing for these types of events can be quite different and is based on expected Client Duty Cycle.

At a sporting event, for example, there are certain areas that require ubiquitous and instant access during the entire event. Ticketing, vendor sales, staff, and press areas generally require the highest amount of access. Of these, the press area is the only one that requires a high level of capacity in the arena itself. For the fans attending the event, only a percentage will be active on the WLAN at any one time. From experience we see a 20 to 30 percent take rate with some well-defined peaks occurring during period breaks. During play, very few fans are accessing the WLAN. However this is changing as applications such as video replay, instant stats, and concession orders from the seat become more commonplace.

Observation and understanding of the requirements of WLAN users and situational requirements guide the development of reasonable design goals. 500 users in a room who require simultaneous access to a single resource is a different design challenge than 1,000 or 1,500 users who only occasionally use the wireless network. Also, be aware that user patterns can and do change with time. This has been seen with the increase in the number of network clients per user. Monitoring network access and keeping good statistics allows wireless engineers to stay on top of user trends on the university campus. Good management platforms such as Cisco Prime Infrastructure are essential for managing the resulting network in real time and monitoring trends in a proactive manner.

Other High Density Considerations

- Choose a high minimum data rate to support increased efficiency, lower duty cycle, and reduce the effective size of the resulting cell.

- Plan for Wi-Fi co-channel interference.
- 5 GHz support is critical for high-density, so determine the channel plan that you will support and how it will be administered.
- Determine the number of channels and cells needed.

Note While an extensive discussion of high density deployment is beyond the scope of this validated design guide, designers interested in deploying the CMX solution in a high-density environment are highly encouraged to read the Cisco High Density Design Guide for additional design recommendations: http://www.cisco.com/c/en/us/products/collateral/wireless/aironet-1250-series/design_guide_c07-693245.html#wp9001186.