High Capacity WLAN Requirements Gathering

The process of requirements gathering should be the first step in designing and implementing a high-capacity WLAN. Through the process of requirements gathering, you can identify network design characteristics, allowing the Wi-Fi network to align with business objectives and meet desired performance levels.

Through this process, identify the following variables:

- Client device types, quantities, and capabilities
- Applications and throughput requirements

Once you have gathered a list of these requirements and reviewed them, you can forecast the required AP capacity for use during network planning and design. You can find a requirements gathering worksheet at the end of this document to aid in this process.

Client Devices

The first step is to identify the client device types that will be supported within the environment, their quantities, and their wireless radio capabilities. The wireless network must serve all client devices simultaneously. Since Wi-Fi leverages a shared RF medium, the network design and configuration must balance requirements between all clients. This often requires a "lowest common denominator" approach to network design and implementation to accommodate the needs of all clients. Understanding the trade-offs and network design options first requires knowledge of the client devices that the Wi-Fi network will support.

Identify specific device models if possible so that you can gather detailed wireless radio characteristics and integrate them into the Wi-Fi network design and configuration. However, in some circumstances, such as public events, you cannot know the specific device models before the event. In such cases, identify generic device categories that are likely to be used and integrate the wireless radio capabilities for such devices into the network design.

Gather the following wireless radio capabilities for client devices:

- **Wi-Fi Radio Type**

  Identify the type of client radio and which Wi-Fi Alliance® certifications[1] it has passed (for example, 802.11b/g/a/n). This will aid in your network design by determining the data rates supported, application throughput capabilities, and if 802.11b support is required on the network. Legacy 802.11a/b/g clients and single-stream 802.11n (1x1:1) clients typically require greater signal strength from the AP to maintain high data rates compared to multiple-stream 802.11n clients. The lack of MIMO (Multiple Input, Multiple Output) capability generally means a lack of advanced digital signal processing techniques, such as maximal ratio combining (MRC), when receiving transmissions from APs and can result in higher error rates and frame loss when signal strength is weak. Understanding these differences will aid in network design and ensure proper coverage and performance for all clients on the network.

  A high-density Wi-Fi network typically disables 802.11b data rates to improve overall network capacity and performance. Frames transmitted at the low DSSS and HR/DSSS data rates consume significantly more airtime than frames transmitted at higher OFDM data rates used by 802.11g and 802.11n. Therefore, elimination of the DSSS and HR/DSSS data rates used by 802.11b clients will increase the overall capacity of the channel by reducing airtime utilization. Elimination of 802.11b data rates also reduces the need for protection mechanisms such as RTS/CTS and CTS-to-Self, which can reduce throughput for all clients on the channel by up to 40%. However, if the network must support a single 802.11b client, then all APs and clients must accommodate this requirement and allow use of at least one 802.11b data rate. The trade-offs associated with balancing the performance and capacity needs of a high-speed network with the need to maintain support for previous generations of low-speed clients are common on Wi-Fi networks. While Wi-Fi standards have continually evolved to provide increasing speed and capabilities, a focus on backwards-compatibility has been central to its increasing adoption, ubiquity, investment protection for consumers, and overall success.

[1] Wi-Fi Certified Products (Wi-Fi Alliance).

It is highly recommended to support only 802.11n clients on high-density Wi-Fi networks whenever possible due to efficiencies in protocol operation that reduce network overhead (discussed in the "Sources of Wi-Fi Network Overhead" section). Also configure a fairly high minimum data rate in most high-density deployments (for example, 18 Mbps) to encourage clients to roam more aggressively to APs with sufficiently strong signal strength. Doing so helps clients use higher data rates, which conserves airtime and increases overall channel and network capacity.

- **Channel Support**

  Identify the wireless channels supported by client devices and design the Wi-Fi network with a channel reuse plan that provides proper coverage and capacity. It is critical that the network design includes only those channels that are supported by the entire client population to avoid creating *dead spots*. Dead spots are the result of a network design in which a client device has no supported channel in a given physical area. Among clients that support dual bands, 5 GHz DFS channels (UNII-2 and UNII-2 Extended) are the least likely to be supported.

  DFS channels are 52-64 and 100-144 in the United States, Canada, Europe, Japan, Australia, and New Zealand.

  Some mobile devices and low-end laptops operate in only a single frequency band to reduce price and to increase battery life. This reduction in functionality means that manufacturers opt for trade-offs in equipment design, such as a single-band 2.4 GHz Wi-Fi chipset. While this allows full functionality within most home environments, it can present significant challenges in multiple-AP enterprise environments where capacity in the 2.4 GHz frequency band is often constrained. As integrated dual-band Wi-Fi chipsets commoditize to nearly equivalent price levels as single-band chipsets and battery performance improvements are realized with the future release of 802.11ac (5 GHz), mobile devices will eventually offer better integration into high-capacity environments. Until that time, you must design Wi-Fi networks with these constraints in mind. It is important to identify the percentage of the client population that only operates in the 2.4 GHz frequency band (single-band devices) and determine the desired client ratio between 2.4 GHz and 5 GHz frequency bands. Then use this information to perform proper network capacity planning and configure client band steering to provide the optimal balance of clients between frequency bands.

  Finally, determine if clients support 40-MHz-wide channels. Wider channels facilitate greater peak throughput but reduce available spectrum for reuse by colocated or adjacent Wi-Fi cells. Many mobile clients and consumer devices do not support 40-MHz channels.

  You can research channel support through local regulatory certification databases, such as the FCC[2], Industry Canada[3], and European Union CE Mark. Alternatively, if you have a particular device model, you can test for channel support by capturing association requests from the device with a wireless sniffer or Aerohive AP operating in promiscuous mode, and then using a protocol analyzer to inspect the *Supported Channels* information element. This information element lists the channels supported by the device. Alternatively, you can attempt to connect it to an AP operating on various channels with various channel widths and verify both the connectivity and data rates that it supports.

- **Transmit Power Output**

  Identify the transmit power output capabilities of client devices for integration into the Wi-Fi network coverage design. Then design your Wi-Fi network to provide high-quality, bidirectional communication between clients and access points with low error-rates, low frame-loss and symmetrical signal quality.

  Poor transceiver capability (low output power, low antenna gain, low receive sensitivity) in some mobile devices requires careful infrastructure planning and high-grade infrastructure hardware to avoid creating asymmetric link quality, often referred to as the "Unbalanced Power Effect." When access points transmit at substantially higher power output than client devices, clients may discover and associate to APs due to a sufficiently high receive signal strength. However, the AP might not be able to receive such a high quality signal due to the lower transmit power output capability of the client. In these situations, it is common to see good

[2] FCC ID Search Form (Federal Communications Commission, Office of Engineering and Technology).

[3] Radio Equipment List (Industry Canada, Certification and Engineering Bureau).

downstream link quality (AP-to-client) with poor upstream link quality (client-to-AP), usually with high error rates and frame loss on the uplink due to poor mobile device transceivers, poor receive sensitivity on the APs, or poor AP antenna gain. High quality enterprise-grade access point hardware with improved *Receive Sensitivity* can help compensate for low client transmit power, but it cannot eliminate this effect completely. In a mixed client environment, the infrastructure must balance the needs of all clients on the network. Therefore, access points should operate at similar power output levels as client devices, to facilitate proper AP selection and association by client devices for high-quality bidirectional communication.

You can research client power output through manufacturer-supplied documentation, local regulatory certification databases as referenced in the previous “Channel Support” section, or by contacting the device manufacturer directly.

Additionally, the IEEE 802.11k amendment includes provisions for radio resource management and dynamic power level negotiation between clients and access points. The forthcoming Voice-Enterprise certification by the Wi-Fi Alliance® includes radio resource management (RRM) capabilities from the 802.11k amendment. However, support for this capability has not been adopted by manufacturers to-date.

- **Maximum Wi-Fi Data Rate**

  Once you identify the radio type (802.11a/b/g/n), including the supported number of 802.11n spatial streams and channel width, determine the maximum Wi-Fi data rate that the device can achieve. Data rates are determined by a number of variables, including modulation, guard interval, coding ratio, number of spatial streams, and channel width. For high-density Wi-Fi environments, a 20-MHz channel width and 800-ns guard interval (GI) should typically be assumed.

  The maximum data rate is important for determining how fast a client can transmit and receive data, which affects the amount of airtime it consumes. In properly designed high-density Wi-Fi networks, clients should have strong signal strength and a high signal-to-noise ratio (SNR), which allows them to use high data rates a larger percentage of the time.

- **Maximum TCP/IP Throughput**

  Due to various sources of network overhead, the maximum Wi-Fi data rate does not represent the actual application throughput. To estimate the maximum amount of TCP/IP throughput that a client can achieve, you must first determine the amount of network overhead either through live network testing under load or through an educated assumption.

  Live network testing is the best and most accurate method of determining the actual throughput capabilities of a device. Whenever possible, you should perform testing with the same client devices that will connect to the network. The actual network equipment and configuration settings that will be deployed in production should be used, and the anticipated network load (number of clients, application traffic) should be emulated to account for network overhead from medium contention.

  However, if you cannot test with the actual devices before deployment, then an educated assumption will still aid in network planning. Historically, Wi-Fi professionals have found, through real-world measurements, that good estimates for network overhead when using TCP-based applications are around 60% for legacy 802.11a/b/g networks and can be as low as 40% for 802.11n networks. You should use a value within this range, with 40% being an aggressive performance estimate and 60% being more conservative. To estimate the maximum throughput for each device, multiply the maximum Wi-Fi data rate by the inverse of the network overhead. For example, an 802.11n single-stream client with a 65 Mbps maximum data rate and a network overhead assumption of 60% yields a maximum TCP throughput estimate of approximately 26 Mbps: 65 x (100% - 60%).

  For UDP and multicast-based applications, always try to perform live network testing to determine the maximum achievable throughput for each client device. However, a reasonable estimate is typically around 40% network overhead, with higher achievable throughput than TCP-based applications. For example, an 802.11n single-stream client with a maximum data rate of 65 Mbps and a network overhead assumption of 40% yields a maximum UDP throughput estimate of around 40 Mbps: 65 x (100% - 40%).

- **Device Quantity**

  Determine the total quantity of each device type that will be in the environment. If you know this and the application throughput requirements for each device, you can estimate the peak network load and forecast the required aggregate AP capacity to support all client devices concurrently. List the quantity for every individual device model or category separately because differences in client device capabilities affect throughput and capacity estimates.

Common client device categories and capabilities:

Device Category	Wi-Fi Radio Type	Channel Support*	Channel Width	Transmit Power Output	Maximum Data Rate 20MHz/40MHz
Feature Phones	802.11b/g	1-11	20 MHz Only	11dBm	54 Mbps
Smart Phones	802.11n 1x1:1	1-11	20 MHz Only	11 dBm	65-72 Mbps
Low-End Tablets	802.11n 1x1:1	1-11, 36-48, 149-161	20 MHz Only	11-14 dBm	65-72 Mbps
High-End Tablets	802.11n 1x1:1	1-11, 36-48, 52-64, 100-140, 149-161	20/40 MHz	11-14 dBm	65-72 Mbps / 135-150 Mbps
Netbooks	802.11n 1x2:2	1-11, 36-48, 149-161	20/40 MHz	11-17 dBm	72 Mbps (Up) 144/300 Mbps (Down)
Low-End Laptops	802.11n 2x2:2	1-11	20 MHz Only	17-20 dBm	144 Mbps
Mid-Range Laptops	802.11n 2x3:2	1-11, 36-48, 149-161	20/40 MHz	17-20 dBm	144/300 Mbps
High-End Laptops	802.11n 3x3:3	1-11, 36-48, 52-64, 100-140, 149-161	20/40 MHz	17-20 dBm	216/450 Mbps
VoIP Handsets	802.11a/b/g	1-11, 36-48, 149-161	20 MHz Only	11-16 dBm	54 Mbps

This chart references channels from the FCC regulatory domain only.

Knowing the client device inventory and their capabilities provides the following input into subsequent Wi-Fi network design and configuration:

- **802.11b Support** – Based on the presence of 802.11b-only client devices
- **Wi-Fi Channel Set** – Based on common channels that all client devices support
- **40-MHz Channel Width Support** – Based on client support of wide channels, decision on DFS band support, and the channel reuse plan (discussed in the "High-Density Wi-Fi Network Design" section)
- **Percentage of Single-Band Client Devices** – For use in network capacity planning
- **AP Wi-Fi Transmit Power Output** – Based on client power output capabilities

Applications

Once you have identified the client devices and their radio capabilities, you must next determine the applications that will be used and their performance requirements (throughput and QoS). A thorough understanding of the application requirements will provide the necessary information to design a successful WLAN that meets organizational needs. Through this process, you can identify the critical and non-critical applications and establish a target application throughput level for each device that is required for successful operation.

The following table lists common application classes, typical bandwidth requirements, and recommended QoS classes. These QoS classes are based on common network best practices but should be tailored to each environment.

Common application classes, throughput requirements, and recommended QoS classes:

Application Class	Required Throughput	QoS Class (Layer 2 / Layer 3)[4]
Web-browsing/email	500 Kbps–1 Mbps	WMM 0 (BE)/DSCP 0
Video Conferencing (example: WebEx[5])	384 Kbps–1 Mbps	WMM 5 (VI)/DSCP AF41 (34)
SD video streaming (example: Netflix[6], Hulu[7])	1–1.5 Mbps	WMM 4 (VI)/DSCP CS4 (32)
HD video streaming (example: Netflix, Hulu)	2–5 Mbps	WMM 4 (VI)/DSCP CS4 (32)
Apple TV[8] streaming	2.5–8 Mbps	WMM 4 (VI)/DSCP CS4 (32)
Apple FaceTime[9]	900 Kbps	WMM 5 (VI)/DSCP AF41 (34)
YouTube[10] video streaming	500 Kbps	WMM 0 (BE)/DSCP 0
Printing	1 Mbps	WMM 0 (BE)/DSCP 0
File Sharing[11]	5 Mbps	WMM 0 (BE)/DSCP 0
E-Learning and Online Testing	2–4 Mbps	WMM 4 (VI)/DSCP CS4 (32)
Thin-client (example: RDP[12], XenDesktop[13])	85–150 Kbps	WMM 4 (VI)/DSCP CS4 (32)
Thin-client (with video or printing)	600–1,800 Kbps	WMM 4 (VI)/DSCP CS4 (32)
Thin-apps (example: XenApp[14])	20 Kbps (per App)	WMM 4 (VI)/DSCP CS4 (32)
Device Backups (example: cloud services)	Uses available bandwidth	WMM 1(BK)/DSCP CS1 (8)
VoIP Call Signaling (example: SIP)	5 Kbps	WMM 3 (BE)/DSCP CS3 (24)
VoIP Call Stream (codec[15] dependent)	27–93 Kbps	WMM 6 (VO)/DSCP EF (46)

These application classes are provided as a general reference only. You should research specific application requirements from the product developer.

[4] The QoS Baseline At-A-Glance (Cisco Systems, Inc., 2005)
[5] WebEx Network Bandwidth White Paper (Cisco Systems, Inc., 2010)
[6] Hunt, Neil, Encoding for streaming (Netflix, 2008)
[7] Hulu Plus System Requirements (Hulu, 2012)
[8] Apple TV (2nd and 3rd Generation): Troubleshooting playback performance (Apple, Inc., 2012)
[9] FaceTime for Mac: About HD video calling (Apple, Inc., 2011)
[10] YouTube Help, Watching Videos, System Requirements (Google, Inc., 2012)
[11] Zimmerman, Tim, *Best Practices for WLAN Site Surveys That Save Money* (Gartner, Inc., 2012)
[12] Remote Desktop Protocol Performance Improvements in Windows Server 2008 R2 (Microsoft, 2010)
[13] XenDesktop Planning Guide: User Bandwidth Requirements (Citrix Systems, Inc., 2010)
[14] ICA Client Bandwidth Analysis (Citrix Systems, Inc., 2001)
[15] Enterprise QoS Solution Reference Network Design (Cisco Systems, Inc., 2005)

Give special consideration to video applications due to their typical requirements for high bandwidth and low latency. Video applications should use a Constant Bit-Rate (CBR) codec to control the performance and user experience, which ensures that the first few clients do not consume a large amount of bandwidth if using a Variable Bit-Rate (VBR) video codec. In addition, consider using multicast transmission for HD video streams of greater than 2 Mbps to increase client capacity in a high-density setting where multiple clients will be receiving the stream. However, multicast transmission can be less reliable than unicast transmission because clients do not acknowledge the data they receive and the server does not retransmit data if it is corrupted in transit. You can enable multicast-to-unicast conversion to improve reliability, but doing so will reduce the number of clients that APs can support concurrently. Be sure to make note of any applications used internally that are based on UDP or multicast, as client performance and throughput capabilities are typically substantially different due to the lower overhead that these protocols use.

Knowing the application requirements provides the following input into subsequent Wi-Fi network design and configuration:

- **Target Application Throughput Levels (per Device)** – Based on critical applications on each device type and application throughput requirements
- **Client SLA Target Throughput** – Based on target application throughput levels; applied to different user groups
- **Per-User Rate Limits** – Based on maximum application throughput requirements plus a recommended additional 20% to accommodate temporary application bursts

Forecasting AP Capacity

Once the client device and application requirements have been identified, you can forecast the required access point radio capacity (and subsequently estimate the number of APs) to aid in Wi-Fi network design. The forecast is derived by estimating the network load (based on airtime) required for each client device to achieve its targeted application throughput level. This is represented as a percentage of airtime that will be consumed per-device on an AP radio, with the aggregate sum being used to forecast the required AP capacity to support all clients concurrently.

First, determine how much airtime the target application throughput level will consume on each device type by dividing the required throughput by the maximum TCP or UDP throughput the device can achieve. For example, a 1 Mbps TCP video application running on an Apple iPad mini that can achieve a maximum 30-Mbps TCP throughput on a 20 MHz channel yields an airtime consumption of 3.33% per device on a particular channel. This means that every iPad mini device running the video application requires 3.33% of the capacity of a single access point radio (assuming no outside noise or interference). Perform this calculation separately for every device type that will be supported in the environment. If you are using 40 MHz channels in the 5 GHz frequency band, perform this calculation separately for both frequency bands with dual-band capable clients. This is required since the airtime consumed on a 40 MHz channel for the same application throughput level is less than a 20 MHz channel.

You can also use the airtime consumption value to estimate of the number of devices that a single AP radio can support (when all devices are of the same type). This can be useful in environments where the client device population is controlled by the organization and client distribution is well known (for example, a classroom environment where all students use the same type of device during the lesson plan). To estimate the number of devices supported on a single AP radio, divide the individual device airtime consumption value into 80%. An individual Wi-Fi channel saturates around 80% of airtime utilization. Therefore, the total capacity of a single AP radio is 80% ÷ the airtime required per device. For example, a single 2.4 GHz AP radio operating at 20 MHz channel width serving only Apple iPad mini devices running a 1-Mbps TCP video application could support a maximum of 24 devices concurrently (80/3.33).

Second, multiply the quantity of devices for each device type by the required airtime for each device to determine the total airtime required for those devices. If you are using 40 MHz channels in the 5 GHz frequency band then you will need to split dual-band capable clients between both frequency bands to accurately determine their airtime consumption. Once the total airtime is determined, divide it by 80% to determine the number of AP radios required to support those devices in the environment. For example, if there are 75 total Apple iPad minis on 2.4 GHz and each one consumes 3.33% airtime, then a total of 3.12 AP radios are required to support all 75 devices concurrently ((75 * 3.33%) ÷ 80%). Perform this calculation individually for every client device type that you expect will be in the environment.

Finally, add the number of AP radios needed to support each client device type to determine the total number of AP radios required in the environment; round up if necessary. Adjust the number of radios if only a percentage of clients will be connected to the WLAN concurrently. If all client devices will be online concurrently, then no adjustment is necessary. Since most deployments use dual-radio APs, divide the adjusted number of AP radios by two to forecast the number of dual-radio access points needed to support the identified devices and applications. If a fractional number results, round up.

It is important to understand that the forecasted AP capacity is only an estimate; the final capacity will likely be slightly different. The forecast is useful as an initial starting point for the Wi-Fi network design. The site survey process verifies your AP capacity forecast and will likely need to be slightly modified due to facility characteristics that require higher capacity in specific physical areas by co-locating APs or coverage in less-common areas such as hallways, stairwells, bathrooms, and elevators.

Through the process of requirements gathering, you are now armed with adequate information about the client, application, and network infrastructure to begin the planning and designing the network deployment.

High Capacity WLAN Requirements Gathering Worksheet

Coverage Area: ____

Step 1 – Client Devices

Identify the client devices that will be used on the network and their wireless radio capabilities, as described in the **Requirements Gathering » Client Devices** section.

Network Overhead Estimate: [ ] 40% (aggressive) to 60% (conservative, recommended)

Max Throughput = Max Wi-Fi Data Rate X (1 – Network Overhead Assumption); Or actual device measurements

Device Model / Category	Wi-Fi Radio	Supported Channels and Width	Transmit Power	Max Data Rate	Max Throughput	Device Quantity
Example: Apple iPad mini	*802.11n, 1x1:1*	*2.4 GHz, 20 MHz 5 GHz, 40 MHz DFS Supported*	*14 dBm*	*65 Mbps, 135 Mbps*	*30 Mbps (20 MHz) 60 Mbps (40 MHz)*	*250*

Input into Wi-Fi Network Design and Configuration

802.11b Support Required	(Circle One)	Yes / No
Wi-Fi Channel Set	(Circle Multiple)	1-11, 12-14, 36-48, 52-64, 100-140, 149-161, 165
40-MHz Channel Width Support	(Circle One)	Yes / No
Target AP Wi-Fi Power Output	(dBm)	

Step 2 – Applications

Determine the applications (or application classes) that each client device identified in Step 1 must support, and the target application throughput level required by those applications, as described in the **Requirements Gathering » Applications** section.

An individual Wi-Fi channel saturates around 80% of airtime utilization. Therefore, the total capacity of a single AP radio is 80% ÷ the airtime required per device.

Device Model/ Category	Frequency Band	Application Class	Target Application Throughput	Max Device Throughput	Airtime Req. per Device (Target ÷ Max)	Max Concurrent Devices per Radio (80% ÷ Airtime)
Apple iPad mini	*2.4 GHz*	*SD Video*	*1 Mbps TCP*	*30 Mbps TCP*	*3.33 %*	*24*
Apple iPad mini	*5 GHz*	*SD Video*	*1 Mbps TCP*	*60 Mbps TCP*	*1.66 %*	*48*

Step 3 – Forecast AP Capacity

Forecast the required AP capacity as described in the **Requirements Gathering » Forecast AP Capacity** section. First, multiply the number of client devices for each device type (identified in Step 1) by the airtime that each client device requires (identified in Step 2). The result of that calculation determines the total airtime required for that group of devices. Then divide the total airtime by 80% to determine the number of AP radios needed to support this device group. Then add the results from all client device types together and multiply by the expected percentage of client devices that will be online concurrently during peak load. From these calculations, you can derive the adjusted number of AP radios required. If you use dual-band APs, divide the adjusted number of radios by 2 to forecast the number of dual-radio access points required to support the identified clients and applications.

5 GHz Band Steering Ratio: ______ (Dependent on RF network design; typically 60-80% on 5 GHz)

Device Model/Category	Frequency Band	Device Quantity (Total X Band Steering Ratio)	Airtime Required per Device	Total Airtime (Quantity X Airtime)	# AP Radios Required (Total Airtime ÷ 80%)
Apple iPad mini	*2.4 GHz*	*75 (30% ratio)*	*3.33 %*	*249.75 %*	*3.12*
Apple iPad mini	*5 GHz*	*175 (70% ratio)*	*1.66 %*	*290.50 %*	*3.63*

Sum # AP Radios:	**Total # AP Radios**	
Multiplied by:	**% Concurrent Clients**	
Equals:	**Adjusted # AP Radios**	
Divided by 2 equals:	**AP Capacity Forecast (Round Up)**	