Why 'One Access Point Per Classroom' Approach is Wrong

A White Paper on Wireless LAN Design

April 23rd, 2014

Keith R. Parsons
Managing Director
Wireless LAN Professionals, Inc.
Wireless LAN Training/Consulting/Design
383 Ridge Road, Orem UT 84057
@KeithRParsons
keith@WLANPros.com
http://WLANPros.com

wirelessLAN PROFESSIONALS

Why ‘One Access Point Per Classroom’ Approach is Wrong

Preface

First, let me say if after doing a going through a proper WLAN design process you end up with a series of classrooms, each with a single Access Point – more power to you! That is not what I am referring to in this article.

I am concerned with the trend in our industry of promoting the quick and easy methodology of One Access Point Per Classroom – and without any real design processes building out a wireless networks with literally one Access Point in each classroom. Compounding these bad habits is the practice of not even doing a post-installation validation survey to find how much co-channel interference actually exists, or what other problems might have been created.

I know many people with Wireless LAN design experience despise the use of blanket statements as real-world methodologies. Other examples of this are things like “One Access Point for every 2,000 square feet of floor space”, or a ”put an Access Point every 18 meters in a grid pattern”, or have a “10% Cell Overlap”. That is exactly what I am attempting to agree with. This “One Access Point Per Classroom” is a blanket statement, and that approach to defining/designing Wireless LANs is wrong.

I am suggesting to always complete proper design processes that include defining requirements, designing to meet those requirements and then validating that the installed solution meets all the requirements.

Promoting a One Access Point Per Classroom “design” is either lazy, ignorant, and/or greedy. The “1 for 1” is a marketing campaign meant to sell Access Points, and is **not** a Wireless LAN design methodology.

K-12 1:1 Confusion

Many in the Wireless LAN industry have taken advantage of the confusion between the up and coming 1:1 initiative within K-12 schools. The premise is based on planning for a minimum of one device per pupil. Many schools are moving towards increasing electronic-based education methods and their infrastructure needs to support at least 1 device per 1 pupil counts.

Additionally, most schools have seen a massive increase in the Bring Your Own Device phenomenon. Even primary schools frequently have the youngest children bringing smartphones or tablets to school.

> *Sidebar: At my grandchildren's elementary school in rural middle Georgia the school is moving to a BYOD model where pupils are strongly encouraged to bring their devices to class for use in the education process. In fact, their school officials sent a letter a few months ago to all parents listing what types of devices would be supported, hinting that these would make great Christmas presents.*

At one school district, I was specifically asked by their school board to design their wireless LAN infrastructure to meet a 5:1 ratio. Their expectation was that the Wireless LAN would someday need to support up to five devices per person. Though I think this is a bit of stretch, we have seen in many of our districts a greater than 2:1 device count today. But even these estimates need to take into account how many devices will actively be used at once, not just carried in a backpack or pocket. But it does widen the discussion of proper Wireless LAN design issues.

> *Note: There are big differences between a device associated and connected via an Access Point, and a device actively sending data across the RF frequency. One is merely associated, the other sharing a load on the channel along with all other Access Points and client radios on the same channel.*

So we know the problem; there are going to be lots and lots of devices using Wi-Fi to access the network in our K-12 schools. They want to prepare for this onslaught, and send out RFPs or contact their VARs, all to get on top of this impending load on their networks.

Sometimes even those writing the RFP's don't understand RF or 802.11 and put this One Access Point Per Classroom requirement, or higher, in their requests. Because other school districts have done it, the buzz-language around 1:1 is easy to plop down in the RFP, especially when you have no frame of reference on RF and 802.11 fundamentals. It is our job as Wireless LAN Professionals to design correctly, even if it includes needing to educate our customers.

Enterprise WLAN Design Processes

Normally – in most enterprise WLAN installations we follow a proscribed process of:

Define – collect information from the customer and fix on a specific set of requirements to design to. These include number of devices, area to be covered, estimated device count growth over time, density, applications in use as best as can be gleaned, etc. See Andrew Von Nagy's presentation from #WLPC – Wireless LAN Professionals Conference for a great process used in definition stage in the links below.

Video Slides

Design – using in-depth knowledge of RF fundamentals, antenna principles, and how the 802.11 protocol works, use both on-site measurements and predictive modeling tools, some still use Access Point-on-a-Stick methods, to arrive at a draft design. But either way, design and measure to make sure the number of Access Point's, power settings, antenna choices, and Access Point placement work together to meet all the design requirements from the definition stage.

> *Note: There exists weeks of training classes to understand these concepts and to learn how to design wireless LANs properly – this is not something to be taken lightly. Never has the term "garbage in-garbage out been more applicable.*

Install – often the easiest of the WLAN infrastructure steps. Have cabling teams place Category 6 or better cable to each Access Point location in accordance with wiring standards and local codes, certify the new cabling runs, confirm switch port configurations, mount Access Points, and test wired-side connectivity.

Validate – this is the critical phase whereby we prove the design actually did meet the design requirements through a site survey, and followed up with active monitoring of network performance during school hours. Think of it as the operational check any technical system should have before handover to the customer. This is a step many don't do – thus just relying on the assumptions made in the design phase and never really knowing how the installed Access Points react with each other and with their surroundings.

Remediate – if any of the Access Point placements, channel or power settings are sub-optimal, or if you need to add/remove Access Points, or to turn off radios. This is the phase where we fix whatever was wrong, so the WLAN does meet all the design requirements.

Yet, somehow in the world of K-12, many if not all of these steps are reduced to simple formulaic processes. One of the more popular techniques is to mistake the 1:1 initiative of one device per pupil, and replace it with One Access Point Per Classroom. This is a very simple 'design' process often used to charge full-scope prices for partial-quality work.

> *Note: One Access Point Per Classroom methodology is NOT a design process at all, but a bastardized shortcut people sometime use to skip all the phases above.*

This simple technique is just that; easy, quick, and understandable without any knowledge or understanding of how 802.11 protocol works, or how RF propagates.

That is the downfall of this simple methodology, and it's greatest weakness – yet at the same time leads to why so many integrators continue to use this defective approach, and why it is frequently over-marketed.

> *Sidebar: Some have noted the proper Wireless LAN design process is expensive and if you use the One Access Point Per Classroom approach save that money. But instead, they spend too much extra money on Access Points, installation costs, cabling costs,*

and switch port costs. These far exceed any design costs on all but the very smallest one-room schoolhouse.

Sidebar: *Others have commented and predicted putting One Access Point Per Classroom will 'future-proof' the school district. This totally ignores the best practices of the Wireless LAN industry of first defining exactly what the needs are today, and in the future before specifying the design requirements. A properly designed Wireless LAN will also meet the future needs of a school, but without needlessly overspending on Access Points and their associated costs.*

Sidebar: *Still others who support this One Access Point Per Classroom approach mention how many 'satisfied customers' they have as some sort of proof the customers didn't over-spend on their Wireless LAN installation. Those school districts probably never did a post-installation validation survey to document whether they had CCI/CCC – and I doubt they had any design parameters more than coverage and client counts.*

RF and 802.11 Fundamentals

Let's return to some of the important basics built in to the 802.11 Wireless LAN protocols for a minute. All devices using Wi-Fi follow this simple set of rules on how they can communicate together. All devices first listen on its specific defined frequency to be sure the channel is clear and that no one is using the frequency at that exact time. If the frequency is clear and empty for a defined period of time, usually something like 9 millionths of a second, then the device grabs a random number and waits another defined set of wait times. When their counter reaches zero, and if the frequency is still clear, the client or Access Point will transmit.

OK, that was a bit much for a single paragraph. But the idea is all Wi-Fi devices on the same frequency have to work together in order to share the available bandwidth. Access Points, laptops, tablets, smartphones, all devices share the frequency.

The easiest of all the design requirements to meet in our WLAN design is Coverage. That is delivering of RF from the Access Point to the client devices. If we want more coverage just turn up the power in the Access Point, and/or add more Access Points. Voilà, you have more coverage. But coverage is rarely of concern in an otherwise properly designed Wireless LAN.

The most difficult of all design requirements to meet is the direct opposite of this; how to control coverage so we can re-use the limited frequencies we have. This is sometimes referred to as minimizing co-channel interference. Again, Andrew Von Nagy does a great job in his posts and re-defines this to "co-channel contention".

Whether we call it co-channel interference CCI, or co-channel contention CCC – it refers to the same thing. How do we design our wireless networks so we can reuse the most scarce resource of all, the RF frequency! – We can always buy more Access Points, and many use this technique whenever there is any issue, they just purchase and install more Access Point's. When in fact they are just exacerbating the issue of CCI/CCC out of desperation and lack of understanding of RF fundamentals. Furthermore, some manufacturers and VARs will shortcut design work with the promise of free Access Points post-installation to make up for any coverage holes, thus further compounding the issue. They must think it easier to do it wrong first, then make up for it later with free product!

The problem with CCI/CCC is that when more clients and Access Points share the same frequency, they share the frequency's inherent capacity as well. Each device essentially gests a smaller piece of the capacity pie available on that frequency. What's more, as client device numbers increase, the negative effects of medium contention increase – namely retransmissions – further degrading the performance and capacity of that frequency and all radios using it.

> *Note: We've learned from professional services experts from Cisco, Aruba, Ruckus, Extreme, and others who have entered the Large Public Venue space the way to get more capacity is to find ways to get frequency reuse. Using the building materials themselves to help block RF. Yes, I said that right. We want to* ***block*** *RF from going where we don't want it to go.*

Marcus Burton wrote a great white paper on the subject of 802.11 protocols and the contention process.

Designing for High Capacity

In order to have a higher capacity Wi-Fi network we need two things:

First – the ability to re-use the same frequency over and over in the building and targeted areas. To get that, we need to design using antenna choices, power selections, and Access Point placement to get control of RF and to focus it where we want. This is one reason why hallways should be a last resort for Access Point placement – they direct RF to places where CCI/CCC is increased.

Second – we need to insure each client device that accesses the Access Point does so at as high a data rate as possible. In order to get more devices to share the same frequency, we need each device to get on and off the frequency as quickly as possible. Slow data rates, and high retry rates will allow fewer devices to use the frequency.

Why One Access Point Per Classroom Approach is Seductive

Back to the problem with using a One-Access Point-Per-Classroom as a starting point.

Yes – I agree it is a simple concept, easy to understand and communicate. Easy to build a Bill of Materials. Easy to install. Easy to sell.

The above alluring description makes One Access Point Per Classroom seductive for both sellers and buyers. For the sellers they get to sell more Access Points. They don't have to explain anything to the school districts. They don't have to worry about actually doing any definition. They don't have to do any design. They don't have to do anything but produce a simple Bill of Materials.

From the end customer's standpoint, they also have it easy. They don't have to think about what they really need. They just count how many classrooms, find a cabling contractor, and there you have it.

The most seductive part of this process is that post install it actually works! This is NOT because it is a good design. It is an artifact of the inherent resilience built into the 802.11 protocols. This is also a problem in many enterprise WLANs. The system gets installed and in the beginning the Wireless LAN works fine. Then as additional loads are placed on it, the wireless network starts to fail. Many of these early adopters of the One Access Point Per Classroom opted for this approach before they had the 1:1 and BYOD clients actually on their campuses. So of course a system with too many Access Points will look good – well at least until the CCI/CCC kicks in and the problems start to rear their ugly heads.

Some even tout having thousands of satisfied customers who went with this simplified One Access Point Per Classroom approach. My reaction here is based on returning to visit many of these same customers who when they actually start getting the client device counts up approaching 1:1 or higher with BYOD the network starts getting slower, much slower. The answer? Remove Access Points – then the network works better.

This is a hold out from other industries like hospitals that frequently placed Access Points primarily in hallways and had too many Access Points. By removing Access Points and lowering the Co-Channel Interference/Contention the throughput increases.

The question is: Did schools waste money by buying more Access Point's than were needed?

If a system you spent too much money on works in the short term, is it still a good financial decision?

Why One Access Point Per Classroom Approach is Wrong

One Access Point Per Classroom methodology is just plain wrong.

It is wrong because it doesn't follow professional processes of Define, Design, Install Validate, Remediate. It skips all the tried and true methods to insure good Wi-Fi for the sake of being easy.

Here is where one problem lies. ***Frequency Reuse***.

In the 2.4GHz frequencies in North America we only have three non-overlapping channels. Channels 1, 6 and 11. If there are two Access Points on the same channel and they can hear each other transmit at a signal level above a defined threshold, then they will share the assigned frequency. So will all the devices associated to them. This is not a guess – this is defined in the 802.11 protocol and hard-coded in both the Access Point's and in all Wi-Fi devices.

So with only three channels it is very difficult to put Access Point's close together and not have them see each other.

One option is to turn the power down so each Access Point doesn't transmit as far, thus shrinking the CCI/CCC domain size. The issue with turning down the power is it may lower the data rates clients use in their communication when not implemented carefully. Thus allowing fewer devices to use the limited frequency. *Higher power means higher SNR, higher SNR means higher data rates and lower retry rates*. This is a good thing – well until the higher power increases the CCI/CCC domain size and causes even lower throughput by more devices all sharing the same limited bandwidth of the frequency.

So in placing one Access Point per classroom, you are almost guaranteeing Co-Channel Contention in the 2.4GHz range. When more devices and Access Point's share the same channel, lower throughput results.

Others react and say, "We will just let the vendor's automated radio management turn the power down on 2.4GHz radios". Empirical evidence of this working properly is hard to come by. In doing validation surveys in hundreds of schools I personally have yet to see this work. Others will then just turn off up to two thirds of all Access Point's 2.4GHz radios, or turn them into monitor mode. But the Access Point itself had capital cost, installation cost, cabling cost, and switch-port cost, and then you turn off some portion of over half of them?

What about 5GHz?

What about 5GHz? There are many more channels available to us in the 5GHz range – up to 22 available 20MHz channels. That is true in many instances. Yet many client devices in use in K-12 environment are not capable of UNII-2/UNII-2e or ISM channels and thus we end up with 8 usable channels of UNII-1 and UNII-3 to choose from.

To repeat – you need to be very sure all of your 5GHz capable devices are able to access UNII-2 and UNII-2e frequencies. If they can't and you have your WLAN using them, these devices will see coverage holes. This can be prevented by limiting your 5GHz channel choices to just UNII-1 and UNII-3.

> Note: *As an industry we need to reinforce the purchase of only devices that support all UNII 5GHz bands. There are negligibly more expensive dual-band UNII-2/2e supported solutions that should be a priority with all purchasing department plans.*

Having eight channels is great! Far better than the mere three channels in 2.4GHz and it does give us a much better paradigm of designing out CCI/CCC from our wireless LANs.

So doing a 5GHz-only design using the One-Access Point-Per-Classroom might be workable, if you plan on not supporting 2.4GHz devices. The other option is to continue to support 2.4GHz devices, but turn off 2.4GHz radios and use as many Band-Steering options to move 5GHz capable devices off the 2.4GHz frequencies. There are a few vendors that support single radio Access Points that could be used to lower the capital expenditure of having Access Points with two radios, but only using one.

> Note: *One very simple method using human nature is to have two separate SSIDs, perhaps 'School' for the 2.4GHz radios, and 'School-Fast' for the 5GHz only radios. Thus allowing natural selection to take place. Those devices that are not capable of*

using 5GHz won't even see the fast SSID, and those who can the humans controlling which Access Point to join will naturally want the fast SSID.

Sadly, even this logic is flawed. It goes back to the Define stage. Do you really **need** to have one Access Point per classroom? Is it part of your design requirements? Why would you want to purchase more infrastructure equipment than is necessary?

Designing to Meet Requirements

I doubt a school would want to design a cafeteria with a commercial kitchen and seating areas for double or triple the predicted size of the student body. That would just be a waste of taxpayer money. No school district would over-design for other infrastructure needs. Can you see building two football stadiums for a single High School? So why would a school district want to purchase and install twice as many Access Points as they need for their 1:1 initiative?

Sure, we need to design for future growth in device counts and network throughput. But remember, the defining criteria in most high density deployments isn't based on Access Point counts, but on how well the WLAN designers can achieve frequency reuse. When the frequency if full, no more traffic can flow no matter how many Access Points you throw at the problem.

We have found after going through the entire wireless LAN process in over 1,500 classrooms in schools from New Jersey to Oregon, and many states in between; we have never needed more than one Access Point per two classrooms in order to meet a very stringent 5GHz coverage goal.

Again, and this deserves repeating; after doing full Define, Design, Install, Validate and Remediate on thousands of classrooms we have NEVER seen a need for more than One Access Point for every Two classrooms! These were normal K-12 classrooms seating from 20-40 students each.

Our shorthand requirement for providing Wi-Fi for a school with a 1:1 initiative, and this includes at least one BYOD per pupil in addition to the one provided by the school district. This shortcut is to design for a minimum of two Access Point's coverage in the 5GHz range at least -65dBm. This guarantees at least four radios are covering every single area of the school at a high RSSI. Two radios in 2.4GHz and two radios in 5GHz.

Below we have included complete list of the other requirements and metrics we design to. But the shorthand version above has proved quite effective as a bellwether metric reflecting the success of other specific requirements.

Association vs Throughput

There are two numbers when looking at Access Point client capacity. The first is if the Access Point has the ability to have enough supported client device associations. These are for the devices that automatically want to join an Access Point, but perhaps don't actually send a lot of data. Think of someone's smartphone in their pocket during class. It is associated, and has an Association Identifier (AID), and is in the Access Point's association table. The Access Point will pick off a packet from the wired network and bridge it over to the smartphone. But it is not doing anything active at the moment.

> *Note: The 802.11 protocol has a max associations per radio of 256 hard coded in the stack. Some manufacturers – Cisco in particular have an even lower maximum limitation of only 200 per radio when using multicast frames.*

The second are those devices that actively transmit and receive data over the Access Point's assigned frequency. They are in the Co-Channel Contention domain and have to wait and 'play nice' with all other devices on the same frequency. These devices should be connected and transmit data as quickly as possible – high data rates and low retry rates are what we are after.

Most enterprise-class Access Points today can handle hundreds of associations. So that shouldn't be an issue in any K-12 design. On the other hand, we are concerned with the number of client devices that are actively transmitting or receiving. This is the concern most have when planning for a 1:1 initiative. They make an incorrect assumption that all student devices will simultaneously be accessing network resources from adjacent classrooms. Remember, network timing is usually measured in millisecond or smaller increments, so "at the same time" frequently needs explaining to the uninitiated.

Predicting Needs

In the telecommunications industry we use some methods of predicting customer behaviors to assist in preparing for anticipated loads. One technique is called an Erlang function. A quick example is a phone company planning on how many lines to bring into an apartment complex. We know that all apartment dwellers won't be accessing the land-line at the same time. So using some averages of number of minutes per day and during peak times, we can anticipate a number of trunk lines to run from the central office to the apartment complex – perhaps as low as 5%-10%. This works great, until the first morning of Mother's Day when everyone wants to call home at the same time. The phone industry doesn't design for Mother's Day – they instead design for a normal day, and install a 'All Lines Are Busy' recording for Mother's Day.

> *Note: You can learn more of the Erlang function and formulas here.*

We can plan around the same thing in schools. Very rarely will adjacent classrooms be truly simultaneously accessing network resources with all possible student devices. This is an anomaly – we don't need to build a network with all BYOD and all student devices in the entire school simultaneously able to connect. We need to design for a reasonable subset that can work at the same time.

Providing Wi-Fi For 1:1 Initiative Needs

In our 'Short Cut' design – we provide for four-radio coverage in all areas. With current Access Points this means we can have 500+ devices associated in any area, way above anything in a K-12 as far as association needs.

This level also gives the ability to have aggregated throughput of four radios in any one area – the actual results here are highly dependent on the devices, applications, and data rates used, but a typical analysis shows a minimum of 100Mb of aggregate throughput in any one area.

We have personally worked on systems where we have more than 180 smartphone devices streaming video from a single 802.11 radio without any problems. But this was using multicast frames. The total frequency utilization was under 40%. Retry rates were <2%. So it is possible to have many devices all streaming video at the same time – it just needs to be the same shared video stream. Obviously this is dependent on having an environment that can support multicast.

As an additional example, at one of our schools they were conducting teacher training on Chromebooks and thought all 150 of them were being load balanced over a large 12-radio array. Later while checking for the session's total throughput noticed the large array had been turned off, and all 150 worked the entire day on a single Access Point.

On the other hand – during our Wi-Fi Stress Test of last year we found we could bring any single Access Point to its knees by streaming multiple unique high-definition videos to no more than 20 tablets per Access Point. Once the frequency was congested, no more data could be sent.

The critical factor is frequency utilization. Some refer to this as airtime utilization. Application types have a huge difference in actual load over the airwaves.

Designing to Meet All Requirements

Another flaw in the One Access Point Per Classroom philosophy is not resultant coverage. There is fantastic coverage with that many Access Points. The actual flaw is never designing for, nor validating against, the co-channel interference.

This is actually fairly easy to measure. Someone qualified to interpret the results should accomplish it at each and every installation.

Conduct a post-validation survey of the school, including a passive survey of all Access Points on all channels. Then in the survey analysis using your survey tool of choice plot everywhere where you have more than two devices at >-80dBm on the same channel. This by definition is Co-Channel Contention. Any Access Point that can see another radio on the same channel will have to defer transmissions – basically they will share the frequency. Thus cutting their throughput in half.

> *Note: Some use -85dBm for the trigger point – but it is really based on the Clear Channel Assessment threshold in the devices. Some Access Point manufacturers allow system integrators to lower this threshold, basically lowering the Access Point's receive sensitivity allowing them only listen to stronger signals.*

It is quite easy to measure and show in a One-Access Point-per-Classroom environment the CCC is quite large – and thus predict the reduction in potential throughput because of too high of Access Point density.

Conclusion

When designing a Wireless LAN for a K-12 1:1 program, do not use a One AP Per Classroom approach. Actually do the work and design the Wi-Fi network properly – considering **all** the design requirements. Not just revert to the 'easy way out'. There is a right way and a wrong way to design Wireless LANs – using the One AP Per Classroom approach is the wrong way!

If your solution provider is making expert design services cost prohibitive, they are doing the school district an injustice. Design should be part of the "Value Add" in Value Added Resellers. There are plenty of skilled Wireless LAN Professionals available in the marketplace. Not using one and resorting to a flawed approach is being disingenuous to their customers.

Review of Design Requirements

Here is a short list of some of the requirements we design for, and measure against when designing Wireless LANs.

First Access Point Coverage
Typically something like −65dBm or better.

Second Access Point Coverage
Typically −70dBm or better (Cisco IP phones want −67dBm). Some people call this '*overlap*' but it's not measured in percent of area, but in dBm for the 'backup' coverage. You can't measure the 'overlap' in percentage, but you can measure 'backup coverage' in dBm.

Signal to Noise Ratio
Typically 25dB or higher in all areas to be supporting VoIP over Wireless LAN

Co Channel Interference
Any radio – but specifically Access Points stronger than −85dBm on same channel equals Co-Channel Interference, or Co-Channel Contention. The area where two or more Access Points would defer to each other, thus sharing the available bandwidth on the frequency.

Channel Utilization
Try to keep under 50% utilization.

Client to Access Point Density
Usually something like 20:1 or 30:1 for heavy laptop users, or for some VoWiFi handsets 7:1. *Check with your Voice handset vendor for their actual recommendations if you are doing a voice deployment*. You actually need to do an Erlang analysis to find how many 'average' minutes of call per hour for this calculation. You need to be able to measure and verify there is appropriate number of Access Points for your client device loads.

Additional Density
Areas where there might be an abnormally high concentration of devices. (Auditoriums, cafeterias, etc.) You need to be able to measure and verify there is

appropriate number of Access Points for the clients in that area. These specific areas need to be highlighted and called out on the floor plans during the definition stage of the WLAN design process.

Data Throughput

First would test for 'gross' data rate, like 1, 2, 5.5, 6, 11, 12, up to 54Mb/Sec. More importantly, check for NET throughput with iPerf to meet design standards. You can check data rates **to** the Access Point (wireless rates) by sending packets to the Access Point. You can check data rates **through** the Access Point by placing the iPerf server directly behind the Access Point. Or check data rates **through** the Network by placing iPerf server back near your core where your servers are.

> *Note: We could enter an entire discussion here of the benefits of turning off low data rates – but we'll leave that for another time.*

Jitter

Less than 5 msec is important for Voice over IP deployments. Check for variation of latency between packets. Measure both upstream and downstream jitter during a two-way conversation call in progress.

Latency

Less than 50 msec – total one-way end-to-end delay. Not to be confused with Round-Trip Time RTT.

Packet Loss

VoIP networks should support less than 1% packet loss.

Access Point Handoff

Typically between 50-100 msec. Try testing on an 'Open' SSID first, then test again with your Authentication system in place. Consider roaming assistance protocols such as 802.11k/v.

Codec

Use a voice codec that has the highest potential MOS score. Don't try to use compressed codecs, we have lots of bandwidth, it's a quality issue with Wi-Fi. But if you are traversing a WAN link, you'll have to re-evaluate your choice of codecs. Sometimes you might not have any control over the codecs being implemented.

Roaming
Data usually has a 'Portability' need for coverage. Portability – loose session, but not IP upon roaming – Like sleeping your laptop as you move to another room. Voice has a 'Mobility' need for coverage. Mobility – can't loose session at all during any roam.

Authentication
Open is the fastest for supporting Voice over Wi-Fi, but has no security WPA-PSK still requires a key exchange upon roaming. WPA-RADIUS might require a full authentication cycle clear to RADIUS Server. Some vendors support key caching or fast secure roaming. Consider roaming assistance protocols such as CCKM or 802.11r.

Beacon Interval
Normally set to 100 time units (approximately10th of a second). This might better fit into the "don't touch this" category. But some vendor installations might need you to confirm this.

DTIM Interval
Check with your handset vendor on a voice deployment. Usually set to 2.

Coverage Areas
Usually designs for Data include areas where laptops go (carpeted or vinyl floors). Voice designs also need elevators, stairwells, bathrooms, parking structures, etc.

Collision Domains
You can't change the Physics of Wi-Fi, it will always be a Shared Medium - Design accordingly

Quality of Service
Choose your Priority Queues appropriately; realize with Wi-Fi we don't have an Absolute Priority, only a statistical advantage; Data **will** sometimes get in front of Voice traffic. Make sure **all** your Access Points and Clients can support the QoS system you choose.

You might want to look to your application vendor to see what they recommend for QoS settings, some applications might work fine without any QOs tuning.

Error Rates
Set parameters and design goals for Retry Rates and CRC rates on your Wireless Network. Measure and be sure to verify before starting the Voice clients on the Wireless Network.

Transport Delays
Understand your data traffic flows. Pay special attention to any potential delays if traffic is tunneled through a controller (if installed) and determine the appropriate deployment method to suit your needs.

End to End Quality of Service
Check to see your QoS set at the client carries those priority bits clear through your network to the receiving client

Protection Mode
When designing around a b/g mixed mode environment the Protection Mode adds overhead to every packet exchange. Remove of all 802.11b clients on your network, as well as within range of your network – if possible. Set expectations accordingly if you must support mixed b/g networks—at least 1⁄2 data rate loss minimum plus additional packet overhead.

Power Management
Use the most efficient Power Mode supported by your clients and Access Points. Some of the older versions were inefficient, especially with respect to Voice frames. Not relevant with most laptops and tablets today, but of high concern with VoIP handsets that need to work an entire shift without a recharge.

Data Rates
Check with your vendor for recommendations for your specific handsets if deploying Voice over IP over Wi-Fi. Some vendors require 1 Mbs data rates be turned on, others recommend 1, 2, 5.5, 6, 9, and 11 be disabled and others still recommend turning off 54, 48, and 36.

SSIDs & VLANs

> ***Note: Remember, the fewer the SSIDs the better!***

When configuring a unique SSID for Voice, remember adding a VLAN to Wi-Fi does **not** break up any broadcast domain, you are still in a collision domain and broadcast domain with all other devices on same frequency within hearing range. Confirm Voice SSID is on correct VLAN. Confirm Voice VLAN is in correct Subnet with correct DHCP scope.

RRM Tuning Thresholds

Look for specific recommended values to set your Auto-RF/RRM/ARM tuning parameters.

Administrative Security

No administration of the Wireless network from the Wireless network. Manage all Access Points and controllers in their own private space. Perhaps look at on-boarding solutions.

Note: Not overspending on too many Access Points might free up some capital to spend on better back-end support for onboarding and authentication support services.

Note: There are many tools available to help measure and report on these metrics. Look to AirMagnet Wi-Fi Analyzer, WildPacket's Omnipeek, Metageek's Eye P.A., Tamosoft's CommView for Wi-Fi or even WireShark to gather and report these specific requirements.

Survey tools like AirMagnet Survey Pro, Ekahau Site Survey, and Tamosoft's TamoGraph can all be used to gather location specific information to analyze.

With any tool – first be very confident in your understanding of antennas, RF propagation, and 802.11 protocols – then of course learn well your tool's properties and techniques.