Whitepaper

7SIGNAL
THE WI-FI PERFORMANCE COMPANY

802.11ac MIGRATION: REAL WORLD BEST PRACTICES

Learn what vendors won’t tell you about 11ac performance in real-world deployments

Table of Contents

Introduction

802.11ac introduces several new characteristics to Wireless LANs and is a key technology in the evolution of faster, more capable networks. However, the protocol gives WLAN Network Administrators yet another set of variables to consider. Truth is, Wi-Fi networks had too many dials to control, even before 802.11ac Wave 1 / Wave 2 came along. The new options with 802.11ac amplify the complexity, leaving many WLAN Administrators dazed and vulnerable to false expectations propagated by vendor hype and side-stepping the performance reality check.

Even hard-core Wi-Fi aficionados find it difficult to explain how to optimize network capacity and client performance in real-world environments during the transition to 11ac. If you are looking for concrete answers about how best to execute an 11ac migration, and not more questions, read on. This paper attempts to cut through the double-speak, and expose what is conveniently left out by 11ac hype mongers. Intended as a reality reset, it also provides specific guidance in order to maximize client performance, network capacity and ROI.

The New 11ac Bells and Whistles

11ac has four new characteristics that contribute to better performance and capacity. They are:

1. 256-QAM encoding scheme
2. 80 MHz and 160 MHz channel bonding
3. Multiple spatial streams
4. Multi-User MIMO (MU-MIMO)

Now let's discuss how they help, and the conditions under which their value can be realized.

256-QAM encoding scheme

802.11n's top encoding scheme is called 64-QAM. 11ac employs a new scheme 256-QAM which packs 33% more data into the signal than 64-QAM. This results in two new data rates (MCS 8 and MCS 9) which achieve higher throughput than the previously highest data rate of MCS 7, in the same amount of airtime. However, the new encoding scheme is more sensitive to disruption by noise in the environment, which means you need a cleaner signal between the client and the AP. In practice, this means you need line of sight access between AP and client from a maximum range of about 20 feet to get the benefits of 256-QAM. In real-world settings, rate control will choose to use 256-QAM only in close proximity of the AP, like 10-20 ft., otherwise it will drop down to lower data rates.

At high data rates 11ac does not penetrate walls any better than 11n, so you are not gaining anything in improved coverage. In fact, in traditional office environments, the distance using 256-QAM is less than it is for 802.11n's top data rate at 5GHz and significantly lower than 802.11n at 2.4 GHz. Add 80 MHz channel bonding, and it only gets worse. An AP's maximum power is typically reduced at higher data rates to avoid non-linearity issues, which means coverage area shrinks dramatically, and receive sensitivity is also reduced. In buildings with heavy concrete walls or a lot of iron and steel, range is even worse and 256-QAM is virtually unusable beyond 15 feet! This is all due to higher signal-to-noise ratios required by these new data rates. When 256-QAM is not viable, 11ac radios automatically downshift to 64-QAM, matching what you already have with 11n. In this case, the bottom line is - no benefit.

The technical explanation for this relates to receive sensitivity. While the top data rate for a single stream of 11n on a 20 MHz channel only requires a receive sensitivity of -64 dBm, 11ac's top data rate using 256-QAM requires a receive sensitivity of -59 dBm. That's a jump of 5 dB. This means clients must use more Tx power to be heard by the 11ac radio in order to get that top data rate, all other things being equal. In practice, since clients don't adjust their power, the client must therefore be closer to the AP. In addition, even more than the APs, clients suffer from transmitter linearity issues with the highest 256-QAM. Clients are battery operated and therefore, compromises are generally needed. This means that when a client switches to 256-QAM, the client output power drops as well. So with 256-QAM, a higher signal level is needed for reception and less output power is available. In other words, beyond the 64-QAM, link budget is reduced 5 dB at the receiver and few dBs at transmitter. Hence 256-QAM is only attainable at relatively short range.

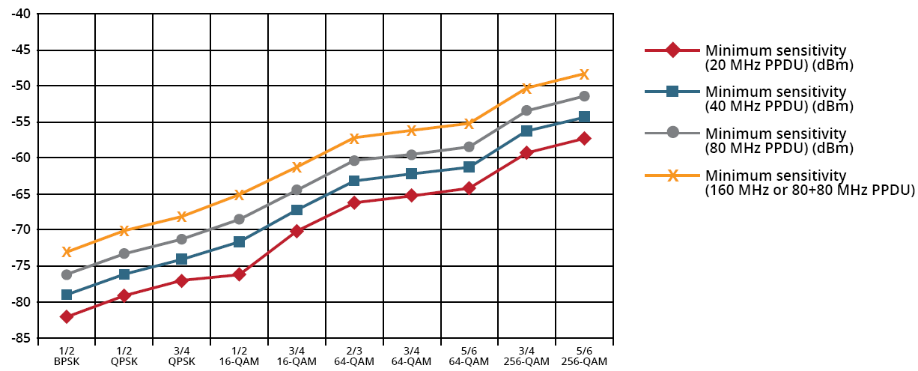

Fig 1: Receive Sensitivity for 11ac encoding schemes and channel widths (chart courtesy of Aruba)

So what does all this mean for 11ac AP placement? To get the benefits of 256-QAM, APs should be placed in open areas, where they can literally be seen by users. Cubicle office areas, hotel lobby areas and outdoors are the best use cases, while the worst use cases are warehouses and buildings with heavy concrete or steel construction.

Of course, only 11ac compatible clients, such as the iPhone 6, can utilize the 256-QAM encoding scheme when it is viable. It is of no value to 11n and other legacy clients except indirectly by relieving the contention for airtime (more on that later).

80 MHz and 160 MHz channels

802.11n was first to introduced the concept of channel bonding. It allowed two adjacent 20 MHz channels to be combined (bonded) into one 40 MHz channel. On the 5GHz band, there are plenty of contiguous non-overlapping 40 MHz channels available, so this made a lot of sense, as it is doubles the amount of data transmitted by each radio. 802.11ac expands on this technique, allowing four 20 MHz channels to be bonded into a single 80 MHz channel. 11ac Wave 2 also provides for 160 MHz contiguous channels and a non-contiguous 80 + 80 configuration.

Channel bonding is the single biggest performance multiplier, and it is the foundation for vendors' claims of 1.3 Gbps speeds for Wave 1, and from 2.3 Gbps for Wave 2 up to 6.7 Gbps when using all 8 spatial streams, should anyone have the nerve to implement them. But let's get real. Can you actually use 80 MHz and 160 MHz channels in your environment? Let's explore.

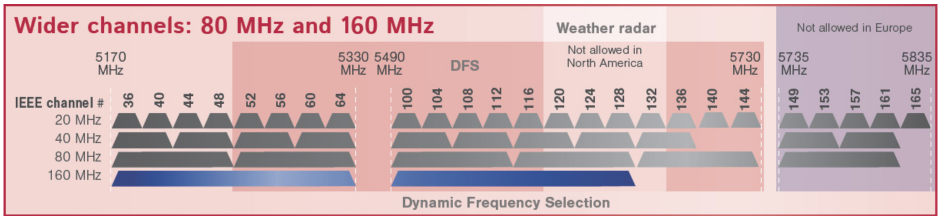

Fig 2: Available contiguous bonded channels (chart courtesy of Meru Networks)

We can discount 160 MHz right away. For most enterprise deployments this is useless. As you can see from the diagram above, in North America there is only one contiguous 160 MHz channel available, and in Europe only two. This renders 160 MHz only useful for point to point WAN links. But what about 80 MHz?

You know from the 2.4 GHz days that planning dense coverage is a bear, especially when you have limited channels. Close proximity of cells using the same channel causes performance degrading co-channel interference. In practice, this means that when a nearby cell is using the channel, it makes the channel busy for other nearby cells on the same channel. Additionally, nearby does not only mean neighboring cells, but due to the nature of the Wi-Fi channel access method CSMA/CS, it also means that cells at a 1-3 cell distance may keep the channel reserved. Channel state is determined based on frame headers, which are coded with a robust 6 Mbps rate (802.11b is even stronger with a 1 Mbps rate). This rate can be decoded at a great distance, much farther than the data fields using higher modulation. Thankfully, that problem is slightly less of a concern today, now that most devices support 5 GHz, and 2.4 GHz is simply considered a "lifeline" service for legacy devices.

5 Ghz Channel Widths	US		Europe	
	Incl. DFS	Excl. DFS	Incl. DFS	Excl. DFS
20 MHz	21	9	20	4
40 MHz	8	4	9	2
80 MHz	4	2	5	1
160 MHz	1	0	2	0

Fig 3: Channel availability by region in the 5 GHz band

However, using 80 MHz channels returns to that same problem. With only 4 or 5 channels available (depending on your region), it is difficult to plan a dense deployment of APs, because you must reuse channels in nearby cells. Without meticulous power management, you will create co-channel interference, which degrades performance and negates much of the expected gain from the bonding. Remember, the top data rates of MCS 8 and MCS 9 are much more susceptible to low signal to noise ratio. The same is also true for bonded channels. Whereas the top rate for a single stream of 11ac in a 20 MHz channel requires a receive sensitivity of -59 dBm, it requires a receive sensitivity of -53 dBm for an 80 MHz channel. That's an increase of 6 dB. As shown with the 64-QAM vs. 256-QAM comparison, once again this means clients must be closer to the AP. So, to get the benefit of 80 MHz bonded channels and 256-QAM, clients need to be within 10-15 feet from the radio.

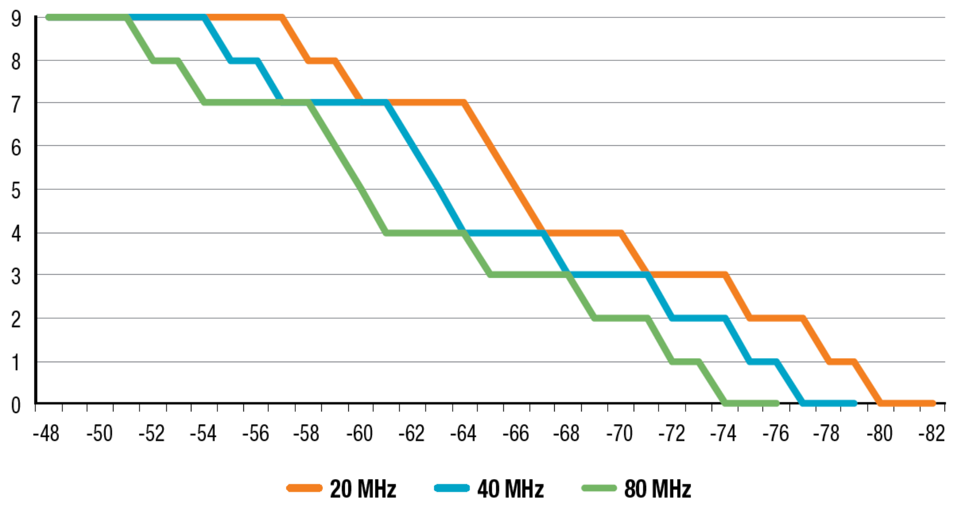

Fig 4: Receive Sensitivity for channel widths and MCS data rates (chart courtesy of Ruckus Networks)

How and where to use 80 MHz

Can you use 80 MHz channels everywhere? No! At home, with only few APs around, 80 MHz enables higher speeds. However in enterprises, one must be more selective about use cases. 80 MHz has tremendous value in two key non-residential applications: mesh backhaul and bridging (point to point and point to multipoint). In both cases you have AP talking to AP, and depending on distances, you can fully utilize the advanced features of 11ac. If you are using 11n for either of these use cases today, this is the first place to start. You can easily double performance of those links overnight. In client access however, the benefits of 80 MHz are more elusive.

For enterprise client access, one of the most practical uses for 80 MHz channels is in isolated high-density Service Areas such as a lecture hall, lobby area, cafeteria or conference center, where most clients have clear line of sight to the AP. Here, it is feasible to buffer the 80 MHz radio from the next major hotspot by surrounding it (including above and below) with radios on other access points using non-overlapping 40 MHz or 20 MHz channels. Bear in mind that a 20 MHz channel overlapping the 80 MHz channel will basically render the 80 MHz channel useless while a client is transmitting on that channel. But it is not strictly necessary to avoid overlapping channels. For example, one AP could be using an 80 MHz channel, and an adjacent AP could be using one of the 20 MHz channels within the 80 MHz range. The 11ac radio can recognize when the 20 MHz channel is in use, and holds off its 80 MHz transmission until it is clear. However, implementing it in a high density network is unwise. Why would you compromise your 80 MHz channel for transmissions on a 20 MHz channel?

Consequently, careful channel selection and maximum channel separation is crucial to reduce the air utilization and noise floor and prevent 80 MHz transmissions from being blocked by overlapping channel use on nearby APs. To the chagrin of some AP vendors, you must also disable automatic channel selection and interference avoidance schemes, because 11n APs may not be able to recognize an 11ac 80 MHz channel correctly resulting in a switch to a channel that overlaps 80 MHz, in order to avoid radar or microwave emissions.

The migration approach for isolated high density Service Areas described above is obviously easiest to accomplish in single story buildings. Compared to a previous configuration based on 20 MHz 11n, a deployment like this, adding or replacing only one 11ac AP in the center, and switching selected 11n AP radios to 40 MHz, could yield an average improvement in client performance of 20-30%, while also increasing aggregate capacity.

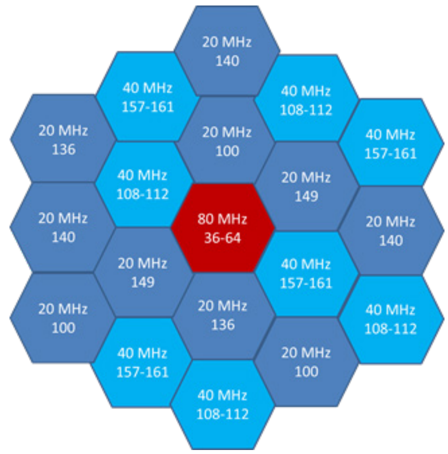

Fig 5: Example 80 MHz channel plan for an isolated high-density service area

Another good way to leverage 80 MHz channels is on the building perimeter, using directional antennas facing inward. Using 80 MHz cells at the building perimeter leaves you with a smaller RF footprint that needs to be buffered by other cells. So, this allows you to use 40 MHz cells more densely. However, bear in mind the line-of-sight requirements noted earlier for optimum performance of the 80 MHz cells. In both deployment scenarios, the more channel separation between cells you can accomplish the better, as this also eliminates adjacent channel interference, thereby lowering the noise floor, which means cleaner signals and a better chance of meeting the required minimum receive sensitivity. Consequently, more separation between the APs sharing the same channel, means less co-channel interference and more available airtime.

Excluding hotspot scenarios, 40 MHz is a generally the better option. With eight 40 MHz channels available in North America and nine available in Europe, 40 MHz is more practical. But this is not a leap over 11n, since 11n also supports 40 MHz bonding. Nevertheless, in most cases you should be looking to migrate your channel plan to 40 MHz, if you have not done so already. For high density deployments where you want high-performance everywhere, such as in offices, residence halls, hospitals and where there is no one central place where everyone aggregates, 40 MHz provides a greater ability to re-use channels which allows for more cells closer together. That means you can scale capacity and keep end user throughputs higher even during a high load. However, it also means that maximum throughputs in lightly loaded networks may be lower than with the wider channels.

How 80 MHz can work against you

While most 11n and 11ac clients support 40 MHz channels, not all 11ac clients support 80MHz channels, and no 11n clients can support it. So when an 11n client connects to an 11ac radio using 80 MHz channels, what happens? The radio drops down to a 40 MHz channel or 20MHz for that connection depending on the capabilities of the client. Obviously this means that the full benefit of 80 MHz can only be realized when all clients are 11ac 80 MHz capable clients. But more importantly it also means that during all communication with non-80 MHz clients, a 60 MHz or 40 MHz slice of channel capacity is going to waste. Unless you implement channel sharing (see cautions above), it cannot be used in another cell, as it is reserved for 80 MHz channel use. Therefore the merits of using 80 MHz are wholly dependent on your client mix, which is a moving target. In fact, if the mix is less than 30% 11ac / 70% 11n, then 80 MHz actually works against you, and you would have gained more capacity by utilizing two 40 MHz channels instead of one 80 MHz channel. This is another indication that in order to maximize capacity and performance as you migrate to 11ac, you must also recognize that your channel plan cannot remain static, it must evolve over time as your client mix shifts in favor of 11ac clients. An 11ac migration is not a set-it-and-forget-it proposition. Even after the APs are deployed, your channel plan will require frequent adjustment, based on your client mix, if you want to get the most from 11ac.

Luckily, 11ac has another feature different from channel sharing which could help you prevent this channel waste. It is especially useful in high density environments once your 11ac client population reaches 50% or more. Beacon frames are normally sent using a data rate that every client can handle, so they can decode the frames to learn the existence of the AP. However, it is possible to set beacon frames to be sent using 256-QAM. These would be unreadable by 11n and legacy clients, so those clients could never discover the AP exists. This action would force only 11ac capable clients to use that AP. This is quite an extreme configuration and it only works in a few scenarios. Devices sharing the channel still need to negotiate channel access and use of airtime between all nearby APs and terminals.

More spatial streams

The 802.11ac standard allows for up to 8 spatial streams. Using a 160 MHz channel and MCS 9 data rates a single radio has a maximum data rate of 6.7 Gbps – another big number being touted by 11ac promoters. Without a doubt, more spatial streams are a good thing. Like channel bonding, they are another linear performance multiplier, and you have probably already seen the results with 1 stream, 2 stream and 3 stream 11n access points.

More streams however, requires more antennas, so just consider what it is going to take to cram 8 Tx and 8 Rx antennas into an AP, and still ensure they perform as well as the antennas in a 3x3 design. Least we mention the increased power requirements. Powering such a beast is sure to exceed the specs for 802.3at, which enterprises are already reluctant to adopt.

From a power perspective, 802.3at defines two types of PSEs (power sourcing equipment):

- Type 1 (over two pairs) – The PSE can supply a maximum of 15.4 W over a voltage range of 44 to 57 VDC using Category 3 cabling or better. This type provides support for legacy installations.

- Type 2 – The PSE can supply 30W (over two pairs) or 60 W (over four pairs) over a voltage range of 50 to 57 VDC using Category 5 or better cabling.

That's why with 11ac Wave 1, we are seeing a number of 11ac AP offerings with only two streams, designed to still operate on legacy 802.3af POE infrastructure. Wave 2 will make the jump to four stream APs, which will undoubtedly require 802.3at POE. Combine this with a 160 MHz channel, and you're now up to 2.3 Gbps. Not too shabby, if you believe a 160 MHz channel is viable. Again staring reality in the eye, 160 MHz makes sense for mesh and bridging applications, but it's unlikely to fly where you have high-density client access and need as many cells as you can manage.

Multi-User MIMO (MU-MIMO)

Multi-User MIMO is an exciting development in 11ac that will first manifest in Wave 2 products. Normally, with 11n and 11ac Wave 1, multiple spatial streams are transmitted over multiple antennas to only one client at a time. With each additional stream comes added throughput.

This is all great when the receiving device is a high-end laptop equipped with multiple receive antennas and can support multiple streams. However, with the explosion of BYOD, the vast majority of mobile devices are not high-end laptops. They are smartphones and tablets which only support one, or at maximum two spatial streams. During transmissions to these clients, the other streams go to waste.

MU-MIMO fixes this – well, it half fixes it. It makes it possible to multiplex airtime in such a way that individual clients hear only their part of the message at any given time, and the rest is nulled to them. This allows each of the three streams to be directed to a different one-stream client simultaneously. So potentially, three clients get serviced in the time it previously took to service one. This has some potential for scaling capacity if done right. The science is not so easy to implement however. When the individual streams are going to different clients, they must be tightly synchronized, so that they don't create a new type of interference between themselves. At the time of writing, MU-MIMO implementations are still in labs and have not seen the light of day, so it is too early to tell how well it works in practice. But even if it works half as well as it looks on paper, it will boost capacity for all those networks suffering from BYOD overload due to smartphones and tablets. Chip vendor Qualcomm is claiming 2x - 2.5x performance improvements, and it will be interesting to see if this holds up in implementations.

What vendors don't tell you however, is the fact that the clients cannot reply to the Access Point in MU-MIMO – this is a downlink only technology that requires a lot of CPU power and resource which low-end clients do not have. Additionally, it is likely that MU-MIMO results in a downshift in encoding scheme, from say 256-QAM to 64-QAM. So it is debatable whether two clients at 64-QAM is better than one at 256-QAM. Another point that is frequently overlooked is this MU-MIMO uses similar techniques to beamforming, and it is therefore mutually exclusive. The two techniques cannot be used concurrently - so you must consider what you lose by not having beamforming. Once again, what this points to is you need small cells with clients close to the AP and ideal RF conditions for MU-MIMO to add value.

The Context for your 11ac Migration

Set client performance goals

How do you know if you even need 802.11ac or whether your network is adequate if you don't set coverage, performance and capacity goals? If you're looking for 100% coverage, at a minimum, you expect everyone to be able to connect on any device, from anywhere without failure or unexpected disconnections, even when roaming. That's a pretty clear goal and easy to validate, but performance and capacity goals are harder to define.

These days well over 75% of enterprise deployments have accomplished pervasive indoor coverage, however users expect much more than mere coverage, they expect consistent speed, videos that stream continuously without buffering interruptions and low latency voice services that never fail. Enterprises are not there yet, only 30-40% of enterprises have a strategy for meeting performance SLAs.

The first question therefore is what performance level is right for your network? That depends on the application mix. You should set a minimum performance level that every user should be able to get at all times, wherever they are, and regardless of how many other users are present.

Below are a few examples of reasonable performance thresholds in different industries. Think about the applications in common use, and consider the mission critical nature of them (you're going to use QoS settings to give those priority over others, but still they need a minimum amount of bandwidth). Once you set a minimum performance level for each user, and you estimate the maximum number of users that may be present and active simultaneously in one place, you have a basis for per AP and per Service Area capacity planning or "network dimensioning".

	Performance SLA Target	VoIP	Remote Desktop	Email	Database	Browsing	Facebook	YouTube	Web Conference	e-Learning	Video Conference	HD Video Streaming	HD Gaming	Imaging
		100Kb	150Kb	400Kb	500Kb	750Kb	1Mb	1Mb	1.5Mb	2Mb	2Mb	3Mb	3Mb	4Mb
Warehouse:	1 Mbps	●		●	●									
K-12:	2 Mbps	●		●		●		●	●	●				
Enterprise:	3 Mbps	●	●	●	●	●			●					
Stadium:	4 Mbps			●		●	●	●				●		
University:	5 Mbps	●		●		●	●	●	●	●	●	●	●	
Hospital:	6 Mbps	●		●	●	●		●	●		●			●

Fig 6: Peak bandwidth requirements for common applications used in different industries

Here's a quick way to set your minimum performance thresholds: Add up the three most demanding applications and divide by two. This assumes users may be doing two things at once e.g. synching email while watching a movie. Now multiply by the maximum number of users you expect to be within a Service Area, and you now have a capacity number to aim for.

Example: Take a hospital wing with 60 beds (lets treat the 60 patients like university students, surfing the web and watching videos) and 12 medical and nursing staff with imaging requirements, and throw in 15 guests (more similar to enterprise users). Patients (60*5) + Guests (15*2) + Staff (12*6) = 402 Mbps capacity required in this Service Area with a per client average requirement of 5.36 Mbps.

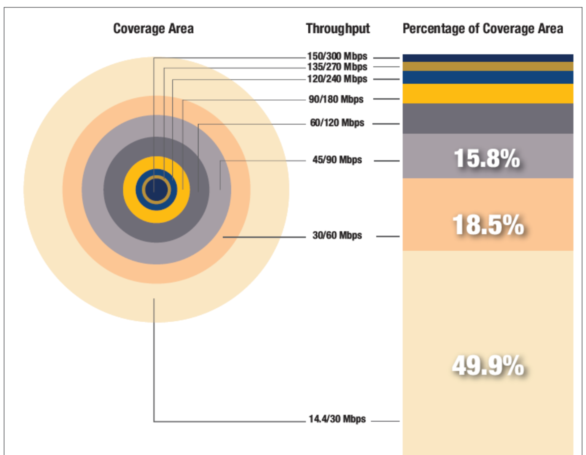

Fig 7: Coverage area for different MCS with 1 stream / 2 stream (chart courtesy of Juniper Networks)

Now, let's assume users are evenly distributed across the coverage area of each AP, which means 50% of users on average will be connecting at the lowest data rate (see Fig 7). And let's assume 40% of those devices are single stream smartphones, not three stream laptops. If all of them are active at the same time, then 15 devices out of 75 would be getting 15 Mbps data rates. If only 50% of users are active simultaneously, that's still ~7.5 users competing for ~15 Mbps of data rate (Multiply by 65% for payload) per AP. You don't know which users they are, so you have to assume they each have an average bandwidth requirement of 5.36 Mbps (payload). So how many APs are needed to service this demand? Only one AP would give them less than 1.5 Mbps each, so you would need four APs to meet the performance thresholds set in this scenario.

As you can see the thing to focus on is performance at the edge of an AP's coverage area. That's where 50% of users are, and they are getting the lowest data rates. If you get that right, anyone closer to the AP is getting more than they need. While four APs may be enough for today, you should always plan the network for growth. Over time as more users join the network, or more devices get connected, performance will drift downward. This is because all connected devices consume airtime and thus drain network capacity. For example a seemingly inactive smartphone in someone's pocket may be periodically synching with email or database applications. So you should also plan for 50% more capacity, in order to avoid a capacity shortfall in the near future. This brings our capacity requirements for the example above up to 603 Mbps and six APs.

This simple network dimensioning example makes many assumptions about the device capabilities mix (# streams), user types and their performance thresholds and the number of active users. They may be different for your network, so build a model, based on your client demographics. But the key message here is to base your capacity calculations on the worst case performance at the cell edge, not the best case scenario with a three stream laptop close to the AP.

Benchmark client performance

Unfortunately, your WLAN management system isn't much use when it comes to performance benchmarking. Knowing how many users are connected and which APs they are connected to says nothing about what they are doing and what their experience is like. You need different tools.

How can you possibly know what the user experience is like without benchmarking performance from the user's perspective? Users will only tell you when the performance is very bad, which is no help at all. You want to know whether or not you are meeting your performance goals and for how much of the time that is the case. Regularly benchmarking the user experience is particularly critical in higher education and large public venues where you have extremely high client density and massive people flows between areas.

Unless you have performance monitoring software and sensors, like the 7signal Sapphire Wi-Fi performance management system, doing this yourself is tedious, but necessary. But here are a few alternative methods: Pick a time of day which your WLAN management indicates is a peak connection time. Then do a site survey with a laptop or mobile device, sampling Internet performance from various Service Areas, using any one of a number of test tools, like IXIA's ixChariot or Ookla's Speedtest.net.

Another idea is to engage your users to do it for you by creating a landing page requesting the performance results and your chosen speed test in an iframe below the results input form. Notify users by email that you want to know what their performance is like, and then tell them what to do. You'll be surprised how many respond. But be sure to make it easy for them to run the test, and give you the results. When the results are in, share them with your users, so they are rewarded for their efforts and willing to do it again.

Monitor your client mix closely

As this guide has shown, how much you are able to realize the benefits of 11ac depends on the client mix and how well the configuration is tuned for that mix. In fact, there are certain features of 11ac that should not be utilized until your 11ac/11n client mix reaches a certain ratio. It is therefore crucial that you keep tabs on the client mix, so you can recognize the optimum time to adjust your channel plan or configuration in order to maximize performance or capacity or both.

There are several ways to do this, but one of the easiest ways is to put 11ac access points in high traffic areas where lots of people are expected to pass by them, in close proximity. Then run regular reports to see the speeds the devices are connecting to those APs. Of course, you're going to catch more handheld devices than laptops, but regardless, you will see the trends, and will be able to project when your mix is likely to reach a certain ratio.

Let's revisit the example of the hospital wing requiring 603 Mbps of capacity. In the real world, you can probably provide that capacity with 6–8 11n APs, depending on the physical area involved. But if you think 1Gbps speeds are required, you might think two or three 11ac access points should be more than enough. Which is right? Well, it depends greatly on the physical environment and on the client mix. In practice, given the hostile RF environment that is typical in hospitals, it is more likely that you'll require the same 6-8 APs if you opt for 11ac over 11n. So what are you gaining? Higher client speeds? Perhaps, but not for everyone, or everywhere.

Your rationale for deployment, your channel plan and optimal access point placement - all hinge on the client mix and how those clients are distributed throughout your facility. Since smartphones and tablets tend to be on a shorter life cycle (2-3 years) than laptops (3-4 years), and seeing as they now make up 40-60% of the connected clients in many enterprise environments, you can expect a fairly rapid transition of the majority client population to 11ac. So, it is crucial you are on top of this and time your 11ac deployments and channel plan changes at appropriate client mix milestones.

General 802.11ac Migration Tips

Where to focus your 11ac investment

As previously mentioned, using 11ac to boost mesh backhaul capacity and bridging is a no-brainer, and it is justified to replace existing 11n APs which you can re-deploy elsewhere. But when it comes to upgrading client access, the best bang for your buck is less clear, since wherever you choose to put 11ac has additional cost implications coming in the form of site surveys, LAN infrastructure and POE upgrades as well as revising your channel plan. Initially, one-for-one replacement is a good way to start, as you should be able to avoid or limit the moves and changes necessary and get some immediate benefit from 11ac for a portion of your client base.

Some people argue that you should concentrate on floor by floor or building by building deployments, so that you can rework your channel plan and cabling if necessary, systematically. While there is merit in this approach from a LAN infrastructure perspective, it is not the fast path to maximizing performance and capacity where it is most needed. Instead, we believe the greatest performance and capacity impact comes from focusing on hitting the high density need areas first. Break up your facility or campus into types of spaces (Service Areas), define what that space needs, then look for similar spaces and replicate the solution. Focus on use/demand categories, not buildings.

With BYOD accounting for over 50% of devices on a corporate network, and nearly 100% of higher education, you have no control over the device capabilities. That means you cannot forget the 2.4 GHz band for many years to come. In fact, you'll want 5 GHz as your capacity channel and 2.4 GHz as your coverage channel. 2.4 GHz is the best band to cost-effectively ensure pervasive coverage across your facility. Use the second radio of a dual-radio 11ac access point to offer 11n at 2.4 GHz. However, since you cannot deploy 2.4 GHz with the same density as 5GHz, many of the second radios on your 11ac APs may end up being turned off and unused. This has been the common practice in high-density 11n deployments. But is this waste really necessary? We think not. For high density deployments, consider saving cost with single radio 11ac APs if your preferred vendor offers them.

How 11ac improves 11n performance

Also keep in mind that the newer 11ac access points are just better than the previous generation of 11n APs. They feature newer lower-latency silicon, more sensitive radios, more linear transmitters, faster processors and better antenna designs. So, even when they are operating in 11n mode, you can expect a 5-10% performance improvement for 11n clients.

As 11n and legacy client devices gradually get retired and replaced with devices capable of 11ac, the remaining clients will enjoy incremental performance improvements as well. The reason for this is that since most of the 11ac clients are able to transmit at higher rates, they spend less time on air. Airtime is the medium of wireless networks. It is a finite resource. Saving airtime on 11ac transmissions means there is more available for 11n and 11a/g clients. However, as good as this sounds, reality may be surprising. The better the Wi-Fi works, the more users who use it. This drives network load up and this has its consequences.

Here's a general indication of the performance improvement you may see with 11n clients connected to 11ac access points as the client population shifts to 11ac.

11ac Population	25%	50%	75%
11n Population	75%	50%	25%
Performance boost for 11n	5%	10%	15%

Fig 8: Potential performance boost for 11n clients as population shifts to 11ac.

There is no substitute for small cells

A consistent message throughout this paper has been the importance of small cells and line of sight, if you want to get the full benefits of 11ac Wave 1 or Wave 2. For many years now, the best practice for maximizing performance and capacity has been to reduce cell size and turn down the power on APs. Bottom line, clients that are far away from the AP have lower performance, because they use a lower MCS data rate, and those using lower data rates effectively bring down the performance of others, because they hog the finite resource of airtime. Reducing cell size means clients are closer to the AP, so they use higher data rates on average, and it tends to result in fewer devices per AP. Keeping AP output power level similar to your clients' power level, or slightly below it, may offer the best compromise for maximum capacity.

Another important aspect to maximizing the capacity and decreasing the cell size takes place at the receiver end. Wi-Fi APs and terminals determine radio channel availability by first listening to it. If they receive frames from other devices with signal levels exceeding the pre-determined threshold, they will back off and wait for their turn. The Wi-Fi standard has determined these signal level thresholds. Generally, this value is referred to as Clear Channel Assessment (CCA) threshold. If this value is increased, more parallel transmission can take place and higher reuse is allowed. This is equivalent to reducing the cell size at the receiver side. If the CCA value is decreased, less parallel transmissions are allowed to be present and cell size would be increased at the receiver. Even though this CCA threshold feature is standardized, it's typically not activated in your equipment or may not even be implemented. If indeed it is implemented, then you should know that the feature is called different names depending on the vendor. For example, Cisco calls this RX Start of Packet (RX SOP) and Aruba has named a similar feature as Cell Size Reduction (CSR). These are advanced settings, which provide potential for capacity increases, but should be implemented with special care.

Cell size reduction strategies apply as much to 11ac as it did for 11n, perhaps more so.

Client Capabilities, What to Expect

As we have stressed repeatedly, everything hinges on the client mix and on the capabilities of those clients. In your estimations of potential performance gains, or required performance thresholds, it is a mistake to assume that a low-budget smartphone or even a high-end tablet will support three spatial streams and 160 MHz channel bonding. Smartphones have been shipping with 802.11ac chipsets since 2013, and according to ABI research smartphones will account for 46% of all 11ac chipsets sold by 2018. Here is a useful chart that illustrates what new devices coming onto the market might be capable of.

	Pre-2014 Devices			Pre-2015 Devices		
Type of device	**Radio Type, Streams**	**Channel Width**	**Max Data Rate**	**Radio Type, Streams**	**Channel Width**	**Max Data Rate**
Smartphone	802.11n, 1-stream	20 MHz	72 Mbps	802.11ac, 1-stream	20/40/80 MHz	433 Mbps
VoIP handset	802.11a/b/g or 802.11n, 1-stream	20 MHz	54 Mbps	802.11a/b/g or 802.11n/ac, 1-stream	20 MHz	87 Mbps
Tablet	802.11n, 1-stream	20/40 MHz	72/150 Mbps	802.11ac, 1-stream	20/40/80 MHz	433 Mbps
Netbook	802.11n, 2-stream	40 MHz	300 Mbps	802.11ac, 2-stream	80 MHz	867 Mbps
Low-end laptop	802.11n, 2-stream	40 MHz	300 Mbps	802.11ac, 2-stream	80 MHz	867 Mbps
High-end laptop	802.11n, 3-stream	40 MHz	450 Mbps	802.11ac, 3-stream	80 MHz	1.3 Gbps

Fig 9: Capabilities of 11ac clients hitting your network in 2015

In addition, it is a myth that 11ac consumes too much battery power on handheld devices. 11ac incorporates a variety of new features to put the radio in a sleep mode when it isn't transmitting. Smartphones are getting bigger, which means the batteries are bigger too, and battery technology continues to improve.

We are already seeing new Samsung® smartphones being released with 2x2 implementations. Apple iPhone is 1x1 but is sure to follow suit in due time. However, most handheld devices for the next year or so will be limited to one antenna and therefore only one stream. But that one antenna will support channel bonding, which as pointed out earlier, is a linear contributor to higher performance. For these reasons it is important not to discount 80 MHz bonding as too difficult to incorporate into your channel plan. Figure out where in your facility 80 MHz can have the most impact and time the migration for when the 11ac client mix crosses the 50% mark. If you really want 11ac to make a big difference, you have to nail the 80 MHz channel plan.

Handheld devices with more than 2x2 antennas are unlikely to hit the streets for many years to come. It is extremely difficult to pack more antennas into these small form factors. And laptops are not likely to get any bigger, so they will hit a wall with 4x4. Therefore access points with more than 3 or 4 streams many not make much sense either, unless MU-MIMO really works well.

Upgrading the Wired Network

Many high-end 802.11ac access points will support two Gigabit Ethernet downlinks from access switches. An important question is, do you really need two links? Pulling cable is expensive. Truth is, except for extremely rare deployments of 11ac, there is little chance of generating anywhere near 1Gbps of traffic with 11ac Wave 1 products. When you begin deploying Wave 2, you could potentially exceed 1 Gbps for short peaks or in special use cases where 80 MHz or 160 MHz bonding is a fit.

Capacity over 1 Gbps in AP downlinks is not going to be an issue for years in dense Wi-Fi deployments, since 80 MHz or 160 MHz bonding is ill-advised in these scenarios. AP/Wi-Fi radios contend so much that aggregate user data speeds/AP will not get even close to those values. In addition, if there is a very short peak throughput actually getting there, investing any significant money to support more than 1 Gbps downlink does not make any sense. Remember too, all the vendor quoted numbers are based on MCS data rates, which are not data payload, or even Ethernet frames. Traffic going over the AP's Ethernet link is stripped of all the radio overheads, which typically constitute approximately 30%.

The practical limitation is driven by limited airtime in areas with multiple APs operating on the same channel. In order to get maximum advertised peak throughputs leading to above 1 Gbps speeds, very wide channels need to be used. At the same time, channel reuse factor drops dramatically and it may be that only three channels are again available. This means that since nearby APs are on the same channel, achievable speeds drop rapidly.

So the recommended strategy is to pull a second cable when you can, but not to delay roll out due to lack of cabling, especially if you are doing in-place replacement of APs without moving them. When you are moving APs or adding new ones, pulling two cables is advised, as this will provide future proofing for 11ac Wave 2 and beyond with 802.11ax. This is also a good time to pull extra cables to nearby APs. If you are pulling new cable, make sure its Category 6 or 7 not Category 5. Cat 6 and 7 are rated for Gigabit throughput over 100 meter spans, Category 5 is not.

Radio downlinks are not the only consideration. Edge switches should be upgraded to 10 GbE uplinks as well. Many switch vendors have been slow to deliver low-cost L2 and L3 switches with 10 GbE uplinks. Unless you're ready to retire access switches prematurely, upgrading the access layer is a costly endeavor. Delay it as long as you can by exploiting EtherChannel or some other link aggregation scheme.

POE is the final cable plant consideration. In almost all cases, full featured (3x3:3) Wave 1 products require an upgrade from 802.3af to 802.3at. Early 11n products could not operate on 802.3af, but then over time newer silicon and AP thermal designs got better, and the power consumption was kept under the bar. The same may happen for 11ac, but the tolerances are more borderline than they were with 11n. Certainly today almost all Wave 1 products require 802.3at for full operation.

Now is the time to commit to 802.3at to get the full benefit of 11ac Wave 1, and future proof for Wave 2 and beyond (802.11ax). If your existing switches support 802.3af, in the short term, supplement with 802.3at power injectors. But insist on high density 802.3at in your new edge switch selection criteria.

Summary of Key Points

1) Upgrade mesh nodes and point-to-point and multi-point bridges to 11ac using 11ac Wave 1 as soon as possible, then repeat when Wave 2 is available. You can redeploy the previous Wave 1 APs for client access.

2) 11ac does nothing to improve 2.4 GHz coverage or capacity, and you cannot discount 2.4 GHz requirements. Use 2.4 GHz as your lifeline coverage band and 5 GHz for capacity scaling.

3) Use 80 MHz selectively, taking care of line-of-sight requirements, and use 40 MHz widely. 40 MHz and 20 MHz allow higher density and more flexible channel reuse.

4) If your network is designed with wideband channels, only 11ac clients can benefit from the bonded channels above 40 MHz. Meanwhile 11n clients are limited to fewer (primary) channels and suffer. 80 MHz efficiency is not realized for the 11n clients. Therefore your client mix should drive channel plan migration.

5) Until your client mix is more than 50% 11ac compatible clients, there is little benefit in replacing 11n APs which are not fully depreciated, especially not 3x3 versions.

6) Begin using 40 MHz channels in 11n now, if you are not doing so already. Very dense deployments may still benefit from using 20 MHz channels, as this allows more cells in closer proximity to one another.

7) 11ac benefits 11n and legacy clients indirectly by freeing up airtime. The more 11ac clients, the more airtime is freed up. And because 11ac silicon, antennas and AP design are just better.

8) Use 80 MHz for high-density Service Areas only where clients have good line-of-sight to the AP and where it is easy to isolate or buffer the Service Area.

9) Monitor your client mix regularly and be prepared to adjust your channel plan several times over the coming years to optimize performance as your client mix shifts in favor of 11ac.

10) Take a close look at your LAN infrastructure and prepare to upgrade edge switches to 10 Gbps uplinks and POE to 802.3at to future proof it for 11ac Wave 2.

11) You probably do not need two GbE downlinks to APs, since total AP load is unlikely to ever exceed 1 Gbps.

12) MU-MIMO holds promise for scaling capacity in environments with high ratios of smartphones and tablets vs laptops, but 1 more stream and 160 MHz, which is of dubious value, is no reason to hold out for Wave 2, and not invest in Wave 1 right now.

7signal's Role in 11ac Migration

As you have now learned, 11ac migration is complicated. There are just too many dials and other variables that leave even Wi-Fi gurus with their heads spinning. Even if you follow the guiding principles we've laid out, optimizing the configuration of your wireless network is a moving target because the client mix and traffic load are both in constant transition. To maximize performance you must be prepared to benchmark performance regularly and constantly monitor the client mix so you can make the necessary adjustments to your channel plan and AP configurations at the right time.

By deploying 7signal's Wi-Fi performance management system, you can monitor network performance automatically 24 by 7, set performance goals, have the network tell you when it is time to change something, and also what to change. Enterprises using 7signal's Sapphire system for 11n, typically see a 30-50% improvement in user performance across the board, even without adding access points or the expense of migrating to 11ac. As we roll out support for 11ac we expect to see the same levels of improvement over what most network administrators can accomplish on their own.

About 7signal

7signal is a U.S.-based network technology company with a mission to make wireless LANs work — reliably. The company's groundbreaking platform provides Wi-Fi performance management and assurance for any vendor-supplied wireless network. 7signal's solutions help ensure that all end-users in hospitals, institutions of higher education and enterprises can connect and stay connected to their critical applications at all times.

7SIGNAL®
THE WI-FI PERFORMANCE COMPANY

Phone: 855-763-9526
Email: info@7signal.com

www.7signal.com
Twitter @7signal
Linkedin.com/company/7signal-solutions