

White Paper: IEEE 802.11ac Migration Guide

This paper is a primer and migration guide for 11ac technology, offering recommendations, best practices, and tips for a successful deployment.

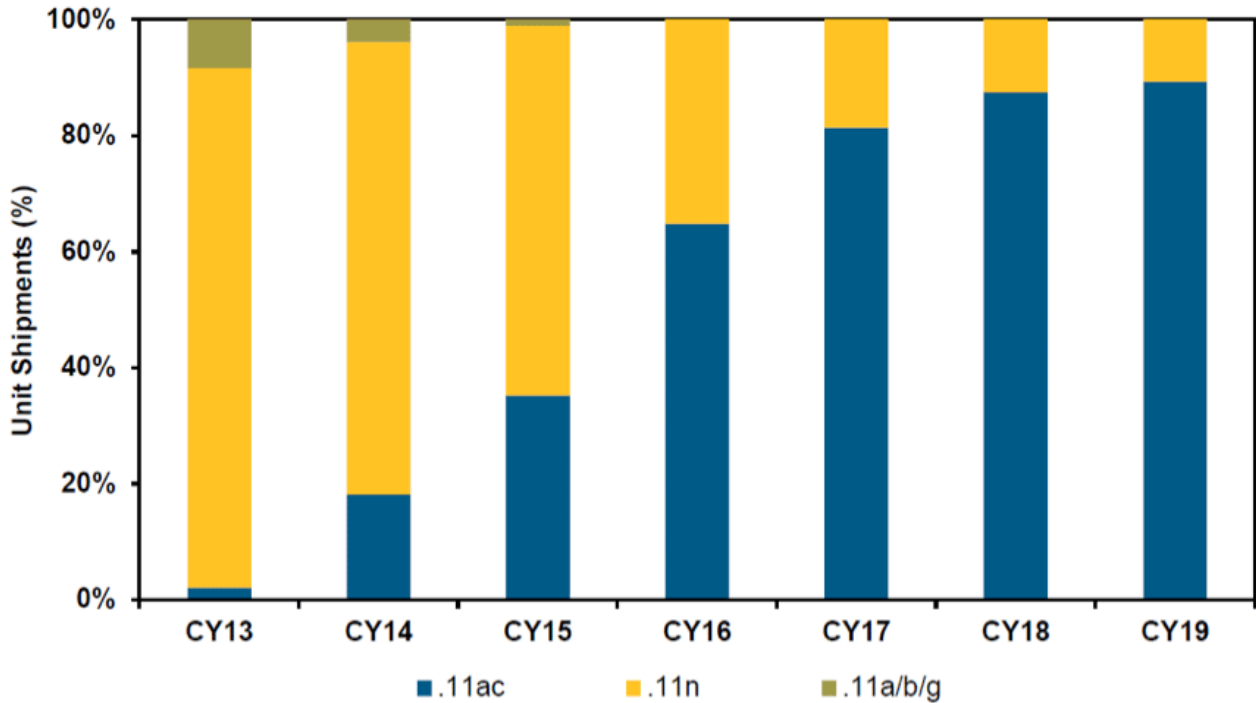
- » TABLE OF CONTENTS
- » Introduction
- » 1-2-3's of 11ac
- » Channel Width
 - » SU-MIMO vs. MU-MIMO
 - » Spatial Streams (SS)
 - » 256 QAM
 - » SNR
 - » Benefits of 11ac
- » Wired Challenges & Considerations
- » 11ac Deployments
 - » Rip-n-Replace & Rolling Upgrades
 - » Out With The Old
- » Planning and Diagnostic Tools
- » Summary
 - » Migration Checklist

Introduction

IEEE 802.11ac (often called just “11ac”) is the rage. Everyone loves trying out sexy new technology, with Wi-Fi being high on everyone’s list. Vendors first launched enterprise-class Wave-1 Access Points (APs) around Q4/2012, and since then, 11ac APs have been flying off shelves at a break-neck pace. The marketing hype around 11ac has been over-the-top due to a competitive market place, and depending on whom you’re listening to, 11ac APs can leap small buildings in a single bound.

According to Infonetics*, 11ac would enjoy ~80% market share for enterprise APs by the end of 2017 and almost 90% by 2019.

Access Points by Technology: 802.11a/b/g vs .11n vs .11ac



*Wireless LAN Equipment and WiFi Phones Quarterly Worldwide and Regional Market Share, Size, and Forecasts: 4Q14

Many owners of 11n infrastructure have been making do with their deployments for many years and have decided to wait for 11ac Wave-2 products to come to market before upgrading. 11n networks, whether of the 2x2:2, 3x3:2, or 3x3:3 variety, can be optimized for high performance in most cases, and are sufficient for the majority of enterprise deployments today. In many situations, it would be fair to say that 11ac is a nice-to-have rather than an essential upgrade. The downside to aging 11n deployments is that many early 11n AP models are no longer supported with new code updates, which limit their security and performance feature sets.

This paper is a primer and migration guide for 11ac technology, offering recommendations, best practices, and tips for a successful deployment. All statements in this paper are in reference to enterprise-class Wi-Fi solutions and may not apply to consumer-class equipment.

1-2-3's of 11ac

11ac is a 5GHz technology, meaning that the IEEE 802.11ac amendment does not specify its use in the 2.4GHz ISM band. Use of wider channels requires more available frequency space, and the 2.4GHz band is limited to a total of 83.5MHz. Any implementations of the 11ac physical layer specification (PHY) in 2.4GHz are proprietary.

11ac technology isn't just about the radios. APs are small computers, each having a CPU, RAM, Flash, etc. With each new generation of radio technology, we also get new software features, some of which weigh heavily on the CPU of APs and/or controllers. Some new 11ac dual-radio APs have large CPUs, often dual-core, plenty of RAM, encryption offload, dual Gigabit Ethernet ports, and many other high-end hardware features.

So what makes 11ac so special that it would replace the aging IEEE 802.11n ("11n")? To answer that question accurately, it's important to understand that 11ac has been launched in two "waves" (called "Wave-1" and "Wave-2"), based on radio chipset capabilities. The chart below shows a brief difference between the technologies implemented into each of the two Waves.

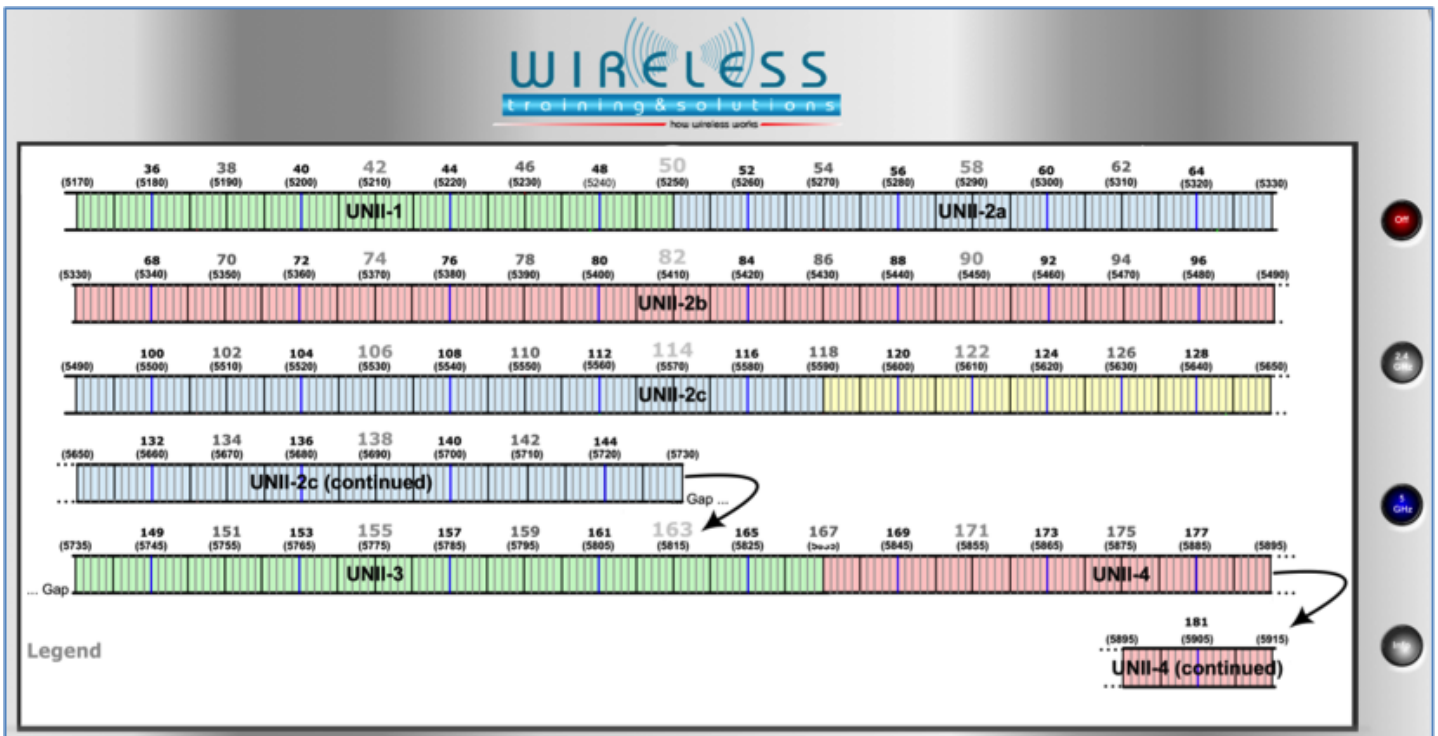
PHY/Feature	802.11n	Wave-1 802.11ac	Wave-2 802.11ac
Channel Width	20, 40 MHz	20, 40 MHz	20, 40, 80, 160MHz
Spatial Streams (SS)	1, 2, 3	2, 3	2, 3, 4
QAM Modulation	64 QAM	256 QAM	256 QAM
MIMO Type	SU-MIMO	SU-MIMO	MU-MIMO
MCS Support	MCS 0-23 for 1, 2, 3 SS	MCS 0-9 for 1, 2, 3 SS	MCS 0-9 for 1, 2, 3, 4 SS
Maximum Data Rate	450Mbps	1.3Gbps	3.467Gbps
TxBF	No	Variable	Yes
Radio Variations	2x2:2, 3x3:2, 3x3:3	2x2:2, 3x3:3	4x4:4*

*Currently expected from leading enterprise Wi-Fi chipset manufacturers.

11ac Wave-1 AP and client device sales have been very successful for the industry, and 802.11ac devices are sure to continue selling well after 11ac Wave-2 devices reaches the market. It takes time for a new line-up of products to be brought to market due to design, manufacturing, and certification. Early 11ac Wave-2 APs will not have DFS certification for some time, and both code stability and feature performance will be unproven.

Channel Width

802.11n devices are capable of supporting either 20MHz or 40MHz channels. 11ac Wave-1 devices support 20, 40, and 80MHz channels, while 11ac Wave-2 devices support 20, 40, 80, and 160MHz channels. 160MHz channels are not currently useful in enterprise deployments due to the lack of contiguous channel space in the 5GHz UNII bands, but with the FCC Report and Order (R&O) 14-30, dated April 1, 2014, changes are proposed that could allow up to four 160MHz, non-overlapping channels in the U.S. For other countries, it depends on regulatory authorities. The channel simulator below, located at WiFiChannelSimulator.com shows the available license-free spectrum since the FCC R&O.



Credit: Wireless Training and Solutions

The simplest way to add throughput to a Wi-Fi network is to double the channel width, provided there are enough reusable, wide channels available. Doubling the channel width roughly means doubling the throughput capability of the channel. Like everything though, the additional throughput comes at a price. As channel width is doubled, the allowable output power is cut in half, across the entire channel. This may not be an issue in some environments, but in others it can create an unnecessary technical challenge. Doubling the channel width also increases the noise floor by 3dB and increases the chance of a collision. For that reason, 80MHz and 160MHz channels are usually dynamic. APs may use protection mechanisms, like RTS/CTS, to “clear” 80 or 160MHz channels. If only a portion of the wide channel is usable, then APs will decrease the channel width of that individual transmission to get as much throughput as possible.

Just because you can use 80MHz and 160MHz wide channels doesn't mean you should. It is recommended to use 20MHz channels in high-density environments such as auditoriums, ballrooms, trade shows, airports, and arenas because they increase channel use efficiency. Low-density/high-throughput environments, such as open office areas, may benefit from 40MHz channels in 5GHz, provided there are enough channels for a reasonable channel reuse plan. If only 1-2 APs will be deployed in a facility (e.g. a branch office), and there is only a minimal amount of interference (modulated and unmodulated), then using 80MHz channels may work well. There is currently no appropriate use for 160MHz channels other than highly directional point-to-point links. If there is a specific area where very high throughput is consistently required, then configuring one AP to use an 80MHz channel in that area may be OK, so long as the nearby APs do not use any part of that 80MHz channel.


SU-MIMO vs. MU-MIMO

All 11n devices are Single-User MIMO (SU-MIMO) enabled, meaning that only one transmission, whether uplink or downlink, can happen simultaneously on a channel. 11ac Wave-1 devices are MIMO (SU-MIMO) enabled while 11ac Wave-2 APs will be endowed with Multi-User MIMO (MU-MIMO) technology.


MU-MIMO is a downlink-only (from AP to client) technology that allows for multiple simultaneous transmissions using Transmit Beamforming (TxBF) technology to increase RF signals in some areas while nulling them in others. Most MU-MIMO APs will be capable of 3 or 4 simultaneous transmissions. MU-MIMO technology enhances MAC efficiency when a 3SS or 4SS capable AP is supporting multiple 1SS capable clients.

Spatial Streams (SS)

A spatial stream is the technology of splitting a data stream into multiple parts (called spatial streams) and then transmitting them simultaneously across multiple radio chains on the same channel.



(a) Single User MIMO, 4 streams

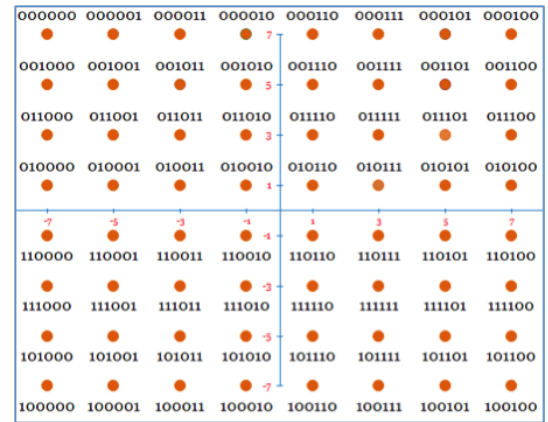


(b) Multi User MIMO, 2 users, 2 streams each

The use of multipath and digital signal processors (DSPs) allows MIMO capable receivers to decode the spatial streams and to rebuild the data stream. 11n and 11ac Wave-1 devices both support up to 3SS, but 11ac Wave-2 devices will support up to 4SS.

256 QAM

Modulation is the means by which data is encoded onto carrier waves. 11n is limited to 64 QAM modulation, while 11ac introduced 256 QAM. 256 QAM is a more sophisticated modulation type, which requires significantly higher SNR to operate. For this reason, it's common to see client/AP connections down-step back to 64 QAM after only 40-50 feet. The chart on the right represents a 64 QAM constellation, where each dot represents 6 bits. A 256 QAM constellation has 64 dots per quadrant, and 8 bits are encoded per dot.



SNR

11ac's higher MCS rates are tied to using 256QAM modulation. The best place to reference the components of an MCS is at MCSIndex.com. Keep in mind that doubling the channel width raises the noise floor by 3dB, and therefore, 80MHz channels will automatically have a 6dB higher noise floor than 20MHz channels. Referencing the chart below, **it takes at least 37dB of SNR to achieve MCS9** (the highest MCS rate for 11ac) for an 80MHz channel. This is an unreasonably high SNR, and without significant channel reuse, the co-channel contention (interference) would be very significant. Below is an excellent chart for mapping MCS rates to required SNR.

MCS Value Achieved by Clients at Various Signal to Noise Ratio Levels (SNR)

Protocol	Channel	1	2	3	4	5	6	7	8	9	10	Modulation Key
802.11b	20MHz	None	None	None	MCS 0	MCS 0	MCS 0	MCS 1	MCS 1	MCS 1	MCS 1	
802.11a/g	20MHz	None	MCS 0	MCS 0	MCS 1	MCS 2	MCS 2	MCS 2	MCS 2	MCS 3	MCS 3	
802.11n	20MHz	None	MCS 0	MCS 0	MCS 0	MCS 1	MCS 1	MCS 1	MCS 1	MCS 2	MCS 2	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11n	40MHz	None	None	None	None	MCS 0	MCS 0	MCS 0	MCS 1	MCS 1	MCS 1	
802.11ac	20MHz	None	MCS 0	MCS 0	MCS 0	MCS 1	MCS 1	MCS 1	MCS 1	MCS 2	MCS 2	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11ac	40MHz	None	None	None	None	MCS 0	MCS 0	MCS 0	MCS 1	MCS 1	MCS 1	
802.11ac	80MHz	None	None	None	None	None	None	None	MCS 0	MCS 0	MCS 0	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11ac	160MHz	None	None	None	None	None	None	None	None	None	None	
SNR in dB		11	12	13	14	15	16	17	18	19	20	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11b	20MHz	MCS 2	MCS 2	MCS 2	MCS 2	MCS 2	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	
802.11a/g	20MHz	MCS 4	MCS 4	MCS 4	MCS 4	MCS 5	MCS 5	MCS 5	MCS 6	MCS 6	MCS 7	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11n	20MHz	MCS 3	MCS 3	MCS 3	MCS 3	MCS 4	MCS 4	MCS 4	MCS 5	MCS 5	MCS 6	
802.11n	40MHz	MCS 1	MCS 2	MCS 2	MCS 3	MCS 3	MCS 3	MCS 3	MCS 4	MCS 4	MCS 4	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11ac	20MHz	MCS 3	MCS 3	MCS 3	MCS 3	MCS 4	MCS 4	MCS 4	MCS 5	MCS 5	MCS 6	
802.11ac	40MHz	MCS 1	MCS 2	MCS 2	MCS 3	MCS 3	MCS 3	MCS 3	MCS 4	MCS 4	MCS 4	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11ac	80MHz	MCS 1	MCS 1	MCS 1	MCS 1	MCS 2	MCS 2	MCS 3	MCS 3	MCS 3	MCS 3	
802.11ac	160MHz	MCS 0	MCS 0	MCS 0	MCS 1	MCS 1	MCS 1	MCS 1	MCS 2	MCS 2	MCS 3	
SNR in dB		21	22	23	24	25	26	27	28	29	30	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11b	20MHz	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	
802.11a/g	20MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11n	20MHz	MCS 6	MCS 6	MCS 6	MCS 6	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	
802.11n	40MHz	MCS 5	MCS 5	MCS 6	MCS 6	MCS 6	MCS 6	MCS 6	MCS 7	MCS 7	MCS 7	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11ac	20MHz	MCS 6	MCS 6	MCS 6	MCS 6	MCS 7	MCS 7	MCS 7	MCS 7	MCS 8	MCS 8	
802.11ac	40MHz	MCS 5	MCS 5	MCS 6	MCS 6	MCS 6	MCS 6	MCS 6	MCS 7	MCS 7	MCS 7	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11ac	80MHz	MCS 4	MCS 4	MCS 4	MCS 5	MCS 5	MCS 6	MCS 6	MCS 6	MCS 6	MCS 6	
802.11ac	160MHz	MCS 3	MCS 3	MCS 3	MCS 4	MCS 4	MCS 4	MCS 5	MCS 5	MCS 6	MCS 6	
SNR in dB		31	32	33	34	35	36	37	38	39	40	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11b	20MHz	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	
802.11a/g	20MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11n	20MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	
802.11n	40MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11ac	20MHz	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	
802.11ac	40MHz	MCS 7	MCS 8	MCS 8	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11ac	80MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 8	MCS 8	MCS 9	MCS 9	MCS 9	MCS 9	
802.11ac	160MHz	MCS 6	MCS 6	MCS 6	MCS 7	MCS 7	MCS 7	MCS 7	MCS 8	MCS 8	MCS 9	
SNR in dB		41	42	43	44	45	46	47	48	49	50	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11b	20MHz	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	MCS 3	
802.11a/g	20MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11n	20MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	
802.11n	40MHz	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	MCS 7	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11ac	20MHz	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	
802.11ac	40MHz	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	802.11 Type Key 802.11b 802.11a/g 802.11n 802.11ac
802.11ac	80MHz	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	
802.11ac	160MHz	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	MCS 9	

Credit: RevolutionWiFi.net

Benefits of 11ac

In typical enterprise Wi-Fi deployments, client devices use no more than 5Mbps on average. Certainly peak throughput can spike well over 100Mbps for a given client, but it's the minority of Wi-Fi client devices that sustain high throughput for long periods of time. This means that most current 11n deployments are often sufficient for low-density office environments and similar deployments. Realizing a good ROI for 11ac is dependent on upgrading to 11ac client devices and optimizing the Wi-Fi network's design, deployment, and configuration. A well-optimized 3x3:3 11n network can outperform a poorly designed 3x3:3 11ac deployment given the same installed client device base.

Where 11ac really shines, and infrastructure upgrades are often warranted, is in high-density and/or high-throughput environments, including:



- Arenas
- Amphitheatres
- Auditoriums
- Ballrooms
- Ballparks
- Convention Centers
- Large Classrooms
- Concert Halls
- Casinos
- Lecture Halls
- Large Meeting Rooms
- Press Areas
- Public Events
- Stadiums
- Trade Shows
- Trading Floors

11ac APs have higher quality radios (e.g. better receive sensitivity), faster CPUs, and will support the latest performance features from its manufacturer. While still unproven, 11ac Wave-2 promises to improve MAC efficiency using MU-MIMO in low-end and mobile device environments.

An important aspect of 11ac deployment is that 11ac APs (both Wave-1 and Wave-2) can seamlessly co-exist with 11n APs (all varieties). It's an industry best practice to locate 11ac APs in high-density/high-throughput areas while relegating 11n APs to low-density/low-throughput areas. It's also common to see specific APs deployed using wide channels while the majority of APs in a deployment are using only 20MHz channels. This allows for efficient spectrum use while also permitting the flexibility of very high throughput in specific areas. A caveat to 11n/11ac co-existence is that 11n and 11ac APs, even from the same manufacturer may support different feature sets due to CPU/RAM limitations in 11n APs. This would mean that leading-edge features could only be implemented in specific areas, and leads to design constraints such as keeping 11n and 11ac as isolated as possible from each other (e.g. different buildings or sites).

Wired Challenges & Considerations

It's typical for enterprise Wi-Fi manufacturers to release a high-end AP first when introducing new technology into the market due to the marketing impact. First generation, high-end APs often require 802.3at (PoE+) to operate at full capacity. It usually doesn't take long before WLAN infrastructure manufacturers follow up with a low-end AP that operates at full capacity on 802.3af (PoE), and then a mid-range AP that barely (or questionably) squeezes into a PoE budget. That's been the consistent go-to-market strategy during 11n and 11ac Wave-1 product launches across a wide variety of enterprise Wi-Fi vendors.

11ac Wave-2 APs are designed for high performance and have a couple of new power-hogging enhancements. The first enhancement is 4 spatial stream (4SS) capability. 4SS capability means having four simultaneously transmitting and receiving radio chains, and while this can offer speed enhancements, it significantly increases the power draw. The second enhancement is CPU and memory. Since we now have the capability to go move more data with 4SS, we need a fast enough CPU and memory to reach the APs potential throughput. CPUs are a main source of power consumption within an AP.

Proliferation of PoE+ has been slow in comparison to the adoption rate of 11ac because of the long (~10 year) buying cycle of Ethernet switches. Fortunately, there are now many new drivers for PoE+ other than just high-end APs. A variety of industrial and commercial devices can also operate within the 30W PoE+ budget. This trend will allow AP manufacturers to worry much less about fitting APs within a 15.4W PoE budget and to focus more heavily on high performance.

There has been much misinformation around the need for >1Gbps Ethernet backhaul links on 11ac APs. Stated plainly, backhaul links greater than 1Gbps are not needed for 11ac – neither with 11ac Wave-1 nor 11ac Wave-2. Recent efforts by standards bodies have produced the possibility of seeing 2.5Gbps and 5Gbps Ethernet in the near future, but these speeds will not be needed to support two-radio, dual-band 11ac Wave-1 or 11ac Wave-2 APs.

An AP's throughput is around 50% of the data rate, best case. Enterprise customers will not use 160MHz wide channels, and 80MHz channels will only be used in very specific cases. For purposes of calculation, using the best-case scenario, and considering both ends of the link as they are 11ac, we could use the following numbers:

$$80\text{MHz channel} \times 4\text{SS} \times 256\text{QAM} + \text{SGI} = 1.733 \text{ Gbps (Data Rate)} \times 50\% = \sim 867\text{Mbps}$$

An important clarification here is that Gigabit Ethernet is full duplex, meaning that it can move 1Gbps uplink and another 1Gbps downlink simultaneously. In contrast, the ~867Mbps mentioned above is unidirectional, does not take into consideration any RF interference, and makes the assumption that there is only one client communicating with one AP (no congestion). This makes these numbers a bit unrealistic to achieve.

Only a short time ago, most clients would pull much more data (downlink) than they would push (uplink). With social media sites being so popular, we see nearly a 50/50 split of uplink/downlink traffic in most networks. This 50/50 split of throughput would essentially cut the traffic in each direction to half of the total (e.g. ~433Mbps uplink, ~433Mbps downlink). However, with bi-directional data flows, the AP would be competing with its own client, and the 802.11 contention process would add overhead (collisions, additional back-off, etc.). This could bring the realistic throughput down to under 400Mbps in each direction, best case, and that's only ~40% utilization of a 1Gbps connection.

Again, this is a perfect throughput storm, and there is practically no chance of reaching this in real-world environments due to:

- Contention while multiple clients try to access the channel (co-channel contention)
- Adjacent Channel Interference (ACI)
- RF interference sources
- Protection mechanisms for backward compatibility
- Mixed PHY client environment or wireless infrastructure
- CPU limitations on the AP
- Poorly performing code on the AP and/or controller
- Poor client drivers

This is not a comprehensive list. There are many other technical issues that can cause less-than-optimal performance. With ~30 client devices, with an assortment of 11n and 11ac capabilities, associated to a 4x4:4 11ac Wave-2 AP (dual-band, with a 2.4GHz 11n radio), aggregate throughput in real-world deployments (e.g. 40MHz x \leq 3SS x 64QAM + SGI) may range from 150-200Mbps, best case, aggregated across both radios. The throughput is completely dependent on the client mix, and it's important to understand that one associated 11a, 11b, or 11g can tank an AP's capacity.

Some Wi-Fi vendors have already begun building dual-5GHz 11ac APs, though it's yet unproven that two 5GHz radios can co-exist without significant throughput loss due to ACI. Even if we could solve the ACI problem, such a configuration would also severely limit configurable channels choices. If two (or more) such radios could co-exist without ACI, then we could more effectively use a 1Gbps backhaul. However, we may still only reach around 80% of a 1Gbps link's capacity (e.g. ~400Mbps + ~400Mbps uplink and ~400Mbps + ~400Mbps downlink) during times of peak utilization considering a nearly unrealistically perfect scenario.

In today's market, where a dual-radio AP houses one 5GHz 3x3:3 11ac radio and one 2.4GHz 3x3:3 11n radio, the highest throughput that could be considered would be ~400 Mbps uplink and downlink for 5GHz 11ac plus an additional ~40Mbps uplink and ~40Mbps downlink for 2.4GHz 11n. This comprises less than 50% utilization of a 1Gbps link in the best case. Due to 802.11 contention with multiple clients being connected to the AP, interference sources (modulated and unmodulated) on both bands, and use of 40MHz channels instead of 80MHz, that calculated "440Mbps" of bi-directional throughput can quickly be cut by 50% or more in each direction.

11ac Deployments

If you've chosen to move forward with an 11ac deployment, you will be faced with one of two deployment options: greenfield or upgrade. Either you are installing a Wi-Fi infrastructure for the first time or you are upgrading an existing deployment. Given the pervasive nature of Wi-Fi technology, it's far more likely that you are considering an upgrade. You may also be facing a fiscal decision on whether your existing equipment has reached the end of its useful and/or cost-effective life.

Rip-n-Replace & Rolling Upgrades

When budgetary constraints allow it, "rip-n-replace" deployments can be exciting...so exciting in fact that some network managers miss the significance of moving from a legacy Single Input Single Output (SISO) system (e.g. 11a, 11b, & 11g) to a Multiple Input Multiple Output (MIMO) system (e.g. 11n and 11ac). These two types of systems are very different, and deployment of 11n or 11ac as a replacement to 11a/b/g systems should always entail a new network design, survey, and validation. It's rare to see 11a, 11b, and 11g APs performing at a level that is acceptable to the customer. Since 11ac APs are essentially the same price as 11n APs (for similar specs), then moving from legacy 11a/b/g APs to 11ac APs makes good financial sense. Moving from an 11n system to an 11ac system may also require a new design, survey, and validation. If the 11n design was optimized, it's possible that many of the AP locations may be reused, if the client density and user throughput requirements (due to application requirements) have remained roughly the same. If there has been a significant increase in client density and/or user throughput requirements, then a redesign is still recommended. Radio quality (e.g. receive sensitivity) of 11ac radios is often significantly better than similarly priced 11n radios, so configuration adjustments may have to be made to tune APs to their environments.

The only scenario where 11ac Wave-2 APs would comprehensively replace 11ac Wave-1 APs would be during the changing of Wi-Fi infrastructure vendors (for a variety of possible reasons). When budgetary constraints do not allow you to replace the entire Wi-Fi infrastructure, then a "rolling upgrade" is the answer. Augmenting your 11n network with 11ac or your 11ac Wave-1 network with 11ac Wave-2 can yield a high ROI while saving a significant amount of money. Rolling upgrades result in the customer simultaneously using a variety of 11n and/or 11ac hardware, sometimes from different manufacturers. The best practice is to separate disparate systems (e.g. via building or site) and then to place 11ac APs in locations that have high-density/high-throughput requirements.

Out With The Old

There are very few scenarios in networking where eliminating a device, or device type, can yield a 10-fold performance increase. Removing legacy 11a/b/g clients and APs from the network can produce just such an increase due to removing the need for some MAC layer protection mechanisms. If an end user wants to maximize ROI and to experience the benefits that 11ac can bring, they must aggressively remove legacy clients and APs.

When clients aren't just computers or mobile devices, but instead are appliances with 11a/g/n, like infusion pumps in healthcare, cash registers in retail, and industrial scan guns in warehousing, then 11ac infrastructure gives little appreciable advantage over 11n, assuming that each is well-designed.

If you have an 11n infrastructure, we recommend upgrading as many client devices as possible (from 11a/b/g/n to 11ac) before an infrastructure upgrade to 11ac. Especially important is to get rid of, and discontinue buying, 2.4GHz-only client devices. To this end, it's important to involve purchasing agents in the technology process to help them understand the consequences of cost-cutting in this area.

Consider the upgrade process for a desktop computer, where you have several components (e.g. motherboard, RAM, hard drive, CPU, cooling fan, etc.) to consider. These components are typically appropriate for each other at the time when the computer is built. When you upgrade, it's common to find that upgrading one component may lead to upgrading most of them. It's often the same with Wi-Fi infrastructure systems. If you upgrade your APs from 11a/b/g or 11n (2x2:2) to 11ac Wave-2 4x4:4, the throughput and feature capabilities of the APs may exceed what the controller is capable of handling. Further, the system may now be considered mission critical, whereas previously it was not. For this reason, it makes sense to evaluate all system components during an upgrade (e.g. WNMS, Controller, APs, Sensors, Software Services) rather than just the APs.

Planning and Diagnostic Tools

Wi-Fi network design is an iterative process that begins with a thorough and accurate assessment of customer requirements and constraints. Once these are understood, the most time-efficient next step is predictive modeling. The more effort and accuracy (e.g. wall loss measurements, understanding how to use the software to its fullest extent, etc.) that goes into this part of network planning, the higher the network's performance will be after installation. AirMagnet Planner & Survey Pro products provide wireless engineers with the tools to efficiently and accurately model and assess a network's ability to deliver the necessary coverage, throughput, and mobility requirements. AirMagnet's Planner offer a variety of modeling features, such as map calibration, wall loss, 2D & 3D visualization, multi-floor modeling, attenuation areas, excluded areas, customizable color palette, and a comprehensive set of AP and antenna parameters. Nearly all features are customizable, and a wide variety of map formats are supported, including CAD files.

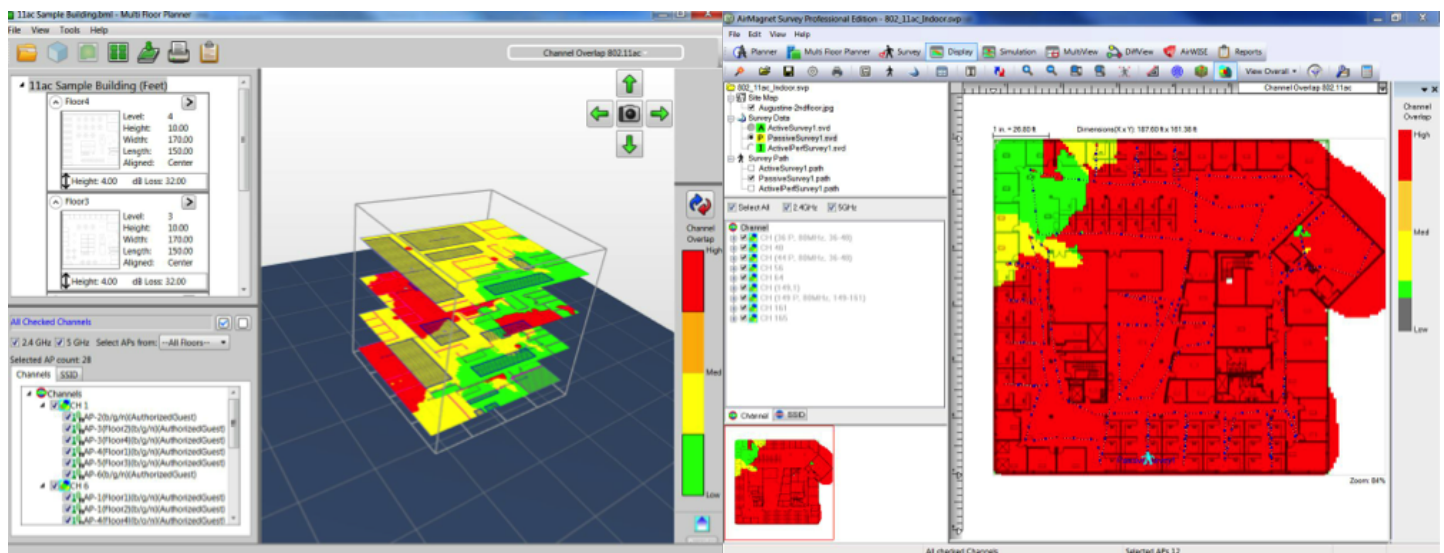


Figure: AirMagnet Planner & Survey PRO tools

When performing a site survey to validate a deployment that needs to support 802.11ac clients, 802.11ac adapters should be used to accurately map and verify those areas that will receive benefit from the newer data rates and channel widths. This is similar to the recommended practice of using smart devices for collecting survey data, when there will be client smart devices in that environment. This will ensure that those areas you have targeted for 802.11ac's benefits are truly seeing the performance improvement.

While an 802.11n-only diagnostic tool can troubleshoot some issues on 11ac networks (most Management and Control frames are sent at 802.11g/n (2.4GHz) or 802.11a/n (5GHz) MCS rates), 802.11ac enabled diagnostic equipment and tools, like the AirMagnet WiFi Analyzer, is necessary to get a complete view of the network and how it is performing. The requirement for an 11ac chipset in diagnostic tools is most often related to the chipset's understanding of the modulation used for 11ac (called Very High Throughput (VHT)) modulated data frames and 80MHz or 160MHz channel widths.

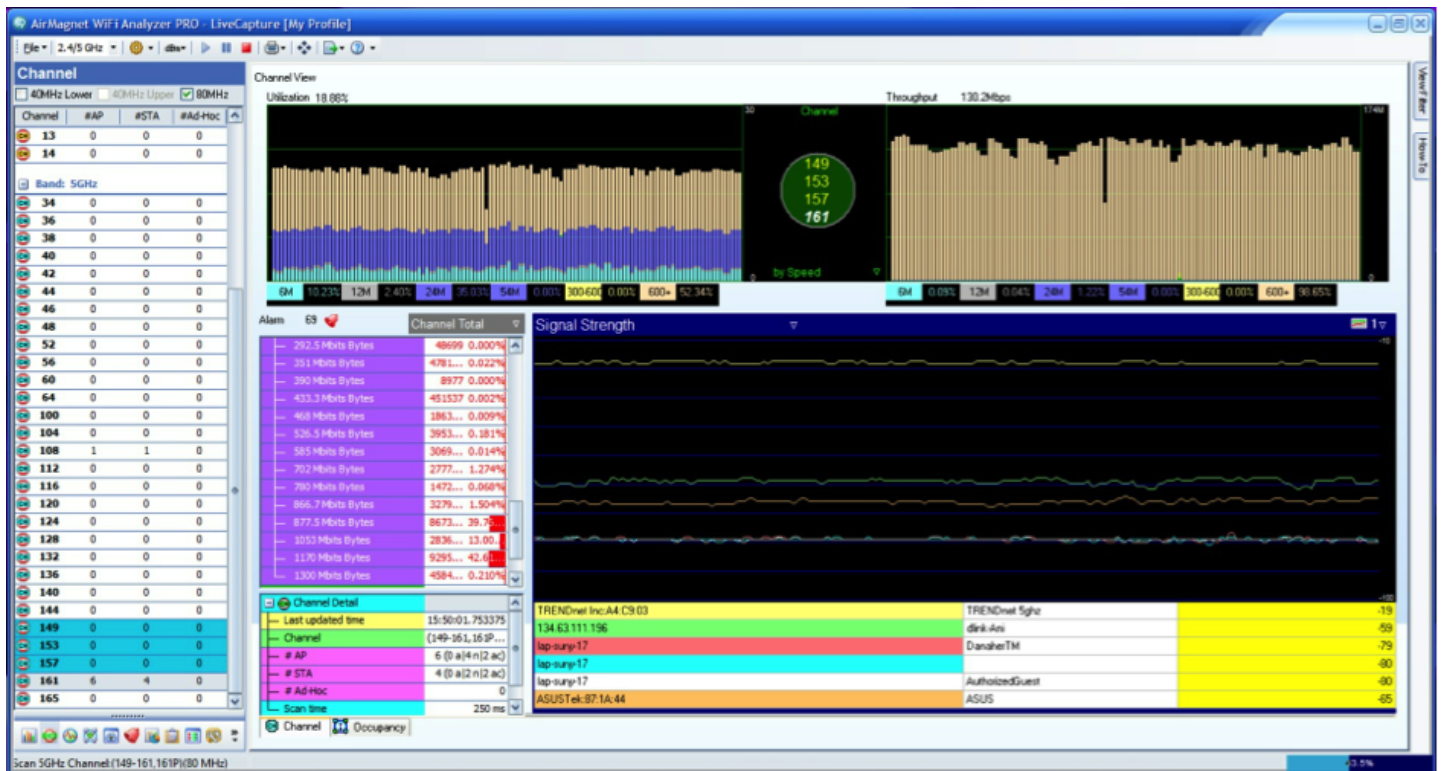


Figure: AirMagnet WiFi Analyzer PRO

At this time, 11ac enabled diagnostic tools are coming into the market. If the diagnostic tool is software-based, rather than a handheld (or similar) appliance, then drivers for 11ac client adapters are now available for many tools on the market. Upgrading to 11ac enabled diagnostics tools, when available, is often money well spent due to 11ac radios' improved quality and receive sensitivity. The better a diagnostic tool can hear, the better it can do its job.

Summary

11ac brings a variety of technical enhancements (as listed below), but it's important to be realistic about its real-world capabilities. Each of these enhancements has the potential to improve communication efficiency in the right environment; however, the marketed benefits of some enhancements may not be realized in some scenarios.

- Wide channels, up to 160MHz
- Downlink MU-MIMO
- 256QAM modulation
- Four spatial streams (4SS)
- Improved rate-over-range
- More powerful hardware
- Higher density handling

In high-density environments, use of 20MHz channels is recommended. MU-MIMO is unproven and provides questionable improvement in downlink throughput due to the protocol overhead and complexity involved. 256QAM may only be useable up to ~50 feet, which may not be helpful in many environments. 4SS is useful only when client devices can support four spatial streams (which excludes mobile devices) and when the environment sufficiently supports spatial stream de-correlation.

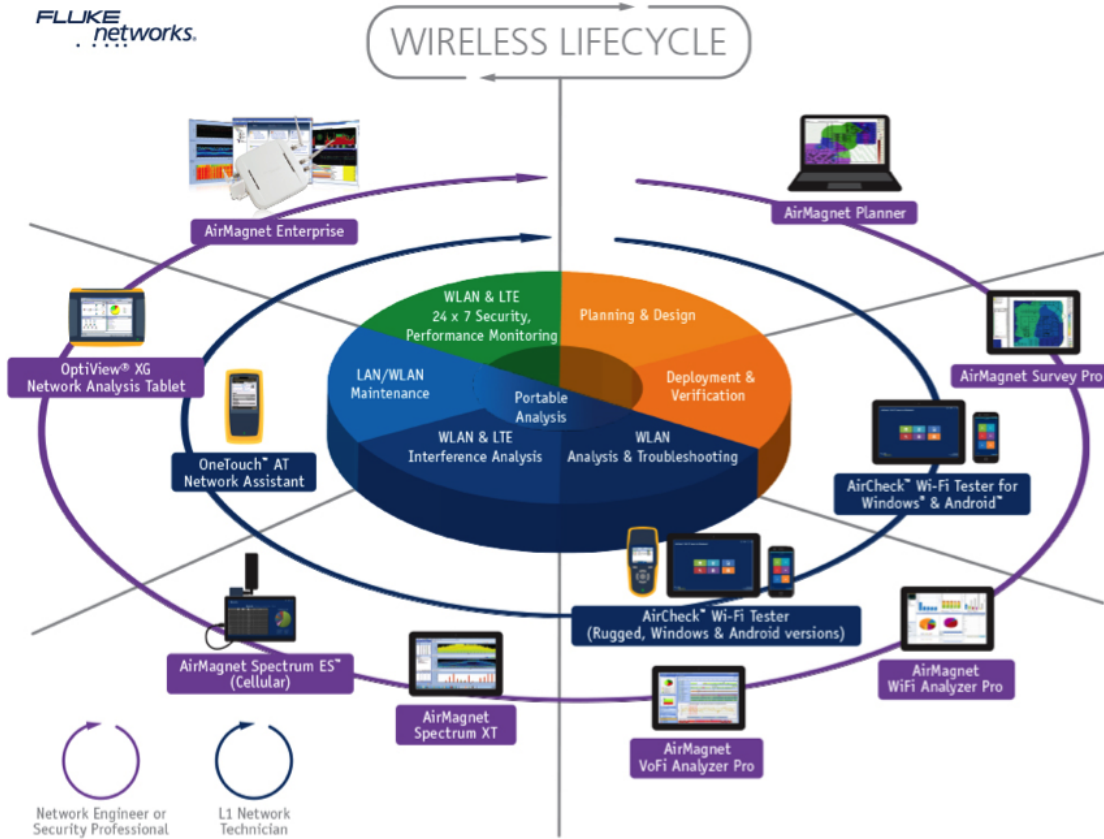
As you consider moving to 11ac, you will be faced with the decision of whether to rip-n-replace (possibly with another vendor) or to upgrade your network in place. Regardless of which approach you take, one key to success is to remember that:

- 11ac and 11a/g are very different technologies, each requiring a different type of network design
- 11ac and 11n use similar technologies, but there have been significant changes in hardware since 11n

In addition to 11ac chipset improvements beyond 11n on 5GHz, dual-radio 11ac APs also have enhanced 2.4GHz 11n chipsets that will be more sensitive, susceptible to interference, and have better rate-over-range.

It is strongly recommended that anyone moving to 11ac consider reducing their use of the 2.4GHz ISM band to the extent possible. Moving away from 2.4GHz will improve user experience, increase network capacity, decrease application failure and support tickets, and significantly reduce TCO of the Wi-Fi infrastructure.

Wi-Fi network design is an iterative process, requiring post-install adjustments and on-going monitoring to achieve maximum ROI. State-of-the-art 802.11ac design, survey, and troubleshooting tools are necessary to assure optimal AP placement and configuration, minimal co-channel interference, maximum performance, and on-going WLAN health. NetScout offers the latest in tools, including the AirMagnet Survey and Planner, Spectrum XT, Wi-Fi Analyzer Pro, and AirCheck Wi-Fi Tester, and is committed to assisting its customers with a smooth migration to 11ac.



Migration Checklist

Consider channel widths

- It is recommended to use 20MHz channels in high-density environments such as auditoriums, ballrooms, trade shows, airports, and arenas because they increase channel use efficiency. Low-density/high-throughput environments, such as open office areas may benefit from 40MHz channels in 5GHz, provided there are enough channels for a reasonable channel reuse plan. If only 1-2 APs will be deployed in a facility (e.g. a branch office), and there is only a minimal amount of interference (modulated and unmodulated), then using 80MHz channels may work well. There is currently no appropriate use for 160MHz channels other than highly directional point-to-point links. If there is a specific area where very high throughput is consistently required, then configuring one AP to use an 80MHz channel in that area may be OK, so long as the nearby APs do not use any part of that 80MHz channel.

Be realistic about access to the highest data rates

- It takes at least 37dB of SNR to achieve MCS9 (the highest MCS rate for 11ac) for an 80MHz channel. This is an unreasonably high SNR, and without significant channel reuse, the co-channel contention (interference) would be very significant.

Evaluate your wired backhaul

- Stated plainly, backhaul links greater than 1Gbps are not needed for 11ac – neither with 11ac Wave-1 nor 11ac Wave-2.

Determine your design methodology

- If there has been a significant increase in client density and/or user throughput requirements, then a redesign is still recommended.
- The best practice is to separate disparate systems (e.g. via building or site) and then to place 11ac APs in locations that have high-density/high-throughput requirements.
- If an end user wants to maximize ROI and to experience the benefits that 11ac can bring, they must aggressively remove legacy clients and APs.

Choose the right tools

- When performing a site survey to validate a deployment that needs to support 802.11ac clients, 802.11ac adapters should be used to accurately map and verify those areas that will receive benefit from the newer data rates and channel widths.
- 802.11ac enabled diagnostic equipment is necessary to get a complete view of the network and how it is performing.