

802.11ac: The Fifth Generation of Wi-Fi

Technical White Paper

August, 2012

Contents

1. Executive Summary	3
2. What Is 802.11ac?	4
2.1 Drivers for 802.11ac.....	4
2.2 How Does 802.11ac Go So Fast?.....	5
2.3 How Do We Make 802.11ac Robust?.....	7
2.3.1 Technology Overview.....	7
2.3.2 Differences Between 802.11ac and 802.11n.....	8
2.3.3 Standards-Based Beamforming.....	9
2.3.4 RTS/CTS with Bandwidth Indication.....	10
2.3.5 All A-MPDUs.....	12
2.3.6 Channelization and 80+80 MHz.....	12
2.3.7 Rate at Range.....	15
2.3.8 Regulatory.....	16
2.3.9 MU-MIMO.....	17
2.3.10 802.11ac Project Authorization Request.....	18
3. When Is 11ac Happening?	18
4. How Does 11ac Affect Me?	19
4.1 Compatibility.....	19
4.2 When to Upgrade to 802.11ac?.....	20
4.3 Radio Resource Management and WIPS Effects.....	21
5. Summary	21
Appendix: What Is 802.11n?	21

1. Executive Summary

802.11ac, the emerging standard from the IEEE, is like the movie **The Godfather Part II**. It takes something great and makes it even better. 802.11ac is a faster and more scalable version of 802.11n. 802.11ac couples the freedom of wireless with the capabilities of Gigabit Ethernet.

Wireless LAN sites will see significant improvements in the number of clients supported by an access point (AP), a better experience for each client, and more available bandwidth for a higher number of parallel video streams. Even when the network is not fully loaded, users see a benefit: their file downloads and email sync happen at low-lag gigabit speeds. Also, device battery life is extended, since the device's Wi-Fi interface can wake up, exchange data with its AP, then revert to dozing that much more quickly.

802.11ac achieves its raw speed increase by pushing on three different dimensions:

- More channel bonding, increased from the maximum of 40 MHz in 802.11n, and now up to 80 or even 160 MHz (for 117% or 333% speed-ups, respectively)
- Denser modulation, now using 256 quadrature amplitude modulation (QAM), up from 802.11n's 64QAM (for a 33% speed burst at shorter, yet still usable, ranges)
- More multiple input, multiple output (MIMO). Whereas 802.11n stopped at four spatial streams, 802.11ac goes all the way to eight (for another 100% speed-up).

The design constraints and economics that kept 802.11n products at one, two, or three spatial streams haven't changed much for 802.11ac, so we can expect the same kind of product availability, with first-wave 802.11ac products built around 80 MHz and delivering up to 433 Mbps (low end), 867 Mbps (midtier), or 1300 Mbps (high end) at the physical layer. Second-generation products promise still more channel bonding and spatial streams, with plausible product configurations operating at up to 3.47 Gbps.

802.11ac is a 5 GHz-only technology, so dual-band APs and clients will continue to use 802.11n at 2.4 GHz. However, 802.11ac clients operate in the less crowded 5 GHz band.

Second-generation products should also come with a new technology, multiuser MIMO (MU-MIMO). Whereas 802.11n is like an Ethernet hub that can only transfer a single frame at a time to all its ports, MU-MIMO allows an AP to send multiple frames to multiple clients at the same time over the same frequency spectrum. That's right: with multiple antennas and smarts, an AP can behave like a wireless switch. There are technical constraints, and so MU-MIMO is particularly well suited to bring-your-own-device (BYOD) situations where the devices such as smartphones and tablets might only have a single antenna.

802.11ac-enabled products are the culmination of efforts at the IEEE and Wi-Fi Alliance pipelines. IEEE 802.11ac delivered an approved Draft 2.0 amendment in January 2012 and a refined Draft 3.0 in May 2012, with final ratification planned for the end of 2013. In parallel, the Wi-Fi Alliance is expected to take an early IEEE draft, most likely Draft 3.0, and use that as the baseline for an interoperability certification of first-wave products in early 2013. Later, and more in line with the ratification date of 802.11ac (that is, after December 2013), the Wi-Fi Alliance is expected to refresh its 802.11ac certification to include testing of the more advanced 802.11ac features. This second-wave certification should include features such as channel bonding up to 160 MHz, four spatial streams, and MU-MIMO. Overall, this arrangement closely follows how 802.11n was rolled out.

Enterprise networks considering an investment in infrastructure Wi-Fi have two excellent choices: (1) buy 802.11n APs, since they deliver a remarkable level of performance, they are available today, and 802.11n is widely deployed in client products, or (2) wait for 802.11ac APs and their state-of-the-art performance. Then there is option (3), which avoids the wait: invest in a modular 802.11n AP such as the [Cisco® Aironet® 3600 Access Point](#) that is readily field-upgradable to 802.11ac.

Meanwhile, consumer-grade products might be launched earlier, in 2012, but before Wi-Fi Alliance certification is available. However, we **strongly** recommend that all enterprises wait for products that are certified by the Wi-Fi Alliance as interoperable (that is, in 2013).

802.11ac, when it arrives, has a few new effects on existing 802.11a/11n deployments even if the deployment is not upgraded to 802.11ac immediately: (1) the wider channel bandwidths of neighboring APs require updates to radio resource measurement, or RRM (and in particular the dynamic channel assignment algorithm), and (2) 802.11a/11n wireless intrusion protection systems (WIPS) can continue to decode most managements frames such as beacon and probe request/response frames (that are invariably sent in 802.11a format) but do not have visibility in data sent in the new 802.11ac packet format.

One thing not to worry about is compatibility. 802.11ac is designed in a deep way to coexist efficiently with existing 802.11a/n devices, with strong carrier sense, a single new preamble that appears to be a valid 802.11a preamble to 802.11a/n devices, and extensions to request-to-send/clear-to-send (RTS/CTS) to help avoid collisions with users operating on slightly different channels.

2. What Is 802.11ac?

First, 802.11ac is an evolution of 802.11n. If you want to learn more about 802.11n, jump to the appendix. If you are already familiar with the channel bonding, MIMO, and aggregation introduced by 802.11n and you don't need a refresher, then read on.

2.1 Drivers for 802.11ac

802.11ac is an evolutionary improvement to 802.11n. One of the goals of 802.11ac is to deliver higher levels of performance that are commensurate with Gigabit Ethernet networking:

- Seemingly “instantaneous” data transfer experience
- A pipe fat enough that delivering high quality of experience (QoE) is straightforward

In the consumer space, the target is multiple channels of high-definition content delivered to all areas of the house. The enterprise has different challenges:


- Delivering network with enterprise-class speeds and latencies
- High-density environments with scores of clients per AP
 - Which are exacerbated by the [BYOD](#) trend such that one employee might carry two or even three 802.11 devices and have them consuming network resources all at once
- The increased adoption of video streaming

802.11ac is about delivering an outstanding experience to each and every client served by an AP, even under demanding loads.

Meanwhile 802.11 is integral to hugely broad range of devices, and some of them are highly cost, power, or volume constrained. One antenna is routine for these devices, yet 802.11ac must still deliver peak efficiency.

The one thing that 802.11ac has in its favor is the evolutionary improvement to silicon technology over the past half-dozen years: channel bandwidths can be wider, constellations can be denser, and APs can integrate more functionality.

Figure 1. How 802.11ac Accelerates 802.11n



2.2 How Does 802.11ac Go So Fast?

Wireless speed is the product of three factors: channel bandwidth, constellation density, and number of spatial streams. 802.11ac pushes hard on the boundaries on each of these, as shown in [Figure 1](#).

For the mathematically inclined, the physical layer speed of 802.11ac is calculated according to [Table 1](#). For instance, an 80 MHz transmission sent at 256QAM with three spatial streams and a short guard interval delivers $234 \times 3 \times 5/6 \times 8 \text{ bits} / 3.6 \mu\text{s} = 1300 \text{ Mbps}$.

Table 1. Calculating the Speed of 802.11n and 802.11ac

PHY	Bandwidth (as Number of Data Subcarriers)	Number of Spatial Streams	Data Bits per Subcarrier	Time per OFDM Symbol	PHY Data Rate (bps)
11n or 11ac	56 (20 MHz)	1 to 4	Up to $5/6 \times \log_2(64) = 5$	3.6 μs (short guard interval)	=
	108 (40 MHz)			4 μs (long guard interval)	
11ac only	234 (80 MHz)	5 to 8	Up to $5/6 \times \log_2(256) \approx 6.67$		
	2x234 (160 MHz)				

Immediately we see that increasing the channel bandwidth to 80 MHz yields a 2.16 times speed-up, and 160 MHz offers a further doubling. Nothing is for free: it does consume more spectrum, and each time we're splitting the same transmit power over twice as many subcarriers, so the speed doubles, but the range for that doubled speed is slightly reduced (for an overall win).

Going from 64QAM to 256QAM also helps, by another $8/6 = 1.33$ times. Being closer together, the constellation points are more sensitive to noise, so 256QAM helps most at shorter range where 64QAM is already reliable. Still, 256QAM doesn't require more spectrum or more antennas than 64QAM.

Then the speed is directly proportional to the number of spatial streams. More spatial streams require more antennas, RF connectors, and RF chains at transmitter and receiver. The antennas should be spaced 1/3 wavelength (3/4") or more apart, and the additional RF chains consume additional power. This drives many mobile devices to limit the number of antennas to one, two, or three.

Collectively, these three speed-ups are significant. From [Figure 2](#) and [Table 2](#), the minimum allowed 802.11ac product is 4.4x faster than the corresponding 802.11n product, and the midtier and high-end wave 1 products are nearly 3x faster, reaching up to 1.3 Gbps PHY data rates. Actual throughput will be a function of MAC efficiency (rarely better than 70%) and the device capabilities at each end of the link.

Figure 2. Evolution of Cisco APs with 802.11 Physical Layer Amendments

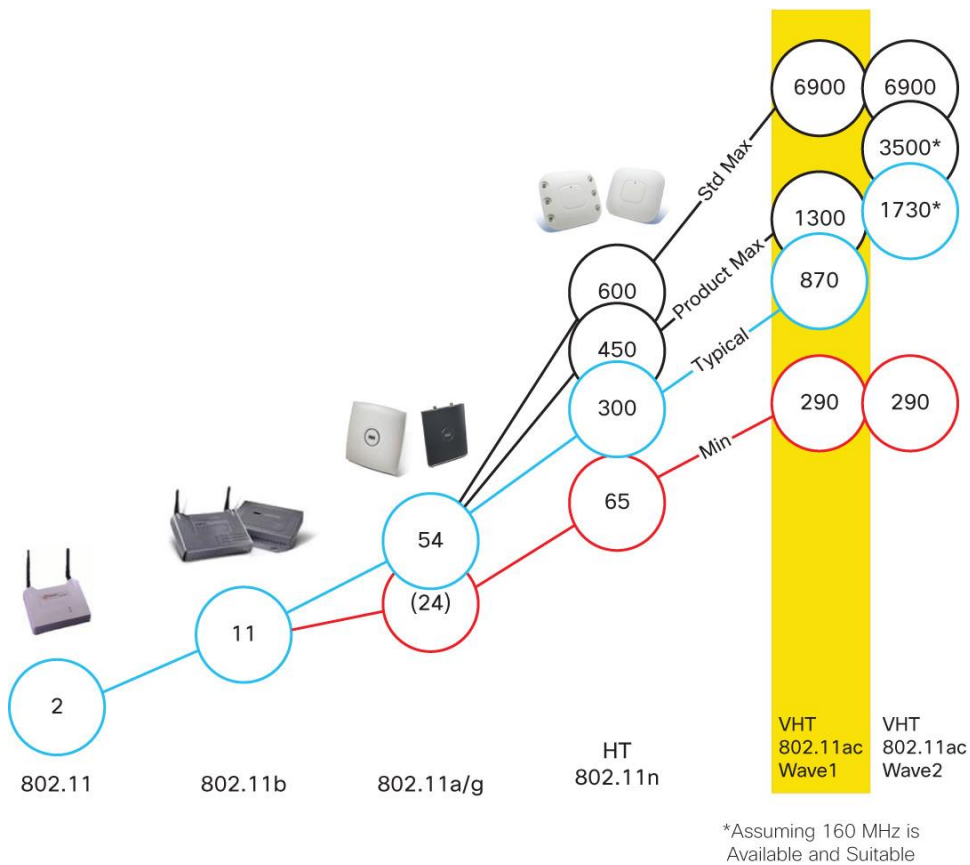


Table 2. Important Data Rates of 802.11a, 11n, and 11ac

Nominal Configuration	Bandwidth (MHz)	Number of Spatial Streams	Constellation Size and Rate	Guard Interval	PHY Data Rate (Mbps)	Throughput (Mbps)*
802.11a						
All	20	1	64QAMr3/4	Long	54	24
802.11n						
Min	20	1	64QAMr5/6	Long	65	46
Low-end product (2.4 GHz only+)	20	1	64QAMr5/6	Short	72	51
Midtier product	40	2	64QAMr5/6	Short	300	210
Max product	40	3	64QAMr5/6	Short	450	320

Nominal Configuration	Bandwidth (MHz)	Number of Spatial Streams	Constellation Size and Rate	Guard Interval	PHY Data Rate (Mbps)	Throughput (Mbps)*
Amendment max	40	4	64QAMr5/6	Short	600	420
802.11ac wave 1						
Min	80	1	64QAMr5/6	Long	293	210
Low-end product	80	1	256QAMr5/6	Short	433	300
Midtier product	80	2	256QAMr5/6	Short	867	610
High-end product	80	3	256QAMr5/6	Short	1300	910
80 MHz amendment max	80	8	256QAMr5/6	Short	3470	2400
802.11ac wave 2						
Low-end product	160	1	256QAMr5/6	Short	867	610
Midtier product	160	2	256QAMr5/6	Short	1730	1200
High-end product	160	3	256QAMr5/6	Short	2600	1800
Ultra-high-end product	160	4	256QAMr5/6	Short	3470	2400
Amendment max	160	8	256QAMr5/6	Short	6930	4900
*Assuming a 70% efficient MAC, except for 802.11a, which lacks aggregation.						
*Assuming 40 MHz is not available due to the presence of other APs.						

2.3 How Do We Make 802.11ac Robust?

The sticker on the box that shows maximum data rate doesn't help us much in the real world, where devices have to contend with interference from non-802.11 devices, preexisting APs that might only use 20 or 40 MHz, multipath fading, few antennas on mobile devices, weak signals at range, and so forth. What makes the raw speed of 802.11ac so valuable are the extensions that help to deliver reliable throughput under realistic conditions.

2.3.1 Technology Overview

By design, 802.11ac is intended to operate only in the 5 GHz band, as shown in [Table 3](#). This avoids much of the interference at 2.4 GHz, including Bluetooth headsets and microwave ovens, and provides a strong incentive for users to upgrade their mobile devices (and hotspot APs) to be dual band so that the 5 GHz band is more universally usable. This choice also streamlines the IEEE process by avoiding the possibility of contention between 802.11 and [802.15](#) proponents. And there is barely 80 MHz of bandwidth at 2.4 GHz anyway.

As we've already seen, 802.11 introduces higher order modulation, up to 256QAM; additional channel bonding, up to 80 or 160 MHz; and more spatial streams, up to eight. There is an alternative way to send a 160 MHz signal, known as "80+80" MHz, discussed later (see [section 2.3.6](#)).

802.11ac continues some of the more valuable features of 802.11n, including the option of short guard interval (for a 10% speed bump) and incrementally better rate at range using the advanced low-density parity check (LDPC) forward error-correcting codes. These LDPC codes are designed to be an evolutionary extension of the 802.11n LDPC codes, so implementers can readily extend their current hardware designs.

Various space time block codes (STBCs) are allowed as options, but (1) this list is trimmed from the overrich set defined by 802.11n, and (2) STBC is largely made redundant by beamforming. 802.11n defined the core STBC modes of 2x1 and 4x2 and also 3x2 and 4x3 as extension modes, but the extension modes offered little gain for their additional complexity and have not made it to products. Indeed, only the most basic mode, 2x1, has been certified by the Wi-Fi Alliance. With this experience, 802.11ac only defines the core 2x1, 4x2, 6x3, and 8x4 STBC modes, but again only 2x1 is expected to make it to products: if you had an AP with four antennas, why would you be satisfied with 4x2 STBC when you could - and should - be using beamforming?

What 802.11ac also gets right is to define a single way of performing channel sounding for beamforming: so-called explicit compressed feedback. Although optional, if an implementer wants to offer the benefits of standards-based beamforming, there is no choice but to select that single mechanism, which can then be tested for interoperability.

Table 3. Primary Ingredients of 802.11ac

Parameter	802.11ac D3.0 Likely Wave 1 Wi-Fi Alliance Certification	802.11ac (Later Draft) Potential Wave 2 Wi-Fi Alliance Certification	802.11ac Complete Amendment
Spectrum	5 GHz (varied support by regulatory domain; nearly 600 MHz in the United States)		<6 GHz excluding 2.4 GHz
Bandwidth	Mandatory: 20, 40, and 80 MHz	Mandatory: 20, 40, and 80 MHz Optional: 160 and 80+80 MHz	
Modulation	Mandatory: BPSK, QPSK, 16QAM, 64QAM Optional: 256QAM		
Number of spatial streams	Mandatory: 2 (non mobile APs*), 1 (others) Optional: up to 3 spatial streams	Mandatory: 2 (non mobile APs*), 1 (others) Optional: up to 4 spatial streams	Mandatory: 1 Optional: 2-8
Forward error correction	Mandatory: BCC Optional: LDPC		
STBC	Optional: 2x1 AP to client		Optional: 2x1, 4x2, 6x3, 8x4
Short guard interval	Optional		
Sounding (a single interoperable protocol)	Optional		
CTS in response to RTS with bandwidth indication	Mandatory		
RTS with bandwidth indication	Optional		
Aggregation	Mandatory: TX and RX of A-MPDU Optional: RX A-MPDU of A-MSDU	Mandatory: TX and RX of A-MPDU TBD: RX A-MPDU of A-MSDU	A-MDPU, A-MDPU of A-MSDU
MU-MIMO	-	Optional	

*Additional requirement introduced by the Wi-Fi Alliance.

Because of the wider channel bandwidths of 802.11ac, it is much more likely that an 80 MHz AP will overlap with another 20 or 40 MHz AP - and similarly an 80 or 160 MHz AP - or even several of them, all potentially on different channels. To enable reliable operation amid this complexity, 802.11ac mandates extensions to the RTS/CTS mechanism, stronger clear-channel assessment (CCA) requirements, and new primary channel selection rules. See [section 2.3.4](#).

802.11ac also introduces a valuable new technology called multiuser MIMO. This is challenging to get right, so is deferred until the second wave of 802.11ac products and is optional. More on this later in [section 2.3.9](#).

2.3.2 Differences Between 802.11ac and 802.11n

802.11ac has avoided the battles of 802.11n and instead has focused on extending the tremendous advances made in 802.11n to deliver the next generation of speed and robustness.

For instance, 802.11n pioneered aggregation through the selective use of A-MPDU, A-MSDU, and A-MPDU of A-MSDU (see Appendix). 802.11ac actually **requires** every 802.11ac transmission to be sent as an A-MPDU aggregate. This is due in part because of the intrinsic efficiency of A-MPDU and for some other reasons too (see [section 2.3.5](#)).

In a further example, 802.11ac extends the 802.11n channel access mechanism: virtual carrier sense and backoff occur on a single 20 MHz primary channel; then CCA is used for the remaining 20 MHz subchannels immediately before transmitting on them.

Given the power of A-MPDU and the 802.11n channel access mechanism, 802.11ac actually didn't need to innovate much in the MAC. Indeed, extensions to the RTS/CTS mechanism are the only new mandatory MAC feature.

802.11n does include many options with reduced value. 802.11ac takes a very pragmatic approach to them. If a "useless" option is used and affects a third-party device, then typically 802.11ac forbids an 802.11ac device (operating in 802.11ac mode) from using the option. If a "useless" option has not been used in 802.11n products or only affects the devices that activate the option, then the feature is not updated for 802.11ac but is instead "left to die."

For instance, there is no 802.11ac version of the "802.11n greenfield" preamble format. 802.11ac only defines one preamble format, which, to legacy 802.11a/11n devices, will look safely like an 802.11a preamble followed by a payload with a bad CRC. This means that legacy devices don't try to transmit over the top of the 802.11ac transmission, nor do they attempt to send a bad payload up the stack.

802.11n introduced a "reduced interframe spacing," which reduces overheads between consecutive transmissions, but experience has shown that A-MPDU solves much the same problem, but even more efficiently. 802.11ac devices operating in 802.11ac mode are not permitted to transmit RIFS (as of Draft 3.0).

802.11n features that are not updated for 11ac (or explicitly forbidden for 802.11ac devices operating in 802.11ac mode) include all the 802.11n sounding options, including extension LTFs, the calibration procedure, antenna selection, PCO, LSIG TXOP protection, unequal modulation, 4x3 and 3x2 STBC modes, MCS32, and dual CTS protection. If you don't know these terms, then no problem, because almost certainly you'll never need to know them.

2.3.3 Standards-Based Beamforming

Any device (with multiple antennas) can beamform to any other device at any time. What 802.11ac adds is the opportunity for the receiver to help the beamforming transmitter to do a better job of beamforming. This is called "sounding," and it enables the beamformer to precisely steer its transmitted energy toward the receiver. 802.11ac defines a single, though optional, protocol for one 802.11ac device to sound other 802.11ac devices. The protocol selected closely follows the 802.11n "explicit compressed feedback" protocol, as follows.

A device, typically an AP, sends a "VHT Null Data Packet (NDP) Announcement" frame. Its only purpose is to contain the address of the AP and of the target recipients. The VHT NDP Announcement frame is immediately followed by a "VHT Null Data Packet" (VHT NDP) intended for those target recipients. Each intended recipient measures the RF channel from AP to itself using the preamble of the VHT NDP and compresses the channel. The first intended recipient responds with the compressed channel information in a VHT Compressed Beamforming frame immediately, and other recipients respond when they are polled by the AP. The VHT NDP Announcement frame, the VHT NDP, and the VHT Compressed Beamforming frame are all similar to features in 802.11n. However, because of some subtle differences, the 802.11ac sounding is not backward compatible with 802.11n devices.

Also, to support this new MU-MIMO feature (see [section 2.3.9](#)), the channel feedback can contain an extra level of detail.

Explicit compressed feedback (ECFB) is known to provide the most precise estimate of the channel that takes into account all the imperfections at transmitter and receiver.

However, ECFB comes with a lot of overhead: the VHT NDP Announcement frame, the VHT NDP itself, and the frame carrying the compressed feedback. For an AP with four antennas, the compressed feedback varies from 180 to 1800 bytes, depending on the number of client antennas and level of compression. Sounding just one single-antenna 80 MHz client takes about 250 microseconds. When devices can transmit at 433 Mbps, this is expensive, since that same time could have instead been used to send an extra 13,000 bytes.

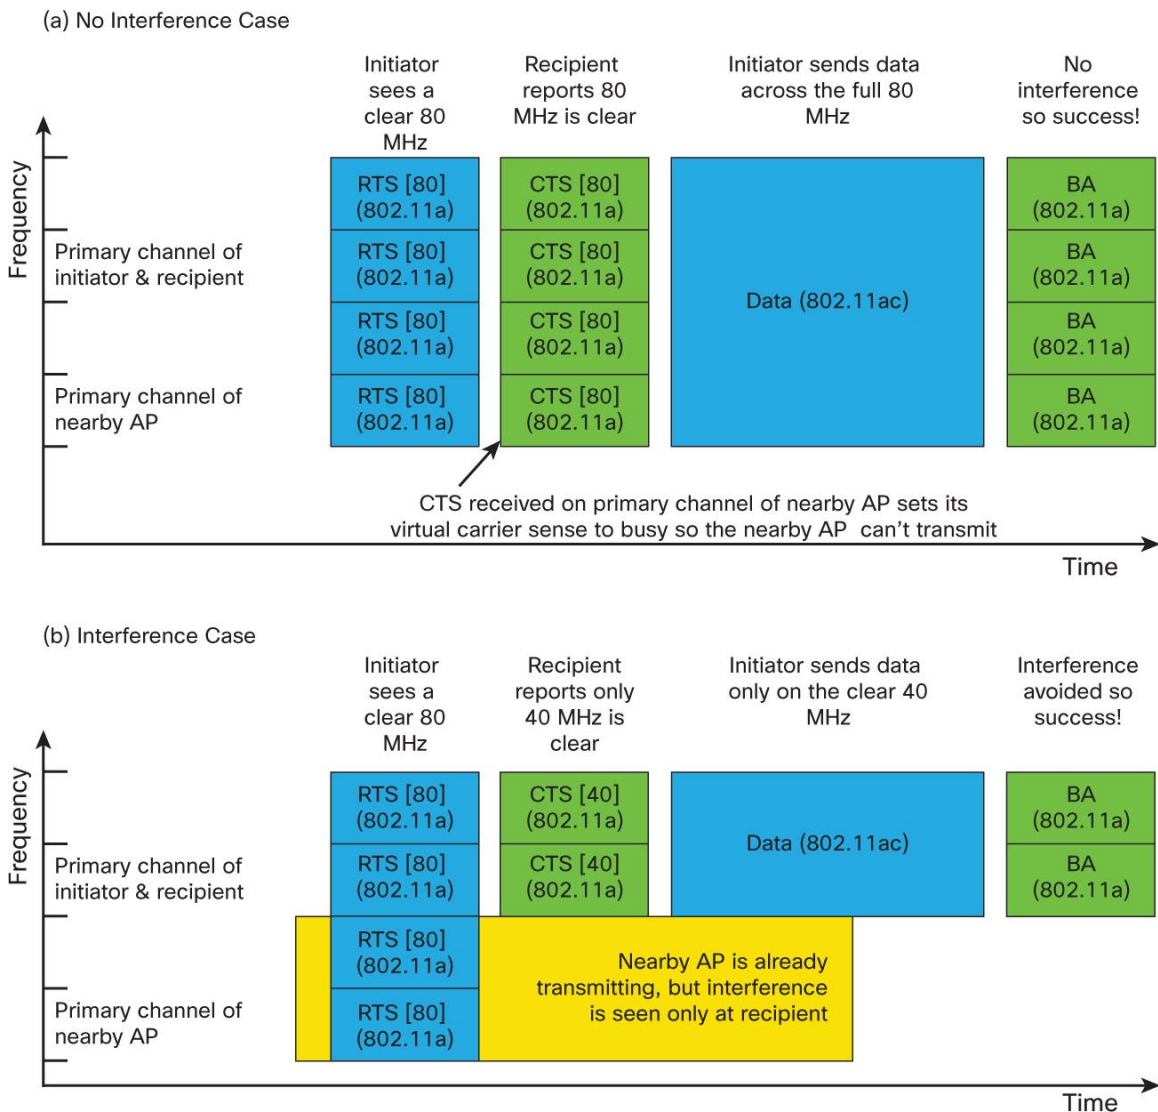
And so technologies that solve the problem of sounding without depending on client assistance (such as Cisco's family of ClientLink technologies) continue to add genuine value. They (1) still help legacy 802.11a/n clients, (2) still help those 802.11ac clients that do not support 802.11ac sounding, (3) still help clients at 2.4 GHz, and (4) can avoid the overhead of standards-based explicit sounding when it is not actually necessary.

2.3.4 RTS/CTS with Bandwidth Indication

An 802.11ac AP operating on 80 MHz (or 160 MHz and so on) should still be capable of allowing 802.11a or 802.11n clients to associate. Thus beacons are sent on one 20 MHz channel, known as the primary channel, within that 80 MHz. The AP and all clients associated to the AP receive and process every transmission that overlaps this primary channel and extract virtual carrier sense from the frames they can decode.

However, the AP could be nearby other uncoordinated APs. Those APs could be preexisting 802.11a or 802.11n APs, and their primary channels could be any 20 MHz within the 80 MHz of the 802.11ac AP. Then the different APs and their associated clients have a different virtual carrier sense, so can transmit at different times on the different subchannels, including overlapping times. With the wide 802.11ac channel bandwidths, this scenario becomes much more likely than with 802.11n.

Figure 3. RTS/CTS Enhanced with Bandwidth Signalling



For this reason, 802.11ac defines an enhanced RTS/CTS protocol. RTS/CTS can be used to find when channel bandwidth is clear and how much, around both the initiator and the responder, as shown in [Figure 3](#).

First, when an 802.11ac device sends an RTS, (1) this initiating device has to verify that the 80 MHz channel is clear in its vicinity, (2) the RTS is normally sent in an 802.11a PPDU format, and (3) the basic 802.11a transmission, which is 20 MHz wide, is replicated another three times to fill the 80 MHz (or another seven times to fill 160 MHz). Then every nearby device, regardless of whether it is an 802.11a/n/ac device, receives an RTS that the device can understand on its primary channel. And every device that hears the RTS has its virtual carrier sense set to busy (see [Figure 3\(a\)](#)). To make the protocol robust, the replication bandwidth of the RTS is reported inside the 802.11a PPDU¹.

¹ Since the 802.11a PPDU format doesn't contain a bandwidth indication, 802.11ac has to play some tricks to maintain backward compatibility. The bandwidth indication is encoded in the scrambling sequence, and also the individual/group bit in the transmitter MAC address in the RTS frame is changed from "individual" to "group." This last change will be visible in sniffer traces.

Second, before the device addressed by the replicated RTS responds with a CTS, the recipient device checks to see if there is anyone transmitting near itself, on its primary channel or on **any other** 20 MHz within the 80 MHz. If a portion of the bandwidth is in use nearby, then the recipient only responds with a CTS on the available and “usable” 20 MHz subchannels and also reports the bandwidth of the replicated CTS inside the CTS’s PPDU. Here “usable” subchannels means the subchannels on which the initiating device is allowed to send something, such as a 20, 40, or 80 MHz (but not 60 MHz) transmission. This is shown in [Figure 3\(b\)](#).

Third, the CTS is sent, like the RTS, in an 802.11a PPDU format, replicated in 20 MHz chunks across the available and useful bandwidth. Again every nearby device receives a CTS that the device can understand on its primary channel.

There are other variations on this protocol, for when the initiator is incapable of switching to a narrower bandwidth on the fly and so forth, but the previous description captures the essence of the enhancement: the recipient can say “these subchannels are busy - don’t use them.”

2.3.5 All A-MPDUs

802.11 defines that every 802.11 PPDU transmission is an A-MPDU, yet the A-MPDU might only contain a single MPDU. Why? The short answer is that it’s complicated.


Long answer: There are three reasons: (1) In 802.11a/n, the duration of the transmission is set by the number of octets and the data rate for the transmission. But for a maximum-length 5.5 ms transmission at 6.93 Gbps, there could be over 4 million bytes, and this takes 23 bits to represent. These bits would be sent at the lowest MCS rate at the start of **every** 802.11ac transmission and so practically would add 4 μ s each time. Instead, the length of an 802.11ac transmission is constrained to be a multiple of the number of data bits per OFDM symbol, and then only the number of OFDM symbols needs to be signaled. Moreover, the number of (assumed to be) 4 μ s-long OFDM symbols is already implicitly available in the legacy portion of the preamble, so this signaling comes almost “for free.”² Then we need a way to completely fill even the last OFDM symbol with data. A-MDPU makes this easy: send the data as MDPU within MDPU subframes in an A-MDPU, then pad the A-MDPU with enough null MDPU subframes to fill up the last OFDM symbol. (2) This same padding mechanism will come in handy for the new MU-MIMO feature. (3) A-MDPU is in general a good idea to increase reliability for long payloads.

2.3.6 Channelization and 80+80 MHz

802.11ac adopts a keep-it-simple approach to channelization. Adjacent 20 MHz subchannels are grouped into pairs to make 40 MHz channels, adjacent 40 MHz subchannels are grouped into pairs to make 80 MHz channels, and adjacent 80 MHz subchannels are grouped into pairs to make the optional 160 MHz channels, as shown in [Figure 4](#). A BSS (that is, AP plus clients) uses the different bandwidths for different purposes, but the usage is principally governed by the capabilities of the clients.

² Just a single bit is needed to disambiguate the number of actual OFDM symbols present if the transmission instead uses the short guard interval and the OFDM symbols are actually 3.6 μ s long.


Figure 4. 802.11ac Channelization (United States)



We see that in the United States, there are 20 to 25 20 MHz channels, 8 to 12 40 MHz channels, 4 to 6 80 MHz channels, and 1 or 2 160 MHz channels. These numbers are ranges because of the evolving regulatory issues surrounding the different spectrum noted in [Figure 4](#).

What if most clients at a deployment are still 802.11n clients with 40 MHz maximum? Does deploying 802.11ac APs mean fewer channels and more interference? As you would expect from an IEEE standard, the answer is a resounding “no.” It is entirely allowed for two 80 MHz 802.11ac APs to select the same 80 MHz channel bandwidth but one AP to put its primary 20 MHz channel within the lower 40 MHz and the other AP to put its primary 20 MHz channel within the upper 40 MHz. What this means is that 802.11n clients associated to the first AP can transmit 20 or 40 MHz as usual, at the same time as 802.11n clients associated to the second AP can transmit 20 or 40 MHz in parallel. What is new in 802.11ac is the ability for any 802.11ac client that sees that the whole 80 MHz as available to invoke a very high-speed mode and to transmit across the whole 80 MHz. This is shown in [Figure 5](#).

Figure 5. Example of Parallel Transmissions with Two BSSs on the Same 80 MHz but with Different Primary 20-MHz Subchannels



The ability to have overlapped APs but different primary channels is made possible by:

- The enhanced secondary CCA thresholds mandated by 802.11ac, which are up to 13 dB more stringent than the secondary channel CCA thresholds defined by 802.11n
- The addition of a bandwidth indication to the RTS/CTS exchange (see [section 2.3.4](#))

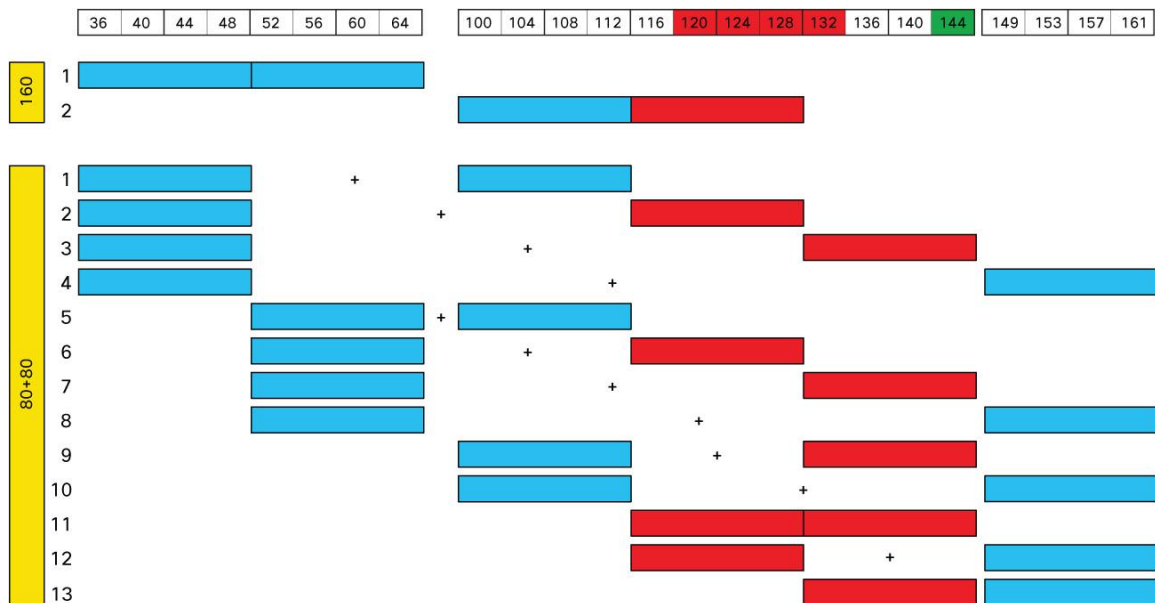
Over time, clients will transition from 802.11n to 802.11ac, so that 80 MHz is used more and more. In this environment APs should change to align their primary 20 MHz channels.

The capability of 80 MHz channels is markedly increased over narrower bandwidths. And this offers a lot of value in many typical scenarios: a few clients transferring a lot of traffic associated to a 40 MHz AP are speed limited to 802.11n's 300 or 450 Mbps. This is true even if the APs on the adjacent 40 MHz are all lightly loaded. With the wider channel, more clients get to transfer their data more quickly, and they can complete their transmissions that much sooner. Overall, less battery energy is consumed, and other clients don't have as long to wait (for better QoS). This discussion all comes under the umbrella of "statistical multiplexing," where more multiplexing is more efficient for bursty traffic.

Since the number of 160 MHz channels is tiny, 160 MHz is unsuited to typical enterprise use. In the home, every 160 MHz channel is subject to difficult radar detection regulatory requirements. Thus 802.11ac also introduces a noncontiguous 80+80 MHz mode. As can easily be imagined from the name, it is the 160 MHz waveform, but transmitted in two separate 80 MHz frequency segments, where each 80 MHz segment can lie on any allowed 80 MHz channel. To make this feasible, it is still a time-division-duplex system, in that APs and clients only ever transmit on 80+80 or receive on 80+80; they are never expected to transmit on one 80 MHz segment and receive on the second 80 MHz segment.

In lightly and moderately used spectrum, this appears to provide vastly more flexibility to avoid interference, as shown in [Figure 6](#). 80+80 MHz has 13 options versus the two options for 160 MHz (ignoring regulatory issues). Unfortunately, an 80+80 MHz device is much more complicated than a 160 MHz device, since the 80+80 MHz device needs twice as many RF chains. A device might operate either as a two-spatial stream 80 MHz device or as a one-spatial stream 80+80 MHz device. In this case, 80+80 MHz allows the use of more spectrum but only uses that spectrum half as efficiently.

Figure 6. Channel Options for 160 MHz and 80+80 MHz



Neither 160 MHz nor 80+80 MHz are recommended for typical enterprises given the currently available spectrum.

From [Figure 4](#), RRM becomes a much more complicated task. It must:

- Avoid channels with radar (if present).
- Uniformly spread the channel bandwidth used by each AP and preferably spread the AP's primary 20 MHz channel too.
- Tend to avoid a channel that overlaps with other 20, 40, 80, 160, or 80+80 MHz APs nearby.
- Within an 80 MHz channel bandwidth (for example), decide whether to align primary 20 MHz channels with other APs or deliberately not to align primary channels. This is not a clear-cut choice:
 - If the primary channels are aligned, then virtual carrier sense works completely, yet all 20/40 MHz traffic (including broadcast, multicast, and data traffic to 802.11a/n devices) is sent in series. During these times, 40 or 60 MHz of bandwidth is unused. Still, if the clients are predominantly 802.11ac, then this is generally the best approach in terms of throughput and airtime fairness.
 - Conversely, if one AP's primary channel is assigned to the lower 40 MHz and another AP's channel is assigned to the upper 80 MHz, then 20/40 MHz traffic can be parallelized (as per [Figure 5](#)). If clients are predominantly 802.11a/802.11n, then this is the better choice. And when the whole 80 MHz is free, as measured by physical carrier sense and/or RTS/CTS with bandwidth indication, then 80 MHz communication between 802.11ac devices is still allowed.

Certainly it is difficult to get the most out of 802.11ac without coordination of AP channel assignment, typically under the aegis of an effective centralized RRM algorithm.

2.3.7 Rate at Range

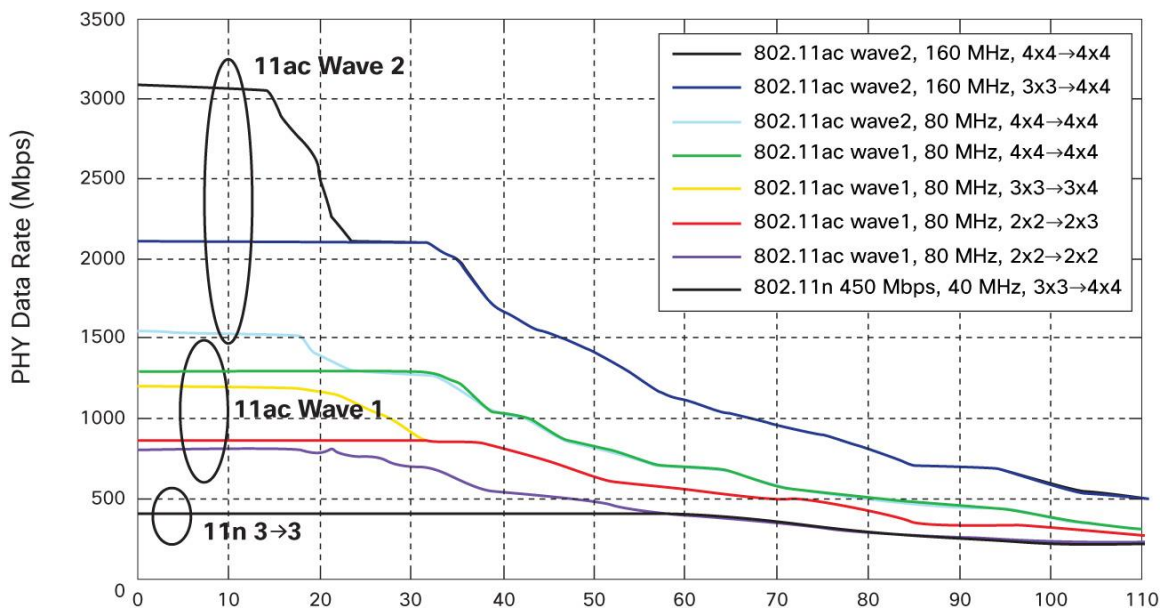
As well as offering higher speeds, 802.11ac also delivers greater robustness than 802.11a or 802.11n.

Consider that, to deliver 450 Mbps, 802.11n has to use 3 spatial streams maxed out at the sensitive 64QAM constellation, and with little multipath immunity: short guard interval and very little coding gain (a rate 5/6 code, so

20% allocated to redundancy). Yet, by going from 40 to 80 MHz, 802.11ac achieves 530 Mbps using only a long guard interval, 16QAM, and rate-3/4 coding (that is, 33% redundancy).

We see this improvement in [Figure 7](#), where 80 MHz links offer higher data rates close in and farther out. In wave 1, different product configurations offer different benefits, but all are a marked step up from 802.11n. Meanwhile wave 2, and especially 160 MHz, offers still greater speeds. However, this improvement is not immediately useful, especially in the enterprise, due to the very limited number of 160 MHz channels that are available.

Figure 7. Simulation of Rate at Range for 802.11ac



2.3.8 Regulatory

Regulatory considerations and 802.11ac intersect in five respects:

- In some regulatory domains, new rules are needed for devices to transmit 80, 160, and/or 80+80 MHz waveforms at all:
 - Effective March 2012, >40 MHz operation is allowed in the United States, the European Union, Australia, New Zealand, Brazil, and South Africa, and no obstacle is expected in multiple other countries.
 - A few countries might only allow 80 MHz or 802.11ac operation after it is ratified by IEEE.
- In some regulatory domains, new tests are needed for devices that 160 and/or 80+80 MHz waveforms across adjacent subbands where the present rules allow this (for example, 5.15-5.25, 5.25, and 5.35 GHz).
- In some regulatory domains, new rules are needed to allow the transmission of waveforms across adjacent subbands where the rules presently don't allow this (for example, below and above 5.725 GHz, also known as channel 144).
- 802.11ac devices (and other unlicensed devices) suffer from reduced access to spectrum containing time domain weather radars in and around 5.6-5.65 GHz.
- Due to the wider bandwidth of 802.11ac, there are strong market desires to open up new spectrum, for instance, in the 5.35-5.47 GHz band (which enables two new 80 MHz channels and one new 160 MHz channel):

- See, for instance, U.S. Act of Congress HR 3630, which empowers the NTIA to study opening up this band to unlicensed use.

Due to the fact that regulations around the world are continually evolving, it is difficult to comment on this topic in detail in this white paper.


2.3.9 MU-MIMO

With 802.11n, a device can transmit multiple spatial streams at once, but only directed to a **single** address. For individually addressed frames, this means that only a single device (or user) gets data at a time. We call this single-user MIMO (SU-MIMO). With the advent of 802.11ac, a new technology is defined, called multi-user multiple input, multiple output (MU-MIMO). Here an AP is able to use its antenna resources to transmit multiple frames to different clients, all at the same time and over the same frequency spectrum. If 802.11n is like a hub, then 802.11ac can be thought of as a wireless switch (on the downlink).

However, MU-MIMO is a challenging technology to implement correctly and won't be available in the first wave of AP products. And even when available, MU-MIMO does come with caveats.

[Figure 8](#) shows one piece of the puzzle. To send data to user 1, the AP forms a strong beam toward user 1, shown as the top-right lobe of the blue curve. At the same time the AP minimizes the energy for user 1 in the direction of user 2 and user 3. This is called “null steering” and is shown as the blue notches. And at the same time, the AP is also sending data to user 2 and forms a beam toward user 2 and forms notches toward users 1 and 3, as shown by the red curve. Then the yellow curve shows a similar beam toward user 3 and nulls toward users 1 and 2. In this way, each of users 1, 2, and 3 receives a strong copy of the desired data that is only slightly degraded by interference from data for the other users.

Figure 8. MU-MIMO Using Combination of Beamforming and Null Steering to Multiple Clients in Parallel



For all this to work properly, especially the deep nulls, the AP has to know the wireless channel from itself to all of the users very accurately. And since the channel changes over time, the AP has to keep measuring the channel, which adds overhead. Some APs might use the higher overhead 802.11ac sounding protocol only, but the greatest benefit of MU-MIMO comes if the AP can minimize the number of explicit sounding exchanges, such as with the ClientLink mechanisms.

Meanwhile, the client is receiving its desired signal distorted by some interference from the signals intended for other users. This interference makes the highest constellations such as 256QAM infeasible within an MU-MIMO transmission.

In summary, MU-MIMO allows an AP to deliver appreciably more data to its associated clients, especially for small form-factor clients (often BYOD clients) that are limited to a single antenna. If the AP is transmitting to two or three clients, the effective speed increase varies from a factor of unity³ (no speed increase) up to a factor of two or three times, according to wireless channel conditions.

2.3.10 802.11ac Project Authorization Request

The 802.11ac project authorization request (PAR) that kicked off 802.11ac includes some throughput numbers: 500 Mbps single-user throughput and 1 Gbps multiuser throughput. These numbers are requirements on the 802.11ac **amendment** (that is, the document), not on individual products. The amendment defines that the minimum product allowed to call itself 802.11ac can operate at 290 Mbps single user and not support multiuser at all.

3. When Is 11ac Happening?

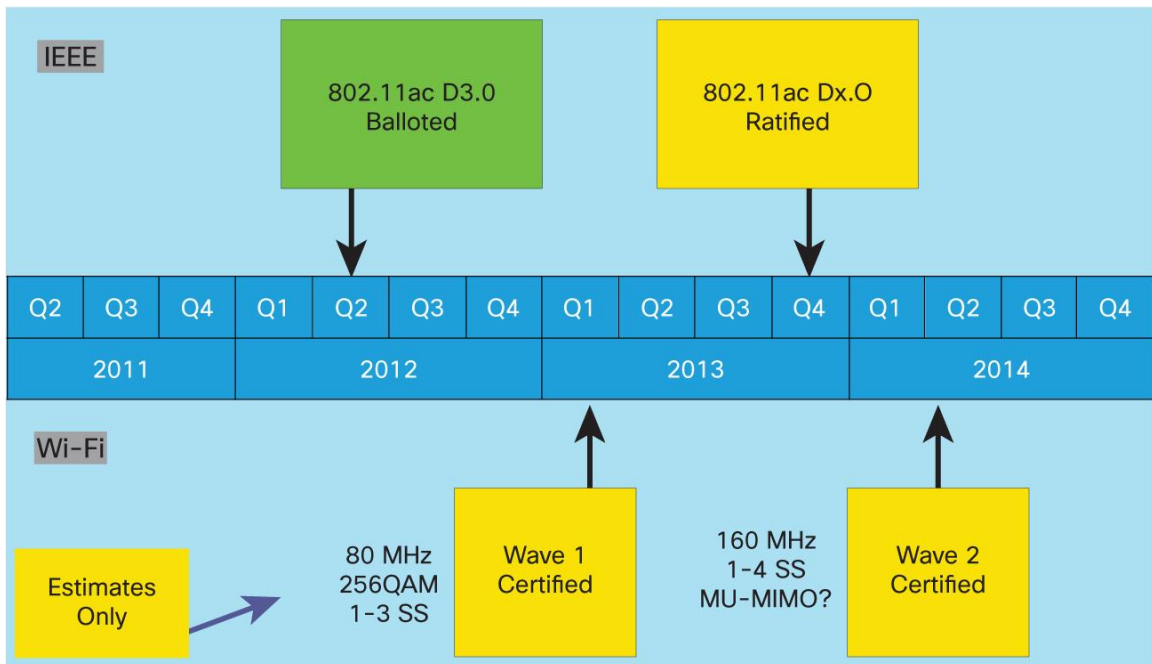
802.11ac is being aggressively standardized, as shown in [Figure 9](#). A mature Draft 3.0 was completed in May 2012, at over 300 pages. It is estimated that the Wi-Fi Alliance will use this draft as the basis for an initial “wave 1” certification in early 2013. 802.11ac products might be released ahead of the certification date, but these are typically consumer-grade products targeting the back-to-school or Christmas markets and don’t come with Wi-Fi Alliance testing of interoperability. Enterprise customers are strongly advised to wait until products certified by the Wi-Fi Alliance become available.

IEEE will continue to refine the 802.11ac amendment based on a continuous improvement process driven by the industry experts. This process creates a sequence of drafts and culminates in the publication of the ratified version. This process is anticipated to complete by the end of 2013.

In parallel, the Wi-Fi Alliance is expected to develop a wave 2 certification that encompasses a wider range of 802.11ac features, such as four spatial streams, 160 MHz operation, and MU-MIMO. The launch date of this is likely to be loosely aligned with the ratification of IEEE 802.11ac, as shown in [Figure 9](#).

³ If the speed-up factor drops below unity, the AP uses SU-MIMO instead.

Figure 9. Timeline for 802.11ac Standardization and Certification



4. How Does 11ac Affect Me?

4.1 Compatibility

One effect **not** to worry about is compatibility.

802.11ac is carefully designed to be maximally forward and backward compatible with 802.11a/n devices. In fact, the 802.11ac design is even simpler and more thorough than 802.11n compatibility with 802.11a devices, as shown in [Table 4](#).

An 802.11ac device must support all the mandatory modes of 802.11a and 802.11n. So an 802.11ac AP can communicate with 802.11a and 802.11n clients using 802.11a or 802.11n formatted packets. For this purpose it is as if the AP were an 802.11n AP. Similarly, an 802.11ac client can communicate with an 802.11a or 802.11n AP using 802.11a or 802.11n packets. Therefore, the emergence of 802.11ac clients will not cause issues with existing infrastructure.

Table 4. Compatibility and Coexistence of 802.11a, 802.11n, and 802.11ac Devices

Receiver Role	Transmitter Receiver	802.11a	802.11n	802.11ac
Intended recipient	802.11a	☑	☑ 802.11n device drops down to 802.11a PPDU	☑ 802.11ac device drops down to 802.11a PPDU
	802.11n	☑	☑	☑ 802.11n device drops down to 802.11n PPDU
	802.11ac	☑	☑	☑

Receiver Role	Transmitter Receiver	802.11a	802.11n	802.11ac
Third-party recipient	802.11a	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> For transmitted HT_MM PPDUs, the third party waits for the packet length as indicated in the legacy portion of the preamble, then an extra EIFS (so no collisions) <input type="checkbox"/> HT_GF PPDUs, preamble should only be transmitted if preceded by MAC protection (for example, RTS/CTS or CTS-to-self) sent using 802.11a format PPDUs (so no collisions)	<input checked="" type="checkbox"/> The third party waits for the packet length indicated in the legacy portion of the preamble, then an extra EIFS (so no collisions)
	802.11n	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/> The third party waits for the packet length indicated in the legacy portion of the preamble, then an extra EIFS (so no collisions)
	802.11ac	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Furthermore, the preamble of the 802.11ac formatted packet is identical to an 802.11a formatted packet, so the CCA mechanism kicks in for third-party both 802.11a and 802.11n devices. As soon as these third-party devices see the 802.11ac preamble, they know the duration of the packet and know not to transmit during that time. Also, since the packet is typically followed by an Ack or Block Ack frame sent in an 802.11a frame, the third-party devices can correctly receive the Ack or Block Ack and then can continue to try to transmit as usual. In the worst case, a third-party device hears the 802.11ac frame but is out of range of the transmitter of the Ack or Block Ack. But even here the third party must wait for an extended duration (called EIFS) to allow time for the Ack or Block Ack to be transmitted without fear of collision.

Because of this preamble-level compatibility, there is no need for 802.11ac devices to precede their 802.11ac transmissions by CTS-to-self or RTS/CTS. The kinds of inefficiencies associated with sending 802.11g packets in the presence of 802.11b devices are completely avoided at 5 GHz.

4.2 When to Upgrade to 802.11ac?

IT administrators are in the fortunate position to be able to pick between two great technologies: (1) 802.11n with A-MPDU, MIMO, beamforming, and speeds from 65 to 450 Mbps within 40 MHz, and (2) 802.11ac with A-MPDU, MIMO, beamforming, and speeds from 290 to 1300 Mbps within 80 MHz.

802.11n is available today and is sufficient for a lot of customer use cases.

802.11ac is the future of wireless LANs, but Wi-Fi-certified 802.11ac APs are not yet available. 802.11ac can provide full HD video at range to multiple users, higher client density, greater QoS, and higher power savings from getting on and off the network that much more quickly.

Most IT administrators deploy new APs at the same time as they fit out a building or retrofit a space. For these, we recommend installing 802.11n APs today, because of the sheer value of 802.11n. Further, for investment protection, it is most desirable to install modular APs that are readily field-upgradable to 802.11ac. As 802.11ac APs become available, these users should automatically start installing 802.11ac APs since the incremental value of 802.11ac exceeds any reasonable price differential.

Also, IT administrators typically upgrade their APs on a three-, four-, or five-year schedule. These IT administrators should continue to upgrade their APs on schedule, since the capability of today's APs significantly exceeds the capabilities of previous generations of APs. Until 802.11ac APs become available, we do recommend that modular 802.11n APs be installed, so as to provide an upgrade path to 802.11ac.

4.3 Radio Resource Management and WIPS Effects

802.11a/11n deployments not upgraded to 802.11ac still have to consider the effect of 802.11ac introduced by neighbors in the normal way and the additional exploits available to attackers.

802.11ac affects RRM since overlapped devices can now transmit over 80 or even 160 MHz. With a software upgrade, it is possible for existing RRM systems to detect the presence of 802.11ac APs from the new 802.11ac fields in beacon frames and extract the affected bandwidth. With this knowledge, the RRM system can mitigate the effect from nearby 802.11ac APs. The RRM system has to work harder since (1) a single overlapped 802.11ac AP affects a wider bandwidth, and (2) the effect on any 20 MHz subchannel depends on whether or not the primary 20 MHz subchannels of the in-network and overlapped APs are aligned (see [section 2.3.6](#)).

Users should verify that their APs are capable of using all the available channels, even those subject to radar detection (DFS) requirements. (Many consumer-grade and some enterprise-grade APs are not certified by regulators to operate on DFS channels. This is unfortunate since in the United States, for example, 63% of 20 MHz channels are DFS channels.)

In general, the wireless intrusion protection system of an 802.11a/11n deployment can detect and mitigate many attacks by 802.11ac devices, particularly when considering naïve attackers. This is because the 802.11ac device communicates using 802.11a/11n format packets when communicating with 802.11a/11n devices, and 802.11ac devices invariably continue to transmit beacons, probe requests, and probe responses at 802.11a rates. However, packets sent using 802.11ac format cannot be decoded by 802.11a/n devices. The recommended countermeasure against such attacks is to provide a sprinkling of 802.11ac APs operating full-time WIPS (for example, one 802.11ac WIPS AP for five or six serving 802.11a/n APs) or a full upgrade of all APs.

5. Summary

802.11ac is an exciting new step for wireless LANs that will be embraced by the industry starting late 2012 (consumer) and from early 2013 (enterprise). 802.11ac is an improved version of 802.11n offering higher speeds over wider bandwidths. 802.11ac is worth having when it is available, and especially when the client mix converges to being dominated by 802.11ac devices. In the meantime, 802.11n offers many of the same technologies, albeit at lower speeds, and is available today. IT administrators looking to invest in wireless LANs in the near term should strongly consider 802.11n APs that are field upgradable to 802.11ac.

Appendix: What Is 802.11n?

802.11n was a major advance over 802.11a. 802.11n introduced several major Medium Access Control (MAC) sublayer and Physical (PHY) layer advances, namely:

- **Multiple input, multiple output (MIMO).** MIMO brings with it a host of benefits:
 - Greater speed without an increase in spectrum consumption using spatial multiplexing (SM). SM splits up the data into pieces and sends each piece along parallel “spatial” channels in a fraction of the time that it would take to send the same data serially. Without SM, 802.11n maxes out at 150 Mbps. With SM, 300 and 450 Mbps are available as long as both transmitter and receiver have at least two and three antennas (and RF chains), respectively.


-
- Greater uplink reliability. Due to multipath, an AP with four antennas receives four copies of a client's signal. Each copy is distorted (constructively or destructively) in four very different ways, so the likelihood that all copies are destructively faded all at the same time is very low. Thus the MIMO equalizer within the receiver can gather all these copies, cleverly combine them, and achieve greater reliability, delivering more predictable data rates and fewer retries. Of course an AP with fewer antennas can't do as well, particularly when the number of spatial streams climbs up toward the number of receive antennas.
 - Greater downlink reliability (maybe). Here 802.11n offers beamforming (with outstanding benefits), space time block coding (useful benefits but inferior to beamforming), and cyclic delay diversity (with very modest benefits). Yet 802.11n offers many incompatible flavors of beamforming, each involving client assistance, and the industry has never put its weight behind any one of them. Thus beamforming is only practically available from techniques that don't expect assistance from the client such as Cisco ClientLink. Beamforming is particularly valuable due to the vulnerability of low-antenna count devices to destructive fading.

This is all described in much greater detail in a [companion white paper](#).

- **Channel bonding.** By doubling the channel bandwidth from 20 to 40 MHz, a single transmission can carry twice as much data in the same time. Actually, the gain is slightly more than two times since the guard band between the two traditional 20 MHz channels can be used as well.
- **Aggregation.** If the PHY is the engine of a car that generates great power, the MAC is like the car transmission, which is responsible for efficiently delivering the power to the wheels.

Before 802.11n, each 802.11a data frame comes with various overheads such as the preamble for the frame, often an acknowledgement frame, and any time gaps between and around these transmissions. When the data size gets smaller than this overhead, then speeding up the data payload doesn't increase the effective speed by that much. The MAC is throwing away the power.

Figure 10. Forms of Aggregation Introduced by 802.11n



802.11n addressed this concern using two aggregation techniques: the “intuitively” named A-MSDU and A-MPDU, which can also be combined together, as in “A-MPDU of A-MSDU.” With aggregation, the data is packed together in a single unit that is sent with one preamble and acknowledged in one transmission. A-MSDU aggregates MSDUs (for example, LLC+IP+TCP+data) at the top of the MAC transmission path, so an individual MSDU in an A-MSDU lacks a MAC header/footer, such as a sequence number or frame check sequence. This is good for efficiency yet makes retries at the individual MSDU level impossible. Meanwhile, A-MPDU aggregates MPDUs at the bottom of the MAC, so each MPDU in an A-MPDU contains its own MAC header. Efficiency is not quite as good, especially for short MSDUs, but if a packet fails to get through the wireless link - for example, from a single isolated bit error - then the other MPDUs can still be received correctly, and only the erroneous packet needs to be retried. This is shown in [Figure 10](#).

- **Channel access for 40 MHz.** Perhaps the fundamental reason for the success of 802.11 is that anyone can install an AP or use a client, regardless of what other 802.11 devices are already nearby, and mostly it all “just works.”

This comes from a MAC design goal that channel access be reasonably efficient and fair to all, regardless of number of devices, distance to AP, device capability, and so forth - put succinctly as “your packet is as important as my packet.” We see the efficiency goal in the range of MAC techniques to reduce collisions, such as physical carrier sense (don’t transmit if you hear a lot of energy) and virtual carrier sense (don’t transmit while someone told you they’d be transmitting or receiving). We see the fairness goal in that each device is only permitted to transmit after meeting the same carrier sense and collision avoidance requirements.

However, 40 MHz brings real challenges to both collision avoidance and fairness since it is either cost-prohibitive or impossible to maintain accurate physical carrier sense and virtual carrier sense on two 20 MHz subchannels in parallel. Instead a “primary” 20 MHz channel is defined with the usual tight requirements on carrier sense and collision avoidance, augmented by a degraded physical carrier sense on the “secondary” 20 MHz. When a device wants to transmit, it performs channel access in the usual way⁴ - all on the **primary** 20 MHz subchannel. Also, immediately before the device can transmit a 40 MHz packet, the device inspects the physical carrier sense state of the secondary channel for a short duration to make sure that the secondary channel is clear too. If clear, the 20 MHz packet is sent; otherwise the device can either (1) transmit a 20 MHz packet on the primary 20 MHz channel or (2) back off again, then recheck if the full 40 MHz is clear. Remarkably, this simple scheme is reasonably fair, and option (1) is reasonably efficient.

Still, in some topologies, devices on the secondary 20 MHz are treated unfairly with respect to the 40 MHz devices, and so 802.11n has additional channel selection rules to try to avoid this scenario in the first place. These rules work pretty well given the large number of 40 MHz channels available at 5 GHz.

802.11n was notorious in the standards community for its slow progress. There were three causes: (1) The process that 802.11n chose to select the winning proposal invited contention. (2) 802.11n was loved to near-death, in that many experts wanted to help and to contribute their technology. It took a long time to work through the contention by adopting many optional modes and then a long time to refine all the optional modes. (3) Operation of 802.11 systems at 2.4 GHz using 40 MHz channel widths in the proximity of 802.15 systems (such as Bluetooth) raised concerns among parts of the 802.15 community.

⁴ That is, the device checks carrier sense and, if the channel is busy, waits until it is clear, randomly backs off a number of slots, and waits while it counts off those slots.




Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

 Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)