



white paper

February 2013

Waves of Wi-Fi

Like most new Wi-Fi technologies, 802.11ac will come to market in phases. Here's a snapshot of the first two phases and a caution for early adopters:

Phase 1

When: First half of 2013

- What: Nominally 1.3 Gbps—80 MHz channels, 256-QAM, up to 3 spatial streams
- Why: First-generation 802.11ac APs are focused on improving Wi-Fi speeds with features that benefit small networks with one or a few APs. 80 MHz channels work well in small sites with a few APs—homes or small offices. In multi-AP installations, wide channels waste spectrum, increase interference, and decrease overall network capacity. Better modulation (256-QAM) is helpful, but limited to very short range connections. Compared with today's 802.11n products, the net gain is minimal.

Phase 2

- When: Late 2013, early 2014
- What: Nominally 3+ Gbps—MU-MIMO, 80/160 MHz channels, 256-QAM, 3+ spatial streams
- Why: In phase 2 APs, we'll begin to see real differentiation with 802.11n. Second-generation 802.11ac APs have much more to offer enterprises facing capacity strain from mobile devices. Though extra wide channels still lack value in enterprises, MU-MIMO will boost overall wireless capacity in mobile-heavy networks by as much as 2x. This network scalability and performance improvement will depend primarily on AP advances, without demanding significant changes to mobile client capabilities—in keeping with cost, size, and power consumption limitations.

802.11ac: Very High Throughput 802.11AC IS THE NEXT-GENERATION OF WI-FI, COMING TO MARKET WITH HIGH EXCITEMENT AND OPTIMISM TO MATCH.

The roar of mobile devices, increasing adoption of all things wireless, and broadening uses for Wi-Fi are all contributing to consumer and enterprise hunger for the next big thing. But, even though 802.11ac is the next generation technology, does it really deserve all this attention, and is it the next big thing?

Introduction

Also known as Very High Throughput (VHT), 802.11ac is currently an IEEE draft amendment that set out to break the 1 Gbps barrier, to improve Wi-Fi's spectral efficiency, and to expand on the capabilities introduced by 802.11n—MIMO radios, wider channel bandwidths, and faster Wi-Fi.

802.11n drastically improved capacity and reliability, which enabled the use of a number of new Wi-Fi applications that were previously too important for unpredictable wireless links. 802.11ac will continue to drive that trend, encouraging the preference for wireless over wired connections at the network edge. 11ac will also enable another surge in the growth of the WLAN. However, VHT deserves further discussion before we get caught in the flurry of eager marketing.

In theory, breaking the 1 Gbps threshold is an exciting accomplishment. An optimist (or journalist) might see 802.11ac data rates (max near 7 Gbps) and salivate in expectation like Pavlov's dog. But, the more realistic perspective is to see 802.11ac enhancements as an engineering target to aim for in the next several years. The near-term gains are somewhat minimal. As with 802.11n, we'll see incremental adoption of 802.11ac capabilities as product manufacturers inch closer to full feature support. Even today, more than 5 years after Draft 2.0 of 802.11n, we still use only a subset of 11n features; but, our real-world utilization of those features is still hindered by other business and engineering forces at play in enterprise networks (battery life, form factor, applications, cost). But lest we get bogged down in negativity, let's see what 802.11ac brings to the table.

802.11ac is built on a set of key enhancements designed to drive data rates very high:

- Wider channel bandwidth: 80 and 160 MHz
- More spatial streams: up to 8
- More efficient modulation: 256-QAM
- Other noteworthy improvements come with 802.11ac:
- 5 GHz only
- Multiple user MIMO (MU-MIMO)
- Simplified transmit beamforming
- Larger aggregated data frames
- And more

5 GHz Only

TADLEA

At first glance, it might seem counter-intuitive to say that 5 GHz only is an "enhancement," but given the massive capacity advantages of the 5 GHz spectrum, by shifting the Wi-Fi focus here and away from the 2.4 GHz junk band, the cumulative benefit will be unmistakably good. This is a strategic, almost political, "limitation" to 802.11ac that will drive more pervasive adoption of 5 GHz by client devices—and consumer wireless equipment—

benefiting the entire ecosystem. In the past, Wi-Fi standards have been focused on 2.4 GHz; 802.11ac marks a new future focusing only on 5 GHz. (See Table 1)

In a clean environment, 2.4 GHz often has more usable range than 5 GHz for low bitrate applications, but capacity is clearly the need of the future; only 5 GHz can deliver it. With wider channels, more spatial streams, and higher-order modulation, we'll need both more bandwidth and cleaner radio environments than 2.4 GHz can provide to actualize 802.11ac's potential. (See Table 2)

Very High Throughput

Much of the excitement about 802.11ac has focused on its impressive maximum data rate—6.933 Gbps. But like 802.11n, 802.11ac has a dizzyingly long list of modulation and coding schemes (MCSs), and the maximum data rate for a specific product will depend on the radio architecture and feature support of that product. 6.933 Gbps is still on the distant horizon of the future. As we dissect the details of the 802.11ac specification, we will be wise to keep 802.11n in the back of our minds as a historical case study. The slow and steady phases of 802.11n feature adoption should help us to set realistic expectations for real-world 802.11ac products—even in the face of impressive specifications.

	Frequency (GHz)		Frequency (GHz)				
Standard	2.4	5	Channel Size(s)	Max Data Rate	Radio Architecture	Max Spatial Streams	
802.11b	Х		25 MHz	11 Mbps	SISO	1	
802.11g	Х		20 MHz	54 Mbps	SISO	1	
802.11a		Х	20 MHz	54 Mbps	SISO	1	
802.11n	Х	Х	20/40 MHz	600 Mbps	MIMO	4	
802.11ac		Х	20/40/80/160 MHz	~7 Gbps	MIMO	8	

TABLE 2

Pand	Papas (MHz)		# of Non-overlap	oping Channels**	
Ballu	nalige (Minz)	TOTALIMITZ	20 MHz	40 MHz	
2.4 GHz	2400 - 2483	83.5	3	1	
5 GHz	5150 – 5875*	495	24	11	

*Not Contiguous, **Varies by country-maximums shown

Wider Channels

One of the easiest ways to gain throughput is to increase channel bandwidth. 802.11ac attempts to capitalize on that fact by introducing much wider channels than 802.11a/b/g/n.

The original 802.11a/g specifications used static 20 MHz bandwidths; then along came 802.11n, enabling 40 MHz channels. Taking the easy road to higher data rates, 802.11ac introduces both 80 MHz and 160 MHz—contiguous 160 MHz or non-contiguous 80+80 MHz—channel bandwidths, providing 4.5x and 9x higher data rates, respectively. 80 MHz support is mandatory for 802.11ac, while 160 MHz is optional, according to the IEEE. (See Table 3)

Higher data rates are significant, but let's keep in mind that Wi-Fi has a broad range of applications, from residential wireless routers to large enterprises, warehouses to hospitality, small public hotspots to large service provider networks. 80 and 160 MHz channels clearly benefit residential use cases the most, where coverage for a single Wi-Fi cell is the key design concern. In enterprises, the network design must account for channel reuse, client device capabilities, applications, regional regulations, and much more. Thus, the uses for very wide channels are less common.

TABLE 3

Channel Size	MCS	1 SS Max Rate	2 SS Max Rate	Bandwidth Increase	Data Rate Benefit
20 MHz	64-QAM, 5/6, 800 ns Gl	65	130	-	-
40 MHz	64-QAM, 5/6, 800 ns GI	135	270	2x	2.07x
80 MHz	64-QAM, 5/6, 800 ns GI	292.5	585	4x	4.5x
160 MHz	64-QAM, 5/6, 800 ns GI	585	1170	8x	9x



Figure 1: 5 GHz Channel Overview

> In 5 GHz, there's plenty of spectrum to justify 40 MHz channels in most cases, but there are still times when 20 MHz channel designs prevail. Specifically, high-density venues that depend on very small "micro-cells" may see greater aggregate capacity from more non-overlapping channels than from fewer, wider channels with poor channel reuse. But the best channel plan for any environment depends heavily on the client devices that will connect to the network and the applications they will support.

> As shown in Figure 1, mobile devices, VoWiFi phones, and tablets with 5 GHz radios have historically been limited to 20 MHz channels. Newer devices are now supporting 40 MHz channels. So if networks are comprised largely of 20 or 40 MHz only devices, those wide channels may go unused and cause unnecessary spectrum waste.

Spectrum regulations also dictate the number of available channels in a region. In the best case scenario, there are twenty-four 20 MHz channels, eleven 40 MHz channels, and five 80 MHz channels. However, if Dynamic Frequency Selection (DFS) channels are not supported, cut those channel numbers to nine (20 MHz), four (40 MHz), and two (80 MHz), respectively. If mission-critical 5 GHz client devices do not support DFS, the DFS channels cannot be used—if they are used, it effectively creates a coverage gap for non-DFS clients. So if DFS channels are thrown out, only two 80 MHz channels are left, which doesn't leave much room for spectral reuse. Aggregate network capacity would suffer as a result. (See Table 4)

Conversely, when 802.11ac proliferates among client devices, if a network is composed primarily of 80 MHz capable stations and DFS can be used with good channel reuse, 80 MHz channels may yield better per-device and aggregate throughput. The key contingency is "if" or "when" most client devices support such wide channels.

Low-power, battery sensitive devices do not receive any maximum range benefits from wider channels. With wider channels, the transmitter's output power must be split over double (20 > 40 MHz) or quadruple (20 > 80 MHz) the channel bandwidth. Effectively this halves (-3dB) or quarters (-6dB) the power density. However, wider channels more than double and quadruple data rates (40 and 80, respectively). If the client is close enough to the AP, it should see a boost in overall capacity with a 40 MHz or wider channel; but to salvage range and power density at the edge of a coverage cell, it will switch back to a low MCS rate on a 20 MHz channel, so overall range will not improve; but, performance at shorter ranges will improve.

The benefits of 80 MHz or 160 MHz channels are not going to be realized in every network. Few enterprise networks will be able to utilize 80 MHz channels, and the use cases for 160 MHz channels are even fewer. On the optimistic side, wider channels are yet

TABLE 4 # of Non-overlapping Channels** Band 20 MHz 40 MHz 80 MHz 160 MHz 5 2 5 GHz - w/ DFS 24 11 5 GHz - no DFS 9 4 2 0*

*Contiguous. 1 non-contiguous potentially available.

TABLE 5

# of Spatial Streams	Max Data Rate: 20 MHz	Max Data Rate: 40 MHz	Max Data Rate: 80 MHz
1	78	180	390
2	156	360	780
4	312	720	1560
8	624	1440	3120

All rates quoted using 800ns GI

> another tool in the designer's toolbox to improve throughput and potentially enable applications that were previously handcuffed to wired-only network connections. Wide channels may be overhyped, but design flexibility—with proper understanding of channel limitations—is always a plus.

More Spatial Streams

Spatial multiplexing is a MIMO feature that enables multiple parallel data streams on the same channel, for an n-fold increase in throughput (n is the # of spatial streams). Prior to 802.11n, Wi-Fi devices were SISO (single input, single output) and were limited to a single simultaneous data stream. 802.11n provided up to 4 simultaneous spatial streams, but 802.11ac expands that to 8 spatial streams. Similar to wider channels, more spatial streams is an easy way to increase maximum data rates for impressive theoretical gains, but real-world applicability is suspect. (See Table 5)

With spatial multiplexing, clients must have as many antennas as spatial streams, which limits the benefits of spatial multiplexing to high-performance workstations. Single-stream portable devices like tablets and mobile phones see no benefit.

Though 802.11n is capped at 4 spatial streams, today's high-end enterprise APs still only support a maximum of 3 spatial streams. Even then, utilizing all 3 spatial streams requires very high quality client connections, excellent RF conditions, and short range.

In order for spatial multiplexing to work, each spatial stream must be distinctly "heard" by the receiver. Spatial streams must be **de-correlated**, or have some type of differentiation, so the receiver can decipher one stream from another. One method to accomplish this is polarization diversity—one horizontally polarized and one vertically polarized transmit antenna. With environmental reflections and/or smart antenna orientation, two spatial streams are relatively easy to achieve reliably. As more spatial streams are desired, new techniques are required to reliably de-correlate parallel data streams, but those techniques are still in development by chip manufacturers and researchers. The end result is that 4+ spatial streams is a goal for the future, not a present reality.

There are other design hurdles that play a part too, such as antenna spacing, which leads to problems with AP size and aesthetics. Also, for every additional spatial stream, another radio chain is needed, which increases power draw (i.e. PoE requirements) or reduces battery life (i.e. for client devices), and increases cost.

256-QAM

In its mission to improve data rates, 802.11ac also introduces more efficient modulation. Like 802.11a/g/n, 802.11ac will continue to rely on OFDM, but VHT devices will have the option of higher-order 256-QAM (quadrature amplitude modulation). Compared with 64-QAM and equivalent coding rates, 256-QAM provides 1.33x (33%) efficiency gains. (See Table 6)

Of course, higher-order modulation techniques increase efficiency, but they also require greater signal quality. Low-order modulation like BPSK (binary phase shift keying) is very simple and reliable at very long ranges. As modulation complexity increases, signal quality (and radio capabilities) must increase with it, so range must generally decrease.

Visually, we can see how modulation complexity changes by looking at a constellation diagram, where each point on the diagram represents a specific bit pattern—four bits for 16-QAM and six bits for 64-QAM. Constellation density is proportional to signal quality requirements: more dense = better SNR required. 16-QAM and 64-QAM are currently used with 802.11a/g/n.

TABLE 6

MCS	Bits per subcarrier	Max Data Rate* (20 MHz, 1 Spatial Stream)	Max Data Rate* (40 MHz, 1 Spatial Stream)
64-QAM, 3/4	6	58.5	121.5
64-QAM, 5/6	6	65	135
256-QAM, 3/4	8	78	162
256-QAM, 5/6	8	-	180

*800ns Gl



........

Figure 2: 16-QAM and 64-QAM Constellation Diagram



	64-QAM			Q			b	0b1b2b3b4b5
	000 100	001 100	011 100	010 100 +7	110 100	111 100 •	101 100 •	100 100
	000 101	001 101	011 101	010 101 +5	110 101	111 101 •	101 101	100 101
	000 111	001 111	011 111	010 111 • +3	110 111	111 111	101 111	100 111
_	000 110	001 110	011 110	010 110 • +1	110 110	111 110	101 110	100 110
	-7' 000 010	-5' 001 010 •	-3' 011 010 •	-1' 010 010 • -1	110 010	+3' 111 010 ●	101 010	100 010
	000 011	001 011	011 011	010 011 • -3	110 011	111 011	101 011	100 011
	000 001	001 001	011 001	010 001 • -5	110 001	111 001	101 001	100 001
	000 000	001 000	011 000	010 000 • _7	110 000	111 000 ●	101 000	100 000

Figure 3: 256-QAM Constellation Diagram

As shown in Figure 2, 256-QAM modulation with 802.11ac increases modulation complexity by another order of magnitude, signifying 8 bits with each constellation point.

In Figure 3, as we evaluate the usefulness of 256-QAM, we must remember that it is not a feature that we simply enable or disable. Such complex and efficient modulation techniques require the right environmental conditions and high quality wireless receivers. With today's products, 256-QAM will only work at very short range (10-20m). So once again, the data rate gains offered by 802.11ac are dependent on the right scenario.

256-QAM is also optional for 802.11ac; first generation devices are expected to support it.

Bigger Frames

As we've seen thus far, 802.11ac techniques increase the physical data rate, improving spectral efficiency. However, as we're able to transfer more data in shorter periods of time, our real-world throughput will always be limited by protocol overhead, of which there are several sources: random backoff, interframe spaces, acknowledgments, errors and retries, and frame headers. Protocol overhead always reduces real-world throughput to some fraction of the signaling rate, often near 50% without enhancements like frame aggregation. (See Table 7)



If wireless channels are a highway, frame aggregation is analogous to carpooling. When each frame (vehicle) has only one data packet (person), we need more cars, which creates more overhead. However, if we transport multiple data packets or frames within the same data unit, we decrease total overhead and increase airtime efficiency.

As wireless efficiency improves, larger payloads can be used without an increase in frame corruption caused by interference (Wi-Fi or non-Wi-Fi), which is sporadic and bursty by nature (see Figure 4). PHY enhancements from 802.11n increased data rates enough to justify the use of frame aggregation without significant interference and retry penalties. 802.11ac continues that trend by significantly increasing PHY rates again, which permits an increase in the maximum frame size.

Table 8 shows there are two types of aggregation:

- aggregate-MSDU (A-MSDU) more Layer-3 packets (MSDU) in the Layer-2 frame (MPDU)
- aggregate-MPDU (A-MPDU) more Layer-2 subframes (MPDU) in the Layer-1 data unit (PSDU)

Real-world 802.11ac products will use frame aggregation commensurate with the connection speed and quality, and the bandwidth demands of the application. As data rates increase or interference decreases, frame size can increase, improving overall throughput. Support for frame aggregation is a critical requirement for high throughput, so all 802.11ac frames must be A-MPDUs.

Multi-User MIMO (MU-MIMO)

In all previous generations of Wi-Fi protocols, a radio could only send one unicast frame at a time to a single receiver. But, in the desire to improve efficiency, 802.11ac seeks to utilize the capabilities of multiple radio chains (i.e. MIMO) to transmit multiple frames simultaneously to different receiving stations. In the past, we had **single-user MIMO**. 802.11ac introduces multi-user MIMO (MU-MIMO).

Wi-Fi enhancements to-date have been exclusively focused on optimizing single user transmissions, which is why we now have sophisticated APs supporting 3X3:3 MIMO. But despite the potential of high-end APs and clients, many clients today especially mobile devices—still support SISO radio designs capable of only one spatial stream. These SISO clients become the airtime bottleneck, requiring more time for each data bit than their high-end MIMO counterparts.

The solution is to utilize an AP's MIMO radio chains to transmit simultaneous downlink frames to multiple stations—up to 4 at a time.

Protocol	Max Data Rate	Throughput
802.11b	11 Mbps	5-6 Mbps
802.11a/g	54 Mbps	20-25 Mbps
802.11n w/ aggregation	450 Mbps	~300+ Mbps
802.11ac w/ aggregation	1.3 Gbps*	Up to 800 Mbps

*1.3 Gbps is the max data rate for the first wave of 802.11ac APs



TABLE 7

TABLE 8

Protocol	Aggregation	Max MPDU Size	Max PSDU Size
802.11a/b/g	Ν	2,304* bytes	~2,336** bytes
802.11n	Y	7,935 bytes	65,535 bytes
802.11ac	Y	11,454 bytes	1,048,575 bytes

*Normally 1500 bytes because of IP packet sizing at Layer-3

**Max MPDU plus typical MAC header

Figure 5 shows the simultaneous downlink data streams with Multi-User MIMO. In order to effectively use MU-MIMO, the AP must be capable of steering individual data streams such that one stream does not interfere with other simultaneous streams. In the IEEE standards, this type of steering is accomplished using baseband signal processing techniques that attempt to create signal peaks (e.g. Signal 1 to Client 1, Signal 2 to Client 2) and nulls (e.g. Signal 1 to Client 2, Signal 2 to Client 1) at the receiver. To do so, it leverages the same channel information that is used to optimize transmit beamforming (chip-based). Unfortunately, because MU-MIMO needs channel feedback from the client device, clients must support MU-MIMO for the feature to be used.



Figure 5: Simultaneous downlink data streams with Multi-User MIMO

With MU-MIMO, the quality of the "steered" signal peaks and nulls depend on environmental channel conditions, but static omnidirectional antennas have limited adaptability to optimize—and isolate—signals for client devices. Smart antennas with dynamic, directional transmit patterns will have a unique advantage with MU-MIMO. With more antenna patterns to choose from, adaptive antennas (i.e. BeamFlex) add value to standardized MU-MIMO by increasing signal optimization (of signal peaks) and isolation (of signal nulls).

Figure 6 demonstrates MU-MIMO with adaptive directional antennas. Even with directional transmissions, signal isolation between clients is never absolute, but the net result is higher SNR (more signal, less interference), which enables simultaneous multi-user streams at higher data rates. Adaptive directional antennas should be a major capacity advantage in network environments with MU-MIMO. However, first-generation VHT chipsets are unlikely to support MU-MIMO.

Transmit Beamforming

Using some of the same feedback and steering paradigms as MU-MIMO, transmit beamforming is also improved in 802.11ac. Even though chip-based transmit beamforming has been standardized by 802.11n for several years, its market acceptance has been very poor. A handful of infrastructure suppliers support it, but client adoption is essentially non-existent, which makes infrastructure support mostly moot. Poor adoption is largely due to the abundance of protocol combinations available in the original 802.11n amendment, which allowed too many sounding options, feedback methods, and feedback formats. (See Table 9)

802.11ac seeks to remedy this by focusing on a single beamforming implementation: null data packet (NDP) sounding and immediate explicit feedback using a compressed beamforming matrix. The end result is that clients and APs alike will be able to focus on a single, compatible, and universal implementation.

TABLE 9										
	Sounding				Feedback Methods and Formats					
		Implicit	Implicit Explicit							
	NDP	Staggered	Staggered			Immediate		Delayed		
		Fleamble	Calibration	-	CSI	Compressed	Non-Compressed	CSI	Compressed	Non-compressed
11n	X	×	х	х	х	х	х	Х	х	х
11ac	х	-	-	-	Х	-	-			

Other Considerations with 802.11ac

Battery Life

For operational purposes, 802.11ac has more power-consuming features and processing requirements than previous generations of Wi-Fi. Because of its radio chains, for example, a 5x5 802.11ac AP would require more power to operate than a 3x3 802.11n AP. To that end, 802.11ac could potentially require more supplied power. However, the higher data rates provided by 802.11ac will allow devices to access the medium, transfer data, and then resume a low power state more quickly than with previous technologies.

With every new product, manufacturing and technology advances are also improving the efficiency inherent in the Wi-Fi chip design, which is also creating power savings. Some of the first chipset manufacturers for 802.11ac are projecting that their newest chips demonstrate better performance and power efficiency than previous 802.11n chips.

Security

The first several generations of 802.11ac products will not require fundamental changes to 802.11 security protocols. However, as throughput increases past the gigabit threshold, inefficiencies with the design of CCMP (currently used with WPA2 and AES) may begin to show up. To meet that need, Galois/Counter Mode (GCM) may supplant CCMP as the cryptographic mode used with AES for Wi-Fi in future VHT networks. Note that this will require hardware upgrades for compatibility, but this need remains far on the horizon.

Backward Compatibility

A fundamental priority for the IEEE is to maintain backward compatibility with previous 802.11 protocols. As a 5 GHz only technology, 802.11ac supports both 802.11a and 11n frame formats and protection mechanisms and is fully backward compatible with both. Wider channels are also compatible with 20 or 40 MHz 802.11a and 11n devices. 802.11ac uses an enhanced version of protection mechanisms (RTS/CTS) to dynamically determine whether all or only some (such as the primary 20, 40, or 80 MHz) of the wider channel is available for transmission.

Timeline

As an IEEE amendment, 802.11ac is currently in Draft 3.0; completion is expected near the end of 2013. 802.11n was a good history lesson to teach us that product manufacturers aren't waiting on a final spec. The Wi-Fi Alliance created a Draft 2.0 certification for 802.11n, and they may do the same with Draft 3.0 of 802.11ac. In 2007, several enterprise products came to market as 802.11n, Draft 2.0. 802.11n was a major change in





> radio technology from previous products, but 802.11ac is more like an extension, so it will be an easier product update for most manufacturers.

Consumer Wi-Fi routers with 802.11ac support have already begun shipping, and a handful of products are now available (October 2012). The first client products are shipping now too, but widespread adoption won't occur until later in 2013. Enterprise APs will become available throughout 2013, some releases being as early as late Q1. However, until the client ecosystem matures and engineering practices are enhanced to enable many of 802.11ac's most noteworthy features, the benefits of an upgrade in early 2013 are minimal.

The details described in this paper are expected to be carried into the final draft, but it is possible that some details may change between Draft 3.0 and final ratification. (See Figure 7)

(Re)Setting Expectations – A Summary

802.11n is an excellent historical case study to help set expectations for 802.11ac. For example, even though 802.11n provides a maximum of four spatial streams, today's 802.11n products hit the ceiling at 450 Mbps (3 spatial streams), not 600 (4 spatial

streams). And even then, achieving 450 Mbps in a real world scenario is rare, requiring nearly perfect RF conditions and short distances. Most enterprise WLAN suppliers continue to sell far more 2x2:2 802.11n APs than 3-stream APs, largely because the higher costs of high-end 802.11n don't often match the level of benefit for their business.

802.11ac is an optimistic protocol with impressive theoretical gains, but the features that improve capacity and enable such high data rates are difficult to use, both because of engineering limitations and because of client device limitations, radio environments, and business factors. 802.11n with MIMO was a massive improvement for Wi-Fi, and 802.11ac will continue to drive wireless adoption.

Perhaps the right perspective is to see 802.11ac as a phased continuation of 802.11n. It sets lofty goals and provides exciting targets to aim for. There's a lot of enthusiasm today about gigabit wireless, but perhaps it is all a little bit premature. If nothing else, industry excitement shows that Wi-Fi has become a bedrock of both consumer and enterprise computing and connectivity. Thankfully, there's no mistaking that 802.11ac gives Wi-Fi an added burst of momentum, even when it is already thriving!



*Draft 3.0 possibly used for a "Wave 1" certification by the Wi-Fi Alliance

Figure 7: 802.11ac Timeline

	Acronyms
A-MPDU	Aggregate MAC Protocol Data Unit
A-MSDU	Aggregate MAC Service Data Unit
CCMP	Counter Mode with Cipher Block Chaining Message Authentication Code (CBC-MAC) Protocol
DFS	Dynamic Frequency Selection
GCMP	Galois/Counter Mode with Galois Message Authentication Code Protocol (GCM with GMAC Protocol)
GI	Guard Interval
HT	High Throughput (802.11n)
DL MU-MIMO	Downlink multi-user MIMO
MCS	Modulation and Coding Scheme
MIMO	Multiple input, multiple output
MU	Multi-user
MU-MIMO	Multi-user MIMO
NDP	Null Data Packet
PLCP	Physical Layer Convergence Protocol
PSDU	PLCP Service Data Unit
QAM	Quadrature Amplitude Modulation
SISO	Single input, single output
SNR	Signal-to-noise ratio
SS	Spatial Stream
SU	Single user
SU-MIMO	Single-user MIMO
TxBF	Transmit Beamforming
VHT	Very High Throughput (802.11ac)
WFA	Wi-Fi Alliance
RTS/CTS	Request-to-Send/Clear-to-Send



Copyright © 2013, Ruckus Wireless, Inc. All rights reserved. Ruckus Wireless and Ruckus Wireless design are registered in the U.S. Patent and Trademark Office. Ruckus Wireless, the Ruckus Wireless logo, BeamFlex, ZoneFlex, MediaFlex, MetroFlex, FlexMaster, ZoneDirector, SpeedFlex, SmartCast, and Dynamic PSK are trademarks of Ruckus Wireless, Inc. in the United States and other countries. All other trademarks mentioned in this document or Website are the property of their respective owners. 803-71287-001 rev 02.

Ruckus Wireless, Inc. 350 West Java Drive Sunnyvale, CA 94089 USA (650) 265-4200 Ph \ (408) 738-2065 Fx

.........

