



802.11ac Wave 2 – Impact on Enterprises

Introduction

With the introduction of 802.11ac Wave 2, enterprise Wi-Fi is entering an inflection point—a turning point where in order to continue advancing one part of your network, other parts must be upgraded.

The two big wins with 802.11ac Wave 2 are four spatial streams and multi-user MIMO, which together can boost aggregate performance and capacity by 30% to 50%. But to get the full benefit, and to lay the foundation for what is to come, there are implications for your power and cable plant.

For years, Wi-Fi engineers have struggled to keep AP power budgets under the 15.4W threshold so enterprises could avoid PoE infrastructure upgrades. However, that fourth stream pushes us over the edge. To enjoy the performance boost of four streams, the PoE infrastructure must be upgraded to PoE+.

The other major advance is MU-MIMO. It is a landmark development which, for the first time, recognizes that not all devices are equal, and addresses the problem in an elegant way. But, it too comes with trade-offs, and the performance gains are greatly dependent on the mix of devices in your network.

Let's explore how Wave 2 can benefit your users, and what it will take to get it.

The Wi-Fi Revolution Continues

ABI Research and IDC both confirm that enterprises have switched over to buying 11ac more quickly than they did when 802.11n was introduced. In the first half of 2016, 802.11ac accounted for almost 60% of enterprise AP shipments.

With 802.11ac Wave 2 access points readily available, it remains to be seen whether enterprises will do the same with Wave 2 products. It largely depends on whether the gain over Wave 1 is perceived to be worth the premium, and whether or not the power, switch, and cable infrastructure is Wave 2 ready.

Ready for Wave 2?

802.11ac Wave 2 adds a fourth spatial stream, and thanks to a new multi-user MIMO scheme, it utilizes the same spectrum more efficiently than ever before.

It sets a new performance benchmark and offers the potential to bring great relief to congested high-density networks. However, its benefits cannot be fully realized until more Wave 2 devices become available.

In 1997 who would have dreamed we would raise throughput by more than 100 times? You have to wonder how much more bandwidth is yet to be squeezed out of the same airspace.

	Standard	The Big Change	20 MHz Max Rate	40 MHz Max Rate	80 MHz Max Rate
2019	802.11ax	OFDMA	~2,000	~ 4,800	~10,000
2017	802.11ac Wave 2+	6-8 Streams	~520 or more	~1,200 or more	~2,595 or more
2015	802.11ac Wave 2	4 streams, MU-MIMO	346.7	800	1,730
2013	802.11ac Wave 1	256 QAM 80 MHz	288.9	600	1,300
2007	802.11n	MIMO 40 MHz	216.7	450	
2003	802.11a	2.4 / 5 20 MHz	54		
1999	802.11b	-	11		
1997	802.11	-	2		

FIGURE 1: TWENTY-YEAR TIMELINE OF 802.11 STANDARDS

802.11ac Wave 2 is just the latest in a long succession of performance advances, with more to come. In 2017 there will be no new standards, but expect more exotic implementations of Wave 2, supporting higher numbers of spatial streams than current models.

Meanwhile, standards experts and RF scientists are hard at work defining what comes next in 802.11ax (dubbed high-efficiency wireless), which is slated for ratification in 2019, though we may see early implementations as soon as 2017.

Huawei engineers have already demonstrated throughput exceeding 10 Gbps in the labs. 802.11ax will use MIMO-OFDMA, a more efficient multiplexing algorithm in which each channel is separated into dozens of smaller sub-channels, each using a different frequency.

OFDMA is already used quite effectively in LTE, and when applied to Wi-Fi, it is hoped that its improved spectral efficiency will quadruple the link speeds of 11ac. It may also improve connection reliability and could reduce interference issues. Anyway, let's return to today and what 11ac Wave 2 does.

Wave 1 and Wave 2 Comparison

Feature	802.11n	802.11ac Wave 1	802.11ac Wave 2
Max Data Rate	450 Mbps	1300 Mbps	3467 Mbps
Modulation	64 QAM	256 QAM	256 QAM
Channel Width	40 MHz	80 MHz	160 MHz 80+80 MHz
Spatial Streams	3 streams	3 streams	4 streams up to 8
Transmit Beamforming	No	Yes	Yes
MU-MIMO	No	No	Yes

FIGURE 2: KEY FEATURES IN 802.11AC WAVE 1 AND WAVE 2

Despite an alluring headline speed rating, the benefits of Wave 2 over Wave 1 are surprisingly not as clear-cut as the difference between 11n and 11ac because there are many dependencies. Let's examine each of the new features:

160 MHz and 80+80 MHz Bonding

Channel bonding up to 160 MHz wide yields impressive headline performance numbers. But don't be fooled—they are misleading, since 160 MHz has few practical uses in client access.

Utilizing 160 MHz wideband channels leaves very few non-overlapping channels for the rest of your channel plan. The best use of wideband channel bonding is not for client access, but wireless mesh and backhaul applications.

At first glance, the limitation of only two contiguous 160 MHz channels in the 5 GHz band seems to make it impractical for almost any application. Dodging such a wide band is an RF planner's nightmare. However, Wave 2 also has the ability to bond two non-contiguous 80 MHz channels. This opens up 19 different permutations for two 80 MHz channels to be combined into a single 160 MHz channel.

The 80+80 bonding feature makes it more practical, but still not easy, to design a channel plan that avoids the two 80 MHz channels while providing good separation between narrower 20 or 40 MHz channels assigned to radios on adjacent APs.

Of course, this is all moot without considering the devices to be supported: First, only Wave 2 chipset-equipped devices can support this 160 MHz channel bonding and there are few of those on the market today; and second, because of the battery-life impact, Wave 2 smartphones or anything smaller may never support 160 MHz. Phablets and tablets will likely be a mixed bag; some will and some won't.

The only devices you can be sure will support 160 MHz bonded channels are high-end laptops. So if that is to be your exclusive user population in a certain location, 160 MHz could be beneficial, provided you can plan around it.

But if you have a mix of laptops and handhelds and they are not all brand new, 160 MHz is not for you. Using wide channels with devices that don't support them is less inefficient than settling for 20 or 40 MHz channels that every device can use fully.

More Spatial Streams

The introduction of multiple spatial stream support (MIMO) and 40 MHz channel bonding was the big performance multiplier for 11n. With each additional stream came increased throughput.

Thus, adding a fourth spatial stream in Wave 2 boosts top-line link speeds by another 30%.

While the standard allows up to eight antenna pairs and eight streams, current Wave 2 APs have only four antenna pairs, allowing up to four spatial streams. For now, the industry has not yet caught up with antenna technology (especially for a smoke detector form factor), to the point that we can cram 16 antennas into a small access point.

In fact, many vendors are still struggling with four antenna pairs. You'll notice most 11ac Wave 2 devices on the market have reverted to ugly rabbit-ear antennas. This is a reversal of the previous trend of sleek AP designs with tamper-proof integrated antennas.

Six- or eight-stream Wave 2 chipsets are not yet available, but this may soon change (Qualcomm, Quantenna, and Marvell all announced plans in 2014). Over the next few years, expect to see more exotic AP variants emerge with six and eight antenna pairs, especially for outdoor, mesh, and backhaul.

PoE and Switch Implications

Unfortunately, the 30% performance increase comes with implications for your power and switching infrastructure. When 11ac Wave 1 first hit the streets, no one wanted to upgrade from 802.3af PoE to 802.3at. Vendors responded with 1x1:1 and 2x2:2 APs to stay under the 15.4W threshold and eventually mastered power management sufficiently well, to power 3x3:3 APs on the same 802.3af PoE infrastructure.

This won't happen again with Wave 2. Operating 4x4:4 takes too much power, and without that fourth stream, what's the point? So vendors are going high-end with Wave 2 AP, and you can expect this trend to continue, with more antennas and streams in the future. 4x4:4 is the new baseline. Future APs will be even more advanced.

So the bottom line is, this time you must upgrade to PoE+ (802.3at), and for most enterprises that really means upgrading a complete layer of access switches for widespread Wave 2 deployment.

Cable Plant Implications

Depending on where Wave 2 APs are deployed, you may also be in for some cable plant upgrades to Cat 6a, Cat 7, or Cat 7a on longer cable drops, if you haven't done so previously.

	Length (meters)	Speed				Power Over Ethernet
		(Mb/s)		(Gb/s)		
		10	100	1	10	
Cat-5	100	X	X			X
Cat-5e	100	X	X	X		X
Cat-6	100 55 for 10 Gb/s	X	X	X	X	X
Cat-6a	100	X	X	X	X	X

FIGURE 3: CABLE DISTANCE AND SPEED RATINGS

Wave 2 APs are capable of saturating their 1 Gbps uplinks, even when both uplinks are bonded using link aggregation to behave like a single 2 Gbps pipe. If cable runs are near to 100 meters long, or they are in an electromagnetically dirty environment, you can expect throughput to fall short of the full 1 Gbps rating. So with Wave 2 you can scrape by with Cat 5 or 5a cabling, but you've hit the end of the line.

Looking beyond 4x4:4 Wave 2, APs will need more than two Gbps uplinks. Vendors may adopt the 2.5 Gbps or 5 Gbps standard, or perhaps jump straight to 10 Gbps copper connections. Either way, support for these standards, at any distance, requires Cat 6a or higher rated cables.

Beamforming

Before 802.11ac, beamforming was a hotly debated topic, with many vendors claiming their proprietary version works best. It was first specified in the standards in 802.11n. However, it was not a mandatory certification requirement, so it never made it to the mainstream. That changed in 802.11ac with a standardized closed-loop method being defined.

Beamforming involves combining the signal from each of the antennas, but applying a specific phase shift to each antenna before adding them. This has the effect of enhancing the signal from certain directions and reducing it from others.

802.11ac radios use a handshaking process with each client to calculate the optimum phase shifts to focus the beam at each device. The result of this beam focusing is to increase the viable range at which devices can operate before they need to step down to lower modulation rates (e.g., QAM 256 step down to QAM 64).

SU-MIMO beamforming and MU-MIMO beamforming are different and mutually exclusive. In MU-MIMO, the AP uses enhanced beamforming techniques to maximize transmission in the desired client direction while minimizing transmission in the direction of undesired clients through null steering.




FIGURE 4: MU-MIMO DIRECTIONAL TRANSMISSION

In SU-MIMO, the beam from all three antennas would be overlaid in the same direction, say the red line, when communicating with device 1, and the yellow line when communicating with device 2.




FIGURE 5: HIGH GAIN AND NULL POINTS IN THE BEAM PROFILE

Figure 5 shows how in MU-MIMO the “sweet spot” of each beam (the point in the beam profile with the highest gain) is focused in one direction, while undesired clients get the weaker part of the signal (represented by the notches in the beam profile). Similarly, the beam for each other stream is focused on a different client, as shown in Figure 4.

SU-MIMO vs. MU-MIMO

The big idea behind multi-user MIMO is this: Instead of directing three or four streams at one client device, let’s split up those streams and direct each one of them at a different device..

Here’s the theory: In standard MIMO, which is now being dubbed single-user MIMO (SU-MIMO), the radio uses all of its antennas to transmit and receive data in multiple streams, but it communicates with only one device at a time.

That’s great, provided all the devices are laptops, which can take advantage of all the streams. But for devices with only

one or two antennas, this system is inefficient. For example, when a 4x4:4 AP communicates with a 2x2:2 tablet or a 1x1:1 smartphone, one or more of those streams is going to waste. What if you could use all the streams all of the time?

Multi-user MIMO aims to do just that, by directing one stream to each of four devices simultaneously (on a 4x4:4 radio). This approach serves more clients at the same time, and uses the spectrum more efficiently. Even still, for now MU-MIMO only works in the downstream direction. The upstream mode was deemed too complicated for the Wave 2 timeline, and has been deferred to 802.11ax. Fortunately, devices generally download much more than they upload, so having downstream MU-MIMO now gives us the lion’s share of its potential benefits.

For MU-MIMO to work, the standard provides a Group ID Management procedure to form client groups. Clients in a given group can be considered together for co-scheduling of transmissions using the MU-MIMO beamforming mechanism.

To be able to perform peak/null adjustments in MU-MIMO beamforming, APs must know about the Tx to Rx antenna / channel matrix for each client in the MU-MIMO group. The standard defines an explicit feedback mechanism for channel learning by which APs transmit sounding packets (Null Data Packet) and clients reply with channel feedback frames.

Obviously this requires MU-MIMO support at both the AP and client. Otherwise, the client has no ability to listen for data on the high-gain beam or to ignore the data on the null-steered beams (which is intended for another client in the MU-MIMO group).

To better understand the real-world impact of MU-MIMO, let’s consider what happens in SU-MIMO when there are multiple clients with mixed capabilities involved and compare it with the same scenario using MU-MIMO instead.

In the simple example below, each column represents a time slice during which a transmission occurs between the AP and one device (SU-MIMO) or more than one device (MU-MIMO).

				
4 x 4 : 4	1 x 1 : 1	2 x 2 : 2	1 x 1 : 1	4 x 4 : 4
4	1	2	1	4
0	3	2	3	0
	  		AIRTIME CONSERVED	
4 x 4 : 4	Phone: 1 x 1 : 1 Phone: 1 x 1 : 1 Tablet 2 x 2 : 2	4 x 4 : 4		
4	4	4		
0	0	0		

FIGURE 6: MU-MIMO VS. SU-MIMO SPECTRUM UTILIZATION OVER 5 TIMESLICES

When MU-MIMO is used, the first transmission window is the same. But in the second transmission window, the AP services two 1x1:1 smartphones and one 2x2:2 tablet at the same time, with no waste. By the third transmission, the MU-MIMO radio services the laptop again, two time slices sooner than SU-MIMO. Thus MU-MIMO saves airtime for other devices.

As of Q4 2016, few MU-MIMO smartphones and tablets have hit the market, since Qualcomm's Snapdragon 820 chipset only found its way into devices in the second half of the year.

With average smartphone refresh cycles currently under three years, by the end of 2017, typical enterprises may see 20% to 30% of mobile devices capable of MU-MIMO. By serving MU-MIMO capable devices more efficiently, more airtime is available for servicing less-efficient legacy devices.

There are caveats: When you enable MU-MIMO, it becomes the default for every connection. Only when the device can't support it does the radio switch back to SU-MIMO. But a MU-MIMO ready 4x4:4 laptop cannot tell the AP "I prefer using 4-stream SU-MIMO." The radio splits the four streams among multiple clients, sometimes giving fewer than four streams to the laptop (depending on demand from other devices). So, when the network is congested, aggregate throughput improves, while individual performance for high-end devices is reduced.

Remember SU-MIMO beamforming and MU-MIMO are mutually exclusive. In practice, MU-MIMO has less range, so it's not good for low-density deployments optimized for coverage, not capacity. Users at the cell edge, which previously could connect in SU-MIMO mode, may find those connections less reliable. For this reason, the majority of vendors have chosen to implement only three streams in MU-MIMO mode and to use the fourth antenna to improve connection reliability.

Where is Fortinet on 11ac Wave 2?

Like other vendors, Fortinet recognizes there is no merit in creating stripped-down versions of Wave 2 APs supporting only two or three streams, as all vendors did with Wave 1. We already have numerous 11ac Wave 1 models fitting that description.

So our latest line of Wave 2 APs—the Universal Access Point (U-AP) series—comes with four antenna pairs and supports four streams in SU-MIMO mode and three stream in MU-MIMO mode. The APs also support link aggregation over dual-redundant GbE uplinks. Fortinet's entire line of switches supports PoE+ in various port density configurations.

For existing Fortinet customers, the U-AP series sets a new benchmark in flexibility by fitting seamlessly into your preferred management framework. The APs may be managed either from a WLAN controller from FortiGate or from the cloud via FortiCloud.

Aesthetics are important to many customers. So we invested extra effort to design an advanced 4x4 antenna system. The U-AP series supports all the new Wave 2 capabilities (160 MHz wide channel bonding, four spatial streams, beamforming, and MU-MIMO) without compromising the sleek, tamper-proof design customers have come to expect from Fortinet.

Conclusion

Despite stunning headline performance figures, Wave 2 will not deliver the gigantic leap in client performance you witnessed when upgrading from 11n to 11ac. Performance gains are more likely to be in the 30% to 50% range and will depend greatly on the mix of laptops, tablets, smartphones, and other equipment such as mobile point-of-sale terminals and medical devices in your environment.

The two big wins with Wave 2 are the addition of a fourth spatial stream and multi-user MIMO. However, to realize the full benefits of Wave 2 and beyond, you must upgrade your power infrastructure to PoE+ (802.3at), which is rated up to 30W.

The optimum timing for Wave 2 will depend on your user base. Without the right device mix, the immediate benefits may be marginal. Therefore, until the majority of your users' devices support Wave 2, selective deployment of Wave 2 APs only in overloaded, high-density areas will give you the best ROI.



GLOBAL HEADQUARTERS
Fortinet Inc.
899 Kifer Road
Sunnyvale, CA 94086
United States
Tel: +1.408.235.7700
www.fortinet.com/sales

EMEA SALES OFFICE
905 rue Albert Einstein
Valbonne
06560, Alpes-Maritimes,
France
Tel +33 4 8987 0500

APAC SALES OFFICE
300 Beach Road 20-01
The Concourse
Singapore 199555
Tel: +65.6513.3730

LATIN AMERICA SALES OFFICE
Paseo de la Reforma 412 piso 16
Col. Juárez
C.P. 06600
México D.F.
Tel: 011-52-(55) 5524-8428