


802.11ac: Next Steps to Next- Generation Wi-Fi

WHAT SHOULD WE DO ABOUT 11AC TODAY

Introduction

In the shifting interest away from 11n and towards 11ac (see Figure 1), products and technologies are changing at a frenzied clip, but end-users are overwhelmed by choices and dizzied by marketing spin. Inevitably, many customers are still weighing their upgrade options, asking questions, and attempting to sort through the timelines, technology, and best practices of 11ac migration. The purpose of this paper is to provide insight to those topics.

Figure 1: Enterprise AP Unit Share of 11n and 11ac. Source: Dell'Oro



In This White Paper:

- .11n vs .11ac
- .11ac Wave 1 vs .11ac Wave 2
- How Should I Prepare my Network?
- Our Advice, Summarized

802.11ac: Next Steps to Next- Generation Wi-Fi

WHAT SHOULD WE DO ABOUT 11AC TODAY

Table 1: 11n and 11ac Product Comparison

	11n			11ac		
Cost Trend	\$300-500	\$600-700	\$1000-1500	\$600-700	\$800-1000	\$1000-1500
Radio Type	Dual Band	Dual Band	Dual Band	Dual Band	Dual Band	Dual Band
Spatial Streams	2	2	3	2	3	3
Value Range	Value	Mid-tier	Premium	Mid-tier	Mid-Premium	Premium
Max Speed (total)	600 Mbps	600 Mbps	900 Mbps	1267 Mbps	1750 Mbps	1750 Mbps
Ruckus Product	R300	7372	7982	R500	R600	R700
Ruckus Cost	\$349	\$649	\$999	\$649	\$799	\$999

What Products are Available?

Vendors are rounding out their product portfolios with 11ac options for both indoor and outdoor as well as value- and premium-focused pricing and performance. Compare some simple product specs and pricing from Table 1.

With just a quick comparison of product options, it's easy to see that you can get similar tiers of 11n and 11ac products for the same cost. The key exception today is the value-focused low-end (sub-\$500) of 11n designed for highly cost-sensitive buyers.

802.11n vs 802.11ac

Much of the industry's original 11ac marketing was focused on stupefying maximum speeds resulting from very wide channels, more spatial streams, and more efficient modulation. The reality is that some of these features are little relevant to actual deployments, making the true benefit of 11ac slightly less than advertisements claim. Nonetheless, 11ac does appear to improve performance in almost all cases and it provides flexibility to use—or not use—new features and design techniques (like wider channels or better modulation), based on the customer's environment and client devices.

Let's take an honest look at some such features and their benefits.

Wider channels

In spec tables, 11n's 40 MHz channels look lousy in comparison with 11ac's 80 and 160 MHz channels. When it comes to speeds and feeds, there's no question this is true. But, there remains a real-world tradeoff when it comes to using wider channels.

First, channel access in Wi-Fi dictates that two APs will interfere with one another if they operate on the same channel. Proper

planning and deployment should seek to reduce this form of self-interference—which quickly cripples Wi-Fi capacity—by spreading APs across the available spectrum. But, by using wider 80 MHz channels, the total number of potential “non-overlapping” channels is reduced to a maximum of 5, and is often fewer in many environments. For the vast majority of deployments, 160 MHz channels would ultimately reduce network capacity as a result of self-interference.

Some areas of the world also have regulations preventing some 5 GHz channels that are necessary to make effective spectrum reuse with 80 MHz channels possible. Some networks, for example, will be constrained to only two 80 MHz channels, making 40 MHz channels more effective by far. In fact, many high density networks should still support 20 MHz channels instead of either 40 or 80 MHz to maximize aggregate network capacity.

As a practical demonstration, see Figure 2. Assuming 4 total 80 MHz channels to reuse (this is common today, worldwide), 1/4 of the APs in a deployment will share the same channel—only a single channel is shown in the 80 MHz example. And even low signal levels can impact neighbor APs several rooms away,

Figure 2: Spectrum Reuse with 40 and 80 MHz Channels




802.11ac: Next Steps to Next- Generation Wi-Fi

WHAT SHOULD WE DO ABOUT 11AC TODAY

which causes large areas of self-interference and unnecessary backoff. However, by using 40 MHz channels, we have at least twice the number of channels to reuse and can drastically reduce detrimental bleed. This allows us to better isolate same-channel APs from neighbors. As you can see, by converting even a small number of APs to non-overlapping channels, we limit contention and increase capacity.

A separate issue with wide channels is that a radio's transmit power is spread over more total bandwidth, reducing power density per unit of frequency. The receiving radio will also listen over a wider portion of radio spectrum, which will increase interference. All else being equal, wider channels require higher signal quality in order to utilize high MCSs (i.e. data rates). Figure 3 shows the minimum receive sensitivity level defined by the 11ac standard for a specific channel width and MCS combination.

Figure 3: Minimum Receive Sensitivity for 11ac Channel Widths



In other words, for a given channel size, how much signal is good enough to use a specific rate. Ignore the specific thresholds, but take note of the fact that 40 MHz channels require roughly 3 dB more than 20 MHz channels for the same MCS. And 80 MHz requires roughly 6 dB more than 20 MHz for the same MCS.

Even though we have to keep these channel width limitations in mind, it doesn't mean wide channels are bad. It means that the specs can be deceiving. Very few real-world situations will provide the advertised gains from wide channels shown on paper. But most average-density networks will see overall gains if there are 80 MHz-capable clients.

More Spatial Streams

The first 11ac products on the market support either 2 or 3 spatial streams; this matches 11n products. Even though 11ac specs allow the use of 8 spatial streams (11n specs allow 4), both technical and business constraints will limit real-world adoption of additional spatial streams. In this way, 11n and 11ac are the same today.


Better Modulation

256-QAM is a more efficient modulation technique introduced by 11ac. It requires 11ac clients as well as very high signal quality for reliable use. In those conditions, it can provide up to 33% efficiency gains over 11n maximums (i.e. 64-QAM). This is one feature that can have some benefits in the right network conditions, with no potential negative impact. But for networks with large coverage areas or slow 11ac adoption, the benefits will prove minimal.

Silicon Improvements

Even though it's difficult to quantify on spec sheets and thus doesn't show up in marketing literature, the benefits of better silicon with 11ac will provide some minor gains to all networks, even with legacy clients. For example, many 11n clients show performance improvements between 5-15% when connected to 11ac APs. Figure 4 shows a test result evaluating the performance (bi-directional TCP throughput) of Ruckus 11n (7982) and 11ac (R700) access points with the same testbed of thirty 11n clients. Though this may not arouse the emotional impulse of buyers, it's yet one more data point in favor of 11ac.

Figure 4: 11ac Enhances Performance for 11n Devices



802.11ac: Next Steps to Next- Generation Wi-Fi

WHAT SHOULD WE DO ABOUT 11AC TODAY

802.11ac: Wave 1 vs Wave 2

As we reach a general consensus on the debate about 11ac vs 11n, the next topic of interest is to identify distinctions in 11ac. By now, most of the market has heard about the concept of 11ac Wave 1 and Wave 2, which generally identify the first and second generations of 11ac products. Unfortunately, there's still plenty of confusion surrounding the issue of "waves," so let's clear some of it up:

- 11ac waves are not defined by a standard, nor is there a standardized technical demarcation that identifies one wave vs another
- The term "waves" can be interchanged with "phases," "steps," "generations," or "_____" <insert favorite synonym>
- Wave 1 generally refers to the first generation of devices supporting initial 11ac features
- Wave 2 generally refers to the second generation of devices supporting additional 11ac features not supported in the first generation of 11ac

So, what's the real difference?

- Wave 1 generally supports
 - 80 MHz channels (all 11ac devices must)
 - 256-QAM (an optional feature for 11ac, but all devices should support it)
 - 2 or 3 spatial streams (hey look, it's 11n again!)
 - Better silicon than 11n
- Wave 2 generally supports
 - Same as Wave 1, plus
 - Possibly 160 MHz channels (real-world relevance for enterprises is small)
 - Up to 4 spatial streams
 - Support for multi-user MIMO (MU-MIMO)

The general problem with the concept of "waves" is that it is a marketing construct with no concrete standard or definition to identify it. For example, a Wi-Fi vendor could release a 2x2 11ac AP today (2014). Down the road, the same vendor could release a refreshed (2nd generation) 2x2 11ac AP with the same "speeds-and-feeds" and maybe some minor low-level feature changes with newer Wi-Fi silicon. So is it Wave 1 or Wave 2? The true answer is: ask someone in marketing.

Table 2: Typical Feature Differences for Next-Generation Wi-Fi

Technology	11n	11ac "Wave1"	11ac "Wave2"
Spatial Streams	Up to 3	Up to 3	Up to 4
Channel Width	20/40 MHz	20/40/80 MHz	20/40/80/160 MHz
256-QAM	No	Yes	Yes
MU-MIMO	No	No	Yes
Max 5 GHz Rates	450 Mbps	1.3 Gbps (80 MHz)	1.7 Gbps (80 MHz) 3.5 Gbps (160 MHz)

As a general classification of products, the concept of "waves" is both necessary and helpful. But, understand the features and benefits of a specific product before you cast a vote for this or that wave. We urge end-users to be discriminating, so let's discuss "Wave 2" features next.

160 MHz Channels

For the same reasons discussed previously with 80 MHz channels—increased self-interference, poor spectrum reuse, overall decrease in capacity—160 MHz channels should be little used in enterprise deployments. There may be niche applications for very wide channels in isolated RF environments, but they will be few. For the enterprise market, 160 MHz support is nearly superfluous. Client support for 160s is also a wildcard.

4 Spatial Streams

Wave 2 APs will support up to 4 spatial streams, but the practical benefits of additional spatial streams are relatively minor when it comes to single-user transmissions (i.e. AP transmitting to a single client at a time). Very few clients will support 4 spatial streams (none do today) due to the added power consumption, though indoor mesh links with two 4x4 APs may take advantage of the added streams. The main benefit of having 4 streams is extra diversity and better spatial control for MU-MIMO transmissions, which makes MU-MIMO more likely to work well.

Multi-User MIMO

One of the most contentious areas of 11ac marketing is a forthcoming feature called multi-user MIMO (MU-MIMO); MU-MIMO has been a core topic fueling uncertainty about when to upgrade to 11ac. On paper, MU capabilities are the perfect solution for high-density networks with mobile devices. But some of the real-world challenges may prevent perfect utilization of its capabilities.

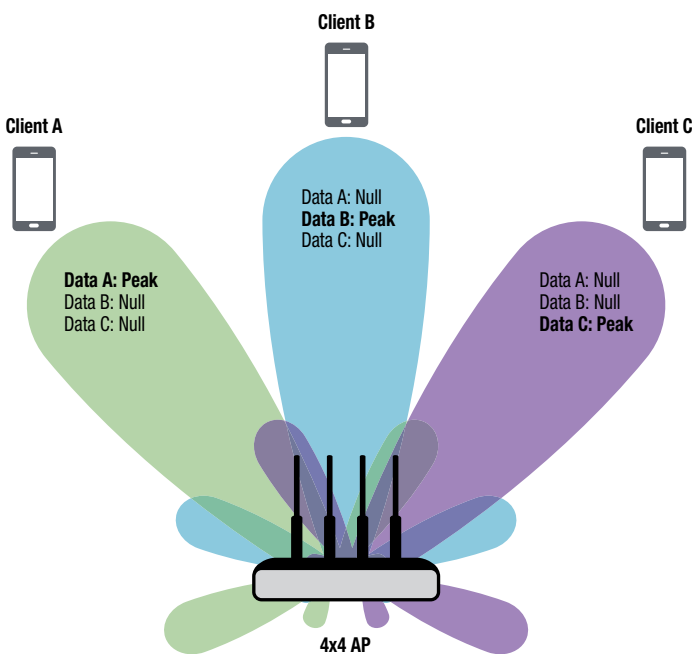
802.11ac: Next Steps to Next- Generation Wi-Fi

WHAT SHOULD WE DO ABOUT 11AC TODAY

The potential benefits of MU technology are very exciting. First, APs typically have several additional spatial streams compared to mobile devices, which only have 1 or 2 streams. MU-MIMO utilizes those extra AP streams most effectively by parallelizing the use of spectrum, instead of merely using extra transmitters for miniscule degrees of link robustness in single-user transmissions (e.g. TxBF, CSD, STBC). MU-MIMO stands to multiply downlink AP capacity by a factor of 2 or 3 when conditions support it.

But it's not a given that MU-MIMO works perfectly in all situations. Figure 5 shows an example of a 4x4 AP using MU-MIMO with three 1x1 mobile clients—note that MU-MIMO requires client support. The first thing you'll notice is that MU-MIMO requires very precise signal control to create peaks and nulls for MU clients with their respective data streams.


Figure 5: Multi-User MIMO Delivering Selective Signal Optimization



The goal is to create a signal peak uniquely for each client and its data. The precision of peaks and nulls will dictate signal quality; where signal nulls (unwanted signals) are imperfect, the result will be interference. Figure 6 shows a view of signal quality graphed over time for each client. This is an idealized scenario where each client receives its data with robust signal quality; simultaneously, the other clients' signals are as near the noise floor as possible.

Keep in mind that Figure 6 represents three different radio perspectives from three different locations over the same period of time and received from the same AP. Wrap your mind around that, and you'll understand why we're hot on this technology.

Figure 6: Optimal Signal Peaks and Nulls with MU-MIMO



The ultimate effectiveness of MU-MIMO depends on expertise with both radio software as well as antenna control. Ruckus' expertise at both of these levels is one reason for our excitement about MU-MIMO, while some other product suppliers have been less optimistic.

How Should I Prepare my Network?

As organizations are planning migration to 11ac, the next major question that comes up is network readiness and integration. After all, what is wireless without wires?

What about PoE?

A sampling of high-end Wave 1 APs will lead you to the conclusion that 802.3at is required. As products were rushed to market in 2013, power budget was sacrificed in the interest of early release, and most APs therefore require 802.3at (PoE+) for full functionality. However, with a more deliberate approach to AP design, Ruckus optimized our AP power efficiency and supports full functionality with 802.3af. This point may be moot for enterprises with pervasive PoE+ Ethernet switches in place already, but for the vast majority still satisfied with Gigabit 802.3af, you can eat your cake and have it too.

The mid-tier of first generation 11ac typically works with full functionality using 802.3af as well.


What about switch backhaul?

Though the complexity of wired network topologies can vary substantially, the general theme with 802.11ac is to ensure that edge switches have sufficient aggregation uplink, which typically means 10Gbps. This is common already today, but it's worth validating during your 11ac migration plans.

802.11ac: Next Steps to Next- Generation Wi-Fi

WHAT SHOULD WE DO ABOUT 11AC TODAY

Figure 7: The Realities of Wave 1 Throughput Demands



Do I need more than a gigabit per AP?

Perhaps the main topic of bewilderment—and misinformation—today is uplink planning for each AP. How much throughput can or will I move through each AP?

Given 11ac's ability to "break the gigabit barrier" in the specs, many end-users immediately assume that they need more than 1Gbps per AP, so they start thinking about link aggregation and running two cable drops per AP. However, if we're critical about the reality of Wave 1 AP capabilities, we shouldn't stress over this. Figure 7 breaks down the real-world realities of 11ac functionality in real deployments with real clients connected.

The descriptions below will provide context for Figure 7.

Quoted MCS – This bar in the graph shows the maximum aggregate data rates supported by an AP. With 1.3 Gbps in 5 GHz and 450 Mbps in 2.4 GHz, the max data rate looks like it threatens 1Gbps Ethernet.

Deployed MCS – When you realize that 450 Mbps in 2.4 GHz requires a 40 MHz channel—which no one uses—then 2.4 GHz is brought down to 216 Mbps. Thus, we can safely say that the "deployed max MCS" is nearly 1500 Mbps ($1300 + 216 = 1516$).

Perfect Throughput – Data rates are not the same as throughput. Wi-Fi has several sources of overhead (e.g. beacons, probes, acknowledgements, errors, retries, frame headers/trailers, etc.), which cause the maximum airtime efficiency to be somewhere near 65% of the data rate. Even then, we still need a perfect scenario to achieve that 65% efficiency:

- No more than 1 or 2 clients per radio (because more clients cause contention and collisions, reducing effective throughput)
- Very little or no interference in the radio environment
- No neighboring networks on the same channel
- Each connected client must
 - Be a 3-stream device (to utilize the highest data rates)
 - Be very close to the AP (to utilize the highest data rates)
 - Have a very clean RF experience
- All clients must also simultaneously send only uplink or downlink traffic (because gigabit Ethernet is full duplex—1Gbps uplink and downlink simultaneously—while Wi-Fi is half-duplex). If some traffic is bi-directional, it decreases the likelihood of gigabit saturation.

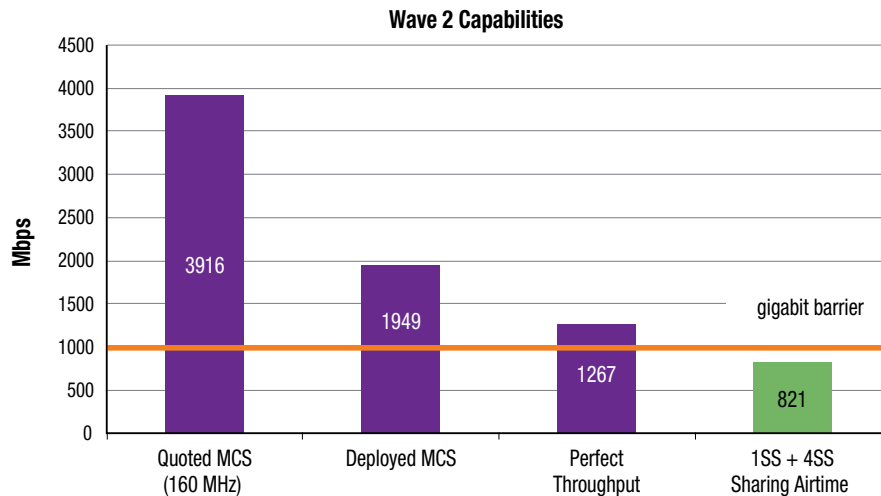
Even if all these conditions align perfectly (which is very highly unlikely), at 65% utilization, we are still operating under line-rate gigabit networks.

1SS + 3SS Sharing Airtime – This metric on the chart simply refers to "perfect conditions" in which there is a 1-stream device and a 3-stream device on each band. However, in most real-world situations, there is a mix of client devices, some with 1, 2, or 3 spatial streams. There are some legacy clients (11a/g/n), some devices farther away from the AP using lower data rates, some amount of interference, and some mixture of bi-directional traffic load. All these factors taken together give us plenty of margin for a single gigabit link.

802.11ac: Next Steps to Next- Generation Wi-Fi

WHAT SHOULD WE DO ABOUT 11AC TODAY

Figure 7: The Realities of Wave 1 Throughput Demands



So how does that compare with Wave 2? Clearly, the max specs for Wave 2 demand additional consideration. Figure 8 assesses the relationship between the gigabit barrier and Wave 2 capabilities.

Quoted MCS – With Wave 2 APs, we have the possibility of 4 spatial streams and 160 MHz channels, which increases our maximum specs considerably. The specs shatter the gigabit barrier.

Deployed MCS – But enterprises are quite unlikely to use either 160 MHz channels in 5 GHz (instead preferring 40 or 80 MHz) or 40 MHz in 2.4 GHz (still preferring 20 MHz). As a result, our deployed data rate is near 1950 Mbps (80 MHz in 5 GHz and 20 MHz in 2.4 GHz). Note that for those rare situations where 160 MHz channels are used either for wireless backhaul or client connectivity, we may need more than 1Gbps support on the wire.

Perfect Throughput – At 65% utilization with a perfect set of clients—1 client per band with 4 stream support, max channel size, perfect RF conditions, maximum data rates, applications that are either all uplink or all downlink, and no other bottlenecks in the application deliver (such as slower WAN links)—we could see Wave 2 APs that break the gigabit barrier by a small margin.

1SS + 4SS Sharing Airtime – This chart again shows two client devices connected to each band, one with 1-stream support and another with 4-stream support, both with ideal conditions delivering throughput equivalent to 65% of data rates. But, even this is generous for real world access point deployments,

which are loaded primarily by mobile client devices with a small admixture of higher-performing 3- or (in the future) 4-stream clients. Most networks also have legacy devices with much lower operating speeds; they also have RF interference, neighbor networks/APs, some combination of bi-directional traffic, errors and retries, and sporadic application load with other application bottlenecks.

Again, these factors combine to reduce the real-world load placed on APs and highly de-risk the likelihood of saturating a gigabit wired link. For those rare networks that see these gigabit-stressing conditions a reality for Wave 2—or for those network planners that prefer to be conservative in the unknowns of “future proofing”—we encourage the use of two cable runs to each AP. Though two cables (for throughput reasons) are unnecessary for Wave 1, Wave 2 may have situations that demand more than a 1Gbps link. Support for link aggregation (LAG), or a similar link-bonding technology should be expected for Wave 2 APs.

802.11ac: Next Steps to Next- Generation Wi-Fi

WHAT SHOULD WE DO ABOUT 11AC TODAY

Our Advice, Summarized

Should I buy 11ac or 11n?

11n still has good shelf life; but, for very similar pricing, 11ac products deliver better performance all around and are mature enough to deserve adoption. Now that prices have come down while both feature parity and stability are nearer to—or better than—11n, it's reasonable to adopt. Two conditions still warrant 11n support:

- Budget sensitive buyers who want enterprise-grade APs at the lowest possible cost
- Highly mission-critical environments that are satisfied with 11n performance and prefer the long history of stability from 11n APs

Should I buy 11ac Wave 1 or Wave 2?

There are still a lot of unknowns regarding the real-world delivery of Wave 2. Will MU-MIMO cure our woes? Will 4x4 APs add much value? If you're faced with a failing network today or "use it or lose it" budget pressure, don't hesitate to upgrade to Wave 1. Without a doubt, you'll see improvements over 11n. However, if your current 11n infrastructure is satisfying your needs, delivering your applications, and providing robust connectivity, don't rush into Wave 1. Wave 2 will deliver better network longevity and remarkable benefits beyond 11n and 11ac Wave 1.

How should I prepare my network for 11ac?

If you agree with our skepticism about 160 MHz channels and very lethargic client support of 4 spatial streams, there's no reason to stress about moving more than 1Gbps of real-world throughput across your network. Don't forget that gigabit Ethernet is full duplex, whereas Wi-Fi is half-duplex. It will take a very unique set of situations in the real world to exceed the capabilities of a full gigabit link per AP, even with Wave 2. If you're on the cautious side, pull two cables and know that you have flexibility down the road.

