Deploying 802.11ac:
Strategies for Gigabit Wireless LANs

A Farpoint Group White Paper

Document FPG 2013-814.1
September 2013

Farpoint Group

802.11ac is here. Yes, of course, the IEEE standard itself is not yet finished, but, as we've seen with earlier wireless-LAN efforts, large numbers of both residential- and enterprise-class products based on relatively-complete (and essentially stable) drafts of the standard are hitting the market now. The Wi-Fi Alliance has established an interim certification program for interoperability, again, as usual, in advance of the completion of the standard itself. Farpoint Group believes that, while the process will clearly take a number of years, 802.11ac will *eventually and entirely* replace 802.11n in all applications. So the question of the adoption of 802.11ac in any given application isn't *if*, but rather *when*, and with what *strategies* applied – the subjects of this Farpoint Group White Paper.

802.11ac Background and Technology

Despite the dramatic increase in throughput at the heart of the new standard, it is best to think of 802.11ac as *evolutionary*, rather than revolutionary. Whereas 802.11n introduced the previously-exotic concept of MIMO (multiple input, multiple output) radio and at-the-time highly-sophisticated modulation techniques, 802.11ac builds upon the success of 802.11n with relatively-straightforward technological advances and extensions. As a consequence, enterprise IT staff will not have a difficult time in understanding the technologies and capabilities of 802.11ac, which, as we are already seeing, are spurring interest and adoption at a remarkable rate (see the Sidebar, "Cisco Systems' 802.11ac Strategy" for more on this).

Assuming all features of the current draft of the standard make it to final approval, and we think such is a good bet, 802.11ac will in theory be capable of reaching a maximum PHY-level throughput of 6.93 Gbps. While we do not believe that products with that level of performance will be common, today's 1.3 Gbps (again, the PHY rate) regardless represents a significant improvement over the 450 Mbps (600 in some cases) maximum data rate of 802.11n. 802.11ac now defines the preferred technology path for chip firms, WLAN system vendors, and end-user organizations of all forms alike. As a consequence, additional innovations yielding advances in throughput across even greater range, support for time-bounded services like voice and video, power consumption, and even backwards-compatible performance (.11ac radios operating in .11n mode) are expected to accrue to .11ac, with .11n implementations now on an essentially terminal trajectory.

Driving the adoption of 802.11ac, however, isn't really a desire for greater per-user throughput, but rather a fundamental requirement for an increase in overall system *capacity*. Consider that, while more throughput is always good, today's enterprise and organizational WLAN environments are being driven by an ever-growing number of users, often with multiple, simultaneously-connected devices per user, running an ever-greater number of applications, many of which have time-bounded requirements that complete the vision of "multimedia" that, interestingly, commenced around the time work was beginning on the first 802.11 standard, more than two decades ago. While an individual user may not have significant throughput requirements at any given moment in time, aggregate demand is skyrocketing in many if not most venues. This phenomenon is a core driver of the adoption of 802.11ac – the faster, and more reliably, that any given user can get his or her bits on and off the air, the more capacity remains for other users. Given that there is always an upper bound on the amount of spectrum available, making the best use of this resource is critical to overall organizational success. Improved spectral efficiency, then, is really the primary goal of 802.11ac – addressing the ever-growing demands for wireless services that have evolved from nice-to-have to essential-to-success.

802.11ac Key Technologies, Features, and Benefits

802.11ac, as previously noted, builds upon technologies and strategies proven in 802.11n, and achieves these new levels of performance via a number of key techniques, as follows:

- *5-GHz. bands only* – Unlike 802.11n, which can be deployed (depending upon implementation, of course) in both the 2.4- and 5-GHz. bands, 802.11ac is 5-GHz. only. This decision was made by those working on the standard based on the fact that 2.4-GHz. spectrum, in most industrialized locations, anyway, is already fairly heavily subscribed, and even oversubscribed in many venues. And, since only 84 MHz. of spectrum is available at 2.4 GHz. (in the United States, anyway; other counties vary but not by much), the wider channels (see below) of 802.11ac wouldn't really be appropriate. Moreover, the 5-GHz. bands have often been underutilized, and, with almost 1 GHz. of spectrum available (with varying regulatory considerations based on locale and specific frequency band, of course), 5 GHz. will provide a very welcoming home for the throughput and capacity enabled by 802.11ac.

- *Wider radio channels* – While 802.11n specifies 20- and 40-MHz. radio channels, 802.11ac expands to 80- and even 160-MHz. channels. More spectrum, of course, is the most important element in increased throughput and capacity. Regulatory changes pending at the FCC will significantly increase the number of 80-MHz. channels possible. It's not expected, however, that 160-MHz. channels, which can be implemented in a single contiguous block, or as two discontiguous 80-MHz. channels, will play much of a role for some time yet, as current 802.11ac chipsets do not support this operational mode. It's also not clear, given how transmit power must be spread across this relatively large space, how much range might be possible given the very high throughput levels specified in this case.

- *Aggressive modulation* – Modulation is the encoding of information on a carrier suitable for transmission. Given that the airwaves are a fairly hostile environment, especially in the unlicensed bands with interference and other inherent challenges always present, modulation techniques also act as a form of error-checking, helping to promote reliable transmission and reception. The core objective in modulation is to encode as many data bits as possible given the variable nature of prevailing radio conditions, so Wi-Fi has the ability to "upshift" and "downshift" performance as appropriate at any given moment in time. Since 802.11ac is building upon the proven MIMO technologies widely deployed in 802.11n, cellular, and other wireless systems, modulation can be quite aggressive, up to 256-QAM, with very good results. This is one of the keys to the higher performance of 802.11ac.

- *Standard beamforming* – Beamforming is a technique where two or more radio transmissions on separate antennas are carefully timed so as to reinforce and even steer (*beamsteering*) a given transmission in a particular direction. This technique is very valuable in creating a more reliable connection between a transmitter and a given receiver, with reliability once again a key to higher throughput. While beamforming has been available in some 802.11n implementations for some time, 802.11ac includes a single beamforming method in the standard itself, meaning we'll be seeing this valuable addition in a wider array of products in the future, with corresponding benefits.

- *Much higher performance* – Like its predecessor WLAN standards, 802.11ac supports a broad range of possible performance levels, extending, as previously noted, all the way to 6.93 Gbps. However, with an increasing market emphasis on handsets and tablets, which inherently have less physical space for antennas and thus benefit from improvements in *single-stream* performance, it's today possible to reach 433 Mbps (see Table 1) with a single-stream 802.11ac implementation. We expect that most smartphones going forward will support at least this level of throughput, a very significant improvement over the 150 Mbps of single-stream 802.11n implementations.

	802.11n		802.11ac		
Mobile Device/Streams	**20 MHz.**	**40 MHz.**	**20 MHz.**	**40 MHz.**	**80 MHz.**
Smartphone (one stream)	72.2	150.0	86.7	200.0	433.3
Tablet (two streams)	144.4	300.0	173.3	400.0	866.7
Notebook (three streams)	216.7	450.0	288.9	600.0	1300.0

Table 1 – Nominal throughput rates comparing 802.11n and 802.11ac. Note both higher throughput and improved spectral efficiency. *Source:* Farpoint Group.

- *Improved price/performance* – But, even more significantly, we are seeing very little, if any, price differential between higher-end enterprise-class 802.11n products and their 802.11ac counterparts. This means that all of that enhanced performance is available at roughly the same price, dramatically improving price/performance, and, of course, further spurring demand. Moreover, 802.11ac products can fall back to 802.11n functionality if required, and we expect the improved technologies included in 802.11ac APs and clients to yield a "better n than n" experience when used in 802.11n mode – that is, .11n performance will be improved by applying 802.11ac products. Thus, we expect that even those sticking with 802.11n for the time being will be purchasing 802.11ac solutions to use in those installations.

Cisco Systems' 802.11ac Strategy

As an example of the rapidly-growing support we're seeing for 802.11ac, we recently spoke with Bill Rubino, Mobility Marketing Solutions Manager at Cisco Systems, about the firm's plans for 802.11ac. Cisco was the first vendor of enterprise-class wireless LANs to introduce a .11ac product, in the form of an add-on module for their 3600 series of 802.11n access points. Cisco, however, has moved well beyond this starting point, with a whole series of new products designed to optimize the capacity inherent in 802.11ac.

"We've coined the term *High Density Experience* to define our introduction of 802.11ac," Bill told us. "With demands for all classes of service and types of traffic continuing their dramatic increase, and with more users and more devices per user, 802.11ac is appearing in enterprise-class products at just the right moment. The rapid deployment of 802.11ac products that we're experiencing indicates we're moving past the early-adopter phase of market growth much more rapidly than was the case with .11n. 802.11ac is already the mainstream technology for many of our customers."

Cisco recently introduced the Aironet 3700 series of access points, which Bill Rubino noted as the "first and only enterprise-class 4x4 MIMO implementation", and which retains the modular design of the 3600 but with 802.11ac now as the primary technology. The modular design can be use to add assurance capabilities, small-cell 3G, or even Wave 2 upgrades when these are available, Bill told us. Farpoint Group believes that the availability of sophisticated products like the Cisco 3700 series of APs is yet another reason that we will likely see more 802.11ac than .11n products shipping – what we call *critical mass* – as early as the end of 2015.

- *Improved battery life* – Power consumption is always a top consideration in wireless communications of any form. 802.11ac implementations take advantage of advances in both the standard as well as improvements in chip and product architecture and production technologies to extend battery life even as throughput and operational duty cycles both increase.

- *Improved location and tracking performance* – While not strictly a part of the standard, we expect the more robust radio signals inherent in 802.11ac to yield significant benefits in real-time location and tracking (RTLS) performance, especially when coupled with relatively dense deployments of APs. Higher density is often desirable regardless and will, we believe, become common, as less range generally corresponds to improved throughput, and thus improved capacity.

- *Multi-User MIMO (MU-MIMO)* – Finally, the so-called "Wave 2" implementations, just beginning to appear in the form of chipsets, will push throughput into the 1.8/3.5 Gbps range, and enable a very exciting capability, *multi-user MIMO*. This technique allows multiple, independent stations to simultaneously receive unique data streams during a single transmit cycle. So, rather than just one station, regardless of its specific throughput requirements, being serviced at any given moment in time by a particular radio in a particular AP, MU-MIMO allows multiple stations to share this capacity during a single transmission period. With many devices utilizing single-stream 802.11ac for reasons of cost or size (as was noted above), MU-MIMO should boost overall capacity in many cases.

If a more technical introduction to 802.11ac is desired, Cisco Systems has produced an excellent tutorial, which can be found here: [http://www.cisco.com/en/US/prod/collateral/wireless/ps5678/ps11983/white_paper_c11-713103.html#wp9000469].

Strategies for 802.11ac Deployment

Still, even with the technological advances and very clear benefits inherent in 802.11ac, it's fair to ask whether a wholesale replacement of 802.11n should be in a given organization's current plans. The answer here is unique to each organization's situation, of course, and depends upon current network utilization trends, budgetary and logistical timeframes, and local 802.11ac client adoption rates. While we advise that investments in 802.11ac, as we'll expand below, should begin now in the vast majority of cases, we are not expecting the wholesale replacement of 802.11n by 802.11ac until into the 2018 timeframe (See Figure 1). Given the large installed base of 802.11n clients, new end-user devices still shipping with 802.11n, and the additional capacity enabled by augmenting 802.11n infrastructure, 802.11n will continue serve quite well in many situations for the next few years.

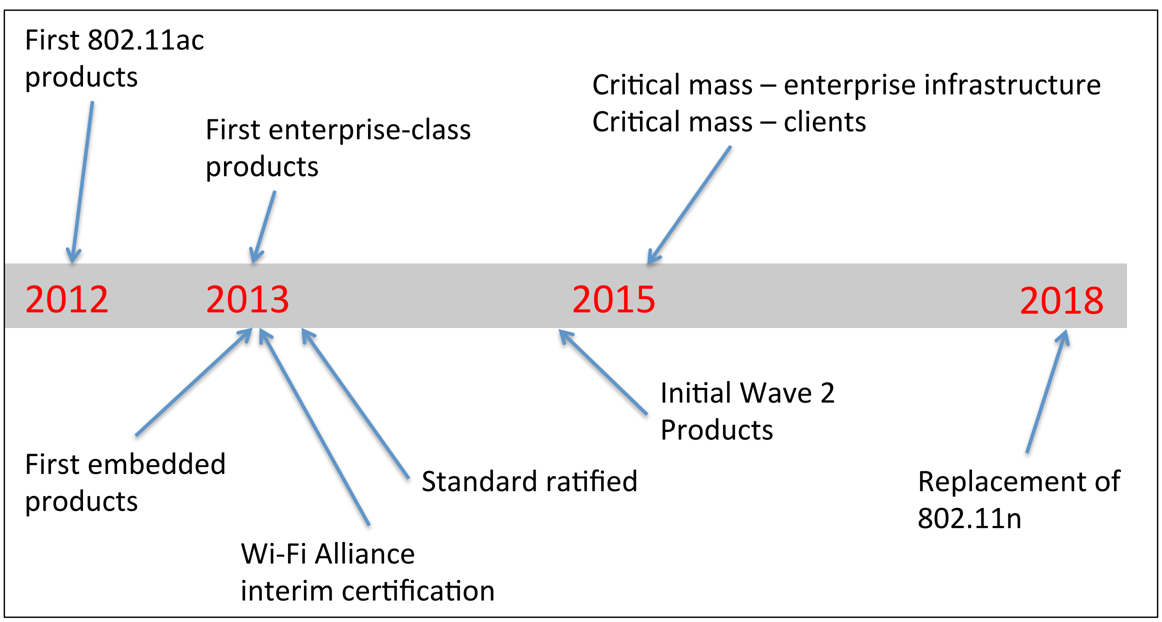

Figure 1 – Forecast timeline for key 802.11ac developments. *Source:* Farpoint Group.

However, given the fact that increasing numbers of clients are indeed shipping with 802.11ac, that the price of 802.11ac APs is roughly the same as that of enterprise-grade 802.11n APs, that the backwards-compatible performance of 802.11ac might in fact be better than that available from an 802.11n AP, and that an eventual upgrade to 802.11ac is inevitable, we suggest that some investment in 802.11ac should indeed be underway at present. This can be justified simply on the basis of providing support for that growing base of .11ac clients, but also to gain experience with the technology in advance of a large-scale deployment – always an advisable strategy.

Regardless, given that the eventual large-scale deployment of 802.11ac is indeed inevitable, a number of possible strategies exist to ease into this eventuality, as follows:

- *Greenfield deployment* – Since venues lacking Wi-Fi coverage of any form still exist (new construction comes to mind, but many older structures also lack WLAN service even today), deploying 802.11ac in these venues, effectively isolated from operation 802.11n (and earlier) WLAN technologies, provides an often ideal environment to gain experience with the technology. Network planners and operations staff can experiment with alternative AP placement, channel configurations, rate-vs.-range behavior, AP density, and other parameters without fear of degrading or even interrupting mission-critical Wi-Fi operations.

- *Limited overlay* – More common, of course, will be a geographically-limited (with respect to extant .11n coverage, anyway) deployment, using available channels in the commonly-underutilized 5-GHz. bands. Note that there is no need to replace existing .11n APs in this case, although many 802.11ac access points will include a 2.4 GHz. radio that can be used to replace .11n (or, as required, .11g) radios operating in this band. The limited-overlay strategy is easily scalable and non-disruptive, and instantly provisions more capacity.

802.11ac Meets Today's Capacity Challenges at the University of New South Wales

Even though it's relatively new on the Wi-Fi scene, 802.11ac is already having a major impact on organizations dependent upon high wireless-LAN reliability, availability, and capacity. A good example here is the University of New South Wales in Australia, a school with over 50,000 students and 5,000 staff. The campus currently has about 2800 APs, 1200 of which are Cisco 3602 models, with 1000 of these equipped with the Cisco 802.11ac add-on module. "This was a very easy way to add 802.11ac to our network," said Greg Sawyer, Manager, Communication Services, at the University.

Meeting growing demands for capacity is precisely the challenge. "Students are perhaps the most demanding early adopters anywhere," Mr. Sawyer told us. "We're seeing more than 20,000 simultaneous wireless users on our network, and over 160,000 unique devices – more than three per student, along with significant voice and video traffic." A simple test also showed the effectiveness of 802.11ac. - using Samsung Galaxy S 4 handsets, roughly *two to three times* the capacity of 802.11n alone was measured.

Mr. Sawyer also told us that the University will add 700 more Cisco .11ac APs this year, and 800 more next year as part of a refresh, with the campus, excluding some residential areas, otherwise going 100% 802.11ac by 2015. That's an adoption rate that well exceeds what was seen in the early days of 802.11n, but one that is clearly required to keep up with the demands of world-class higher education today.

- *Replacement of 802.11n* – Wholesale, rip-and-replace upgrades to 802.11ac will, at least for the next year or two, remain relatively uncommon. 802.11n has a relatively long life ahead of it, although we usually recommend purchasing 802.11ac APs from this point forward to enable cutover to this technology as capacity demands continue to grow rapidly. See the Sidebar, "802.11ac Meets Today's Capacity Challenges at the University of New South Wales" for more on this topic.

Regardless of which of these paths is taken, four additional planning elements are of vital concern in devising an .11ac deployment strategy:

- *Wired network audit* – Even with WLAN throughput quoted at 1.3 Gbps, today's gigabit-Ethernet ports should prove more than adequate, as effective throughput will normally be in the 500-600 Mbps (although we have seen cases where throughput has exceeded 800 Mbps, such is not likely to be common). 10 Gbps (or greater) uplinks from the switch, however, should be specified, as should 802.11at power over Ethernet. Pulling two Cat-6 cables to each AP location is also a good idea; wire is relatively cheap, but installation labor is relatively expensive, and two Ethernet links will smooth the way to deploying significantly-faster.11ac implementations, starting with Wave 2, in the future.

- *Capacity planning* – It's a good idea to review the entire wired network value chain for any potential capacity bottlenecks, including backhaul links to ISPs. A review of network management logs should easily illuminate trends in demand for bandwidth in a given location and additional useful data, and a review of applications requirements will also point the way to successful, robust 802.11ac deployments. It's also best not to assume that higher throughput implies less of a need for APs – more APs will, instead, undoubtedly be required, and, again, denser deployments provision more capacity, especially in high-demand locations.

- *Assurance functions* – A review of intrusion detection, spectral analysis, and related integrity and availability assurance functions is also important. It might be desirable to deploy a few 802.11ac APs to monitor for unauthorized usage of the technology before IT management is ready to proceed with a full .11ac rollout. 802.11n-based sensors or APs will not be able to detect 802.11ac activity.

- *Wave 2* – Lastly, there's a good deal of concern being expressed about the availability of Wave 2 functionality, with a common question being whether waiting for Wave 2 before deploying any 802.11ac service is advisable. We do not believe that such a path is viable in most cases. Not deploying 802.11ac means that 802.11n will continue to define the upper bound on capacity in a given venue – this is not necessarily bad, but demand for capacity may grow faster that .11n's ability to service these demands. As Wave 2 is not yet formally defined in the standard or by the Wi-Fi Alliance, it may be some time before such functionality is universally available or interoperable. We also expect that Wave 2

APs will be able to be added to existing 802.11ac installations transparently and thus with very little concern.

Finally, keep in mind that there will be significant differences in terms of both performance and management features in vendor product offerings, and that such items as radio resource management can have as much impact on overall system performance as .11ac itself. As we have discussed in previous reports, the architecture and implementation of a given enterprise-class wireless LAN system is perhaps even more important than the choice of a specific radio standard in determining ultimate mission success.

Conclusions and Recommendations

We have, in this Farpoint Group White Paper, made the case that 802.11ac is likely to see a more rapid adoption than previous IEEE WLAN standards. The technology is evolutionary, easing adoption concerns. Improvements in reliability, throughput, and capacity are inherent and already being seen. Price/performance is significantly advanced over 802.11n. And one more point: 802.11ac likely represents the end of the line in terms of major new additions to the 802.11 family. 802.11ad, which operates in the 60-GHz. bands, offers higher throughput today, but with at least some restriction on usable range. Barring the availability of large chunks of additional unlicensed spectrum, 802.11ac will be with us for a very long time indeed.

The bottom line, then, is the time to deploy 802.11ac is *now*. Certainly, 802.11n is not obsolete, and will continue to serve us well for many years. But we suggest here that additional deployment of 802.11n limits the ability of a given organization to meet demands for service that are very clearly growing – and rapidly. A staged, overlay-based upgrade strategy should be easily manageable by essentially all IT staffs, and it's not too early to begin down this path today. The products are here – and the time is now.

A final point: as we have often said, *networks are cheap; people are expensive*. Those networks, then, exist to improve the productivity of the users who depend upon them. And there is no greater way to improve the productivity of those users than to make sure they have access to the IT resources they need with maximum availability and minimal latency. 802.11ac is, we believe, key to reaching – and maintaining – that goal.

Ashland MA USA
508-881-6467
www.farpointgroup.com
info@farpointgroup.com

The information and analysis contained in this document are based upon actual testing and publicly-available information sources believed to be correct as of the date of publication. Farpoint Group assumes no liability for any inaccuracies that may be present herein. Revisions to this document may be issued, without notice, from time to time.