802.11ac Action Plan:
A Network-Readiness Checklist

A Farpoint Group Technical Note

Document FPG 2012-779.2
April 2013

We should state, right up front, that products based on the upcoming IEEE 802.11ac wireless-LAN (WLAN) standard are indeed going to dominate the enterprise wireless local-area network landscape. But as is often the case with new technologies, the hype surrounding the initial availability of .11ac products is running, to our thinking, anyway, a bit ahead of reality. And that's OK, as the schedule ultimately involved here enables enterprise network planners, managers, and operations staff sufficient time to devise a framework for evaluating the features and benefits of 802.11ac, and especially for establishing a strategy and plan to roll out .11ac infrastructure within their organizations – the subject of this Farpoint Group Technical Note.

802.11ac: What's New and Important

The IEEE 802.11ac standard, the latest in a long line of enhancements to both PHY and MAC that can trace their roots back to the start of the project in 1991 and the first standard in 1997, is still under development and completion is not anticipated until late 2013. Similarly, the Wi-Fi Alliance, which assumes the role of specifying and certifying interoperability, is now just beginning work on 802.11ac certification. Continuing with what has become a tradition, however, and as was the case with both 802.11g and .11n, "draft" products are appearing in the marketplace well in advance of the completion of work by either organization noted above. Demand here is being driven, of course, by the fundamental appeal of *gigabit-class* wireless LANs – residential-class products claiming 1.3 Gbps are already on the market, and the standard as anticipated will specify numerous levels of PHY throughput from 6.5 Mbps all the way to 6.93 *Gbps*. While we do not anticipate seeing any products even approaching such lofty levels of performance anytime soon (see Table 1), and ignoring the fact that the historical and regardless significant gap between peak PHY rates and realizable Layer-7 performance will continue to exist (our own tests of early .11ac products to date bear this out), the appeal remains obvious. But while 802.11ac will indeed eventually replace today's 802.11n networks, such is a number of years off. We'll return to this point below, but suffice it to say that we continue to encourage the purchase of 802.11n solutions for the time being.

	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 – Comparative peak performance in Mbps of 802.11n and 802.11ac for common numbers of spatial streams and channel sizes (in MHz.). We expect that single- and two-stream clients will dominate at least initially, as will 40- and 80-MHz. channels. *Source:* IEEE.

802.11ac itself is more evolutionary than revolutionary, building on what were indeed revolutionary technologies in 802.11n, most notably MIMO. Among the key technological innovations in 802.11ac are wider radio channels (80 MHz., and eventually to 160 MHz. in some implementations), standard beamforming, improved modulation

efficiency (potentially resulting in much higher throughput), and even "multi-user MIMO" the ability to transmit individual, unique data streams to multiple receivers simultaneously. But, to keep things simple here, 802.11ac is primarily about improving *reliability*, *throughput*, and *overall capacity* – just as was the case with 802.11n.

What Should the Enterprise Do Now?

It's important to point out here that improvements in WLAN performance in enterprise applications will rarely be evaluated on the basis of end-to-end throughput alone. Indeed, while we can expect three-stream, 80-MHz. 802.11ac to yield on the order of 500 Mbps at Layer 7, organizational deployments are much more concerned with overall reliability and especially *capacity*. This means meeting the diverse needs of a large and increasing number of users with a similarly growing number of devices per user (driven in large measure by today's broadening shift to BYOD) and with a robust mix of application demands, and not simply provisioning the maximum possible point-to-point throughput. And many if not most .11ac-equipped client devices over the next few years will support only one or two spatial streams – still a big boost in potential performance, but certainly not pushing the limits of either 802.11ac or existing gigabit-Ethernet switch ports.

Note also that simply deploying two 802.11n APs (or .11n radios within a single AP) on distinct 40-MHz. channels yields roughly the same effective capacity as, and only slightly less spectral efficiency than, a single 80-MHz. .11ac radio, all other parameters being equal. Given the lack of .11ac clients today, and, as we expect the limited availability of these clients to persist for some time (see "An 802.11ac Timeline", below), continued deployments of 802.11n are therefore recommended. As is always the case, however, such deployments must be evaluated in terms of return on investment (ROI), this being of course financial but primarily in terms of *enhanced end-user productivity*. As we have often remarked, *APs are cheap, and users are expensive*. Continuing with deployments of 802.11n resulting in improved productivity *today*, then, is likely the best use of funds dedicated to the wireless LAN for the time being, rather than saving these for an eventual upgrade to 802.11ac. As a general rule, we recommend that if ROI, however quantified, can be demonstrated for additional deployments of 802.11n over the next two years, then such deployments should proceed without delay.

But since the deployment of 802.11ac essentially everywhere is indeed inevitable, it is also vital that a readiness plan for the enterprise network *overall* be in place in advance of the availability of products based on 802.11ac. The network-readiness checklist thus required should include all of the following activities:

- *Perform a wired-network audit* – We do not expect that 802.11ac as typically deployed (three stream, 80-MHz. channels) will exceed the capacity of ports on gigabit-Ethernet switches. We do, however, recommend that future switch deployments include at least ten-gigabit uplinks – from the switch towards the network core, not for connecting APs – so as to avoid traffic bottlenecks at the switch itself. New cable installations where required should pull at least two

uplink cables to every switch unless fiber is used. And while we expect many .11ac APs to operate on 802.3af power over Ethernet (PoE), we recommend that switches for new deployments provision 802.3at just in case and for future contingencies. And, finally, it's important to look for potential performance problems in the network core and especially in backhaul to the Internet. Current network-management console information can be invaluable in monitoring and analyzing any potential issues here. Keep in mind that, no matter when one upgrades to 802.11ac, traffic demands can be expected to continue to grow at a significant pace regardless. Capacity modeling and planning, then, should extend well into the future beyond any anticipated initial deployment or upgrade.

- *Establish a coexistence/migration plan* – Except in rare cases, there should be no need for a rip-and-replace approach to an 802.11ac upgrade. Indeed, Farpoint Group recommends a *phased, staged, non-disruptive* strategy wherever possible. We suggest that a good place to start with .11ac will be a deployment in a new physical space, one not previously covered by Wi-Fi. This will create an opportunity for careful evaluation, experience, and tuning as required. The optimal deployment strategy regardless will likely involve setting aside a 40- or perhaps 80-MHz. channel just for 802.11ac. Continuity for existing .11n clients is vital, and, as we doubt that optimal performance will be achieved by enabling backwards compatibility for .11ac APs (i.e., provisioning .11n and .11ac on the same AP and channel simultaneously), existing .11n APs can be left in place. Finally, given the increasing use of direct-forwarding architectures, wherein traffic to and from APs need not flow through a controller, controller upgrades may not be required. However, as controllers continue to evolve in terms of new features, upgrades here may regardless be desirable in many deployments.

- *Examine current Wi-Fi channel utilization* – The potential for some re-farming of channel allocations must be considered. This may very well be automatic through the adaptive radio management features of a particular Wi-Fi management console, but we still recommend verification with appropriate tools. We do not anticipate any significant issues here in most cases, however, as the 5 GHz. bands have been notoriously underutilized in most venues. But it may be desirable to operate 802.11ac in 40 MHz. channels in some locations, especially during initial deployments – throughput will be lower in this case, but the potential for disruptions minimized. And we would suggest that 160-MHz. channels should not be considered at this time due what we believe will be an extended period before compatible client devices become available, as well as the very limited number of possible channels and their potential impact on existing channels. See Figure 1 for more on Wi-Fi channels in the 5 GHz. bands.

- *Consider operations and network-management requirements* – Farpoint Group believes that management-console functionality is today *critical* to the success of any given enterprise-class Wi-Fi deployment (see Figure 2). We therefore recommend meeting with one's vendor to discuss both the anticipated 802.11ac upgrade, as well as what features will be available to smooth and enhance

ongoing operations. .11ac also creates a great opportunity, we believe, to consider deploying unified wired/wireless management – another interesting topic for discussions with vendors, as well as to review current and required assurance, compliance, and analytics functionality.

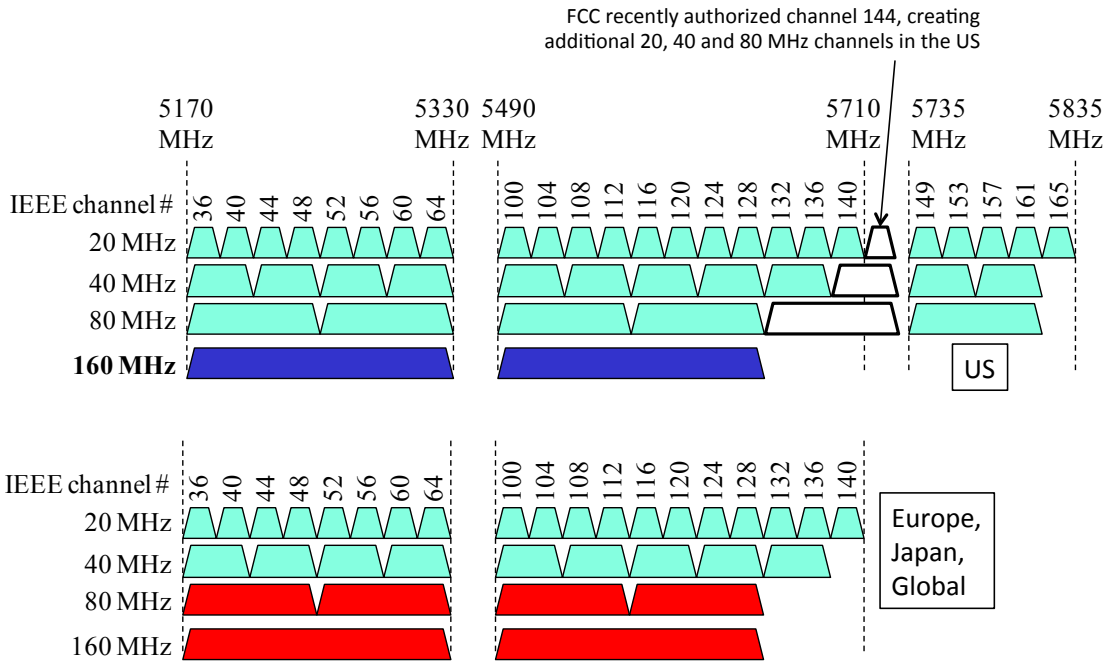

Figure 1 – Channel plan for the 5-GHz. unlicensed bands. Note the limited availability of 80-MHz., and especially 160-MHz., channels. *Source:* octoScope.

- *Apply a dense-deployment strategy* – We have advocated deploying APs with an emphasis on capacity rather than coverage alone since the early days of enterprise-class wireless LANs well over a decade ago. And, as we believe that optimal 802.11ac performance will be obtained by keeping the distance between endpoints as short as possible, a dense-deployment strategy continues to be our recommendation with 802.11ac. Even with the advances inherent in 802.11ac, wider channels may result in less effective range due to the spreading of RF power across more spectrum, and a greater potential for interference. Ever-increasing device density with correspondingly-increasing demands for service will regardless continue to motivate dense deployments irrespective of any specific WLAN technology.

- *Begin a budgetary analysis* – Finally, and this goes without saying, updating budgetary models is also advisable. While some current modular .11n/.11ac designs have appeared at the upper end of current AP prices, we do not anticipate that mainstream, dedicated, enterprise-class 802.11ac APs will be much more expensive than .11n APs. We also expect that investments in improved management consoles will yield positive benefits in the operating expense calculation. We continue to advise capital expenditures that minimize OpEx.

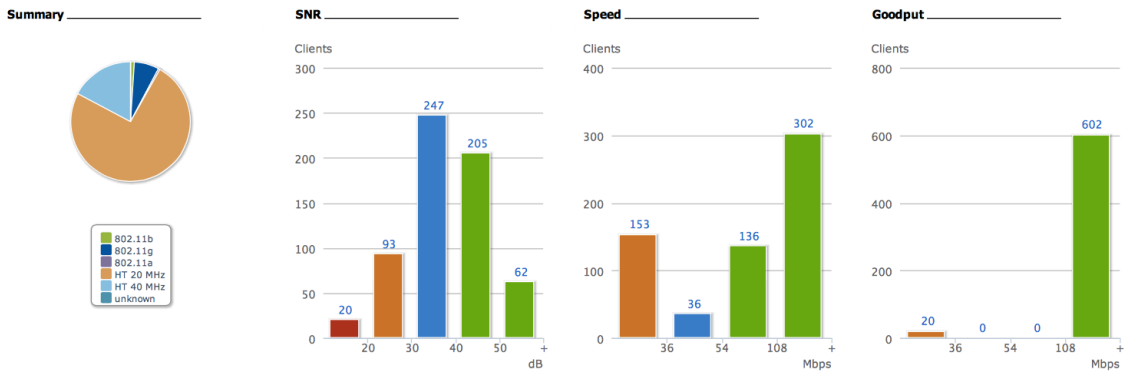

Figure 2 – An example of the information obtained from a management console. This data is invaluable in tuning and monitoring any WLAN deployment, and will be critical in optimizing initial installations of 802.11ac. *Source:* Aruba Networks.

An 802.11ac Timeline

While, as noted above, some residential-class, pre-standard 802.11ac products are now available, Farpoint Group believes that it will be at least mid-2013 before we start to see meaningful numbers of both enterprise-class infrastructure and client products. These will primarily be two-and three-stream, 40- and 80-MHz. APs, and one- and two-stream client adapters.

A full list of our predictions for product and market timing can be seen in Figure 3. Of greatest significance is our belief that 802.11n products will continue to outsell 802.11ac until 2015, and that we will not see wholesale replacement of 802.11n networks until around 2018. Indeed, we believe that 802.11n, properly deployed, will continue to provide service well into the next decade in many venues.

The 802.11ac Action Plan

As the requirements for 802.11n are well-established today, the purpose of this document is to present a checklist of important activities required to prepare for the eventuality of 802.11ac no matter when a given enterprise begins deployment of this next step in the evolution of what has become the default and even primary access for many if not most users in organizations everywhere. A quick review of our suggestions can be found in Summary of Recommendations sidebar, below. None of these items represents any significant departure from what successful IT operations have been doing for years, but such planning today has a greater importance than ever given the mission-critical status that wireless has achieved in the enterprise.

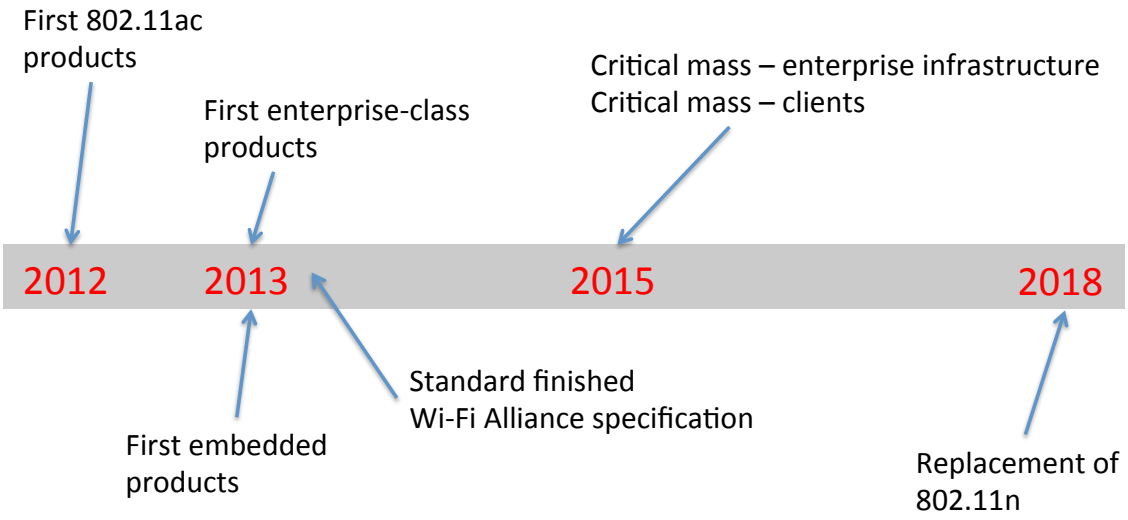

Figure 3 – A projected timeline for 802.11ac. While 2013 will see significant product announcements and the completion of required technical work, sales of .11ac clients and APs exceeding those of .11n ("critical mass") will not occur until 2015, and wholesale replacement of 802.11n will be quite rare before 2018. *Source:* Farpoint Group.

Even as we anticipate significant performance improvements across the dimensions noted in this document, we also expect that enterprise-class vendors will take advantage of the advent of 802.11ac to introduce advances in their overall system architectures and implementations. We must stress here, as we have in previous documents, that elements such as AP and WLAN system architecture and management will be just as important to organizational success as basic technologies like 802.11ac. This is a very important point: while it's possible, for example, to add-on an 802.11ac radio to some 802.11n access points, the best price/performance and perhaps even the best performance overall is much more likely to be achieved with enterprise-class products engineered end-to-end for 802.11ac. These will, again, begin to appear during 2013.

Nonetheless, we are looking forward to meaningful improvements in overall wireless performance, especially in terms of capacity, with 802.11ac. We must caution again that extending range will be desirable in only a restricted set of cases, and that dense deployments focused on capacity will continue to be a key to success. And we expect to see large numbers of single- and two-stream .11ac clients as handsets and tablets continue on their path to becoming the dominant access devices in the majority of enterprise applications. The deployment of 802.11ac for gigabit-class service, then, will be relatively rare for at least the next few years.

Regardless, 802.11ac is clearly now the trend for the future. While upgrades will take some time, Farpoint Group believes that it is important to begin preparing the organizational network for the arrival of this new technology. Properly planned and executed, this transition can be easy for network operations staff and smooth and transparent for users.

Summary of Recommendations

- ✓ **Deploy 802.11n in any situation where a demonstrable need and business case exists today**
- ✓ **Plan switch upgrades as required with 10 Gb uplinks**
- ✓ **Consider pulling two Cat 6 cables to new switch locations**
- ✓ **Plan for an 802.11at PoE requirement**
- ✓ **Perform a budgetary analysis**
- ✓ **Explore network management requirements and opportunities with selected vendor**
- ✓ **Verify core-network and backhaul capacity**
- ✓ **Consider a greenfield deployment of 802.11ac, rather than a replacement of 802.11n, as a first step with the technology**
- ✓ **Augment existing 802.11n deployments with 802.11ac on separate channels**
- ✓ **Consider deploying 802.11ac initially using 40-MHz. channels so as to avoid disrupting operations and the current channel plan**
- ✓ **Ignore the use of 160-MHz. channels for the time being**
- ✓ **Continue to pursue a dense-deployment strategy – *capacity* remains the overriding objective**

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.