White Paper

Implementing 802.11ac

Revolution or evolution?

The increasing demand for wireless bandwidth is driving the development of new standards to increase capacity and throughput and address congestion. The next standard to the market is 802.11ac, which is due to be ratified in early 2014. Some products supporting the new standard have already been released, and more will follow in the next few months. Enterprises need to decide when to implement 802.11ac and how to plan and implement the transition at the appropriate time. The decision will vary from one organisation to another and will depend on both their immediate needs and their long term vision.

To enable their organisation to make the right decision at the right time, network engineers need to understand what the technology offers and the various options available for implementation. This White Paper explains the technology behind the 802.11ac standard, the potential benefits it offers and the factors that should be considered when planning the future development of their wireless network.

Table of contents

The need for greater wireless capacity

Increasing user demands for mobility and connectionless protocols have led to tremendous growth in requirements for wireless capacity within the enterprise. Whether using a company supplied laptop or tablet or their own devices (BYOD), users expect the performance of mobile applications to keep pace with their wired counterparts. They want to use any device to access any application, anywhere, without experiencing any lag or delay in performance.

Growth in wireless use for business and personal use is reflected in increased equipment sales. While PC sales are sluggish, worldwide consumer Wi-Fi customer premises equipment shipments were more than 43.3 million at the end of 1Q 2013 according to ABI Research, a 16.8 per cent increase over the last quarter of 2012. At the same time we are seeing the growth of voice over IP (VoIP) across the enterprise, driven by increasingly mature technology and a desire to reduce costs. By having a single network infrastructure, organisations can – in theory at least – reduce capital expenditure and create a single infrastructure that is easier to maintain and manage.

However, using the same wireless infrastructure to support voice as well as data imposes a penalty in the amount of traffic that can be supported. When VoIP is introduced, more bandwidth hungry applications such as voice and video will be driven over the wireless infrastructure. 4G roll-out will also necessitate an unavoidable increase in the number of VoIP calls made from portable devices, as this is the only type of communication offered by up-to-date networking technology.

These additional demands on the wireless LAN (WLAN) are creating increased pressure to redesign and upgrade wireless infrastructure to provide more bandwidth at higher rates. One alternative to increasing bandwidth is to increase compression of the audio signal, but the trade-off is that any packet loss will have a much more significant impact on call quality. This increasing demand from users is also increasing congestion in the 2.4GHz spectrum, leading to interference and a detrimental effect on the user experience.

Introducing a new standard: 802.11ac

In an attempt to address congestion and the need for higher throughout speeds, the IEEE is developing the new 802.11ac standard for WLAN technology. The standard is still in development, but is expected to be ratified by February 2014, and will be backwards compatible with 802.11n.

So how does the new standard intend to deliver all these benefits? First, it moves wireless traffic to the 5GHz band, instead of 2.4GHz. At the moment this is less busy, so there should be significant benefits to users.

Second, it claims to provide 'faster throughput at greater distance' – which is why it is also known as the very high throughput or VHT amendment. In essence this means providing higher bit rates over a WLAN connection, improving spectral efficiency and building on the techniques introduced in 802.11n through providing:

- wider channels
- higher modulation and coding
- beamforming
- multi-user MIMO
- more spatial streams.

Wider channels

The current 802.11n wireless protocol introduced 40MHz channels, a significant improvement over the 20MHz channels in earlier standards. In theory, with 802.11n you can use up to 14 channels. In practice, to avoid interference, you can only use three or four. If you have conflicting channels, network performance will drop off significantly.

One of the key ways that 802.11ac gains its speed is by using 80MHz wide channels, and in a second phase this will increase further to 160MHz channels. Achieving these higher data rates, however, comes at a cost: fewer available channels in the 5GHz band.

Wi-Fi CERTIFIED ac products: KEY BENEFITS

- **Higher data rate** – Delivering data rates up to 1.3 Gbps – more than double that of a typical Wi-Fi CERTIFIED n network.
- **Higher capacity** – More devices can be simultaneously connected to a Wi-Fi CERTIFIED ac network without reducing performance to address congestion challenges.
- **Lower latency** – Wi-Fi CERTIFIED ac products can deliver a higher-quality user experience with applications such as gaming or streaming music, where even the slightest delay can have a detrimental impact.
- **Efficient power usage** – Enhancements in Wi-Fi CERTIFIED ac mean less power consumption when transmitting data.

Figure 1: the benefits of 802.11ac. Image courtesy of the Wi-Fi Alliance

Exactly how this will work depends on your country, as there are a wide variety of rules on how the 5GHz range can be used *(see Table 1)*. In the United States, it means 802.11ac will have at most five available 80 MHz channel selections; currently there are three available. When 802.11ac second-wave appears, it will have at most two available 160 MHz channel selections, and probably only one will be available. In contrast, there are a total of 13 non-overlapping 20 MHz channels available.

In Europe, 802.11ac has four available 80 MHz channels, and will have two available 160 MHz channels in the second wave. By comparison, there are 19 non-overlapping 20 MHz channels available. We need to keep in mind that 160 MHz wide channels will be optional, even with the ratified standard, and typically the maximum channel bonding will be 80MHz wide channels. However this still significantly reduces the number of non-overlapping channels available in the 5GHz UNII bands.

	INCLUDING DFS*		EXCLUDING DFS	
Channel size	US	EUROPE	US	EUROPE
40 MHz	6	9	4	2
80 MHz	3	4	2	1
160 MHz	1	2	-	-

** DFS = Dynamic Frequency Selection – for avoiding interference with weather radar*

Table 1 - Available 802.11ac channels

Note that without using DFS, in Europe the available 80 MHz channels drops to 1 and in the US it drops to 2 so DFS support in APs and clients is going to be a necessity to deploy 802.11ac effectively.

Higher modulation and coding schema

802.11ac introduces higher order modulation using 256QAM. This increases the number of bits that can be encoded in a single symbol and can provide up to a 33 per cent improvement in bit rates. However, it requires a change in transmitter and receiver design, making the RF design of the system more challenging.

Beamforming

Beamforming is what allows 802.11ac routers to deliver a wireless signal straight to a device rather than bathing the entire surrounding area in the signal intended for that device. Although it is supported in the previous-generation 802.11n, the new standard is more efficient — in part because it only includes one method of beamforming rather than supporting several possible options.

Multi-user MIMO – more spatial streams via more antennae

MIMO, or Multiple Input Multiple Output, means sending and receiving more than one signal at a time. 802.11ac will use multi-user MIMO to support simultaneous transmissions to multiple clients, provided they are spatially separated, which maximises utilisation of the RF band.

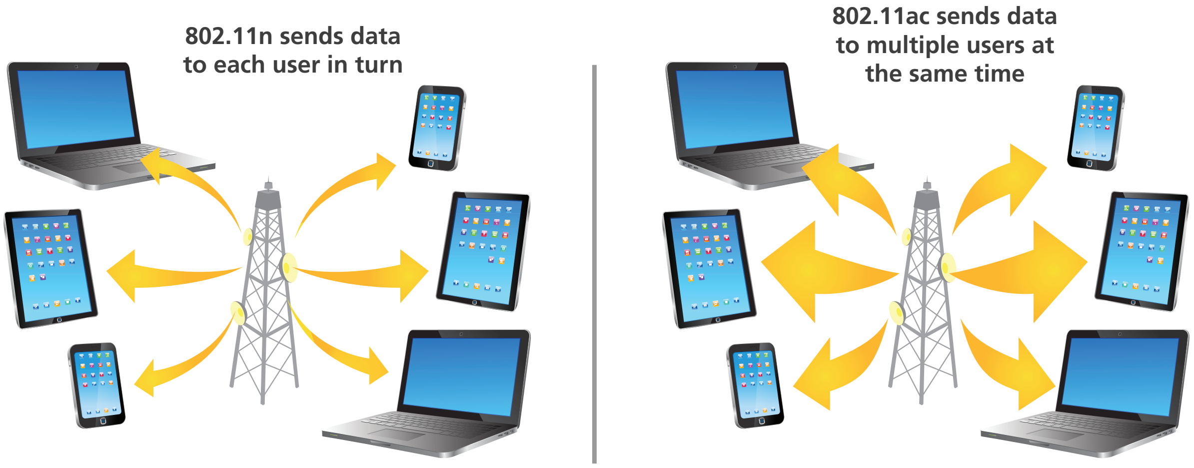

Figure 2: Multi-user MiMo.

In theory 802.11ac can handle up to four spatial streams per client, and each device continuously receives the full bandwidth on offer. In comparison, an 802.11n access point (AP) can only send and receive data from one device at a time. This means that that as more devices use an access point, the slower data transfer becomes as the router transmits a burst to each device in turn before returning to the first devices.

	802.11n	802.11ac
Frequency band	2.4GHz and 5GHz	5GHz
Channel width	20 and 40MHz	20, 40, 60, 80MHz (option 160MHz)
Spatial streams	1 to 4	1 to 8 (up to 4 per client)
Multiple user MIMO	No	Yes
Single stream max. client data rate	150Mb/s	433MB/s (if 80MHz channel)

Table 2: Comparison of 802.11n and 802.11ac protocols

GCMP security protocol

The security protocols used with 802.11ac will in most respects be the same as those used with 802.11n. 802.11ac data rates are incompatible with WEP and TKIP, and so devices operating at very high throughput rates will largely use AES-CCMP, the security protocol required for WPA2 certification.

However, 802.11ac also permits use of GCMP – the Galois/Counter Mode Protocol. Like CCMP, this provides data authentication and encryption. While CCMP encrypts data broken into blocks and then authenticates blocks chained together in sequence, GGMP uses a technique called Galois Field Multiplication to authenticate each block individually. This means that GCMP can encrypt data blocks in parallel instead of in sequence.

When processing data passed at very high rates, the ability to encrypt and authenticate in parallel becomes increasingly important to reduce latency. Each GCMP authentication is also computationally faster than CCMP authentication.

The first wave of 802.11ac products may continue to use CCMP, at least for link layer security, but second wave products may well use GCMP to keep up with high density WLANs, MU-MIMO clients and wider channels that enable data rates approaching 7Gbps. So administrators should start looking for GCMP support in WLAN planning, surveillance and diagnostic tools ready for second wave 802.11ac products and also for 802.11ad, which is best suited for high throughput communications between nearby devices (preferably in the same room) such as HD video transmission to wall mounted wireless displays.

Phased introduction

Current 802.11ac access points and devices support single user MIMO, using beamforming to send the signal efficiently to one device at a time. The next phase will see chips that can send signals to multiple devices at once. This offers several benefits:

- faster devices can save power because they do not need to keep the radio on for as long
- the signal can be made more efficient by only sending it in the direction of devices transferring data, thus reducing interference
- voice calls will be prioritised, but if there is another connections that will not interfere with a voice call because data is being sent in a different direction, it can be sent at the same time
- 802.11ac can send targeted signals to each device. A lower bandwidth device (e.g. a phone) will get a lower bandwidth signal, but unlike 802.11n this will not force higher bandwidth devices such as tablets and notebooks to drop down to the same bandwidth.

The first wave of 802.11ac will include 80MHz channels and 3x3 APs. The next wave will bring 160MHz channels, MIMO configurations greater than 3x3 and multi-user MIMO. Physical layer connection rates will eventually reach 6.9GHz.

In theory, most initial implementations of 802.11ac should enable speeds of up to 1.3 Gbps as well as better coverage than 802.11n. In practice, while significantly faster than 802.11n, 802.11ac can only reach Gigabit-per-second speeds in laboratory conditions.

In the real world, its range is also likely to be more limited than the older 2.4GHz 802.11n and 802.11g technologies, and it can only deliver as much broadband as the slowest link in the network.

However, 802.11ac means that user throughput (in bits per second) will increase. This higher throughput will increase the capacity of 802.11ac APs. Because a user can download files and upload, say, email attachments at faster transmission rates, they will use less time on shared RF channels, so more users who transmit at higher rates can access the AP.

Choosing when to implement the new standard

A number of products supporting the draft 802.11ac specification have already been released. The Wi-Fi Alliance has begun certifying programmes for the new standard and it is expected that certification and equipment will be rolled out in at least two phases, with the first around now and the second in a year or more's time to include all of the specification's enhancements.

Network engineers have a number of implementation options. They may decide to roll out a pure 802.11ac WLAN site in one location, whereas in another they may need to provide backwards compatibility for existing devices with 802.11n and possibly earlier protocols, even if it is a brand new site.

We believe most people will have a hybrid or mixed network for some time in order to support existing and new user devices. This will require design and planning capability to cover both the 802.11ac and 802.11n standards and ensure that users obtain the best performance with both.

Whether your organisation chooses to implement the new standard as soon as it is finalised, or prefers to wait until more equipment is available, it is important to begin planning your approach to 802.11ac now. The chances are that you will need to add new capacity to your LAN over the next few months. Do you begin to implement 802.11ac at that point while maintaining backwards compatibility, or continue to add 802.11n capacity while waiting for more equipment supporting the new standard to become available?

Even if you plan to delay implementation, you will still need to prepare for 802.11ac. Network engineers should therefore begin to familiarise themselves with 802.11ac now so that they can take an informed decision when they next need to increase the capacity or expand the coverage of their WLAN.

They should also consider upgrading the capacity of their Ethernet access and uplink networks. For example, if the AP links are currently 100MB, they will need to be upgraded to 1GB; if 1GB, consider upgrading them to 2GB.Aggregation links need to be sized to allow for all the 802.11ac APs that they will have to accommodate.

Planning for 802.11ac implementation

There are five key factors to consider when planning for 802.11ac:

- throughput
- capacity
- channel allocation
- impact of using DFS channels
- impact of older standards

Measuring throughput

The most important factor when planning a hybrid network is accuracy, particularly when carrying out a site survey. 802.11ac should enable better performance and require fewer APs. However, it is more complex to implement than with earlier standards because there are fewer, wider channels, beamforming has to be considered and there are regulatory variances, making RF management and control essential. This means that signal strength is not a true indicator of WLAN performance; the only true performance indicator is throughput.

The optimum solution for deploying 802.11ac properly and taking advantage of the improvements it offers is therefore to carry out active site surveys as well as Iperf surveys. This enables engineers to measure and map actual end-user performance using an 802.11ac adapter and thus design and deploy 802.11ac networks with a high degree of accuracy.

Assessing capacity

Using a network performance planning tool that supports both existing and new protocols also helps engineers to assess whether there is sufficient capacity in the WLAN. Users are increasingly accessing high bandwidth applications such as Skype and organisations need to assess readiness for VoIP, so it is important to have a means of determining whether the network is reaching capacity.

By using a planning tool which provides visualisations of key performance factors such as channel width, channel overlap and MCS coverage, network engineers can quickly determine areas where high throughput can be achieved and hence high client density can be supported.

Planning channel allocation

It is important to develop a channel application plan when planning for 802.11ac. The wider channels introduced with 802.11ac increase the likelihood of co-channel interference, which will adversely affect performance.

802.11ac designates one sub-channel in a bonded channel as 'primary'; this is the channel that is used for transmission at a specific bandwidth. A planning tool needs to show where primary and secondary channels interfere with each other to enable engineers to adjust channel allocations and AP locations to maximise performance.

Assessing the impact of DFS channels

The 5GHz band used by 802.11ac contains channels with Dynamic Frequency Selection (DFS) capabilities to avoid using the same frequency range as radar. The AP has to vacate the channel it is operating on if it detects radar, which negatively impacts performance. Here a planning tool incorporating a spectrum analyser will help by enabling the network engineer to detect and measure any RF signal on each channel to determine if DFS channels are available or occupied. It also enables engineers to find any non-WiFi interference, preventing expensive network redesigns and ensuring a clean environment for 802.11ac deployment.

Impact of slower transmission rates

Engineers need to ensure that 802.11ac performance is not negatively impacted by the slower transmission rates of 802.11a and n clients. By using a coverage map they can visualise regions where legacy clients can be supported, while a throughput survey using an 802.11ac client validates whether the WLAN can provide the required user performance.

Fluke Networks offers the ability to detect, analyse and troubleshoot 802.11ac APs using currently supported 802.11n adapters. This provides key metrics such as the number of 802.11n and 802.11a clients present in the network, the APs these clients are connected to and network channel utilisation by 802.11n and 802.11a clients. By decoding 802.11ac management frames in real time, engineers can detect VHT capabilities of the AP and thus troubleshoot performance issues in 802.11ac networks resulting from the presence of legacy clients.

About Fluke Networks

Fluke Networks is the world-leading provider of network test and monitoring solutions to speed the deployment and improve the performance of networks and applications. Leading enterprises and service providers trust Fluke Networks' products and expertise to help solve today's toughest issues and emerging challenges in WLAN security, mobility, unified communications and datacenters. Based in Everett, Washington, the company distributes products in more than 50 countries.

For more information on our wireless solutions, visit **www.FlukeNetworks.com/wlan**

Solutions from Fluke Networks

Fluke Networks offers its AirMagnet product line to help solve the BYOD challenge. It spans the entire WLAN lifecycle, ensuring security, performance and compliance. Automatically discovering employee-owned mobile devices it assesses their impact on the corporate network, reduces unwanted side effects and facilitates trouble-free and appropriate use.

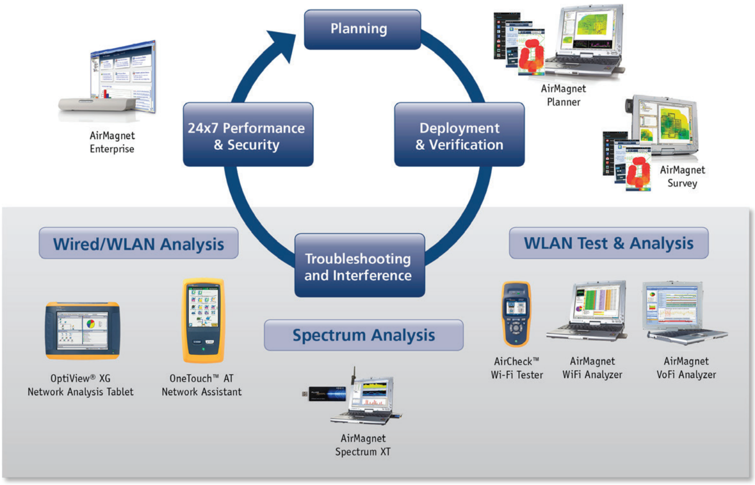

AirMagnet Wireless Solutions

AirMagnet enables predictive modeling of enterprise WLANs, provides advice on AP placement and channel allocations, and runs what-if analyses on the impact of BYOD growth. After WLAN deployment, AirMagnet measures actual coverage and verifies true end-to-end performance and provides an Android app. to visualise coverage.

More information at **www.flukenetworks.com/wlan**

OptiView XG® – Automated network and application analysis

The OptiView XG is the first tablet specifically designed for the Network Engineer. It automates root-cause analysis of network and application problems allowing the user to spend less time on troubleshooting and more time on other initiatives. It is designed to support deployment of new technologies, including unified communications, virtualization, wireless and 10 Gbps Ethernet. The result is that new initiatives get up and running faster and network stay productive even in these days of smaller teams.

More information at **www.flukenetworks.com/xg**

OneTouch AT™ – Client to Cloud troubleshooting in 60 seconds

The OneTouch™ AT Network Assistant greatly reduces troubleshooting time through a streamlined, three-step approach:

1. The unique AutoTest replaces multiple tools and an hour of troubleshooting time.
2. A powerful set of network performance measurements to troubleshoot wired and Wi-Fi networks.
3. It enhances team collaboration through a simple web-remote interface and easy-to-use inline packet capture capabilities.

More information at **www.flukenetworks.com/OneTouchAT**

Fluke Networks operates in more than 50 countries worldwide. To find your local office contact details, go to **www.flukenetworks.com/contact**

Corporate Office:
Fluke Networks
P.O. Box 777 Everett, WA USA 98206-0777
1-800-283-5853
e-mail: **info@flukenetworks.com**

European Office:
Fluke Networks
P.O. Box 1550, 5602 BN Eindhoven
Germany **0049-(0)682 2222 0223**
France **0033-(0)1780 0023**
UK **0044-(0)207 942 0721**
e-mail: **sales.core@flukenetworks.com**