

The Network Impact of 802.11ac

Understanding Network Design Considerations for High-Speed Wi-Fi in a Mobile-First Network



Table of Contents

Executive Summary	3
Introduction	3
802.11ac: The Latest High-Speed Wi-Fi.....	3
Personalization of Mobility	4
The Network Impact	6
Protect Your Network's Future.....	7
Why is This Problem Different?	7
Two Architecture Approaches.....	8
Distributed Network Processing—Rightsizing for Mobility	9
Summary	11
About Aerohive	12

Executive Summary

Two significant industry trends are converging, and the result will greatly affect how wireless networks are designed. The first trend is high-speed Wi-Fi technology, based on the 802.11ac protocol. This protocol supports data rates to 1Gbps and beyond. The second trend is the need for personalizing users' mobile experiences to deliver anytime, anywhere access without wreaking havoc on network operations and cost controls.

Before integrating 802.11ac into your wireless network (WLAN), you must have clear network design objectives and a plan for incorporating high-speed technology, like 802.11ac, without creating additional expense or complication. This paper discusses two approaches for efficient WLAN design. Both let you integrate 802.11ac into your WLAN and both can help lay the foundation for future, high-speed technologies. Learn more about the advantages and disadvantages of these two approaches.

Introduction

Emerging high-speed Wi-Fi, based on the 802.11ac protocol, will have a huge impact on network architectures. At the same time, many IT organizations are struggling with an equally significant trend—personalizing the mobile experience. By understanding each and the relationship between them, you can better inform your strategy for incorporating current and future high-speed technologies in your network architecture.

802.11ac: The Latest High-Speed Wi-Fi

The 802.11ac protocol gives mobile devices revolutionary advances in throughput and capacity. Often referred to as “Gigabit Wi-Fi,” its ability to support WLAN data rates in excess of 1Gbps actually exceeds the capacity of typical wired desktop connections. That means that smartphones and tablets are beginning to enjoy faster connectivity than powerful desktop computers.

Forrester Research predicts that 59% of all data traffic will move from wired to wireless connections by 2017¹. This means that wireless connections are becoming users' *primary* connection. Industry experts also predict that the number of Wi-Fi-connected devices will continue to grow exponentially as machine-to-machine communications begin to proliferate. For WLAN architectures, 802.11n deployments will be quickly supplanted by 802.11ac.

¹ Source: Forrsights Mobility Survey, Q2 2013

But 802.11ac in its current (December 2013) commercially available form is not the end of the road. Increased channel widths and improved antenna technology will push the current standard forward, possibly as soon as 2015 by some estimates. For example, multi-user MIMO (MU-MIMO) is a set of multiple-input and multiple-output technologies that allow multiple transmitters to send separate signals—and multiple receivers to receive separate signals—simultaneously in the same band. Effectively, this capability increases the number of data streams that can be handled and/or the capacity of each stream for higher network performance. An adjacent technology, beamforming, focuses signals to specific client devices, instead of broadcasting a signal over a wide area. This means that more data can reach the target device. If the Wi-Fi client supports beamforming, the transmitter and device can exchange information about each other's location.

Higher raw data rates have significant implications for the network architecture. As these data rates will continue to rise, future-proofing the network architecture becomes even more important. You want to be able to easily adopt new technologies without constantly re-designing and adding cost to the enterprise network.

Personalization of Mobility

When increasing data rates and high throughput intersect with the realities of efficient mobile workforce management, things really get interesting. Mobile applications are deployed to improve employee productivity and business agility. That means the network must provide anytime, anywhere access to applications on any device, corporate issued or BYOD, and ensure an optimized experience for each mobile user regardless if they're a guest, employee, or contractor. Not an easy task.

IT traditionally approached this challenge by creating multiple wireless networks, each with its own service set identifier (SSID). These distinct SSIDs were an attempt to account for:

- Every possible combination of device type and ownership: Whether a user's Bring Your Own Device (BYOD) platform or corporate-issued
- Each user type: Employee, guest, contractor, etc.
- Separate applications: Email, Internet browsing, intranet, etc.

The approach of using SSIDs to control application traffic and the user experience is incredibly complex. In addition to magnifying the management challenge, the SSID approach becomes even more complicated when organizations have more than one operating site. Figure 1 illustrates the complexity.

The Network Impact of 802.11ac

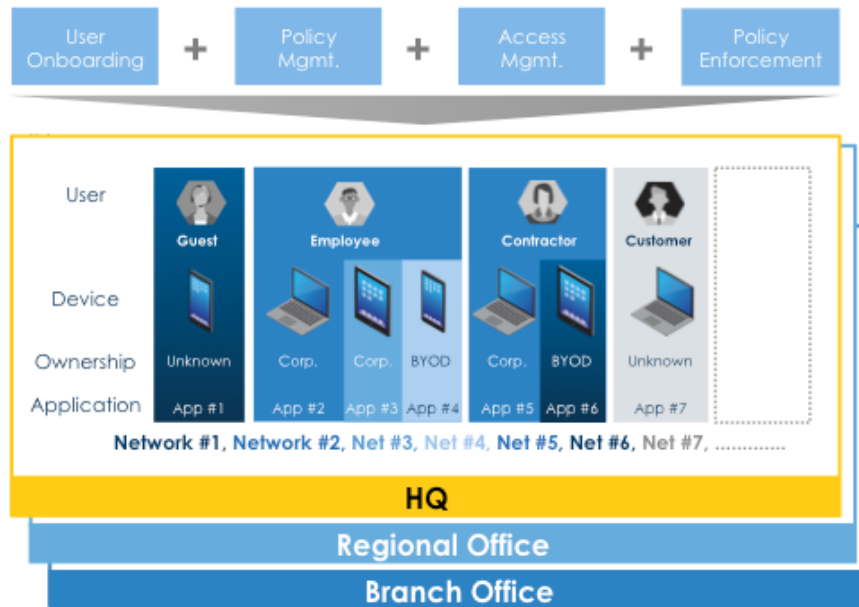


Figure 1: Manually enforcing policy to optimize the user experience is complex.

Modern Wi-Fi systems approach this issue by dynamically applying policy in real-time to the specific user or “personalizing” mobility to the user. This approach uses technology to accommodate all of the differences between individuals, devices, and applications automatically on a simple infrastructure. This level of automation is fundamental for a “mobile-first network”—one in which wireless traffic comprises the critical network mass and bandwidth priority—and for simplifying enterprise networking.

Personalizing is based on key attributes, such as user types, device types, location, application types and other factors (sometimes referred to as “context”), which are then used to automatically determine and enforce network policy. Automating policies in this way vastly simplifies network deployment and operation. It also avoids the need to manually create separate networks for managing every possible combination of attributes. Figure 2 shows how personalization can simplify network architecture.

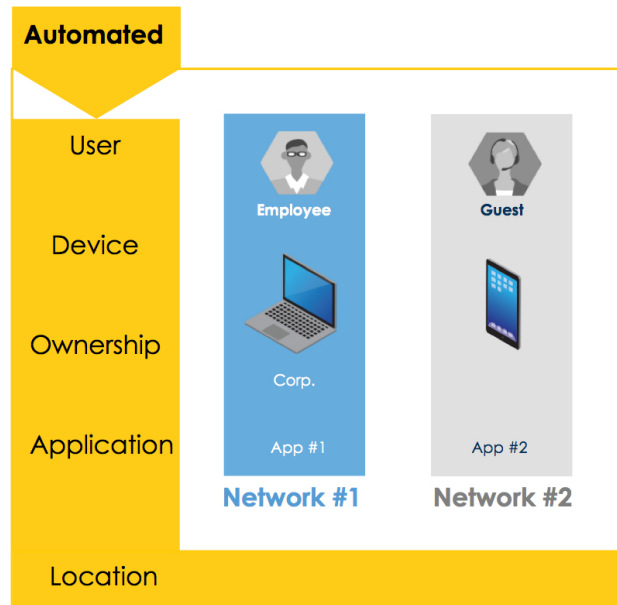


Figure 2: Much simpler. A network that performs personalization requires minimal SSIDs. The infrastructure automatically discovers context and applies policy in real time.

Personalization simplifies enterprise networks. However, it also requires significant processing capability, which affects the network infrastructure. Because of this fact, architectural approaches have changed to increase processing capabilities and future-proof against technology obsolescence.

The Network Impact

Network data rates *will* continue to increase, and 802.11ac is only the latest innovation. Mobile workforces and BYOD, combined with huge capacity and throughput improvements, fundamentally affect network architecture. A poorly designed network results in:

- Cost overruns as networks become overbuilt in an attempt to accommodate high throughput and the necessary processing power
- High complexity, particularly for organizations with multiple remote locations
- High support requirements as newer, faster technologies are incorporated into the network

Protect Your Network's Future

The processor-intensive nature of modern mobility is likely to soon obsolete most existing WLANs. To automatically personalize application service to a user requires the access network to perform three primary tasks:

1. **Understand the user context:** The access network must recognize the user's identity, device, and ownership of that device, access location, time of day, request for application access, and any other discernable characteristics. Any additional information that can be gleaned about the user enhances context and allows for a more personalized experience. There are several ways to accomplish this, but the most reliable and accurate approach incorporates deep packet inspection.
2. **Understand application traffic flows:** The access network must be able to identify the application traffic coming from the user's device in real time. Context can vary over time. The user might initiate a VoIP call one minute, send email the next, then try to access Netflix after that. Packet inspection cannot just statistically sample current user activity. The network must make real-time determinations on actual use without throttling traffic.
3. **Automatically respond:** The access network must be able to automatically make real-time policy decisions based on user context. It then must apply proper quality of service and security to applications to optimize the experience for the specific device being used. All at Gigabit rates! Since 802.11ac is already pushing data rates beyond 1Gbps, and future technologies promise even more throughput, being able to do this now and in the future is critical.

Why is This Problem Different?

How is this challenge different from the age-old "controller as bottleneck" hype? Wi-Fi technology has evolved significantly since enterprise WLANs became important. Legacy WLAN architectures have not evolved appreciably since 2002. Then, WLANs were designed with a centralized controller box and lightweight access points (APs) at the edge. This was because, at the time, enterprise LAN speeds were higher than WLAN speeds, so AP connectivity to the controller did not create bottlenecks. Moore's Law favored the centralized controller box—high-speed processors and encryption co-processors could handle encrypting and decrypting packets at wire speed.

But mobility means that time, place, application usage, device type, and ownership are dynamic and can't be processed by hardware. Personalization requires intelligent decisions to be made in software, on the fly—based on traffic flows and sometimes even packet by packet. Now add high throughput to every client, thanks to 802.11ac, and the magnitude of processing power required becomes clear.

Another aspect of Moore's Law is chipset miniaturization. For example, instead of keeping a big chip and adding more capabilities or capacity, you keep the capabilities the same and use a smaller chipset. As a result, 802.11ac is implemented on tiny, low-power chipsets that didn't even exist when WLAN controllers were invented. And for that matter, neither did today's smartphones, tablets, and embedded systems. Not only are today's mobility requirements more processor-intensive, there will be far more devices that operate with Gigabit throughput over the air.

Two Architecture Approaches

If the wireless network will be the primary network for mobile users with high 802.11ac data rates, where will the bottlenecks be? Knowing this is critical before making significant investment in your network architecture.

Legacy WLAN Controller: The Over Build Strategy

In Figure 3, we depict how legacy WLAN controller architectures make policy application decisions for mobile users.




Figure 3: Legacy WLAN controller makes mobility decisions by centralizing processing at a single point

In Figure 3, the red line represents the sender. The controller makes the decision to apply a specific policy to a traffic flow based on user context. Decision processing happens at a single point—the controller. This is a perfectly valid approach and will work. The disadvantages with this implementation becomes clear when trying to optimize service levels and retain today's network investments when new technology emerges. IT architectural decisions are not

isolated to today's needs. They must also account for future needs. From a return on investment perspective, this approach has two primary disadvantages:

- **Increases complexity and limits scalability:** Purchasing a centralized controller requires a deep understanding of traffic usage in order to “right-size” the system. If you buy controller capacity for today's usage and traffic, scalability will be limited. As a result, overall network throughput will be compromised and the end user will bear the brunt. If you over-buy controller capacity, you can ensure a good user experience but the solution costs increase significantly. In addition, you will need to make the decision for every enterprise location, increasing complexity and deployment costs.
- **Exponentially increases costs:** Many organizations are preparing their networks to take advantage of 802.11ac data rates when they arrive. With a legacy network architecture, “future-proofing” is possible. However, IT departments will have to predict their future capacity needs in order to right-size controller hardware. Because it's impossible to precisely predict the future, you will want to build in appropriate over-capacity to account for future processing needs based on 802.11ac. All future capacity must be purchased upfront to avoid hardware upgrades before it has out-lived its depreciation cycle. Once again, predicting future usage for every remote location simply makes decisions more difficult and costs higher.
- **Introduces risk:** With technology, there is always the additional risk that a wireless networking breakthrough during the controller hardware lifetime could make it obsolete. For example, new software requirements could make the capacity unusable. The risks of over-building are well understood, as IT departments learned that lesson while deploying wired networks. Today, much wired switch capacity that has been deployed is being left unused because WLANs have replaced wires as the primary access network.

Distributed Network Processing—Rightsizing for Mobility

A second approach is to employ a distributed system. In a distributed system, there is no single point that makes *all* personalization decisions. Each AP can make real-time decisions based on a user's context. The AP becomes responsible for personalizing network performance and policy to the users that are connected to it. These APs are designed to manage more than 100 users—more than enough to handle traffic in the area it is covering with Wi-Fi signal². In

² This assumes that the Wi-Fi network is designed properly and the correct number of APs is used to cover exceptionally high user density areas.

a distributed Wi-Fi deployment, parallel decisions are carried out simultaneously at the edge of the network without overwhelming any single processor.




Figure 4: Distributed processing reduces requirements on each location and helps ensure that processing loads will not exceed capacity

Figure 4 illustrates how a decision is made in a system that distributes personalization decisions to local APs. In this system, policies are centrally managed through the cloud. Traffic-handling decisions are made locally, at the network entry point, which is the AP. These APs coordinate with each other to ensure that correct policies are applied to users and their traffic, even if they wander from AP to AP.

A distributed approach is ideal for a mobile-first user base for many reasons:

- **Migrate to new technology at your pace:** Rightsize your WLAN and only buy what you need when you need it. With a distributed system, there is no need to predict the future. No need to overbuild any part of your network with excess capacity “just in case.” Simply integrate new technologies where you need them, when you need them.
- **Reduce capital expenses:** A distributed system is additive. Every AP added to the system adds processing power. The more Wi-Fi coverage you need, the more APs you naturally need. That’s the only decision you must make. Because no single AP controls all decisions, the system is natively redundant without additional expense.
- **Optimize user experiences:** With no bottlenecks or single points of failure, you easily optimize the user experience. Everyone receives a customized user experience when they attach to the network, whether you have 1, 5, 500, or 50,000 users.

The Network Impact of 802.11ac

- **Future-proof your infrastructure:** Finally, a distributed system provides a simple, future-proof migration path to newer, faster Wi-Fi technologies. If a new Wi-Fi technology arrives, just replace the APs where higher-speed technology is required. Legacy APs can be easily moved to less mission-critical areas, such as guest lobbies or commons areas, where they can continue to provide service.

Summary

Managing a network for your company is a complex proposition. Managing a mobile-first community of users is exponentially more complex. Forward-thinking organizations are implementing wireless access networks that automate security and quality of service policy decisions to personalize access based on a users' context. A network that automatically personalizes wireless service vastly simplifies the task of optimizing a mobile workforce. Typically, this level of automation also requires significantly greater processing power. And controller-based architectures designed to achieve higher power often lead to:

- High costs, due to overbuilding the network for overcoming network bottlenecks and predicting future needs
- Increased complexity, due to having to make diverse capacity decisions and adding hardware in multiple locations
- Supportability issues when newer Wi-Fi technologies are introduced into the network

Mobile-first networks require architecture built from the ground up for WLAN access. Distributing processing and coordinating across APs:

- Saves money, by avoiding overbuilding and allowing you to pay for only what you need when you need it
- Simplifies enterprise network complexity, by automating policy decisions and personalizing each user's experience
- Future-proofs the WLAN by allowing new high-speed network technology to be easily integrated without obsoleting other network infrastructure.

About Aerohive

People want to work anywhere, on any device. IT needs to enable them—without drowning in complexity or compromising security, performance, reliability, or cost. Aerohive's mission is to Simpli-Fi access networks with a cloud-enabled, self-organizing, service-aware, identity-based infrastructure that includes innovative Wi-Fi, VPN, branch routing, and switching solutions.

Aerohive was founded in 2006 and is headquartered in Sunnyvale, Calif. The company's investors include Kleiner Perkins Caufield & Byers, Lightspeed Venture Partners, Northern Light Venture Capital, New Enterprise Associates, Inc. (NEA) and Institutional Venture Partners (IVP). For more information, please visit www.aerohive.com, call us at 408-510-6100, follow us on Twitter @Aerohive, subscribe to our blog, join our community, or become a fan on our Facebook page.



Corporate Headquarters

Aerohive Networks, Inc.
330 Gibraltar Drive
Sunnyvale, California 94089 USA
Phone: 408.510.6100
Toll Free: 1.866.918.9918
Fax: 408.510.6199
info@aerohive.com
www.aerohive.com

International Headquarters

Aerohive Networks Europe LTD
The Court Yard
16-18 West Street
Farnham, Surrey, UK, GU9 7DR
+ 44 (0) 1252 736590
Fax: + 44 (0) 1252 711901