

White Paper


The Network Impact of 802.11n



Aerohive Networks, Inc.
3150-C Coronado Avenue
Santa Clara, California 95054
Phone: 408.988.9918
Toll Free: 1.866.918.9918
Fax: 408.492.9918
www.aerohive.com

The Network Impact of 802.11n

802.11n is delivering on promises to bring revolutionary advances in throughput and capacity to the wireless LAN. For the first time in the history of the development of IEEE 802.11 networks, wireless LAN speeds are comparable to commonly used wired technologies. Now that wireless LAN users have access to speeds well in excess of 100 Mbps, wireless LANs can no longer be treated as an afterthought. Overlay approaches of years past that treated wireless access as subordinate to the wired LAN are no longer feasible given the speed parity and the expectation of users that 802.11 is the default method of connecting to a network. With the adoption of 802.11n, concerns over the capabilities of wireless LAN technology have become concerns that affect the whole network, not just the first hop across the radio.



Before successfully migrating to 802.11n, there are several questions that must be considered. These questions include:

- How do 802.11n APs connect to the existing core network?
- How are 802.11n APs powered?
- What changes are required to the network backbone to support the higher speeds of 802.11n?

802.11n and the Core Network

Early wireless LANs were designed and built primarily as convenience networks, and were commonly implemented as overlays to an existing wired network. Early technologies were slow and seen by users as unreliable and insecure. IT managers resisted the use of wireless LANs and kept mission-critical data away from the wireless LAN. At the start of wireless LAN revolution, design was minimal and had little impact on the core network, and almost no attention was paid to traffic flows on the wireless network because they were small and low-speed.

802.11n provides the final step in the evolution of wireless LANs from convenience technology into core technology for network connectivity. Performance has increased from a few megabits per second to data rates that exceed common Fast Ethernet desktop connection speeds. The industry has worked to enhance reliability, enabling wireless LANs to become the primary connection technology, even for power users that demand bulletproof dependability for mission-critical applications. The “all-wireless office,” once a glimmer in the eyes of the dedicated engineers bringing 802.11n to market, offers a future network where users replace “plug and play” with “open and play” – that is, users walk into an office, open a laptop, and begin working.

802.11n APs and the Core Network

Yesterday’s 802.11a/b/g AP maximum data rate is 54 Mbps. Like all contention-based medium access protocols, the 802.11 MAC is not perfectly efficient. Radio-based protocols have intrinsic overhead to ensure reliable transmission, and the actual TCP throughput experienced by users will be less than the peak transmission rates used by the MAC.

802.11n achieves high speeds both by increasing the data rate used to transmit frames, as well as improving the efficiency of airtime utilization. Readily available 802.11n APs can reach data rates of 300 Mbps, while improved efficiency gains can make the TCP throughput to the user reach as high as 200 Mbps.

Total network throughput is affected by co-channel interference, which reduces the capacity available to APs with overlapping coverage areas. The 2.4 GHz band consists of 83 MHz of spectrum, and needs to support legacy 802.11b/g access. Therefore, most deployments of 802.11n in the 2.4 GHz band retain the three non-overlapping 20 MHz channels traditionally used for 802.11b/g. The following table shows how the different technologies compare, which offers rough guidance on the performance of the network as a whole.

Technology	Maximum Data Rate	TCP Throughput
Channel Bonding + 2 x Spatial Stream .11n (w/ Short Guard Interval)	300Mbps	~150Mbps
Single Channel 2 x Spatial Stream .11n (w/ Short Guard Interval)	144 Mbps	~72 Mbps
Single Channel 1 x Spatial Stream .11n (w/ Short Guard Interval)	72Mbps	~36Mbps
802.11a or 802.11g	54Mbps	~22Mbps

*Note: These throughput numbers are quantities expected on well-designed networks with good signal-to-noise ratios and a preponderance of large frames.


A dual-radio 802.11n AP configured for maximum throughput will configure one radio in the 5 GHz band to use a 40 MHz channel and a second radio in the 2.4 GHz band using a backwards-compatible 20 MHz. To prevent a bottleneck from forming at the interface to the core network, the uplink must support a total throughput of approximately 220 Mbps, though support for 802.11g would drop the required throughput slightly. These speeds are well in excess of a single Fast Ethernet link, which has important implications for the connection of APs to the network.

If gigabit links are unavailable, Fast Ethernet link aggregation allows network users to experience some of the benefits of the higher speed of 802.11n without requiring that network administrators perform a forklift upgrade to gigabit wiring closet switches. By using the both Ethernet ports on Aerohive HiveAP 300s to connect to two fast Ethernet switch ports, the 172-222 Mbps TCP throughput can be shared across two, aggregated 100 Mbps full-duplex ports.

Controller-Based WLAN Architectures and 802.11n

To make wireless LANs into an “enterprise ready” technology, the second wave of innovation focused on building WLAN controllers for management. Controllers are responsible for enabling mobility, enforcing user policy, and managing radio resources. In most network architectures, controllers are also responsible for handling data flows sent and received by wireless client devices. Networks built on controllers are often called “overlay” networks because all data destined to or from the wireless LAN must be handled by a controller’s forwarding process. In the early days of 802.11a/b/g enterprise wireless LANs, building a controller with the forwarding capacity to handle a network of low-speed access was not particularly difficult. As networks migrate to the high speeds of 802.11n (and other standards beyond), the double-switching of traffic forces the network core to be significantly higher-performance than necessary.

Data Flows Using a WLAN Controller



Controllers provide many functions by virtue of being in the data forwarding path for the entire network. Providing security for an entire network is easier to implement as a centralized process on the controller, but doing so requires that the controller's cryptographic capabilities match the demands of an entire network today as well as in the future. Because all data traffic to and from every wireless client flows through the controller, it must traverse the core network twice: once from the client through the core and on to the controller, and then again from the controller to its destination. Return traffic makes the same detour to the controller, again requiring multiple hops. As 802.11n drives higher speed at the network edge, the controller connection to the core is also under pressure to grow. Many high-end controllers now have multiple 10G Ethernet interfaces to handle the expected traffic load from 802.11n clients.

Tweaks to this architecture have begun to appear, in large part because the controller backplane has become a restriction on network growth. Common high-end security processors are limited to 10 Gbps of CCMP throughput, which may only be the throughput of 50 APs. Centralized controller-based forwarding may offer tens gigabits of throughput when connected to a robust core network, yet suffer from bottlenecks in security processing when 802.11 security functions are performed in the controller. As the deployment of 802.11n accelerates, the limitations available in centralized forwarding and security constrict the growth of networks in a silicon straitjacket.

In addition to the significant investment in core network equipment, controllers also undermine existing security infrastructure. Traffic which is placed in an encrypted tunnel between the AP and the controller in the core is hidden from any existing threat mitigation, intrusion detection and prevention, and quality of service engineering. To address the invisibility of traffic between the AP and controller, network managers must either deploy duplicate security systems behind the controller, or license network security features from the controller vendor that duplicate existing systems on the wired network.

PoE (Power over Ethernet)

Power over Ethernet (PoE) has emerged as the most elegant way of powering legacy 802.11a/b/g access points. PoE can be integrated into the distribution switch or can be provided by separate PoE injectors in the wiring closet. The first standard for delivering PoE is 802.3af-2003, which supplies 15.4 watts over a 100m Category 5 cable run. 802.3af was succeeded by 802.3at-2009, which supplies 25.5W. There are a variety of benefits to PoE including:

- **Cost** – Most APs are located in places, like ceilings, where power outlets are not typically found. The cost of running power into the plenum can be extremely expensive and often requires specialized labor and electrical conduit installation. In comparison, network cabling cost is minimal.
- **Mobility** – If devices are moved, new power runs are not required.
- **Resiliency** – Many wiring closets and data centers are equipped with battery-backed power, and are unaffected by power outages. Users with laptops and other mobile devices are able to access the network and any resources in the data center regardless of the state of utility power.

Dual radio 802.11n APs have higher power requirements than legacy APs because of the additional transceivers in the radio front-end and the Digital Signal Processing built into 802.11n. Most 802.11n APs are equipped with faster processors and more storage to take advantage of new protocol features as well.

- Network administrators generally decide to live within the 15.4 watt budget offered by 802.3af. Aerohive responded to this decision by introducing the HiveAP 100 series, which offers a combination of robust 2-radio 802.11n performance within the 802.3af constraint. When additional performance is required and the 3-radio HiveAP 300 series is used, most customers will be able to use the full capacity because 802.3af is designed to accommodate power loss over the cable run, and almost all real-world cabling runs are significantly shorter than the 100m maximum. Taken together, the higher power supply and the lower loss provide sufficient power to enable a HiveAP 300 series to operate at full capabilities .
- HiveAPs can be powered via a power adapter, which is particularly useful in a wireless mesh that does not rely on Ethernet backhaul.
- Aerohive's "Smart PoE" enables use of multiple PoE sources for network resilience. Dual-Ethernet HiveAPs may be connected to multiple power devices, whether on the same switch or to two different switches.

Aerohive's flexible power options overcome the issues with legacy switches in the wiring closet by decoupling the upgrade of wiring closet switches from the deployment of 802.11n.

New Technology Demands New Architecture

Transitions in the network technology have driven dramatic shifts in the networking business because new technologies have swept away obsolete architecture. Hubs replaced repeaters, only to be replaced by switches. Controllers swept aside individually-managed autonomous APs because they addressed the management problems of convenience 802.11a/b/g networks. Controllers were built with enough forwarding power to handle low speeds, and the controller was not considered a serious point of failure in comparison with the inexpensive autonomous APs it replaced. 802.11n exacerbates the underlying problems with the architecture, making the penalty in core upgrades to handle double-switching quite expensive.

802.11n has doubled the traffic-forwarding requirement on controllers, and as future generations of 802.11n technology roll out, requirements will double again. With each of these doubling steps in network capacity, controllers must be prepared to handle twice as much user traffic. Controller vendors depend on security processors to perform high-speed encryption and decryption at gigabit rates, but security processor vendors are losing the race to keep up with user demands for data service. Controllers strain to keep up with the load; even with 10-gigabit interfaces, controllers are constrained by security processors offering a fraction of that speed. In extreme cases, the encryption capacity of the controller may be less than 5 Mbps for each radio it supports. Such a slow data rate cannot


adequately support the 802.11a/b/g networks of five years ago, let alone grow to support the next generation of 802.11n.

Fundamental questions stemming from increasing bandwidth, potential bottlenecks and user expectations of high availability are not able to be answered given the current state of security processors. Controllers cannot grow as fast as user traffic is.

Aerohive's Cooperative Control architecture was designed for the high-bandwidth future when all applications view the wireless LAN as its primary transport. Therefore, Cooperative Control was built on the proven heritage from wired networking:

- Data traffic flows from wireless clients to the access point, then to the client's destination in a direct, open path, exactly like existing wired clients, and leveraging the existing infrastructure and investment that has been built out for the wired clients.
- Traffic does not flow across the core unless it must. Control traffic between APs is localized and flows only between access points that are in the same RF neighborhood, and user traffic does not need to traverse the core multiple times.
- Data traffic is distributed across the network without bottlenecks into and out of a single device. Traffic destined to local resources like printers and work-group servers traverses only the wiring closet and never touches the core.
- Wireless traffic is no longer opaque to the rest of the network, enabling the WLAN to benefit from security and QoS schemes already deployed. VoIP traffic destined for the wireless LAN can be treated with appropriate priority and handled carefully through the core, not just on the radio link.
- Policy enforcement can be provided at the edge of the network, where it is most effective, instead of in the core.

Data Flows Using Cooperative Control



What is Next Beyond 802.11n?

Wireless LAN data rates have exploded from 11 Mbps in 2002 to 300 Mbps today with second-generation 802.11n chips. Additional spatial stream capabilities in the chipsets will push data rates to 450 Mbps next, followed by the 802.11n “speed limit” 600 Mbps per radio. To cope with the continued user demand for higher data rates, the IEEE 802.11 working group has launched two task groups charged with developing gigabit wireless standards. Ethernet speeds followed a similar growth curve from 10 Mbps shared networks to 100 Mbps switched networks, which drove demand for gigabit Ethernet switches followed by 10G Ethernet. As network traffic streams course through the network, they join together and surge through the core. Network designers can only prevent the torrent of traffic from overwhelming the core by carefully distributing traffic at the edge. Just as with Ethernet in the previous decade, keeping local traffic local and only using expensive core technologies when necessary will be the key to enabling the 802.11 adoption curve.

Aerohive’s Cooperative Control architecture is reminiscent of past enabling innovations in the wired network. To fully provide scalability, wired networks needed to dynamically react quickly to local conditions based on local knowledge. Achieving the same rapid reaction in wireless LANs is incompatible with centralized core decision making. Cooperative Control gives wireless clients a network with high resiliency and throughput similar to what wired clients enjoy. Aerohive’s Cooperative Control approach allows networks to start small and grow as user traffic needs to, enabling IT managers to plan for new innovations and reduce the risk to the network by making performance and deployment more predictable.

Realizing the Benefits of 802.11n Without Costly Upgrades

Aerohive has demonstrated an approach to the 802.11n upgrade that enables migration with a minimal impact upon the existing network. By leveraging existing switching infrastructure at the edge, a costly switch upgrade can be avoided. By utilizing the same network forwarding resources used by wired clients, there is no impact to the network backbone as network utilization migrates from the wired to 802.11n wireless network. Aerohive-based wireless LANs grow organically in response to user demand because each HiveAP is designed to handle the load from its clients without support hardware such as controllers, with all of their accompanying complex network engineering, and licensing.

802.11n enables significant productivity benefits for mobile computing, while opening up new wireless applications; Aerohive enables the simplest-to-deploy and most cost effective 802.11n solution to get you there.