White Paper

The Network Impact of 802.11n

Aerohive

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 promises to bring revolutionary advances in bandwidth, throughput, and reliability to the wireless LAN. It will also have a perhaps unforeseen impact on the wired network, as well. Bandwidth, throughput and reliability will now become an end-to-end network issue, particularly in cases where the WLAN is deployed as network infrastructure rather than as an overlay.

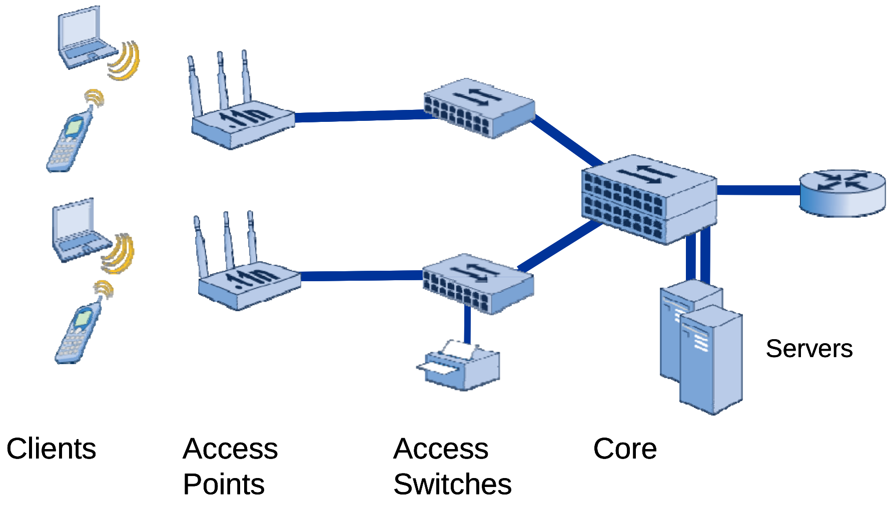

Because of differences between 802.11n and existing a/b/g wireless deployments, there are several new questions that must be considered in order to make migration seamless. These questions include:

- How do 802.11n APs connect to existing switches?
- How are 802.11n APs powered?
- What is the impact to the network backbone?

Switch Ports and 802.11n APs

Today’s 802.11a/b/g AP maximum data rate is typically 54 Mbps; because of the high amount of overhead that is intrinsic to WLANs, the actual TCP/IP throughput is usually significantly less than that. For example, with .11a and g, throughput is around 22 Mbps. Today’s 802.11n access point *data rates* can go as high as 300 Mbps for a single radio, with a peak TCP/IP throughput of around 150-160 Mbps. Because of the potential for co-channel interference, it is very unlikely you will be running an 802.11n AP at full capabilities – that is, with channel bonding in both 2.4 and 5 GHz spectra. In the 2.4 GHz range it is unlikely channel bonding will be used in order to maintain a 3 non-overlapping channel plan for deployment; in fact it’s most likely that in the 2.4 GHz range, an 802.11n radio will be supporting legacy 802.11g devices so channel bonding is unlikely to have any benefit anyway. Consideration of the data rate and actual performance of 802.11 a, g and n in different configurations is important, however, as it can effect 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)	**144Mbps**	**~ 72Mbps**
Single Channel 1 x Spatial Stream .11n (w/ Short Guard Interval)	**72Mbps**	**~ 36Mbps**
802.11a or 802.11g	**54Mbps**	**~ 22Mbps**

Please note – These throughput numbers assume large packets and good signal-to-noise ratios.

Based on this chart, a dual-radio 802.11n AP supporting .11n clients, running one radio in the 5 GHz range and one radio in the 2.4 GHz range simultaneously will realize a maximum throughput of approximately 222 Mbps. This presumes that the .11n radio is running on the 5 GHz band with channel bonding for a throughput of 150 Mbps. In the 2.4 GHz band, supporting .11n clients, maximum throughput is about 72 Mbps. The reason that the 2.4 GHz band is so much slower is that it is not feasible to run channel bonding in the 2.4 GHz range given that there are only 3 non-overlapping channels; 4 in Europe and many other parts of the world. If the radio running in the 2.4 GHz band were supporting .11g clients, the total AP throughput would drop to 172 Mbps. The result is a maximum TCP throughput of approximately 222 Mbps for an all-802.11n environment and 172 Mbps for an 802.11n + 802.11g environment. Armed with this information we are now ready to have a discussion around the wiring closet switch.

Many vendors, particularly those that also sell switching products, will insist that a gigabit link is necessary to support any kind of 802.11n AP, because this theoretical maximum throughput is higher than a switched fast Ethernet 10/100 port can handle without oversubscription. This argument can seem compelling because if an AP is running channel bonding, a single radio could push 150 Mbps in a 2 radio system. This would appear to far outstrip a 100 Mbps Ethernet link. While all of us are very aware that 100 BaseT Ethernet is capable of 100 Mbps full duplex, it is easy to forget that radios can either send or receive, but cannot do both at the same time. This means that instead of a duplex operation, radios have a simplex operation. Therefore the full 150 Mbps throughput discussed could, theoretically, run into a 10/100 port if the send and receive traffic were balanced. The reality of most enterprise clients, however, is that about 70 to 80% of the traffic is downstream towards the client, due to the mismatch in traffic from mail and web. That said, there are some notable exceptions, like network backup, that are largely upstream.

The chart below considers a 2 radio AP plugged into different switch media.

Wireless	Wired	Wired Bandwidth Available
802.11a + 802.11bg	**100BaseT**	Non-blocking in all cases
2 x 802.11n	**100BaseT**	Blocking, but still better performing than 802.11abg
2 x 802.11n	**2 x 100BaseT** (Link Aggregation)	Fully non-blocking if traffic is 80% downstream and 20% upstream
2 x 802.11n	**1000BaseT**	Non-blocking in all cases

Ethernet is clearly optimal from a pure performance perspective, Aerohive advocates link aggregation as the method of choice to quickly realize the benefits of 802.11n without having to consider a forklift upgrade to gigabit wiring closet switches. Using the two Ethernet ports that come standard on Aerohive 802.11n APs connected to two fast Ethernet switch ports, the 172-222 Mbps TCP throughput can be shared across two, aggregated 100 Mbps ports. In this manner, the enterprise can easily enjoy excellent WLAN performance without having to migrate switch infrastructures.

PoE (Power over Ethernet)

Power over Ethernet (PoE) has emerged as the most elegant way of powering today's 802.11a/b/g access points. PoE can be integrated into the distribution switch or can be provided via a separate PoE injector in the wiring closet. The current standard for delivering PoE is 802.3af, which allows for 15.4 watts per Ethernet link port over a 100m cable run. 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. The cost of running Cat 5 cable, however, is minimal and can be done by anyone.
- **Mobility** – New power runs are not required as you move devices from one area to another.
- **Resiliency** – Network devices that operate via PoE are not affected by power outages, assuming that the switch powering the AP has battery backup.

Dual radio 802.11n APs may have higher power requirements due to the additional transmitters, receivers and Digital Signal Processing, or DSP, that 802.11n demands. Several potential solutions exist for this, including some that are unique to Aerohive innovations:

- The majority of switches deliver greater than 15.4 Watts in the 802.3af spec. In addition, the majority of cable runs for APs are far less than the 100m maximum in the same spec. So in the vast majority of cases a single PoE switch port is sufficient to fully power an Aerohive HiveAP with two 3x3 MIMO 802.11n radios. For those cases where the switch delivers the lowest limit of the PoE spec and/or the cable is at its maximum 100m length, Aerohive's HiveAPs, using Smart PoE circuitry are able to automatically measure the power delivered and, if necessary, reduce the power being consumed by the HiveAP. The HiveAP first turns off the unused gigabit Ethernet port. If there is a further requirement to reduce power consumption, the HiveAP reduces both .11n radios to 2x3 MIMO mode, without reducing the transmit power. Together these power reduction procedures ensure a HiveAP can work with a single (802.3af) PoE connection.

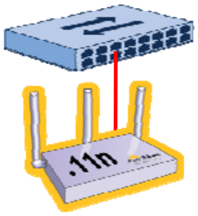

10/100 or GigE with standard 802.3af PoE

Automatic PoE power detection

- Aerohive 802.11n HiveAPs have two PoE ports and the Smart PoE circuitry is able to draw power from two PoE switch ports. This is more than adequate to supply the necessary power to the HiveAP, and, in most cases, is able to provide PoE redundancy. This can be done whether the two ports connect to the same switch or to two different switches.

10/100 or GigE & 802.3af PoE

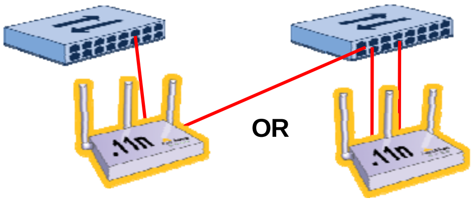

Dual 802.3af power

- New specifications, PoE+ and 802.3at, which are intended to meet the higher demands of 802.11n APs, and devices like video surveillance cameras with motorized zoom and swivel capabilities, will allow for 30watts of power delivery per port. The new 802.3at standard is due out in late 2008, although pre-standard high power switches and power injectors are already available.

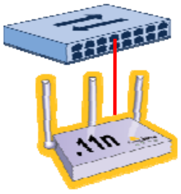

Gigabit switches with PoE+ or power injector

Full PoE+ power and Gibabit throughput

- HiveAPs, like most APs, can also be locally powered via a power adapter. If the HiveAPs are being deployed as part of a wireless mesh without any wires at all, then this would be the only available power option.

Aerohive has overcome the possible issues with legacy switches in the wiring closet. By decoupling the upgrade of the switch from the deployment of 802.11n, enterprises can avoid the costly upgrade to gigabit Ethernet until wired bandwidth demands it.

Network Backbone Impacts of 802.11n

Wireless LANs have typically functioned primarily as convenience networks, and therefore could conceivably be implemented as overlays to an existing wired network. WLANs have also suffered from the perception of being slow and/or unreliable, and often not been used to carry mission critical data as a result. Seen in this light, the WLAN has had little impact on the core backbone of the enterprise network, nor has the design of WLAN traffic flow been a primary concern.

This paradigm changes with the advent of 802.11n. As performance, throughput, and reliability increase, the WLAN becomes a viable alternative/companion to the wired network for high bandwidth and mission-critical applications. At this point, it becomes very important to understand the flow of traffic over the WLAN. There are two types of traffic which must be considered, and they can be intertwined - data traffic, the real data sent and received by the wireless clients; and control traffic, the method by which the access point enables mobility, RF control and maintains user policy as clients roam. These traffic flows will depend primarily on the architecture used for the wireless system, and can result in widely varying actual WLAN performance. Configurations that were acceptable in a low-throughput convenience network may not even function with higher performance demands. And as the WLAN performance begins to become a primary consideration of the network performance overall, these configurations may have unforeseen impacts on the wired network as well.

Controller-Based WLAN Architectures and 802.11n

If the WLAN access points utilize a centralized controller, that controller is most often in the data path for ***all*** the wireless traffic generated through the access points that are within the controller’s domain. *All* data traffic to and from *all* wireless clients ultimately flows into and then out of the controller. By definition, the controller is also the source and destination for all access point control traffic. Instead of taking the most direct open path between the source and destination, as a typical wired network would, traffic is “double-switched,” going from the wireless client to the AP and then to the controller; then from the controller on to its ultimate destination. The same double-switched path is taken for traffic in the other direction. Some controller vendors have started tweaking this architecture by enabling local forwarding even though there is still a central controller; however these hybrid approaches are in their infancy and force serious compromises in functionality in order to be enabled. For example, approaches may break inter-subnet roaming or bypass policy enforcement normally done by the controller. Often these approaches are not recommended for real world deployment even by the same companies that sell them. For the rest of this section, then, we will concentrate on the primary deployment model used by all major wireless LAN vendors – centralized, controller-based data forwarding.

Centralized data forwarding seemed to work well in convenience 802.11 a/b/g wireless networks when network speeds were much lower, but as more 802.11n bandwidth is utilized this centralized approach has a huge impact on backbone links, traffic engineering, and the controller itself. *Imagine if you were to build a routed/switched network today in which all of the data and control traffic generated by your network users flowed into and out of a single appliance.* With 802.11n enabling bandwidth intensive applications traditionally reserved for the wired network to now be routed over the WLAN, this will become, in essence, how controller-based architectures will work.

This issue will only be magnified as 802.11n becomes more widely deployed throughout the network. Ironically, many enterprises will deploy 802.11n to reap the over-the-air reliability benefits it provides, yet a controller-based configuration is essentially built around a single point of failure.

There are additional impacts of a controller-based architecture. While some of these concerns are not intrinsic to 802.11n, they may not have been obvious in 802.11a/b/g deployments due to the lower performance expected or the non-mission-critical nature of such deployments.

- Traffic tunneled between the AP and the controller is rendered “invisible” to 3rd party threat mitigation, intrusion detection/prevention, firewalls and traffic engineering.
- Increases in the client access bandwidth brought by 802.11n are aggregated and multiplied onto a very few backbone links, creating bottlenecks on the routers/switches and ports in the path to and from the controller. To illustrate this issue many controllers have multiple 10 Gigabit Ethernet interfaces to handle the traffic load, with no mention of what you can hook that up to.

New network designs are required to avoid single points of failure for the controller as well as the paths to and from the controller.

Data flows using a WLAN Controller

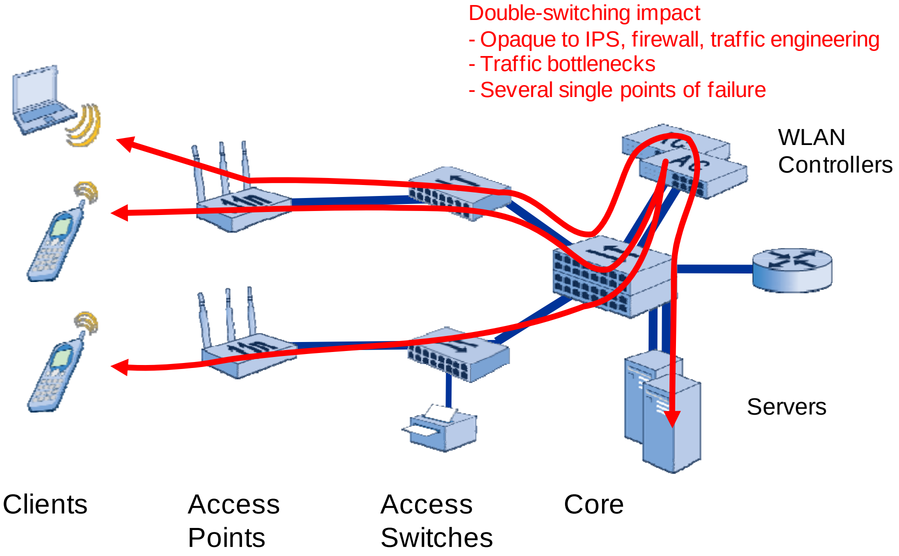

Higher Speed Demands Better Architectures

Like its wired predecessors that used repeaters before switches, migration to higher speeds requires architectures that can handle them. The controller-based approach to WLAN systems with its double-switching of client data, creation of a bottleneck to the surrounding wired network, and inherent single point of failure was tolerable for 802.11 a/b/g convenience environments but will not scale with the traffic multipliers and availability expectations brought on by 802.11n.

The fundamental questions of increasing bandwidth, potential bottlenecks and availability issues are not answered by creating even bigger single points of failure with even larger controllers. Luckily, the wired network already being used provides the model for approaching the issue.

Aerohive's Cooperative Control architecture offers all of the wireless functionality promised by 802.11n, but without the necessity of a controller. The results include:

- Data traffic flows from wireless clients to the access point, then to the client's destination in a direct, open path, just like the wired clients, leveraging the existing infrastructure that has been built out for the wired clients.
- Control traffic is localized and flows only between access points that are in the same RF neighborhood.
- No required "double-switching," tunneling or single points of failure. Traffic is just as it is on your wired network.
- The data traffic from higher speed radios is distributed across the network and not bottlenecked in to and out of a single device. For example, traffic destined to local resources like printers and workgroup servers may only traverse the wiring closet, never hitting 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.
- Policy enforcement can be provided at the edge of the network, where it is most effective, instead of at the controller.

Data flows using Cooperative Control

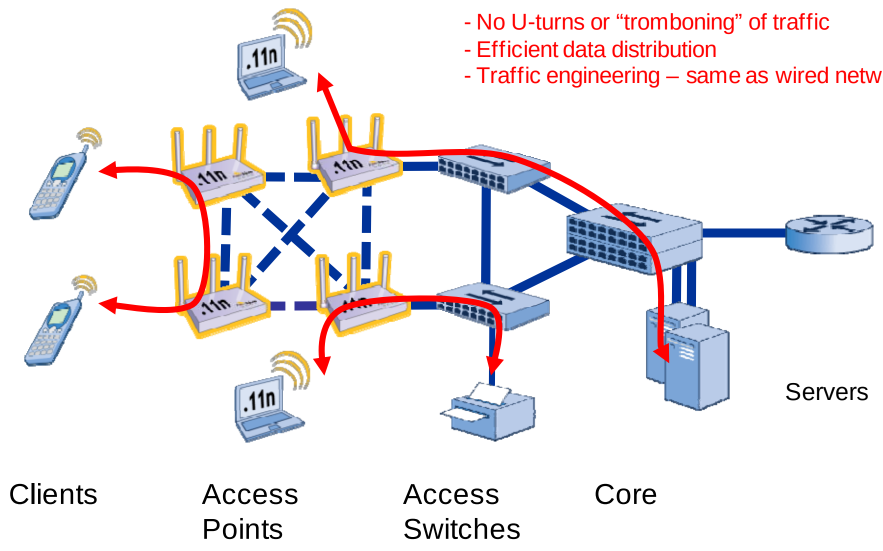

- No U-turns or "tromboning" of traffic
- Efficient data distribution
- Traffic engineering – same as wired network

Are we done at 802.11n?

As we contemplate the backbone impacts of 802.11n today, we need to also think about upcoming performance innovations on the horizon. We need only look back a few years to see such performance increases, as when 802.11g took 802.11b from 11 to 54 Mbps data rates, or when APs started supporting concurrent 802.11a and g, resulting in 2 x 54 Mbps data rates. Plans exist today for additional spatial stream capabilities in the chipsets, which will improve 802.11n performance from 300 Mbps to 450 Mbps, and eventually 600 Mbps per radio. To say that 802.11n will be the last word on performance improvements would be foolhardy. But we have been down a very similar adoption curve before with wired clients having moved from 10 Mbps shared to 100 Mbps switched networks, necessitating Gigabit and 10Gb switching and routing infrastructures to keep up. Key to this adoption curve was keeping traffic localized through switching and routing and having multiple backbone paths. The fact that the possibility for such growth was built into the architectures will ease the migration forward to high speed WLAN networks.

Aerohive's Cooperative Control architecture is reminiscent of past enabling innovations in the wired network, including dynamic routing and Spanning Tree. By eliminating the controller, Aerohive allows wireless clients to leverage similar network design benefits, including bandwidth and resiliency, that wired clients enjoy. To the IT manager, the impacts of 802.11n and beyond can be seen in the identical light as an upgrade of a wired client's bandwidth. Aerohive's Cooperative Control approach makes it easier to plan for such innovations, and their impact on the network as a whole more predictable.

We've also seen that while the drawbacks of centralized controller-based approaches can be overlooked in convenience 802.11a/b/g architectures, they will provide unavoidable roadblocks when faced with the significant increases in bandwidth promised by 802.11n. The domino-like impact to both the controller and the network devices in-between may be spun by switch or controller vendors, but the fact is that it will have very real implications on higher speed network design based on what has already been learned. The migration from "fat APs" that provided no coordination or centralized administrative control was to a purely centralized model. Now, with Aerohive, that model evolves to capture the benefits of both cooperative control and centralized administration without the necessity of a controller or the negative impacts it has on data and control traffic flows.

Realizing the Benefits of 802.11n Without Costly Upgrades

Aerohive has demonstrated an approach to the .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.

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.