The Promise and Pitfalls of 802.11n

Next generation wireless LANs offer opportunities but pose new challenges.

802.11n sounds like a panacea - delivering wireless data networking rates that are four to six times faster than earlier 802.11a/g networks and also improving transmission range. But with any new technology, opportunity is accompanied by a few new challenges.

This white paper discusses the pitfalls that might arise and how to overcome those issues such as increased power output, multipath signal propagation and "smart" antenna designs which affect range and performance. Also, existing wired LANs may not support aggregate 802.11n traffic volumes and 802.11n's extended coverage range and new frame format may introduce new security issues.

Table of contents

FLUKE networks

Executive summary

Businesses and consumers everywhere have grown excited about 802.11n wireless LAN technology. 802.11n is a set of draft IEEE standard specifications used for designing products that deliver wireless data networking rates that are four to six times faster than earlier 802.11a/g networks and also improve transmission range. 802.11n products accomplish this by using multiple transmit and receive antennas, spatial multiplexing, channel-bonded 40MHz operation and frame aggregation (see box, "Technically Speaking," page 3). Because early products typically deliver data-connect rates of 300Mbps and at least 100Mbps of actual throughput, 802.11n is the fair-haired favorite to displace wired Ethernet connections and enable all-wireless LAN access networks. All-wireless networks are likely to emerge first in organizations with highly mobile user populations that run bandwidth-intensive streaming and interactive multimedia applications.

Though the final IEEE 802.11n standard will not be formally ratified until year-end 2009 , there are pre-standard consumer- and enterprise-class products available from nearly all the major wireless suppliers today. These products comply with Draft 2.0 of the emerging standard, and vendors are banking on that specification not changing significantly. For an added measure of reassurance, the Wi-Fi Alliance, the prominent WLAN industry consortium, has already certified numerous products based on 802.11n Draft 2.0 for basic interoperability.

802.11n sounds like a panacea. However, as with any new technology, opportunity is accompanied by a few new challenges. Among the pitfalls that might arise in early deployments:

- New variables, such as increased power output, multipath signal propagation and "smart" antenna designs, affect range and performance. Consequently, they also affect a business's associated AP placement, RF management strategies, optimization and troubleshooting strategies.
- The capacity of existing wired LANs must be sufficient to support aggregate 802.11n traffic volumes.
- Some 802.11n WLANs require greater sustained power than today's standard power-over-Ethernet (POE) equipment supplies to operate at their maximum potential.
- New security considerations come into play associated with 802.11n's extended coverage range and new frame formats.

The remainder of this paper will provide some industry context and background, then address each of these challenges in turn.

Early applications

To date, pre-standard 802.11n networks have been installed most frequently in universities, which are technologically competitive and where highly mobile student users are likely to have the latest 802.11n client connections in their laptops. Some university and business campuses have also used the wireless mesh technology supported by many early 802.11n products for backhaul applications: wireless LAN traffic can be forwarded over the air from access point to access point (AP), both indoors and outdoors, until reaching a wired Ethernet switch at the edge of the wireless network. This eliminates the need to connect every AP to an Ethernet switch using physical cabling, saving significant labor costs, switch-port costs and time.

In addition to universities, 802.11n is also prevalent in healthcare organizations. Medical environments frequently need the greater bandwidth to download high-resolution medical images and to allow mobile caregivers to update electronic patient charts as they move from room to room. Meanwhile, the Dell'Oro Group research firm expects that by the end of this year, nearly 40 percent of wireless routers shipped will support 802.11n, rendering the technology more mainstream.

Performance monitoring and testing tools are an especially important component for early enterprise installations. One reason is that 802.11n has yet to be proven on a very large scale while supporting a large mix of client types and various combinations of the vendor-optional features allowed by the standard.

Basic operations

802.11n can operate in either the 2.4GHz or 5GHz unlicensed spectrum bands and has been designed to be backward compatible with earlier 802.11 networks. This means that existing 802.11a/b/g clients can associate with new 802.11n APs and continue to send and receive packets.

Technically speaking

802.11n uses a number of new technologies and concepts to derive Ethernet-like performance for mobile users. The primary advances used to improve the performance and range of wireless LANs in 802.11n include the following:

- Multiple input, multiple output. MIMO indicates that multiple transmitting and receiving antennas are in use. With more receive antennas, the system experiences combining gain, whereby more copies of the same signal are transmitted, increasing the signal-to-noise (SNR) and strengthening the signals.
- Multipath. Radio signals bouncing off walls and other obstructions create copies of the original signal. 11n benefits from multipath in that multiple reflections of the RF signal arrive at the receiver at slightly different times, increasing the receiver's ability to recover the message information from the signal.
- Spatial multiplexing. A technique whereby multiple antennas separately send different flows of individually encoded signals called spatial streams over the air in parallel, "multiplexing" the signals to shove more data through a given channel. At the receiving end, each antenna sees a different mix of the signal streams. In order to decode them accurately, the receiving device needs to separate the signals backout (or "demultiplex" them).
- Channel bonding. Blending two 20MHz channels to 40MHz-wide channels to double the effective data rate.
- MAC/frame aggregation. Bundling multiple frames together to reduce transmission overhead.
- Block acknowledgements (ACKs). Acknowledging a block of packets, instead of individual packets, identified by a beginning and ending sequence identifier. Block ACKs improve spectral efficiency.

Depending on how the 802.11n network is configured to support legacy clients and newer 802.11n clients, though, there may be a performance tradeoff for the 802.11n clients. For example, running 802.11n in mixed mode—with 802.11a/b/g and n clients all served by 802.11n APs and sharing a single band—will decrease the throughput improvement for 802.11n clients by approximately 30%, according to recent industry test measurements. The primary reason is that in mixed-mode operation, clients must decode 802.11a/b/g preambles tagged onto the 802.11n preamble in order to "speak" to the 11n AP. The use of the 802.11a/b/g preamble introduces significant overhead and slows down the overall network. In addition, mixed-mode deployments can't leverage the 40MHz option available in pure 802.11n networks, because legacy clients weren't designed to operate across 40MHz-wide channels. This situation further reduces available capacity.

All this being said, however, most 802.11n use is expected to wind up in the 5GHz band, which offers many more non-overlapping channels to better accommodate channel bonding and to support greater numbers of users. The 5GHz spectrum supports up to 27 non-overlapping channels worldwide and 12 in the U.S.; the ultimate number is dependent on each country's regulations and the frequencies supported by each vendor's products. By contrast, the 2.4GHz band supports just three non-overlapping channels worldwide.

For the foreseeable future, the 5GHz band will be less cluttered than 2.4GHz, which is home to a number of wireless and microwave devices already. Many early implementers are thus deploying in the 5GHz band, where their traffic should be less prone to interference and resulting performance problems. Some have eased 11n into their environments by deploying it in 5GHz as a wireless backhaul technology, described earlier, while using already widely deployed 2.4GHz 802.11g/b clients for user access.

Performance variables

As noted, the promise of 802.11n for liberating users while accommodating their multimedia application and communications needs is great. But there are many variables governing how a given system will perform, particularly at the scale expected for 802.11n all-wireless networks. Some of these variables are the same as those that have affected earlier wireless LAN iterations: the mix of client devices in use, construction materials of the building, whether floors are open areas or divided by lots of walls, and RF management techniques and tools provided by the vendor and third parties.

In addition, with 802.11n, the number of transmitting and receiving antennas working in parallel (called multiple input, multiple output, or MIMO) and the number of data streams supported by the system play a significant role. The phrase "NxN" is used to describe the number of antennas at each end of the 802.11n transmission. The minimum configuration required by the emerging IEEE standard is 2x2: two transmitting and two receiving antennas operating concurrently.

Early simulated and real-world performance tests documented by WLAN chipmaker Atheros® Communications show about a 20% increase in average performance in two-stream systems using a 2x3 MIMO configuration compared to a 2x2 MIMO configuration in a 20MHz-wide channel. They also show that average uplink throughput rates when transmitting across 40MHz channels (two, 20-MHz bonded channels, which the 11n standard allows) can be up to 40% greater in the 2x3 configuration than in the 2x2 configuration over distances of 30 to 40 feet and 20% greater in the 60- to 100-foot range .

Still, large-scale testing of enterprise 802.11n networks has yet to be done. It's an unknown how 802.11n performs on a grand scale with thousands of APs supporting fully loaded multimedia applications, because this environment hasn't completely unfolded yet.

RF management considerations

In some ways, RF management has matured while, at the same time, new RF management considerations have emerged along with 802.11n.

For example, many vendors of WLAN systems that use microcell architectures (those that tune APs to alternate channels in a checkerboard-style layout), now have embedded tools in their centralized controllers that automatically adjust power output in the face of interference or move traffic to alternate channels if necessary.

There are also so-called "fourth-generation" WLAN networking products that place all APs on a single channel by giving them all the same MAC address, which corresponds to that of a centralized controller. In these configurations, the controller—rather than the client devices—makes load-balancing decisions based on network-wide conditions about which AP a client should associate with and when the client should shift to another AP. Such architectures profess to eliminate co-channel interference, at least with other devices in the operator's own network.

There is still interference to worry about from other devices and nearby networks, particularly in multi-tenant buildings and particularly as 802.11n extends coverage to twice the range of earlier WLANs at a given speed. There are emerging technologies that allow the network administrator to set up barriers around a given building to more or less define the perimeter of the wireless network, but these have challenges in offices located in the interiors of multi-tenant dwellings.

Access point placement and migration

There is a bit of debate about whether physical site surveys are required for wireless LANs. The point of a site survey is to figure out how many APs to install and where to place them in order to provide a minimum throughput rate with adequate coverage throughout the building. Fortunately, extremely sophisticated tools are emerging that allow enterprises to perform site surveys electronically by feeding an application information about the layout of the building and its construction materials and then programming in the desired coverage, minimum data rate and received signal strength indication (RSSI). Not doing such a survey – either physically or electronically or, more typically, both – results in coverage holes.

Some of the more recent tools are very accurate in specifying where to places APs to get the desired coverage and throughput. Note, though, that it is not uncommon that there might be environmental conditions unaccounted for in the building's blueprints that require an occasional visit to a physical site.

There are other approaches, too, to determining AP placement, though they are losing muster. Especially before the advent of 802.11n, one argument has been that APs are inexpensive enough that enterprises come out ahead by skipping the site survey and tossing up an AP here and there as additional capacity is needed. On the other hand, new 11n APs are currently about 2 to 3 times as expensive as 11a/b/g APs (most are in the $1300 range). So installing them willy-nilly isn't quite the no-brainer with 11n as it might seem to be with older WLANs.

For the same reasons, when migrating older APs to newer APs, a simple AP swap-out could be an expensive approach. Doing so could be overkill, anyway, because 11n improves performance and range to the degree that it is possible to achieve the same performance with fewer APs. Or, the enterprise can opt to gain still better performance with the same investment in a system tuned to address the 802.11n system behavior.

Another consideration: Unlike its predecessors, 802.11n relies on multipath—the combination of an original transmitted signal plus duplicates created from reflection off obstacles during transmission—to enhance performance. The effects of multipath, however, will change the optimum layout of APs.

As noted, suppliers have begun offering tools that allow network administrators to import drawings of a building's layout into a predictive modeling tool and specify building construction materials for a high degree of accuracy on propagation models. Data sampling with automated tools further enhances accuracy of optimum placement.

And what about making use of both the 2.4GHz and 5GHz bands? Using both while transitioning from earlier WLANs can mitigate the performance impact of legacy 802.11 networks on the 802.11n infrastructure. One option is to run 802.11n traffic and 802.11b/g traffic in different bands. Most vendors sell dual-radio APs with one 2.4GHz radio and one 5GHz radio. Assigning all 802.11n traffic to the 5GHz channels and 802.11b/g traffic to the 2.4GHz channels (the only spectrum 11b and g operate in) helps maximize the 11n infrastructure's performance while continuing to serve 802.11g/b clients as usual.

In installations with 802.11a clients, which run in the 5GHz band, the 11a clients will communicate with 11n APs at 11a speeds, or 54Mbps, but, again, this impacts the speed of 11n clients.

Capacity of the wired network

Many enterprises currently have 10/100Mbps switches installed in their wiring closets. These have been adequate to aggregate traffic from 802.11a/b/g networks, which support actual throughput of up to about 22Mbps. However, aggregating traffic from dual-radio APs that support about 200Mbps per client (100Mbps per radio) requires faster upstream connections so as not to create a performance bottleneck. In fact, many 802.11n APs themselves support 1Gbps wired uplinks. As the 802.11n network begins to crank at full capacity, 10/100Mbps switches will likely need an upgrade to 1Gbps speeds. It stands to reason, then, that the proverbial domino effect will drive migration to 10-gigabit cores to accommodate the aggregation of 1Gbps traffic from the wiring closets.

Similarly, some of the WLAN vendor system architectures were originally designed with APs that relay each wireless transmission all the way back to a centralized controller, often located in the data center, for applying quality-of-service (QoS) priority markings and forwarding. Some of the vendors are re-thinking this architecture in light of 802.11n, whose traffic volumes could clog the Ethernet switches along the way. Perhaps more important, though, real-time interactive traffic requires a more peer-to-peer forwarding mechanism to minimize delays and sustain high-quality sessions.

Power implications

Deriving maximum performance from some 802.11n systems requires upgrading the existing power infrastructure. Today's IEEE POE standard for delivering power over Ethernet cabling, 802.3af, currently specifies power output of 15.4 Watts, sustainable at 12.95 Watts

over 100 meters. Some 802.11n access points, however, require greater wattage to run both 2.4GHz and 5GHz radios for maximum performance. The IEEE is working on a 30-Watt upgrade to PoE, called 802.3at, which might be ratified as early as 2009.

In the meantime, there are some options. One is to install APs with a single 802.11n radio in them, as single-radio devices are likely to operate within the power budget, and forfeit some capacity. Enterprises seeking a dual-radio implementation, with 802.11n capabilities in both the 2.4GHz and 5GHz bands, should be sure to investigate the vendors who say they can "do it all" with an 802.3af power infrastructure. Some may automatically disable some AP services, require two ports on a PoE switch and/or two cables, or sacrifice range to stay within the power budget and maintain performance. These tradeoffs might be perfectly acceptable, as long as the enterprise is aware that it is making them.

Another option is to use a power injector compatible with pre-standard 802.3at specifications, becoming available by a number of sources. Local powering of an 802.11n AP is an option, too, if a power source is available.

Monitor impact of 802.11n APs with EtherScope™ Network Assistant

The Fluke Networks EtherScope Series II Network Assistant excels at troubleshooting problems in a switched network environment. The EtherScope analyzer features Switch Scan to monitor network activity. Monitor the switches along the path from the network core to the 802.11n access point to understand the impact of an 802.11n AP on the capacity of the wired network. Measure and trend network utilization by switch port and slot. Use this information to configure the LAN for optimum throughput.

Security

802.11n systems require the use of industry-standard 802.11i security authentication and encryption. As such, they can be deemed inherently more secure than their predecessors, some of which used weak encryption methods such as Wired Equivalent Privacy (WEP) or even no security at all.

Still, there are a couple of new risks to be aware of. 802.11n is capable of transmitting twice the distance of earlier WLANs. This means its tendency to leak RF signals outside a building is potentially greater. If 802.11i is accurately deployed and the internal wired network is properly secured, leaked signals shouldn't be a problem. However, if even one AP has been misconfigured, it could be vulnerable to outside rogue devices.

Rogue devices operating in promiscuous mode can passively listen to all network packets passing by, regardless of destination address, which is a greater threat when signals travel outside the corporate perimeter where hackers might go undetected. Users of promiscuous devices can gather sensitive information such as user credentials or credit card information if the data hasn't been properly encrypted. The WLAN devices do this without emitting any signal of their own, so they are undetectable to wireless intrusion detection/prevention systems (WIPS). This is why regular audits that check security configurations are important.

The latest WIPS products are able to detect rogue (unauthorized) 802.11n devices while monitoring the airwaves and take appropriate action, depending on whether or not the rogue is connected to the enterprise network. One WIPS issue that rears its head with 802.11n, though, is that when using channel bonding to transmit across 40MHz channels, it will take a WIPS twice as long to scan the frequencies for malicious patterns as it did to scan earlier 20MHz channels. The situation effectively doubles the time a hacker has to penetrate a given frequency until the scanner makes its way around to that frequency again, from about 4 seconds to about 8 seconds. Driver exploits, which are 1- or 2-packet attacks, could be conducted in this time frame.

Driver vulnerabilities make a system open to administrative access by a hacker, and the WLAN industry's rush to get 802.11n technology to market has resulted in some vulnerable code. Tools are becoming available, however, that use a database of known wireless vulnerabilities to assess the versions of installed drivers and identify systems and specific drivers that are at risk to wireless driver exploit attacks.

Finally, IEEE 802.11n introduces a mechanism to acknowledge a block of packets, instead of individual packets, identified by a beginning and ending sequence identifier to improve network efficiency. At the time of this writing, the block ACK mechanism is not protected; an attacker can spoof one of these messages and create an enormous window within which frames can be sent with no ACK. In this way, they could potentially create an 802.11n denial-of-service (DoS) vulnerability.

Quickly identify and locate rogue devices

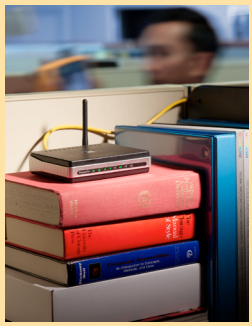

Fluke Networks EtherScope and OptiView analyzers help you address wireless security vulnerabilities. Walk your site to detect and identify unauthorized rogue devices and unprotected access points, then use the locate feature to hunt them down. Periodically survey the network for changes that could indicate a security breach. These portable analyzers discover active networks, mobile clients and access points. Drill down into devices to see configuration details. Troubleshoot WLAN connectivity, authentication and performance issues.

Troubleshooting and optimization

In legacy WLANs, the mobile client usually associates to the closest AP, which won't necessarily be the case in all 11n networks. As a result, applications and interference will behave differently with 11n. Monitoring tools are needed to discover and troubleshoot where and why WLAN network performance is sub-optimal and proactively address problems that could affect mobile application performance.

A number of organizations wish to add voice over IP (VoIP) to 802.11n wireless LANs, which will have plenty of capacity to support it. Such organizations will need automated tools to detect degradation of voice calls in the presence of data. While voice doesn't occupy much bandwidth, it is sensitive to delay, particularly that incurred when users roam among APs. In traditional microcell enterprise WLAN environments that place a population of APs on alternating channels, inter-AP roaming requires a handoff and reauthentication of the user's credentials.

The good news is that the IEEE recently ratified new extensions to the 802.11 standard, including 802.11r for Fast Handoff and 802.11k for Radio Resource Measurement. Along with both standard and vendor-enhanced QoS features for packet prioritization, call admission control and airtime fairness, these improvements will help ease the delay problems associated with roaming and data interference. Still, they are relatively new features; both pre-deployment and ongoing testing is necessary to ensure the quality of real-time conversations.

Summary

802.11n is greatly anticipated for the freedoms it will bring to highly mobile users of collaborative applications, many of which are multimedia in nature. To deliver four- to six-fold performance improvements or two-fold range improvements over earlier WLAN iterations, 802.11n specifications have been designed with a number of new properties. While backward compatible with earlier WLAN standards from an interoperability standpoint, the differences in the 802.11n standard alter the RF propagation landscape a bit. And running 802.11n in mixed mode with legacy clients in a single spectrum band will degrade the throughput improvements of newer 11n clients by about a third.

As a result, new issues rear their heads as enterprises move to 802.11n, most notably in the following areas:

- RF management
- AP placement and migration
- Wired infrastructure capacity
- Power requirements
- Security
- Troubleshooting and optimization

Accurately planning, optimizing, securing and troubleshooting 802.11n networks, both in the presence of 802.11a/b/g networks and in Greenfield deployments, requires new tools that are up to the job. Because 802.11n might become a wholesale replacement for wired LAN cabling over time, 802.11n networks will probably grow quite large. So tools that can simulate the environment, sample data, and scan on a grand scale are in order.

Wireless lifecycle management

Keeping control of a wireless network, that by design is constantly changing, requires an end-to-end management approach focused on each phase of the wireless lifecycle. The wireless lifecycle involves distinct and interrelated phases: pre-deployment and expansion planning, installation and verification, troubleshooting, and management and optimization.

To navigate through each phase successfully and efficiently, a network manager needs tools that provide features and functions specific to the unique requirements of each phase of this wireless lifecycle. Fluke Networks offers a suite of wireless solutions that deliver complete network visibility to help you successfully manage your network's wireless lifecycle.

For more information about 802.11n, and Fluke Networks wireless solutions, visit **www.flukenetworks.com/wireless**

NETWORKSUPERVISION

Fluke Networks
P.O. Box 777, Everett, WA USA 98206-0777

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

Printed in U.S.A. 9/2008 3365154 H-ENG-N Rev A