802.11n: The End of Ethernet?

Version 2.0, Sept 14, 2009

AUTHOR:

Paul DeBeasi

(pdebeasi@burtongroup.com)

TECHNOLOGY THREAD:

Wireless and Mobility

48970

BURTON GROUP 7090 Union Park Center · Suite 200 · Midvale, Utah · 84047
P 801.566.2880
F 801.566.3611
www.burtongroup.com

Table of Contents

Summary of Findings

Bottom Line: The Institute of Electrical and Electronics Engineers (IEEE) 802.11n standard represents the beginning of the end for wired Ethernet as the dominant local area network (LAN) *access* technology in the enterprise. Over the next few years, improvements in system silicon, radio design, network control, wireless security, power management, and cost will improve 802.11n to the point where it will begin to erode the switched Ethernet market.

Context: Enterprise mobility is driven by the demand for seamless access to information anytime, anywhere, and from any device. Mobility can facilitate business continuity, improve collaboration, simplify teleworking, and increase employee retention. In addition, many younger workers are pushing enterprises to embrace mobility solutions. Employees want the convenience of being untethered.

Takeaways:

- Pervasive mobility
  - The definitive and unalterable competitive advantage that 802.11 has over Ethernet is the ability to provide pervasive mobility.
- Performance
  - Gigabit Ethernet provides greater throughput than 802.11n; however, for many enterprise users, 802.11n throughput is good enough.
  - 802.11n latency and jitter are much greater than that of Gigabit Ethernet; however, 802.11n latency and jitter are well below most latency and jitter budgets.
- Security
  - 802.11n involves greater security threats than wired Ethernet in the areas of eavesdropping, Denial of Service (DoS), and MAC/PHY intrusion.
  - Technology exists to make an 802.11n network almost as secure as Ethernet, but that security comes at a greater cost than wired Ethernet network.
- Management
  - 802.11n requires network managers to provide new tasks such as site planning and spectrum management.
  - Network managers must use new tools to manage WLANs such as spectrum analyzers and WLAN monitoring equipment.
- Staffing
  - 802.11n is an entirely new technology replete with new terminology.
  - Personnel must learn what these terms mean in order to deploy, maintain, and upgrade their network.
- Cost
  - The costs for 802.11 and Ethernet deployments vary significantly and depend on many enterprise-specific factors.
  - Therefore, one should not assume that wireless is more expensive than wired Ethernet (or vice versa).
- Market Impact
  - The growth of the Ethernet switching market will begin to slow due to WLAN substitution.

  - Ethernet-switching vendors that rely on a wireless original equipment manufacturer (OEM) partner for WLAN systems are vulnerable to losing their partners through acquisition.
  - Vendors with unified wireless-wired solutions such as Cisco, HP ProCurve, and Siemens-Enterasys will have an advantage over pure-play wireless vendors such as Aerohive, Aruba, Extricom, Meru, Ruckus Wireless, and Xirrus.

- Recommendations: Enterprises should consider 802.11n an appropriate LAN access substitute for wired Ethernet when:
  - The number of mobile device users is growing
  - The enterprise uses mobile applications
  - Fast Ethernet throughput is good enough
  - The enterprise deploys Voice over Internet Protocol (VoIP)
  - Moves/adds/changes (M/A/C) are made frequently
  - The risk of deliberate DoS attack is low to moderate
  - Ethernet cable installation is difficult

Analysis

This is the second in a series of four *Network and Telecom Strategies* reports on 802.11n. The first, "802.11n: Beyond the Hype," analyzes 802.11n technology. The third report, "802.11n: Enterprise Deployment Considerations," analyzes 802.11n deployment issues. The fourth report, "802.11n: Impact on WLAN Management," analyzes 802.11n impact on WLAN management tools.

This report compares 802.11n to Gigabit Ethernet for use in local area network (LAN) *access* applications (refer to the Reference Architecture technical position, "Local Area Networks," for a description of an access network). More specifically, the report compares both technologies in six dimensions: mobility, performance, security, management, staffing, and cost. Many of the comparisons between 802.11n and Gigabit Ethernet apply equally well when comparing 802.11 versus Ethernet. Therefore, this report will use the following conventions:

- The terms "802.11" and "wireless" will refer to any of the 802.11 technologies (.11b, .11g, .11a, and .11n).
- The terms "Ethernet" and "wired" will refer to any of the 802.3 technologies (10 Mbps Ethernet, Fast Ethernet, Gigabit Ethernet, and 10 Gigabit Ethernet).
- The terms "802.11n" and "Gigabit Ethernet" will only be used when specifically referring to "802.11n" and "Gigabit Ethernet," respectively.

Pervasive Mobility

802.11 enables pervasive mobility and Ethernet doesn't. Pervasive mobility is the ability of workers to remain connected to the LAN regardless of where they are located within the enterprise. One can analyze the differences between 802.11 and Ethernet with respect to performance, security, manageability, cost and impact on staff. However, the definitive and unalterable competitive advantage that 802.11 has over Ethernet is the ability to provide pervasive mobility. Employees want the convenience that being untethered provides.

Growing consumer use of 802.11 will drive the demand for pervasive mobility. 387 million Wireless Fidelity (Wi-Fi) chipsets were sold in 2008, and over 30% of those chipsets were installed in consumer electronics devices. Teen mobile communication usage is soaring, thus driving the next generation of employees to expect pervasive mobility in the workplace. For more information, refer to the *Network and Telecom Strategies* report "Enterprise Mobility at a Crossroads."

Laptops

Laptops and notebook computers outsell desktop PCs. Virtually all laptops now come equipped with 802.11n technology. Laptops enable enterprise user mobility. Along with that mobility comes the expectation of wireless connectivity. Laptop users expect to access e-mail, browse the Internet, and access corporate LANs. As the number of laptops in the enterprise grows, so will the demand for wireless LANs (WLANs).

Voice over WLAN

The number of business calls on mobile cellular phones now exceeds that on wired desktop phones. The growing demand for Voice over WLAN (VoWLAN) in the enterprise will drive the use of 802.11n as the preferred LAN access technology. According to ABI Research, by 2014, 90% of smartphones will support 802.11. As the demand for VoWLAN grows, so does the need for 802.11n.

VoWLAN systems have already found wide acceptance in manufacturing, medical, and retail environments. In virtually any enterprise, some workers will need to have access to their phones regardless of where they are located within the building. Examples include technical support personnel who move from location to location as problems arise, or management personnel who may be away from their offices much of the day as they attend meetings with customers and other employees.

In many cases, these mobile workers already own mobile cellular phones. So long as there is a strong mobile cellular signal within the building, including all locations where these employees may be working at any given time, a mobile cellular phone might be acceptable. But using mobile cellular phones to communicate with other employees within the same building can be costly. With the introduction of both Voice over Internet Protocol (VoIP) and WLANs within the enterprise, it begins to make sense to implement VoWLAN for calls made from within the enterprise. This will eliminate the need for these internally mobile workers to use mobile cellular phones. Refer to the *Network and Telecom Strategies* report "Voice-Enabled WLANs: Is the Network Ready?" for further information on VoWLAN.

Fixed Mobile Convergence

The growing demand for fixed mobile convergence (FMC) in the enterprise will drive the use of 802.11 as the preferred LAN access technology. FMC can be defined very simply as a combination of wired and wireless telecommunications services, but it is much more than that. Beyond just the bundling of separate fixed and mobile service offerings, FMC can exploit synergies between wireless and wireline network technologies to deliver important new capabilities, ultimately creating the appearance of a single system and delivering a consistent user experience independent of the underlying wired or wireless network transport.

For example, mobile cellular phone users can automatically be switched to a WLAN whenever they move within close proximity of a WLAN access point (AP), and then have their calls routed over fixed wireline networks. Internet Protocol private branch exchanges (IP-PBXs) can provide call control and enterprise telephony applications for both wired desktop and wireless phones. And FMC can apply to multimedia (e.g., video and messaging) as well as just voice communications.

Now that FMC products or services that permit seamless roaming between different networks are emerging, enterprises will begin to deploy enterprise FMC solutions further driving the demand for a mobility infrastructure. Refer to the *Network and Telecom Strategies* report "Fixed Mobile Convergence: Aggregation or Aggravation?" for further information on FMC. Also, refer to the *Network and Telecom Strategies* report ”Mobile Unified Communications” for information on the related topic of mobile unified communication (UC).

Performance

Although Gigabit Ethernet is the performance winner over 802.11n in all three categories—throughput, latency, and jitter—the more interesting question is: “Does 802.11n provide *good enough* performance for most enterprises?”

Throughput

Industry testing has shown that enterprise 802.11n APs can achieve 150+ Mbps aggregate throughput under real world test conditions (see Cisco/Intel test and Network World test). However, it is not uncommon for 15 or more users to associate with the same AP. Therefore the average throughput per user will vary as the number of users per AP varies.

Conversely, Gigabit Ethernet switches provide up to 1,000 Mbps of dedicated full duplex throughput per user. Dedicated throughput means that as the number of users per device increases, the throughput per user stays constant. The backplane for some Gigabit Ethernet switches may block traffic above a threshold level, and thus some switches may limit per user throughput. Also, many enterprises still deploy Gigabit Ethernet on their trunk links. Thus, if several Gigabit Ethernet users simultaneously need to use the trunk, users will experience less than 1,000 Mbps on their access link. Therefore, Gigabit Ethernet can also be subjected to bandwidth limitations, although not anywhere near as much as for 802.11n.

Gigabit Ethernet’s dedicated throughput is significantly faster than 802.11n’s shared throughput. For engineering and scientific applications that transfer files that are hundreds of megabytes in size, greater throughput can make a significant difference in user experience. However, does greater throughput really matter to most enterprise users who are not transferring enormously large files?

To answer this question, Burton Group calculated the one-way file transfer time of moderately sized files (e.g., 2, 4, 6, and 8 MB) for wired and wireless LANs. We used two AP capacity plans—10 users per AP and 20 users per AP—and assumed 25 Mbps, 100 Mbps, 150 Mbps, and 1,000 Mbps for 802.11g, Fast Ethernet, 802.11n, and Gigabit Ethernet throughput, respectively. We found that although 802.11n was slower than Gigabit Ethernet, the download time difference was negligible (see Figure 1). Even with 20 users per AP, the file download times ranged from two to eight seconds—still satisfactory for most users.

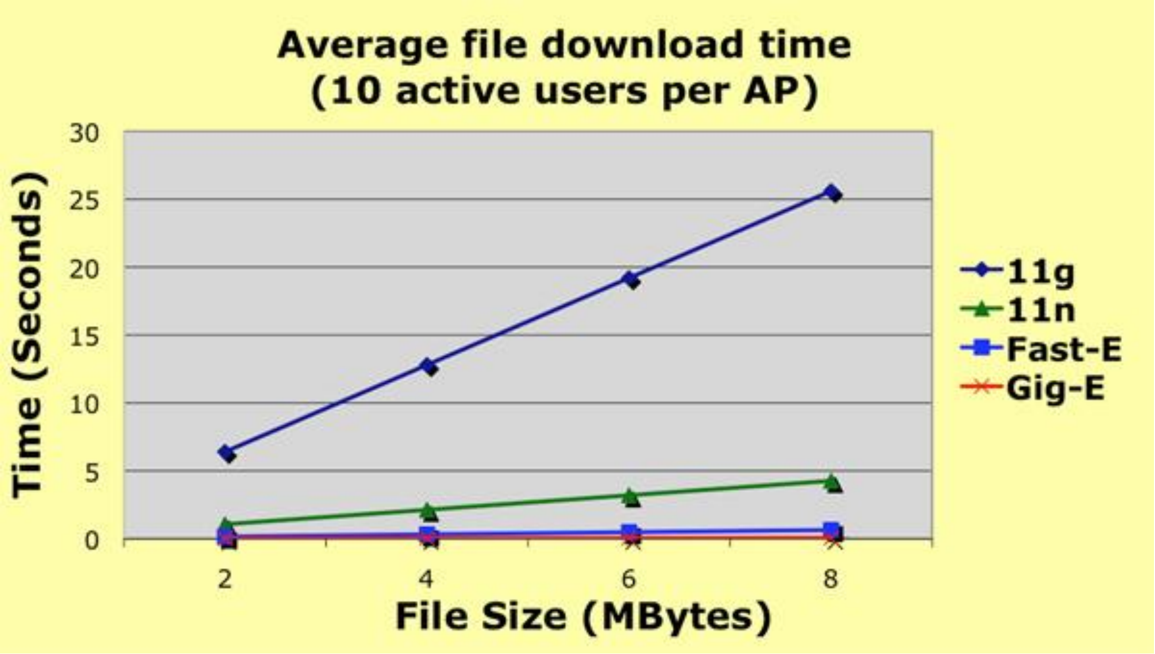

Figure 1: *Average File Download Time: Wired vs. Wireless LANs*

Gigabit Ethernet products have matured to the point that there is little additional throughput that can be squeezed out of them, because most products are already operating close to wire speed. In contrast, 802.11n is early in its lifecycle and will certainly enjoy significant performance improvements in 2010 and 2011 as vendors improve their systems (e.g., with four transmitter chains). In addition, 802.11n prices will fall as volumes increase, thereby enabling higher-density AP deployment. Higher-density deployment will increase aggregate network capacity and improve the average throughput per user (because there will be fewer users per AP). Note that network design and co-channel interference will significantly affect the achievable aggregate network capacity. Refer to the *Network and Telecom Strategies* report "Demystifying Radio Management" for more information.

Latency

Latency and jitter are important in applications such as VoWLAN. Latency is the one-way delay between a sender and a receiver. It can be caused by a number of factors including voice signal sampling time, packet creation delay, network delay, and propagation delay. The International Telecommunication Union (ITU) recommends that total one-way latency not exceed 150 milliseconds (ms). Latency that exceeds this metric will cause awkward pauses and collisions of spoken words. Refer to the *Network and Telecom Strategies* report "Voice-Enabled WLANs: Is the Network Ready?" for more information about VoWLAN.

As shown in Figure 2, 802.11n latency is much greater than that of Gigabit Ethernet. But this difference will have little impact on latency-sensitive applications such as VoIP because the absolute value is well below the VoWLAN latency budget (i.e., 150 milliseconds, equal to 150,000 microseconds), especially for the small packet sizes that are commonly used for VoIP communication. Therefore, 802.11n latency should be more than good enough for most latency-sensitive applications.

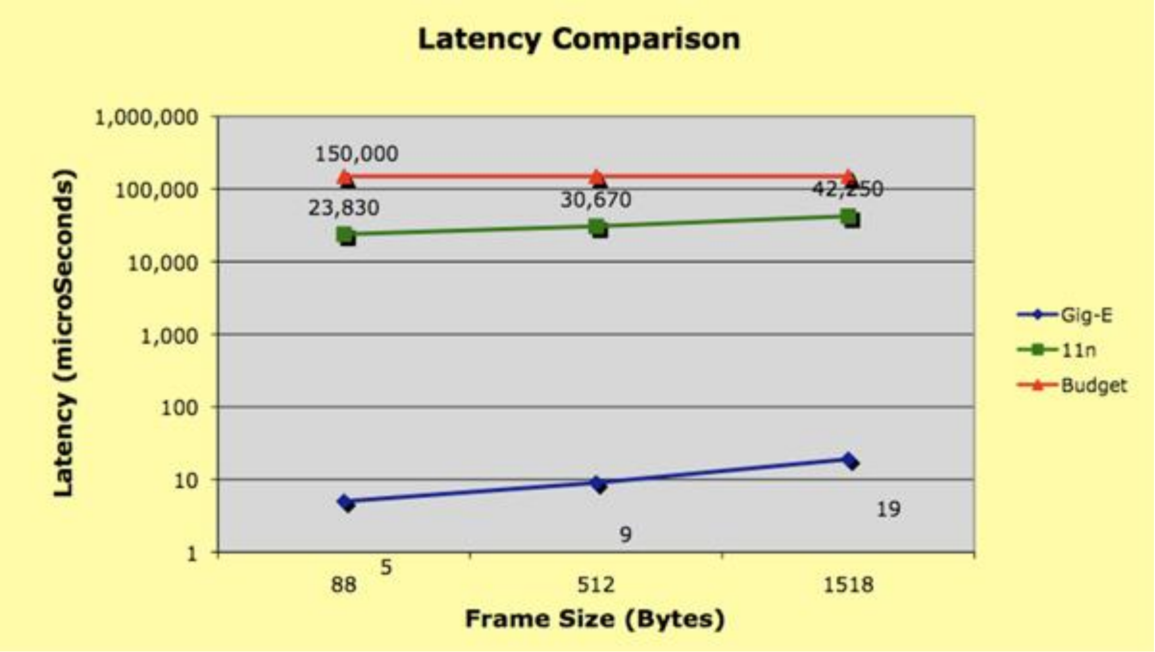

Figure 2: *Latency: Gigabit Ethernet vs. 802.11n (Source Data: Network World tests)*

Jitter

Jitter is the amount of variation in the arrival times of VoIP packets. Jitter buffers that smooth the playback at each phone can accommodate some jitter. However, excessive jitter (beyond that which can be buffered) causes packets to be dropped, which degrades the quality of the voice. The total jitter should not exceed 30 ms.

As shown in Figure 3, 802.11n jitter can be more than an order of magnitude larger than that of Gigabit Ethernet. But who cares? Again, the difference will have little impact on jitter-sensitive applications such as VoWLAN, because the absolute value is well below the VoWLAN jitter budget. Therefore, 802.11n jitter should be good enough for most jitter-sensitive applications. See the "Latency and Jitter" section of this report for more information.

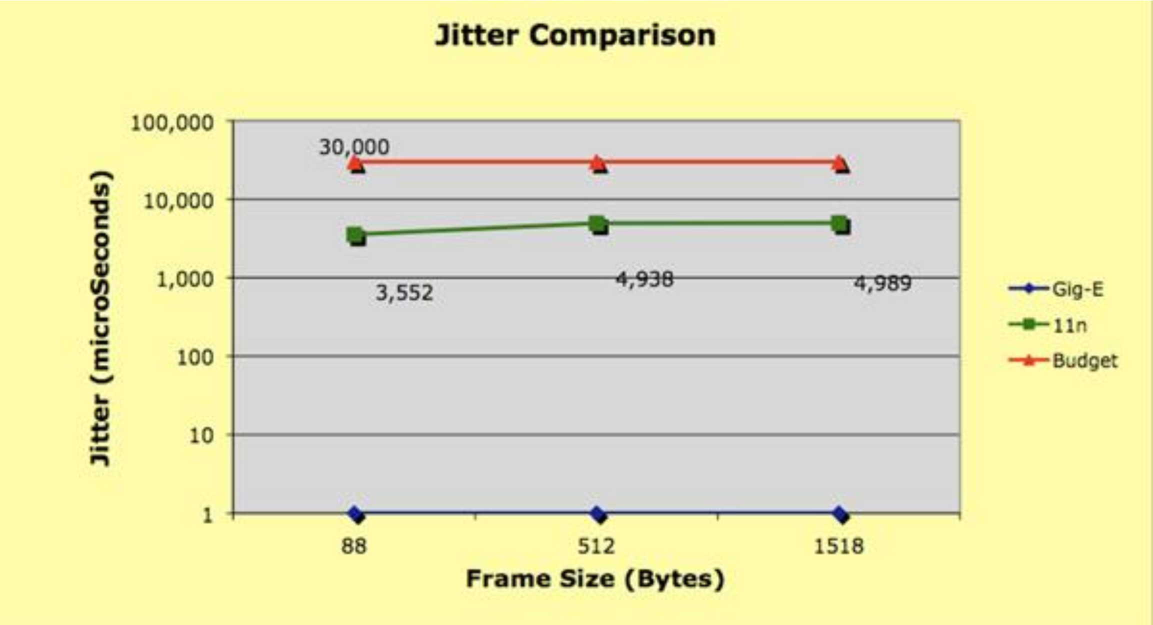

Figure 3: *Jitter: Gigabit Ethernet vs. 802.11n (Source Data: Network World tests)*

Security

Perhaps the biggest concern for many enterprises is securing a WLAN. This concern is reinforced by highly publicized crimes such as the TJX wireless security breach, where the intruders took advantage of the weak Wired Equivalent Privacy (WEP) to steal more than 45 million credit card numbers. The fact that wireless signals are invisible, and impossible to completely control, means that, unlike a wired LAN, a network manager never fully knows where the network is deployed. This situation is analogous to having an enabled RJ-45 port installed in the visitor parking lot. Even worse, the location of the port may move over time without the network manager even knowing that it ever existed.

Network managers have had to deal with network security threats since the first networks were created. These attacks can be directed at the LAN itself—the Physical (PHY) layer, or medium access control (MAC) layer—or at higher layers (see the *Security and Risk Management Strategies* Reference Architecture technical position "Network Intrusion Detection and Response"). Network security threats that occur above the MAC layer, such as firewall port scans and ICMP flood attacks, can occur over WLANs as well as wired LANs. In addition, virtually any LAN attack against Gigabit Ethernet can also be targeted at wireless LANs. Refer to the *Network and Telecom Strategies* Methodologies and Best Practice document "Securing WLANs in the Enterprise" for further information on WLAN security.

Eavesdropping

Wired LAN eavesdropping is difficult for an intruder to achieve while inside—and virtually impossible to achieve while outside—an enterprise facility. An intruder must physically connect eavesdropping equipment to the network somewhere inside the facility. Also, switched Ethernet normally restricts traffic on the access link to packets destined for individual users. Therefore, eavesdropping equipment on an access link has very limited visibility into network traffic.

Conversely, a WLAN uses radio waves that can propagate in multiple dimensions (above, below, beside) throughout the enterprise facility. Therefore, an intruder can be on the floor above (or even outside the facility) and still eavesdrop on the network. The intruder has no physical connection, so it is difficult to discover. Lastly, a WLAN is a shared, broadcast technology, so eavesdropping equipment, once connected, can easily see traffic destined for many users.

Network administrators must deploy Layer 2 encryption (e.g., Wi-Fi Protected Access 2 [WPA2]) or Layer 3 encryption (e.g., IP security [IPsec] or Secure Sockets Layer [SSL]) in order to maintain privacy on WLANs. Unfortunately, WLANs have three types of Layer 2 encryption schemes—WEP, WPA, and WPA2. Many enterprise managers maintain all three schemes because upgrading to WPA2 can be time-consuming and costly, sometimes requiring new hardware. This results in additional network complexity because many enterprises attempt to isolate the three security schemes by using different service set identifiers (SSIDs) and virtual LANs (VLANs).

Denial of Service

Any event that prevents authorized users from performing appropriate functions may be considered a denial-of-service (DoS) event. DoS events can occur within any component of the information technology (IT) infrastructure or even outside of IT. UDP floods (directed at the enterprise Internet connection) and radio frequency (RF) jamming (directed at the enterprise WLAN) are types of DoS attacks. In the context of this report, a DoS attack is one that prevents operation of the LAN (rather than a DoS attack aimed at upper layers). Refer to the *Security and Risk Management Strategies* overview, "Beyond Denial of Service: Is Availability a Security Issue?" for more information about DoS.

When Ethernet LANs were built using hubs instead of switches, a faulty station or broadcast storm could bring down an entire LAN. Even switched LANs can be susceptible to broadcast storms if the spanning tree protocol is not configured correctly. A DoS attack against a wired LAN must be initiated from within the facility because the intruder must physically connect to the network. However, these DoS attacks against wired LANs are rare and preventable.

In contrast, WLAN DoS attacks are easy to launch. In fact, the simple act of reheating a meal in the microwave oven can unintentionally prevent the WLAN from operating. The increased range of 802.11n compared with 802.11g may also cause unintentional interference from neighboring WLANs. Deployment of 802.11n in the 5 GHz band can help to reduce the likelihood of unintentional DoS interference because far fewer products on the market operate in this band. However, intruders can launch DoS attacks while outside the facility by using a directional antenna to aim interference at the target WLAN. Unlike the older broadcast storms that, once launched, were forwarded throughout the wired network, wireless DoS attacks are localized near the source of the interference.

Network Intrusion

Network intrusion refers to a situation or sequence of events with consequences induced by an unauthorized attacker. Firewall port scans and malicious AP associations are examples of network intrusion attacks. Network intrusion results in unauthorized network traffic that may be targeted to exploit vulnerabilities on systems or associated with malicious code (e.g., worms and Trojan horse programs), or it may be traffic that violates the organization's acceptable use policy.

A network intrusion detection and response system (NIDRS) can help to protect against attacks. An NIDRS (sometimes simply called intrusion detection system [IDS]) device is a mechanism that detects unauthorized network traffic and responds to identified security events. Most NIDRS devices rely on some type of signature detection mechanism, but new techniques are being used to improve their effectiveness and discrimination. An NIDRS operates *above* the MAC layer so it can be used for both wireless and wired LANs (refer to the *Security and Risk Management Strategies* report "Network Intrusion Detection and Response: More Than Just Speed Bumps on the Network?").

A wireless intrusion prevention system (WIPS)—also referred to as a wireless intrusion detection system—can monitor for rogue APs and unauthorized devices, maintain policy adherence in the air and on the APs, and look for anomalous or suspicious behavior on the WLAN. An overlay WIPS solution relies on dedicated, distributed, hardware sensors that look like APs. The sensors continuously monitor multiband channels and report anomalies back to a centralized management console. Overlay WIPS solutions can be costly to implement and maintain. Alternatively, enterprises can use an integrated WIPS solution from WLAN system vendors that integrates a sensor function into an AP. However, most of these integrated solutions do not provide continuous monitoring and thus may not catch some intrusion attempts.

If companies can meet their WLAN security needs with wired-side management tools or integrated solutions from WLAN system vendors, a WIPS may not be worth the investment. However, without an overlay WIPS solution, it is impossible to gain an independent, security-aware view of the airspace. Refer to the *Security and Risk Management Strategies* report "Wireless LAN Intrusion Detection Systems: Something's in the 'Air'" for more information on WIPS.

Security Comparison

Information protection is something you do, not something you buy. It requires well-defined processes and effective use of technologies—all based on a sound understanding of the business that the organization performs and how it performs that business. (See the *Security and Risk Management Strategies* report "A Systematic, Comprehensive Approach to Information Security.") However, some significant differences between wired and wireless LANs introduce new attack vectors and the need for new technologies to mitigate risk.

802.11n involves greater security threats than wired Ethernet in the areas of eavesdropping, DoS, and MAC/PHY intrusion. These threats require an IT manager to buy products, develop procedures, educate personnel, and continuously monitor the network. Technology exists to mitigate eavesdropping and MAC/PHY intrusion threats. DoS threats are much easier to detect than to mitigate. However, for most enterprises, the risk of deliberate DoS attack is low. (For a Burton Group definition of "risk," see the *Security and Risk Management Strategies* report "Risk

Management: Concepts and Frameworks.") The result is that an 802.11n network can be made almost as secure as Ethernet, but that security will come at a greater cost than for a wired Ethernet network.

Security threats	Ethernet	802.11
High potential for eavesdropping		√
High potential for DoS attack		√
Intrusion: Vulnerable to network layer (and above) attacks	√	√
Intrusion: Vulnerable to MAC/PHY layer attacks		√

Table 1: *Security Threats: Ethernet vs. 802.11*

Management

The modern IT environment includes a constantly changing, heterogeneous collection of servers, networks, and clients. Network management in such a dynamic multivendor environment is difficult and is complicated by the lack of widespread adoption of advanced standards. Refer to the *Network and Telecom Strategies* overview, "Architectural Overview of Network Management," for additional information on network management.

Moves/Adds/Changes

Moves/adds/changes (M/A/C) can be costly. If an enterprise is reconfiguring office space, then data and voice cabling will likely need to move, cabling will need to be re-terminated, and patch panels may need to be reconfigured (see the Reference Architecture technical position, "Building Wiring," for further information on building wiring). In addition, IT staff will need to change the telephony directory to reflect the new locations of the employees' phones, possibly requiring new numbers. With 802.11n, employees who use laptops and VoIP phones can move throughout the enterprise without incurring any of these M/A/C costs.

Network Management Tasks

In addition to the tasks that confront a network manager for a wired network, the inherent differences between 802.11 and Ethernet present a new set of challenges (see Table 2). For example, 802.11 requires network managers to perform site planning, spectrum management, and to monitor for poor coverage and RF interference. These additional tasks require wireless expertise and the use of new tools.

Required network management tasks	**Ethernet**	**802.11**
Incident reporting and resolution	√	√
AP and switch provisioning	√	√
Account management	√	√
Data gathering, processing, and reporting	√	√
Network modeling and capacity management	√	√
Intermediate distribution frame cable management	√	
Power and cooling management	√	√
Desktop M/A/C	√	
RF interference and poor coverage resolution		√

Table 2: *Required Network Management Tasks: Ethernet vs. 802.11*

Tools

Given the increasing deployment of WLANs, network managers must find a way to address these new tasks. Though the management software provided by the WLAN equipment vendors can answer some of these needs for their own homogeneous product environments, other critical performance and security information may be missing. This problem is exacerbated when WLAN products from multiple vendors exist in the same network.

To meet the demand for tools that help network managers address these issues, a number of companies have developed various WLAN management tools. (Refer to the *Network and Telecom Strategies* report "802.11n: Impact on WLAN Management" for additional information.) These tools enable the management and monitoring of heterogeneous WLANs, and provide the information necessary for performance, security, and configuration management and capacity planning. Network managers must use new tools that are purpose-built for the WLAN, such as:

- **Spectrum analyzer tools:** Monitor the RF layer of the WLAN.
- **WLAN monitoring tools:** Collect performance information and provide protocol analysis.
- **WLAN management tools:** Assist in managing device configuration and network security.

As vendors integrate wireless functionality with wired switches and routers, network management systems will evolve to simultaneously support both Ethernet and 802.11, thus simplifying the transition to a WLAN. However, for the foreseeable future, an 802.11 LAN will require the purchase of new tools, an investment in employee training, and the acceptance of greater network management complexity compared with any Ethernet LAN.

Staffing

Enterprise IT departments have been deploying and maintaining Ethernet networks for over twenty years. Assuming a conversion is warranted, the transition from a Fast Ethernet to a Gigabit Ethernet network does not require a significantly new learning curve for IT personnel.

This is not true for the transition from Ethernet to 802.11. The transition from a wired to a wireless network will impose a significant burden on IT personnel.

Compared to Ethernet, 802.11 is an entirely new technology with new terminology, such as spatial multiplexing and dBm. Personnel must learn what these terms mean in order to deploy, maintain, and upgrade the network. IT personnel must create their own design and management best practices because no such standard practices exist for WLANs. Unlike Ethernet, the 802.11 transmission medium is radio waves and the RF signal propagation can therefore be affected by factors such as weather, facility construction, and electromagnetic interference. Although wireless products are becoming more adaptive to RF signal interference, failures and poor performance can still occur. Therefore, IT personnel must understand the propagation characteristics of their facilities and must think about wireless interference in multiple physical dimensions.

Compared to Ethernet, 802.11 introduces new security threats and new network management tools. IT personnel must learn about the new threats, and how to use the new tools. Unlike Gigabit Ethernet, wireless bandwidth is a more limited resource and IT personnel must ensure that the bandwidth is allocated to ensure adequate performance. It is very likely that as the WLAN becomes more pervasive, new mobile applications will emerge to take advantage of the wireless network. IT personnel must therefore be prepared to deploy applications such as location services, VoWLAN, and FMC.

Table 3 summarizes the impact that 802.11 has on IT staff when transitioning to wireless, compared against simply maintaining a wired Ethernet LAN. However, the low cost of consumer Wi-Fi equipment enables enterprises to easily purchase it for training purposes. In addition, the pervasive deployment of Wi-Fi in the home enables staff to gain experience with wireless technology. Lastly, organizations such as the Certified Wireless Network Professional (CWNP) provide vendor-neutral training and certification for IT professionals.

Staff impact	Ethernet	802.11
Must learn a different communication technology		√
Must create design and management best practices		√
Must understand facility propagation characteristics		√
Must deploy new network management tools		√
Must defend against new security threats		√
Must manage new mobility applications (e.g., location)		√

Table 3: *Staff Impact: Ethernet vs. 802.11*

Cost

The costs for 802.11 and Ethernet deployments vary significantly and depend on many enterprise-specific factors. Therefore, one should not assume that wireless is more expensive than wired Ethernet (or vice versa). Enterprises should consider all of the factors listed below and perform their own detailed cost comparisons.

Many factors affect wired Ethernet costs including:

- **Cable installation costs:** Facility age, design, and construction may significantly increase costs for new cabling. Older Ethernet cabling (such as Category 3) already in place will need to be upgraded to new cabling (such as Augmented Category 6) in order to support Gigabit Ethernet and beyond. This upgrade can be costly.
- **Frequency of M/A/C:** If an enterprise frequently reconfigures office space, data and voice cabling will likely need to be moved and re-terminated, and patch panels may need to be reconfigured. In addition, IT staff will need to change the telephony directory to reflect the new locations of employee phones, possibly requiring new numbers.
- **Equipment costs:** Wired Ethernet equipment will vary in cost depending on whether Fast Ethernet or Gigabit Ethernet is deployed to the desktop. Aggregation switches that aggregate traffic from the wiring closet must operate at higher speeds than wiring closet switches. Gigabit Ethernet LAN access deployment may require deployment of switches with 10 Gigabit Ethernet trunks, further increasing the total cost for Gigabit Ethernet deployment.

Many factors affect 802.11 costs, such as:

- **Design and installation costs:** Wired LANs typically adhere to structured wiring standards such as Telecommunications Industry Association/Electronic Industries Alliance (TIA/EIA)-568-B. These standards define structured wiring design, cable types, connector types, cable distances, testing methodologies, and so on. Such standards lower wired LAN costs by enabling use of standard components, creation of best practice methodologies, and establishment of standard design and installation training programs. Unfortunately, such design and installation standards do not exist for WLANs. The result is that design methodologies (such as for site survey and RF deployment design) can vary from integrator to integrator, resulting in additional costs.
- **Equipment costs:** Wireless network equipment includes APs and controllers. Equipment costs vary from vendor to vendor and depend on whether 802.11g or 802.11n APs are deployed. The number of APs depends on whether they are pervasively deployed (e.g., enterprise-wide VoWLAN versus conference room wireless) and also on the network capacity design (e.g., 10 users per AP versus 20 users per AP). Also, the number of required controllers will depend on controller architecture (e.g., centralized versus distributed). See the *Network and Telecom Strategies* reports, "Wireless LAN Systems: Ready for the Future?" and "Demystifying Radio Management" for more information on WLAN architectures.
- **Network management costs:** Wireless often requires that IT personnel learn new technology, tools, and tricks, resulting in increased operational costs. IT personnel may need training on wireless problem diagnosis, such as the use of spectrum analyzers to diagnose interference issues. WLAN system element management software (e.g., from Aruba Networks) may need to be integrated with Ethernet element management software (e.g., from Cisco Systems).
- **Security costs:** Wireless networks require that additional security tools, such as a WIPS, be purchased, deployed, and maintained. In addition, many enterprises will deploy 802.1X security systems to complement WPA2 security. 802.1X requires enterprises to integrate 802.X client software, an Extensible Authentication Type (e.g.,

EAP-FAST [Flexible Authentication via Secure Tunneling]), WLAN system management software, and an Authentication Server. This can be time-consuming and costly.

Market Impact

Enterprises are increasingly deploying 802.11 for network access. The emergence of 802.11n, and related 802.11 standards, will accelerate that trend so that 802.11n will eventually become the dominant enterprise LAN technology. As this happens, the growth of the Ethernet switching market will begin to slow due to WLAN substitution.

Ethernet switching vendors that rely on an original equipment manufacturer (OEM) partner for WLAN systems—such as Brocade Networks (Meru Networks OEM), and Alcatel-Lucent (Aruba Networks OEM)—are vulnerable to losing their partner through acquisition. For example, when Cisco Systems acquired Airespace in 2005, both Nortel Networks and Alcatel had to scramble to find a new OEM partner. The same thing could happen to other wired Ethernet switching vendors. Although this vulnerability existed before, 802.11n will make the stakes much higher because it will be pervasively deployed throughout the enterprise. Ethernet switch vendors must ensure continuous wireless product delivery or else risk losing market share in one of the biggest growth markets in the networking industry.

802.11n is also a more complex technology than earlier versions of 802.11. Will customer support for a wired Ethernet switch vendor be able to handle a tricky wireless interference problem? Probably not. Ethernet switching vendors are at a competitive disadvantage because they rely on their OEMs' wireless expertise rather than their own core competencies.

Finally, as 802.11 controllers become more embedded in traditional Ethernet switching products, it will become more difficult to purchase separate wired and wireless products. This means that Ethernet switching products will evolve to become converged wired-wireless products. Ethernet switching vendors will need to integrate wireless into their entire solution portfolios, including their network management systems, consulting services, and customer support. Vendors with unified wireless-wired solutions such as Cisco, HP ProCurve, and Siemens-Enterasys will have an advantage over pure-play wireless vendors such as Aerohive, Aruba, Extricom, Meru, Ruckus Wireless, and Xirrus.

Recommendations: When to Use 802.11n

The premise of this report is that the arrival of the Institute of Electrical and Electronics Engineers (IEEE) 802.11n standard marks the beginning of the end for wired Ethernet as the dominant LAN access technology in the enterprise. Over the next few years, refinements in system silicon, radio design, network control, wireless security, and power management will significantly improve 802.11n to the point where it will begin to erode the switched Ethernet market. Therefore, the recommendations in this section provide guidelines that answer the question: "When is 802.11n an appropriate LAN access substitute for wired Ethernet?"

When the Number of Mobile Device Users Is Growing

As the number of mobile devices (e.g., laptops, netbooks, and smartphones) grows, so does the demand for wireless access to e-mail, the Internet, and LAN resources, such as file and printer servers. Many laptop users deploy consumer Wi-Fi APs in their homes and expect the same wireless access in their workplaces. 802.11n can satisfy the demand for wireless access.

When the Enterprise Uses Mobile Applications

Many enterprises rely on mobile applications to support daily operations. For example, a hospital may need to maintain location awareness for mobile assets such as ventilators, or a retailer may need to track the location of inventory. 802.11n provides a foundation for deploying mobile applications.

When Fast Ethernet Throughput Is Good Enough

802.11n APs now offer throughput speeds to well over 150 Mbps. In addition, 802.11n prices will fall as volumes increase, thereby enabling higher-density AP deployment. Higher-density deployment will increase aggregate network capacity and improve the average throughput per user. Even though 802.11n uses shared media, its throughput improvements suggest that 802.11n APs will eventually approach the performance levels provided by Fast Ethernet by using multiple radio APs and improved radio design.

When the Enterprise Deploys VoIP

802.11n can support latency- and jitter-sensitive applications, such as VoIP. Rather than deploying VoIP over wired Ethernet, the enterprise should consider 802.11n. Use of 802.11n will enable mobile workers to make phone calls while moving through the enterprise. In addition, 802.11n will provide a foundation upon which an enterprise can deploy FMC and unified communication solutions in the future.

When Moves/Adds/Changes Are Made Frequently

Moves/adds/changes (M/A/C) can be costly. If an enterprise frequently reconfigures office space, data and voice cabling will likely need to be moved, cabling will need to be re-terminated, and patch panels may need to be reconfigured. In addition, IT staff will need to change the telephony directory to reflect the new locations of the employees' phones, possibly requiring new numbers. With 802.11n, employees who use laptops and VoIP phones can move throughout the enterprise without incurring any of these M/A/C costs.

When the Risk of Deliberate DoS Attack Is Low to Moderate

WLAN DoS attacks are easy to launch. Intruders can launch DoS attacks while outside the facility by using a directional antenna to aim interference at the target WLAN. DoS attacks and unintentional interference can be detected by WIPSs. For most enterprises, the risk of *deliberate* DoS is low and is outweighed by the benefit of wireless mobility.

When Ethernet Cable Installation Is Difficult

For some enterprise facilities, cable installation may be very challenging. An old facility with poor access to pathways for horizontal cable runs may make cable installation cost prohibitive. In some cases, 802.11n APs may be the only viable solution to provide network access.

The Details

This section provides an overview of the differences between wired and wireless local area networks (LANs) as well as detailed test information.

Why Wireless Is Different

Wired and wireless networks are fundamentally different. A Gigabit Ethernet wired connection between a client and the network operates at a data rate of 1,000 Mbps. The connection speed is fixed, as is the physical location of the connection. Ethernet switches provide a dedicated, full duplex connection that requires physical access to the switch port and, as a result, can offer predictable quality of service (QoS) because there is no contention for bandwidth with other devices.

A wireless connection, on the other hand, varies in data rate speed from 1 to 300 Mbps (with effective throughput approximately one half the data rate), based on the wireless technology chosen, signal strength, and signal quality. Unlike a wired connection, the radio frequency (RF) link is a shared, half-duplex connection. The three-dimensional (3D) nature of radio allows the signal to pass through walls and ceilings, rendering it susceptible to interception and interference. Other devices, such as microwave ovens or cordless telephones, may share the same unlicensed frequency band and cause interference. In multi-tenant office buildings, wireless LANs (WLANs) in adjacent offices can also be a major source of contention, interference, and potential security problems. Another challenge associated with RF networks is that many network managers are unfamiliar with the characteristics of radio networks and therefore resort to guesswork to adjust and configure the access points (APs).

When a wired LAN is deployed it usually behaves predictably. The switches, routers, and communication links that comprise a wired LAN have well-defined operating characteristics. If a wired LAN is deployed according to well-established best practices then that network should operate as expected. A WLAN, however, is dynamic and unpredictable because radio waves (the physical LAN medium) are carried in the air, as opposed to within shielded copper or fiber cabling. Radio waves can bounce off of some objects and penetrate others. Radio waves are also extremely susceptible to problems such as noise, co-channel interference, and multipath reflections.

Testing: General Information

The following sections provide the test topology, methodology, parameters, and results that were used to create the "Analysis" section of this report. The following terminology and reference information applies throughout this report:

- Request for Comments (RFC) 2285, "Benchmarking Terminology for LAN Switching Devices," includes definitions of common benchmarking terms. In particular, the RFC defines ILOAD (intended load) and SUT (system under test). The tests below use RFC 2285 terminology.

- RFC 2544, “Benchmarking Methodology for Network Interconnect Devices,” describes tests that may be used to characterize network device performance.
- RFC 2889, “Benchmarking Methodology for LAN Switching Devices” extends the methodology for benchmarking network interconnecting devices define in RFC 2544.
- RFC 3550, “RTP: A Transport Protocol for Real-Time Applications,” is used in the Gigabit Ethernet jitter tests.

Throughput

This section provides throughput information for 802.11n and Gigabit Ethernet.

802.11n

As of July 2009, most enterprise WLAN vendors have shipped second-generation 802.11n equipment, and many have published performance test results. Two tests of note are the Cisco-Intel test and the Network World test.

Cisco and Intel tested over-the-air (OTA) performance in a real-world environment. The test environment was modeled after a typical enterprise office and included access points, client devices, desks, cubes, file cabinets, and load-bearing walls. The Cisco-Intel test demonstrated that an 802.11n AP could achieve an average throughput of 182 Mbps with peak throughput of 195 Mbps in an office environment. Refer to “Cisco and Intel: Collaborative 802.11n Leadership and Testing” for further details.

Network World tested four vendors: Bluesocket, Siemens, Aerohive and Motorola. In contrast to the Cisco-Intel over-the-air office test, Network World placed each AP in its own shielded radio frequency (RF) chamber and used the VeriWave WT-90 system to test each AP under various traffic conditions. The Network World test demonstrated that enterprise APs could achieve well over 150 Mbps aggregate throughput. In some cases, aggregate throughput per AP reached 259 Mbps. Refer to Network World’s test results and test methodology for additional information.

Both tests provide justification for the claim made in the “Throughput” section of this report that enterprise 802.11n APs can achieve an aggregate throughput of at least 150 Mbps.

Gigabit Ethernet

The “Analysis” section of this report makes the claim that enterprise-class Gigabit Ethernet switches could provide up to 1,000 Mbps of dedicated full-duplex throughput per user. To justify this claim, Burton Group asked VeriWave to test an enterprise-class Gigabit Ethernet switch. Figure 4 shows the test topology for the Gigabit Ethernet test. According to VeriWave, traffic is transmitted in the direction of the arrows. The test client port identifiers and Internet

Protocol (IP) addresses are indicated in the boxes. The Gigabit Ethernet vendor name and model number were not disclosed.

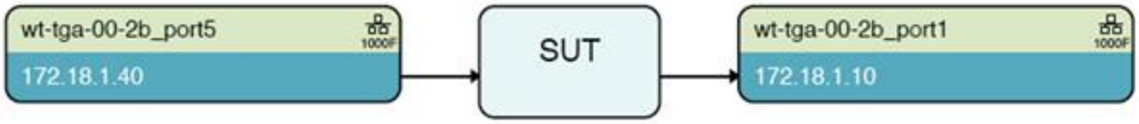

Figure 4: *Gigabit Ethernet: Throughput Test Topology (Source: VeriWave)*

The following text is taken from the VeriWave test report and describes the testing methodology:

> The test is performed by associating test clients with the SUT ports, performing any desired learning transmissions, and then generating test traffic between the test clients. The test then calculates throughput according to the procedure specified in RFC 2544. Proprietary signatures and tags are inserted into the test traffic to ensure accurate measurement results.
>
> A binary search algorithm is used to obtain the throughput, by finding the ILOAD resulting in the highest forwarding rate for which the packet loss ratio is less than the acceptable threshold. The Search Maximum and Search Minimum parameters may be used to constrain the search algorithm. The Starting Point is the starting value of the offered load and its value must be greater or equal to the Search Minimum and less than or equal to the Search Maximum. By default, the search algorithm will start at 50% of the theoretical throughput calculated for the test topology.
>
> The test is repeated for each frame size, and also if the number of trials is greater than 1. The results are recorded separately for each combination of frame size and trial number.

Table 4 describes the test configuration parameters.

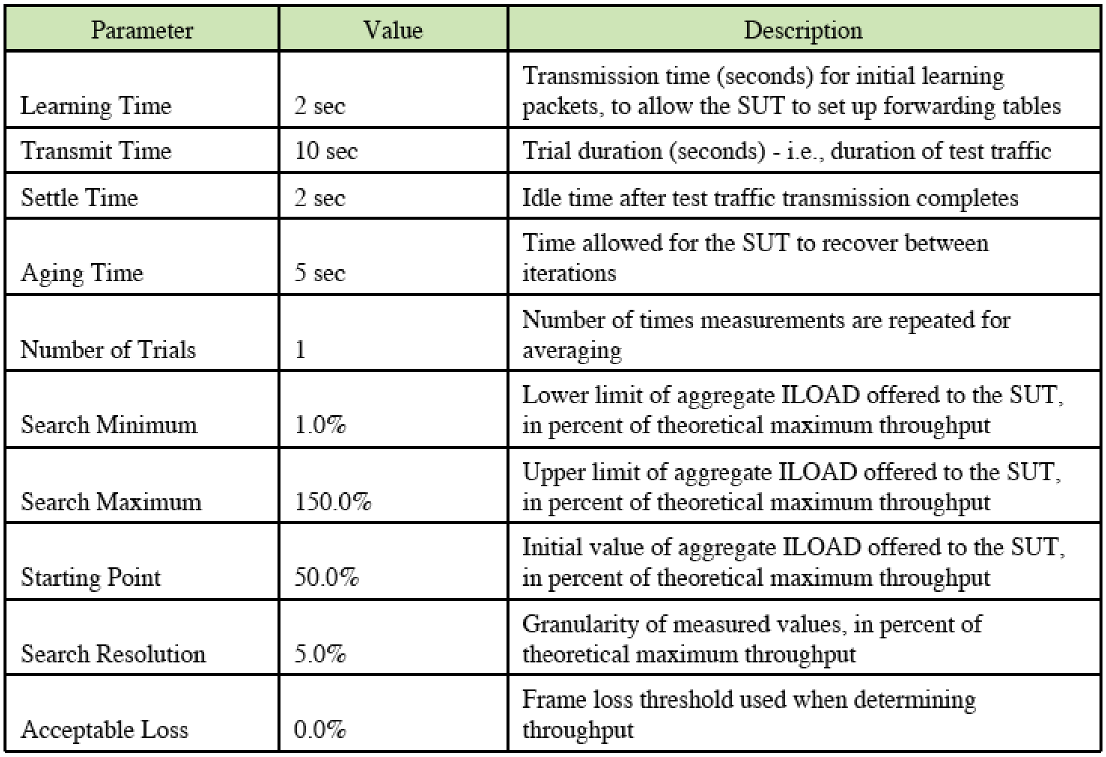

Parameter	Value	Description
Learning Time	2 sec	Transmission time (seconds) for initial learning packets, to allow the SUT to set up forwarding tables
Transmit Time	10 sec	Trial duration (seconds) - i.e., duration of test traffic
Settle Time	2 sec	Idle time after test traffic transmission completes
Aging Time	5 sec	Time allowed for the SUT to recover between iterations
Number of Trials	1	Number of times measurements are repeated for averaging
Search Minimum	1.0%	Lower limit of aggregate ILOAD offered to the SUT, in percent of theoretical maximum throughput
Search Maximum	150.0%	Upper limit of aggregate ILOAD offered to the SUT, in percent of theoretical maximum throughput
Starting Point	50.0%	Initial value of aggregate ILOAD offered to the SUT, in percent of theoretical maximum throughput
Search Resolution	5.0%	Granularity of measured values, in percent of theoretical maximum throughput
Acceptable Loss	0.0%	Frame loss threshold used when determining throughput

Table 4: *Gigabit Ethernet: Throughput Test Conditions (Source: VeriWave)*

The following text is taken from the VeriWave test report and describes the binary search options (see Table 5):

> The maximum, minimum, starting point and search resolution of aggregate ILOAD values are calculated in percent of the theoretical maximum frame rate for the particular frame size

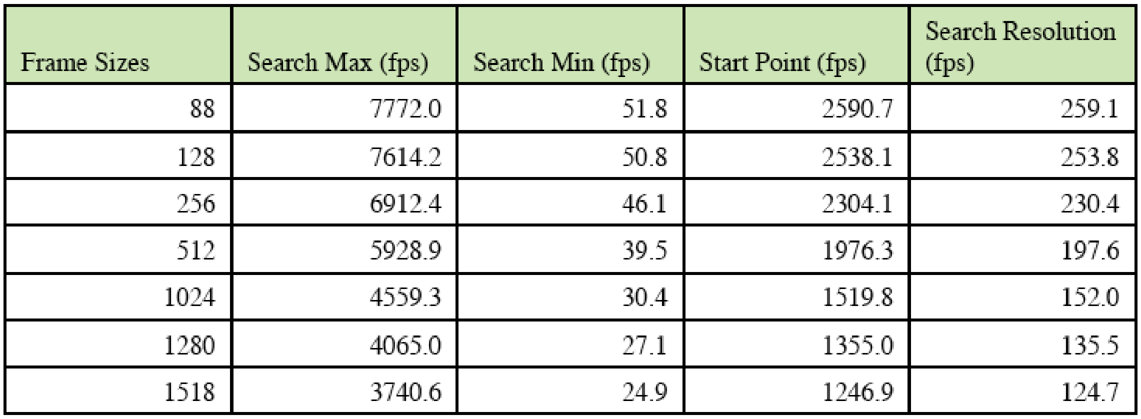

Frame Sizes	Search Max (fps)	Search Min (fps)	Start Point (fps)	Search Resolution (fps)
88	7772.0	51.8	2590.7	259.1
128	7614.2	50.8	2538.1	253.8
256	6912.4	46.1	2304.1	230.4
512	5928.9	39.5	1976.3	197.6
1024	4559.3	30.4	1519.8	152.0
1280	4065.0	27.1	1355.0	135.5
1518	3740.6	24.9	1246.9	124.7

Table 5: *Gigabit Ethernet: Throughput Binary Search Options (Source: VeriWave)*

The detailed test results are provided in Table 6.

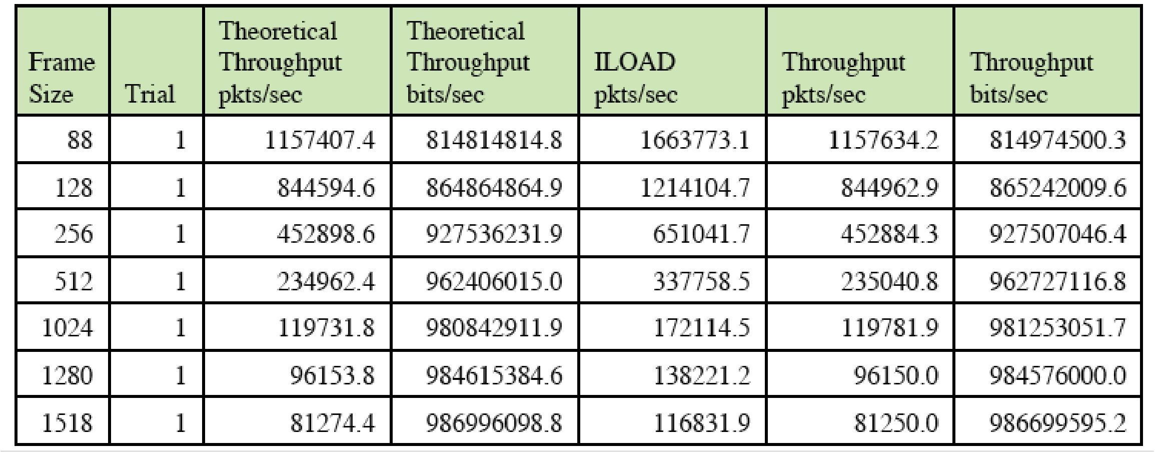

Frame Size	Trial	Theoretical Throughput pkts/sec	Theoretical Throughput bits/sec	ILOAD pkts/sec	Throughput pkts/sec	Throughput bits/sec
88	1	1157407.4	814814814.8	1663773.1	1157634.2	814974500.3
128	1	844594.6	864864864.9	1214104.7	844962.9	865242009.6
256	1	452898.6	927536231.9	651041.7	452884.3	927507046.4
512	1	234962.4	962406015.0	337758.5	235040.8	962727116.8
1024	1	119731.8	980842911.9	172114.5	119781.9	981253051.7
1280	1	96153.8	984615384.6	138221.2	96150.0	984576000.0
1518	1	81274.4	986996098.8	116831.9	81250.0	986699595.2

Table 6: *Gigabit Ethernet: Throughput Results (Source: VeriWave)*

Latency and Jitter

The latency test measures the delay incurred by frames passing through the SUT. The test also measures the amount of jitter, which is the variation in latency over many frames. Latency and jitter are key performance metrics that determine how well the SUT can handle traffic that is sensitive to the delay between source and destination, such as voice or real-time video. This test measures latency and jitter according to RFC 2544 and RFC 3550, respectively.

802.11n

As described in the "802.11n Throughput" section of this report, Network World tested four vendors using the VeriWave WT-90 system (refer to the test results and test methodology for additional information). After reviewing the latency and jitter results of the four vendors referenced in the "Analysis" section of this report, Burton Group chose to use the latency and jitter results from Motorola. Of those four vendors, Motorola has the largest enterprise WLAN market share and is therefore most representative of the 802.11n products in the enterprise WLAN market. Table 7 shows the selected test results used in this report.

Frame Size (bytes)	Latency (microseconds)	Jitter (microseconds)
88	23,830	3,552
512	30,670	4,938
1,518	42,250	4,989

Table 7: *802.11n: Latency and Jitter Test Results (Source: Network World)*

Gigabit Ethernet

The "Latency" and "Jitter" sections of this report make the claim that enterprise-class Gigabit Ethernet switches could provide very low jitter and throughput. To justify this claim, Burton

Group asked VeriWave to test an enterprise-class Gigabit Ethernet switch. Figure 5 shows the test topology for the Gigabit Ethernet latency and jitter test. According to VeriWave, traffic is transmitted in the direction of the arrows. The test client port identifiers and IP addresses are indicated in the boxes. The Gigabit Ethernet vendor name and model number were not disclosed.

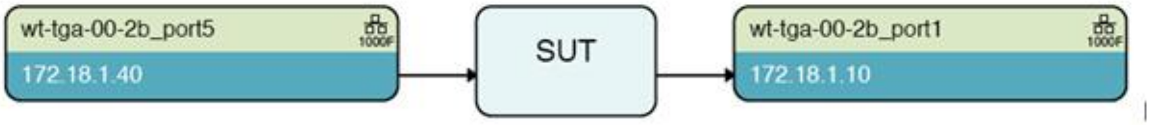

Figure 5: *Gigabit Ethernet: Latency and Jitter Test Topology (Source: VeriWave)*

The following text is taken from the VeriWave test report and describes the testing methodology:

> The test is performed by associating test clients with the SUT ports, performing any desired learning transmissions, and then generating test traffic between the test clients. Proprietary timestamps inserted in each test traffic frame are then used to calculate the minimum, maximum and average latency as per RFC 2544, as well as the smoothed inter-arrival jitter according to RFC 3550. The results are recorded separately for each combination of test conditions, as well as for each trial if multiple trials are run.
>
> Different intended loads (ILOADs) and frame sizes can be set up, to understand how latency varies with different types of traffic in a real environment. Each combination of ILOAD and frame size is tested separately. Test traffic may be configured to flow either from Ethernet to wireless, from wireless to Ethernet, or from wireless to wireless. If multiple APs are involved in the test, the ILOAD is divided evenly across the APs; if multiple clients are associated with an AP, the ILOAD for that AP is divided evenly between the clients.
>
> Latency measurements are made accurately even in the presence of frame loss. However, the ILOAD should be set such that no frame loss occurs; otherwise, buffer occupancy delays can obscure actual SUT datapath delays. The throughput test may be used to determine this traffic level.

Tables 8 and 9 show the test conditions and the test configuration, respectively, for this test.

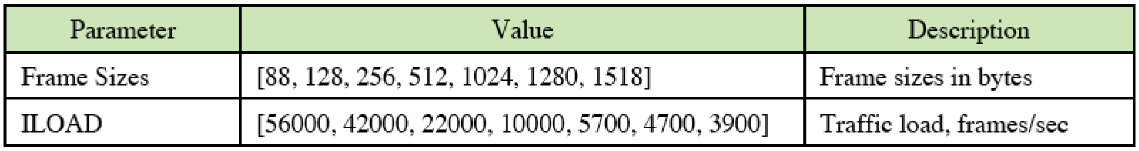

Parameter	Value	Description
Frame Sizes	[88, 128, 256, 512, 1024, 1280, 1518]	Frame sizes in bytes
ILOAD	[56000, 42000, 22000, 10000, 5700, 4700, 3900]	Traffic load, frames/sec

Table 8: *Gigabit Ethernet: Latency and Jitter Test Conditions (Source: VeriWave)*

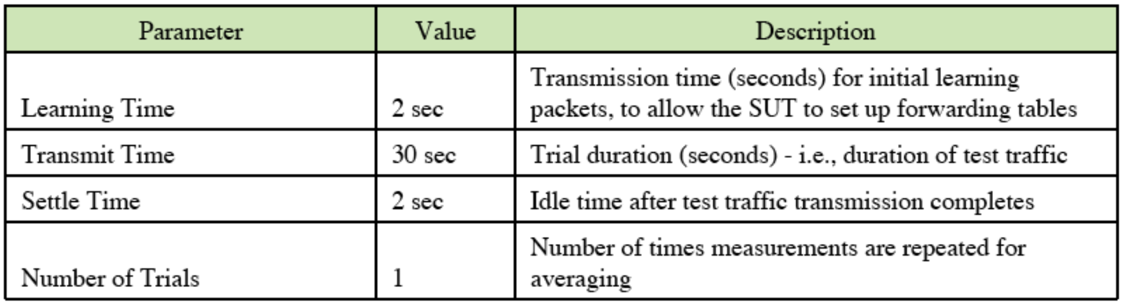

Parameter	Value	Description
Learning Time	2 sec	Transmission time (seconds) for initial learning packets, to allow the SUT to set up forwarding tables
Transmit Time	30 sec	Trial duration (seconds) - i.e., duration of test traffic
Settle Time	2 sec	Idle time after test traffic transmission completes
Number of Trials	1	Number of times measurements are repeated for averaging

Table 9: *Gigabit Ethernet: Latency and Jitter Test Configuration (Source: VeriWave)*

Table 10 shows the test results.

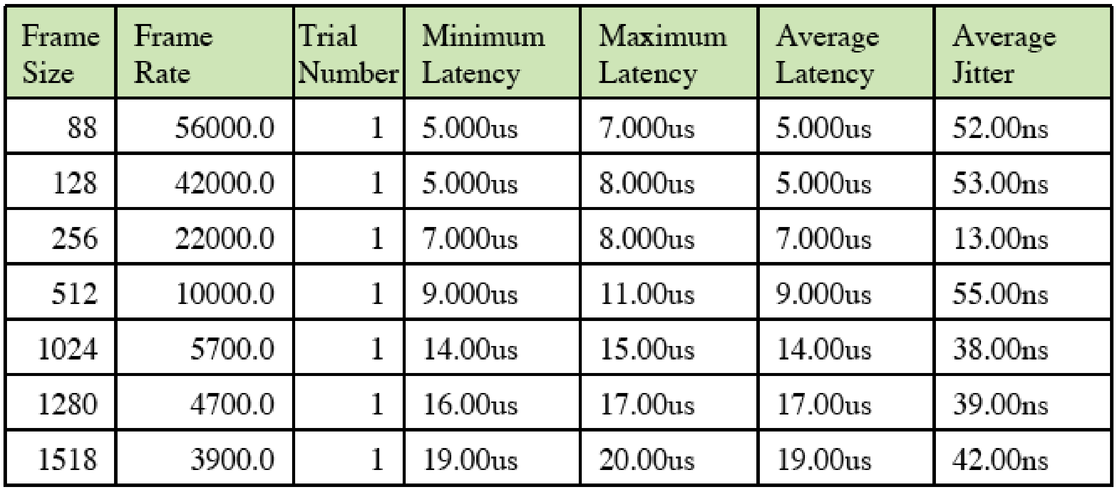

Frame Size	Frame Rate	Trial Number	Minimum Latency	Maximum Latency	Average Latency	Average Jitter
88	56000.0	1	5.000us	7.000us	5.000us	52.00ns
128	42000.0	1	5.000us	8.000us	5.000us	53.00ns
256	22000.0	1	7.000us	8.000us	7.000us	13.00ns
512	10000.0	1	9.000us	11.00us	9.000us	55.00ns
1024	5700.0	1	14.00us	15.00us	14.00us	38.00ns
1280	4700.0	1	16.00us	17.00us	17.00us	39.00ns
1518	3900.0	1	19.00us	20.00us	19.00us	42.00ns

Table 10: *Gigabit Ethernet: Latency and Jitter Test Results (Source: VeriWave)*

Conclusion

The Institute of Electrical and Electronics Engineers (IEEE) 802.11n standard represents the beginning of the end for wired Ethernet as the dominant local area network (LAN) *access* technology in the enterprise. Over the next few years, refinements in system silicon, radio design, network control, wireless security, and power management will improve 802.11n and its successor products to the point where they will begin to erode the switched Ethernet market.