



WHITE PAPER

Top Twenty Questions About 802.11n

What you should know about wireless

Author: Joe Epstein, Sr. Director of Technology

INTRODUCTION

Wireless local area networks have come a long way from their introduction nearly a decade ago as a convenience for home and small office networks. The largest companies and institutions run their daily business over large-scale enterprise 802.11n WLANs, dependant on the mobility, cost savings, and reliability that they can provide. The following twenty questions will help guide your investigation into different 802.11n solution offerings, highlighting where and how wireless networks can replace wired networks and focusing on the up-front and hidden costs of each.

Does the 802.11n network provide stability similar to that of wired networks?

Wireless networks provide the mobility and reduction in capital costs that wired networks simply cannot provide. However, many do so by sacrificing the stability and reliability that comes with wires. These wireless networks are *microcell-based*, which means that they frequently tune power levels up and down and shift channels as a part of their “adaptive” or “automatic” operations, thus eliminating stability, predictability, and repeatability. It is important to find out if a network offering you are looking at requires these outdated microcell techniques, or whether it is capable of offering a fixed-coverage solution that provides the same quality network for each device every day.

Can the wireless network have 99.99% availability?

As the role of wireless networks transform from being simply for convenience to becoming the mainstay for business-critical applications, IT organizations need to know that the network will be available for their users’ needs. Networks based on wireless LAN virtualization can have wired-like reliability.

Is the network able to be built out with 30% less access points?

A major weakness of microcell 802.11n is that it requires significantly more access points to seal up potential coverage holes and provide voice or video services. The cost of the additional access points, cable installations, switch ports behind them, and resulting increase in the required size and cost of controllers and licenses can make switching to wireless networks cost-prohibitive. Make sure to know how many access points you will need in the future to scale up microcell networks to voice, video, and real-time applications, even if you do not have an immediate need, so that you can understand the hidden costs.

Does the network provide ‘switched’ wireless, or does it behave like a hub?

The underlying Wi-Fi specification and IEEE 802.11n standards provides the choice for access points to behave simply, like a hub, or to use advanced techniques to act like a switch. The hub-based mode of operation throws all devices into the same wireless network (basic service set), fighting among themselves for limited resources and leading to the network slowing down or collapsing under density or load. The switched-based mode of operation puts each device

into its own virtual wireless network—such as the Meru *Virtual Port™*—keeping each device segregated and bounding contention to allow the network to scale to high densities, all over the same physical access points.

Can the network remain in control under high client density?

Being able to handle high and variable client density is crucial for any network where mobility or the size of the user population is a factor. Hub-based WLANs not only see their throughput crushed under density, but suffer from clients getting increasingly unequal throughput—some getting twice as much as others, or more. This leads to user complaints, and makes using wireless for many critical applications difficult if not impossible. Therefore, the question becomes whether throughputs can remain tightly bounded between the fastest and slowest client, and high overall, as the density ramps up.

How does its 802.11n performance compare to the leading performers?

Not all 802.11n offerings perform alike. In fact, the standard leaves a tremendous amount of leeway that can lead to both fast and slow implementations—benchmark speeds can differ by 40Mbps easily. Furthermore, the disparities are exaggerated when voice or video is introduced into the network. It is important to know how each offering compares to top performers. Look for validated, third-party tests, and skip the vendor-published reports.

Can the network provide wireless redundancy for every square foot?

Most wireless offerings provide redundancy of the wireline resources, such as high-availability modes, multiple controllers, and controller load balancing. But what about redundancy of the air itself: what can the offering do ensure that wireless access is redundant? Does it use half-measures, such as trying to stretch the range of neighboring access points to cover for one with access problems? Or does it allow for full channel layering, where wireless channels are deployed simultaneously for every square foot?

What does the network do to handle voice and real-time applications?

Voice and real-time applications are becoming major uses of wireless LANs, and the technology underlying the wireless network needs to be able to handle these services from day one. Does the network provide for seamless, interruption-free calls as the user hands off between access points? Can the network handle high-density voice without compromising or being hindered by high data performance? Are there public test results that back these claims?

Does it allow seamless mobile high-definition video?

Video over wireless has all of the real-time challenges of voice and the higher bandwidth requirements of data streams. Especially with 802.11n, the wireless technology must be

optimized to handle these unique aspects of video. Does the network support secure, high-bandwidth wireless multicast efficiently? What can the network do to keep the audio and video streams synchronized? Can high-definition video be handed off seamlessly as users move between access points, or does the video break and either buffer or simply drop out during that time?

How does the network bring down the soaring costs of wireless network operations?

A 2009 industry report on wireless LANs has shown a 63% increase in the cost of operating a WLAN, with a 69% growth in the number of employees dedicated to managing these networks in the last two years. It is not enough to reduce capital layout by replacing wired ports with wireless; you must also prevent the explosion of operations costs. Existing WLAN management tools, patterned off of switch-and-router management techniques, are not able to make a dent in the expense of solving problems that are uniquely wireless. New WLAN operations solutions replace WLAN management with proactive wireless intelligence centered around service assurance, rather than “speeds and feeds” reporting.

How well is the network client agnostic?

The differences between various wireline client interface cards are rather limited, but this is not so for wireless. 802.11n provides a bewildering array of different features and methods of implementing and optimizing, leading to wide swings in possible performance and reliability from each wireless card. The wireless network needs to be immune to these differences—one misbehaving or poorly-tuned client must not bring down the performance or stability of the rest of the users on the network. And all of this must happen using only standard clients, without requiring proprietary software or special features. Does the network deliver on its promise without relying on bundled proprietary extensions, which are not present on every device and may not always be at the latest versions?

Can the network predict when trouble is about to occur?

The more used a wireless network is, the more data is pulled into the network management system. Existing network management tools make this data available in graphs, charts, and reports, but leave it to the administrator to set traps or sort through the expansive set of charts to see if a problem might be occurring—a time-consuming proposition at best that leaves major holes in the oversight of the network. Instead, the network must be able to mine the data itself, using management signatures that infer and diagnose when problems are occurring or about to occur—before they become a problem.

Does the network proactively assure that the service provided fits the needs of the application?

When it comes to providing wireless as a strategic IT service, passively monitoring the quality of a wireless network simply does not cut it. When a critical application stops working, too many users suffer until IT finally is made aware of the. To solve this, the network and the applications running on it needs to be tested, vetted, and validated continuously. This concept of *service*

assurance makes sense only when tied directly into the network, when the network can test itself, proactively rooting out problems before they become the morning's crisis.

What true wireless problems can be detected and diagnosed remotely without sending people on-site?

When someone calls the helpdesk at 9am on Monday, reporting the problem they suffered at 4pm on Friday that may have caused them to go home early, how does the network help that get resolved? Does an IT worker have to go to the site to try to reproduce the problem, with sniffers and coverage analyzers? Or has the network already collected all of the necessary information, without even having to know that a trouble ticket might be coming, allowing the IT worker to simply roll back time in the diagnostic dashboard and see precisely what events were happening to that user's machine?

How does the management platform scale for large networks?

Wireless networks are only growing larger, more diverse, and spread out over wider geographic areas. Branch offices, remote campuses, teleworkers and road warriors all add to the size and complexity of the wireless network that needs to be monitored. And because the challenges of managing these networks are uniquely wireless, more granular, accurate, and better data is needed all of the time. Antiquated wireline techniques such as SNMP simply cannot scale to get the resolution of data for such broad networks without causing heavy scale issues on the core management infrastructure. And fully distributed management means that there is no centralized insight, and important problems and trends are completely missed. Look for a management and network operations suite that ditches the old *pull* model and replaces it with a highly-optimized, low-bandwidth *push* model designed to scale to far larger networks with granularity down to the second.

Can the network do physical-layer security?

Authentication and encryption (WPA2, 802.1x, AES, TLS, PEAP, etc.) cover the security at the packet layer and above. But the physical layer—the air itself—remains exposed, and transmissions continue to radiate out in all directions from the network, still with important information, such as device addresses, network advertisement in Beacons, and user identities in certificates unprotected by encryption. Physical wireless LAN security opens up the possibility of stopping transmissions at the RF layer itself. Look to physical layer security to address tough security demands in sensitive network installations.

Can the network stop zero-day passive attacks from outside the physical perimeter?

Wireless data breach remains the major security concern in retail and financial wireless networks. Because the signals continue to radiate through walls and outside of buildings into the surrounding area, any attack that depends on passively collecting data—whether encrypted or not—and then processing it to extract secrets can be successfully mounted simply by parking in the lot outside of the building in question and directing a high-gain antenna towards the building. No amount of tuning power levels down or using directional antennas inside the

building can stop this, because any change the administrator makes to the network can be countered by the attacker using a slightly more powerful antenna in the attack. Instead, physical layer security with technologies such as Meru's RF Barrier allow the signals to be proactively stopped at the perimeter, rendered into the useless Wi-Fi equivalent of "hello" messages when it reaches the attacker and devoid of any secret information.

Is security compromised when the network is running voice, video, or heavy data loads?

Part of wireless security is the ability for access points to scan for unknown, or rogue, access points and clients. Most enterprise-grade access points are dual-band, having one radio serve the 2.4GHz band and the other serve the 5GHz band. For an access point to scan, it needs to use one of the two radios to change channels and passively observe on another channel. How well a network can support scanning while simultaneously serving voice, video, or heavy data loads is a crucial factor. Beware of access points that, for each band, can only serve clients or scan—and be especially aware of access points that automatically disable security scans whenever voice, video, or significant data loads are present on the radio. Doing so leads to easily predictable blind spots that attackers can exploit.

Does the FIPS offering follow the principle of separation, or are too many non-security features embedded in the security boundary?

Government and other high-security networks look to the national FIPS 140-2 standard to ensure that the wireless network is capable of providing the advanced security for their needs. Products supporting FIPS 140-2 require a year-long certification program which locks down the version of the system and firmware running within the certification boundary. Some vendors place the entire wireless LAN system within the FIPS boundary, mixing in both security and non-security features into the same environment, violating security best practices for design and tying the hands of administrators who want to upgrade their WLAN feature set but are forbidden from doing so because the latest and greatest firmware is still in its year-long post-release certification cycle for FIPS. Look for architectures where the security processing is separated out from the underlying WLAN system and placed into a special locked-down FIPS gateway, allowing the underlying system to be changed, upgraded, or reconfigured as needed without compromising security or requiring using year-old versions of WLAN firmware. Also, consider going for the highest grade of FIPS 140-2 security possible—FIPS 140-2 Level 3 provides significant deployment cost benefits over the lower-grade FIPS 140-2 Level 2.



**advanced
network
systems**

For More Information,
Contact Us:

800.639.6757
sales@getadvanced.net

www.getadvanced.net