Special Report: Wireless Everywhere

Pervasive WLANs will inevitably break Ethernet's hold on the desktop. Is it time to set your people free?

By Frank Bulk

IT'S AN OVERSIMPLIFICATION to say that 802.11n heralds the era of the wire-free office—though with top speeds of up to 300 Mbps, it's clearly a catalyst for cutting the cords that tether users to their desks. Yet there's no question that within a few years, Wi-Fi will become the new network edge for companies interested in saving money, attracting top talent, and increasing security.

Of course, pure-play wireless LAN vendors have been saying for a while now that wired Ethernet to the desktop is dead, despite lingering concerns about reliability, the suitability of WLANs for telephony, the complexity of managing mixed wireless and wired networks, branch office and teleworker support … and, oh yeah, the fact that the legacy infrastructure is chugging along just fine .

Business technology managers have long weighed these factors against the most touted benefit of Wi-Fi: increased productivity. The efficiency studies are many and the refrain generally the same: Wireless keeps information at employees' fingertips, enables quicker decision making, reduces downtime, and enables collaboration. But in today's tight economic environment, the savings picture is just as compelling. Intel estimates—and we agree—that moving to a largely wireless network can reduce capital costs 40% to 50% and operational costs 20% to 30%. Luc Roy, VP of enterprise mobility at Siemens Enterprise Communications, cites a Canadian government customer that's saving $500 per event for moves, additions, and changes.

With the rising price of all modes of travel, teleworking is looking mighty attractive as well, and IT can now extend wireless to remote sites. Aruba Networks recently announced an access point, developed with Avaya and called the Mobile Remote Access Point, that can use any broadband connection to provide secure access to business resources for both data and voice. All the employee needs is a single- or dual-mode phone, or a softphone on a wireless laptop. Remote and branch offices also are obvious places to take advantage of all-wireless access (see story, p. 35), especially as management tools emerge for monitoring mixed-vendor WLANs (see story, p. 28).

Cisco Systems, Motorola, and others now offer 3G interfaces that can provide backups for branch offices and locations with minimal WAN connectivity, or for failover of critical applications. And WLAN security can beat that of most wired LANs—yes, you read that right. Sites looking into desktop virtualization should do fine on an all-Wi-Fi network as well, thanks to the small packet sizes inherent in virtual desktop infrastructures.

Motorola sees 802.11n as an inflection point in the industry and has adopted the slogan "Wireless

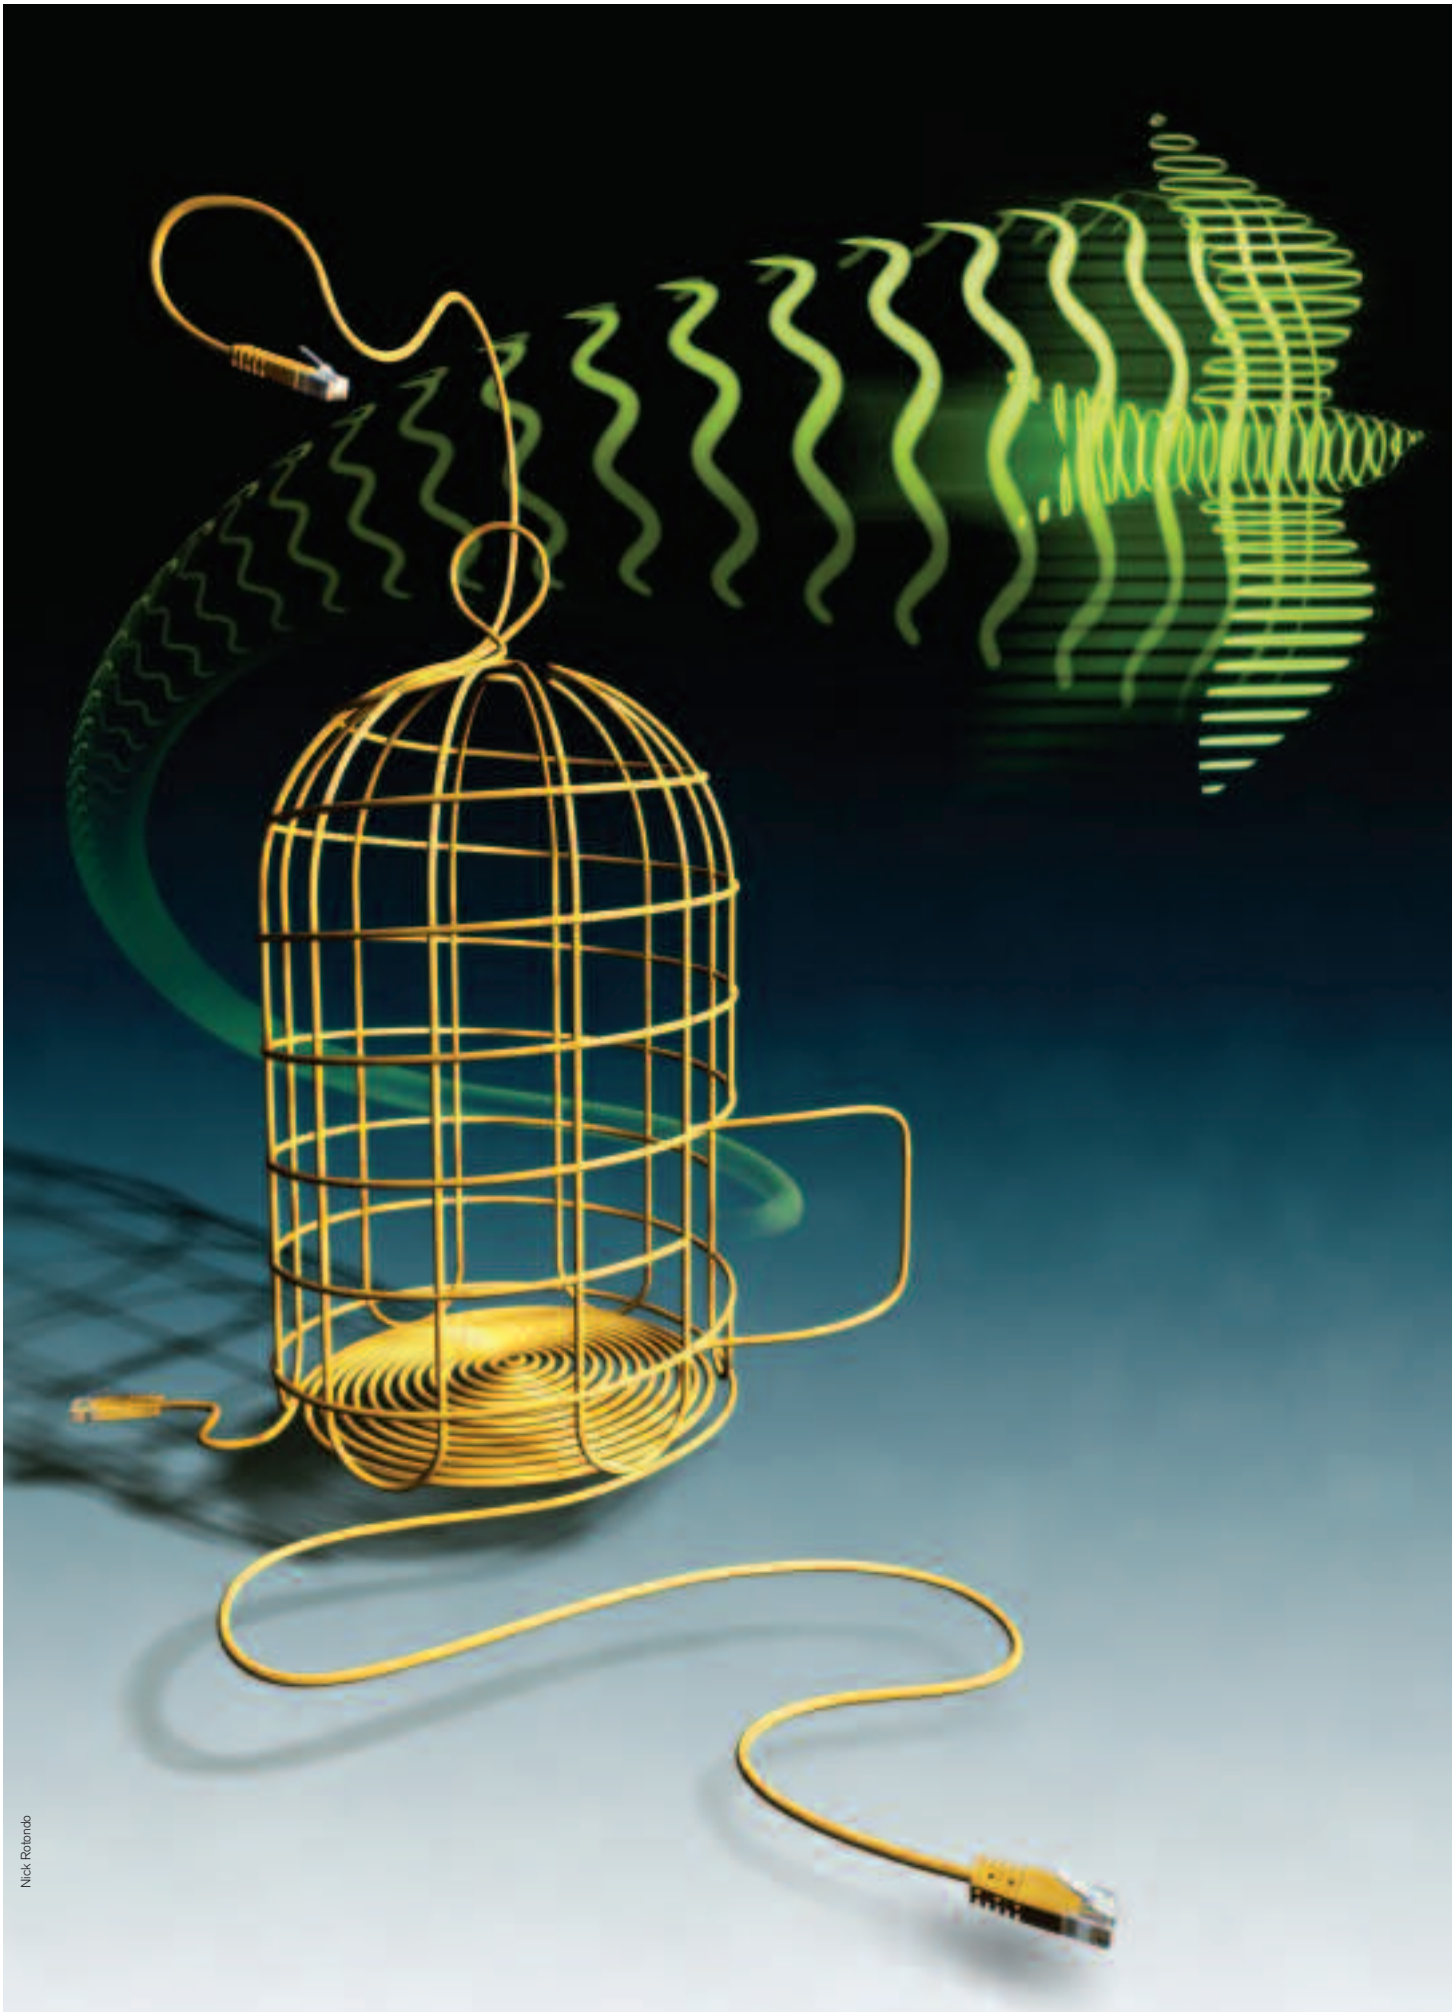

Nick Rotondo

by default and wired by exception." It and other vendors are practicing what they preach, deploying ubiquitous WLANs in their own offices.

Should you follow suit?

Although wireless vendors such as Motorola are happy to promote the wire-free office concept, Ethernet switch sellers, including Cisco and Hewlett-Packard, approach the concept with caution. That's not surprising: Switch vendors stand to lose big money as we move away from Ethernet to the desktop. Even if companies pay the manufacturer's suggested retail price for enterprise-class 802.11n gear, it's still much less expensive per user than a new 10/100/1,000-Mbps switch deployment with $250-per-drop wiring costs.

DIG DEEPER

THINKING ABOUT VO-FI? Our report examines the integration of VoIP, cellular telephony, and Wi-Fi. Purchase this *InformationWeek* Analytics Report at:

informationweek.com/1171/analytic_vofi.htm

See all our Analytics at informationweekanalytics.com

CAREFUL STEPS

But don't feel too bad for Cisco—no enterprise WLAN vendor is claiming to replace wire at the core or distribution layers, and besides its wire-side dominance, Cisco owns more than half of the enterprise WLAN market with its wireless gear set, originally from Aironet and later supplemented with its Airespace acquisition. Chris Kozup, manager for mobility solutions at Cisco, emphasizes that the company is making the most of its leadership in both wired and wireless with a "unified" network approach that blankets the office with Wi-Fi while keeping a few wired ports at every workstation. Nice if you can afford it. Cisco is clearly cautious in its pronouncements regarding the all-wireless office. Don't look to the WLAN gear leader to be in front of this charge.

No. 2 switch vendor HP, which mixes some of its own Wi-Fi gear with licensed technology, is also approaching the all-wireless office carefully. Andre Kindness, Americas security and mobility solution manager for ProCurve networking, says HP's customers are driving that stance. Companies are looking to reduce their operational costs through a consistent management system that covers both wireline and wireless and provides product longevity, Kindness says. However, such management doesn't yet exist. Cisco talks about a unified network, but it's not yet providing integrated management. HP openly discussed the problem of inconsistent management tools between wired and wireless networks, and we see it making the most credible progress of any of the "we do both wired and wireless" players. Other vendors looking to cover these bases include Nortel Networks, which says it's developing its

Impact Assessment: 802.11n And The Wire-Free Office

	Benefit	Risk
IT organization	802.11n delivers speeds that exceed 100 Mbps to the desktop while enhancing reliability and coverage. Capital and operational savings over wired connections if the legacy infrastructure is fully depreciated and can be mostly turned off.	The 802.11n standard has not yet been ratified by the IEEE; moving early may mean higher prices and immature chipsets. Because wired networks are not going away anytime soon, it's another set of equipment to manage.
Business organization	Unchaining an end user's tools (laptop and phone) from her desk should increase productivity by enabling anyplace access.	A poorly functioning wireless network can lead to user frustration. Some employees may feel uncomfortable taking communication tools out of offices and won't readily change work patterns.
Business competitiveness	Companies are all about doing more with less. A wire-free office based on 802.11n will cost less than a comparable wired LAN, freeing up cash. Security is improved as well.	Offices with WLANs for primary network access are relatively rare, and productivity increases aren't guaranteed. This is one competitive difference that can be easily duplicated.

Bottom Line

While the 802.11n standard isn't yet ratified, the Wi-Fi Alliance's certification should put most worries to rest. Because the 802.11n market is so nascent, businesses can gain a first-mover advantage thanks to operational savings. Still, Wi-Fi is still more a black art than wired Ethernet. That said, the productivity, cost savings, and flexibility benefits are compelling.

own 802.11n gear—essentially shunning its OEM partner, Trapeze Networks—and Enterasys, Extreme Networks, Foundry Networks, and Juniper Networks, all of which are OEMs or resell wireless products.

Meanwhile, overlay vendors such as Aruba, Motorola, and Trapeze treat the wired network as more or less a dumb transport for their wireless traffic. It makes for easier sales to the wireless-oriented parties in IT organizations, but this stance leaves those who must manage both with a less-than-easy feeling.

Another angle enterprise switch vendors play is to suggest that all-wireless is a better fit for the remote or branch office, rather than main sites, appealing to interest in this architecture while protecting their wireside revenue. Most also deliver some variation on the message that IT should be about "providing flexibility to the business"—in other words, preserving wired connectivity where it exists and delivering wireless where it's wanted. Tim Purves, CTO of the Henry Ford Health System in Detroit, says it's his department's aim to "align technology with business workflow processes." While that's a familiar mantra, if those processes are tied to immobile approaches that ignore the productivity increases and workflow improvements possible via a pervasive wireless network, IT must step up and champion a new way forward.

Fortunately, not all enterprise switch vendors are stonewalling. Trent Waterhouse, VP of marketing for Enterasys, says his company sees wireless as a strategic component of its business and is evaluating WLAN players with an eye toward an acquisition. Juniper is shopping around, too; it was spurned by Meru Networks, which also acts as an OEM for Foundry, on at least one occasion, say industry sources. No matter—Aerohive, Bluesocket, Colubris, and Xirrus stand out as attractive acquisition targets for enterprise switch vendors that lack their own wireless products. Trapeze might be a good fit for Nortel, if it decides to turn to its former OEM partner rather than build its own 802.11n gear.

MAKE THE MOVE

Truly transforming the workspace extends beyond installing access points and providing laptops, to physical reconfiguration. Take Capital One's Future of Work pro-

AirWave Aims To Manage Mixed WLANs

THE WIRELESS NETworking market is changing at a dizzying pace, with once-cutting-edge hardware sets becoming legacy gear in the blink of an eye. The result: Many companies—especially those working toward all-wireless offices—find themselves proud owners of mixed-vendor wireless LANs.

To address management of these networks, AirWave, which was acquired by Aruba Networks in January for $37 million, brings a new mix of tools to version 6.0 of its multivendor Wireless Management Suite, or AWMS. While the suite doesn't have all the pieces required to completely replace vendor-specific tools, it's the closest we've seen to a heterogeneous WLAN management product.

We tested a beta version of AWMS 6.0 at our Syracuse University Real-World Labs. The new code handled a fair portion of the campus' Lightweight Access Point Protocol (LWAPP)-based production AirOrange Network. With AWMS 6.0 managing or monitoring IOS-based point-to-point bridges, LWAPP and Aruba 802.11n access points, Radius servers, a Cisco Wireless LAN Solution Engine, a variety of network switches, and hundreds of clients, it didn't take long to see the value proposition. We also gave the product full configuration control over several test devices in a nonproduction lab environment. AWMS supports new 802.11n access points from Cisco Systems and Aruba, and support for 802.11n gear from Meru and others is due soon. As in previous AWMS versions, the complete list of wireless products supported reads like a who's who in the wireless industry, including mesh and WiMax product sets from 3Com, Tropos, and almost everyone in between.

RIGHT ON SCHEDULE

One of the most valuable features in any wireless management system is the ability to schedule tasks. In this regard, AWMS shines. New to 6.0 is the ability to specify standard date/time formats, and we found greater flexibility in planning downtime, a must-have capability given that most production WLANs have become critical resources.

Also new in Wireless Management Suite 6.0, AirWave provides help-desk functionality to go with its hardware management capabilities. For example, if a wireless user can't authenticate to a secure wireless network segment, first-level responders with role-appropriate AWMS access can capture symptoms and feedback with screen shots and annotations that can be integrated with systems like Remedy Service Desk or used within AWMS.

When we induced failures during testing, AWMS revealed a wealth of information, enabling better and quicker escalation. And for large companies with several wireless networks spread out across multiple geographic areas, enhancements made to the Master Console architecture allow distributed AWMS systems to be controlled from a central location.

gram. The financial services firm's 360-acre, eight-building campus almost doubled the number of employees it could house, from 650 to 1,100, by adopting the concept of hoteling. Rather than being assigned a specific location, employees who participate in this optional program have access to a generic cubicle, as well as conference rooms and open areas. Space is essentially overbooked. Each employee is assigned a telephone number that flows to a Cisco voice-over-IP phone and/or BlackBerry. The WLAN is the primary medium for network access.

"Today, work is what you do, not a place you go," says Rob Alexander, Capital One's CIO. "The wireless and mobile technologies we provide through our Future of Work environment provide our associates greater flexibility in how and where they work, which in turn improves collaboration and productivity." Employees are happy, and the company saves big on facilities.

Intel takes a similar approach at its Jones Farm, Ore., campus. This location serves almost 6,000 employees using Cisco wireless gear. Intel started with an overlay network for wireless access, but as Wi-Fi caught on, it's become the first choice for employees. In addition to Centrino-based laptops (of course), Intel also uses Cisco Wi-Fi phones for voice services, as well as softphones and dual-mode devices.

Cisco has its own initiative, called the Connected Workspace. In line with its preferred converged approach, wireless is deployed everywhere, but wired ports for high-bandwidth communications needs, such as backups and video streaming, also are available. Still, the company has cut its need for copper by 60%. "The Connected Workspace encourages collaboration and reduces real estate and infrastructure costs, while accommodating different work styles," Cisco's Kozup says.

Aruba and Motorola have been the most vocal vendor supporters of the wire-free office. With no wired revenue to lose, they can only gain by stealing away dollars that would normally be spent on their competitors' Ethernet switches. With 802.11n offering comparable performance to a wired network, but with added mobility, they have a strong argument.

Of course, the wireless office is like the paperless of-

AWMS 5.0's VisualRF module was quite clunky. In 6.0, VisualRF has improved; for example, the SVG format has given way to Flash for expanded browser compatibility. AirWave also heeded the call for easier floor-plan imports by adding native support for bulk import of CAD files, and AP placement may be carried over if provisioned properly.

AirWave provides robust client location services, and recognizing increased use of radio frequency identification, the Wireless Management Suite now supports tracking of AeroScout RFID tags when used with Cisco LWAPP controllers.

As with other network management tools, bringing AWMS to a functional state requires discovering devices and building profiles, policies, and the general management framework for a given environment. AWMS is fairly intuitive in this regard, and it was a swift process getting our devices found and managed.

Still, while AirWave has made great strides with Wireless Management Suite 6.0, it isn't yet a full replacement for vendor-specific management platforms. Searching for a client device in AWMS is just as easy as doing it in Cisco's Wireless Control System—but if you need to push a configuration template to an LWAPP controller that's denying the client WLAN access, AWMS can't help. Same with configuring access control lists available through Wireless Control System's proprietary mechanisms—AWMS isn't there yet.

That said, for networks that use multiple wireless systems, and for those who aren't satisfied with the management products offered by their WLAN vendors, AirWave's Wireless Management Suite 6.0 provides an attractive alternative. The Professional Edition license lists for $36,995 and supports as many as 1,000 devices, including APs, controllers, routers, and switches. The master console is an additional $14,995.

—LEE BADMAN (lbadman@nwc.com)

THE UPSHOT

» **CLAIM:** AirWave wants to lure customers away from vendor-specific WLAN management systems with its Wireless Management Suite 6.0. Heterogeneous environments may benefit from an overarching management console that provides enterprise-class features and supports a broad range of devices.

» **CONTEXT:** Rival tools from wireless hardware vendors tend to manage only their own products—and sometimes not very well. If you have WLAN gear from multiple sources, your headaches are magnified. A vendor-neutral management system is the best medicine, if you can afford it.

» **CREDIBILITY:** Configuration scheduling, help-desk-related data gathering, and powerful visuals all contribute to the effectiveness of a wireless management system. AirWave nails these but needs to go beyond SNMP-controlled capabilities to be a full replacement for vendor-supplied systems.

fice—though electronic documents and e-mail have become the main forms of information storage and redistribution, there's still paper exchanged in the postal mail. In the same way, wireless will become the primary connection only at the access layer. " 'All wireless' is a bit of a misnomer," says Kozup. There will still be cables, but they'll reside predominately in the distribution and core layers of the network, unseen by the average user.

SECURITY MATTERS

The security breach at TJ Maxx parent TJX, where attackers took advantage of a wireless connection secured only with Wired Equivalent Privacy to capture credit card information on tens of millions of the retailer's customers, remains fresh in many CIOs' minds. The fact that the key element in that equation is "secured only with WEP" is a detail easily ignored by the security paranoid.

Done right, Wi-Fi can be deployed with greater security than wired networks, which often leave ports unprotected in cubicles and conference rooms. Because security concerns have long been a drag on WLAN adoption rates, it's now standard form to use 802.1X to ascertain a connection's user credentials and the Advanced Encryption Standard to encrypt traffic until it reaches a wireless controller in the data center or at the network edge. Those still using a VPN overlay on an open wireless network, take note: Unless you have specific application requirements or hardware limitations, now is the time to move to 802.1X with AES.

A wireless network's greatest vulnerability is in performance-degrading interference or denial-of-service techniques, some facilitated by options in the 802.11n standard. Your wireless infrastructure management system may be able to pinpoint the source of malicious traffic, or else a product from an overlay wireless intrusion-prevention system vendor like AirDefense, AirMagnet, or AirTight can do that and more. Work on the 802.11w standard is progressing to offer management frame protection, among other capabilities, to fill gaps.

PEOPLE, GET READY

If you have some sentimental attachment to the copper feeding your desktop, consider that your future workforce has spent the past four years in a wireless oasis. Most colleges and universities provide Wi-Fi in a substantial portion of their classrooms and public spaces, some in their dorms. Freshly minted graduates expect mobility when they step into the workforce, and that starts with Wi-Fi access in the office.

If businesses want to attract young talent, staying on the cutting edge isn't optional. To see how close we can come to going wire-free, we broke down wireless communication into three areas: data, voice, and video.

Conventional office applications account for the majority of data access. Whether e-mail, productivity suites, or line-of-business applications, data apps consume the largest amount of a knowledge worker's time and have been successfully mobilized, in and out of the office.

Wireless voice is often thought of in terms of cellular services, but voice over Wi-Fi, or Vo-Fi, increasingly is considered a key application for wireless networks. CIOs are generally cautious about running voice over their enterprise WLANs, for good reason: Unless the wireless network was engineered with voice in mind, whether it be first- or third-generation gear, poorly implemented quality-of-service functions and a weak sig-

Copper Costs Lots Of Pretty Pennies

IN A DECEMBER REPORT, Gartner analyst Ken Dulaney predicted that by the end of 2011, 70% of all new worldwide voice and data client-to-LAN connections will be wireless. The firm also estimated that $100 billion will be wasted over the coming five years following outdated network design principles.

Included in that figure is Gigabit Ethernet to the desktop.

The takeaway is that all organizations need to ask a fundamental question: "What is our strategic platform for network access connectivity?" If the answer isn't "wireless," you need to take a hard look at up-front and ongoing costs—and possibly reconsider your stance. In our experience, wiring Ethernet commonly hovers around $250 per drop, though site-specific considerations such as a historical building or union labor can double or triple that cost. Once you've paid for the copper wire, there are edge Ethernet switches to consider at anywhere from $50 to $100 per port.

According to Motorola, which admittedly has a stake in this game, a wired network costs $88 per user per year for maintenance and support, compared with $12.51 for a WLAN. While that might be a bit optimistic, there's no way around the fact that purposing a single wired Ethernet port and cable to serve many clients via an access point does in fact translate to significant savings.

While you'll still pay a premium for 802.11n APs—Cisco's new 11n 1250 AP was priced at $1,299 for the dual-radio version we tested recently— we expect prices to drop over the coming year.

—FRANK BULK

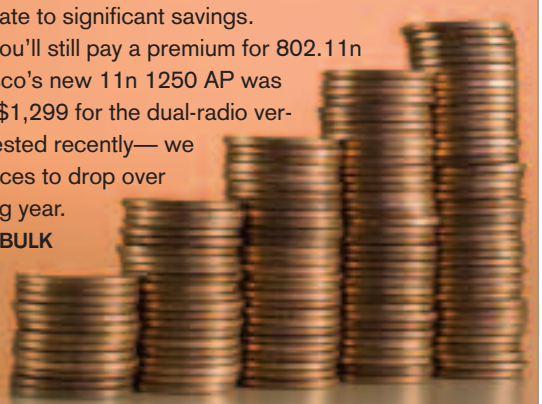

nal will lead to disappointed users. All the major WLAN infrastructure vendors have spent considerable time working with enterprise-class Vo-Fi providers, such as Cisco, Polycom (formerly SpectraLink), and Vocera, developing deployment guides to assist VARs and IT groups with configuring the WLAN for QoS.

Wireless video, which generates much higher traffic volumes than voice, requires special consideration as well. Although we don't see enterprises deploying Cisco's TelePresence over Wi-Fi anytime soon, video-based corporate training and closed-circuit television for both inside cameras and those mounted in the parking lot are here now.

Not all apps can be neatly siloed into voice, video, and data. Environmental controls and security monitoring can also be performed wirelessly, eliminating time-consuming and expensive installations. Services such as location and presence increase productivity and security. We're in the midst of a Rolling Review covering location systems, and we like what we see; check out our findings at *informationweek.com/rollingreviews*.

Voice Options For The Wire-Free Office

	PROS	CONS
Softphone on laptop	Works anywhere with wired or wireless connection; wide PBX support	Laptop must be on to take a call; requires headset or earpiece
Mobile cellular phone	Form factor and experience well-understood by users; variety of providers and pricing plans	Indoor coverage typically challenging; no PBX integration; may not be acceptable for regulated industries
Voice-over-Wi-Fi handset	Truly portable voice option; coverage wherever your WLAN reaches	Few PBX vendors offer Vo-Fi systems; requires strict attention to RF design
e-FMC phone that supports cellular and Wi-Fi	Best of both worlds	Usually requires integration and sophisticated handsets; nascent market

NEED FOR SPEED

Throughput is the first consideration when it comes to network connectivity, and 802.11n delivers: Both vendor and independent tests have shown that peak rates upward of 130 Mbps are achievable in good conditions. Advanced antenna designs, spatial streams, and multiple input/multiple output (MIMO) technology mean 11n also offers better coverage and improved radio frequency reliability and consistency. Access points can be spaced farther apart, if desired, but the better signal may more effectively be used to achieve higher access rates. Multipath, which previously degraded signal quality, is now used to good effect by MIMO to reduce the effects of fading and interference.

There are other benefits of 802.11n. First, it's essentially the fourth generation of the 802.11 standard, yet despite the evolution, each revision is backward compatible on both clients and access points, albeit at lowest common denominator rates. Companies can upgrade gradually because 802.11n clients work with 802.11a/b/g APs, and vice versa.

Second, as the market developed, amendments have been added to address deficiencies in the original 802.11 specification. The most significant are 802.11i, which deals with security, and 802.11e, which introduced quality-of-service features. Architectural approaches also have broadened. First-generation access points were standalone, with little to assist IT in terms of scalability, RF management, and Layer 3 roaming. Startups generally swung to the opposite extreme and centralized everything, leading to what pundits called "thin" APs.

With development of 802.11n and its higher traffic rates, a more sensible distributed approach, first used by Colubris in 2005, has evolved. The management plane remains centralized, as is common in any enterprise service framework, but the control and data planes can be placed at the core, edge switch, or access point. Motorola calls this "adaptive AP," while Trapeze has taken the moniker "Smart Mobile." Even Aruba, with its emphasis on centralized data flows, provides flexibility as described earlier with its Mobile Remote Access Point. Even if the WAN link is interrupted, connections stay up and local traffic will continue to be switched locally.

With 802.11n just around the corner, early adopters whose 802.11b/g gear is nearing end of life face a conundrum: Pay top dollar for 802.11n, stick with b/g, or add 802.11a support to their access points by buying new gear or moving to a different vendor. While 802.11a buys some advantages, at this point we recommend sitting tight until prices, AP maturity, and/or standard

adoption are such that you feel comfortable upgrading to 802.11n. In fact, Aruba has a new marketing pitch: Buy its 802.11a/b/g APs today, and buy a key later to activate 802.11n. This approach helps customers split their costs over time—and assures Aruba market share.

It doesn't help purchasing decisions that the 802.11n standard isn't complete. Working group approval is tentatively scheduled for March 2009, many months past predictions. Vendor adoption of the draft 2.0 spec, along with all the pre-standard chipsets already in use, make it highly unlikely that a final standard that's incompatible with existing products will be adopted. Nevertheless, we can't argue the logic of waiting. Second-generation standards-based 802.11n products, even if functionality equivalent, will have many of the bugs and kinks—for example, 802.3af Power over Ethernet support—worked out. Prices will drop, and processes regarding site planning, installation, and maintenance will be better defined.

Enterprise network administrators also are concerned about reliability. Will that unforgiving terminal session or enterprise application drop every time the microwave goes on in the cafeteria? There remain a plethora of wireless supplicants, and connectivity is still not as certain as with Ethernet. With proper device selection and configuration, connectivity bugs can be minimized, but there's still room for improvement. Most users will trade a few connectivity blips for mobility. Some won't.

THERE'S MORE TO 802.11N THAN JUST SPEED

Note that 802.11n is more than just a catalyst for the wirefree office: It's also making mesh-based office networks viable. Mesh is the extension of connectivity using wireless rather than wired backhaul. Each access point, or node, connects back with another node until it reaches a gateway, which is a node with a wired connection. To date, mesh setup has been performed simplistically using a WDS (wireless distribution system), but that's insufficient for the magic sauce that most vendors add to their products.

One issue with mesh is that with a single radio for backhaul, each hop in a mesh network has a reduced level of performance. When wireless networks operated at peak speeds of 30 or so megabits per second, half or a quarter of that would be less than acceptable. But with speeds ranging from 120 Mbps to 140 Mbps with 802.11n, even a few hops would give most users an adequate experience. Vendors prefer to use a separate radio, traditionally operating at 5 GHz (previously 802.11a, now 802.11n) , for backhaul, while using 802.11b/g at 2.4 GHz, or a second 5 GHz radio.

Mesh can help reduce cabling requirements across the board, but it's more commonly used to extend connectivity to areas that are inaccessible for wiring, or for outdoor service. Mesh could be better thought of as a fall-back or high-availability mechanism in case the wired backhaul fails. In that situation, the AP would associate with one of its radios to a near-by access point and transport its client traffic via that new wireless backhaul. Aerohive communicates this concept most effectively with its line of products, but most enterprise WLAN vendors support some kind of mesh offering.

Also mentioned earlier are the advantages of the distributed access point. They provide flexibility and remote survivability. Again, this is not directly related to 802.11n as a standard, but these features are being fleshed out by vendors, helping make the case that the wire-free enterprise is an attainable goal.

Talk Vo-Fi To Me

VOICE SERVICES ARE A SIGNIFICANT sticking point for companies considering an all-wireless office. Voice over Wi-Fi is still in its infancy. Fixed-mobile convergence using a dual-mode handset is another approach, but that technology and market are no more mature than Vo-Fi. Standardizing on cellular voice services is also a possibility, but that requires good in-building coverage, something rarely found in North America without investing in special gear, such as femtocells.

Until the voice question is definitively answered, organizations have good reason to be cautious.

Even if Vo-Fi isn't yet on your radar, your WLAN needs to be designed to support quality of service and multiple applications (data, voice, and video). While wired networks need the same type of planning, the added latency of the air and that fact that Wi-Fi is a shared medium require specific attention. And QoS won't help you if the radio in the employee's Vo-Fi phone can't connect. Building the wireless network to saturate all locations, including hallways, elevators, stairwells, and utility closets, is critical to providing complete in-building coverage. For that reason, Vocera, well known for its *Star Trek*-like voice badges, uses AirMagnet's VoFi Analyzer in all its installations. Depending on your building's location and setup, outdoor coverage may be required as well. If you already perform post-installation surveys, that's a good start, but remember that the radios in Vo-Fi phones are weaker than the clients built into most laptops. **—FRANK BULK**

GOTCHAS REMAIN

Despite all the performance and other benefits of 802.11n, there are still questions about reliability, performance, legacy devices, integration into the existing wire-centric infrastructure, and market dynamics.

RF remains a black art, and although MIMO makes Wi-Fi more reliable, it's still no guarantee that interference won't interrupt. For starters, good planning is required, perhaps using a tool such as Cisco's Spectrum Expert (formerly Cognio), which identifies possible sources of interference. There are also architectural approaches to address the reliability problem. Meru's newest 802.11n access point, the AP400, was designed for robustness. Its four built-in radios can operate simultaneously, on different channels; interference on one channel or band doesn't prevent a client from roaming to another radio. Another approach, used by Ruckus Wireless and Xirrus, is to employ directional antennas. These approaches are still considered a bit unconventional, but they're worth watching.

If aggregate performance is a key issue, legacy clients that operate only in 802.11b mode may need to be replaced or upgraded. That's not always possible with older Vo-Fi handsets, portable scanners, and other application-specific devices. In these cases, moving nonlegacy clients to the 5-GHz band, where there's great channel selection and support for multiple 40-MHz channels, may be prudent. That way, the legacy clients won't impact the peak-performance capabilities of the 802.11n-capable gear.

The challenge of consistent network management between wired and wireless networks is also vexing. Even Cisco, which leads in market share in both segments, doesn't have a management interface between both platforms. As HP points out, enterprises aren't eager to layer on a different set of intrusion-detection and -prevention systems, security, and network-access control tools for the wireless environment. For now, you will need different sets of tools for managing wired and wireless networks, so for organizations that do both, back-end support costs will rise, not fall.

Frank Bulk is an InformationWeek contributing technology editor. He works for a midwest-based telecommunications company. Write to him at fbulk@nwc.com.

Verticals Challenged

THE EARLIEST Wi-Fi adopters—higher education, healthcare, and retail—are not necessarily the first to go "all wireless," or even the first to trial 802.11n gear. Still, it's worth looking at what these old hands are up to.

Higher ed has been a steady adopter of Wi-Fi. Colleges enjoy a young and mobile user base and have a desire to differentiate, and with 802.11n, IT groups can address the challenges of demanding applications and dense wireless usage in lecture halls. And it doesn't hurt on the marketing front, either: Morrisville State College in central New York made sure prospective students knew that the school was one of the first in the nation to adopt 802.11n. But wireless has also become substitutive: Only 10 to 15 years ago there was a significant move to pull Ethernet cable "port per pillow." For schools that have deployed Wi-Fi in the dorms, the wired network is seeing a significant drop in usage. Nowadays students can't be bothered to rummage behind a stack of books for the Ethernet jack.

Healthcare might have the highest percentage of mobile knowledge workers, so it's interesting to see the somewhat polar views of wireless connectivity: Some healthcare IT pros consider wireless a reliable means of delivering services, while others point to the life-and-death nature of their business and prefer cable. What all agree on is that wireless can enable new services, such as providing continuous monitoring of patients while moving them between rooms, or speeding up the process of locating medical equipment. With 802.11n there's the possibility that x-ray images, typically large in size, can be retrieved from mobile devices and viewed, right in the room. The new standard also gives additional headroom to data in relation to voice, and extra capacity for non-medical services such as patient or guest access.

The retail environment has long used wireless and is the source of Symbol's past dominating market share in the enterprise Wi-Fi industry. While 11n provides minimal benefit to the multitude of legacy client devices such as handheld scanners, better coverage and potentially new marketing mechanisms, such as streaming location-specific advertisements to tablet-equipped shopping carts or ceiling-mounted TV monitors, is catching interest. There's also the more mundane element of CCTV for security cameras.

The hospitality industry, encompassing hotels, convention centers, and casinos, has seen widespread use of wireless for its employees and guests. What's unique about this environment is that the number of guest users, and the square footage that those guests occupy, often dwarf the facility's own workforce and work space many times over. Once the wireless network has been proven to work well for guests, is there any need for employees to use anything different? In fact, for renovations and new sites, wireless becomes the primary mode of connectivity.

Professional services such as real-estate and legal as well as smaller financial services such as insurance, tax processors, or loan agents, are obvious targets for wire-free offices. Besides the desire of professionals to be mobile, their offices are often leased spaces that change with business conditions and opportunities. They likely lack an IT person to run cables and operate switches, but once a Wi-Fi access point is put into place, they're free to work anywhere around the office. **—FRANK BULK**

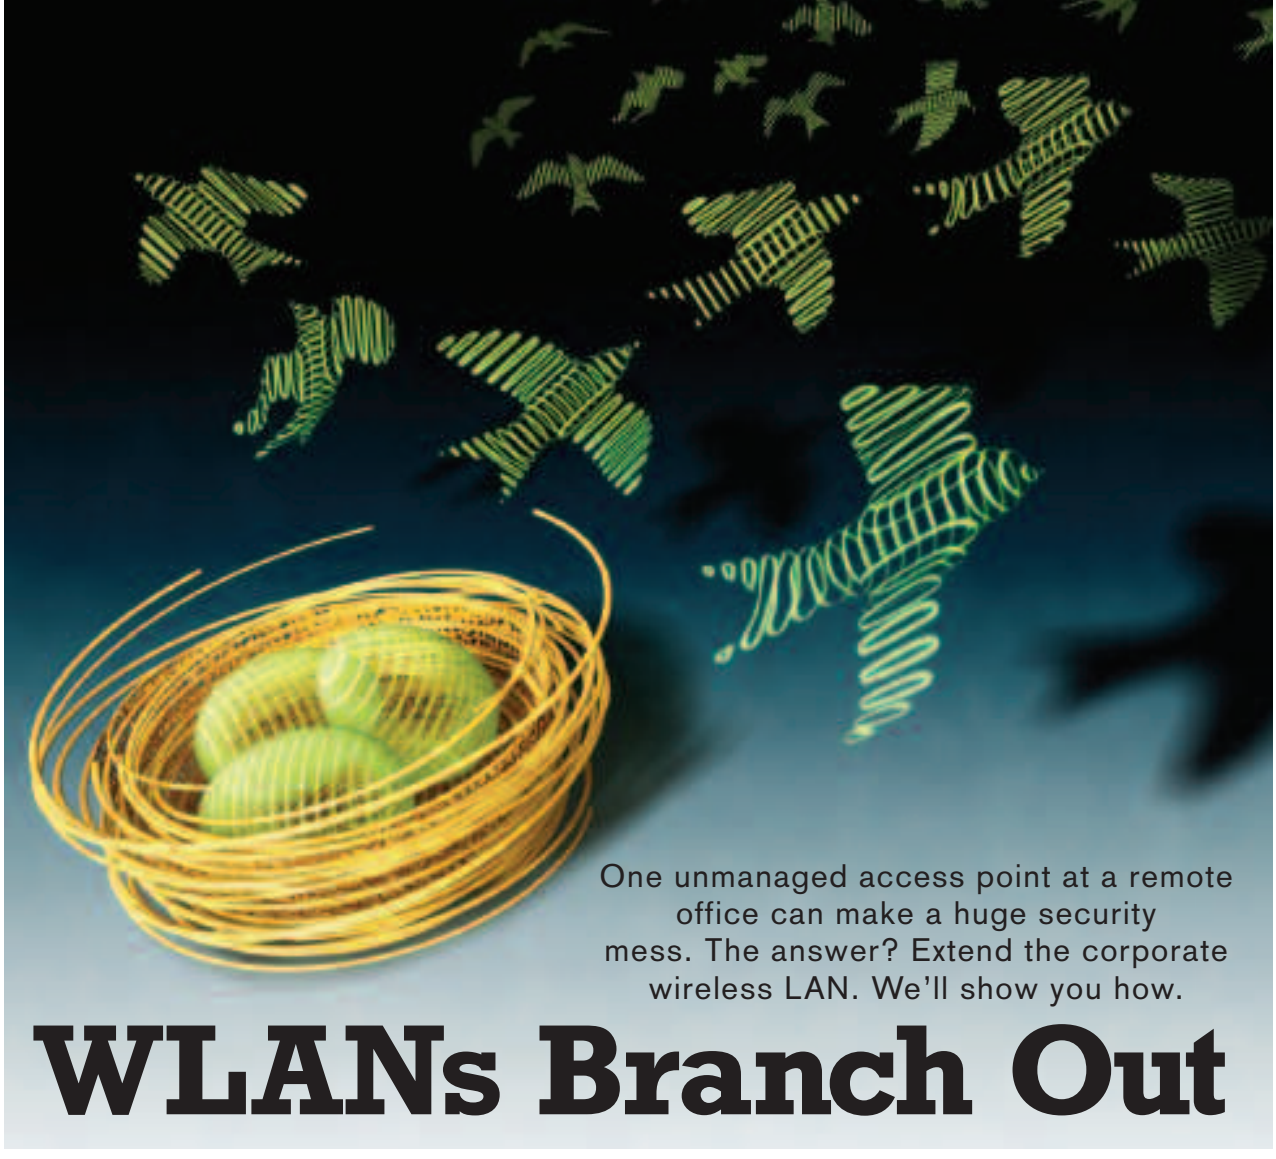

One unmanaged access point at a remote office can make a huge security mess. The answer? Extend the corporate wireless LAN. We'll show you how.

WLANs Branch Out

By Richard S. Dreger Jr. and Grant P. Moerschel

REMOTE USERS CAN FEEL MARGINALIZED if they don't have the same technology amenities that employees at headquarters enjoy, and they won't take design complexity, management overhead, or security risk as an excuse. A prime example is a branch office that deems itself underserved because "everyone else has wireless." Employees might just pitch in to buy a $50 access point and believe they're doing the corporate IT folks a favor by solving the "problem" themselves.

Of course, security is only as strong as its weakest link, so that $50 rogue access point could neutralize thousands of dollars' worth of sophisticated, layered access controls. Put simply, an open AP connected to the corporate network is tantamount to placing an Ethernet jack in the parking lot. Even when the device is configured with Wired Equivalent Privacy, it's vulnerable. Armed with a high-gain antenna and a proximate location to the target, an attacker can inject and/or collect 802.11 data frames and recover static WEP keys and passphrases used by the "helpful" employee who's attempting to secure his unauthorized device.

To make matters worse, once someone gains access to the remote office's network and obtains a valid IP address, the intruder could appear, at least from a network perspective, to be an authorized corporate user. Unless you have network access con-

Nick Rotondo

trol or core firewalling in place, the attacker may well gain access to all local and WAN-connected corporate assets via the branch-office connection.

With the advent of enterprise-class 802.11n systems, the remote WLAN equation becomes even more complex. The upside is that 802.11n will greatly increase the throughput rates of each AP radio while enhancing IT's ability to identify rogue devices. The downside—besides the enormous cost premium that 11n gear commands—is that it will be even easier for wireless users to saturate available WAN bandwidth.

The best answer for geographically diverse organizations now may be to bite the bullet and provide enterprise-class 802.11 WLAN coverage at branch offices. While you could just stick lightweight access points at remote sites, link them to the controller at your main office, and call it a day, problems with subpar connectivity and bandwidth hogging make this a poor choice. Better are scaled-down WLAN controller appliances from companies such as Aruba Networks, Cisco Systems, and Motorola-Symbol that can support as many as six access points while providing many of the sophisticated capabilities available in controllers that scale to well over 1,000 APs.

Alternately, manufacturers such as Aruba and Cisco offer enhanced systems designed to extend corporate WLAN standards to branch offices while addressing the bandwidth constraints inherent in WAN connectivity. Aruba's Remote Access Points and Cisco's Hybrid Remote Edge Access Points use standard lightweight APs loaded with specialized firmware that integrates seamlessly with centralized WLAN controllers, letting branch offices enjoy the functionality and security provided to headquarters without the need to deploy local WLAN controllers—or have advanced IT resources on site to maintain them.

DESIGN TIME

The dominant WLAN architecture secures wireless access using a strong controller and lightweight managed APs. This centralized approach makes it easy to create WLAN profiles to provide tailored wireless access to diverse groups. For companies using multiple controllers, the addition of a wireless network management system can bring all WLAN infrastructure components into a single management interface.

There are a few major design requirements to keep in mind when rolling out any 802.11 wireless service to remote sites:

›› Remote office security implementations must be consistent and interoperable with the HQ setup;

›› User authentication should be uniform. For example, if your main site uses EAP-TTLS, remote offices should do the same;

Impact Assessment: Remote-Office WLANs

	Benefit	Risk
IT organization	Extending the main WLAN to branches allows for consistent, centrally managed WLAN profiles, including security controls, authentication schemes, and monitoring.	The lack of on-site IT resources can make troubleshooting challenging, particularly when a remote WLAN device is malfunctioning.
Business organization	A standard corporate wireless setup on mobile devices makes accessing resources easy as users move from office to office.	Any risks associated with the implemented WLAN security control framework will be exacerbated by extending it to small sites.
Business competitiveness	Convenient Internet access for guests and secure access to corporate resources for employees promote a professional image and increase productivity, even at small remote sites.	It's arguably more risky from a competitiveness standpoint not to offer Wi-Fi to employees and guests.

Bottom Line

Deploying WLANs to remote locations adds some complexity and sophistication to the design, implementation, and ongoing management of a wireless network. However, the cost of not building a system that can scale to all sites may ultimately be compromised security if remote users implement their own APs.

›› Data encryption should be consistent as well. If the WLAN security policy states that CCMP/AES encryption is required, then WPA2 should be used at all locations;

›› Wireless intrusion prevention systems should be used to enforce a "no rogue AP" policy. This security control prevents purposeful—and clueless—network attachment of unsanctioned access points. A well-designed WIPS can disable rogue APs by shutting down their copper LAN switch ports, temporarily tar-pitting their RF resources, and helping IT locate the device to facilitate physical removal;

›› Web-portal-based guest access should be made available to accommodate visitors. Captive portal functions like this should always use secure authentication protocols, such as HTTPS.

›› Finally, role-based access control (RBAC) bandwidth throttling,on a per-group or -user basis, should be considered when available.

KEEP IT SIMPLE? NOPE

So why not deploy the same lightweight APs at all branch offices that are used in HQ, and manage them with central controllers? To answer this question, we must understand at a high level how the centralized WLAN model works, paying particular attention to traffic and data flows.

When an access point first powers up, it must obtain an IP address and information about the controller with which it needs to communicate. This IP address must be reachable, which means that the remote office must be able to route back to the controller. Once the remote AP has this information, it creates a tunnel—using GRE (Aruba), LWAPP (Cisco), or another format—back to the controller to obtain updated WLAN configuration information, firmware, and settings.

It is a feasible setup, but because the hardware has not been optimized to communicate over a WAN link, inefficiencies and failure points greatly diminish the appeal of this option. Specifically:

›› No controller = no access: If connectivity to the controller is lost, as from a WAN failure, all WLAN users may be immediately disconnected. The split-MAC architecture used by the basic AP dictates that virtually all WLAN traffic must be encapsulated in a GRE, LWAPP, or other tunnel packet for transport back to the controller for processing. When the connection to the controller fails, the AP cannot by itself process the WLAN information and will begin searching for a backup controller. When this occurs, wireless clients are dropped and can access neither remote nor local resources. Note that some WLAN vendors, including Colubris and Trapeze Networks, do build resiliency into their basic APs by leveraging a design model called "distributed forwarding" to push more switching intelligence back out to the AP. This approach has its own pros and cons, however, additional discussion of this architecture falls outside the scope of this article.

›› Poor bandwidth conservation: Most basic lightweight access points are configured to tunnel all traffic back to their target controllers. Thus, traffic destined for any device on the network—even those at the same site as the AP—must first traverse the WAN to get to its destination. For traffic originating and terminating at the remote site there is no local switching option, so the data frame must cross the WAN twice just to exchange one data frame. Stated simply, basic lightweight APs are not smart enough to selectively forward traffic based on source and destination information.

›› Controller NAT issues: Security issues aside for the moment, basic lightweight APs cannot properly communicate with a NAT-enabled controller, even if static network address translation is being used. The problem is that a controller's IP address is learned as part of the connection process. If NAT is active, the controller has both a public (global) and a private (local) address. There's no way for the controller to provide both addresses to the access point because the controller knows only its local address, not its global NAT address. This means that the remote site must be tied directly back to the HQ network via a VPN or point-to-point link that eliminates the need for NAT.

Clearly, while the basic AP architecture model can be made to work, it's not an elegant solution to the remote office WLAN problem. The basic AP was designed and optimized for LAN deployments that could leverage fast local connections back to the controller. If you're lucky, branches are connected back to corporate via MPLS or dedicated lines, but many make do with the Internet and a VPN for remote employees who require privileged access to critical corporate systems and sensitive data.

A better route is either a scaled-down controller appliance at the branch site or a product, like Aruba Remote Access Points, or RAP, or Cisco Hybrid Remote

Lighten Your WAN Load

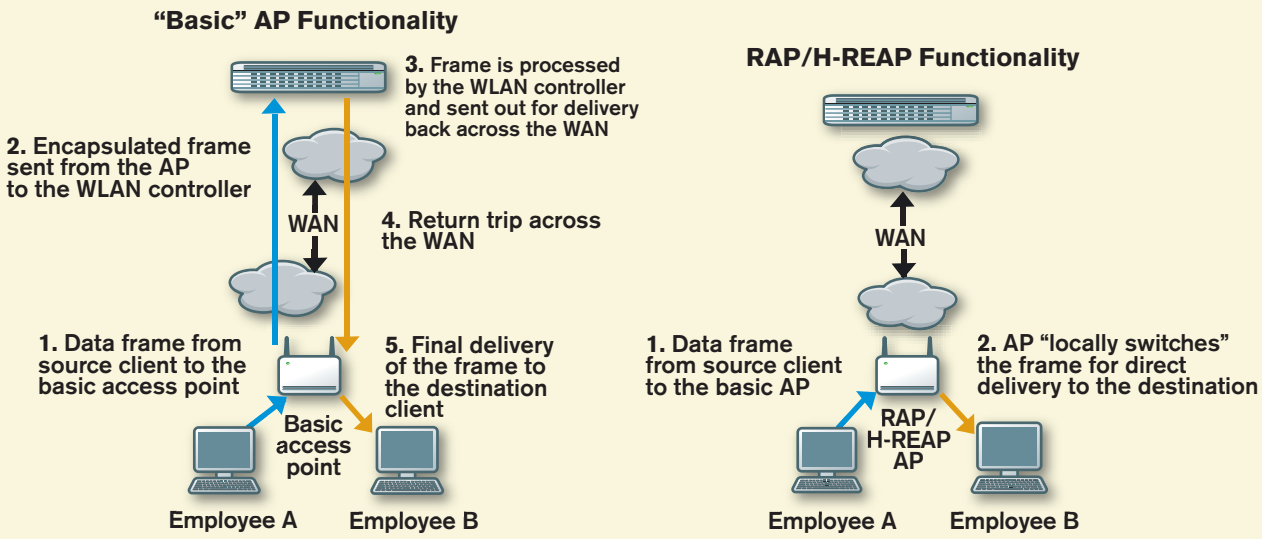

The local switching model allows for a single point of control, while not forcing traffic to traverse the WAN. By locally switching traffic, there's substantial improvement in performance and a decrease in latency.

Edge Access Points, or H-REAP, that extends central WLAN management to remote sites. Both the controller appliance and remote AP options provide WLAN deployment consistency, direct remote troubleshooting capabilities, wireless intrusion prevention system (WIPS) capabilities to enhance RF visibility, and the knowledge that approved corporate WLAN security standards are being extended to all locations within the organization.

Each architectural variation has pros and cons, which we'll discuss. What's consistent is that they extend your security policies and control standards throughout the organization.

A BETTER WAY

Two smart scenarios for extending WLAN access: Deploy a lightweight AP configured with RAP or H-REAP, or place a lightweight controller and APs on site. When a basic AP and its controller are co-located within a campus' high speed LAN, it's acceptable—often advantageous, even—for the AP to forward virtually every 802.11 frame to the controller for evaluation. In a WAN environment, with bandwidth constraints and latency issues, this approach is not usually practical. Aruba's RAP and Cisco's H-REAP technologies, choosing two as examples, reintroduce intelligence to the traditionally "dumb" lightweight AP model and are designed to reduce the required round trips to the controller by selectively switching packets locally if they're destined for local devices.

Setting up a RAP or H-REAP AP requires only configuring the device to provide LAN awareness so local switching can be performed when possible. For example, if a wireless user at a remote office simply needs to print a document to a local printer, this awareness would enable the AP to send the job within the LAN instead of over the WAN to a controller and back again.

Picture a situation where Mary and Bob are in the same office, and Mary's PC needs to connect to Bob's PC. If the local AP is joined to the HQ controller via the WAN in a conventional manner (that is, with no local switching intelligence) the volume of WAN transactions is enormous, as illustrated by the "Basic AP Functionality" shown in diagram, above. By using RAP or H-REAP functionality and adding some additional intelligence into our switching decisions, data flows can be optimized. Note that these flows don't take into account the return traffic from Bob's system back to Mary's, which only increases network resource usage.

As with the other design options, there are a few considerations to keep in mind with RAP or H-REAP:

›› Smart bandwidth use: As shown in our diagram, RAP/H-REAP-enabled APs are intelligent about how they handle traffic, providing for better performance and minimized WAN bandwidth use.

›› Poor roaming: RAP/H-REAP APs are not designed to provide fast Basic Service Set transitions be-

Bottom Line: Branch Office Wireless

With the growth of enterprise-class 802.11 WLANs an old design problem has reemerged: How to provision remote offices to support corporate IT standards in a cost-effective manner.

These sites don't just lack the IT infrastructure or support personnel found in larger headquarters offices. A dearth of local authentication servers, robust wired network components, application servers, and high-speed Internet connectivity can often frustrate IT administrators and impact the employee's user experience. If this problem becomes too pronounced, it can lead to remote office staff feeling marginalized and taking matters into their own hands, compromising security.

The solution to preventing an uncontrolled WLAN situation is to give employees what they want, but on terms compliant with corporate policy. When implemented correctly, employee satisfaction will be high, because they get the resilient, secure mobility they need with fewer problems and less downtime. In return, administrators get the security, functionality, and manageability they require in a cost effective fashion.

It seems that the only way to properly answer the ubiquitous question of "which option is best" is with the equally trite answer of "it depends." As described in our main article, there are well-engineered solutions that have been designed explicitly for the purpose of providing manageable, extensible, and cost-effective WLAN coverage to remote or branch offices. RAP/H-REAP offerings provide organizations with the means to outfit very small sites with a few access points that intelligently direct traffic to optimize WAN bandwidth. This option effectively eliminates the deployment of basic APs in most circumstances, especially given the nominal cost of implementing the RAP/H-REAP functionality.

Alternately, a lightweight controller provides larger branch offices with the ability to provision as many as six access points (typically) with the additional benefit of having a local controller to terminate AP tunnels. The cost of this setup is a little higher, but it provides additional robustness and configuration options.

IT must seize control of the wireless remote site situation or risk losing control of the network. A sound approach that satiates the need for mobile network access for remote office employees and their guests while maintaining corporate standards for authentication, encryption, high availability, and cost is within your grasp. Organizations that do not proactively manage remote site connectivity may well find that the problem resolves itself in a host of unsavory ways.

tween access points. This is of particular concern for a site containing multiple RAP/H-REAP devices that needs to support wireless VoIP or other latency-sensitive applications.

›› Authentication resiliency: While RAP and H-REAP do allow for some local authentication robustness, typically, if access to the controller is lost, new users are not able to authenticate using common 802.1X/EAP mechanisms such as PEAP or EAP-TLS. Depending on your authentication approach, some configuration can be done on the APs themselves to provide a certain degree of local authentication of users to help bridge outages where WAN resources might not be available. This would help local wireless users gain access to the WLAN and still be able to contact local systems and resources. Vendor support for local authentication and EAP types varies, so ensure compatibility prior to deployment.

›› Flexible remote offices: When an AP is in RAP/H-REAP mode, it has the ability to traverse NAT and provide remote users with access to corporate resources. This approach can be extended to road warriors, who could plug a preconfigured AP directly into a customer network and interface securely with corporate systems. In this regard, Aruba RAP devices have a bit of an advantage because they leverage standard IPSec VPN protocols to secure the connection between the AP and controller and to allow either wireless or wired access directly to the AP for tunneling back to the controller. The Aruba system has other benefits as well, such as the ability to traverse remote captive portals like those commonly found in hotels. In contrast, Cisco's H-REAP uses the LWAPP protocol, which authenticates WLAN components with public certificates and encrypts communications with AES. Though LWAPP has been criticized by academics to some degree because of its vulnerability to spoofing and DoS attacks, it is reasonably secure in private environments. It can be somewhat less enticing in a public, shared setting, however.

Overall, the RAP/H-REAP architecture provides a compelling solution for extending enterprise WLANs to remote sites requiring from one to three APs. IT gains a range of well-balanced features and design options that facilitate consistency, extend centralized management, and provide reasonable WLAN visibility into remote offices. However, even with all of these advantages, there are certain applications in which there is no substitute for a local, dedicated controller.

LOCAL CONTROL

Where RAP and H-REAP are tailored for small sites needing three or fewer APs, a controller-based system is more viable for larger locations and those in need of high performance for applications such as voice over Wi-Fi. One design option, depending on the size of the

office being served, is the placement of a suitably sized WLAN controller at each site. When deciding on whether or not to place a controller at a remote office, consider the following issues:

›› Centralized management: As the number of controllers grows, so too does the overhead required to maintain consistency and properly monitor the various systems. Vendors typically offer customized wireless network management systems that provide a unified way to create WLAN profile templates, manage multiple controller settings, and centralize alerts across a geographically diverse enterprise. There is a cost associated with purchasing and configuring such software, so this should be factored into the overall equation.

›› Scalability: Controller capacity typically starts at five or six access points and can grow into supporting many hundreds or thousands of devices. The per-AP cost decreases precipitously with larger controllers as economies of scale begin to kick in. Many remote sites only need three or four APs, but may not wish to just automatically go with a SOHO controller. Consider such factors as WAN latency, QoS, and local traffic filtering to make an informed decision.

›› Quality of service: RAP/H-REAP systems are not designed to support the fast roaming required to optimize secure voice communications. If a remote location requires a high-level of performance, along with multi-AP roaming, the use of a local controller may be required to support even a nominal number of access points.

›› WAN latency: A slow WAN connection, or high congestion on the remote office LAN, can cause high-millisecond latency. If RAP/H-REAP devices have slow communication (> 100ms) back to their controllers, they can become temporarily "disconnected" and cut over to local-switching mode. As network issues are resolved, the devices reestablish connections with the controllers and switch their states again. This scenario may lead to thrashing, which in turn can cause user connectivity problems and impact the accessibility of the WLAN.

›› Local resiliency: A local controller makes network operations less dependent on the WAN connection. RAP/H-REAP devices are designed to be flexible, offer direct authentication options, and perform local switching to help compensate for a lost WAN connection; however, they are not as flexible as a controller located just off a local high-speed LAN. A local controller can facilitate direct firewalling, fast and secure roaming, EAP-offload, VPN termination, and a host of other features directly at the remote location.

For most sites, the correct architecture decision can be made by simply determining the number of required access points. Generally speaking, if the remote location is very small, requiring one or two APs, then the RAP/H-REAP approach is almost always the best way to go. If the office is a bit larger and requires five or more access points, then placing a controller on site is probably the right solution.

For those offices falling into the three-or-four-AP gray area, a more thorough review of the site's requirements and capabilities is required to select the best approach. Cost considerations may also play a significant role in the decision-making process. Cost will vary based on the number of deployed sites, the presence of existing infrastructure, the availability of technical support teams, and other business factors. When calculating costs remember that RAP/H-REAP APs still count towards the total number of APs that a controller can manage. So in addition to the cost of the APs, you'll need enough available AP capacity on your central controllers to manage all the remote access points. Fortunately the cost-per-AP decreases as the size of the controller increases.

Rick Dreger (CISSP, CWNE) and Grant Moerschel (CISSP, CWSP, CCSP) are co-founders of WaveGard, a vendor-neutral technology consulting company focused on providing outstanding solutions to help secure the IT enterprise. For more information please visit us at www.wavegard.com **or contact the authors directly at** info@wavegard.com.