



Considerations for Deployment of an Enterprise Wireless Network

By Joel Shore

Sponsored Exclusively By:



This Special Advertising Section Produced By:





Table of Contents

- 3** WLANs in today's world
- 4** History lesson
- 4** Fat, thin, or fit
- 5** Routing advantage
- 6** Layer by layer
- 6** Routing is the answer
- 7** Scalability and availability
- 7** Security considerations
- 8** Virtualization and service levels
- 8** Telephony as productivity booster
- 9** RF management
- 9** Conclusion

Wireless LANs are fast becoming an essential tool for businesses, large and small. Once seen as lacking the features, bandwidth, and security needed to withstand the demands of corporate use, their widespread deployment is now a fact of life, especially with the advent of wireless telephony applications. But as a corporation's wireless appetite grows, the need to provide seamless access for mobile users, and support for new applications, coupled with IT's requirement for comprehensive management and security, has become a challenge.

An enterprise WLAN is more than the willy-nilly deployment of access point devices. It requires a network design that is scalable — in numbers of users and geographically — and must provide seamless connectivity as mobile users change location, from office to office and even city to city. It must be ready for imaginative new uses, such as wireless telephony (VoIP or wireless VoIP,) and high-speed wireless data services. And from the network administrator's view, the WLAN must appear as a unified continuum managed as a single entity, not as a collection of discrete access points unaware of each other's existence. Finally, the administrator must be able to create logical groups of users — virtual groups — whose access rights and WLAN capabilities are easily defined and managed.

To achieve these goals, traditional switch-based network architecture, a self-limiting technology, is giving way to a new approach, one conceived specifically with large-scale WLANs in mind. That approach is router-based WLANs that operate at Layer 3 of the Open Systems Interconnection (OSI) reference model.

This paper contrasts switch- and router-based WLAN methods and examines the business, technical and financial advantages that can be achieved by moving to newer WLAN technologies.

About the Author

Joel Shore is Editorial Director of Reference Guide, a professional-services firm that provides market analysis and custom editorial services to technology vendors, and solutions strategies to businesses. His opinion column in Accela Communications' *Solutions Integrator* newsletter appears biweekly.



Previously, Joel was Editor-In-Chief of ITworld.com and host of the television webcast *IDC Live*. He was the cofounder and longtime director of the *Computer Reseller News* Test Center, the award-winning publisher of product reviews. Prior to CRN, Joel was a senior systems analyst at a major bank and director of point-of-sale systems for two national retail chains. A frequent speaker at industry events, he is the author of more than a dozen books on personal computing, and has appeared on numerous television and radio programs.

WLANs in today's world

Corporations have, at last, discovered that WLAN technology is robust and increasingly secure, with a widening array of solutions that target large companies. These products, designed to run on the corporate network and support hundreds of access points and thousands of wireless users, bear little resemblance to the basic home-networking products sold in the major consumer-electronics retailers.

Major market research firms confirm the growing presence of WLANs in corporate America:

- In an August 2004 report, market research firm Dell'Oro Group projects that sales of WLAN products in the U.S. will grow 20% in 2004 to \$2.1 billion for the year from 2003. The fastest growth is occurring in the enterprise, not consumer segments: Dell'Oro predicts that shipments of enterprise-class wireless access-point devices will increase by 75% in 2004, and continue to grow at an average annual rate of 47% through 2008.
- Research firm IDC predicts that by year-end 2004, more than 75% of all enterprises will have deployed wireless networks.
- According to CIO magazine, 68% of CIOs consider wireless networking to be important

Indeed, there is now a race underway among corporations and other institutions to deploy enterprise-class WLAN capability. To meet the burgeoning demand, several vendors offer solutions that target the high-capacity enterprise WLAN. Among them are Airespace in San Jose (www.airespace.com); Aruba Wireless Networks in Sunnyvale, Calif. (www.arubanetworks.com); Chantry Networks in Waltham, Mass. (www.chantrynetworks.com); Cisco in San Jose (www.cisco.com); MobileAccess in Vienna, Va., (www.mobileaccess.com); and Trapeze Networks in Pleasanton, Calif. (www.trapezenetworks.com).

No longer merely a convenience; an increasing number of businesses and institutions would simply cease to function without WLANs. For these organizations, availability, scalability to thousands of users, and the ability for users to roam over great distances are essential. So too are robust security and the capability to create virtual user groups regardless of physical location or movement from one place to another. Consider several forward-thinking WLAN deployments:

New Mexico's Bernalillo County Metropolitan Court, which prohibits the use of cell phones, instead uses an enterprise-class WLAN throughout its 10-story facility. The court's security force, whose safety could be endangered by the 900MHz cellular phones that frequently drop calls in the concrete-and-steel building, now use WLAN handsets that are assured of a dropout-free signal. Even lawyers, who must stay in contact with their offices,

are issued a WLAN phone handset each morning when they arrive. Prospective jurors can use their notebook computers over this same WLAN, no longer losing a day of productivity while waiting in the jury pool.

Passengers aboard the Washington State Ferries system enjoy full Internet access, thanks to a WLAN implementation that supports seamless roaming throughout the 10 routes served by the ferries 29 vessels. Even seven miles from shore, far beyond the reach of any cellular telephone tower, the ferries 75,000 daily passengers access their e-mail and make VoIP telephone calls using the ferries wireless 802.11 network.

Cornell University's wireless network serves thousands of students and faculty, and nearly 400 buildings and outdoor stadiums spread across more than 750 bucolic acres. By replacing its push-to-talk, walkie-talkie cellular phones with 802.11 wireless handsets, the university's security force is eliminating hundreds of costly, monthly, cell phone bills and can be assured a guaranteed signal, even from within Cornell's many historic buildings with two-foot-thick stone walls that are impenetrable by cell phone signals.

Other environments are turning to WLAN technology to overcome logistical obstacles, reduce costs, and improve productivity. They include hospitals that equip medical staff with 802.11-enabled tablet PCs, hotels that offer Internet access to guests while restricting them from the hotel chain's intranet, and multinational corporations whose executives enjoy the same network access performance whether they are in their office or halfway around the world.



Chantry
NETWORKS

**We Made A Promising
Technology Live Up To
Its Promise**

Experience Everything WLANs
Can Do — With Chantry Networks

Delivering on the promise of true wireless mobility

History lesson

A WLAN is essentially a traditional wired network that eliminates the last length of cable to the desktop. Wireless networks don't typically exist on their own; they are often implemented as adjuncts to an existing corporate wired network. That eventually may change with the latest improvements to WLAN technology and security.

The concept of a WLAN is hardly new. As far back as the early 1990s, offerings less-sophisticated and far slower than today's solutions were available from several companies. One of these suppliers, Telesystems SLW, a Canadian firm, was granted U.S. patent 5,046,066 on Sept. 3, 1991 for a "wireless local area network" based on spread-spectrum radio technology. Developed for military use in World War II, spread-spectrum, as its name suggests, carves up a signal into pieces that are each transmitted over a spectrum of different radio frequencies. The technique simultaneously enhanced security and reliability. Thanks to built-in redundancies, the original signal could be re-assembled at the receiving end even if some of the transmitted pieces were garbled or missing.

Modern WLAN technology is built on the very same spread-spectrum principles. Wireless Fidelity (Wi-Fi) WLAN solutions are based on the 802.11 family of standards, developed under the auspices of the Institute of Electrical and Electronics Engineers (IEEE). The best-known standard, 802.11b, ratified in September 1999, operates in the 2.4-GHz frequency spectrum and supports data-transmission speeds up to 11M bit/sec. 802.11g, which is backward-compatible with 802.11b, is rapidly supplanting 802.11b. Ratified in June 2003, 802.11g operates in the same frequency range, but supports speeds up to 54M bit/sec. A third protocol, 802.11a, supports speeds to 54M bit/sec, but operates in the 5.4-GHz spectrum and is compatible with neither 802.11b nor 802.11g.

In January 2004, IEEE announced a plan to create a new standard, 802.11n, for wireless WANs, operating up to five times faster than 802.11g and able to operate at greater distances. Completion of this work is slated for late 2006, with ratification likely sometime in 2007. To protect legacy wireless investments, it is likely the 802.11n standard will be backward-compatible.

Access points represent the outer edge of the corporate network infrastructure, bridging the wired network to wireless computers, PDAs, phones and other devices. Mounted on walls or ceilings (and often hidden above the tiles of a suspended ceiling), the newest access points can draw their operating power directly through the Ethernet cable, eliminating the expense and labor of installing a 120-volt a/c. outlet at each device. Known as Power over Ethernet (PoE), "injector" devices, usually located in wiring closets, insert, or inject the appropriate direct-current voltage onto the network cable connected to each access point.

Access points are now almost universally compliant with the 802.11a, b, and g, standards, and mark the end of the wired network. For this reason, to understand the various considerations of deploying an enterprise-worthy WLAN, it makes sense to examine access points first and work back toward the center of the network.

Fat, thin, or fit

Access points can be built to include a widely varying array of capabilities. Traditionally, these have been categorized as "fat" and "thin." While both have their proper place in the universe of wireless networking, a third type, known as "fit," is gaining acceptance.

Fat access points are so-nicknamed because they perform a broad series of functions. Naturally, they include 802.11 radio capability, but they also handle encryption, user authentication, mobility and management. These devices also handle many fundamental network functions, including routing, IP tunneling, and network-address translation. Some access points include VPN endpoint capability. Fat access points — and there are likely to be hundreds in a corporate headquarters, hospital or multi-structure campus setting — are each connected to a network switch port in the nearest wiring closet.

Each access point is an autonomous device, handling its chores unaware that other access points are performing identical networking chores. And owing to their autonomous nature, each device requires separate management, adding significant complexity to the IT network administrator's job. There must be a better way.

Enter the so-called "thin" access point. Endowed with fewer capabilities than their "fat" siblings, these devices are little more than radio transceivers. All of the aforementioned networking tasks are pulled out, consolidated in specialized controllers, or in WLAN switches, which are installed in each wiring closet.

The advantages of this architecture are clear. Management and administration of the overall wireless infrastructure is simplified. Thin access points are less expensive, and let economies of scale accrue as more are deployed. And there's no need to assign and configure an IP address every time an additional thin access point is installed. This is more true in that thin devices do not have to be identified to back-end servers as they do in the fat case.

Alas, thin access-point architecture is not perfect, either. Certainly, consolidating access-point intelligence into a WLAN switch is an improvement: It facilitates central management of the access points connected directly to it. But this solution is handcuffed by the segmentation and manageability challenges intrinsic to switch-based technology. Suitable for a small business, this topology does not scale well to large corporations.

Switches are designed for deployment in distinct subnets, individual neighborhoods that together form the overall network. Once an enormous technological leap that provided a means for expanding networks as employee ranks swelled, subnets are essentially physical walls that, among other things, separate employees in different locations even though they might work in the same corporate department. Also, the switch constitutes a single point of failure, potentially jeopardizing availability: Lose the switch and all of the access points connected to it cannot exist on their own. Although fault-tolerant configurations are available, they add significant complexity and cost.

Thin access points each require a direct connection to the switch in order to operate. And because the switch, by definition, is an isolated subnet, the thin access points connected to it are merely an appendage hanging off the network, not a full-fledged member. As the number of switches increases, the task of network management grows rapidly in complexity. In the end, while adding many easily managed thin access points would seem to reduce overall management overhead, this is not the case. Instead of managing a large number of access points, the administrator winds up managing a large number of isolated switches.

Although consolidated WLAN switches solve some shortcomings associated with fat access points, the requirement for a direct connection to their access points necessitates deploying them in wiring closets throughout the enterprise. While this can be done, this has done little to simplify administration and maintenance.

Few corporations, however, are willing to adopt such a deployment strategy. To achieve demonstrable economies of scale and ensure physical security, corporate IT departments, not surprisingly, prefer to install these controllers — regardless of the quantity required or the location of access points — in one location, ideally within the confines of the corporate data center. Doing so reduces management overhead and lets network administrators replace a faulty unit without the need to track down its location and travel to it.

To do this, controllers and access points must have the ability to find and communicate with each other over the network itself instead of via direct connections. Switches, no matter how specialized they are, can't do this. But routers can.

Routing advantage

The limitations of switch-based wireless networking are clear. So what is the proper solution for a forward-thinking company? It's the router-based WLAN. The reason is sheer simplicity: Switches connect isolated groups of devices, routers connect groups of networks.

Instead of employing switch-based WLAN controllers, routers leverage the seamlessness of the Internet itself, or at least the Internet's underlying and monumentally scalable IP . Switches, because of their inherent architectural limitations, impose geographical segmentation that adds complexity to LAN management while simultaneously inhibiting flexibility and growth.

With a router-based WLAN architecture, it is possible to:

- Manage access points centrally while breaking through the geographical and distance limitations of switched architectures. With a routed solution, deploying access points anywhere in the enterprise network becomes possible, down the hall, throughout a campus, miles away in a branch office, or around the world.
- Deploy plug-and-play wireless routers and access points at any location in the network, across multiple switch or router hops or even across WAN links. Fit access points need not be connected directly to the controller, a scheme necessary with thin access points. Depending on the vendor, a plug-and-play access point is just that: Connect it to the network and the WLAN router instantly recognizes it.
- Scale a WLAN from a single department to a worldwide company. No less scalable than the Internet itself, IP routing unshackles the WLAN from the scalability and manageability restrictions associated with a switch-based architecture.
- Define virtual WLAN groups, leading to a high degree of flexibility. Members of the same corporate department defined as a virtual group all operate under the access-rights and



**The leading provider of
secure integrated mobility
management solutions
for WLANs**

Delivering on the promise of true wireless mobility

service-level policies even though they are located in offices hundreds of miles away from each other. Explored further in depth in this report, virtualization provides a platform that lets the IT department more intelligently manage policies and administrative domains.

- Assign different operating priorities by application or task. Applications, such as voice over WLAN (VoWLAN) phones, require a high priority, or QoS, to ensure a smooth, unbroken phone conversation. Conversely, a data-transfer task that can withstand delays of a few milliseconds can be assigned a lower QoS.

These clear advantages of router-based wireless networking arises from the standards all networking devices employ, the seven-layer OSI reference model. Switches operate at Level 2, the Data Link Layer. Routers, by contrast, operate at Layer 3, the Network Layer. To understand the technical advantages of implementing router-based networking, it's useful to take a step back and explore overall network architecture.

Layer by layer

In the late 1970s, the International Organization for Standardization (ISO) and the International Telegraph and Telephone Consultative Committee (CCITT) began parallel but independent projects to define a standard for network architecture. In 1983, the two independent bodies merged their work, leading to the Basic Reference Model for Open Systems Interconnection, commonly called the OSI Reference Model. Published in 1984 simultaneously by the ISO as standard ISO 7498 and by the CCITT as standard X.200, the model was intended to serve as a foundation for establishing international networking protocols. That never fully happened because of the sudden rise of an unrelated protocol suite, TCP/IP. Nevertheless, the OSI Model has become the universal standard for network infrastructure.

The seven layers of the OSI model, in ascending order, are physical, data link, network, transport, session, presentation and application. Each layer communicates with the layer or layers above or below it.

Layer 2, the data link layer, handles point-to-point addressing through the transit of data across a physical link. Data packets are encoded and decoded into bits. Divided into two sublayers, the media access control layer controls how a computer gains access to data and permission to transmit it while the logical link control layer controls frame synchronization, flow control and error checking. Devices operating at Layer 2 sort physical addresses, but are not particularly smart.

Layer 3, the network layer, is responsible for end-to-end addressing. It provides switching and routing capabilities, creating logical paths for transmitting data from one network node to another. Layer 3 also handles addressing, internetworking, error processing, congestion control and packet sequencing. Devices operating at Layer 3 are smarter than Layer 2 devices; they incorporate routing functions that calculate the optimal route for sending packets to their destination.

Because they operate at Layer 3 and not at Level 2, which handles component building blocks, routers are powerful devices.

Routing is the answer

Plugging access points into a corporation's network switches is the longstanding traditional method for implementing wireless networking capability. Attach the access points, configure a few parameters, and you have a wireless infrastructure, albeit a rudimentary one, ready to grant access to computers and PDAs outfitted with wireless transceivers.

Although this method certainly works, it falls down in terms of scalability. This simplest of architectures is likely to collapse under its own weight as the number of users swells, and as mobile employees roam greater distances, perhaps to branch offices located hundreds of miles away. Administration is certain to be complex, and security may be jeopardized.

To illustrate, consider a business housed in a three-story building. With few wireless users, a WLAN might be implemented as a single subnet that lets wireless computer users roam throughout the building. As the number of users increases, performance gradually degrades. To expand its wireless capacity, the company parallels the architecture of its wired network and installs a separate Layer 2 switch-based wireless subnet with several access points on each floor.

While this solution might be adequate for roaming computers where IP addresses can be assigned dynamically as the user moves from one zone to another, it just won't do for wireless telephony, perhaps the top justifier for implementing a WLAN in the first place.

With VoWLAN, a telephone number is bound to a specific IP address. But in a Layer 2 WLAN, as the user wanders from one subnet to another or even from one fat access-point zone to another within the same subnet, the handset's IP address changes and causes the call to be dropped. In other words, the physical phone device is assigned a new IP address, yet the call itself remains associated with the IP address in effect when the call was initiated — IP persistence. In this Layer 2 environment, ensuring that calls are not dropped requires nothing short of a network infrastructure overhaul.

These subnet-related shortcomings do not exist in a Layer 3 WLAN implementation. Because access points and controllers provide seamless roaming throughout the network, IP re-addressing simply doesn't need to occur, and calls therefore are never dropped. Not only is roaming VoWLAN now possible, the Layer 3 solution essentially future-proofs the network. Business solutions yet to come— perhaps wireless video- or data-streaming devices — are likely to require the same IP persistence. In addition to user woes imposed by the subnet solution, each of those subnets requires separate administration. In essence, in a Layer 2 architecture, the global WLAN cannot be managed as a single continuum as it is in a Layer 3 WLAN.

From this simple example, it's clear that a Layer 2 WLAN is inadequate. Scale this scenario to a large corporate headquarters, then to its far-flung branch offices, and the advantage of seamless roaming that's possible only with Layer 3 wireless networking becomes even more clear.

Scalability and availability

One fact is clear when it comes to corporate networks— they get bigger. More users, more devices, more bandwidth, more management. From the previous example, it's easy to see that a switched, or subnet Layer 2 architecture does not scale well.

Freed of the limitations imposed by a Layer 2 subnet architecture, a Layer 3 router-based WLAN solution offers considerable advantages as the enterprise adds more users and access points.

With plug-and-play controllers and access points deployed over Layer 3, no direct connection between them is required. Consequently, access points and WLAN controllers can be installed anywhere and find each other. This capability assures high availability. Should one controller go offline for any reason, access points immediately seek and find another on the network, regardless of its location. This isn't the case with the thin access points used in a Layer 2 WLAN implementation — because they require a direct controller connection, they cease to function if their host controller goes offline.

Because router-based Level 3 WLAN controllers need not be in close proximity to access points, they can be located anywhere, perhaps within the confines of the corporate data center. Once connected to the corporate IP backbone, they are online and available. The router-based WLAN scenario has an added benefit that is a direct benefit to the IT budget, eliminating the need to scatter IT staff throughout the enterprise to install controllers in geographically far-flung wiring closets.

Security considerations

It's only natural to think that the very openness of a wireless network would be at odds with the need to enforce stringent security. But the perception that wireless networking is not secure is obsolete. While it is true that the original 1997 802.11 WLAN specifications did not address security to a level that businesses could embrace, products based on newer standards provide ample security, so long as they are properly installed and policies enforced.

A comprehensive wireless security implementation must exist for both authentication and encryption.

Authentication is obvious: Just who is this person attempting to access the WLAN, and does he or she have the right to be here?

Encryption, although not as obvious at first, turns out to be equally critical. With the open nature of wireless networking, and wired networks, an intrepid soul bent on mischief can potentially intercept packets as they traverse the network. If the information was sent "in the clear," reassembling the original conversation, transaction, document, or other content becomes possible. Encrypting all activity ensures that content cannot be re-assembled. Encryption does not prevent packet theft, but without a way to unencrypt the data, the exercise is one of futility. In other words, it is encryption — even in the face of data theft — that provides privacy.

Although a variety of options are available for wireless encryption, two bear examination, Wi-Fi Protected Access (WPA) and 802.11i.

WPA was designed to replace Wired Equivalent Privacy, the original wireless encryption standard that proved to be not

Chantry
NETWORKS

first

- Routed Wireless LAN (WLAN)
- Carrier-ready, massively-scalable WLAN infrastructure
- Subscriber mobility management solution
- Seamless mobility VoWLAN solution

Delivering on the promise of true wireless mobility

robust enough for business use. WPA provides stronger data encryption and user authentication.

An interim measure, WPA has been superseded by the IEEE 802.11i WLAN security specification, ratified in June 2004. This standard includes Advanced Encryption Standard and provides the most robust encryption available outside the U.S. military. Its key caching feature allows quick re-attachment to servers. It also offers pre-authentication, essential for roaming across several network access points in applications such as encrypted VoWLAN.

Virtualization and service levels

With all the advantages that router-based WLANs offer, it might be natural to think of them as too accessible, too open, and too universal. After working to implement a WLAN that offers seamless mobility and management, the need to control access and service levels is essential. Advanced WLAN architectures solve this issue through virtualization.

Consider the following scenarios: A business wants to allow visitors to use their WLAN-capable laptop computers to access the Internet but not the company's internal intranet. A hospital needs to ensure that medical professionals but not administrative staffers have access to patients' medical records. And that hospital also plans to allow patients to use their own computers for Internet access. Another company plans to roll out wireless VoIP telephony and needs to guarantee a high QoS.

Virtualization is a means for implementing multiple policies and access privileges for various groups of users and devices.

Although partitioning groups through virtualization would seem to be at odds with a WLAN's physical seamlessness, they actually work together, providing capabilities that simply cannot be built with Layer 2 switch-based technology. For example, a multi-building medical complex spread over several acres should allow a physician the same capabilities and access rights no matter where he is. And the same goes for the aforementioned admissions staffer: same WLAN, different constituencies.

Virtualization supports multiple separate virtual subnets that operate simultaneously on the same physical Wi-Fi infrastructure, regardless of geographical location, without requiring complex VLAN configurations. With this capability, a seamless WLAN can be partitioned in a secure manner, providing differing wireless QoS levels to staff, management, business partners, visitors and even grouped classes of devices, such as voice handsets.

Telephony as productivity booster

VoWLAN is the convergence of two previously discrete technologies, VoIP and WLAN. Although VoWLAN (sometimes called

voice over Wi-Fi, or VoFi) is relatively new, it is poised for explosive growth. Cornell University and New Mexico's Bernalillo County Metropolitan Court, both cited earlier in this document, are using the technology today, enhancing staff safety and reducing expenses dramatically.

From a business viewpoint, the advantages of VoWLAN are clear. Wireless IP telephone handsets function in locations where cellular phones are often useless, inside steel and concrete buildings, aboard ships, and in geographic areas beyond the reach of cellular towers. With a WLAN deployed, even the most remote warehouse can be guaranteed excellent VoWLAN operation. Staff security is maintained and productivity enhanced.

For companies with a WLAN infrastructure already deployed, the incremental cost of implementing VoWLAN and purchasing VoWLAN handsets is minimal. By eliminating cellular or push-to-talk phones and their costly monthly charges, and replacing them with VoWLAN handsets, significant savings can be achieved.

Despite these clear business benefits, acceptance by users is perhaps the last hurdle to widespread acceptance. Implementing voice over a WLAN raises significant issues regarding QoS levels: whereas a lag time, or latency, of several seconds in accessing one's e-mail via a WLAN is acceptable and almost certainly imperceptible, that is not at all the case with voice. Any latency of more than a few milliseconds is easily noticed and quickly becomes annoying. The goal, of course, is zero latency when roaming, meaning that the WLAN sounds and feels like the traditional wired telephone system, service level that the cellular phone system does not meet.

Virtualization is the technique for assuring an adequate service level, a technique that cannot be achieved in a switched WLAN environment.

Simply put, the data packets that constitute VoWLAN in a routed WLAN implementation can be assigned a higher priority, placing them ahead of other traffic. Voice users, defined as a virtual group, benefit from this higher priority, assuring that their conversations are transported without the delay that might otherwise lead to poor sonic quality, dropouts or loss of the call. Additionally, differing levels of security can be specified by service level — payroll transactions might be subjected to highly robust encryption while applications requiring continuous bandwidth, such as VoWLAN, might have a level of encryption of security requiring less processing overhead.

As Wi-Fi standards evolve, so, too, do those that impact VoWLAN technology. In fact, a single all-encompassing standard for VoWLAN doesn't yet exist. Instead, various aspects of VoWLAN are based on different standards. QoS is covered by 802.11e, while 802.11f covers interaccess point communication,

and fast handoff of authenticated users. Another standard, 802.11e, prioritizes voice packets over data. Currently lacking is a standard specifying how calls are handled and how users should be re-authenticated as they roam from one access point to another. This and other VoWLAN-specific security issues are being addressed by a new task force, IEEE 802.r.

RF management

As the number of WLAN users in a given area increases, installation of additional access points becomes necessary. But as more access points are deployed, their overlapping radio frequency (RF) signals interfere with each other, leading to impaired performance of the network. Other factors also contribute to RF interference. To minimize these RF issues, two separate best practices, a site survey and ongoing RF management, are essential.

A site survey is a good first step for determining the optimal number and location of access points in an enterprise WLAN. Although valuable for the initial deployment, the survey represents only a snapshot in time; it is of no help in keeping access points optimally tuned as the number of users, arrangement of furniture, and other general environmental factors change over time. These continually shifting general factors include the movement of people and the installation of RF technologies, such as Bluetooth, microwave, possibly cordless keyboards, and even fluorescent lights.

Industry-specific sources of RF interference also can be largely overcome through dynamic RF management. In hospitals, these sources include the movement of electronic equipment or lead-lined curtains. In a warehouse, steel pallet racks, conveyor systems, fork-lift trucks and the movement of pallets of inventory are common aggravating factors.

To assure optimal performance, advanced enterprise WLANs employ dynamic RF management. These self-monitoring tools continually adjust and balance the transmission power and channel assignments of neighboring access points to ensure that adequate coverage and service levels are provided while minimizing overlapping or colliding signals.

Conclusion

With enterprise networks continuously growing, the classic solution of running more cable from more wiring closets to more desktops is expensive and counterproductive. To service more users quickly and economically, an increasing number of companies are turning to wireless networking as their salvation. In other scenarios, where cables cannot reach or where they are not permitted, the WLAN may be the only means to provide network services.

Traditional WLAN solutions, based on OSI Reference Model Layer 2 switching, are inappropriate for a large network. These switch-based solutions create isolated subnet pools that increase network management complexity and eliminate the possibility of enterprise-wide roaming and session persistence, which places an unnecessary burden on users. It also requires geographically dispersed controller devices and a direct connection to access points, resulting in lost service if the controller should go offline.

Router-based wireless networking, which operates at Layer 3 of the OSI reference model, offers plug-and-play ease of installation, and, with no direct connection from access point to controller required, permits aggregated controller installation within the corporate data center. This eases administration and avoids IT expenses.

Free of subnet barriers, the Level 3 WLAN becomes a contiguous environment, allowing users to roam at will without fear of losing connectivity or requiring re-authentication. New types of applications such as VoWLAN telephony are possible, and they provide higher levels of service at more locations and at a lower cost than cellular or push-to-talk wireless handsets. Router-based Level 3 wireless networking is the forward-looking solution.

© 2004 Network World, Inc. All rights reserved.

[To request reprints of this special report contact networkworld@reprintbuyer.com](mailto:networkworld@reprintbuyer.com)