WIRELESS

Implementing Wireless "Switches"

C. Michael Disabato

Michael Disabato covers wireless technologies and mobility for Burton Group. He has more than 25 years' experience in consulting, applications development, networking, security, wireless and technology assessment.

This new product category isn't what it may appear. Understanding WLAN switches is key to getting your 802.11 implementation right.

Now that Wi-Fi Protected Access (WPA) has resolved the known security vulnerabilities of wireless local area networks (WLANs), and the confusion over radio technology selection has been cleared up, many in the wireless industry have turned their attention to the architecture of WLANs and how they should be engineered and deployed. These thoughts have spawned several arguments within the WLAN segment.

The leading argument is about where the network intelligence should reside. Traditional systems put all the intelligence in the access point (Figure 1). This results in economies for smaller networks, but leads to significant redundancies in large networks.

The second argument concerning distributed intelligence is: If all the intelligence shouldn't go in the access point, how much should? Most distributed system vendors only put the radio and enough intelligence in the access point to allow it to communicate with the central controller. "Enough" may mean a MAC-layer communications processor or a full IP stack with encryption functions.

The third argument has been about WLAN designs. Using radio frequencies (RF) at Layer 1 means giving up the deterministic nature of switched, wired networks. Further, mobility changes the delivery model for network services and forces the addition of new services to accommodate address and session management, security and policy enforcement.

FIGURE 1 WLAN Systems

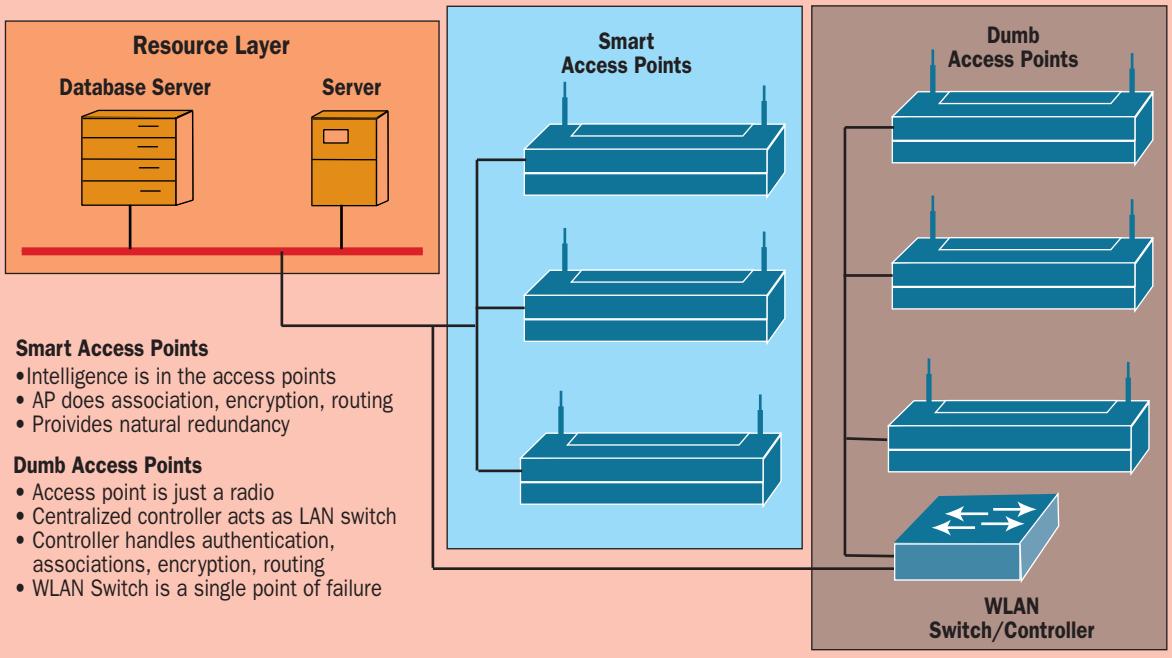

A site survey is typically out-of-date before it's complete

In September 2002, Symbol Technologies attempted to resolve these issues with the announcement of its Mobius WLAN switch. Far from clarifying the issues, Symbol fired the first shot in a new war for the wallets of network managers who are considering WLANs. Seven start-ups and five established network vendors now are vying for a share of the WLAN switch market. Clearly, too many vendors are chasing too few dollars, and consolidation in this market is inevitable.

While the marketing hype about these devices continues, the most misleading concept is that they are truly wireless switches. Nothing could be further from the truth. The RF link is still a shared medium, and barring a major shift in radio technology, will be so for the foreseeable future.

TABLE 1 WLAN Design/Deployment Approaches

Engineering approach:	"Normal" approach:
1.Determine a pilot group	1. Determine a pilot group
2.Perform a site survey	2. Select equipment
3.Determine coverage areas	3. Hang a few access points
4.Select equipment	4. See how well it works
5.Install access points	5. Hang a few more access points
6.Verify coverage areas	6.Tune network (maybe)
7.Install access points for general coverage	7. Respond to user complaints
8.Tune network as usage patterns emerge	8. Return to Step 5 until there are no more complaints
9.Add access points as coverage requirements increase	
10.Re-tune network	

Common Features

A quick glance at the current offerings on the market reveals that all these systems have adopted a fairly uniform set of features:

- Power over Ethernet (PoE) using the 802.3af standard.
- Support for 802.11a, 802.11b, and 802.11g.
- Support for Wi-Fi Protected Access (WPA) and 802.11i (AES).
- 802.1X for WLAN authentication.
- Support for the Lightweight Directory Access Protocol (LDAP), Microsoft Active Directory, Novell directories and RADIUS servers.
- Layer 2 radio connections.
- Connections for non-WLAN devices (printers, scanners, etc.).
- Rogue access point detection.
- Seamless roaming.
- Command line interface (CLI) and graphical user interface (GUI) for system management.
- Support for VLANs.
- Backbone links using either copper or fiber at 10/100/1000 Mbps .
- Policy-based security and QOS management, including "guest" access.
- A Secure Sockets Layer (SSL)-secured website available for guest registration or use in a hotspot environment.
- Load balancing and self-healing.

If the feature set is similar among vendors, what should the selection criteria be? Since the primary business case for these systems is reduced cost of ownership, the answer is: RF management.

Designing For Service Levels

With more users and time-sensitive applications such as voice moving to the WLAN, wireless network designs must support service levels, rather than just providing coverage. However, creating a WLAN that can support service levels implies creating a more deterministic environment than is typical for WLANs. The ability to meet service levels further implies that the network manager can control bandwidth, deliver defined quality of service and provide policy management that allows the wireless network's characteristics to be tailored in response to enterprise business needs.

Designing WLANs to meet performance objectives requires a higher density of access points than simply designing to provide coverage. These additional access points provide a seamless "RF umbrella" that delivers uniform service levels to each mobile device regardless of location, but they complicate the channel assignment problem. The increased number of access points used to form "microcells" and the attendant channel assignment process creates a level of complexity that cannot be handled by today's planning tools. WLAN switches automate the channel assignment process and adjust power levels in each access point, automating a tedious process.

Site Surveys

A site survey provides a one-time snapshot of the RF environment that is dependent on the range of the RF analyzer and typically becomes out-of-date before it is completed. Using current WLAN management tools, network managers have no way of knowing if any rogue access points are connected to the network unless they get lucky, another site survey is performed or users report performance problems. Also, the labor-intensive nature of site surveys fails to scale as the network grows and multiple physical locations must be supported.

In the past, products from AirMagnet, Network Instruments and Finisar have helped the network manager plan and manage a WLAN. However, these were primarily post-installation tools, because they worked by measuring the RF signal strength from each access point. Further, they are labor-intensive, as the tool needs to be moved around to find the limits of coverage. An accurate

set of as-built blueprints or a global positioning system (GPS) receiver is also necessary for these tools to be of value.

As a result of the limitations of site survey and planning tools, until now there have been two methods of designing and implementing WLANs, shown in Table 1.

While the engineering approach involves deterministic processes such as network planning, installation verification, deployment, management and optimization, it is time-consuming and labor-intensive. The network resulting from the "normal" approach, on the other hand, suffers from lack of management, unpredictable coverage areas and a varying quality of user experience. Clearly, a quantitative approach that does not involve large amounts of labor is necessary. This becomes even more evident when the network manager realizes this is a cyclic process that must be repeated as new applications are added, the network expands and more access points are required. This is where WLAN switches provide the most value.

Network Management

Network management (along with security) has always been one of the last areas addressed whenever a new technology has been introduced, and it's no different with WLANs. Early access points had just enough configuration management to get the network up and running, but as more access points were added, the ability to push out configuration data *en masse* was sorely lacking. In addition, there were no integrated RF management tools, which left WLAN managers scrambling to relate the information presented by third-party management tools to the installed network.

WLANs are organic entities that grow and change in response to their use, since mobility changes the social dynamics of the enterprise and the way people work. WLAN switch monitoring tools can verify service levels for each microcell based on requirements established during the design phase. They also can determine loading on an access point and balance the load on the system to reach required service levels and indicate where more access points may be needed.

Rogue Access Point Detection

Rogue access points are the "next big security risk" for WLANs. Detecting rogue access points is a continuous process that must indicate the physical location of the unauthorized device and the port to which it is connected, so that appropriate action may be taken to disable the device. Walking around (essentially doing a "war walk") with a WLAN analyzer, installing a network of passive detectors, or using a system that leverages the installed base of access points can detect rogue access points, but each suffers from limitations.

War-walking is labor-intensive and not all that accurate. First, the rogue device must be turned on and operating for it to be detected. Second, the person who installed the rogue access point can readily identify when a sweep is being performed and can simply turn off and hide the equipment until the sweep is over. Like site surveys, war-walking provides a snapshot of the network at a specific moment in time, and given the dynamic nature of WLANs, that snapshot is obsolete as soon as it is taken.

Using a system of passive monitors or using existing access points as passive monitors can be very expensive. Further, existing access points cannot function as passive monitors while they are being used for communicating with mobile devices. On a heavily used WLAN, this approach can suffer from the same problem as war-walking: inconsistent monitoring.

Antennas, Controllers And Switches

Throughout this article, the term "WLAN switches" has been used to represent any device that implements centralized management and control of a WLAN infrastructure. In actuality, there are two types of devices that provide this functionality. Switches provide ports to allow antennas and other wired equipment to be directly attached. Controllers do not have the capability to attach any wired devices.

Some of the vendors allow connection over a routed environment that permits installation of lightweight antennas in remote locations and controls them from a central site over the wide-area network (WAN). While this may seem attractive at first, the lack of intelligence in the lightweight antennas means that all traffic must be routed to the controller and then back to the remote site. Since lightweight antennas do not contain a MAC layer, all traffic processing (encryption, authentication, association control, etc.) must be done in the controller. This can have serious ramifications for network efficiency, because the maximum throughput will be limited to the bandwidth of the WAN connection (Figure 2).

Switches can provide capabilities for medium-sized installations that require some wired connections while providing the advantages of centralized management. Controllers will work well in environments where there is a mix of standard (smart) access points and lightweight access points.

Currently, each vendor has a proprietary method of controlling the antennas and effecting changes in the RF environment. In an effort to establish a standard protocol for switch-access point (lightweight or standard) communications, Airespace, Legra, and NTT DoCoMo successfully petitioned the Internet Engineering Task Force (IETF) to establish a working group that would be responsible for the development of the Lightweight Access Point Protocol (LWAPP). The LWAPP will drive the commoditization of the antenna market, which will be followed by a reduction in the cost of these devices. Though

There are really two types of devices in this category: Switches and controllers

Most enterprises will mix "smart" and "dumb" access points on their networks

FIGURE 2 Controllers vs. Switches

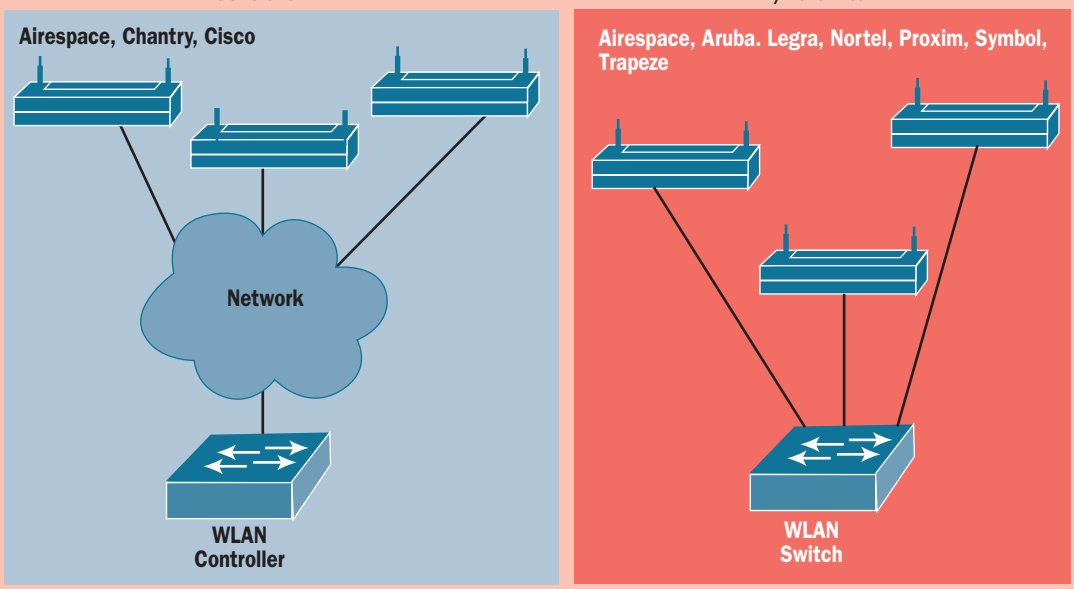

Cisco opposed the formation of this working group, the company will watch the developments there and may eventually sign onto the standard.

The trend will be to place standard access points in some areas and lightweight access points (antennas) in others, depending on business requirements and the network infrastructure. For example, regional or branch offices will use standard access points controlled from a central location to eliminate the need to backhaul traffic, which would be necessary if lightweight antennas were used with a centrally located switch.

Recommendations

In spite of the common set of features exhibited by these systems, implementation, especially RF management, varies with the product. When investigating a WLAN switch or controller for purchase, several questions need to be answered:

- Can configuration for the system be performed on a global basis, or must each access point be individually configured? Global policy management is far preferable, since large networks would require an excessive amount of labor to configure each access point individually.
- Is the management interface complete? The rush to market has caused quality problems with early device shipments. Several products that have started shipping still use early versions of their software. User interfaces are not always complete, and network operations may not be totally stable. Enterprises should wait until these systems reach a level of maturity that can support their reliability requirements, and systems should be tested in as close to a production environment as possible.
- How is RF monitoring performed? Can the access points handle traffic delivery and monitoring simultaneously, or is a separate overlay network required to perform RF monitoring? An overlay network will add expense and complexity to the deployment.
- Is RF monitoring performed continuously or at scheduled intervals? To rapidly respond to changes in the RF environment, RF monitoring should be a continuous process. If measurements are taken at discrete intervals, most transient events will be missed and uncorrected.
- RF management should include access point transmit power, load balancing, dynamic channel assignment, detection of coverage holes, interference detection and access point failures.
- Rogue access point detection must indicate the location of the rogue. The wired network connection (switch and port) also is helpful to assist in deactivating the rogue. Another method of neutralizing rogue access points is to transmit disassociate requests to all devices associated with the rogue device.
- Managing the RF link is critical to meeting service levels. A highly desirable feature of Wireless LAN switches is the ability to match access point location to physical building plans, thus creating a set of "as-built" network diagrams. WLAN switches also can convert design plans into configuration data for access points and other system elements and push that configuration data out to all devices automatically. This automation will be increasingly critical as enterprise WLANs increase in size.
- What is the cost of the complete system? Several vendors are charging on an *a la carte* basis, where the switch is priced separately from its support software. In at least one case, the cost of the software necessary to achieve the minimum

functionality is nearly half the cost of the switch itself.

Conclusions

That WLAN switches constitute a new product category would seem self-evident. However, if the products of each vendor are decomposed into their elements, it becomes apparent that these are Layer 2 or Layer 3 switches with software that performs functions specific to WLANs. Over time, the term "WLAN switch" will fade from use, as will some of the vendors who specialize in these systems. Many of the functions that are touted as unique are being built into network edge switches or can be handled by dedicated appliances.

Pricing for these systems is relatively high, compared to normal switches, assuming that customers will initially pay a premium for the improved functionality and lower operational expenses associated with WLAN switching. However, that will not last more than 18–24 months. During that time, expect to see the price of antennas and access points fall to near-commodity levels and some of the RF management capabilities embedded in standard access points.

As competition heats up and some of the current players exit the market, prices will fall. As a result, network managers must be careful in their selection process to ensure the vendor they pick will be around over the next few years□

Companies Mentioned In This Article

Airespace (www.airespace.com)

AirMagnet (www.airmagnet.com)

Cisco (www.cisco.com)

Finisar (www.finisar.com)

Legra (www.legra.com)

Network Instruments (www.networkinstruments.com)

NTT DoCoMo (www.nttdocomo.com)

Symbol Technologies (www.symbol.com)

Over the next 18–24 months, prices—and some vendors—will fall