

Wireless LAN Architectures

Distributed Access Points Vs. Centralized WLAN 'Switches'



Introduction

This white paper addresses the current industry debate about the relative merits of a centralized WLAN architecture using a WLAN switch versus the more common distributed 802.11 access point architecture.

Before we get started, I will provide a little background on LXE, for those of you possibly not familiar with us. LXE has been selling wireless LAN products for data collection applications for about 30 years. By ‘data collection applications’, we are referring to applications for inventory management, logistics, distribution, warehousing and the like. LXE was the first company to develop RF networking specifically for these applications. Our original RF networks used licensed UHF frequencies, and later moved into the unlicensed ISM band at 900 MHz. Those networks all used proprietary protocols, and LXE manufactured both the client radios and the networking infrastructure products. When the 802.11 standards were developed in the late 1990s, LXE recognized that we could not compete with some of the high volume vendors who began to enter the marketplace. Today, LXE does not manufacture any WLAN access points. LXE resells a limited line of industrial access points that are developed and manufactured by other vendors.

Definitions

To avoid confusion, I will start with a short definition of the terms ‘distributed access point architecture’ and ‘centralized WLAN switch architecture’.

Distributed Access Point Architecture

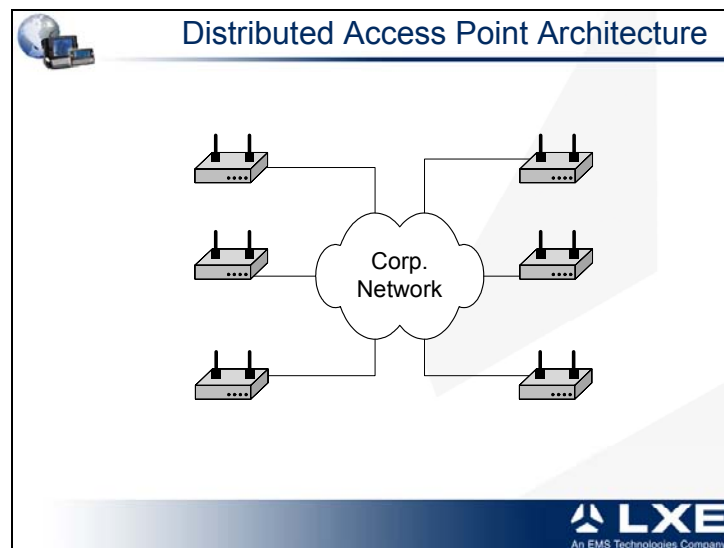
By ‘distributed access point architecture’, we are referring to the original design of 802.11 standard based access points. These access points put all of the WLAN functionality into a single piece of hardware.

Distributed access points implement the complete 802.11 specification. They provide the wireless to Ethernet layer 2 bridging function. They also implement the WLAN-specific security, including access control and encryption, as well as QoS functions. All “enterprise” class access points also provide some higher layer functionality, such as protocol filtering, address filtering, access control lists and configuration tools.

All enterprise class access point vendors also market tools to provide centralized management of their access points. While centralized management capability is probably a requirement for every organization with a modest to large wireless network, these centralized tools are not, strictly speaking, required to operate the network.

The great beauty of the distributed access point architecture is that the AP is the only network component required to provide wireless LAN capabilities. The distributed architecture was originally perceived as a great benefit relative to existing WLAN systems. Existing narrow-band (or UHF) and 900 MHz systems in the mid to late 1990s relied on a ‘network controller’ working in conjunction with several ‘radio frequency

units' to provide WLAN coverage. The new 802.11 architecture simplified all that by combining the functionality of the network controller and the radio frequency units into a single device. With the new architecture, all that was needed to create a WLAN was to plug an access point into an available switch port, install a client radio in your computer, and you were up and running. In most cases, it was not even necessary to configure either the AP or the client.

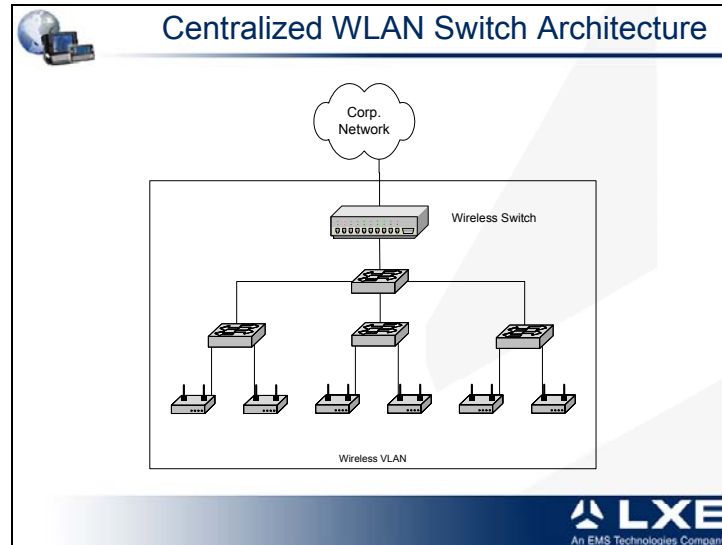


The distributed access point architecture is intended to be extremely simple. Wireless becomes an integral part of the overall network. Access points connect directly to distribution layer switch ports and wireless traffic is switched most efficiently to its destination.

Centralized WLAN Switch Architecture

The Centralized WLAN switch architecture requires two devices to provide wireless network. The access point is still an edge device and provides the wireless to Ethernet layer 2 bridging function. All other functionality can be moved off the access point and implemented in the WLAN switch. Each WLAN switch vendor makes different choices about how to divide the functionality between the switch and the access point. Some vendors choose to implement all functionality, other than the bridging function, in the WLAN switch. Others may determine that some functionality 'belongs' at the network edge and they put that functionality in the access points. Data encryption and access control are examples of functions that are sometimes left in the access point.

One essential characteristic of all Centralized WLAN Switch architectures is that all traffic either to or from the wireless network must pass through the switch. This allows the WLAN switch to be the traffic cop and to completely control the flow of wireless traffic.



A centralized WLAN switch architecture treats the wireless LAN as an overlay to the Ethernet network. The centralized WLAN switch requires a separate network dedicated to the wireless traffic. This separate network is probably implemented as a VLAN so that existing switching can be utilized. The separate network is required because all wireless traffic must pass through the wireless switch before it can be sent to its final destination.

There are several variations of the Centralized WLAN switch architecture as well. This diagram shows two levels of switching between the access points and the Wireless Switch. This architecture is required when the wireless switch provides a limited number of switch ports. Some wireless switches are provided with larger numbers of ports, and would not require the second layer of switching. In some cases, for small wireless networks, it may be possible to connect a sufficient number of access points directly to the wireless switch.

Some vendors also recommend that wireless LAN and Ethernet traffic share ports on the wireless switch. The shared port architecture is made to allow the vendor to claim higher throughput to the Wireless LAN. For example, take a case where the wireless switch has only two 100 Mbps Ethernet trunk ports. If each access point can generate 6 Mbps of traffic (assuming 802.11b access points), then a single 100 Mbps trunk port could accommodate 16 access points. To be able to claim more access points per wireless switch, both ports are dedicated to the wireless traffic. While using both wireless switch ports for the wireless traffic allows for more throughput, it is less secure than architecture diagramed here because it allows unauthenticated traffic to have access to the corporate backbone.

There are two observations that come immediately to mind upon considering the two definitions of distributed access point architecture and centralized WLAN switch architecture just described:

- First, the centralized WLAN switch architecture sounds a lot like the narrow-band and 900 MHz architecture mentioned in the previous slide. It is, and for good

reason. The vendors who first developed products using the centralized WLAN switch architecture were all vendors who had previous experience with narrow-band and 900 MHz systems. The centralized WLAN switch architecture is sometimes portrayed as a 'new' WLAN architecture, but it is not really. It is new only to the 802.11 standards-based marketplace. Vendors and users of previous generations of wireless LAN products are very familiar with this architecture.

- The second observation for anyone who has investigated the variety of WLAN networking products on the market today is that the line between a distributed architecture and a centralized one is often difficult to distinguish. Some WLAN switch vendors put as much functionality into the access points as do the distributed access point vendors. Other WLAN switch vendors require multiple switches to implement a network of any size. (A distributed, centralized WLAN switch architecture?) At the same time, distributed access point vendors are developing centralized management products that provide much of the functionality of some WLAN switches.

For the purposes of this presentation, if all WLAN traffic *must* pass through a second device in addition to the access point, we will call that a centralized WLAN switch. If the WLAN traffic passes directly to the Ethernet distribution layer from the access point, we will call that a distributed access point architecture.

Focus On Functionality

Most vendors who have been in the wireless data networking business for more than a few years have experience with both the distributed access point architecture and the centralized switch architecture. LXE has sold products using the basic centralized architecture in our early narrow-band and 900 MHz systems. The 802.11 systems we currently market are all based on the distributed access point architecture.

Our experience has taught us that a wireless LAN can be operated based on either architecture, and that the architecture itself is not the primary concern. The real issue is the functionality delivered by the specific vendors being considered. The architecture contributes to the manner in which certain functionality is implemented, but is never a gating factor. More important is the wireless LAN vendor's vision, understanding of the marketplace, and their ability to execute on their plan.

As you investigate wireless LAN products today, you will find almost all vendors claim to be able to solve all problems. Mostly, these are vision statements. Actual implementation can be quite different – and often very difficult to determine.

The big difference between the older centralized systems we used to sell and the newer version now being promoted for 802.11 is the level of functionality delivered. With the older systems, basic connectivity was the overriding issue. Connectivity is now taken for granted, and the attention of WLAN users has turned to issues such as throughput, coverage, security, network management and others. And, of course, price is always a consideration.

We will now briefly discuss these issues, focusing on issues related to the network architecture. Note that these issues are approached roughly in order of importance to a warehouse or distribution center application – except for price, which may be of greater importance in almost any market.

Wireless Network Coverage

There are several issues that impact the coverage of a wireless LAN. Network architecture is not one of them.

Radio receiver sensitivity can greatly impact the distance a user can roam from an access point and still maintain connectivity. Some wireless vendors publish receiver sensitivity specifications for their radios, others don't. In any case, the data should be available. If the vendor does not publish it, ask for it. If you do get it, the more negative the value, the better. For 802.11b radios, look for receiver sensitivity in the -90 to -95 dB range for one Mbps transmission. While receiver range can be affected by many factors, a general rule of thumb is that an eight dB improvement in receiver sensitivity approximately doubles the range of the radio.

Antenna selection has perhaps a greater impact on network coverage than does receiver sensitivity. Keep in mind that antennas are passive elements. That is, they do not add any power to the signal being transmitted. The role of the antenna is to shape the pattern of radiated RF energy. In general, the higher the gain of an antenna, the more focused the radiated RF pattern. High gain antennas can be used to provide coverage further from the access point, but only in narrowly defined areas. Lower gain antennas do not propagate the signal as far, but they cover a broader area. Different antennas are used in different locations depending on the desired coverage pattern.

Be aware of access points that include captive antennas. These access points often cannot accommodate antennas with different coverage patterns. This limits the flexibility of these access points, and they are not suitable for many applications. This warning applies to both centralized wireless switch systems and to distributed access point systems.

Site survey expertise normally comes from the vendor who installs the wireless network. This vendor is frequently not the manufacturer of the wireless LAN hardware. Look for a vendor who has experience with environments similar to yours. An installer with a lot of experience in an office environment will not be the best choice to install a wireless LAN in a warehouse or outdoors in a container yard.

Some wireless vendors now advertise tools for automating site surveys. Mostly, these tools are connected with the centralized wireless switch systems, but not always. Automated site surveys are no help at all in a warehouse. The tools are not capable of modeling a warehouse environment sufficiently well to depend on them for a site survey. The tools are designed for use in an office environment. Even then, they depend on

installing more access points than required. Automated site survey tools are most useful for verifying coverage after a network is in operation.

The environment has the biggest impact on network coverage, and we unfortunately do not often have any control over it. The nature of a warehouse application is that stocking levels will change, and the goods being stocked may also change. These two factors have a huge impact on RF propagation. This also makes doing a site survey in a new (and empty) facility very difficult. These are some of the reasons that you want your installer to have experience with your environment. A warehouse that houses mostly paper goods will require different access point placements than a similar warehouse filled mostly with glass or one filled with metal parts.

Container yards and warehouses that contain a lot of metal also suffer from a lot of multi-path – or RF signal echoes – that can greatly affect coverage. Some radios deal with multi-path better than others. Antenna diversity is also useful in high multi-path environments.

None of the issues mentioned above is dealt with any better by a centralized wireless switch architecture than by a distributed access point architecture. In both architectures there are products with a variety of design goals. Those products intended to meet the lowest price points or intended do-it-yourself installation tend to perform poorly in terms of network coverage. Products intended for the enterprise, and particularly those designed with industrial applications in mind provide better network coverage.

Almost every wireless LAN vendor has a low-end access point that features captive antennas and is advertised as easy to install and configure. These access points also often have less expensive radio components with poor receiver sensitivity. These products are meant to appeal to customers who focus on hardware acquisition costs. These products may not be the most economical in the long run. Their poor network coverage characteristics may require more units be deployed, thus driving up the cost to install and administer the network.

Wireless Security

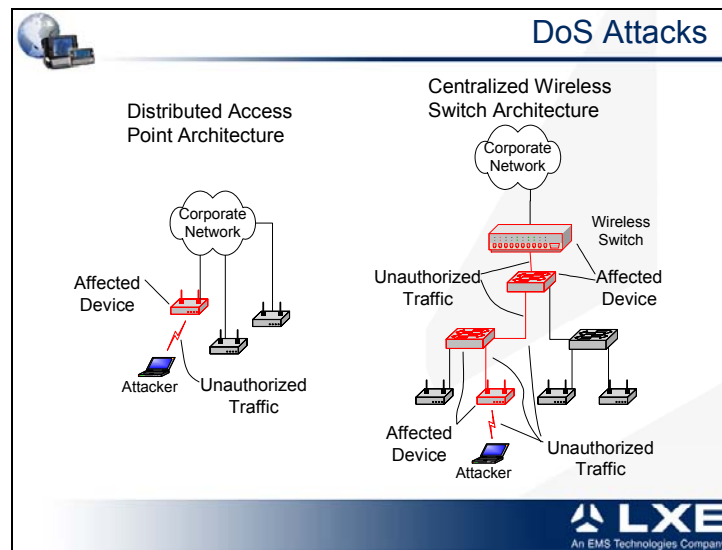
Wireless security has become the largest inhibitor to wireless LAN adoption today. When we talk about wireless security, we generally include three distinct components: User authentication, data privacy and data validation. Wireless LAN products are gradually moving to a security model based on the WiFi Protected Access (or WPA) model. In the WPA model, user authentication is accomplished using 802.1x. Data privacy is accomplished by encrypting traffic using the Temporal Key Integrity Protocol (or TKIP). Message Integrity Check (or MIC) is used to provide data validation. We expect the WPA model to satisfy most customer concerns over wireless network security.

Products using the distributed access point model have no choice but to implement all security functions in the access point – that is, at the network edge. Wireless switch vendors can choose where to implement security functions. Some wireless switch vendors implement all security in the access point. Some implement all security at the

switch. Others divide the security functions between the access point and the switch – for example by using the wireless switch for user authentication, but maintaining the data encryption and integrity check functions at the access point.

We believe that the authenticator function, at least, belongs at the edge. If an intruder is attempting to access the network, it is best to catch them as early as possible. If they can be stopped before reaching the Ethernet, so much the better. For example, let's look at how a denial of service attack might affect the network:

DoS Attacks



In a distributed access point network, a DoS attack will affect only the single access point under attack. Since the access point blocks all unauthorized traffic, other network components will be isolated from the attack. The attack may disable the one access point responding to the attacker, but other access points and network devices will be unaffected.

In the wireless switch network, where the switch performs that authentication function, the access point cannot block the unauthorized traffic. Instead, it must forward all traffic through the Ethernet network to the wireless switch. The wireless switch will recognize the traffic as being unauthorized, and will block it from the rest of the corporate network. But all devices from the access point through to the wireless switch will be required to carry the traffic, and will be affected by the attack.

Physical Security

One security issue frequently cited when arguing against the distributed access point architecture is that placing security at the network edge poses a physical security threat. The idea is that it is fairly easy to steal an access point. Since distributed access points contain encryption keys and other security settings, this creates a serious vulnerability. A thief could take an access point back to their lab, and take their time to extract several

pieces of important information. Not only could they retrieve the security settings of the AP, they could also discover MAC address of other network devices.

An intruder might also use as stolen access point as a rogue AP. By installing the stolen access point on their own network, but in close proximity to the target network, they might entice legitimate clients to connect to their network and so capture login information.

This security threat, while real, is a minor concern in industrial applications. Access points in industrial applications are typically located between 35 and 100 feet off the ground, in ceiling trusses or on light poles. These locations are very difficult to get to. Anyone trying would undoubtedly attract someone's attention.

Network Latency

Response time is a significant performance criteria for industrial wireless LAN users. While industrial users don't typically care about throughput, they do care that they don't have to wait for a response from the application host. Many times these users consider a response time of greater than one second to be unacceptable.

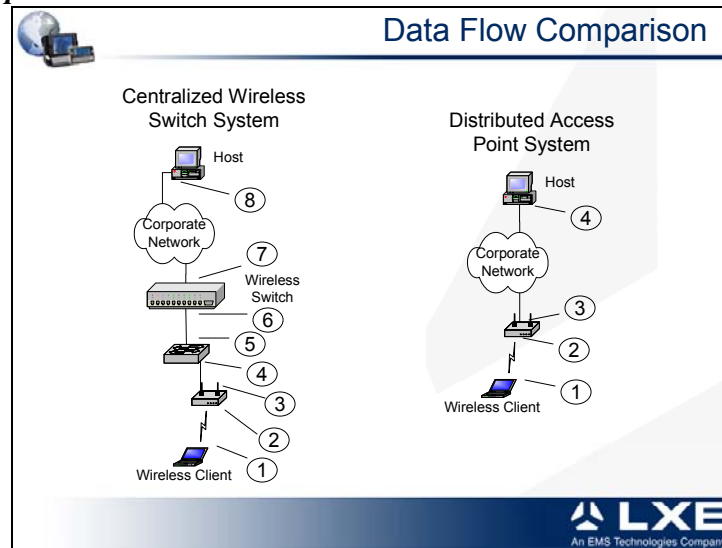
The distributed access point architecture provides the most efficient routing of wireless traffic. A data packet passes directly from the client, through the access point to the Ethernet switching fabric.

On a centralized switch network, however, data from the client passes through the access point and is directed to the wireless switch before it can delivered to its final destination. This can mean passing through several extra network components along the way, and can add appreciable latency to the network.

As a general rule, it is best not to run a centralized wireless switch on a network that involves a wide-area network link between the wireless clients and the wireless switch. This configuration can add too much latency to the network. Of course, it also puts extra traffic on portions of the network.

As usual, there are exceptions to this rule. It might make sense to put the wireless switch on the other side of a wide-area link if the application host is also located across that same link. In this case, the data is going to traverse the wide area link anyway, so it does not matter which end of the link the wireless switch is located on. This assumes a typical data collection application, where the only wireless clients are the data collection clients and only one application is run on those clients.

Data Flow Comparison



A centralized wireless switch architecture will always cause any given data packet to pass through more network interfaces than will a packet in a distributed access point architecture. In the network shown here, a data packet will traverse eight network interfaces (plus however many interfaces are included in the corporate network cloud) before reaching its final destination.

The same packet on a distributed access point system will traverse only four network interfaces (plus the cloud) before reaching its final destination. Each interface traversed will add a small amount of latency to the network, so even without considering the effects of a wide area link, the switch network will create additional latency in the network.

The packet processing done by the wireless switch may add latency to the network as well. In any wireless network, there is some processing done on every data packet in bridging the packet for the wireless network to the wired network. This processing includes the layer 2 bridging function, authentication, encryption and decryption. In a distributed access point architecture, all of this processing is done at the access point. In a centralized wireless switch network, some, and sometimes all, of this processing is done at the wireless switch. The processing power of the switch becomes an issue in this case. Since it must process traffic coming from many access points, it needs a significant amount of processing power. This issue may become even more critical if new encryption schemes are deployed.

The centralized wireless switch network we show here adds one layer of network switching between the access points and the wireless switch. This is a common requirement for centralized switch networks, but some centralized wireless switch networks will require more layers of switching, and some include enough network switching in the wireless switch that access points can be directly connected.

The point is, that in any application where latency is a concern, it is best to minimize the number of devices a data packet must traverse. Centralized wireless switch networks always add at least one device (the wireless switch) to the network.

WLAN Management

These days, it seems most of the really pressing security issues are being addressed, and that network management is becoming the next big frontier to be conquered by all wireless LAN vendors. Wireless switch vendors will claim this is where their products really stand out. By definition, the wireless switch systems provide a centralized point of control that provides an ideal platform to provide all network management functions. It is just a matter of adding software to provide access point configuration, performance monitoring, fault detection, accounting, security policy enforcement and other network management functions.

The distributed access point architecture does not lend itself quite so naturally to centralized network management. With a distributed architecture, it is necessary to add a separate management component to the network.

In reality, however, there is nothing in the architecture that makes one type of wireless network inherently more manageable than another. It really comes down to the software available, whether that software runs on the centralized wireless switch, or whether it runs on a management specific component added to the network. Many vendors make impressive claims about how their networks can be centrally managed, but these same vendors have delivered only very rudimentary capabilities. This comes right back to the central point: Be extremely careful to evaluate which features your wireless LAN vendors are actually delivering. If advertised features are ‘coming soon’, how confident are you that the vendor can deliver on their promises? How long has the vendor been promising these particular features?

The centralized wireless switch systems also suffer from a scalability problem. Most wireless switches will handle only a limited number of access points. An enterprise will be faced with deploying multiple ‘wireless clusters’, each centered around a wireless switch. An enterprise deployment, therefore, requires a centralized management point to manage all the wireless switches in the network. In some implementations, one wireless switch can be designated as the ‘master’ and can propagate configuration and security profiles to the other switches in the network. In other implementations, this functionality is still missing.

If we look to the Ethernet world for a model, all hardware vendors provide some level of management for their equipment, but there are also a large number of 3rd party management applications that address a wide variety of issues. There are also large management platform vendors who can tie together products from multiple vendors and can provide a system-wide management view.

For the most part, these products are lacking in the wireless world. Management protocols tend to be a combination of SNMP and proprietary protocols. In the distributed

access point architecture marketplace, there is only one vendor who has had any success at providing cross-platform management capabilities. And that vendor can manage access points from only a half-dozen manufacturers.

In the centralized wireless switch marketplace, there are no 3rd party management tools. These wireless switch architectures are very closed and proprietary. Some wireless switch vendors attempt to provide a subset of management functionality for access points from other vendors. But the most heavily marketed wireless switch products are totally closed and can be used only with access points marketed by the wireless switch vendor.

There has been a move made to standardize the management protocols used by wireless switch products. This movement has been dubbed the “Lightweight Access Point Protocol”, or LWAPP. If LWAPP ever gets off the ground, there may be some hope for interoperability between wireless switches. But the success of LWAPP is far from certain. So far, the industry is divided in its backing for LWAPP.

Network Throughput

For data collection networks, throughput is not typically a major concern. Most data collection applications demand only 10 to 20 kilobits per second per access point. For these networks, it really does not matter much if the network can deliver 6 megabits per second, or if it can deliver only 4½ megabits per second.

However, many enterprises are beginning to look forward to the day when they will have new data collection applications available for the warehouse. These new applications might have higher throughput requirements. There is also the possibility that new wireless applications will be added to warehouse network. One commonly talked-about application is wireless voice over IP – or 802.11 IP telephones. Latency, which we already discussed, is the most important factor for deploying a useable Wireless IP telephony solution, but network throughput is also an important consideration.

The biggest factor in determining the throughput of a wireless network is the radio used. Each 802.11 family has a theoretical maximum throughput. No real-world implementation actually achieves this theoretical maximum, but some come closer than others. For example, with 802.11b radios, the best performing radios will deliver approximately 6 Mbps of throughput as measured using FTP file transfers.

Although throughput is typically governed by radio performance, the rest of the wireless network must be able to keep up. This means the access points must be able to process packets at the rate received by the radio. Each packet requires bridging between the wireless interface and the Ethernet interface. Packets may also require filtering, VLAN tagging or encryption or decryption.

As you can see, the issue here is not so much the throughput in bits per second, but rather the number of packets per second that must be processed. When throughput figures are quoted for a wireless system, these are always measured using large file transfers. This results in maximum sized packets, and so does not particularly tax the systems ability to

process packets. A more difficult test would be to flood the network with 100 to 200 byte packets. This packet size is more typical of packets encountered in a data collection application. But, then again, the most data collection applications generate so few of these packets that this never becomes an issue.

Throughput Vs Architecture

In a distributed access point network, the access point is responsible for all packet processing. This is a pretty easy target to hit using modern embedded microprocessor technology. Even with the 802.11a and 802.11g networks, which employ wireless signaling rates as high as 54 Mbps, modern microprocessors can easily keep up. The current generation of 'enterprise' class access points is even designed to handle the increased processing that may be demanded to comply with the upcoming 802.11i security standard. The 802.11i standard is expected to specify AES encryption to replace the current 802.11 WEP standard. Even though AES is designed to be computationally efficient, it still requires significantly more processing power than the RC4 algorithm, which forms the basis for WEP.

The downside to the distributed access point architecture is that it is difficult to upgrade if a future wireless standard requires more processing power than today's access points can deliver.

The centralized wireless switch architecture, on the other hand, moves a good deal of the processing off the access point and onto the wireless switch. This relieves the burden on the access point, but means that the wireless switch must be able to process packets from multiple access points at line speeds. This is a much more difficult than processing packets from one or two radios.

In theory, this type of performance should be achievable. Traditional Ethernet routers and switches are capable of processing millions of packets per second. But Ethernet routers and switches are built with custom designed hardware and ASICs, and run optimized operating systems. The current generation of wireless switches is typically built on off-the-shelf, general purpose computing platforms and commercial operating systems.

As a result, the throughput capabilities of a wireless switch limit the number of access points that can be connected to it – which, in turn, dictates the number of wireless switches required in the network.

However, it is easier to upgrade a single wireless switch than it is to upgrade 30 to 40 access points if additional processing power should become necessary. Some vendors expect to be able to upgrade their wireless switches simply by adding an expansion card. Even if it becomes necessary to replace the entire unit, that is relatively easy – being just a matter of replacing one device in the wiring closet with another.

An enterprise must consider, though, the type of upgrades that might be required in the future. For the next 3 to 4 years, it is likely the only processing power issue to be

encountered will be related to the new 802.11i security standard. The current generation of distributed access points is already equipped to handle those processing demands. Looking further into the future, it is likely there will be new, higher speed wireless protocols coming. But any higher speed wireless will require new radios. And radios are a component of the access point regardless of the architecture. Upgrading radios will require changes to the access points whether using a wireless switch or using distributed intelligent access points.

In the wireless switch architecture, packet processing is divided between the wireless switch and the access point. Wireless switch vendors often cite this as an advantage of the switch architecture. They claim that it is less expensive to

WLAN Costs

I have left cost for last, not because it is least important, but because it is perhaps the most difficult topic to come to make any conclusive statements about. But one thing is certain – everyone cares about costs!

The cost of a network can be considered from many different perspectives. Here we will attempt to look at two components of cost: Acquisition costs and operating costs.

Some wireless switch vendors have attempted to make the case that their systems are more economical to purchase, because they can rely on less expensive access points. This is just marketing. The purchase price of any product is governed by market factors, not by the manufacturing costs. Two systems with identical specifications can have vastly different acquisition costs if one vendor offers better support than another, or if one vendor has a more favorable reputation than the other, or even if one vendor spends more on marketing than the other.

Or, put another way, if one product appears to be less expensive than another, it is probably because the first vendor is trying to take market share away from the second. It is not surprising that many of the wireless switch systems on the market today appear to low-priced compared to distributed access point systems from the market leaders. It is because many of the wireless switch vendors are newcomers to wireless networking and have not yet established a good reputation, or have not demonstrated enough longevity to build confidence in their business.

Put yet another way, those vendors who can charge a premium for their products, because of their service and market reputation, do charge a premium. Those who can't, don't.

Operating costs, also, are not determined by the centralized switch vs. distributed access point architecture. Some important factors that contribute to operating costs are the management capabilities delivered with the network, the product's reliability and vendor support, and the flexibility of the network to accommodate changes. None of these things are more generally attributed to one type of architecture over the other.

Conclusion

There are many factors that need to be considered when shopping for a wireless LAN. Some (but probably not all of them) are listed here. Network architecture is not a critical consideration. What is critical is that you make sure the network does what you want it to do. Make sure you understand what is actually delivered today, and what is merely promised for the future. Make sure the vendor can support the wireless network in a suitable fashion. Make sure you have confidence the vendor can deliver on their promises. Do they have a track record of providing the support you require? If there is any 'planned' (not delivered) functionality that is important to you, do you believe the vendor will be able to deliver as promised?