

Network Design – Best Practices for Deploying WLAN Switches

A New Debate

As wireless LAN products designed for the enterprise came to market, a debate rapidly developed pitting the advantages of standalone fat access points (APs) against the benefits of WLAN mobility switches paired with slimmed-down APs. As WLAN systems have evolved, the industry has rallied around the WLAN mobility switch architecture, recognizing its scalability, security, performance, and cost of ownership advantages.

But just as that first debate has been settled, a new one is emerging. This time, at issue is where to deploy the WLAN switch, and three primary options have emerged. The WLAN switch can reside:

- 1) In the wiring closet, with directly attached APs;
- 2) In the distribution layer, with L2 switches between it and the APs; or
- 3) In the core, with L3 and L2 switches between it and the APs.

Key Issues

To understand the relative advantages of each option requires a closer look at the WLAN features impacted by these architecture choices.

Security of WLAN Switch-to-AP Link

While the security focus for WLANs is typically on encrypting the air link, IT also needs to take care that the wired portion of the network doesn't contribute to security vulnerabilities of the WLAN. The more devices that reside in the path between the AP and the WLAN switch, the greater the number of places a hacker can insert a malicious device. Similarly, if the WLAN switch is linked to the AP via one or more intermediate switches, then the AP can be replaced by a hacker's client device or AP and used as a point of attack on the wired network. A hacker can use the inserted device to launch an attack or snoop and duplicate wired traffic. If the WLAN switch hosts APs directly, then that link cannot be usurped for another purpose – the switch will recognize the absence of the expected AP and disable the port, shutting down that link as a potential point of attack.

Scalable Throughput

Given the limited bandwidth that the shared medium of wireless LANs provides, no WLAN device should itself ever become a potential bottleneck. Distributing the WLAN switching across multiple wiring closets increases the overall computational capabilities of the WLAN system. With a distributed architecture, no single centralized device is responsible for processing all the WLAN traffic. When WLAN switching is centralized in one device designed for the data center, that switch quickly becomes a bottleneck for the WLAN. Performing tasks such as terminating users' Transport Layer Security (TLS) tunnels needed for 802.1X authentication and classifying WLAN traffic for prioritized delivery are processor intensive and better served by a distributed architecture.

Impact of Wireless Traffic Data Path on Wired Backbone

In a wireless system based on WLAN switching, all WLAN traffic must pass through a WLAN switch to ensure security, even if the users communicating are associated with the same AP. The WLAN switches have the needed horsepower to enforce traffic separation between users; APs don't.



Consider the data path implications of where the WLAN switch is deployed. Putting a WLAN switch in the core means that even traffic between users on the same AP will traverse the wired backbone twice, increasing the traffic load on the backbone and the potential for congestion. In a distributed architecture, the WLAN switches in the wiring closet enforce policies on traffic between users associated with the same AP, without that traffic ever leaving the wiring closet. This design significantly contributes to the scalability of the WLAN system without increasing the traffic burden on the rest of the wired network.

System Redundancy

Deploying WLAN switches in the wiring closets provides stronger system redundancy for the WLAN. This deployment ensures that APs with dual 10/100 Mbps ports can be connected to two different mobility switches. Having two Ethernet ports provides power redundancy and continuous operation in the event of a WLAN switch failure.

In contrast, the centralized WLAN switch and all intervening power over Ethernet (PoE), Layer 2, and Layer 3 devices become single points of failure for the wireless network. In addition, core WLAN switches either cannot support redundant configurations between themselves or to the APs, or they complicate distribution of wireless services.

Delivery of Quality of Service (QoS)

To provide QoS, a system must both classify and mark traffic and enforce traffic treatment according to those markings. Given their relative processing strength, WLAN switches have the horsepower to classify, while APs are the appropriate place to treat traffic. When traffic classification occurs in a centralized WLAN switch, it has no assurance that the intervening switches will preserve that classification or enforce the classification when congestion occurs. When classification takes place in a WLAN switch in the wiring closet, with a direct connection to the APs, the appropriate latency, jitter, and prioritization metrics are enforced.

Configuration, Deployment, and Replacement of APs

The ease with which APs are configured and deployed or replaced will depend on the WLAN switch placement in the network. APs that are directly connected to a WLAN switch can be fully configured and managed by that switch. The direct physical connection means the WLAN switch automatically knows which AP is where, and the AP configurations reside on the switch for download to the APs. This benefit of automatic configuration will apply not only for the initial deployment but also for APs that are added to increase capacity or service new WLAN coverage areas and for a replacement AP when one fails. APs directly attached to WLAN switches can be designed to be so simple they act like light bulbs and can be replaced without touching the configuration of the wireless switch or the management application.

With a centralized WLAN switch, the AP is physically separated from the switch, so the system needs a mechanism for addressing each AP and identifying which one is where. This design will necessarily increase the time needed to configure wireless switches and APs throughout the WLAN, since some AP-dependent information, such as its MAC address, will need to be documented. Replacing an AP will also be more complicated, since the new AP's information, such as its MAC address, will be required for a management application to send it the correct configuration information.

Number of Managed Devices

WLAN switch designs will impact the number of new devices to be managed. With either a centralized or distributed architecture, the number of APs will be constant – what will vary is the number of other managed network devices. At first glance, it might appear that the core design, where the WLAN switch resides in the data center, introduces fewer managed devices. However, this perspective overlooks a key fact – all APs require power, so PoE must be added for all APs. The deployment of WLAN switches in the core or distribution layers requires separate PoE, increasing the number of managed devices in the network and the number of software management interfaces. With the wiring closet approach, the WLAN switch and PoE are integrated, resulting in fewer managed devices and a single management interface to control both the WLAN switch and the power.

PoE

Since all WLANs require some change in the wiring closet to provide power, deploying WLAN switches in the wiring closet eliminates the additional hassle of determining a compatible power source for APs when

centralized WLAN switches are used. In addition, providing PoE in the WLAN switch allows IT to maintain intelligent PoE control, simplifying configuration and troubleshooting since all WLAN errors, including power errors, will be recorded in a single record log. Adding WLAN switching to the core or distribution layer necessarily means that PoE will not be integrated and will therefore require separate configuration, management, and troubleshooting.

Cost Implications – Capital, TCO

Ultimately, the design advantages associated with distributing WLAN switching in the wiring closet yield a lower overall cost. The combination of fewer devices, a single management application, zero-configuration APs, and easier troubleshooting with integrated PoE lowers the capital and support costs of the wireless LAN.

Traffic Isolation – VLANs/subnets

All WLAN designs must address the issue of traffic separation and isolation among users. In the wired network, IT has a significant investment in subnets and VLANs for separating traffic and isolating users into groups. Where the WLAN switches reside has a direct implication on how wireless user traffic is separated for broadcast control and policy enforcement.

Many advocates of data-center-based WLAN switches, often called appliances, require the creation of new “wireless” subnets, implemented just for users to log into when on a wireless device. Providing one shared subnet for all wireless users, however, makes traffic isolation and differentiation of users’ traffic nearly impossible. Users who would typically reside in different subnets when on the wired LAN now share a common subnet when using wireless devices. As a result of this traffic mixing, this approach requires the added use of VPNs to achieve traffic isolation. For more information on the use of VPNs to secure a WLAN, read the Trapeze Networks white paper “Using IPsec VPNs to Secure the Air.”

This design creates several problems. As the network grows and requires more appliances, those new “wireless” subnets need to be distributed to all the appliances. In addition, IT has to figure out how to support roaming of VPN sessions across the appliances as users move throughout the enterprise. Either the system must disconnect a user’s VPN and restart it when users roam, or the user’s traffic must be immediately recognized by the new appliance and securely tunneled back to the originating appliance. Both options create significant performance implications in a scaled deployment. Bear in mind that many proponents of centralized WLAN switches are advocating placement in the data center not because that location has inherent benefits but because deploying VPNs across distributed WLAN switches, and maintaining those sessions as users roam across subnet boundaries, is very difficult.

A distributed WLAN switch design can more easily make use of existing subnets, and some distributed WLAN switches can dynamically support VLANs as needed, removing the requirement to distribute VLANs on all backbone router ports. The distributed design removes interoperability concerns between VPN servers and clients since it uses the existing VPN infrastructure, and it enables IT to avoid the added cost of testing and deploying new VPN servers. It also solves the issue of roaming, because it provides a Layer 2 path from all WLAN switches to the VPN server that stays intact regardless of the user’s location.

Design Myths

Myth #1: I can deploy WLAN switching without touching my wiring closets.

All WLAN deployments require power for APs, so all WLANs necessitate a change in the wiring closet to add new Ethernet links with PoE. Though you may have an extra Ethernet port, deploying PoE separate from the WLAN switch actually complicates installation and troubleshooting, increasing the number of devices in the network and compromising system integration. If PoE is already deployed in the wiring closet, issues regarding power compatibility and adequate wattage remain. Often, earlier PoE deployments were aimed at supporting voice-over-IP phones and therefore lack sufficient wattage for APs. In addition, the 802.3af standard for PoE has passed so recently that few vendors are shipping standard-compliant equipment and interoperability issues will emerge because the specification itself has several options that vendors will adhere to differently.

Myth #2: Putting WLAN switching in my data center will impact my network less.

Regardless of where the WLAN switching resides, APs will require Ethernet and likely PoE. So core deployments mean IT touches the network in more, not fewer, places. IT will change the core to add the WLAN switch, the closet to add power, the ceilings to deploy the APs, and points in between to provide an Ethernet link. In addition, the core could be further impacted by the need to support mobility. Depending on the implementation, core deployments may also require the creation of new “wireless” subnets and, to enable roaming of multiple private groups over the WLAN, network access for all of those separate subnets must be available in the data center. To maintain those groups in the WLAN and enable full mobility, IT has to change the configuration of the router ports and wiring closet switches to support all the groups throughout the enterprise.

Myth #3: WLAN switches are all the same – where I put them in my network won’t impact their functionality.

Network design and placement of the WLAN switch will dramatically impact the performance of the wireless system. When intervening devices sit between the WLAN switch and the AP, key features such as QoS, performance scalability, and security are compromised. IT cannot be sure that intervening switches will classify wireless traffic the same way the WLAN switch does. Also, a single device must process all wireless traffic, and intervening devices can be replaced by hackers’ devices and used to attack the network. Directly connecting APs to the WLAN switch eliminates all these issues and ensures the highest level of service, scalability, and security.

Myth #4: Putting WLAN switching in my wiring closets will be more expensive.

At first glance, it appears that buying one core device will be cheaper than several wiring closet devices. But in reality, core- or distribution-based WLAN switches have to be complemented by PoE in the wiring closet, so the number of WLAN-related devices increases in these designs. In addition, core and distribution designs cannot easily support automatic configuration of APs – the APs must somehow be addressable, which means the AP must be configured prior to deployment, increasing the cost to deploy.

Best Practices

Directly connecting APs to WLAN switches provides the highest security and scalability, simplest configuration, and greatest system functionality. This design ensures a high-capacity system that delivers user-specific services and security. As a result, this configuration is and will remain the best practice for deploying WLAN switching.

Putting WLAN switches in the distribution layer, with intervening L2 switches, fits a couple of applications. First, this design enables IT organizations to support a remote location that requires just one AP. In this case the rest of the WLAN with directly connected APs will support the necessary capacity and high-level functionality, and IT can accept the less-integrated approach in an isolated area of the network. Also, IT may use this design on an interim basis as a starting point for locations where they are just beginning a WLAN deployment. It provides IT with a phased approach to building out the wireless network, using fewer APs and providing just WLAN coverage initially rather than designing for capacity. This design will yield limited functionality, scalability, and physical security and so will meet IT’s needs only in the early stages of WLAN use. As the WLAN grows, IT will migrate to locating WLAN switches in the wiring closet and directly connecting APs to avoid spending additional “throw away” dollars on separate PoE devices and other interim equipment.

Some IT organizations will still decide to deploy centralized WLAN switches in the core, and these groups should make sure the architecture will meet their needs. IT will need a means for controlling APs across IP router boundaries, since intervening equipment will now certainly include L3 devices. IT will still have to make changes in the wiring closet, and they’ll need to look for PoE devices that support management and troubleshooting. Further, IT should keep in mind that the dollar savings anticipated by using existing wiring closet ports to support APs will likely be spent on the additional costs for managed PoE, AP configuration, and backbone changes needed to deploy the WLAN.

Summary – Table Comparing WLAN Switch Network Design Options

	WLAN switch in wiring closet with directly attached APs	WLAN switch in distribution layer with L2 switches between it and APs	WLAN switch in core with L2 and L3 switches between it and APs
Security of WLAN switch-to-AP link	High	Medium	Low
Scalable throughput	High	Medium	Low
Impact of wireless traffic data path on wired backbone	Low	Medium	High
System redundancy	High	Low	Low
Delivery of QoS	High	Medium	Low
Ease of rollout, configuration and replacement of APs	High	Low	Low
Number of managed devices	Low	Medium	Medium
Ease of PoE integration and management	High	Low	Low
Cost	Medium	Low	Medium
Ease of traffic isolation – VLANs/subnets	High	Low	Low



www.trapezenetworks.com

Corporate Headquarters: 5753 W. Las Positas Blvd., Pleasanton, CA 94588 Phone 925.474.2200 Fax 925.251.0642

EMEA Headquarters: Olympia 10c, 1213 NP Hilversum, The Netherlands Phone +31 (0) 35.64.64.420 Fax +31 (0) 35.64.64.429

Trapeze Networks, the Trapeze Networks logo, the Trapeze Networks flyer icon, Mobility System, Mobility Exchange, MX, Mobility Point, MP, Mobility System Software, MSS, RingMaster, Trapeze Access Point Access Protocol and TAPA are trademarks of Trapeze Networks, Inc. Trapeze Networks SafetyNet is a service mark of Trapeze Networks, Inc. All other products and services are trademarks, registered trademarks, service marks or registered service marks of their respective owners. © 2003 Trapeze Networks, Inc. All rights reserved.