

White Paper |

Enterprise



Scaling Enterprise Wireless LAN Deployments

Keerti Melkote

ARUBA[®]
networks

Introduction

Enterprise wireless LANs (WLANs) have expanded rapidly over the past few years, moving from small hotspot style deployments in conference rooms and other common areas, to pervasive enterprise-wide deployments that span the campus, branch office, telecommuters and even nomadic remote offices. As these wireless networks grow in size, their scalability is primarily determined by the underlying architecture for of the wireless LAN and its interworking with the wired architecture.

Enterprise WLAN design has evolved from a distributed to a centralized model. It is clear that centralized WLAN architectures are here to stay and will be the dominant method of building enterprise wireless networks. However, not all centralized architectures are created equal. Customers are faced with two architectural options even with centralized architectures. One option is to embed centralized WLAN capabilities into the existing network infrastructure. This requires an upgrade to the fixed, or wired, edge of the network to address the challenges associated with mobility. The other option is to create a new mobile edge that extends beyond the existing fixed edge and allows users to connect from any location, at any time. A mobile edge requires an overlay network model that delivers mobile connectivity across the corporate network and the public Internet.

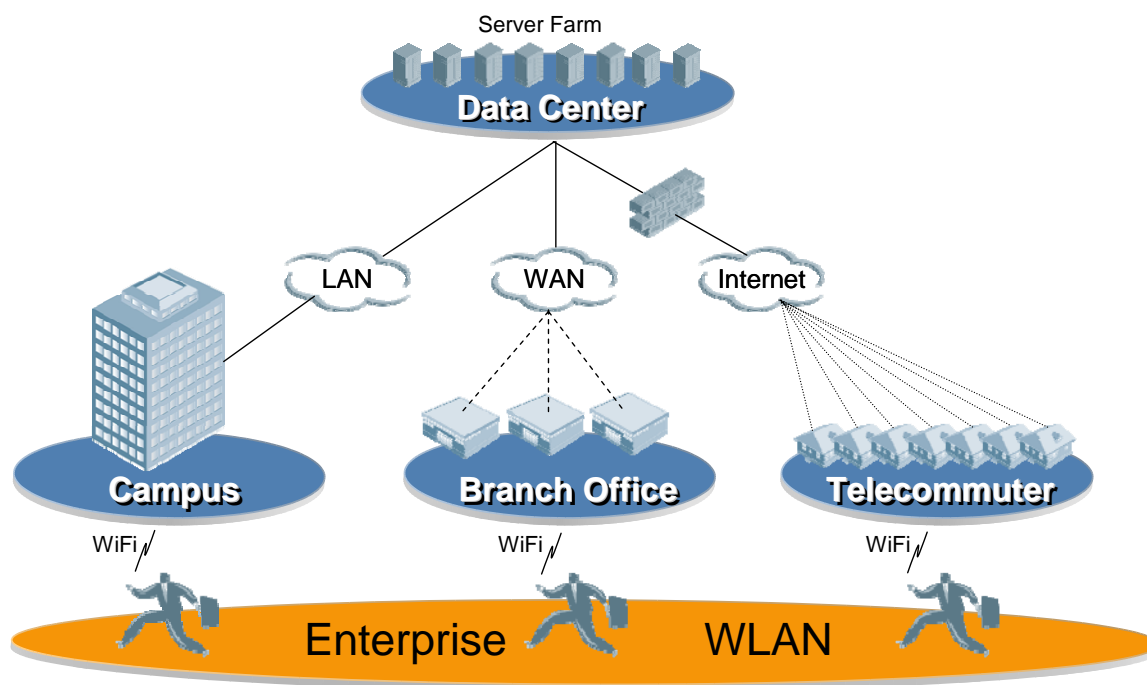


Figure 1. Mobile Edge Architecture – Common User Experience across LAN, WAN and Internet

Determining which products and solutions available today can address this fundamental architectural difference can be difficult since most of the industry rhetoric seems similar. One key area of differentiation is scalability. Traditional scalability metrics of centralized wireless LAN architectures have focused on controller throughput and the number of thin access points supported by centralized WLAN controllers. While these are

important metrics, real-world experience in deploying high-end enterprises has yielded fresh insight into scaling requirements for wireless LANs. The challenges of scaling enterprise wireless LANs can be categorized by three key enterprise WLAN categories

- Campus wireless LANs that have hundreds to thousands of users and devices,
- Branch office wireless LANs that have between ten and one hundred users and devices, and
- Telecommuter and nomadic office wireless LANs that have between one and ten users.

Scaling Campus Wireless LANs

As the enterprise workforce becomes increasingly mobile, user counts on campus wireless LANs are constantly on the rise, and with the proliferation of Wi-Fi-equipped personal handheld devices, device counts are increasing even more rapidly. The key challenges of scaling a campus wireless LAN are caused by the density of users and devices, instantaneous loads caused during peak hour usage, and the mobility of users between different areas on the campus. The associated technical challenges relate to the scaling of RF capacity, AAA services and VLAN architecture for mobile networks.

Scaling RF Capacity with Multi-channel RF Architecture

All centralized WLAN architectures today incorporate some level of RF management functionality, which is designed to automate the site survey process. However, in most implementations, RF management is limited to pre-planning, and makes use of heavy duty RF planning software. Other vendors claim to eliminate the entire planning process by moving to a single channel architecture. Both approaches leave much to be desired when it comes to delivering high-capacity wireless LANs.

In the first instance, planning access point placement based on building materials and other RF planning models is fundamentally flawed because the RF characteristics are dynamic and change constantly. This results in a failure to adjust to ambient RF conditions or, worse yet, in sub-optimal results, when assumptions regarding building materials and other variables are flawed. Single channel architectures, while eliminating the planning problem, introduce an issue related to client density. When all clients are operating on the same channel, co-channel interference increases significantly, resulting in poor performance.

Multi-channel RF architectures are inherently better suited for high density usage since they utilize all available channels in the spectrum to reduce co-channel interference. However, multi-channel architectures must be completely automated from a deployment standpoint. New techniques such as Adaptive Radio Management (ARM) are emerging in the industry to completely automate the deployment of multi-channel RF architectures and reduce co-channel interference. This leads to much higher RF capacity and better RF performance of WLAN networks.

As density increases, enterprises are employing strategies to migrate to 802.11a on the 5GHz band which has 4-5 times more capacity than the 2.4GHz band. The 5GHz band is also inherently much cleaner with respect to interference, yielding better and more consistent channel performance. The 2.4GHz band will continue to be the first choice for equipment manufacturers of most handheld mobile devices such as Vo-Fi phones, PDAs, dual-mode phones, barcode scanners and active RFID tags because of the greater maturity, lower cost, and lower power demands of 802.11b/g silicon. However, laptop manufacturers have finally caught up and are now implementing new power management efficiencies and adding support for 802.11a. The newer laptops with 802.11a/b/g network interface cards auto select, and, wherever possible, opt for, the 5GHz band. This, in turn, is resulting in a hybrid approach, using the 5GHz band for laptops and the 2.4GHz band for other handheld devices.

In addition, enterprises are increasing using a four-channel architecture in the 2.4GHz band instead of the traditional three channel approach, as the extra channel yields additional capacity, especially valuable in dense deployments.




Figure 2. Multi-channel RF Delivers Up To 3 Times the Capacity of a Single-channel RF Architecture

Scaling AAA Services with Hardware Acceleration of 802.1X Authentication

Even with additional RF capacity and a successful 802.11 association, devices in large enterprise networks may still be unable to connect to the network. This is often the result of heavy loads on the back-end authentication, authorization and accounting (AAA) server. This situation is being compounded with the implementation of new authentication practices as part of 802.11i.

802.11i, which requires all users and devices to authenticate to the WLAN using the 802.1X authentication protocol, is established as an industry standard best practice for securing enterprise WLANs, 802.11i The

National Institute of Standards and Technology (NIST), responsible for setting government standards has, in fact, mandated the use of 802.11i in securing WLAN networks.

Traditionally, in centralized WLAN architectures, the controller only serves as an authenticator in the 802.1X authentication process. The actual AAA transaction of verifying a username and password combination is carried inside an encrypted TLS tunnel between the wireless client and the AAA server. Typical tunnel types used today are PEAP and EAP-TLS, with PEAP as the dominant method.

The introduction of 802.11i forces AAA servers to take on an even greater computational burden.

The AAA server is given the responsibility of both terminating encrypted authentication network protocols such as EAP-PEAP, as well as generating the encryption keys that are used by WLAN clients and access points for secure wireless 802.11 communications.

As user density and the number of login requests per second goes up, the backend AAA server's ability to process cryptographic information with consistent response times while simultaneously authenticating and authorizing users becomes a bottleneck. Users in heavily loaded wireless networks end up with slow, variable response times during network login, and may even experience network disconnects due to timeouts. Customers who have experienced this problem end up having to set up multiple AAA proxy servers to scale AAA processing capacity in the network. The extra proxy servers and associated network redesigns increase network complexity and add both capital and operational expense.

Solutions to this problem are emerging from some centralized wireless LAN vendors whose WLAN controllers are architecturally capable of absorbing the fixed, but immense, overhead of the 802.1X authentication process. These controllers incorporate purpose-built hardware encryption processors to terminate the PEAP/TLS tunnels and centrally compute the crypto keys for secure wireless communications, offloading the back-end AAA server from this significant processing burden and leaving it free to perform the tasks of AAA. This approach, known as AAA FastConnect, results in over 1,000 authentications per second – a tenfold increase – eliminating the issue of slow connect times and failed login attempts.

BEFORE: TRADITIONAL AAA ARCHITECTURE WITH CENTRALIZED WLANS

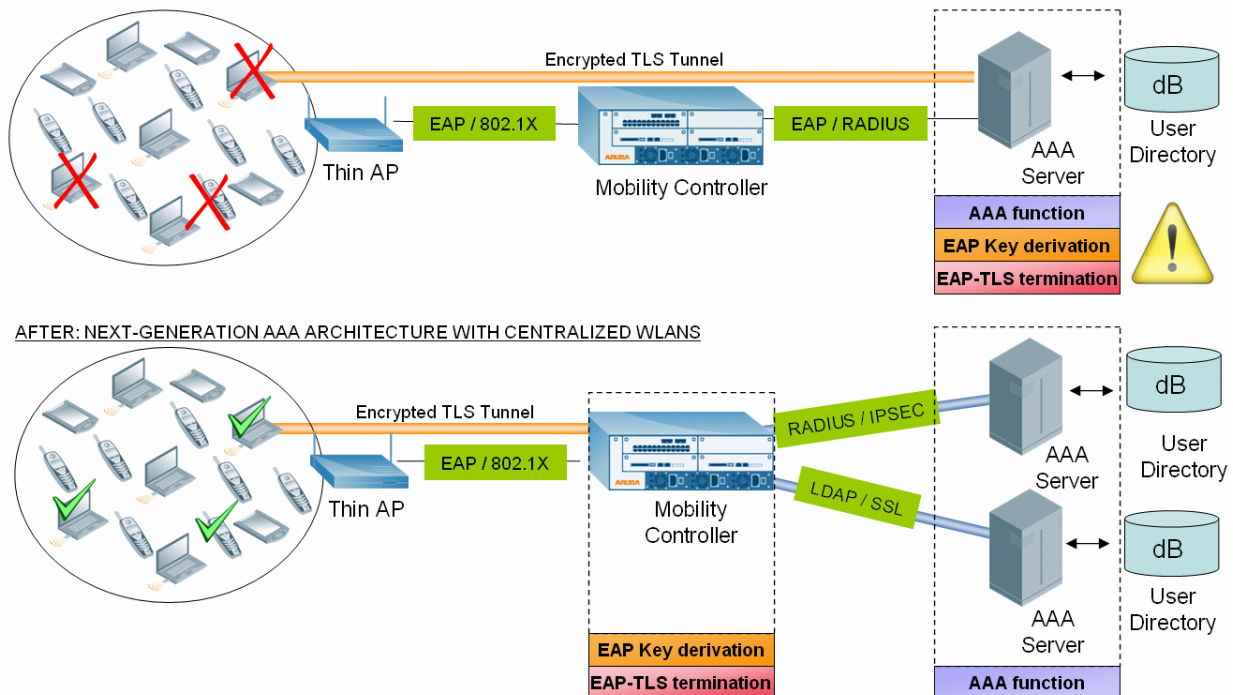


Figure 3. Before and After Comparison with Hardware Acceleration of AAA Services

AAA FastConnect not only results in faster and more predictable connect times, but also greatly simplifies the integration of secure WLANs with various back-end servers. In traditional AAA architectures, back-end AAA servers must be upgraded to handle 802.11i security since centralized controllers are just a pass through relay in the authentication phase. With AAA FastConnect, a mobility controller can interoperate directly with a AAA server using RADIUS or LDAP since all AAA related 802.11i security requirements are absorbed into the mobility controller itself. Furthermore, RADIUS packets can be encrypted in an IPsec tunnel, while LDAP transactions can be encrypted in SSL to keep the entire AAA transaction encrypted end-to-end. Such flexibility is not possible with traditional AAA architectures. This enables the entire WLAN to operate as a secure overlay, without requiring any additional investment to upgrade or add security to the wired network and cost-effectively solving the scalability problem.

Scaling Virtual LANs (VLAN) for Mobile Networks

Once a user successfully associates and authenticates with a wireless network, the next layer up in the protocol stack requires assignment to a virtual LAN (VLAN) based on role or some other factor, and then dynamic acquisition of an IP address on the assigned VLAN via DHCP. The VLAN serves as the interconnect/interworking mechanism between mobile and fixed networks. As illustrated below, traditional fixed models of VLAN network design don't work for mobile networks; they are still important, but need to be altered to accommodate mobility.

Traditionally, VLANs have been used as a fundamental building block for all network designs. Their primary use today is to scale enterprise IP networks by mapping each VLAN to an IP subnet, thus limiting broadcast domains. Traditional fixed network design principles are port-based, where each port is implicitly assumed to serve a single device, such as a desktop computer or a VoIP desk set. Based on this assumption, as a rule of thumb, most fixed network designs placed no more than 200 ports, mapped to an equivalent number of IP addressed devices, into a given broadcast domain and or single VLAN. It is quite typical to map every port on the floor of a building to a single VLAN, and then map multiple floors to their own VLANs and IP subnets, thus creating a deterministic design that scales even to very large networks.

With the emergence of VoIP, VLANs are being used as a QoS mechanism as well. All VoIP traffic on a floor is separated onto its own VLAN, and the entire VLAN is tagged as high priority. More recently, VLANs are being used as a security mechanism, where, for example, guest users are placed on a dedicated guest VLAN or machines infected with viruses and worms are placed on a quarantined VLAN. These VLANs are secured from the rest of the network with VLAN access control lists (VACLs) that determine their network access levels.

Given the above model of fixed network usage of VLANs, how do mobile networks map into this framework? The default method of mapping wireless users to VLANs is to associate an SSID with a VLAN. However, as mentioned above, the scalability of this model is very limited since a VLAN typically cannot handle more than 200 users very effectively. Increased broadcast traffic on large VLANs not only causes performance problems and consumes precious over-the-air bandwidth, but also drains battery life on mobile devices. Therefore, the next step enterprises often take is to segment wireless traffic into multiple VLANs using floor-based VLAN assignments similar to the fixed network model. This approach results in the creation of a new set of mobile VLANs that parallels the fixed network data, voice, guest and quarantine VLANs. The addition of so many VLANs creates a VLAN explosion and excessive network complexity. Further, this approach fails at a fundamental level because it does not take mobile network usage patterns into account.

As campus user density increases and mobile network usage becomes prevalent, it becomes very difficult to predict the number of users that might be associated on any given floor at any given point in time. Under these circumstances, users may get associated and authenticated to the wireless LAN, but be unable to receive an IP address because the VLAN IP address space has been completely exhausted at the DHCP server. At this point, enterprises resort to flattening the IP subnets serving mobile users by increasing the VLAN size to accommodate the transient loads. Again, this results in the undesired effects of large broadcast domains.

Another approach is to over-provision the number of VLANs for each floor in anticipation of transient peak loads, but this leads to an even greater explosion in the number of VLANs required. Over-provisioning of VLANs not only wastes resources, but also raises operational costs and vastly increases complexity when network troubleshooting.

A new model is required to simplify and scale the VLAN architecture for mobile networks.




Figure 4. VLAN Pooling Simplifies Dense, Mobile Campus WLAN Deployments

A new mechanism, called VLAN Pooling, delivers the flexibility of VLAN-based network planning without any of the negative side effects discussed above. In VLAN Pooling, multiple VLANs form a VLAN pool, and all VLANs belonging to the VLAN pool are available at any location on the campus. VLAN assignment is performed dynamically at the time a user logs into the network and is based on current user loads on the different VLANs that form the VLAN pool. As an example, if a campus network has 1,000 users that can connect anywhere on the campus at any point in time, a total of 5 VLANs are required, based on the 200 users to a VLAN rule-of-thumb. These 5 VLANs are placed in a VLAN pool and made available at all points in the campus network. When a user logs in, they are assigned to one of the VLANs, typically the least used VLAN, based on current user counts of each VLAN in the VLAN pool. This results in even loading of all VLANs and ensures that every user gets an IP address and successfully connects to the network every time.

In the mobile network, VLANs are used only to limit the broadcast domains and map to IP subnets. Security and QoS assignments are done on a per-user basis and determined by the role of a user, not by their association with a particular VLAN. Security policy enforcement in a VLAN pool is performed using a built-in identity-based stateful firewall integrated in the mobility controller. There is no need to use VLANs for security purposes since the stateful firewall enforces policies on a per-user basis. QoS policy enforcement is also performed by the mobility controller using built-in VoIP application layer gateways (ALGs) that recognize VoIP call flows and run call admission control (CAC) algorithms based on the number of active calls in the air. QoS information is signaled to the fixed network via DSCP and 802.1p tags.

The VLAN pooling model greatly simplifies mobile network design. It retains the familiar VLAN construct, but uses it in an innovative way to minimize disruption and meet all the mobility, security and QoS needs of mobile users and devices.

Scaling Branch Office Wireless LANs

The primary challenges associated with branch office WLANs are the cost and complexity of deploying WLANs in a large number of branch offices, centrally managing a large number of branch offices distributed across a Wide Area Network (WAN) and keeping users connected to the branch WLAN even when the WAN link goes down.

Self-configuring Mobility Controllers for Automated Large-scale Deployments

Branch offices typically lack skilled IT personnel to set up and operate secure WLAN networks. Yet, users expect a consistent and secure mobility experience regardless of their location. To deliver a consistent user experience at the lowest operating cost for a branch WLAN, mobility controllers must provide simple self-configuration. This capability allows for the mobility controller to be centrally provisioned and drop shipped to a branch location for plug-and-play operation.

Self-configuring mobility controllers dynamically obtain an IP address from the branch firewall/router or broadband access provider using a built-in DHCP client or a PPPoE client. Upon obtaining the IP address from the network, the local branch controller automatically synchronizes its configuration with a centrally located configuration server (master mobility controller). This capability allows a non-technical employee to bring up a secure WLAN by simply plugging the mobility controller into the branch network, eliminating the cost and hassle of sending skilled IT staff to branch offices.

Automating configuration of branch office wireless LANs drastically cuts the total cost of deployment and is a critical first step in enabling a large-scale branch WLAN deployment.

Centralized Management for Thousands of Distributed Branches

As the number of branch office WLANs increases, a scalable centralized management system becomes imperative to keep operating costs from skyrocketing. The centralized management system must not only deliver a full suite of network management tools, including configuration, monitoring, reporting and troubleshooting, but also be able to scale to thousands of distributed branches without requiring the use of multiple management systems.

A hierarchical approach is the only way to scale to such large numbers and maintain a single point of centralized management. A single master mobility controller in a hierarchical model manages several

hundred local mobility controllers. A master controller delivers global configuration updates and monitors all local controllers under its control. Multiple master controllers are monitored together from a single mobility management system in order to deliver a global view of the entire system, resulting in unprecedented scalability for distributed mobile enterprises.

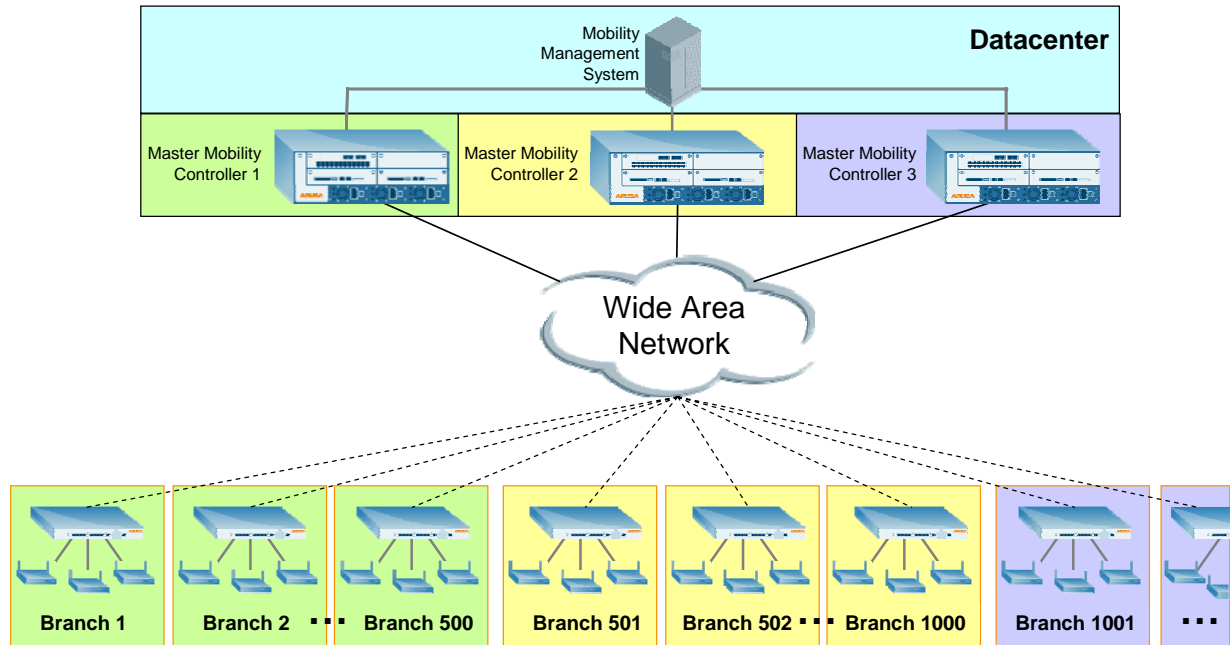


Figure 5. Hierarchical Management Model Scales to Thousands of Distributed Branch Offices

Resilient, Secure WLAN Connectivity for Continuity of Branch Operations

Secure WLAN connectivity using WPA or WPA2 requires an always-on connection to back-end AAA services, which tend to be centralized in a datacenter across the WAN. If the WAN link to the centralized AAA server goes down, local branch office operations are impacted; users are unable to print to local printers or access local file services at the branch offices. These limitations are unacceptable for most enterprises. The wireless LAN must be resilient and continue to offer secure connectivity, even when the WAN link to the central AAA server is down. The user expectation is for continuous local branch network operation.

Some work-around options available today require administrators to manually configure local AAA services on branch office routers as a failover option when the WAN link goes down. However, this configuration is not synchronized with the central AAA server and requires significant manual administration. If automatic synchronization of local and central AAA servers is an option, it requires a single vendor solution, which is often unfeasible, especially when corporate mergers and acquisitions have brought together a host of department-level AAA service implementations and user directories that must be federated together to create a seamless AAA infrastructure. Furthermore, locking down a branch to a specific AAA server defeats the

very purpose of mobility— allowing any user to show up at any location in the distributed enterprise and enjoy a consistent user experience.

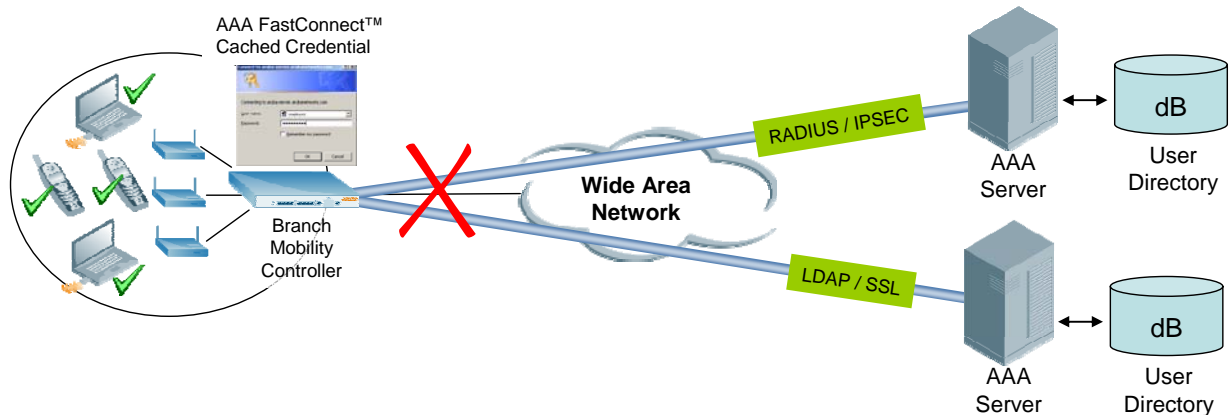


Figure 6. AAA FastConnect Credential Caching – Users can connect even when WAN link is down

AAA FastConnect, discussed earlier as a means of scaling campus AAA services, also offers a unique solution to the challenge of scaling branch AAA services. With AAA FastConnect, branch office controllers have the ability to locally cache user credentials the first time a user logs into the network. All subsequent connection attempts by the user are locally authenticated by the branch office mobility controller. This enables resilient access to the WLAN even when the WAN link goes down. Since AAA FastConnect caches user credentials in the form of an encrypted cookie, they are completely secure, even if the branch office controller is compromised. In addition, AAA FastConnect has the ability to interface with all industry-leading AAA servers, making it easy to integrate mobile users into federated and multi-vendor AAA infrastructures.

Scaling Telecommuter Wireless LAN Connectivity

Telecommuters increasingly demand access to corporate VoIP and data resources from their home offices. The requirement is for a simple and secure solution that users can just plug into their home networks to gain instantaneous, secure access to the corporate network over the Internet. However, telecommuter wireless LAN deployments have depended on either difficult-to-manage, stand-alone enterprise access points or completely unmanaged, highly vulnerable consumer access points.

Similar to the telecommuter WLAN requirement, there is an ever-increasing need for nomadic offices, which require setting up temporary network that last for a few weeks, a few days, or even just a few hours. This is a very common and critical requirement in the construction industry, where access to corporate resources from remote building sites. Tradeshow are another example of a nomadic office where multiple users need secure access to corporate resources from the show floor.

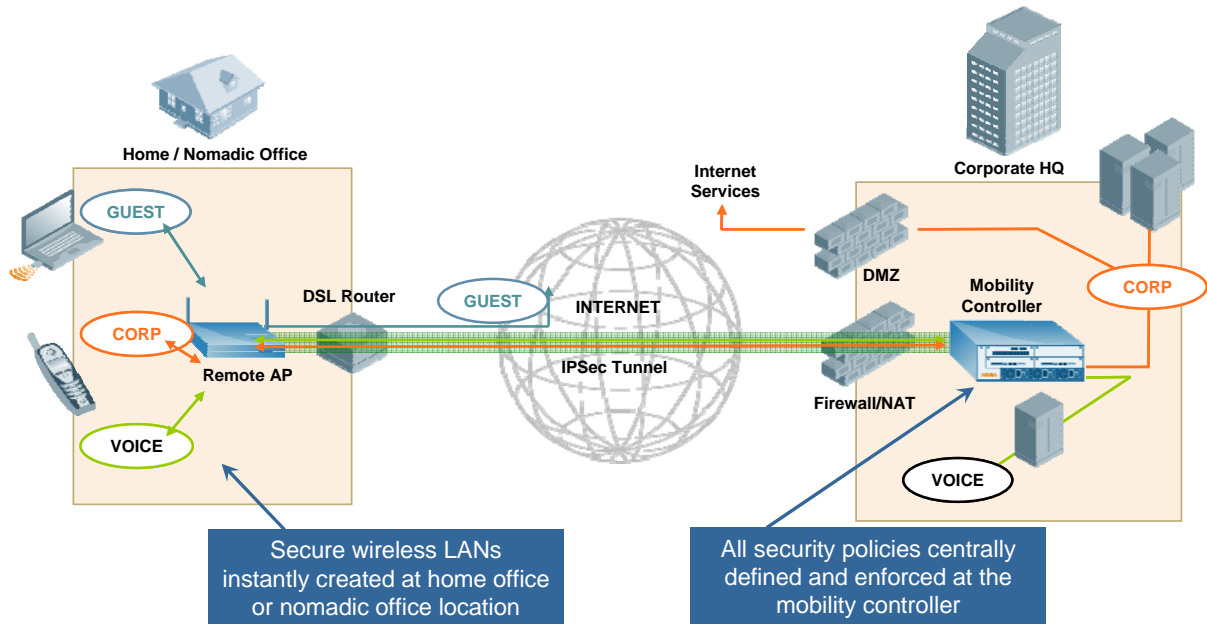


Figure 7 – Remote APs Instantly Create Secure Enterprise WLANs for Telecommuters

Remote Access Points (AP) deliver the benefit of securely and easily extending enterprise WLANs, to home offices and nomadic office locations. Remote APs are plug-and-play devices that require only very basic one-time provisioning by the IT department. Once provisioned to discover the central mobility controller over the Internet, remote APs allow mobile workers to take the enterprise wireless LAN with them wherever they go, securely accessing corporate VoIP and data services from any location. Large deployments of remote APs are possible at the lowest operational and capital costs since they are simple, secure and plug-and-play.

Conclusion

As workforce mobility becomes pervasive, enterprises are increasingly considering large scale deployment of secure wireless LANs. Enterprises are faced with two architectural choices; either extend the fixed edge of the existing network, or create a new mobile edge that spans the LAN, WAN and Internet. The mobile edge architecture not only delivers ubiquitous and secure mobile access, but also delivers unprecedented scalability. Unique capabilities, such as AAA FastConnect, VLAN Pooling, and Remote AP, created based on the needs of actual, large scale enterprise wireless LANs, are essential to delivering a reliable, cost-effective and operational enterprise wireless network.

Acronyms

AAA	– Authentication, Authorization and Accounting
ALG	– Application Layer Gateway
DHCP	– Dynamic Host Control Protocol
DSCP	– DiffServ Code Point
EAP	– Extensible Authentication Protocol
IPSEC	– Internet Protocol Security
LDAP	– Lightweight Directory Access Protocol
PEAP	– Protected Extensible Authentication Protocol
PPPoE	– Point-to-Point Protocol over Ethernet
RADIUS	– Remote Authentication Dial-in User Service
SSL	– Secure Sockets Layer
TLS	– Transport Layer Security
ToS	– Type of Service
VLAN	– Virtual Local Area Network
VACL	– VLAN Access Control Lists
VoIP	– Voice over IP
VoFi	– Voice over IP over Wi-Fi or Voice over Wi-Fi
WAN	– Wide Area Network
WLAN	– Wireless Local Area Network
WPA	– Wi-Fi Protected Access

About Aruba Networks, Inc.

Aruba securely delivers the enterprise network to users, wherever they work or roam, with user-centric networks that significantly expand the reach of traditional port-centric networks. User-centric networks integrate adaptive WLANs, identity-based security, and application continuity services into a cohesive, high-performance system that can be easily deployed as an overlay on top of existing network infrastructure. Adaptive WLANs deliver high-performance, follow-me connectivity so users are always within reach of mission-critical information. Identity-based security associates access policies with users, not ports, to enable follow-me security that is enforced regardless of access method or location. Application continuity services enable follow-me applications that can be seamlessly accessed across WLAN and cellular networks. The cost, convenience, and security benefits of user-centric networks are fundamentally changing how and where we work. Listed on the NASDAQ and Russell 2000® Index, Aruba is based in Sunnyvale, California, and has operations throughout the Americas, Europe, Middle East, and Asia Pacific regions. To learn more, visit Aruba at <http://www.arubanetworks.com>.

© 2007 Aruba Networks, Inc. All rights reserved. Specifications are subject to change without notice.

Aruba Networks, BlueScanner and RFprotect are trademarks of Aruba Networks, Inc. All other trademarks or registered trademarks are the property of their respective holders.

WP_SCA_US_071217

