

White Paper |

Government



Developing a Sound Security Policy for Mobility at the Department of Defense

Robert Fenstermacher and Jon Green

ARUBA[®]
ARUBA
networks

Introduction

As mobility becomes more prevalent in DoD operations, there has been an aggressive stance taken to mandate security of these networks. While the DoD represents a typical enterprise network in many respects, it also has unique requirements that set it apart from commercial deployments. Due to the nature of information that these networks transport and the stringent availability requirements for many mobile applications, an aggressive security posture has been established, one that only the most capable networking vendors have been able to comply with.

While a properly implemented mobile security policy will offer simple deployment and management and make wireless access even more secure than wired access, an improperly implemented or insufficient plan can lead to disaster.

The popularity of wireless technology and an increasingly mobile workforce are leading to a new type of network access layer – the “Follow Me Enterprise.” This user-centric architecture securely and reliably delivers enterprise networks to users, wherever they work or roam – at headquarters, remote sites or nomadic offices when users travel. However mobility, including wireless technology, has the potential to expose networks to intruders, leak sensitive data, and subject the network to virus and worm outbreaks. Proper planning avoids these issues without negatively impacting capital or operational expenses. This white paper provides clear recommendations for building effective security policies in the DoD that comply to stringent DoD Directives and National Institute of Standards and Technology (NIST) requirements. It also goes beyond compliance to look at the manageability of a mobile network and evaluate the pros and cons of two architectural approaches to wireless security – a centralized approach and a distributed approach.

Architectures for Mobility

There are three major network architectures available for building wireless LANs, although for the purpose of security this can be narrowed down to just two: Distributed and Centralized. A distributed architecture, as the name implies, distributes security functions to multiple devices while a centralized architecture collapses security functions into one device. A distributed architecture may consist of standalone “fat” access points, where the AP itself contains all functionality for wireless LAN operation. A distributed architecture may also consist of a controller with “thin” APs when the security functions of the wireless LAN are broken up between multiple devices. For example, if an AP performs encryption, the controller performs authentication, and a firewall performs access control, this is a distributed system from a security standpoint. A centralized system, on the other hand, places all

security functions in a single unit. In the example just given, encryption, authentication, and access control would all be done by a single controller in the centralized architecture. A centralized architecture is always made up of “thin” APs and a central controller. These architectures will be revisited in each section below to provide comparison and contrast between the capabilities of each.

Locking the Air

The first step in any wireless security policy is to lock down the radio spectrum against threats. This step must be done even if a wireless network is not actually deployed. For these locations that have no-wireless policies in place, a complete strategy for mitigating wireless threats must be implemented. Once a wireless network has been deployed, monitoring for attacks becomes an additional need.

Rogue APs

The very existence of wireless technology is a threat to security of the wired network. Internal demand for mobility is so great that many people, if not provided with wireless access, will install it themselves. Consumer-grade access points are inexpensive and easy to set up, and it only takes moments for someone to install one of these “rogue” APs in an office. Connected to the wired network, rogue APs become instant portals into the rest of the network, bypassing firewalls and other security systems. Putting an automated system in place to find, classify, and disable rogue APs is a critical requirement of a wireless security policy. This must be done for all points in the organization’s network where rogue APs could potentially be installed, including remote locations.

Uncontrolled clients

A second category of threats to the network is that of uncontrolled client devices. Many end-user devices such as laptops, PDAs, and mobile phones come equipped with wireless interfaces. When these devices are not properly secured, they can become a security risk with intrusion or loss of confidential information possible. As one example,

Windows XP can be configured to bridge a wired network interface together with a wireless interface. If this happens, an attacker may be able to use the bridged connection as a gateway into the network. As another example, many mobile phones support Bluetooth for connection to other wireless devices. If the mobile phone is not configured with correct security features, an attacker could wirelessly tap into the phone and download address books, stored email, and other information that could reveal business contacts or business plans.

Active Attacks

If a wireless LAN has been deployed, it must be monitored and protected against attack by malicious persons. Attacks range from simple RF jamming up to sophisticated “man in the middle” attacks where an attacker inserts himself into the communication path and is able to add, delete, or modify data in

transit. The proper use of encryption and authentication, discussed later, mitigates many of the risks, but a wireless intrusion prevention system is necessary for detecting and preventing the remainder. At a minimum, a wireless intrusion prevention system will identify active denial of service attacks so that valuable time is not wasted troubleshooting connectivity problems with the wireless LAN when the actual problem is an attacker. For organizations that intend to reprimand attackers, wireless intrusion prevention systems provide valuable forensic evidence of what activities took place.

Architectural Differences in Intrusion Prevention

Two approaches exist to locking down the air. In an integrated approach, all intrusion prevention functions, including rogue AP and uncontrolled client management, is included in the same system that provides wireless LAN access. The second architecture takes a distributed approach, where a separate dedicated system is used for wireless intrusion detection. Of these two approaches, the integrated approach is considered superior for the following reasons:

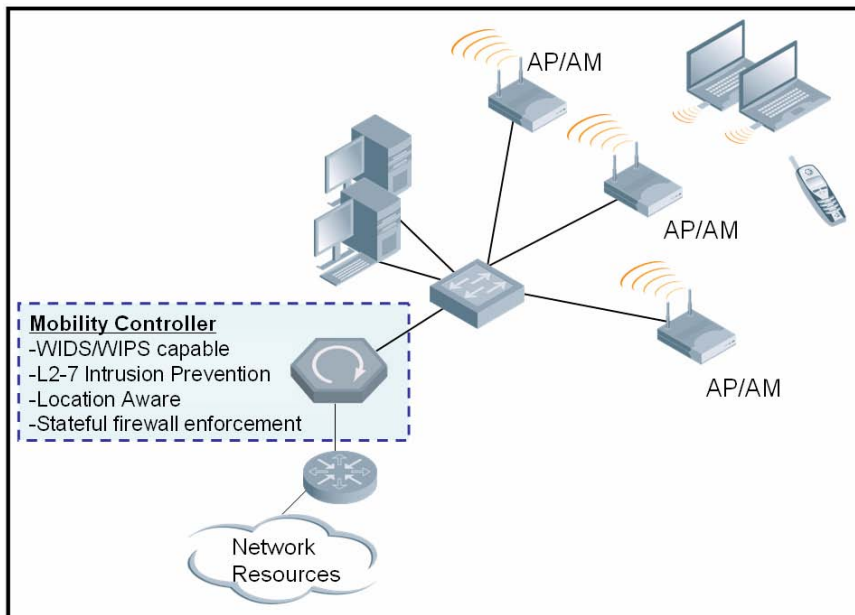
L2-7 Intrusion Prevention - The system monitoring for threats is also in the data path for wireless clients. This gives the system visibility from the RF layer up to the application layer as opposed to a distributed approach which has Layer2 scanning only. If a valid client is the source of an attack, the attack can be prevented rather than just detected and reported.

Enforcement techniques – For Intrusion Prevention systems that integrate a mobility controllers containing stateful firewall functionality, users can be automatically blacklisted if they violate a security policy. Overlay systems simply don't have this level of dynamic control over the client.

Integrated Management – With an integrated solution, the same mobility management console is used for WIDS, greatly simplifying overall management of the mobile network and reducing training costs. For organizations that are starting with a new wireless policy this is also compelling because the management system remains the same as wireless access is added.

Location – Integrated systems provide location context to the identified threat, showing exactly where that device is located on a building floor plan.

Reduced Capital Requirements - Access points used for wireless access are also sensors for the wireless intrusion prevention system. This saves on cabling and deployment costs since a single unit can do both jobs.



Integrated WIDS

IDS Compliance with DoD Directive (DoDD) 8100.2

The DoD's recent mandate on secure wireless access and Intrusion Detection Systems (IDS), DoD Directive (DoDD) 8100.2, released on June 2, 2006, provides guidance on the requirements for any wireless device that is connected to the DoD Global Information Grid and specifies that all such systems should be capable of delivering integrated IDS in addition to other security measures. A mobility architecture must, therefore, be certified by NIST as having achieved FIPS 140-2 Level 2 validation for 802.11i, as well as provide IDS, all in an integrated system.

In fact, 8100.2 goes a step further and requires that Wireless IDS be implemented for all DoD networks, whether wireless or wired LAN. These systems must continuously scan for and detect authorized and unauthorized devices 24 hours a day, seven days a week, as well as have the ability to sense a rogue device's location

Government agencies that are evaluating a network upgrade for use in Federal networks need to closely examine each solution to ensure they meet all elements of directive 8100.2

Keeping the Bad Guys Out: Authentication

Because radio waves travel outside their desired coverage area, it is critical to ensure that only valid, authorized users obtain access to wireless networks. The way to accomplish this is through authentication – a process that validates that a user is who he or she claims to be, and that the user is authorized to be on the network. Authentication typically consists of providing a username and password, or some other credential, to the mobility system. The mobility system checks this information against a database, such as Microsoft's Active Directory, and either grants or denies access based on the outcome. There are multiple ways to accomplish authentication, but without question the most secure method for wireless networks is the 802.1x protocol. This standard, widely implemented by equipment vendors and operating systems, provides a flexible framework for authenticating multiple types of users and devices through multiple types of credentials. 802.1x is incorporated into WPA (Wi-Fi Protected Access) versions 1 and 2 from the Wi-Fi Alliance, a certification found on all enterprise-grade wireless equipment and many consumer products as well.

8100.2 references specific requirements for authentication in a WLAN. Starting in 2007, new acquisitions of wireless technology must be 802.11i compliant and WPA2 Enterprise certified, implement 802.1X access control with EAP-TLS mutual authentication, and a configuration that ensures the exclusive use of FIPS 140-2 minimum overall Level 1 validated AES-CCMP communications. Properly-implemented authentication should include the following:

802.1x access control with EAP-TLS. PEAP, TTLS and TLS are 802.1x methods that include encrypted tunnels. Encrypted tunnels inside 802.1x function just like common SSL web sites used for ecommerce or sending usernames and passwords. Although an intruder can monitor the exchange over the air, data inside the encrypted tunnel cannot be intercepted. Do not use non-tunneled EAP types such as EAP-MD5 or LEAP.

Always perform mutual authentication to ensure that clients only communicate with valid networks. Authentication of the network is done using a digital certificate. Upon joining the network, the client is presented with a server-side digital certificate. If the certificate is trusted by the client, authentication will continue. If the certificate is not trusted, the process will stop. Do not disable server-side certificate checking on the client – if this is done, any access point can claim to be a valid enterprise access point and cause the client to provide login credentials. For this reason, more Secure EAP types that use a server side digital certificate such as PEAP, TTLS, and TLS should be used.

Implement a strong password policy. If a username and password are used to obtain access to the wireless network, it should not be easy for an attacker to guess usernames and passwords. The best form of wireless security uses one-time passwords such as SecurID or other token products. If one-time passwords are impractical to use, strong passwords consisting of eight

or more characters and a mixture of alphanumeric and special characters should be used. Most popular network operating systems provide policies to enforce strong password usage automatically.

Consider doing two-stage authentication, authenticating the computer as well as the user, if the client operating system allows this feature. For example, on a Microsoft Windows network, the computer can be authenticated as a valid domain member first, and then the user can authenticate as a valid user. If both steps do not take place, the wireless system can block access to the network.

AAA-FastConnect – Methods exist to speed connection times and greatly simplify the integration of secure WLANs with various back-end servers. With AAA FastConnect, a mobility controller can interoperate directly with an 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.

When it comes to authentication, architecture of the mobility system makes a difference. With a centralized system, a single device or small number of centralized devices act as the 802.1x authenticator, meaning that only a small number of devices need to be recognized by the authentication server. This results in greater scalability for the system, since less administrator time needs to be spent managing multitudes of entries in a RADIUS server. In addition, wireless roaming is enhanced with a centralized system, since a single centralized device holds all information about authentication, encryption, and mobility. When a user roams between multiple wireless APs, a centralized system can more quickly re-authenticate the client since that client was previously authenticated.

Hiding in Plain Sight: Encryption

After authentication, the second most important factor for solid wireless security is encryption. If an intruder cannot make use of intercepted data because it is encrypted, then there is no need to worry about how far the radio signals travel

Strong Encryption

FIPS 140-2 Level 2, the Federal Information Processing Standard for cryptography specifies encryption requirements for IT systems that are used for Sensitive But Unclassified (SBU) information. FIPS 140-2

identifies client-to-controller wireless encryption with AES-CCM (Advanced Encryption Standard-Counter Mode & CBC-MAC) as defined by the IEEE 802.11i standard. The Wi-Fi Alliance WPA2 certification includes AES-CCM as an encryption component, along with 802.1x authentication previously described. Thus, by installing the appropriate WPA2 wireless network, organizations can immediately get the benefits of strong authentication and encryption at the same time. In addition, with WPA2 the encryption keys are dynamic, meaning that each user on the network has a different encryption key that changes each time the user authenticates. This prevents one authorized user from intercepting the communications of another authorized user, and also makes the possibility of key “cracking” extremely remote.

Encryption Architecture

The architecture of the mobility system is extremely important to doing encryption properly and safely. Per DoDD 8100.2,

"encryption for unclassified data in transit via WLAN-enabled devices, systems, and technologies must be implemented end-to-end over an assured channel and be validated by NIST as meeting requirements per FIPS 140-2 Overall Level 1 at a minimum. If WLAN infrastructure devices which store keying information are used in public unprotected environments, then those products must meet FIPS 140-2 Overall Level 2".

The primary safety concern involves the passing of encryption keys across wired networks in a distributed system. Whether the system involves distributed “fat” access points or a controller with thin access points that implements encryption on the access points, encryption keys must be passed from a secured system (the authentication server) to the access point across a wired network. This introduces security risks to the network:

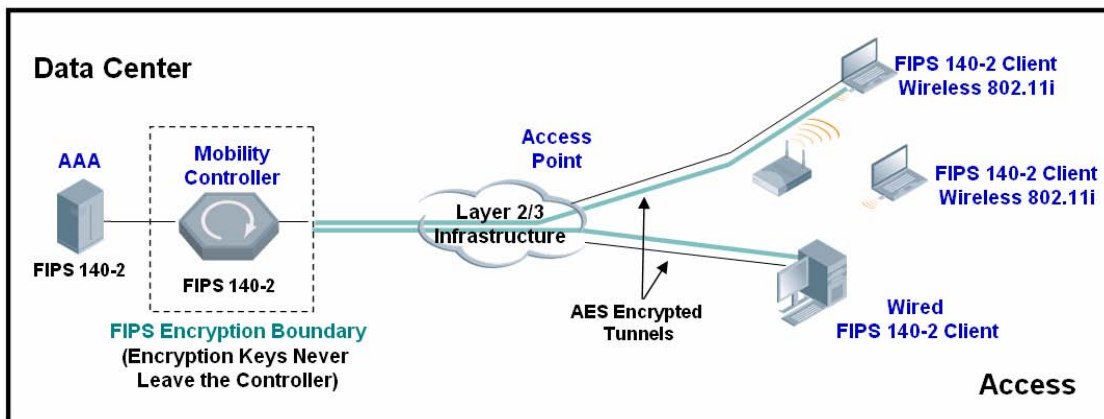
An intruder or malicious user on the wired network could intercept encryption keys and use them to wirelessly monitor private communication. Wireless makes an attractive means for such eavesdropping attacks, since it is impossible for the eavesdropper to be detected. Were the same attack conducted on a wired network, the use of ARP poisoning would give away the presence of the eavesdropper.

If a wireless access point in a distributed system performs encryption/decryption, then that access point must be a trusted device in the network infrastructure. But access points are not locked inside datacenters and wiring closets – they must be close to the users in order to provide wireless service. One attack targets “thin” access points that perform encryption: If the protocol running between a thin access point and a controller is understood, because the protocol is a published standard or through reverse-engineering, then an attacker can build a

software replica of an AP. The simulated AP will contact the controller and will be given a configuration, after which time it is treated as a peer of other APs in the system. If this simulated AP provides wireless service to users, it is now capable of performing a “man in the middle” attack where data can be deleted, added, or changed.

In a distributed system, FIPS validation of the solution is difficult because keys are stored in APs, meaning that vendors must go through the FIPS validation process for their APs *and* control channel.

Centralized architectures get around these risks by performing all encryption and decryption in a controller, where no encryption keys are distributed to the AP. A mobility controller is typically located in a physically secure area such as a data center, and is often in the same room with the authentication server. With the controller as the encryption boundary, and encryption keys never leaving the data center, there is no risk of interception by an unauthorized user. Access points in such a system are untrusted devices – an attacker building a software simulator of such an AP would obtain a channel and power assignment as part of configuration, but would have no extra privileges on the network. Even if a user authenticated through the imposter AP, no man in the middle attack would be possible since encryption is maintained all the way to the controller.



Wired and Wireless FIPS compliant encryption

As shown in the figure above, a centralized system follows the 8100.2 requirement by encrypting “end-to-end,” from the client to the controller. A centralized encryption and integrated security boundary obviates the need to FIPS validate the APs as well as the control channel between the APs and the mobility controller. This has allowed vendors such as Aruba Networks to FIPS validate an entire mobility solution and not just point products that, when put together, will not comply. For these reasons, a central encryption model is preferable to a distributed approach, which will typically

increase both the costs and time to market for using Commercial Off-The-Shelf (COTS) technology within the Federal marketplace.

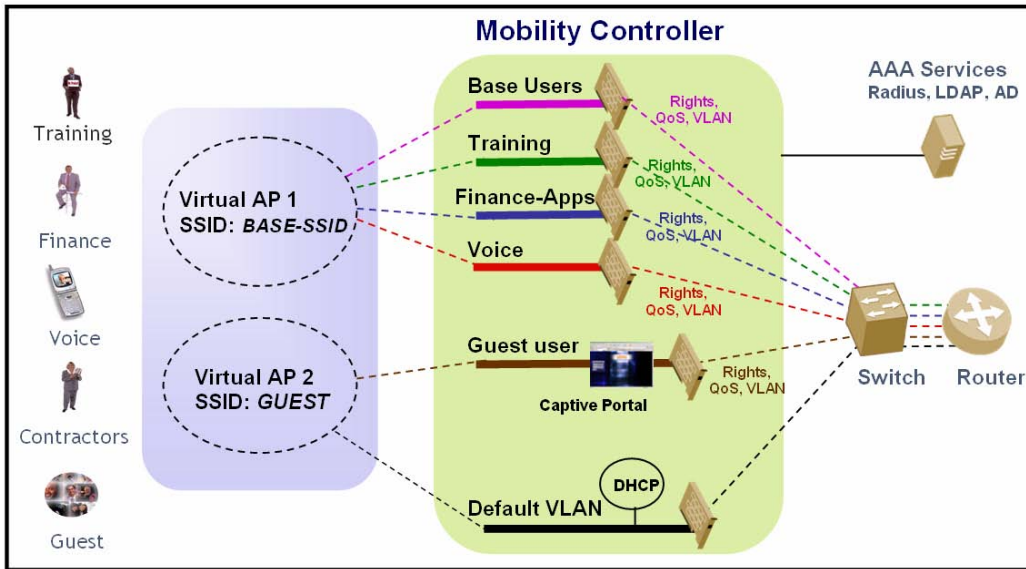
People Not Ports: Identity-based Security

Multiple types of users with multiple types of devices may be found on a typical mobile network. Mobile networks are unique when it comes to securing them, because mobile users and devices, by definition, do not connect to the network through a fixed port. For this reason, the network must identify every user and device that joins the network. Once this identity is known, custom security policies may be applied to the network so that only access appropriate to the business needs of the user or device is provided. A key concept is applying policies to people – or devices – rather than ports. In a mobile world, fixed ports are no longer a reliable indicator of the type of user connected.

Instead, identity must be used. Identity is learned through the authentication process, during which the device or the user provides some type of identifier, normally a username. Once identity is learned, it is mapped to the business role of that user or device. The business role may be determined through membership in specific departments or groups, security clearance, or the actual business position of a user. Role information is normally contained in the enterprise authentication system, such as Active Directory in a Microsoft Windows environment. Some examples of roles and their associated security requirements include:

- A member of a department training division, who needs access to the Internet and to internal web based databases. A member of a training division has no practical need to communicate with servers in, for example, Finance.**
- A temporary hospital staff user who should be disallowed access to the base-network from the hospital. Location-based access control must be enforced on the edge of the network.**
- An outside visitor, who needs access to specific applications on the Internet only during daytime business hours.**
- A legacy bar code scanner device that must send purchasing data as well as retrieve inventory updates. This device would communicate only with a specific server using specific protocols.**
- A converged device that supports both voice and data. Such a device should provide appropriate service and QoS requirements over a single SSID with or without WMM support.**
- A voice-over-WLAN handset that needs to communicate using the SIP protocol to a SIP gateway. The voice handset supports only WEP encryption and cannot perform a secure form of authentication.**

All these users have different privilege levels that must be enforced. In addition, data traffic from these users must be kept separate and isolated so that a user with lower privilege cannot intercept data from a more privileged user. Finally, devices with lower security standards, such as the voice handset, must not be permitted to open security holes in the network by nature of their lower security standards. Identity-based security is the mechanism through which all of these problems are solved.



Identity-Based Mobile Access Control

Architectural Requirements for Identity-based security

When implementing user centric security, the architecture of the mobility access system is important. Because user identity is the key factor when making access control decisions, it must be impossible for a user to assume the identity of another user. Three components of the system must be aware of each other and, ideally, integrated into the same system in order to provide the necessary level of security:

Authentication - supplies the system with identity information. Authentication must be done in a secure manner, such as through 802.1x.

Encryption - provides confidentiality and integrity of data. When using WPA2 for wireless access, encryption provides an extra benefit for identity-based security: Because the encryption key itself is derived during authentication, data that decrypts successfully can be assured of coming from the authenticated user and only the authenticated user.

Authorization - enforces identity-based policies. When the system knows who the user is (through authentication), and knows that received data came from that user (through encryption), it can then reliably perform identity-based authorization and policy enforcement.

In a distributed system, as shown in the diagram below, authorization is performed by an external firewall. The firewall is not aware of user identity, because it does not perform authentication. Additionally, the firewall does not perform encryption and decryption of user data, so it cannot be sure that data claiming to come from a user actually came from that user. This makes the external firewall unreliable for performing identity-based security. The firewall applies rules to IP addresses rather than to users – this makes it suitable for macro-level global policy enforcement, such as enforcing policies that apply to all wireless users. But without user identity and assurance of non-tampering with user data, it cannot perform identity-based security.



A centralized system, in contrast, implements all three functions described above in a single system. Because these functions are integrated and aware of each other, identity-based security can be provided. Even an authenticated user who tries to fool the system by injecting crafted packets or changing an IP address cannot gain excess privilege on the network.



Planning for Global Mobility: Remote Access

On a fully leveraged mobile network, users are not only mobile within a headquarters location. They also move between different sites, telecommute from home, and work in off-site locations such as partner offices, hotels, and public hotspots. An organization's security posture cannot weaken just because users are not at the corporate headquarters – it must be uniform throughout the network regardless of how or where it is accessed. This is especially important for emergency preparedness providing a high-level of network access and wireless continuity to users during quick-hitting, spontaneous emergency situations, like natural disasters and terrorist attacks.

Uniform Authentication Infrastructure

The first step towards effective global security is establishing a uniform authentication infrastructure. Wherever a user travels, the user should be required to authenticate to the network. But users cannot be forced to manage multiple user identities, accounts, and passwords. A single set of access credentials should provide for authentication at any location – ideally this is the same set of credentials used to login to the user's own workstation. Authentication servers should be set up to coordinate with one another by replicating user information. Alternatively, the mobility controller systems should be set up to understand domain names, realms, and other regional identifiers so that authentication requests can be routed to the correct set of authentication servers. Using this principle, a user can travel to any enterprise location in the world and be granted access to appropriate network resources.

Consistent Access Methodology

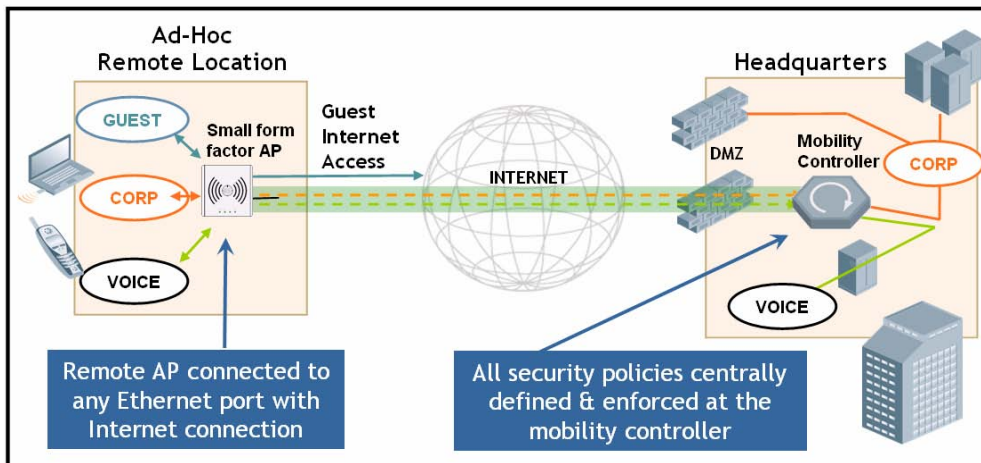
Second, the access method needs to be consistent wherever the user roams. Users do not want to reconfigure systems as they move from headquarters to branch offices to their homes. This means that the same wireless SSID (Service Set Identifier) should be present in all locations with the same authentication and encryption policies present. All locations should use the same authentication infrastructure. A user should be able to start an email application at the corporate office, put a laptop in sleep mode, go home, start up the laptop again, and have the email application continue to work without intervention from the user and without the user needing to start a separate VPN client. When this happens, support help desk calls go down dramatically.

Remote Voice over WLAN Access

Third, the solution needs to take voice mobility into account. Many organizations are evaluating voice over WLAN (VoWLAN) technology today with expected large-scale deployments in the near future. One of the key benefits of this technology will be the ability to use it wherever wireless LAN service is available. When employees travel to remote locations, voice mobility will allow their VoWLAN handset to continue operation just as it would in the normal work location. Specifically, the mobile network

infrastructure must provide quality of service control, secure transport of voice traffic back to a corporate telephony server, and consistent authentication and encryption schemes throughout the global network.

The mobility system must be architected properly to support global mobility. A traditional distributed system normally treated each office location as a separate network, possibly with different authentication services, different SSIDs, and different security policies. Telecommuters and traveling employees were serviced using Internet-based VPNs, with VPN client software installed on laptop computers. Notably, any device without support for VPN client software could not join the network – this includes voice handsets, some PDAs, and any client operating system not supported by the VPN vendor. With a centralized architecture, global mobility is treated just like intra-office mobility. Wireless access points are placed in any location where wireless access is desired and a wired internet connection is available – the corporate office, branch offices, retail outlets, and home offices. Traveling users may carry small “personal” access points with them and connect these to ubiquitous wired Ethernet ports commonly available wherever business travelers may be found. The Access Point configures on its own (zero-touch) immediately, tunneling back to a central mobility controller. Because the architecture is centralized, the access points are not responsible for service delivery, security enforcement, or authentication. Instead, a network of mobility controllers actually provides the network services, while the access points serve as secure wireless portals to make the connection to the mobility controller. In a centralized architecture, there is no need for VPN clients. Instead, WPA2 serves as the common security framework for global mobility. All access to the network is authenticated using 802.1x and encrypted using AES.



No-Touch Remote Office Setup with Remote AP

Defensive Networks: Knocking Out Malware

The old model of data networks was a number of PCs connected to an office LAN with an Internet connection through a firewall. Attached to the firewall might have been an intrusion detection system, a VPN concentrator, and perhaps service appliances such as an anti-virus gateway or a web proxy. These devices formed the security perimeter around a company's information resources. Today, mobility has become so prevalent that the security perimeter is rendered ineffective. The sale of laptop computers has now surpassed the sale of desktop computers in most organizations, meaning that more and more people are being equipped with mobile computing. These laptops leave the office with its associated perimeter protection on a regular basis, many times connecting to the Internet through unprotected and untrusted networks. When the user returns to the office and connects the laptop to the network, any malicious software that found its way onto that laptop is now inside the firewall and is free to spread to other unprotected devices in the network.

This problem is not unique to wireless – it is caused by mobility in general. However, the prevalence of wireless hotspots makes the problem appear more often. Thus, addressing client security is a necessary component of any wireless security policy. Client security can be addressed in two major ways:

Client integrity assessment. With this method, the client is monitored for compliance with various enterprise policies. One policy may be that anti-virus software must be installed, enabled, updated within the past three days, and the system scanned within the past week. Another policy must be that personal firewall software is installed and enabled. When a system attempts to join the wireless network, the integrity agent signals the current policy compliance state to the network. If the device is out of compliance, it is quarantined from the network and

optionally redirected to a remediation server that automatically forces updates to bring the system back into compliance.

Network-based services. With this method, all client devices are quarantined from each other and all client traffic is inspected and passed through network-based service appliances such as anti-virus gateways and intrusion detection systems. This technique is particularly useful for client devices that cannot or do not have client integrity agents installed. Examples of such devices would include barcode scanners, PDAs, voice handsets, certain client operating systems, and laptop computers belonging to visitors and contractors.

Strategies for Guest Access

Guest access is often a requirement for a wireless network to provide Internet access for visitors. These guests are better able to carry out their jobs when instant access to timely information is available. But guest access must be provided in a way that does not pose a security risk to the network, and must also not allow unauthorized persons to steal network access. Controlled guest access increases network security, since guests with authorized access will not cause a security breach by plugging their laptop into an internal network port. There are two pieces to guest access: authentication and policy control.

Authentication forces a guest user to prove to the system that he or she is authorized to use the network. This prevents outsiders, such as “wardrivers”, from using the organization’s network. There are several popular strategies for providing guest authentication:

Open access. Guest access is available to anyone who can receive the wireless signal. This is often used in isolated buildings on large plots of land where wireless signals would not easily reach an outsider. It is also used by some organizations that are not concerned with outsiders using their Internet access. In general, it is not a recommended strategy from a security perspective.

Common guest password. A guest network is set up with a single username and password for guest access. The guest information is posted on conference room walls or otherwise made available. Visitors needing Internet access will be given this username and password, and will use it to login to a web-based portal system. This is a good option for a low-maintenance guest system, but it does not provide any individual accountability for activities on the network. Many organizations are willing to accept this trade-off in exchange for simplicity.

Provisioned guest access. With this scheme, each guest user is given a unique time limited username and password. This may be done by a receptionist when the guest checks in, or may be requested ahead of time by an employee through an automated system. This method

provides the best security and accountability, but is also the most work to set up. Once set up, however, this system is for the most part self-maintaining and does not require the ongoing involvement of IT resources.

Whatever guest authentication method is chosen should be implemented globally so that employees and visitors have a common experience at any work location.

Policy control for guest users manages what resources the guest is able to access, when they are able to use the network, and what quality of service they receive from the network. Of these, the most important is access control. Guest users must be prevented from accessing internal corporate resources while still being provided with Internet access. For liability reasons, it is also desirable to restrict what protocols and even what destinations a guest user may communicate with. For corporate guest users, the only protocols needed may be HTTP for web browsing, POP3 for email, and IPSEC/PPTP for VPN access. Outgoing email using SMTP should be blocked to prevent the network from becoming a spam relay, and peer-to-peer file sharing should also be blocked to limit legal liability. In addition to protocol control, guest traffic may be limited by time of day so that it is not available outside of normal working hours. Guest traffic may also be bandwidth limited so that guest users cannot consume excess amounts of network capacity.

The network architecture must provide identity-based security in order for guest access to be implemented safely and effectively. Without identity-based security, there is potential for guest users to communicate with internal network resources, since tight access control cannot be performed. With identity-based security, guest users are placed into a guest role with an associated guest access policy, while employees are placed into an appropriate internal role. While the two classes of users share the same wireless infrastructure, no crossover is possible.

Summary

Wireless technology is mandatory to support new mobility requirements in the DoD. The technology has been rapidly changing over the past several years with major innovations in architectures and the creation of new compliance requirements. Furthermore, DoD directives and NIST requirements have quickly constrained the choice of vendors and mobility architectures. This has led to confusion regarding best practices for setting security policies and deployment. This white paper, while not providing exhaustive coverage of all options, has provided current best practices along with a discussion of how these practices can be implemented using different wireless architectures. The best security approach for mobility must meet all standards requirements and will require layering the following security functions:

-
- **Wireless intrusion protection**
 - **Authentication**
 - **Encryption**
 - **Access control**
 - **Client security**

Implementing these best practices will protect against unauthorized and uncontrolled wireless as well as the malicious hacker determined to gain access to the network. By implementing these best practices in a global wireless policy, mobile wireless networks will provide strong security with the productivity benefits brought about by mobility.

About Aruba Networks

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. Aruba Networks, BlueScanner and RFprotect are trademarks of Aruba Networks, Inc. All other trademarks or registered trademarks are the property of their respective holders. All rights reserved. Specifications are subject to change without notice.

WP_GVSECPOL_US_071217



1322 Crossman Ave. Sunnyvale, CA 94089-1113
Tel. +1.408.227.4500 | Fax. +1.408.227.4550 | info@arubanetworks.com
<http://www.arubanetworks.com>