

Building Global Security Policy for Wireless LANs

Jon Green, CISSP
Aruba Networks

ARUBA[®]
networks

www.arubanetworks.com

Building Global Security Policy for Wireless LANs

Jon Green, CISSP
Aruba Networks

Table of Contents

Introduction	1
Chapter 1 Lessons Learned: What Doesn't Work	2
Chapter 2 Architectures for Mobility	5
Chapter 3 Locking the Air	6
Chapter 4 Keeping the Bad Guys Out: Authentication	8
Chapter 5 Hiding in Plain Sight: Encryption	10
Chapter 6 People, Not Ports: Identity-based Security	12
Chapter 7 Planning for Global Mobility: Remote Access	16
Chapter 8 Defensive Networks: Knocking Out Malware	18
Chapter 9 Strategies for Guest Access	20
Chapter 10 Putting It All Together: Sample Security Policy	22
Summary	30
About Aruba Networks, Inc.	31

Introduction

As wireless devices become more and more common in today's enterprise networks, now is a good time for CIOs and IT managers to plan their strategy for overall control, deployment, and management of this important technology. Security is one component of that strategy, and it is a big one. While a properly implemented wireless security policy makes wireless more secure than wired networks, 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 connectivity model where users connect over wireless networks wherever they go – at the corporate office, working from home, or traveling on the road. Mobility, including wireless technology, has the potential to expose corporate networks to intruders, leak sensitive data, and subject the enterprise network to virus and worm outbreaks. Proper planning avoids these issues without necessarily costing a lot of money. This white paper first explains wireless security techniques that have failed, and then provides clear recommendations for building effective security policies. Also presented are pros and cons of two architectural approaches to wireless security – a centralized approach and a distributed approach. At the end of this document, a sample wireless security policy is provided.

Chapter 1

Lessons Learned: What Doesn't Work

A number of “Best Practice” guides for securing wireless networks have been written over time. While many of the more recent ones are useful, a number of them advocate older techniques that have been proven ineffective.

Policies Without Enforcement

A written IT security policy is necessary in any size organization, but it is meaningless without a way to check compliance. Too many companies write a security policy banning all wireless devices, then fail to monitor for their use. Users demand mobility, and experience shows that if wireless networks are not provided by the IT department, users will install consumer-grade equipment themselves. Typically this consumer-grade equipment has no security turned on by default, and most users do not bother with additional configuration steps to turn on even basic security. These “rogue” access points (APs) effectively open an organization’s network to anyone in the parking lot.

Some organizations establish “no wireless” policies and do periodically scan for unauthorized equipment. However, this eats up valuable personnel time as a network administrator walks through the building with a laptop or other wireless scanner. If an AP is detected, the administrator must then spend additional time to determine if that AP is inside the building, or if it belongs to another nearby business.

RF Engineering

A common question heard from organizations looking to deploy wireless is, “How do I ensure that the wireless signal doesn’t travel outside the building?” Some security analysts recommend using special directional antennas to accomplish this, or recommend using “decoy” access points with antennas pointed outside the building as a way to defeat would-be intruders. Both techniques are costly, complex, and do not work. Radio signals are invisible and travel in unpredictable ways after bouncing off reflective surfaces such as file cabinets and whiteboards. In addition, an attacker can use a high-gain directional antenna to transmit and receive signals from far away, even when a standard laptop wireless card does not detect a usable signal. Wireless networks should be installed with the assumption that anyone can be within radio range of the network, and security should be adjusted appropriately.

SSID Cloaking

Some APs offer the ability to “hide” the broadcast of the Service Set Identifier (SSID), also known as the “network name.” Some wireless security best practice guides in the past have advocated doing this, with the idea that the SSID can be used as a password. In theory, if an attacker doesn't know the SSID name in advance, he can't connect to the network. In reality, it is simple to learn the SSID by simply monitoring the normal process of an authorized client joining the network. SSID cloaking is not harmful, but it should never be treated as a security technique.

MAC Address Filtering

A common wireless practice for consumer-grade equipment is to turn on MAC address filtering. With this feature, only computers on the “approved” list are allowed on the wireless LAN. Unfortunately, MAC address filtering is ineffective because it is trivial for an attacker to impersonate a valid computer by changing the MAC address of his or her computer. MAC address filtering also does not scale in enterprise networks, since the address database must be updated each time a computer is bought, replaced, or eliminated.

WEP

WEP (Wired Equivalent Privacy) is the original wireless encryption standard provided for 802.11 wireless LANs. WEP is widely recognized as being ineffective as an encryption protocol on multiple fronts. Using modern attack tools, WEP can be cracked in one minute or less, rendering the interior network open to intruders. Two types of WEP networks may be deployed: static WEP with pre-configured keys, and dynamic WEP with 802.1x authentication. While dynamic WEP provides scalability benefits in an enterprise setting, both forms of WEP are equally weak and are unsuitable for use today. Where application needs require WEP to be used, network access should be extremely restricted using firewall policies to allow the minimum access required.

Recently some vendors have begun providing so-called “WEP cloaking” or “WEP shielding” products. These are designed to be used in conjunction with a WEP network to defeat attackers by injecting “decoy” traffic into the air that confuses WEP cracking tools, thus making WEP safer for use. Attack tools were quickly modified to defeat these products, and thus they do not measurably improve security of WEP networks.

Cisco LEAP

Cisco invented LEAP (Lightweight Extensible Authentication Protocol) as a way to provide authentication and dynamic encryption keying for wireless networks before standards existed to provide those services. While LEAP served a purpose in the past, it has now been broken and can be exploited by attackers to break into a wireless network. The use of strong password policies may make such attacks more difficult, but ultimately this protocol should be retired and replaced with stronger security standards that have been subject to widespread peer review. It should be noted that the proprietary replacement for LEAP, known as EAP-FAST, also has known security failures and should be avoided in preference to standards-based protocols such as PEAP and EAP-TLS.

What Doesn't Work: Conclusion

Today there are modern approaches to wireless security that render the above techniques obsolete. Rather than sacrifice security with stopgap measures that provide only partial protection, organizations can deploy standards-based technology that provides solid protection for mobile networks.

Chapter 2

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 an external 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 re-visited in each section below to provide comparison and contrast between the capabilities of each.

Chapter 3

Locking the Air

The first step in any wireless security policy is to lock down the radio spectrum against threats. This step must be taken even if a wireless network is not actually deployed, to prevent against uncontrolled wireless devices that may be brought in. Once a wireless network has been deployed, monitoring for attacks against that network becomes an additional need.

Rogue APs

The very existence of wireless technology is a threat to security of the wired network. Employee 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 an employee to install one of these “rogue” APs in an office or cubicle. Connected to the corporate wired network, rogue APs become instant portals into 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 branch and remote offices.

Uncontrolled clients

A second category of threats to the enterprise 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 corporate 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 malicious attacks. 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 wireless LAN connectivity problems when the actual problem is an attacker. For companies that intend to prosecute attackers under the law, wireless intrusion prevention systems provide valuable forensic evidence of what activities took place.

Wireless Intrusion Prevention

The technology used to monitor and prevent these types of threats is called a Wireless Intrusion Detection System (WIDS) or Wireless Intrusion Prevention System (WIPS). Two architectural approaches exist to locking the air. In a centralized architecture, all intrusion prevention functions, including rogue AP and uncontrolled client management, are included in the same system providing WLAN access. With the distributed approach, a separate dedicated system is used for wireless intrusion detection. Of these, the centralized integrated approach is considered more cost-effective and secure for the following reasons:

- 1) 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.
- 2) 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. If a valid client is the source of an attack, the attack can be prevented rather than just detected and reported.
- 3) Where WIPS technology is deployed to enforce a “no-wireless” policy, either a stand-alone distributed system or a centralized integrated system can do the job. However, with the centralized integrated system, the same equipment can later be used to provide WIPS + WLAN access if the organization decides to change the no-wireless policy. Thus such a system should be given serious consideration even if WLAN access is not currently planned.

Chapter 4

Keeping the Bad Guys Out: Authentication

Because radio waves travel outside their desired coverage area, it is critical to ensure only valid, authorized users obtain access to wireless networks. This is accomplished with authentication—a process that validates that a user is who he claims to be and 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 grants or denies access based on the outcome. There are multiple ways to accomplish authentication, but 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 the Wi-Fi Alliance’s WPA (Wi-Fi Protected Access) versions 1 and 2, a certification found on all enterprise-grade wireless equipment as well as many consumer products.

Authentication must be done right. Done incorrectly, it can be the single biggest flaw in wireless security. Recommendations for properly-implemented authentication include:

- Use 802.1x EAP methods that include encrypted tunnels. These include PEAP, TTLS, and TLS. Encrypted tunnels inside 802.1x function just like common SSL web sites used for e-commerce or sending 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—accomplished by way of a digital certificate—to ensure that clients only communicate with valid networks. 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 valid and cause the client to provide login credentials. One EAP type, EAP-FAST, does not use a server-side digital certificate and thus does not perform mutual authentication unless extra labor-intensive steps are taken. For this reason, EAP-FAST should be avoided in favor of more secure EAP types such as PEAP, TTLS, and TLS.

- Lock down 802.1x client settings. Many 802.1x supplicants provide options for validating server certificates, for trusting only specific authentication servers, for trusting only specific certificate authorities (CAs), and for allowing the end user to add new trusted servers and certificate authorities. To achieve the best security, always use the most restrictive settings. The server certificate must always be validated. The client should trust only a specific set of certificate authorities—and for the strongest security, these should be well-run internal CAs rather than public CAs. The client should only authenticate against specific RADIUS servers. Finally, the end user should not be permitted to allow new trusted authentication servers or CAs.
- Implement a strong password policy. It should not be easy for an attacker to guess a username and password used to obtain access to the wireless network. The best form of wireless security uses one-time passwords such as SecurID or other token products. If one-time passwords are impractical, use strong passwords consisting of eight or more characters and a mixture of alphanumeric and special characters. 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.

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 acts 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 system scalability, 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.

Chapter 5

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. The state of the art for wireless encryption is AES-CCMP (Advanced Encryption Standard-Counter Mode & CBC-MAC Protocol) as defined by the IEEE 802.11i standard. The Wi-Fi Alliance WPA2 certification includes AES-CCMP as an encryption component, along with 802.1x authentication previously described. Thus, by installing a 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.

The architecture of the mobility system is extremely important to doing encryption properly and safely. 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:

- 1) An intruder or malicious employee on the wired network could intercept encryption keys and use them to wirelessly monitor other employee’s 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. One recent case where wireless eavesdropping was used involved a large company about to conduct a major acquisition of another company. An employee was able to intercept email communication of a senior executive over a wireless network, and the employee then used this information for financial gain. The malicious employee obtained the wireless encryption key by first monitoring an access point from the wired side of the network.

- 2) 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. The situation becomes worse when fast-roaming schemes such as Proactive Key Caching are used, since encryption keys from one access point are pushed out to all other APs in the system in order to speed up roaming. If one of the APs is the intruder’s simulator, it now has encryption keys for the entire network.

Centralized architectures get around these risks by performing all encryption and decryption in a controller. This 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 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.

Chapter 6

People, Not Ports: Identity-based Security

Multiple types of users with multiple types of devices may be found on wireless networks. 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. If this sounds like something the industry has been promising through Network Admission Control (NAC), you would be correct. Identity-based security is the first and most important component of NAC, and wireless actually has a unique advantage over wired networks in implementing NAC, since authentication is a native part of wireless.

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 an enterprise user database, such as Active Directory in a Microsoft Windows environment. Some examples of roles and their associated security requirements include:

- A member of the sales department, who needs access to the Internet and to internal web-based sales databases. A member of the sales department has no business need to communicate with servers in, for example, the human resources department.
- An outside visitor, who needs only access to specific applications on the Internet only during daytime business hours.
- A POS (Point of Sale) handheld device in a retail environment that must send credit card data as well as download inventory and price updates. This device would communicate only with a specific server using specific protocols.

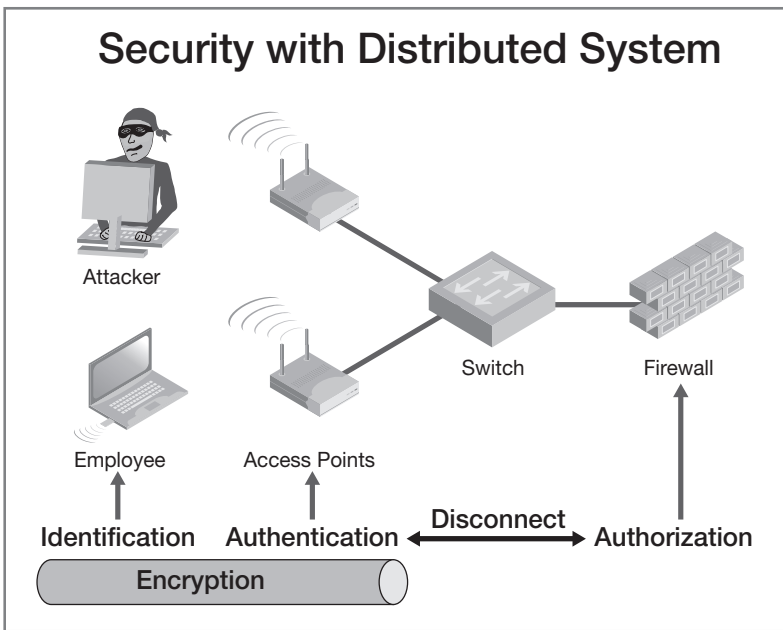
- A public PC-based kiosk for use by the general public. This device would be permitted to do web browsing, but would be denied all other network access.
- A voice-over-WLAN handset that needs to communicate using the SIP protocol to a SIP gateway. The voice handset supports only WEP encryption and cannot perform a secure form of authentication.

All these users and devices 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.

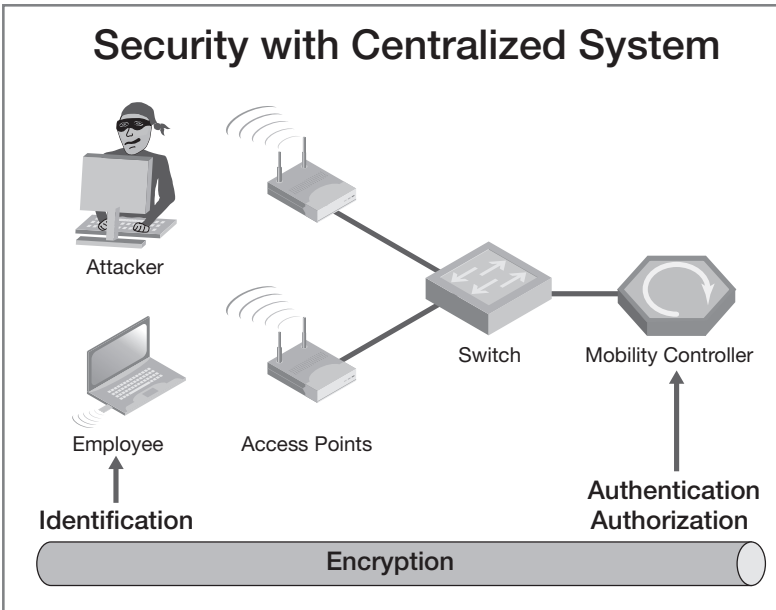
When implementing identity-based security, the architecture of the mobility access system is important. Because user identity is the key factor when making access control decisions, it must be impossible for a user to assume the identity of another user. Three components of the system must be aware of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

- 1) Authentication, which supplies the system with identity information. Authentication must be done in a secure manner, such as through 802.1x.
- 2) Encryption, which 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.
- 3) Authorization, which 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.



Think of this difference using the analogy of an airport. A distributed system is like a domestic airport, where typically a check of your identity is made only once as you pass through a security checkpoint. There, your identification (authentication) is matched against a boarding pass (authorization). But nothing stops you from printing a fake boarding pass with a name that matches your identification; this is possible because authentication and authorization are not linked together, and there is no way to validate the authorization token. In addition, once inside the security checkpoint, you are free to exchange boarding passes with anyone else, and board their flight instead of your own. Contrast this with an international airport, where your identification is checked as you board the aircraft – this is also the time when your authorization token is checked against a database. Here, your identification provides your name, your boarding pass must match your name, and additionally the boarding pass will be checked against a computer system to make sure it is valid. Linking these security steps together at the same point provides a much stronger security system, and the same is true with centralized wireless architectures.

Chapter 7

Planning for Global Mobility: Remote Access

Users are not only mobile within a corporate headquarters location. They also move between different office 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 wherever users access the network.

The first step towards effective global security is establishing a uniform authentication infrastructure – this gets into the realm of single sign-on and identity federation. 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 network 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.

Second, the access method needs to be consistent wherever the user roams. Users do not want to reconfigure systems as they move from the corporate office 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.

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 sometime in 2009 or 2010. 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 – 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. 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.

Chapter 8

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 in the enterprise space has now surpassed the sale of desktop computers, meaning that more and more employees are being equipped with mobile computing. These laptops leave the company 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 corporate 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:

- 1) Client integrity control. This is another element of Network Admission Control (NAC) where agent software is loaded on each client device, either as a permanent software install or as a temporary “dissolvable” agent that terminates after running its scan. The agent software monitors the system 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 might 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.

- 2) Network-based services. This is a third element of NAC, where 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. Depending on equipment capabilities, it may be possible for only certain types of traffic from certain clients to be passed through scanning appliances – for example, HTTP traffic may be scanned for malware while SIP traffic may not.

At the same time, patch management of the client device is also critical when it comes to network driver software. A number of well-publicized attacks against popular operating systems were performed by sending malformed data directly to a wireless workstation, where flaws in the driver software for the wireless hardware allowed buffer overflows and in some cases, remote code execution. The same attacks have been carried out against wired systems that were on public networks. Client devices should always be updated with the latest driver software, and particular attention should be paid to security related announcements from network device vendors.

Chapter 9

Strategies for Guest Access

Guest access is often one of the first requirements for a wireless network. Many companies want to provide Internet access for visitors in conference rooms, lobbies, and other meeting areas. Visitors are better able to carry out their work when instant access to timely business information is available. But guest access must be provided in a way that does not pose a security risk to the corporate 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 Internet access as a free connection. There are several popular strategies for providing guest authentication:

- 1) 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 companies who are not concerned with outsiders using their Internet access. In general, it is not a recommended strategy from a security perspective.
- 2) 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 to employees. 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.
- 3) 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.

Chapter 10

Putting It All Together: Sample Security Policy

The following is a sample security policy that ties together recommendations and best practices discussed in this paper. It may be easily cut and pasted from this document and adapted to your organization's needs

1 Purpose

This policy establishes standards that must be met when wireless communications equipment is connected to <Company Name> networks. The policy prohibits access to <Company Name> networks via unsecured wireless communication mechanisms. Only wireless systems that meet the criteria of this policy or have been granted an exclusive waiver by Information Security are approved for connectivity to <Company Name>'s networks.

2 Scope

This policy covers all wireless data communication devices (e.g., personal computers, cellular phones, PDAs, etc.) connected to any of <Company Name>'s internal networks. This includes any form of wireless communication device capable of transmitting packet data. Wireless devices and/or networks without any connectivity to <Company Name>'s networks do not fall under the purview of this policy.

3 Policy

3.1 Approved equipment

- 3.1.1 All wireless LAN access must use corporate-approved products and security configurations.

3.2 Monitoring of uncontrolled wireless devices

- 3.2.1 All company locations where permanent data networks are installed will be equipped with sensors and systems to automatically detect, classify, and disrupt communication with unapproved wireless access points.
- 3.2.2 All company locations where permanent data networks are installed will be equipped with sensors and systems to automatically detect the presence of wireless devices

forming a connection between the network and any wireless network. This would include laptops that are serving as a bridge between wired and wireless networks.

- 3.2.3 In company locations where wireless LAN access has been deployed, wireless intrusion detection systems will also be deployed to monitor for attacks against the wireless network. The wireless intrusion detection system shall be integrated with the wireless LAN access system whenever possible.

3.3 Authentication of wireless clients

- 3.3.1 All access to wireless networks must be authenticated.
- 3.3.2 The Company's existing strong password policy must be followed for access to wireless networks.
- 3.3.3 The strongest form of wireless authentication permitted by the client device must be used. For the majority of devices and operating systems, WPA or WPA2 with 802.1x/EAP-PEAP must be used. WPA2 is preferred wherever possible.
- 3.3.4 Where 802.1x authentication is used, mutual authentication must be performed. Client devices must validate that digital certificates presented by the authentication server are trusted and valid. Under no circumstances may clients disable validation of server certificates and blindly trust any certificate presented. EAP methods that do not support certificate-based mutual authentication may not be used.
- 3.3.5 EAP methods that exchange authentication credentials outside of encrypted tunnels may not be used. These methods include EAP-MD5 and LEAP.
- 3.3.6 When legacy devices that do not support WPA or WPA2 must be used on a wireless network, they will be isolated from all other wireless devices and will be restricted to the minimum required network access. Violations of the configured rules, indicating that an intrusion has taken place, must cause the device to be immediately disconnected and blocked from the network.
- 3.3.7 Any Company user with an account in a Company user database shall be able to authenticate at any Company location where wireless access is present.

3.4 Encryption

- 3.4.1 All wireless communication between Company devices and Company networks must be encrypted. Wireless networks providing only Internet access for guest users are exempted from this requirement.
- 3.4.2 The strongest form of wireless encryption permitted by the client device must be used. For the majority of devices and operating systems, WPA using TKIP encryption or WPA2 using AES-CCM encryption must be used. WPA2 with AES-CCM is preferred wherever possible.
- 3.4.3 Client devices that do not support WPA or WPA2 should be secured using VPN technology such as IPSEC where allowed by the client device.
- 3.4.4 The use of WEP requires a waiver from Information Security. Client devices that require the use of WEP must be isolated from all other wireless devices and will be restricted to the minimum required network access. Violations of the configured rules, indicating that an intrusion has taken place, must cause the device to be immediately disconnected and blocked from the network.

3.5 Access control policies

- 3.5.1 Access to corporate network resources through wireless networks should be restricted based on the business role of the user. Unnecessary protocols should be blocked, as should access to portions of the network with which the user has no need to communicate.
- 3.5.2 Access control enforcement shall be based on the user's authenticated identity, rather than a generic IP address block. This is also known as "identity-based security."
- 3.5.3 The access control system must be implemented in such a way that a malicious inside user is unable to bypass or circumvent access control rules.
- 3.5.4 Access control rules must use stateful packet inspection as the underlying technology.

3.6 Remote wireless access

- 3.6.1 Telecommuting employees working from remote locations must be provided with the same wireless standards supported in corporate offices.
- 3.6.2 Employees should be discouraged from connecting Company computers though consumer type wireless equipment while at home in lieu of Company-provided equipment.

3.7 Client security standards

- 3.7.1 Where supported by the client operating system, the wireless network will perform checks for minimum client security standards (client integrity checking) before granting access to the Company network. Specifically:
 - 3.7.1.1. All wireless clients must run Company approved anti-virus software that has been updated and maintained in accordance with the Company's anti-virus software policy.
 - 3.7.1.2. All wireless clients must run host-based firewall software in accordance with the Company's host security policy.
 - 3.7.1.3. All wireless clients must have security-related operating system patches applied that have been deemed "critical" in accordance with the Company's host security policy.
 - 3.7.1.4. All wireless clients must be installed with Company-standard wireless driver software.
- 3.7.2 Clients not conforming with minimum security standards will be placed into a quarantine condition and automatically remediated.
- 3.7.3 Client operating systems that do not support client integrity checking will be given restricted access to the network according to business requirements.

3.8 Wireless guest access

- 3.8.1 Wireless guest access will be available at all facilities where wireless access has been deployed.
- 3.8.2 All wireless guest access will be authenticated through a web-based authentication system.
- 3.8.3 A single username/password combination will be assigned for all guest access. The password for guest access will be changed monthly and distributed to local facility managers.
- 3.8.4 Wireless guest access is available from the hours of 7:00 until 20:00 local time.
- 3.8.5 Wireless guest access is bandwidth limited to 2Mb/s per user.
- 3.8.6 Guest access will be restricted to the following network protocols:
 - HTTP (TCP port 80)
 - HTTPS (TCP port 443)
 - POP3 (TCP port 110)
 - IKE (UDP port 500)
 - IPSEC ESP (IP protocol 50)
 - PPTP (TCP port 1723)
 - GRE (IP protocol 47)
 - DHCP (UDP ports 67-68)
 - DNS (UDP port 53)
 - ICMP (IP protocol 1)

4 Definitions

Terms	Definitions
802.11	A set of Wireless LAN/WLAN standards developed by the IEEE LAN/MAN standards committee (IEEE 802). Also commonly referred to as “Wi-Fi.”
802.11i	An amendment to the IEEE 802.11 standard specifying security mechanisms for wireless networks.
802.1x	A framework for link-layer authentication specified by the IEEE.
AES-CCMP	Advanced Encryption Standard-Counter with CBC-MAC Protocol. A wireless encryption protocol specified by IEEE 802.11i. Currently regarded as the strongest form of wireless encryption.
EAP	Extensible Authentication Protocol. A series of authentication methods used inside 802.1x to achieve wireless authentication.
IEEE	Institute of Electrical and Electronics Engineers. An international professional organization dedicated to the advancement of technology related to electricity. The IEEE is one of the main standards bodies associated with networking technology.
IETF	Internet Engineering Task Force. Develops and promotes Internet standards, in particular those of the TCP/IP protocol suite.
IPSEC	IP Security. An IETF standard for protecting IP communication by encrypting or authenticating all packets.

Terms	Definitions
LEAP	Lightweight Extensible Authentication Protocol. A proprietary protocol supported by Cisco Systems that acts as an EAP method within 802.1x. LEAP was proven insecure in 2003 and does not comply with current security standards.
PEAP	Protected Extensible Authentication Protocol. A tunneled EAP method that uses a server-side digital certificate for server authentication and a username/password for client authentication.
Stateful Packet Inspection	A filtering or firewall technology that keeps track of the state of network connections, such as TCP streams, traveling across it. Only packets which match a known connection state will be allowed, while others are rejected.
VPN	Virtual Private Network. A method of building private networks on top of public networks such that the private network is protected and separate.
WEP	Wired Equivalent Privacy. This is the encryption protocol specified in the original version of IEEE 802.11. It is now deprecated and does not meet current security standards.
Wi-Fi	A set of product compatibility standards for wireless LANs based on IEEE 802.11. The Wi-Fi term is managed by the Wi-Fi Alliance. Products carrying Wi-Fi certification have passed a series of compatibility tests.
WLAN/Wireless LAN	A type of wireless system based on the IEEE 802.11 series of protocols.

Terms	Definitions
WPA	Wi-Fi Protected Access. WPA implements the majority of the IEEE 802.11i standard, and was intended as an intermediate measure to take the place of WEP while 802.11i was prepared. WPA is designed to work with all wireless network interface cards. Products displaying the WPA logo have passed a certification program run by the Wi-Fi Alliance.
WPA2	Wi-Fi Protected Access version 2. WPA2 implements the full IEEE 802.11i standard, but will not work with some older network cards. Products displaying the WPA2 logo have passed a certification program run by the Wi-Fi Alliance.

Summary

Wireless technology is a fact of life in today's enterprise networks. The technology has been an area of rapid change over the past several years, which has led to confusion regarding best practices for 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 wireless is a layered approach consisting of the following layers:

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

Organizations implementing these best practices will be well protected against unauthorized and uncontrolled wireless as well as the malicious hacker bent on network intrusion. By implementing these best practices in a global wireless policy, organizations will find that wireless networks provide stronger security protection than current wired networks, with the economic benefits brought about by mobility.

About Aruba Networks, Inc.

People move. Networks must follow. Aruba securely delivers networks to users, wherever they work or roam. Our unified mobility solutions include Wi-Fi networks, identity-based security, remote access and cellular services, and centralized multi-vendor network management to enable the Follow-Me Enterprise that moves in lock-step with users:

- **Follow-Me Connectivity:** Adaptive 802.11a/b/g/n Wi-Fi networks optimize themselves to ensure that users are always within reach of mission-critical information;
- **Follow-Me Security:** Identity-based security assigns access policies to users, enforcing those policies whenever and wherever a network is accessed;
- **Follow-Me Applications:** Remote access solutions and cellular network integration ensure uninterrupted access to applications as users move;
- **Follow-Me Management:** Multi-vendor network management provides a single point of control while managing both legacy and new wireless networks from both Aruba and its competitors.

The cost, convenience, and security benefits of our unified mobility solutions 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>

© 2008 Aruba Networks, Inc. *AirWave®*, *Aruba Networks®*, *Aruba Mobility Management System®*, *Bluescanner*, *For Wireless That Works®*, *Mobile Edge Architecture®*, *People Move. Networks Must Follow®*, *RFProtect*, *The All Wireless Workplace Is Now Open For Business*, *Green Island*, and *The Mobile Edge Company®* are trademarks of Aruba Networks, Inc. All rights reserved. All other trademarks are the property of their respective owners.

WPB_SEC_US_080723



www.arubanetworks.com

| Tel. +1 408.227.4500 | Fax. +1 408.227.4550