
WHITE PAPER

WI-FI: SECURE ENOUGH FOR FEDERAL GOVERNMENT?

TECHNOLOGY, POLICY, AND REAL-WORLD RISK

aruba
a Hewlett Packard
Enterprise company




TABLE OF CONTENTS

| | |
|---|----|
| INTRODUCTION | 3 |
| SUMMARY OF RELEVANT POLICIES | 3 |
| HOW WI-FI SECURITY WORKS: AUTHENTICATION AND ENCRYPTION | 4 |
| ACCESS CONTROL FOR MULTI-USE NETWORKS | 9 |
| CONCLUSION | 13 |

INTRODUCTION

One of the few places that pervasive Wi-Fi is not found these days is in US Federal Government office buildings and military bases. Government IT departments explain this lack of modern technology by pointing to Information Assurance (IA) departments who block their planned deployments because of security concerns. IA departments, on the other hand, point to unclear rules, regulations, and policies around Wi-Fi use which prevent them from making informed risk decisions.

Wi-Fi is secure enough for government use and is approved by policy. This paper will provide a detailed, technical explanation of how modern enterprise Wi-Fi networks are secured. It will cover generic industry capabilities as well as Aruba-specific innovations that decrease risk for customers. The paper will also highlight different government policies governing Wi-Fi use.

SUMMARY OF RELEVANT POLICIES

Some government policies affect Wi-Fi use; how they are applied depends on the specific agency or department, the specific use case for the Wi-Fi network, and in some cases, geographical location. This section presents a summary of each policy and describes how it applies to Wi-Fi. For further details, refer to the specific policy.

FIPS 140-2

IT products which implement cryptography must be validated under FIPS 140-2 to be used by any US Government entity. Wi-Fi, specifically 802.11i/WPA2, makes use of AES-CCMP and a key derivation function, both of which are compliant with FIPS 140-2. Since 2015, all FIPS 140-2 validated products must provide 112 bits of security strength. Applied to Wi-Fi, this means that 802.1X authentication must be used; authentication based on pre-shared keys must use very large and carefully-constructed keys to be compliant.

Common Criteria Wireless LAN Access System Extended Package

The National Information Assurance Partnership (NIAP) recognizes a Common Criteria extended package for Wireless LAN Access Systems (<https://www.niap-ccevs.org/Profile/Info.cfm?id=376>). Products evaluated under this package are approved for use by the US Government for “sensitive but unclassified” applications, and also for classified applications when deployed as part of the NSA Commercial Solutions for the Classified program.

DoD 8420.01

The Department of Defense recently updated DODI 8420.01 on November 3, 2017. This policy covers rules for deployment of Wi-Fi technology in any DoD installation for unclassified use and mandates adherence to the Commercial Solutions for Classified Program for all classified networks. The policy requires the installation of a Wireless Intrusion Detection System (WIDS), mandates that two-factor authentication is used, and mandates specific forms of 802.1X authentication. Wi-Fi products must also be validated under FIPS 140-2. Systems that meet these requirements may connect Wi-Fi clients directly to unclassified DoD networks such as NIPRnet.

DoDIN APL

Previously known as the Unified Capabilities Approved Products List (UC APL) the Department of Defense Information Network Approved Products List (DoDIN APL) is administered by DISA and, for many DoD components, serves as a centralized list of products that are approved for deployment in DoD networks. Products on the list (found at <https://aplits.disa.mil>) are tested by one of three DoD test labs for both Information Assurance (IA) as well as Interoperability and Performance (IO). For Wi-Fi, three categories exist in the Wireless Local Area Network System. One, called “WLAS,” contains products that are approved to be deployed as Wi-Fi access systems. The other category, “WIDS,” is for wireless intrusion detection systems. Aruba Networks offers products listed in a variety of categories, including Wireless LAN Access System (WLAS), Wireless Intrusion Detection System (WIDS), Network Access Control (NAC), VPN Gateway, and Network Management System (NMS).

Other DoD Policies

The Marine Corps maintains “USMC Cybersecurity Directive 005” related to the use of portable electronic devices. This policy contains a significant amount of information on Wi-Fi networks. Wi-Fi is authorized in US Marine Corps networks. The policy is FOUO and thus not available on the public Internet.

Commercial Solutions for Classified

The NSA Commercial Solutions for Classified (CSfC) program publishes a series of Capability Packages (<https://www.nsa.gov/resources/everyone/csfc/capability-packages/>) that describe how composed systems of commercial products may be used to protect classified information. Both the

Campus WLAN CP (<https://www.nsa.gov/resources/everyone/csfc/capability-packages/#wlan>) as well as the Mobile Access CP (<https://www.nsa.gov/resources/everyone/csfc/capability-packages/#mobile-access>) include wireless components. Under CSFC, Wi-Fi is specifically authorized for use in classified environments.

CNSS

The Committee on National Security Systems is tasked with the protection of National Security Systems (NSS). CNSS policies can be found at <https://www.cnss.gov/CNSS/issuances/Policies.cfm#>. Policy 17, "Policy on Wireless Systems," describes procedures that must be followed when wireless networks are used near or within NSS. This policy also refers to CNSS policy 15, which describes cryptographic protection that must be applied in NSS and forms the basis for the use of commercial cryptography in Commercial Solutions for Classified.

CNSS Advisory Memorandum (CNSSAM) TEMPEST/1-13, otherwise known as "Red/Black Installation Guidance," was updated in January 2014 and specifies the distance that must be maintained between classified and unclassified information systems. The document is FOUO and thus not available on the public Internet. Page 8 of the document defines "low transmitter power" as 100mW or less, and Table 3 on page 9 of the document mandates a separation of 1 meter between a low-power transmitter and a classified information system. With typical Wi-Fi access point deployment on a ceiling, one meter of separation is normally easily achievable.

NIST

The National Institute of Standards and Technology (NIST) has produced Special Publication 800-153 (<http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-153.pdf>) which provides guidelines for securing wireless LANs. While the guidance in this publication is not binding for any Federal entity, it does provide best practices. References to NIST SP800-53 are provided at the end of the document.

HOW WI-FI SECURITY WORKS: AUTHENTICATION AND ENCRYPTION

The security of Wi-Fi is rooted in two interrelated functions: authentication and encryption. The strength of the entire security system rests on doing these two things correctly. All enterprise Wi-Fi vendors who carry the Wi-Fi Alliance "WPA2 Enterprise" certification are capable of basic authentication and encryption sufficient to meet the needs of enterprise users. Some vendors add additional capabilities to meet and

exceed the requirements of government users, who typically have more stringent security requirements than the average enterprise user.

Wi-Fi Supports Different Security Levels

Wi-Fi networks and devices are commonplace, giving people broad experience with the technology. Unfortunately, use of Wi-Fi in private homes and public hotspots often gives a false impression of the level of security that Wi-Fi can provide. Wi-Fi in public hotspots typically uses no security at all – the network is open, and anyone can connect. Wi-Fi in private homes is often protected using a pre-shared key.

Anyone with access to the key can connect. These are appropriate security levels for a small private network, but this type of security is not appropriate for internal corporate or government networks. For these higher security networks, Wi-Fi must be configured for WPA2-Enterprise. WPA2-Enterprise mandates that all users and devices must authenticate themselves using the 802.1X protocol before gaining any access to network services. The authentication process is also used to generate one-time dynamic encryption keys that protect the privacy and integrity of data as it travels over wireless links. Together, authentication and encryption provide Wi-Fi with better security than the average wired Ethernet network.

A Detailed Look at 802.1X

The 802.1X standard defines a method of port-based authentication for networks. In the context of Wi-Fi, a "port" is a logical concept that is applied to a wireless client or station. 802.1X is a framework for how authentication is performed – within the framework, some Extensible Authentication Protocol (EAP) methods may be used. The choice of EAP method depends on the desired security strength, whether or not encryption keys need to be generated as part of the authentication process (this is always true for Wi-Fi), and the type of authentication credentials to be used. In enterprise settings, two types of authentication credentials are common: username/password and X.509 certificates. For government networks, only certificates are generally acceptable, because they provide a known strong level of security. Certificates, when combined with technologies like smart cards or crypto tokens, can provide two-factor authentication: a PIN or password is entered (something you know) to unlock the smart card, followed by the certificate and private key (something you have) being used to perform 802.1X authentication. Smart cards and tokens provide an additional benefit – they cannot be copied.

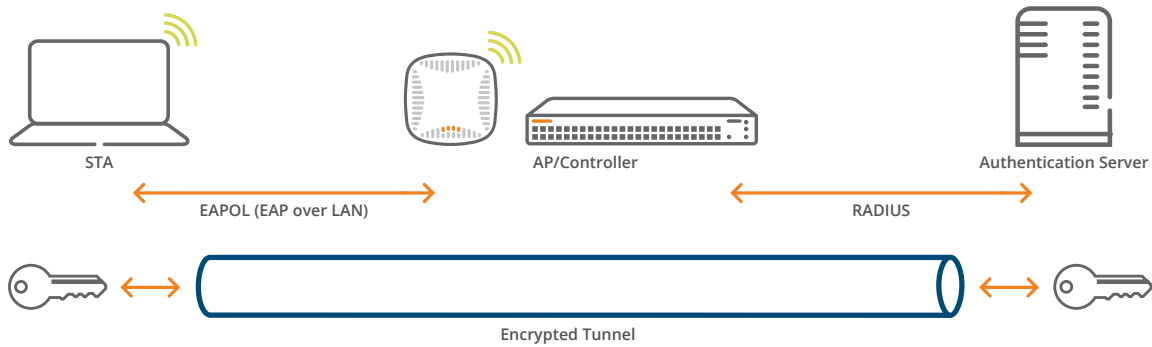


Figure 1: 802.1X Framework

Three components make up an 802.1X system, as shown in Figure 1: The station (or client), the authenticator (typically an AP-controller combination in an enterprise network), and an authentication server. EAPOL messages are sent over the wireless interface, and those EAPOL messages are translated into RADIUS messages on the wired network between the AP/controller and the authentication server.

Encryption keys are negotiated between the station and the authentication server through a key derivation process. Keys are derived independently by both the station and the authentication server – they are never transmitted over the air and are thus never at risk of interception. It is important to note that before authentication, the station has no access to the wired network. Only well-formed EAP messages are passed through, and only to a defined set of authentication servers. All other network traffic is blocked. Additionally, wireless clients cannot communicate with other wireless clients on the same access point until authentication has been completed.

The EAP method approved for use in all government agencies and departments is EAP-TLS, which is defined in RFC 5216. EAP-TLS performs mutual, certificate-based authentication using the TLS protocol. First, the client validates the authentication server by checking whether the certificate

presented by the server is trusted. The process is nearly identical to that carried out by web browsers when communicating with a secure website over HTTPS. If the client trusts the authentication server, it will present its own certificate and then prove that it has possession of the corresponding private key. Standard cryptography is used to perform these steps. RSA2048 is most common and is always used when CAC or PIV cards are employed. For stronger security – particularly for classified applications – ECDSA may be used. This entire process will be familiar to anyone who has examined a packet capture of an HTTPS session between a client and a government website that requires CAC or PIV authentication – the steps and protocols involved are nearly identical.

The security of cryptography, and of protocols that use cryptography, is normally measured regarding the number of bits of strength that it provides. The number of bits corresponds to a symmetric cipher (e.g., AES) using the same number of bits for its key. FIPS 140-2 requires a minimum of 112 bits of strength for unclassified use cases. CNSS policy requires a minimum of 128 bits to protect SECRET, and 192 bits of strength to protect TOP SECRET. The Wi-Fi WPA2 authentication process using EAP-TLS, as described above, is capable of achieving all of these strengths, as shown in Table 1.

| TABLE 1: EAP-TLS SECURITY STRENGTH | | | |
|------------------------------------|---------------|-------------------------------------|----------|
| Use Case | TLS Version | Authentication Credential | Strength |
| Unclassified (SBU) | 1.0, 1.1, 1.2 | RSA 2048-bit certificate | 112 bits |
| SECRET | 1.2 | ECDSA certificate using p256 curve* | 128 bits |
| TOP SECRET | 1.2 | ECDSA certificate using p384 curve | 192 bits |

* Note: As of September 2015, CSfC no longer permits the 128-bit level to be used. All CSfC deployments must use the 192-bit level.

Support for TLS 1.0 and RSA2048 is common in all enterprise Wi-Fi equipment. Support for TLS 1.2 and ECDSA, however, is not common and only products designed for high-security deployments will have these features.

Authentication seeks to establish whether or not the peer possesses a valid credential – in the case of EAP-TLS. The process determines that the peer has a certificate that is issued by a trusted certificate authority, that the certificate is not expired or revoked, and that the peer proves possession of a private key that corresponds to the certificate's public key. More advanced Wi-Fi systems also permit an authorization process to take place. Authorization determines what an authenticated client is allowed to do. For example, a military Wi-Fi system may be configured to authenticate any Wi-Fi user with a CAC, but to provide access to a local facility's internal network only if the CAC corresponds to an account in a local Active Directory domain. Or, an authorization step may consult a database that returns threat condition information, and the system may deny access to certain low-priority users and devices during a time of elevated threat. Authorization is independent of authentication but may take place simultaneously.

EAP-TLS has fewer weaknesses than other EAP methods (e.g., password-based authentication with EAP-PEAP-MSCHAPv2.) Specifically, EAP-TLS will never reveal a credential that would allow an attacker to impersonate an authorized user. EAP-TLS relies only on public-key cryptography and PKI – a certificate indicates an identity (and a certificate is intended to be exposed), but the private key proves that the certificate belongs to a particular user or device. The private key is never exposed – only a cryptographic signing operation is performed. For as long as RSA and ECDSA remain secure, EAP-TLS authentication is also secure. There are, however, remaining risks:

- First, if a client were to misconfigure the 802.1X supplicant software so that it did not validate the authentication server's certificate, this would be a critical security flaw. A client could be "authenticated" by a rogue authentication server and subsequently connected to a rogue Wi-Fi network. This is in theory no more dangerous than connecting to a public Wi-Fi hotspot at a hotel or airport, but the difference is that the user may think he

or she is connected to a secure internal network, and this may make social engineering attacks more effective. It is thus important that the 802.1X supplicant software on all client systems be configured properly to validate server certificates presented to it during authentication. For client systems that are part of an Active Directory domain, this setting can be enforced using Group Policy.

- Another remaining risk is theft of credentials. If a certificate and private key are stored on a hard disk or in unprotected flash memory of a mobile device, they could be at risk of copying if an attacker gains control of the client device. If an attacker gains possession of a certificate's private key, he can impersonate the original user at-will. This is why most government Wi-Fi policies require two-factor authentication and some form of hardware-protected key storage (smart card, USB token, or a mobile device with hardware-integrated keystores). With hardware-protected key storage, an attacker would need to retain control of a physical device or card to use the credential, and also would have to know the password or PIN to unlock the device. Stealthy and undetectable copying of the private key would not be possible in this scenario.

Encryption for Confidentiality and Integrity Protection

Once authentication is complete, each station will derive a set of encryption keys to be used for confidentiality and integrity protection of wireless data. The key derivation process begins with a Pairwise Master Key (PMK) that is independently calculated by the station and the authentication server as a result of random data used during the authentication process. The PMK is then used, in a process called the 4-way handshake, to derive other keys; these keys are used to protect network traffic. This process is completely unlike WEP (Wired Equivalent Privacy) from the early days of Wi-Fi. WEP is an inherently broken encryption technique where a pre-shared key entered by the user is used directly for over-the-air encryption. Breaking a WEP key once means that all future communication can be compromised. WPA2, on the other hand, generates new keys every time authentication takes place, and during a long-lived session will also periodically re-key. Attacks against WEP are completely ineffective against WPA2.

WPA2 relies on an encryption protocol called CCMP (Counter Mode with CBC-MAC), which uses AES as the underlying cipher. CCMP is depicted in Figure 2. The protocol provides Authenticated Encryption with Associated Data (AEAD). This is a block cipher mode that, in a single protocol, provides confidentiality, integrity protection, and authenticity assurance. NSA mandates AEAD ciphers in the CSfC program; AES-CCM is one such cipher, while AES-GCM is the other. As seen in Figure 2, two different keys are used for CCMP. The first key is a data integrity key which is used by the CBC-MAC portion of the cipher. The second is a data encryption key, used by AES in counter mode to encrypt network traffic. Both keys are 128 bits and are derived from the PMK during the 4-way handshake. These keys are known as Pairwise Transient Keys (PTKs), the name “transient” indicating that they are temporary keys that live for a certain period and then are destroyed.

As it stands today, WPA2 encryption provides 128 bits of security strength – more than adequate for sensitive/unclassified use, and good enough for protection of SECRET under CNSS policy. However, in an actual classified application, NSA now requires 192 bits of security strength in a Commercial Solutions for Classified deployment. This effectively means that WPA2 encryption cannot be used for CSfC. Recognizing the gap, NSA has authorized use of Wi-Fi at the 128-bit level as long as the second CSfC encryption layer provides 192 bits of strength. An amendment to IEEE 802.11 ac (the newest gigabit Wi-Fi standard) authored by Aruba and co-sponsored by Cisco permits Wi-Fi devices to use ciphers

compliant with Suite B – that is, ciphersuites that achieve the required 192 bits of strength. At the time of this writing, no commercially-available client devices or operating systems support the Suite B ciphersuites, but this is likely to change in the future.

Risks of using WPA2 encryption are minimal, as long as keys remain secret. There are no known practical attacks against WPA2 encryption as long as the network and the client use a correct implementation; indeed, a break of Wi-Fi encryption would effectively mean a break of all commercial cryptography. In 2017, the widely-publicized KRACK attack revealed a widespread implementation flaw that required patches to be deployed. KRACK, however, was not a break of Wi-Fi encryption.

Theft of keys, however, represents a real risk. An attacker with possession of a PMK for a single session can obtain transient keys for that session for as long as it lasts. An attacker with possession of multiple PMKs for multiple sessions can decrypt all Wi-Fi communication for an entire network. This is where physical security matters! An attacker must be prevented from obtaining physical control of hardware where key material is present. The authentication server must be protected from compromise because the PMK is derived from the authentication server before being transmitted to the Wi-Fi infrastructure. Wi-Fi controllers must be physically protected; typically 802.1X processing happens on the Wi-Fi controller so that it will have access to the PMK. Finally, depending on architecture, Wi-Fi access points may need to be physically protected.

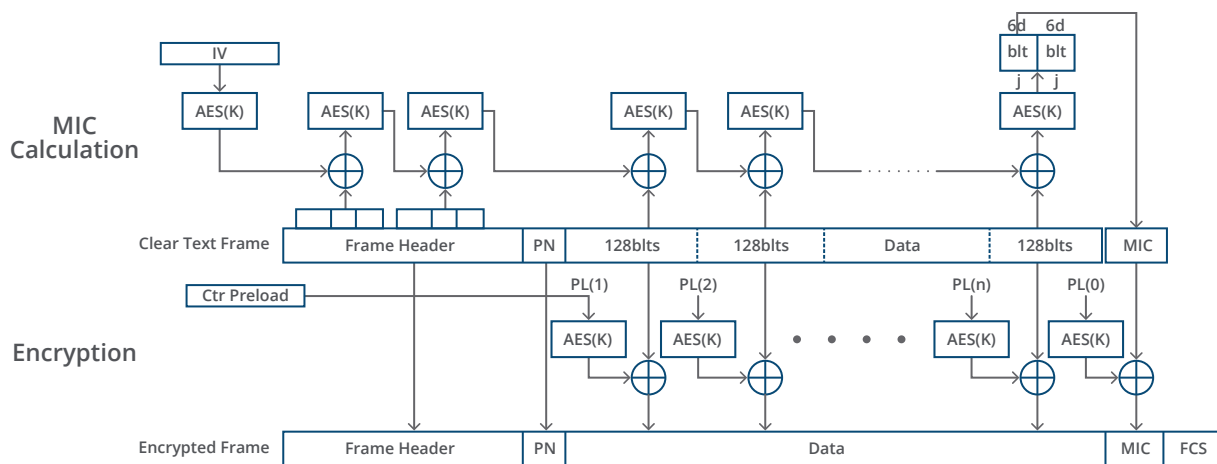


Figure 2: AES-CCMP

Aruba's security architecture is different than all other vendors. In the default configuration, known as tunnel mode, Aruba access points do not perform encryption/decryption and thus do not contain any encryption keys. The access points receive wireless frames from the radio interface, encrypted, and immediately package these encrypted wireless frames into an IP tunnel to the mobility controller. Once at the mobility controller, the IP tunnel packet header is removed and what remains is an encrypted 802.11 Wi-Fi frame. The controller then processes this frame, decrypting it and turning it back into a standard routable IP packet. Access points never had access to encryption keys, and thus were unable to process the Wi-Fi traffic locally. The implication is that an attacker who gains physical control of an Aruba AP, even one who replaces the AP's firmware with custom malicious code, will be unable to break into Wi-Fi sessions that pass through that access point. All Wi-Fi encryption is between the client and the mobility controller; the AP is simply a pass-through device. Mobility controllers must be physically protected, but access points do not.

In non-Aruba Wi-Fi deployments, access points must be physically protected against compromise. To highlight a worst-case scenario, one vendor authenticates Wi-Fi access points with a certificate and private key that are stored in flash memory, with no hardware protection of the key. Anyone who can obtain physical access to the AP for a brief period can copy this certificate and key, and later use it to "become" an AP by connecting to the controller using an AP emulator. Because these APs perform encryption/decryption for Wi-Fi traffic, keys are sent to the access point. Furthermore, to enable fast roaming to work, encryption keys for sessions on neighboring access points are also shared with surrounding APs. Thus, by stealing a single long-lived credential, one can get access to numerous encryption keys for active Wi-Fi sessions, enabling an attacker to decrypt these sessions and recover their traffic. For this reason, APs that performs encryption/decryption must be installed inside tamper-proof or tamper-evident enclosures, and regular inspection schedules must be put in place to ensure APs have not been accessed by unauthorized persons.

Anyone who can obtain physical access to the AP for a brief period can copy this certificate and key, and later use it to "become" an AP by connecting to the controller using an AP emulator. Because these APs perform encryption/decryption for Wi-Fi traffic, keys are sent to the access point. Furthermore, to enable fast roaming to work, encryption keys for sessions on neighboring access points are also shared with surrounding APs. Thus, by stealing a single long-lived credential, one can get access to numerous encryption keys for active Wi-Fi sessions, enabling an attacker to decrypt these sessions and recover their traffic. For this reason, APs that performs encryption/decryption must be installed inside tamper-proof or tamper-evident enclosures, and regular inspection schedules must be put in place to ensure APs have not been accessed by unauthorized persons.

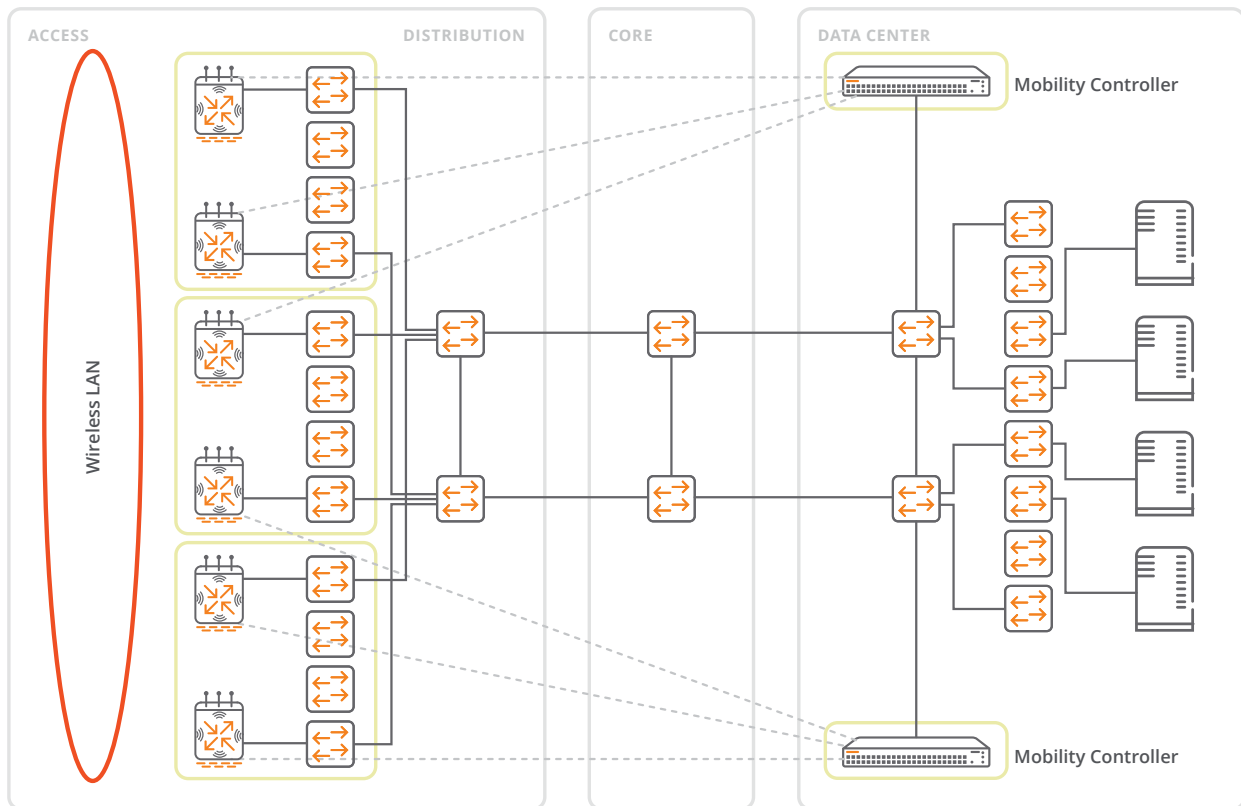


Figure 3: Aruba Wi-Fi Architecture

ACCESS CONTROL FOR MULTI-USE NETWORKS

Many network administrators approach Wi-Fi from the perspective of port-based security, where a given physical port is assigned to a particular security domain or VLAN.

Port-based security also lends itself to “air-gapped” networks, where the complete physical separation between networks is possible. In the Wi-Fi world, things are more complex. A port-based security approach will have a new Wi-Fi SSID created for each service or security domain, and physical separation often implies multiple Wi-Fi access points in the same area. Both approaches lead to sub-optimal radio frequency (RF) spectrum utilization, and air-gapped Wi-Fi networks lead to significant expenses for duplicate equipment along with associated installation and cabling costs. It is therefore wise to consider new approaches when deploying multi-service Wi-Fi networks.

The Need for Separation

Controlling access to IT resources is difficult and complicated in a large organization. Generally, access control is performed at a server or application level, closest to the point where information is stored. But it is very easy to make mistakes in access control configuration or to have vulnerabilities present in applications that inadvertently expose data. For this reason, access control at the network level is also typically employed as defense-in-depth. Here, access control is typically coarse-grained; employees use an employee network while guests use a guest network. A third network may be added to support “devices” – a catch-all term for everything from barcode scanners to printers. In non-military organizations, these three networks are typically attached to the same physical infrastructure, with VLAN tagging providing separation. The same construct is then extended into the Wi-Fi domain through the creation of three wireless SSIDs – an employee SSID supporting WPA2-Enterprise, a guest SSID that is open and uses captive portal authentication, and a “devices” SSID that may use legacy Wi-Fi security such as TKIP or even WEP. Within each network, however, most organizations apply a one-size-fits-all access policy, allowing application servers to deal with fine-grained access control.

Network segmentation also augments a defense-in-depth posture by giving access rights only where needed – if a corporate server contains a security vulnerability, the risk posed by that vulnerability is reduced if only members of a particular group can reach the server.

In military and government environments, the concept of physically separate networks is commonplace – there are different networks based on classification levels (e.g., NIPRnet, SIPRnet) and on communities of interest (e.g., CENTRIX, BICES). High-security organizations such as the military and intelligence community also deploy and configure firewalls between different segments within their internal networks much more frequently than other types of organizations.

These firewalls are used to provide fine-grained control over who has access to resources within the network. IT administrators in these organizations are then understandably leery about connecting multi-service Wi-Fi networks into their previously isolated and controlled networks. The typical response is to insist on air-gapped, separate Wi-Fi networks for each service. But as previously discussed, such an approach is not efficient – and with the security capabilities of Aruba Wi-Fi networks, such an approach is not necessary.

Aruba Firewall Basics

All Aruba Mobility Controllers integrate a full stateful firewall, designed to enforce security policies and separation between different types of Wi-Fi users. The firewall delivers role-based access control (RBAC) by first putting users into roles, and then enforcing a set of firewall rules per role; this is shown in Figure 4. Each user gets a separate copy of firewall rules so that traffic separation can be as granular as a single user or device. The mobility controller can act as a traditional L3-4 firewall, enforcing stateful access control lists, and also as a next-generation firewall with Layer 7 application intelligence, recognizing over 1500 enterprise applications through deep packet inspection. The firewall is NIAP-accredited under the Common Criteria Traffic Filtering Firewall Extended Package (TTFW-EP).

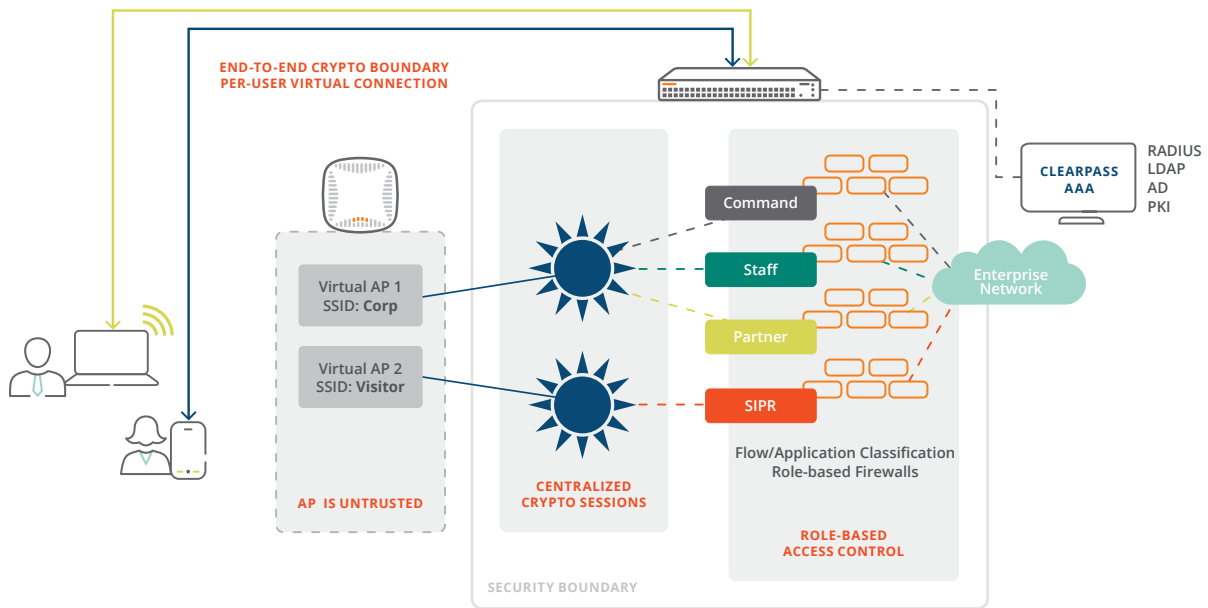


Figure 4: Role-Based Access Control

The key to the Aruba firewall’s security is centralized encryption, discussed in the previous section. It is critical to ensure that wireless users cannot manipulate their network traffic in such a way as to bypass firewall rules or obtain the rights of a user from a different role. Centralized encryption ensures that there is a solid, unbreakable linkage between network traffic coming from a user and the role and firewall policies that are applied to that traffic at the mobility controller. A user/device can be uniquely identified by a cryptographic session, with the key for that session residing in only two places – the mobility controller and the end-user device. Regardless of any tampering that a user may do with his or her network traffic (e.g., changing MAC address, changing IP address, inserting 802.1q VLAN tags, etc.) the mobility controller will still recognize the traffic as coming from that particular user, based on the traffic decrypting under a given key. And because the mobility controller is the same device that both holds the cryptographic key and also enforces firewall rules, there is no possibility of the user’s traffic bypassing any rules.

Separation Within a Single Security Zone

There is no precise definition of what a security zone is, but we can adopt the following for the purpose of this discussion:

A security zone is a grouping of devices and network services under common administrative control, and with similar degrees of risk in the event of a breach.

Using this definition, a corporate IT department might group all enterprise networks with access to internal servers, and all devices that are managed by corporate IT, into the same security zone. A “bring your own device” (BYOD) security zone might exist in parallel, which consists of employees of the organization but using personally owned devices instead of corporate-issued devices. The BYOD zone may have access to a restricted set of corporate resources. A third security zone may consist of visitors – people who are not employees of the organization.

Often, network access is granted equally within a security zone – an employee in the Finance department may have the same network access as an employee in Human Resources. Increasingly, however, it is becoming more attractive to segment users within the same security zone, as a means of greater protection against malware and advanced persistent threats. To perform this type of micro-segmentation, two pieces are required: 1) The network must know the identity and role of the user, and 2) A security profile must be built for each role to indicate what types of network traffic are allowed for users in that role. Wi-Fi provides an ideal way to learn the identity of network users since all users on a secure Wi-Fi network must authenticate before connecting. Aruba's ClearPass Policy Manager (CPPM), acting as an authentication server, can query multiple backend identity management systems such as Active Directory, LDAP, SQL databases, and other systems to determine into which role a Wi-Fi user should be placed. Commonly, this is done by group membership in Active Directory, with each group having unique security policies applied. Thus, for example, members of the IT group may be given the ability to use protocols such as SSH and SNMP on the network, while members of other groups are not. Once the user's role is determined by ClearPass, it can be passed back to an Aruba mobility controller where appropriate firewall policies are applied.

Separation between Multiple Security Zones

The risk is higher when providing separation between different security zones. Within the same security zone, segmentation is a defense-in-depth technique, and the consequence of misconfiguration is a loss of depth rather than a complete loss of security. When providing separation between different security zones, a misconfiguration could have much more serious consequences. To be clear, the separation between multiple security zones is an entirely reasonable and solvable problem using Aruba Wi-Fi solutions, but it does require more care.

At a basic level, it would be entirely possible to set up a network supporting multiple security zones in the same way one would set it up for separation in a single security zone. Users would be placed into roles, and firewall policies would be assigned to ensure that the two classes of users could not access unauthorized resources. Three specific changes from that approach are recommended, however:

1. It is typically desirable to use different IP address space for different security zones (e.g., employees versus guests). Firewall rules in an Aruba mobility controller are one line of defense, but very often there are other

firewalls in the the network that may not be identity-aware but that base their access control decisions on source IP address. It would be unwise to have users from a different security zone sharing the same IP address space. Further, use of "foreign" or non-routable IP address space for less-trusted security zones can serve as another line of defense, since internal systems would have no way to route traffic back to that address space. A best practice then is to place less-trusted security zones into their own VLAN within the mobility controller (note that this VLAN will not necessarily extend outside the mobility controller) with unique IP address space.

2. The traffic forwarding behavior for less-trusted traffic should be different than that used for trusted network traffic. The exit point of each VLAN in a less-trusted zone should be a different router than is used for trusted traffic. In some cases, the exit point of the traffic could be a tunnel (GRE or IPsec) that takes traffic across a trusted backbone and exits at a gateway or router that is dedicated to this security zone.
3. A different SSID should be used for different security zones. Most often, this is out of necessity – different security zones are often using different Wi-Fi encryption and authentication methods. An internal network will typically be WPA2, while a guest network, for example, will be open with Web-based portal authentication. The additional problem is that Wi-Fi clients sharing the same SSID can receive broadcast and multicast traffic encrypted using group keys. It is possible to filter such traffic, but this again requires careful configuration. Allowing a less-trusted user to receive broadcast traffic from an internal network could reveal sensitive information about network topology and other devices on the network. Best practice then is to place less-trusted traffic on its own SSID. This can be a tricky recommendation too many SSIDs will lead to performance problems, so the number of SSIDs must be weighed against the need for separation. If going beyond 4-5 SSIDS, consider combining similar security zones into one, and implement filtering carefully to prevent traffic bleed.

Where risk is deemed to be too high for the architecture described in this section, Aruba MultiZone may be an option. It is more costly, but provides outstanding security separation. MultiZone will be described in the following section.

Separation between Multiple Security Classifications

Transmission of classified information over Wi-Fi networks is only possible today using a dual-tunnel CSfC architecture.

The two dominant architectures that involve Wi-Fi are the Mobile Device Capability Package (MACP) and the Campus WLAN Capability Package (WLANCP). In the case of MACP, the Wi-Fi layer is not involved in the security of the system; Wi-Fi is a third layer of encryption. It is thus possible to mix multiple classification levels over the same Wi-Fi network since security separation is provided at higher layers of the network stack. Aruba recommends treating the MACP network as a security zone, using the architecture in the previous section. This keeps the classified Wi-Fi traffic off the same Wi-Fi network used for unclassified traffic, which can provide some security benefits.

For organizations deploying the WLANCP architecture, Wi-Fi is one of the two required encryption layers. Under WLANCP 1.0, equipment that removes the outer layer of cryptography

must be treated as being at the same level as the cleartext data – in other words, a Wi-Fi controller carrying SECRET network traffic must itself be treated as SECRET, even though traffic passing through the controller is still encrypted using IPsec. This rule made it impossible to use WLANCP with multiple networks carrying different classification levels since the Wi-Fi controller could not simultaneously carry traffic from different classification levels.

The new WLANCP 2.0 publication allows for a “shared gray network”, as shown below in Figure 5. This means that traffic of multiple classification levels can be carried over the same Wi-Fi network. However, for unclassified traffic, the CSfC two-tunnel architecture must still be used – unclassified devices would need to connect to a Wi-Fi network and then establish an IPsec tunnel to a VPN concentrator to get network access. That is a drastic design change from the unclassified Wi-Fi networks already deployed, and so is not drop-in compatible with existing Wi-Fi networks.

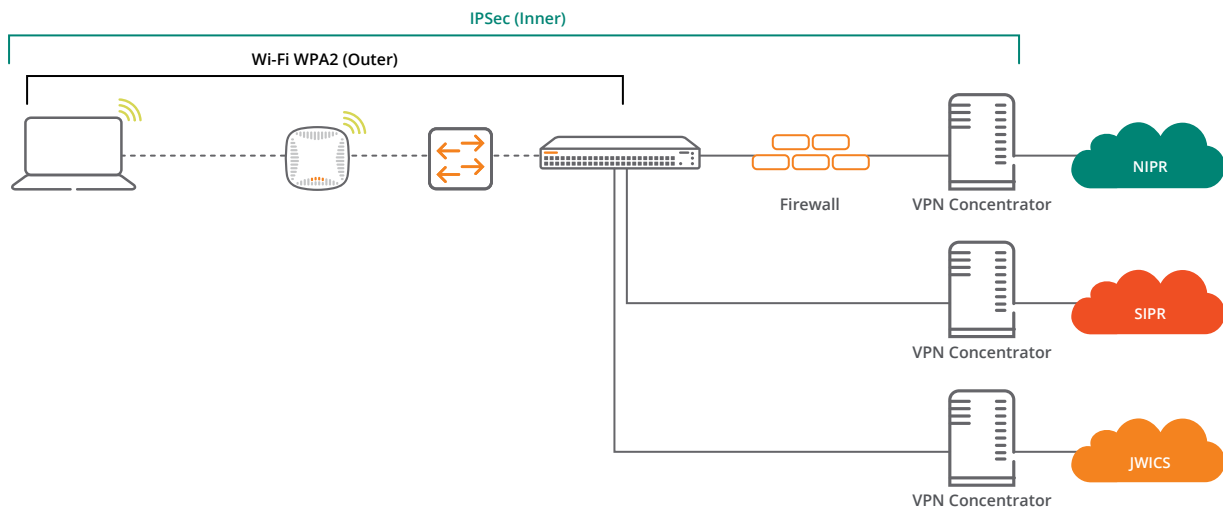


Figure 5: WLANCP 2.0 Shared Gray Network

A new capability from Aruba, called MultiZone, provides a potential solution. With MultiZone, shown in Figure 6, different Wi-Fi SSIDs are serviced by different mobility controllers.

Aruba's centralized encryption architecture allows this feature to provide equivalent security to physically separated networks. Because access points do not perform encryption or other security functions, the access point can be shared between different classification levels. Controllers have visibility only into traffic corresponding to their SSID, so one controller cannot see another controller's user information. This architecture thus supports both the WLANCP 1.0 requirements as well as the WLANCP 2.0 requirements.

CONCLUSION

Given the proper care and design, Wi-Fi is undoubtedly secure enough for Federal Government. Not only is it compliant with applicable policies, it includes technology that makes it demonstrably secure against attack. Wi-Fi in Federal Government does require extra care – for example, selecting components that are FIPS 140-2 validated and evaluated under Common Criteria. Both those are not exotic, expensive components – the leading enterprise Wi-Fi providers in the market provide solutions that meet these requirements.

Wi-Fi is here to stay, and Wi-Fi is poised to replace a significant portion of wired campus networks. Now is the time to determine how Wi-Fi fits into your own organization's security framework.

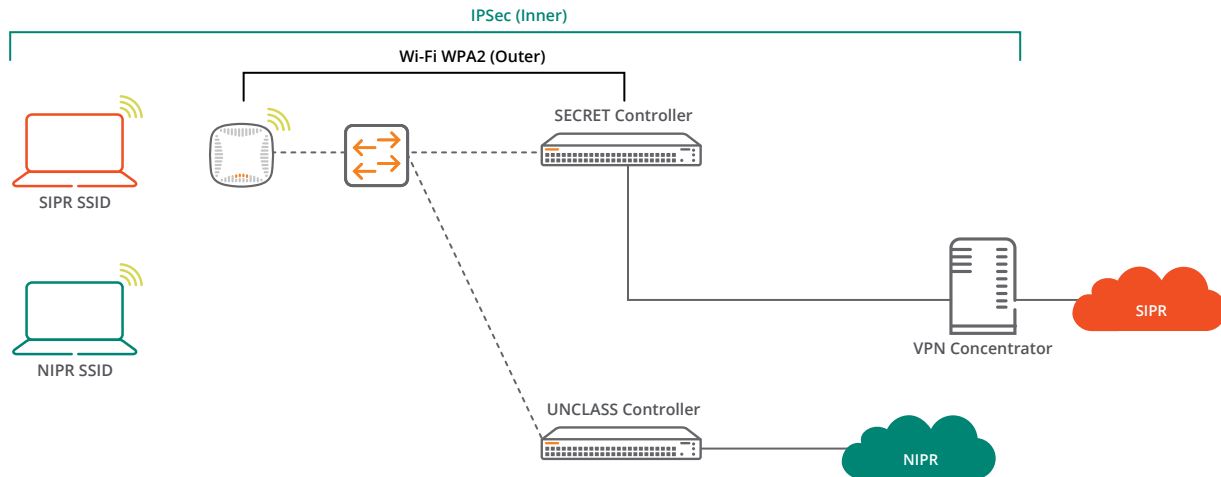


Figure 6: Aruba MultiZone