

Network Working Group
Request for Comments: 5418
Category: Informational

S. Kelly
Aruba Networks
T. Clancy
LTS
March 2009

Control And Provisioning of Wireless Access Points (CAPWAP)
Threat Analysis for IEEE 802.11 Deployments

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Abstract

Early Wireless Local Area Network (WLAN) deployments feature a "fat" Access Point (AP), which serves as a stand-alone interface between the wired and wireless network segments. However, this model raises scaling, mobility, and manageability issues, and the Control and Provisioning of Wireless Access Points (CAPWAP) protocol is meant to address these issues. CAPWAP effectively splits the fat AP functionality into two network elements, and the communication channel between these components may traverse potentially hostile hops. This document analyzes the security exposure resulting from the introduction of CAPWAP and summarizes the associated security considerations for IEEE 802.11-based CAPWAP implementations and deployments.

Table of Contents

1. Introduction4
1.1. Rationale for Limiting Analysis Scope to IEEE 802.115
1.2. Notations6
2. Abbreviations and Definitions7
3. CAPWAP Security Goals for IEEE 802.11 Deployments8
4. Overview of IEEE 802.11 and AAA Security8
4.1. IEEE 802.11 Authentication and AAA9
4.2. IEEE 802.11 Link Security11
4.3. AAA Security11
4.4. Cryptographic Bindings12
5. Structure of the Analysis13
6. Representative CAPWAP Deployment Scenarios14
6.1. Preliminary Definitions14
6.1.1. Split MAC15
6.1.2. Local MAC15
6.1.3. Remote MAC15
6.1.4. Data Tunneling16
6.2. Example Scenarios16
6.2.1. Localized Modular Deployment16
6.2.2. Internet Hotspot or Temporary Network17
6.2.3. Distributed Deployments18
6.2.3.1. Large Physically Contained Organization ...18
6.2.3.2. Campus Deployments18
6.2.3.3. Branch Offices18
6.2.3.4. Telecommuter or Remote Office19
7. General Adversary Capabilities19
7.1. Passive versus Active Adversaries20
8. Vulnerabilities Introduced by CAPWAP22
8.1. The Session Establishment Phase22
8.1.1. The Discovery Phase22
8.1.2. Forming an Association (Joining)23

- 8.2. The Connected Phase23
- 9. Adversary Goals in CAPWAP24
 - 9.1. Eavesdrop on AC-WTP Traffic24
 - 9.2. WTP Impersonation and/or Rootkit Installation25
 - 9.3. AC Impersonation and/or Rootkit Installation26
 - 9.4. Other Goals or Sub-Goals26
- 10. Countermeasures and Their Effects27
 - 10.1. Communications Security Elements27
 - 10.1.1. Mutual Authentication27
 - 10.1.1.1. Authorization27
 - 10.1.2. Data Origin Authentication29
 - 10.1.3. Data Integrity Verification29
 - 10.1.4. Anti-Replay29
 - 10.1.5. Confidentiality29
 - 10.2. Putting the Elements Together30
 - 10.2.1. Control Channel Security30
 - 10.2.2. Data Channel Security30
- 11. Countermeasures Provided by DTLS30
- 12. Issues Not Addressed By DTLS31
 - 12.1. DoS Attacks31
 - 12.2. Passive Monitoring (Sniffing)32
 - 12.3. Traffic Analysis32
 - 12.4. Active MitM Interference32
 - 12.5. Other Active Attacks32
- 13. Security Considerations32
- 14. Acknowledgements32
- 15. References33
 - 15.1. Normative References33
 - 15.2. Informative References33

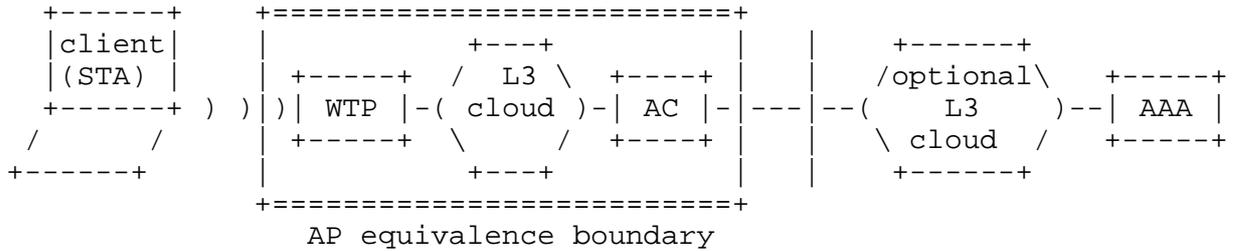


Figure 2: WLAN Deployment utilizing CAPWAP

For our purposes, it is useful to think of the entire CAPWAP system as a sort of "AP equivalent", and the figure above illustrates this concept. With this model in mind, our ideal is to ensure that CAPWAP introduces no vulnerabilities that aren't present in the original fat AP scenario. As we will see, this is not entirely possible, but from a security perspective, we should nonetheless strive to achieve this equivalence as nearly as we can.

1.1. Rationale for Limiting Analysis Scope to IEEE 802.11

Although CAPWAP derives from protocols that were designed explicitly for management of IEEE 802.11 wireless LANs, it may also turn out to be useful for managing other types of network device deployments, wireless and otherwise. This might lead one to conclude that a single security analysis, except for minor per-binding variations, might be sufficient. However, based on a limited number of additional related scenarios that have been described as likely candidates for CAPWAP thus far, e.g., use with Radio Frequency Identification (RFID) sensors, this does not seem to be a simple, binary decision.

Contrasting RFID and IEEE 802.11 deployment scenarios, IEEE 802.11 users typically authenticate to some a back-end AAA server, and as a result of that exchange, derive cryptographic keys that are used by the STA and WTP to encrypt and authenticate over-air communications. That is, the threat environment is such that the following typically holds:

- o The user is not immediately trusted, and therefore must authenticate.
- o The WTP is not immediately trusted, and therefore must indirectly authenticate to the user via the AAA server.
- o The AAA server is not necessarily trusted, and therefore must authenticate.

- o The medium is not trusted, and therefore communications must be secured.

RFID tags, on the other hand, typically do not have the same authentication, confidentiality, and integrity requirements as the principals in a WLAN deployment, and are not, generally speaking, well suited to environments in which similar threats exist. As a result of the combination of WLAN security requirements and the Medium Access Control (MAC)-splitting behavior that epitomizes the IEEE 802.11 binding for CAPWAP, there are definite security-related considerations in the WLAN case that simply do not exist for an RFID-based adaptation of CAPWAP.

Now, there certainly are similarities and overlapping security considerations that will apply to any CAPWAP deployment scenario. For example, authentication of the AC for purposes of WTP device management operations is probably always important. Even so, the threats to RFID are different enough to suggest the need for a separate security analysis in that case. For example, since RFID tags commonly deployed today implement no over-air authentication, integrity, or confidentiality mechanisms, the need for similar mechanisms between the WTP and AC for RFID tag data communications is not clearly indicated -- that is, the threats are different.

We have limited visibility into the varied ways in which CAPWAP might be adapted in the future, although we may observe that it seems to provide the basis for a generalized scalable provisioning protocol. Given our currently limited view of the technologies for which it might be used, it seems prudent to recognize that our current view is colored by the IEEE 802.11 roots of the protocol, and make that recognition explicit in our analysis. If newly added bindings turn out to be substantially similar to IEEE 802.11, then the associated binding documents can simply reference this document in their security considerations, while calling out any substantive differences. However, it does appear, based on our current limited visibility, that per-binding threat analyses will be required.

1.2. Notations

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Abbreviations and Definitions

- o AAA - Authentication Authorization and Accounting
- o AC - Access Controller
- o AES-CCMP - Advanced Encryption Standard - Counter-mode CBC MAC Protocol
- o AP - (wireless) Access Point
- o CAPWAP - Control And Provisioning of Wireless Access Points
- o Cert - X509v3 Certificate
- o DIAMETER - proposed successor to RADIUS protocol (see below)
- o DoS - Denial of Service (attack)
- o DTLS - Datagram Transport Layer Security
- o EAP - Extensible Authentication Protocol
- o MitM - Man in the Middle
- o PMK - Pairwise Master Key
- o PSK - Pre-Shared Key
- o RADIUS - Remote Authentication Dial-In User Service
- o STA - (wireless) STAtion
- o TK - Transient Key
- o TKIP - Temporal Key Integrity Protocol
- o WEP - Wired Equivalent Privacy
- o WLAN - Wireless Local Area Network
- o WTP - Wireless Termination Point

3. CAPWAP Security Goals for IEEE 802.11 Deployments

When deployed for use with IEEE 802.11, CAPWAP should avoid introducing any system security degradation as compared to the fat AP scenario. However, by splitting the AP functions between the WTP and AC, CAPWAP places potentially sensitive protocol interactions that were previously internal to the AP onto the Layer 3 (L3) network between the AC and WTP. Therefore, the security properties of this new system are dependent on the security properties of the intervening network, as well as on the details of the split.

Since CAPWAP does not directly interact with over-air or AAA protocols, these are mostly out of scope for this analysis. That is, we do not analyze the basic AAA or IEEE 802.11 security protocols directly here, as CAPWAP does not modify these protocols in any way, nor do they directly interact with CAPWAP. However, by splitting AP functionality, CAPWAP may expose security elements of these protocols that would not otherwise be exposed. In such cases, CAPWAP should explicitly avoid degrading the security of these protocols in any way when compared to the fat AP scenario.

4. Overview of IEEE 802.11 and AAA Security

While this document is not directly concerned with IEEE 802.11 or AAA security, there are some subtle interactions introduced by CAPWAP, and there will be related terminology we must touch on in discussing these. The following figure illustrates some of the complex relationships between the various protocols, and illustrates where CAPWAP sits:

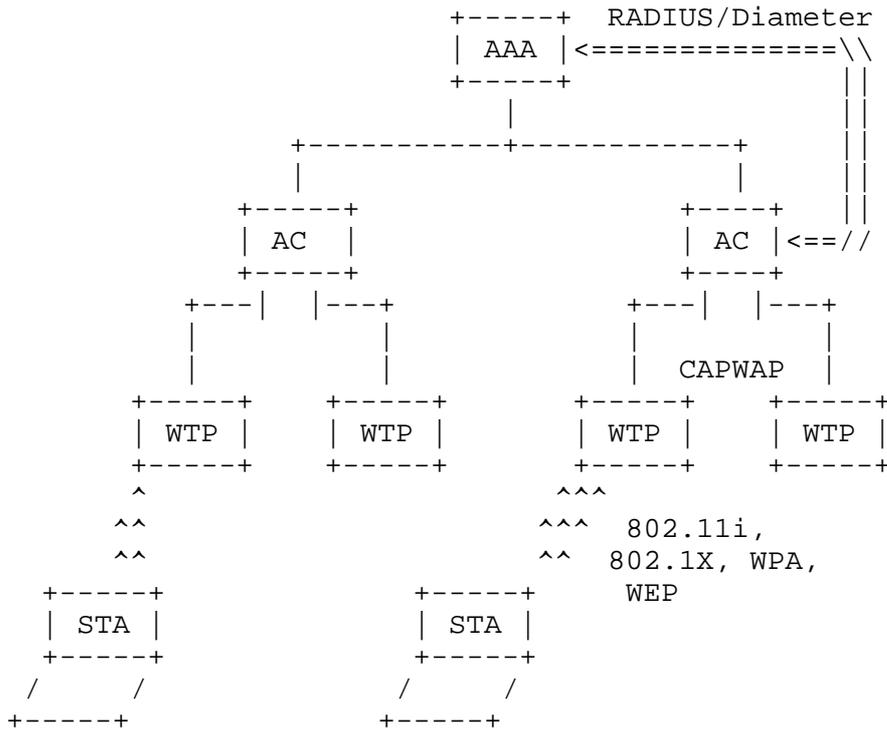


Figure 3: CAPWAP Protocol Hierarchy

CAPWAP is being inserted between the 802.11 link security mechanism and the authentication server communication, so to provide supporting context, we give a very brief overview of IEEE 802.11 and AAA security below. It is very important to note that we only cover those topics that are relevant to our discussion, so what follows is not by any means exhaustive. For more detailed coverage of IEEE 802.11-related security topics, see e.g., [80211SEC].

4.1. IEEE 802.11 Authentication and AAA

IEEE 802.11 provides for multiple methods of authentication prior to granting access to the network. In the simplest case, open authentication is used, and this is equivalent to no authentication. However, if IEEE 802.11 link security is to be provided, then some sort of authentication is required in order to bootstrap the trust relationship that underlies the associated key establishment process.

This authentication can be implemented through use of a shared secret. In such cases, the authentication may be implicitly tied to the link security mechanism, (e.g., when Wired Equivalent Privacy (WEP) is used with open authentication), or it may be explicit, e.g., when Wi-fi Protected Access is used with a Pre-Shared Key (WPA-PSK).

In other cases, authentication requires an explicit cryptographic exchange, and this operation bootstraps link security. In most such cases, IEEE 802.1X is used in conjunction with the Extensible Authentication Protocol [RFC3748] to authenticate either the client or both the client and the authentication server. This exchange produces cryptographic keying material for use with IEEE 802.11 security mechanisms.

The resulting trust relationships are complex, as can be seen from the following (simplified) figure:

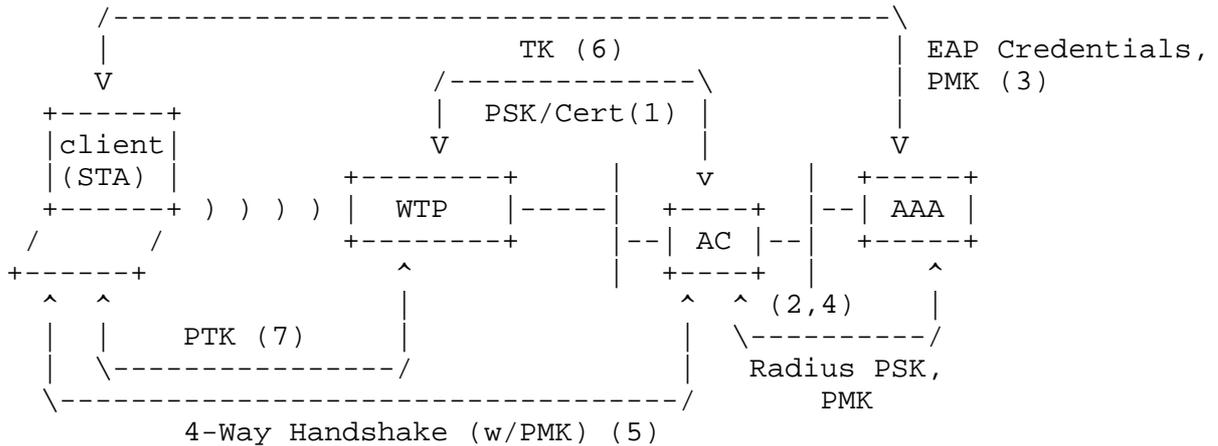


Figure 4: Trust Relationships

The following describes each of the relationships:

1. WTP and AC establish secure link
2. AC establishes secure link with AAA server
3. STA and AAA server conduct EAP, produce PMK
4. AAA server pushes PMK to AC
5. AC and STA conduct 4-way handshake (producing TK, among other keys)
6. AC pushes TK to WTP (if decentralized encryption is used)
7. WTP/STA use TK for IEEE 802.11 link security

4.2. IEEE 802.11 Link Security

The current CAPWAP binding for IEEE 802.11 primarily supports the use of IEEE 802.11i [80211I] security on the wireless link. However, since IEEE 802.11i does not prohibit the use of WEP for link security, neither does CAPWAP. Nonetheless, use of WEP with CAPWAP is NOT RECOMMENDED.

If WEP is used with CAPWAP, the CAPWAP system will inherit all the security problems associated with the use of WEP in any wireless network. In particular, 40-bit keys can be subject to brute-force attacks, and statistical attacks can be used to crack session keys of any length after enough packets have been collected [WEPSEC]. As of late 2008, such attacks are trivial, and take mere seconds to accomplish.

Newer link security mechanisms described in IEEE 802.11i, including TKIP and AES-CCMP, significantly improve the security of wireless networks. It is strongly RECOMMENDED that CAPWAP only be used with these newer techniques.

The only slight insecurity introduced by CAPWAP when using it with IEEE 802.11i is the handling of the KeyRSC counter. When performing decentralized encryption, this value is maintained by the WTP, but needed by the AC to perform the 4-way handshake. The value used during the handshake initializes the counter on the client. In CAPWAP, this value is initialized to zero, allowing the possibility for replay attacks of broadcast traffic in the window between initial authentication and the client receiving the first valid broadcast packet from the WTP. This slight window of attack has been documented in [RFC5416].

4.3. AAA Security

CAPWAP has very little control over how AAA security is deployed, as it's not directly bound to AAA as it is to IEEE 802.11. As a result, CAPWAP can only provide guidance on how to best secure the AAA portions of a CAPWAP-enabled wireless network.

The AAA protocol is a term that refers to use of either RADIUS [RFC3579] or Diameter [RFC4072] to transport EAP communications between the authenticator and the AAA server. Here the authenticator is the AC, thus the AAA protocol secures the link between the AC and AAA server. Use of non-unique or low-entropy long-term credentials to secure the AC-AAA link can severely impact the overall security of a CAPWAP deployment, and consequently is NOT RECOMMENDED. Each AC should have a mutually authenticated link that provides robust data

confidentiality and integrity. The AAA protocols and keys used SHOULD be consistent with the guidance in [RFC4962].

Since CAPWAP does not directly interact with AAA, it does not affect the overall security of a AAA network. In fact, by decreasing the number of devices that communicate with the AAA server, we can actually decrease its exposure and risk of attack.

4.4. Cryptographic Bindings

One key shortcoming of both the EAP and AAA models is that they are inherently two-party models. In CAPWAP deployments, we rely on a variety of security associations when an IEEE 802.11 WLAN client session is established. These include:

- o WTP-AC (CAPWAP Authentication)
- o AC-AAA (AAA Authentication)
- o STA-AAA (EAP Method Execution)
- o STA-AC (AAA pushes keys to AC)
- o STA-WTP (AC pushes keys to WTP)

Each of these security associations involve a pairwise trust relationship, and none result from a multi-party key agreement protocol such as Kerberos. In particular, just because an STA trusts a AAA server who trusts an AC who trusts a WTP doesn't necessarily mean that the STA should trust the WTP. The WTP may be compromised and using his credentials to maintain a trust relationship with an AC, while advertising false information about the network to an STA.

Protection against attacks like these can be very difficult, while maintaining scalable trust relationships with other entities in the network. Since multiple protocols are being used, multi-party keying to derive end-to-end trust relationships is infeasible.

Since CAPWAP is a collection of pairwise trust relationships, in order to claim CAPWAP is secure, we must believe in the transitivity of these trust relationships. Its hierarchal nature mitigates the domino effect, but any compromised device in the hierarchy can affect those below it in the hierarchy.

5. Structure of the Analysis

In order to conduct a comprehensive threat analysis, there are some basic questions we must answer:

- o Exactly what are we trying to protect?
- o What are the risks?
 - * What are the capabilities of would-be attackers?
 - * What might be goals of would-be attackers?
 - * What are the vulnerabilities or conditions they might attempt to exploit?
 - * For various attackers, what is the likelihood of attempting to exploit a given vulnerability given the cost of the exploit versus the value of attack?
- o What potential mitigation strategies are available to us?
- o What kinds of trade-offs do these mitigation strategies offer, and do they introduce any new risks?

This is a very simplistic overview of what we must accomplish below, but it should give some insight into the manner in which the discussion proceeds.

As a preliminary, we describe some representative CAPWAP deployment scenarios. This helps to frame subsequent discussion, so that we better understand what we are trying to protect. Following this, we describe a threat model in terms of adversary capabilities, vulnerabilities introduced by the CAPWAP functionality split, and a representative sampling of adversary goals. Note that we do not spend a lot of time speculating about specific attackers here, as this is a very general analysis, and there are many different circumstances under which a WLAN might be deployed.

Following the development of the general threat model, we describe appropriate countermeasures, and discuss how these are implemented through various means within the CAPWAP protocol. Finally, we discuss those issues that are not (or cannot be) completely addressed, and we give recommendations for mitigating the resulting exposure.

6. Representative CAPWAP Deployment Scenarios

In this section, we describe some representative CAPWAP deployment scenarios, and in particular, those scenarios that have been taken into consideration for the current CAPWAP protocol security design. For clarity, we first provide some preliminary definitions, along with descriptions of common deployment configurations that span multiple scenarios. Following this is a sampling of individual deployment scenarios.

6.1. Preliminary Definitions

The IEEE 802.11 standard describes several frame types, and variations in WLAN architectures dictate where these frame types originate and/or terminate (i.e., at the WTP or AC). There are three basic IEEE 802.11 frame types currently defined:

- o Control: These are for managing the transmission medium (i.e., the airwaves). Examples include RTS, CTS, ACK, PS-POLL, CF-POLL, CF-END, and CF-ACK.
- o Management: These are for managing access to the logical network, as opposed to the physical medium. Example functions include association/disassociation, reassociation, deauthentication, Beacons, and Probes.
- o Data: Transit data (network packets).

There are a number of other services provided by the WLAN infrastructure, including these:

- o Packet distribution
- o Synchronization
- o Retransmissions
- o Transmission Rate Adaptation
- o Privacy/Confidentiality/Integrity (e.g., IEEE 802.11i)

The manner in which these (and other) services are divided among the AC and WTP is collectively referred to in terms of "MAC-splitting" characteristics. We briefly describe various forms of MAC-splitting below. For further detail, see [RFC5415] and [RFC5416].

6.1.1. Split MAC

Split MAC scenarios are characterized by the fact that some IEEE 802.11 MAC messages are processed by the WTP, while some are processed by the AC. For example, a Split MAC implementation might divide IEEE 802.11 frame processing as follows:

WTP

- * Beacons
- * Probes
- * RTS, CTS, ACK, PS-POLL, CF-POLL, CF-END, CF-ACK

AC

- * Association/Reassociation/Disassociation
- * Authentication/Deauthentication
- * Key Management
- * IEEE 802.11 Link Security (WEP, TKIP, IEEE 802.11i)

The Split MAC model is not confined to any one particular deployment scenario, as we'll see below, and vendors have considerable leeway in choosing how to distribute the IEEE 802.11 functionality.

6.1.2. Local MAC

Local MAC scenarios are characterized by the fact that most IEEE 802.11 MAC messages are processed by the WTP. Some IEEE 802.11 MAC frames must be forwarded to the AC (i.e., IEEE 802.11 Association Request frames), but in general, the WTP manages the MAC interactions. Data frames may also be forwarded to the AC, but are generally bridged locally.

6.1.3. Remote MAC

Remote MAC scenarios are characterized by the fact that all IEEE 802.11 MAC messages are forwarded to the AC. The WTP does not process any of these locally. This model supports very lightweight WTPs that need be little more than smart antennas. While Remote MAC scenarios are not addressed by the current IEEE 802.11 protocol binding for CAPWAP, the description is included here for completeness.

6.1.4. Data Tunneling

Regardless of the approach to MAC splitting, there is also the matter of where user data packets are translated between wired and wireless frame formats, i.e., where the bridging function occurs. In some cases, user data frames are tunneled back to the AC, and in others, they are locally bridged by the WTP. While one might expect Remote MAC implementations to always tunnel data packets back to the AC, for Local and Split MAC modes the decision is not so clear. Also, it's important to note that there are no rules or standards in place that rigidly define these terms and associated handling. Data tunneling is further discussed for each scenario below.

6.2. Example Scenarios

In this section, we describe a number of example deployment scenarios. This is not meant to be an exhaustive enumeration; rather, it is intended to give a general sense of how WLANs currently are or may be deployed. This will provide important context when we discuss adversaries and threats in subsequent sections below.

6.2.1. Localized Modular Deployment

The localized modular model describes a WLAN that is deployed within a single domain of control, typically within either a single building or some otherwise physically contained area. This would be typical of a small to medium-sized organization. In general, the LAN enjoys some sort of physical security (e.g., it's within the confines of a building for which access is somehow limited), although the actual level of physical security is often far less than is assumed.

In such deployments, the WLAN is typically an extension of a wired LAN. However, its interface to the wired LAN can vary. For example, the interconnection between the WTPs and ACs can have its own wiring, and its only connection to the LAN is via the AC's Distribution System (DS) port(s). In such cases, the WLAN frequently occupies its own distinct addressing partition(s) in order to facilitate routing, and the AC often acts as a forwarding element.

In other cases, the WTPs may be mingled with the existing LAN elements, perhaps sharing address space, or perhaps somehow logically isolated from other wired elements (e.g., by Virtual LAN). Under these circumstances, it is possible that traffic destined to/from the WLAN is mixed with traffic to/from the LAN.

Localized deployments such as these could potentially choose any one of the MAC-splitting scenarios, depending on the size of the deployment, mobility requirements, and other considerations. In many

cases, any one of the various MAC-splitting approaches would be sufficient. This implies that user data may be bridged by either the WTP or the AC.

6.2.2. Internet Hotspot or Temporary Network

The Internet hotspot scenario is representative of a more general deployment model one might find at cafes, hotels, or airports. It is also quite similar to temporary WLANs, which are created for conferences, conventions, and the like. Some common characteristics of these networks include the following:

- o Primary function is to provide Internet access
- o Authentication may be very weak
- o There frequently is no IEEE 802.11 link security

Sometimes, the Local MAC model is used. In such cases, the AC may be "in the clouds" (out on the Internet somewhere), and user data traffic will typically be locally bridged, rather than tunneled back to the AC. Some IEEE 802.11 management traffic may be tunneled back to the AC, but frequently the authentication consists in simply knowing the Service Set Identifier (SSID) and perhaps a shared key for use with WEP or WPA-PSK.

In other cases, a Split MAC model may be used. This is more common in airport and hotel scenarios, where users may have an account that requires verification with some back-end infrastructure prior to access. In these cases, IEEE 802.11 management frames are tunneled back to the AC (e.g., user authentication), and stronger IEEE 802.11 link security may be provided (e.g., RSN), but user data is still typically locally bridged, as the primary goal is to provide Internet access.

A third variation on this scenario entails tunneling user data through a local AC in order to apply centralized policy. For example, in a hotel or airport scenario, there is no reason to provide direct access between WLAN users (who typically are within a private address space), and in fact, allowing such access might invite various sorts of malicious behavior. In such cases, tunneling all user data back to the (local) AC provides a security choke point at which policy may be applied to the traffic.

6.2.3. Distributed Deployments

The distributed deployment model describes a more complex WLAN topology that features network segments that are typically somehow spatially separated, although not necessarily so. These segments might be connected in a physically secure manner, or (if they are widely separated) they might be connected across potentially hostile hops. Below we discuss several subgroups of this model.

6.2.3.1. Large Physically Contained Organization

One distributed deployment example involves a single large organization that is physically contained, typically within one large building. The network might feature logically distinct (e.g., per-department or per-floor) network segments, and in some cases, there might be firewalls between the segments for access control. In such deployments, the AC is typically in a centralized datacenter, but there might also be a hierarchy of ACs, with a master in the datacenter, and subordinate ACs distributed among the network segments. Such deployments typically assume some level of physical security for the network infrastructure.

6.2.3.2. Campus Deployments

Some large organizations have networks that span multiple buildings. The interconnections between these buildings might be wired (e.g., underground cables), or might be wireless (e.g., a point-to-point Metropolitan Area Network (MAN) link). The interconnections may be secured, but sometimes they are not. The organization may be providing outdoor wireless access to users, in which case some WTPs will typically be located outdoors, and these may or may not be within physically secure space. For example, college campuses frequently provide outdoor wireless access, and the associated WTPs may simply be mounted on a light post.

For such deployments, ACs may be centrally located in a datacenter, or they may be distributed. User data traffic may be locally bridged, but more frequently it is tunneled back to the AC, as this facilitates mobility and centralized policy enforcement.

6.2.3.3. Branch Offices

A common deployment model entails an enterprise consisting of a headquarters and one or more branch offices, with the branches connected to the central HQ. In some cases, the site-to-site connection is via a private WAN link, and in others it is across the

Internet. For connections crossing potentially hostile hops (e.g., the Internet), some sort of Virtual Private Network (VPN) is typically employed as a protective measure.

In some simple branch office scenarios, there are only WTPs at the remote site, and these are managed by a controller residing at the central site. In other cases, a branch site may have its own subordinate controller, with the master controller again residing at the central site. In the first case, local bridging is often the best choice for user data, due to latency and link capacity issues. In the second case, traffic may be locally bridged by the WTPs, or it may be forwarded to the local subordinate controller for centralized policy enforcement. The choice depends on many factors, including local topology and security policy.

6.2.3.4. Telecommuter or Remote Office

It is becoming increasingly common to send WTPs home with employees for use as a telecommuting solution. While there are not yet any standards or hard rules for how these work, a fairly typical configuration provides Split MAC with all user data tunneled back to the controller in the organization's datacenter so that the WTP is essentially providing wireless VPN services. These devices may in some cases provide their own channel security (e.g., IPsec), but alternative approaches are possible. For example, there may be a stand-alone VPN gateway between the WTP and the Internet, which forwards all organization-bound traffic to the VPN gateway.

Similarly, it is becoming increasingly common for travelers to plug a WTP into a hotel broadband connection. While in many cases, these WTPs are stand-alone fat APs, in some cases they are configured to create a secure connection to a centralized controller back at headquarters, essentially forming a VPN connection. In such scenarios, a Split MAC approach is typical, but split-tunneling may also be supported (i.e., only data traffic destined for the headquarters is tunneled back to the controller, with Internet-bound traffic locally bridged).

7. General Adversary Capabilities

This section describes general capabilities we might expect an adversary to have in various CAPWAP scenarios. Obviously, it is possible to limit what an adversary can do through various deployment restrictions (e.g., provide strict physical security for the AC-WTP link), but such restrictions are simply not always feasible. For

example, it is not possible to provide strict physical security for various aspects of the telecommuter scenario. Thus, we must consider what capabilities an adversary might have in the worst case, and plan accordingly.

7.1. Passive versus Active Adversaries

One way to classify adversaries is in terms of their ability to interact with AC/WTP communications. For example, a passive adversary is one who can observe and perhaps record traffic, but cannot interact with it. They can "see" the traffic as it goes by, but they cannot interfere in any way, and they cannot inject traffic of their own. Note that such an adversary does not necessarily see all traffic -- they may miss portions of a communication, e.g., because some packets traverse a different path, or because the network is so busy that the adversary can't keep up (and drops packets as a result).

An active adversary, on the other hand, can directly interact with the traffic in real-time. There are two modes in which an active adversary might operate:

Pass-by (or sniffer)

- * Can observe/record traffic
- * Can inject packets
- * Can replay packets
- * Can reflect packets (i.e., send duplicate packets to a different destination, including the to the packet source)
- * Can cause redirection (e.g., via Address Resolution Protocol (ARP)/DNS poisoning)

Inline (Man-in-the-Middle, or MitM)

- * Can observe/record traffic
- * Can inject packets
- * Can replay packets
- * Can reflect packets (with or without duplication)
- * Can delete packets

- * Can modify packets
- * Can delay packets

A passive adversary could be located anywhere along the path between the AC and WTP, and is characterized by the fact that it only listens:

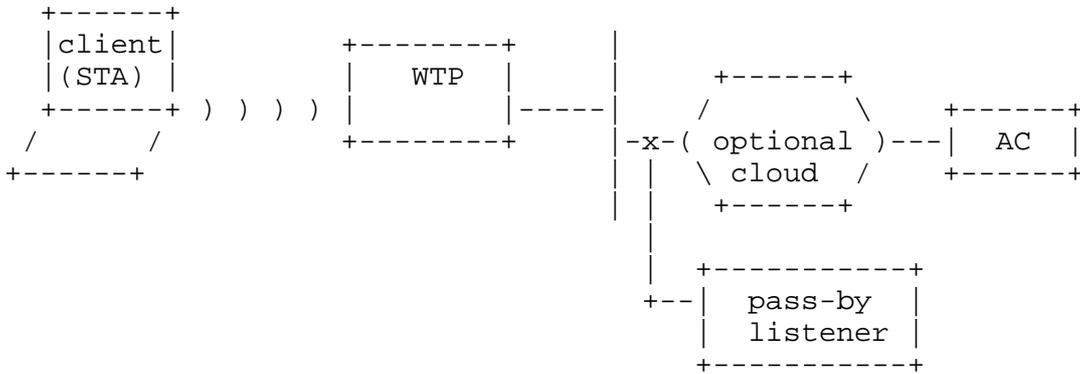


Figure 5: Passive Attacker

An active adversary may attach in the same manner as the passive adversary (in which case it is in pass-by mode), or it may be lodged directly in the path between the AC and WTP (inline mode):

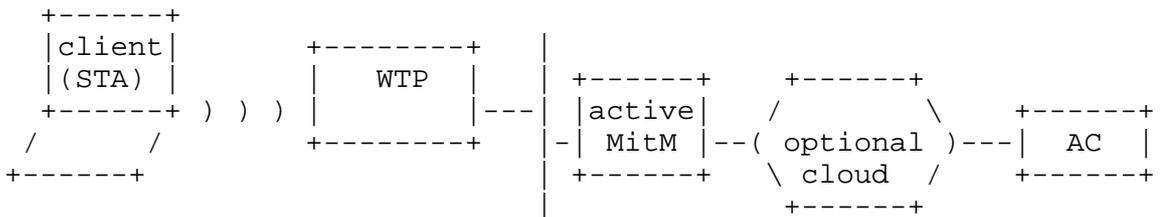


Figure 6: Active Man-in-the-Middle Attacker

If it is not inline, it can only observe and create traffic; if it is inline, it can do whatever it wishes with the traffic it sees.

It is important to recognize that becoming a MitM does not necessarily require physical insertion directly into the transmission path. Alternatively, the adversary can cause traffic to be diverted to the MitM system, e.g., via ARP or DNS poisoning. Hence, launching an MitM attack is not so difficult as it might first appear.

8. Vulnerabilities Introduced by CAPWAP

In this section, we discuss vulnerabilities that arise as a result of splitting the AP function across potentially hostile hops. We primarily consider network-based vulnerabilities, and while in particular we do not address how an adversary might affect a physical compromise of the WTP or AC, we do consider the potential effects of such compromises with respect to CAPWAP and the operational changes it introduces when compared to stand-alone AP deployments.

Functionally, we are interested in two general "states of being" with respect to AC-WTP communications: the session establishment phase or state, and the "connected" (or session established) state. We discuss each of these further below. Also, it is important to note that we first describe vulnerabilities assuming that the CAPWAP communications lack security of any kind. Later, we will evaluate the protocol given the security mechanisms currently planned for CAPWAP.

8.1. The Session Establishment Phase

The session establishment phase consists of two subordinate phases: discovery, and association or joining. These are treated individually in the following sections.

8.1.1. The Discovery Phase

Discovery consists of an information exchange between the AC and WTP. There are several potential areas of exposure:

Information Leakage: During Discovery, information about the WTP and AC hardware and software are exchanged, as well as information about the AC's current operational state. This could provide an adversary with valuable insights.

DoS Potential: During Discovery, there are several opportunities for Denial of Service (DoS), beyond those inherently available to an inline adversary. For example, an adversary might respond to a WTP more quickly than a valid AC, causing the WTP to attempt to join with a non-existent AC, or one which is currently at maximum load.

Redirection Potential: There are multiple ways in which an adversary might redirect a WTP during Discovery. For example, the adversary might pretend to be a valid AC, and entice the WTP to connect to it. Or, the adversary might instead cause the WTP to associate

with the AC of the adversary's choosing, by posing as a DNS or DHCP server, injecting a spoofed Discovery Response, or by modifying valid AC responses.

Misdirection: An adversary might mislead either the WTP or AC by modifying the Discovery Request or Response in flight. In this way, the AC and/or WTP will have a false view of the other's capabilities, and this might cause a change in the way they interact (e.g., they might not use important features not supported by earlier versions).

8.1.2. Forming an Association (Joining)

The association phase begins once the WTP has determined with which AC it wishes to join. There are several potential areas of exposure during this phase:

Information Leakage: During association, the adversary could glean useful information about hardware, software, current configuration, etc. that could be used in various ways.

DoS Potential: During formation of a WTP-AC association, there are several opportunities for Denial of Service (DoS), beyond those inherently available to an inline adversary. For example, the adversary could flood the AC with connection setup requests. Or, it could respond to the WTP with invalid connection setup responses, causing a connection reset. An adversary with MitM capability could delete critical session establishment packets.

Misdirection: An adversary might mislead either the WTP or AC by modifying the association request(s) or response(s) in flight. In this way, the AC and/or WTP will have an inaccurate view of the other's capabilities, and this might cause a change in the way they interact.

Some of these types of exposure are extremely difficult to prevent. However, there are things we can do to mitigate the effects, as we will see below.

8.2. The Connected Phase

Once the WTP and AC have established an association, the adversary's attention will turn to the network connection. If we assume a passive adversary, the primary concern for established connections is eavesdropping. If we assume an active adversary, there are several other potential areas of exposure:

Connection Hijacking: An adversary may assume the identity of one end of the connection and take over the conversation. There are numerous ways in which the true owner of the identity may be taken off-line, including DoS or MitM attacks.

Eavesdropping: An adversary may glean useful information from knowledge of the contents of CAPWAP control and/or data traffic.

Modification of User Data: An adversary with MitM capabilities could modify, delete, or insert user data frames.

Modification of Control/Monitoring Messages: An adversary with MitM capability could modify control traffic such as statistics, status information, etc. to create a false impression at the recipient.

Modification/Control of Configuration: An adversary with MitM capability could modify configuration messages to create unexpected conditions at the recipient. Likewise, if a WTP is redirected during Discovery (or joining) and connects to an adversary rather than an authorized AC, the adversary may exert complete control over the WTPs configuration, including potentially loading tainted WTP firmware.

9. Adversary Goals in CAPWAP

This section gives an sampling of potential adversary goals. It is not exhaustive, and makes no judgment as to the relative likelihood or value of each individual goal. Rather, it is meant to give some idea of what is possible. It is important to remember that clever attacks often result when seemingly innocuous flaws or vulnerabilities are combined in some non-intuitive way. Hence, we don't rule out some goal that, taken alone, might not seem to make sense from an adversarial perspective.

9.1. Eavesdrop on AC-WTP Traffic

There are numerous reasons why an adversary might want to eavesdrop on AC-WTP traffic. For example, it allows for reconnaissance, providing answers to the following questions:

- o Where are the ACs?
- o Where are the WTPs?
- o Who owns them?
- o Who manufactured them?

- o What version of firmware do they run?
- o What cryptographic capabilities do they have?

Similarly, snooping on tunneled data traffic might potentially reveal a great deal about the network, including answers to the following:

- o Who's using the WLAN?
- o When, and for how long?
- o What addresses are they using?
- o What protocols are they using?
- o What over-the-air security are they using?
- o Who/What are they talking to?

Additionally, access to tunneled user data could allow the attacker to see confidential information being exchanged by applications (e.g., financial transactions). An eavesdropper may gain access to valuable information that either provides the basis for another perhaps more sophisticated attack, or which has its own intrinsic value.

9.2. WTP Impersonation and/or Rootkit Installation

An adversary might want to impersonate or control an authorized WTP for many reasons, some of which we might realistically imagine today, and perhaps others about which we have no clue at this point. Examples we might reasonably imagine include the following:

- o to facilitate MitM attacks against WLAN users
- o to leak/inject or otherwise compromise WLAN data
- o to give an inaccurate view of the state of the WLAN
- o to gain access to a trusted channel to an AC, through which various protocol attacks might be launched (e.g., hijack client sessions from other WTPs)
- o to facilitate Denial-of-Service attacks against WLAN users or the network

9.3. AC Impersonation and/or Rootkit Installation

For reasons similar to the WTP impersonation discussed above, an adversary might want to impersonate an authorized AC for many reasons. Examples we might reasonably imagine include the following:

- o to facilitate DoS attacks against WLANs
- o to facilitate MitM attacks against WLAN users
- o to intercept user mobility context from another AC (possibly including security-sensitive information such as cryptographic keys)
- o to facilitate selective control of one or more WTPs
 - * modify WTP configuration
 - * load tainted firmware onto WTP
- o to assist in controlling which AC associates with which WTP
- o to facilitate IEEE 802.11 association of unauthorized WLAN user(s)
- o to exploit inter-AC trust in order facilitate attacks on other ACs

In general, AC impersonation opens the door to a large measure of control over the WLAN, and could be used as the foundation to many other attacks.

9.4. Other Goals or Sub-Goals

There are many less concrete goals an adversary might have which, taken alone, might not seem to have any purpose, but when combined with other goals/attacks, might have very definite and undesirable consequences. Here are some examples:

- o cause CAPWAP de-association of WTP/AC
- o cause IEEE 802.11 de-association of authorized user
- o inject/modify/delete tunneled IEEE 802.11 user traffic
 - * to interact with a specific application
 - * to launch various attacks on WLAN user systems

- * to launch protocol fuzzing or other attacks on the AC that bridges between IEEE 802.11 and 802.3 frame formats

- o control DNS responses
- o control DHCP/BOOTP responses

Anticipating all of the things an adversary might want to do is a daunting task. Part of the difficulty stems from the fact that we are analyzing CAPWAP in a very general sense, rather than in a specific deployment scenario with specific assets and very specific adversary goals. However, we have no choice but to take this approach if we are to provide reasonably good security across the board.

10. Countermeasures and Their Effects

In the previous sections, we described numerous vulnerabilities that result from splitting the AP function in two, and also discussed a number of adversary goals that could be realized by exploiting one or more of those vulnerabilities. In this section, we describe countermeasures we can apply to mitigate the risks that come with the introduction of CAPWAP into WLAN deployment scenarios.

10.1. Communications Security Elements

10.1.1. Mutual Authentication

Mutual authentication consists in each side proving its identity to the other. There are numerous authentication protocols that accomplish this, typically using either a shared secret (e.g., a pre-shared key) or by relying on a trusted third party (e.g., with digital credentials such as X.509 certificates).

Strong mutual authentication accomplishes the following:

- o helps prevent AC/WTP impersonation
- o helps prevent MitM attacks
- o can be used to limit DoS attacks.

10.1.1.1. Authorization

While authentication consists in proving the identity of an entity, authorization consists in determining whether that entity should have access to some resource. The current IEEE 802.11i CAPWAP protocol binding takes a rather simplistic approach to authorization,

depending on the authentication method chosen. If pre-shared keys are used, authorization is broad and coarse: if the device knows the pre-shared key, the device is "trusted", meaning that it is believed to be what it claims to be (AC versus WTP), and it is granted all privilege/access associated with that device class.

When using pre-shared keys, some granularity may be achieved by creating classes, each with their own pre-shared key, but this still has drawbacks. For example, while possession of the key may suffice to identify the device as a member of a given group or class, it cannot be used to prove a device is either a WTP or an AC. This means the key can be abused, and those possessing the key can claim to be either type of device.

When X.509v3 certificates are used for authentication, much more granular authorization policies are possible. Nonetheless, the current IEEE 802.11i protocol binding remains simplistic in its approach (though this may be addressed in future revisions). As currently defined, the X.509v3 certificates facilitate the following authorization checks:

- o CommonName (CN): the CN contains the MAC address of the device; if the relying party (AC or WTP) has a method for determining "acceptability" of a given MAC address, this helps prevent AC/WTP impersonation, MitM attacks, and may help in limiting DoS attacks
- o Extended Key Usage (EKU) key purpose ID bits: CAPWAP uses specific key purpose ID bits (see [RFC5415] for more information) to explicitly differentiate between an AC and a WTP. If use of these bits is strictly enforced, this segregates devices into AC versus WTP classes, and assists in preventing AC/WTP impersonation, MitM attacks, and may also help in limiting DoS attacks. However, if the `id-kp-anyExtendedKeyUsage` `keyPurposeID` is supported, this mechanism may be on a par with pre-shared keys, as a rogue device has the ability to claim it is either a WTP or AC, unless CN-based access controls are employed in tandem.

It should be noted that certificate-based authorization and zero configuration are not fully compatible. Even if the WTPs and the ACs are shipped with manufacturer-provided certificates, the WTPs need a way to identify the correct AC is in this deployment (as opposed to other ACs from the same vendor, purchased and controlled by an adversary), and the AC needs to know which WTPs are part of this deployment (as opposed to WTPs purchased and controlled by an adversary).

The threat analysis in this document assumes that WTPs can identify the correct AC, and the AC can identify the correct WTPs. Analysis of situations where either of these do not hold is beyond the scope of this document.

10.1.2. Data Origin Authentication

Data origin authentication typically depends on first authenticating the party at the other end of the channel, and then binding the identity derived from that authentication process to the data origin authentication process. Data origin authentication primarily prevents an attacker from injecting data into the communication channel and pretending it was originated by a valid endpoint. This mitigates risk by preventing the following:

- o packet injection
- o connection hijacking
- o modification of control and/or user data
- o impersonation

10.1.3. Data Integrity Verification

Data integrity verification provides assurance that data has not been altered in transit, and is another link in the trust chain beginning from mutual authentication, extending to data origin authentication, and ending with integrity verification. This prevents the adversary from undetectably modifying otherwise valid data while in transit. It effectively prevents reflection and modification, and in some cases may help to prevent re-ordering.

10.1.4. Anti-Replay

Anti-replay is somewhat self-explanatory: it prevents replay of valid packets at a later time, or to a different recipient. It may also prevent limited re-ordering of packets. It is typically accomplished by assigning monotonically increasing sequence numbers to packets.

10.1.5. Confidentiality

Data confidentiality prevents eavesdropping by protecting data as it passes over the network. This is typically accomplished using encryption.

10.2. Putting the Elements Together

Above we described various security elements and their properties. Below, we discuss combinations of these elements in terms of CAPWAP.

10.2.1. Control Channel Security

The CAPWAP control channel is used for forming an association between a WTP and AC, and subsequently it allows the AC to provision and monitor the WTP. This channel is critical for the control, management, and monitoring of the WLAN, and thus requires all of the security elements described above. With these elements in place, we can effectively create a secure channel that mitigates almost all of the vulnerabilities described above.

10.2.2. Data Channel Security

The CAPWAP data channel contains some IEEE 802.11 management traffic including association/disassociation, reassociation, and deauthentication. It also may contain potentially sensitive user data. If we assume that threats to this channel exist (i.e., it traverses potentially hostile hops), then providing strong mutual authentication coupled with data origin/integrity verification would prevent an adversary from injecting and/or modifying transit data, or impersonating a valid AC or WTP. Adding confidentiality would prevent eavesdropping.

11. Countermeasures Provided by DTLS

Datagram TLS (DTLS) is the currently proposed security solution for CAPWAP. DTLS supports the following security functionality:

- o Mutual Authentication (pre-shared secrets or X.509 Certificates)
- o Mutual Authorization (pre-shared secrets or X.509 Certificates)
- o Data Origin Authentication
- o Data Integrity Verification
- o Anti-replay
- o Confidentiality (supports strong cryptographic algorithms)

Using DTLS for both the control and data channels mitigates nearly all risks resulting from splitting the AP function in two.

12. Issues Not Addressed By DTLS

Unfortunately, DTLS does not solve all of our CAPWAP security concerns. There are some things it just cannot prevent. These are enumerated below.

12.1. DoS Attacks

Even with the security provided by DTLS, CAPWAP is still susceptible to various types of DoS attack:

- o **Session Initialization:** an adversary could initiate thousands of DTLS handshakes simultaneously in order to exhaust memory resources on the AC; DTLS provides a mitigation tool via the HelloVerifyRequest, which ensures that the initiator can receive packets at the claimed source address prior to allocating resources. However, this would not thwart a botnet-style attack.
- o **Cryptographic DoS:** an adversary could flood either the AC or WTP with properly formed DTLS records containing garbage, causing the recipient to waste compute cycles decrypting and authenticating the traffic.
- o **Session interference:** a MitM could either drop important session establishment packets or inject bogus connection resets during session establishment phase. It could also interfere with other traffic in an established session.

These attacks can be mitigated through various mechanisms, but not entirely avoided. For example, session initialization can be rate-limited, and in case of resource exhaustion, some number of incompletely initialized sessions could be discarded. Also, such events should be strictly audited.

Likewise, cryptographic DoS attacks are detectable if appropriate auditing and monitoring controls are implemented. When detected, these events should (at minimum) trigger an alert. Additional mitigation might be realized by randomly discarding packets.

Session interference is probably the most difficult of DoS attacks. It is very difficult to detect in real-time, although it may be detected if exceeding a lost packet threshold triggers an alert. However, this depends on the MitM not being in a position to delete the alert before it reaches its appropriate destination.

12.2. Passive Monitoring (Sniffing)

CAPWAP protocol security cannot prevent (or detect) passive monitoring. The best we can do is provide a confidentiality mechanism.

12.3. Traffic Analysis

DTLS provides no defense for traffic analysis. If this is a concern, there are traffic generation and padding techniques designed to mitigate this threat, but none of these are currently specified for CAPWAP.

12.4. Active MitM Interference

This was discussed in more limited scope in the section above on DoS attacks. An active MitM can delete or re-order packets in a manner that is very difficult to detect, and there is little the CAPWAP protocol can do in such cases. If packet loss is reported upon exceeding some threshold, then perhaps detection might be possible, but this is not guaranteed.

12.5. Other Active Attacks

In addition to the issues raised above, there are other active attacks that can be mounted if the adversary has access to the wired medium. For example, the adversary may be able to impersonate a DNS or DHCP server, or to poison either a DNS or ARP cache. Such attacks are beyond the scope of CAPWAP, and point to an underlying LAN security problem.

13. Security Considerations

This document outlines a threat analysis for CAPWAP in the context of IEEE 802.11 deployments, and as such, no additional CAPWAP-related security considerations are described here. However, in some cases additional management channels (e.g., Simple Network Management Protocol (SNMP)) may be introduced into CAPWAP deployments. Whenever this occurs, related security considerations MUST be described in detail in those documents.

14. Acknowledgements

The authors gratefully acknowledge the reviews and helpful comments of Dan Romascanu, Joe Salowey, Sam Hartman, Mahalingham Mani, and Pasi Eronen.

15. References

15.1. Normative References

- [80211I] "IEEE Std 802.11i: WLAN Specification for Enhanced Security", April 2004.
- [80211SEC] Edney, J. and W. Arbaugh, "Real 802.11 Security - Wi-Fi protected Access and 802.11i", 2004.
- [8021X] "IEEE Std 802.1X-2004: Port-based Network Access Control", December 2004.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4118] Yang, L., Zerfos, P., and E. Sadot, "Architecture Taxonomy for Control and Provisioning of Wireless Access Points (CAPWAP)", RFC 4118, June 2005.
- [RFC5415] Calhoun, P., Ed., Montemurro, M., Ed., and D. Stanley, Ed., "Control And Provisioning of Wireless Access Points (CAPWAP) Protocol Specification", RFC 5415, March 2009.
- [RFC5416] Calhoun, P., Ed., Montemurro, M., Ed., and D. Stanley, Ed., "Control And Provisioning of Wireless Access Points (CAPWAP) Protocol Binding for IEEE 802.11", RFC 5416, March 2009.

15.2. Informative References

- [RFC3579] Aboba, B. and P. Calhoun, "RADIUS (Remote Authentication Dial In User Service) Support For Extensible Authentication Protocol (EAP)", RFC 3579, September 2003.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowitz, "Extensible Authentication Protocol (EAP)", RFC 3748, June 2004.
- [RFC4072] Eronen, P., Hiller, T., and G. Zorn, "Diameter Extensible Authentication Protocol (EAP) Application", RFC 4072, August 2005.
- [RFC4962] Housley, R. and B. Aboba, "Guidance for Authentication, Authorization, and Accounting (AAA) Key Management", BCP 132, RFC 4962, July 2007.

[WEPSEC] Petroni, N. and W. Arbaugh, "The Dangers of Mitigating Security Design Flaws: A Wireless Case Study", January 2003.

Authors' Addresses

Scott G. Kelly
Aruba Networks
1322 Crossman Ave
Sunnyvale, CA 94089
US

EMail: scott@hyperthought.com

T. Charles Clancy
DoD Laboratory for Telecommunications Sciences
8080 Greenmead Drive
College Park, MD 20740
US

EMail: clancy@LTSnet.net

