

Network Working Group
Internet-Draft
Expires: April 19, 2004

M. Mani
Avaya Inc.
B. O'Hara
Airespace
L. Yang
Intel Corp.
October 20, 2003

**Architecture for Control and Provisioning of Wireless Access
Points(CAPWAP)
draft-mani-ietf-capwap-arch-00**

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on April 19, 2004.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

While conventional wisdom has it that Wireless Access Points are strictly Layer 2 bridges, such devices today perform some higher layer functions of routers or switches in wired Infrastructure in addition to bridging the wired and wireless networks. For example, in 802.11 networks, Access Points can function as Network Access Servers. For this reason, Access Points have IP addresses and can function as IP devices.

This Document analyzes WLAN (Wireless LAN) functions and services; and describes a flexible balance of such AP (Access Point) functions as allowed in the Standards and practiced in the industry, to be meaningfully split between lightweight Access Point (LAP) framework and AP Controllers or AR (Access Router) framework managing them.

Table of Contents

1.	Introduction	4
1.1	Conventions used in this document	4
1.2	CAPWAP Purpose and Scope	4
1.3	Document Organization	4
2.	Terminology	5
3.	Overview	7
3.1	The IEEE 802.11 in Brief	7
3.2	CAPWAP Motivation	9
3.3	AP to AR Network Topology Considerations	10
4.	CAPWAP Component Architecture	12
4.1	WLAN functions and Services	12
4.1.1	Access Point Functions and Services	14
4.1.2	Access Controller Functions and Services	14
4.1.3	Other Conventional WLAN Functions and Services	15
4.1.4	Architectural Trends	15
4.2	CAPWAP Network Topology	16
4.2.1	Functional Distribution of WLAN Services	16
4.2.2	AP to AR Topologies	16
4.3	CAPWAP Security	18
4.3.1	WLAN Security	19
4.3.2	Mutual Authentication of AP and AR	20
4.3.3	Path Security of AP and AR	21
4.4	AP Provisioning	21
4.4.1	AP Identity	21
4.4.2	AP Configuration	21
4.4.3	Access Router Availability	21
4.5	Access Point Service Management	22
4.5.1	Monitoring	22
4.5.2	Control	22
4.6	Access Router Discovery	23
4.6.1	Access Router Availability	23
4.6.2	Capabilities Negotiation	23
4.7	Summary	24
5.	Prior Work	25
5.1	DIRAC	25
5.2	ForCES	26
5.2.1	Similarities in Objectives and Architectural Considerations	26
5.2.2	Overlap in Topology Considerations	27
5.2.3	Differences in Design Approach	27

5.2.4	Differences in the Functionality Controlled	27
5.2.5	Similarities in Security Requirements	27
5.2.6	Difference in Operation Scope	28
5.2.7	Comparison in Protocols	28
6.	Security Considerations	29
7.	Acknowledgements	30
	References	31
	Authors' Addresses	32
	Intellectual Property and Copyright Statements	33

1. Introduction

1.1 Conventions used in this document

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](#) [7].

1.2 CAPWAP Purpose and Scope

The purpose of CAPWAP work is to define the framework reflecting the architectural trend that delegates and aggregates selected WLAN functions and services from APs to ARs to enhance WLAN resource management. On the basis of such definition CAPWAP aims to provide a secure protocol to enable AP-to-AR communications and AP provisioning & management.

1.3 Document Organization

Overview section describes the IEEE 802.11 WLAN architecture and services in brief followed by AP-AR network topological considerations leading to CAPWAP motivation.

Subsequent section describes the CAPWAP architecture and its components.

The section that follows discusses related research work and an applicable standards topology.

The document concludes with Security Considerations which are also discussed in Architecture.

2. Terminology

LWAP: Lightweight Access Point

AB/AR: Access Bridges/Routers

AC: Access Controllers

AP: access point

BSS: basic service set

ESS: extended service set

SSID: service set identifier

WLAN: wireless local area network

RSN: robust security network

TSN: transition security network

PMK: pair-wise master key

PTK: pair-wise transient key

TK: temporal key

GMK: group master key

GTK: group transient key

KCK: key confirmation key

KEK: key encryption key

PSK: pre-shared key

WEP: wired equivalent privacy

Throughout the document the terminologies of AR (Access Router), AC (Access Controller) and AB (Access Bridge) are used synonymously in contexts of allowable network topology arguments. In other cases the distinction is called out explicitly.

However, at the outset AC is to be assumed the generic term for the entity with which an AP registers or associates - which terms will be qualified later in the CAPWAP architecture sections ([Section 4](#)).

AR is called out in the context that the Access Controller that an AP associates with is over an allowed L3 cloud between the wired network backend of APs and the ACs.

Access Bridge (or WLAN switch) is called out in the context of such network cloud or connectivity may be over a predominantly L2 network.

It may be observed in following sections that the proposed architecture chooses to stay agnostic and equivalent to either Network Protocol and focuses on the interface and generic encapsulation that shall allow for both. The Protocol MAY end up specifying merely for IP.

3. Overview

Prior to setting out on details, a snapshot of WLAN standards are in order to put the CAPWAP motivation and standardization benefits in perspective, particularly when the required interfaces appear in the landscape bordering L2 and L3 standards scope.

3.1 The IEEE 802.11 in Brief

The IEEE 802.11 standard for wireless local area networks [1] specifies a MAC protocol, several PHYs, and a MAC management protocol. Each of these operates over the air, between two or more 802.11 devices. 802.11 also describes how mobile devices can associate together into a basic service set (BSS), the rough equivalent of a single broadcast domain or a segment of a bridged Ethernet LAN. A BSS is identified by a common service set identifier (SSID) or name. An SSID is an arbitrary byte string, up to 32 bytes long, though most implementations utilize ASCII strings for readability. 802.11 also describes the functionality of a specific device, called an access point (AP), that translates frames between mobile 802.11 devices and hosts on a wired network. When more than one AP is connected via a broadcast layer 2 network and all are using the same SSID, an extended service set (ESS) is created. An ESS is also similar to a single broadcast domain, where a mobile device associated with one AP can successfully ARP for the address of a mobile device associated with any other AP in the ESS. Within an ESS, a mobile station can roam from one AP to another through only layer 2 transitions coordinated by the 802.11 MAC management protocol. Higher layer protocols, including IP are unaware that the network attachment point of the mobile device has moved.

The 802.11 working group is currently proceeding on work related to layer 2 security and quality of service. The 802.11i task group is addressing the security issues of the original 802.11 standard in the areas of authentication and encryption. This work refers to other standards, including 802.1X Port Based Access Control [14] and the Extensible Authentication Protocol [9]. The 802.11e task group is addressing layer 2 quality of service items through extending the access method, frame definitions, and MAC management protocol of the original standard. This work refers to the 802.1Q [15] standard.

802.11 PHYs are wireless, by definition, and principally use radio technology for communication. An aspect of an 802.11 WLAN that is not addressed by the standard is the necessity to manage the self-interference of one AP when operating on a radio channel equal to or near the radio channel of another AP within reception range. Managing self-interference within the WLAN involves both measurement of the level of interference, as well as control of the transmit

power and transmitting channel in each of the APs. Work currently in process in the 802.11k task group is addressing the issue of radio resource measurement, which will provide the information on level of interference, among other things.

Some definitions of 802.11 terminology is in order, since it is unique to the 802.11 standard.

- * "Distribution" is the service of forwarding MSDUs for an associated station by an AP. As it is described in 802.11, distribution by an AP is providing sufficient information to enable a frame received from an associated station to be successfully delivered to its proper destination. For the most part, this involves translating the frame format from 802.11 to Ethernet (typically) and removing any SNAP encapsulation that was applied to the 802.11 frame, due to its lack of an equivalent to the Ethertype field. This is similar to standard bridging, except that 802.11 APs are not 802.1D bridges. APs typically do not implement spanning tree protocols or algorithms. They are considered to be edge devices, connected only to leaf nodes with no further bridging taking place down stream from them. This is not always a valid assumption and can sometimes result in unanticipated bridging loops.
- * "Integration" is a concept unique to 802.11 that is a result of the underlying architecture. 802.11 considers that the individual APs that make up a WLAN, an extended service set (ESS) in 802.11 terminology, are connected by a closed system, called the distribution system. Only frames that are "within" the ESS are carried by the distribution system. This includes frames that are moving from one AP to another for delivery to a mobile station, frames received from outside the ESS for delivery to a mobile station, and frames from a mobile station to be delivered outside the ESS. Connecting the closed distribution system to the outside world is a "portal". The portal is the single point at which the distribution system exchanges frames with the network outside of the ESS.

The problem with the 802.11 architecture, or maybe just with the AP implementations, is that AP implementations do not adhere to this architecture. An AP typically implements both the distribution and integration services, and the portal function, inside the skin of the AP. In this sense, every AP is its own, isolated ESS and no APs actually implement the architecture described in the standard. When a set of APs is connected together to create a WLAN, what is actually created is a set of independent ESSs that happen to communicate, in spite of the implementation.

In addition to the 802.11 standard, the 802.11 working group produced a recommended practice for inter-AP communication, 802.11F [12]. The recommended practice describes the use of a new application protocol, the Inter-AP Protocol (IAPP), carried on UDP or TCP. It permits APs to exchange information about roaming mobile devices, including an envelope for general context transfer purposes, and to push layer 2 keying information to neighboring APs in preparation for the roaming of mobile devices to those neighboring APs. The recommended practice specifies the use of 802.2 XID frames for updating layer 2 devices when a mobile device's point of attachment to the network has changed due to a roaming event. 802.11F also specifies the use of RADIUS for the authentication of one AP to another and, along with portions of the IAPP protocol, to establish secure IAPP packets exchanged between participating APs.

The IAPP is not applicable to this architecture, though it may be implemented in the access controller for communication with other access controllers as 802.11 intended it to be used between individual APs. It is not applicable within the CAPWAP architecture because, presumably, the communications defined by CAPWAP would be internal to the access controller and not require such a protocol to be utilized.

3.2 CAPWAP Motivation

As evidenced over the past few months, there is overwhelming support in the market for a new WLAN architecture. This architecture moves much of the functions that would reside in a traditional access point (AP) to a centralized access router (AR). Some of the benefits that come out of this new architecture include:

- o Ease of Use: By centrally managing a WLAN as a system rather than as a series of discrete components, management and control of the WLAN is much easier
- o Increased Security: Having a centralized AR enforce policies and being able to detect potential threats across a much larger RF domain increases the security of the network.
- o Enhanced Mobility: By terminating the WLAN "management" protocol in the AR, these messages may be used as "mobility triggers", providing mobility across an RF domain without the need for any client software.
- o Quality of Service: By allowing the centralized AR manage the RF links, offers systemic perspective to perform efficient load balancing across multiple Access Points - thus increasing the efficiency of the wireless network. It also offers scope to have

higher layer applications influence roaming and placement policies in a streamlined manner.

All of the above can be providing by terminating the 802.11 management frames in the AR. This approach is also commonly referred to as Split AP, where the real-time components of the 802.11 protocol are handled in the Access Point, while the access control components of the 802.11 protocol terminate in the Access Router.

Having a module in the AR that understands 802.11 management frames and 802.11 WLANs will provide much better control and optimization of the WLAN operation than will an abstract, protocol-agnostic control module. Adding support to CAPWAP/LWAPP for other wireless technologies then becomes a task of encapsulating the new frames and adding a new control module to the AR to handle the new technology. Presumably, the LWAPP protocol and CAPWAP architecture will need little, if any change.

3.3 AP to AR Network Topology Considerations

APs and ARs are linked directly as required by some architectures. Among such classifications

1. ARCH0: The classic AP is at one of the spectrum interfaces to the Infrastructure Network cloud with no specific connectivity to a controller. In this case the AP can be considered to have a self-contained controller possibly communicating with other APs in the ESS to form a WDS.
2. ARCH1: APs which defer all WLAN functions other than real-time services ([Section 4.1](#)) create a vastly different paradigm of vertical (real-time frontend AP and aggregated backend AC) functional distribution calling for a trust model between the two and a discovery process of AC by AP. The latter (discovery) is accentuated when the connectivity is through a cloud and there's potential for m-to-n correspondence of AP-AR.
3. ARCH2: APs which tend to shift some normally real-time functions as well to the backend with benefits such as extending OTA (over-the-air) protection for AP-AR thus allowing for an extended Trust Model for client data.
4. ARCH3: There's the case which carries (3) to render the AC as a single "AP-switch" treating all connected APs as smart antennae.

While, at the outset, the architectures seem at wider variance, the varied market requirements of

1. deployment scope
2. scalability
3. performance and
4. end-end security demands

seem to allow for all such architectures to have a role with varying scope and limitations. This further underscores the argument to provide a negotiable interface protocol.

4. CAPWAP Component Architecture

Given the preliminary outline of the three primary architecture types (and a fourth variant) in [Section 3.3](#) the predominant architectural components are presented in three perspectives:

1. Functional & Service-based (WLAN standards)
2. Architectural Split
3. Topological

This is required as a means to realize the way the three aspects are inter-dependent.

The Figure 1 illustrates the basic outline of communications architecture between AP & AC.




Figure 1: Basic Communications Framework

4.1 WLAN functions and Services

The IEEE 802.11 standard [1] says very little about the functionality

required of an AP. There is some discussion of the AP at a block diagram level, in the General Description in clause 5 of the standard. There, an AP is described as containing functional blocks for 802.11 station services and for distribution system services. Station services consist of the following four services:

- a) Authentication
- b) Deauthentication
- c) Privacy
- d) MSDU Delivery

Distribution system services consist of the following five services:

- a) Association
- b) Disassociation
- c) Distribution
- d) Integration
- e) Reassociation

There are additional services that are required of an AP, that are described in the MAC Layer Management Entity (MLME) in clause 11. These additional management services are

- a) Beacons
- b) Synchronization
- c) Power Management

Other functionality that is not described, except implicitly in the MIB, is control and management of the radio-related functions of an AP. These include:

- a) Channel Assignment
- b) Transmit Power Control
- c) Clear Channel Assessment

- d) Radio Resource Measurement (work currently under way in IEEE 802.11k)

The 802.11h [13] amendment to the base 802.11 standard specifies the operation of a MAC management protocol to accomplish the requirements of some regulatory bodies (principally in Europe, but expanding to others) in these areas:

- a) RADAR detection
- b) Transmit Power Control
- c) Dynamic Channel Selection

4.1.1 Access Point Functions and Services

The services that MUST be in a lightweight AP are those that are directly related to the real-time aspects of the 802.11 MAC protocol and those related to the radio nature of an 802.11 AP. These functions are:

- a) Privacy
- b) MSDU Delivery
- c) Beaconsing
- d) Synchronization
- e) Power Management
- f) Channel Assignment
- g) Transmit Power Control
- h) Clear Channel Assignment
- i) Radio Resource Measurement
- j) RADAR detection

4.1.2 Access Controller Functions and Services

The functions that MAY be moved from the lightweight AP and located in the AR are those dealing with the management and control aspects of an 802.11 AP. These are the distribution system services, in

addition to authentication and deauthentication services. These functions are:

- a) Authentication
- b) Deauthentication
- c) Association
- d) Disassociation
- e) Reassociation
- f) Distribution
- g) Integration
- h) Dynamic Channel Selection
- i) Dynamic Control of transmit power

4.1.3 Other Conventional WLAN Functions and Services

"Heavy" Access Points being the bridge to the wired world MAY (and normally do) also support various services and protocols that provide seamless connectivity of WLAN clients to the wired network such as

- a) Port and Protocol-based VLANs
- b) SNMP
- c) QoS (DiffServ and 802.1Q) mapping
- d) IP routing
- e) DHCP relay/server
- f) RADIUS client/proxy
- g) MobileIP (client proxy)

Based on the definition of lightweight access points these services SHOULD qualify for offloading to the AR.

4.1.4 Architectural Trends

4.2 CAPWAP Network Topology

The CAPWAP network topology primarily consists of the WLAN topology and the AP-AC (AP-AR) topology.

The WLAN topology is straightforward and is as described in Overview section. This is not of much current interest as the relevant portal variants of WLAN are applicable equally to all new AP-AC topologies.

4.2.1 Functional Distribution of WLAN Services

Functional distribution of WLAN services described in earlier sub-sections are partly an artifact of the architecture types ARCH0-3. However, they may result in AP-AC topological constraints. Such constraints include direct connectivity to the AC being required and in most cases mandate L2 connectivity.

4.2.2 AP to AR Topologies

CAPWAP assumes that the AR and AP are within the same administrative domain, i.e. they are owned/controlled by the same entity. CAPWAP makes no topological assumptions beyond these, meaning there are several topologies which must be considered for our purposes.




Figure 2: Directly Connected




Figure 3: Switched Connections




Figure 4: Routed Connections

4.3 CAPWAP Security

The CAPWAP architecture spans more than the topology over the air. IEEE 802.11 (and now 802.11i) describes the single-hop security over-the-air.

The resulting security scheme only protects the data frames (multicast, broadcast and unicast) between stations and AP.

This leaves a security gap in the CAPWAP topology between AP and AC (AR).

As discussed earlier security of control and management traffic between the AP and AR subsystem, needs to be secured, failing which control of AP can be compromised.

In addition there may be explicit requirements to secure the data flow between AP and AR segment. This is end-end traffic between WLAN stations and their WLAN/wireline destinations invariant to CAPWAP considerations. Protection of this traffic in this segment may incidentally ensue in architectures such as in ARCH2 and ACRH3.

4.3.1 WLAN Security

802.11 provides layer 2 authentication and privacy services. Severe deficiencies have been documented in the mechanisms of the original standard. The current task group, 802.11i, is completing work on an amendment to the standard that addresses these deficiencies. The requirements for 802.11i are more difficult than simply providing the desired level of protection for the information carried by 802.11 frames in new equipment. 802.11i must also provide a mechanism that can be used by equipment already deployed, to eliminate the deficiencies of the original standard to which the equipment was built.

WLAN Security offers over-the-air single-hop MAC-layer frame security for data frames between Mobile Stations and AP's. This is built on top of 802.1X-based authentication and Session and Data-encryption Key Exchanges derived thereof.

4.3.1.1 Authentication - EAP over LAN (802.1X)

802.11i specifies an extensible authentication method, based on negotiation between the AP and mobile device that occurs during the association process. After successfully negotiating the particular authentication method to be used, the mobile device is allowed to associate, but must immediately complete the negotiated authentication before any data exchange will be permitted.

The default mechanism for authentication defined by 802.11i is to piggyback on the 802.1X standard, using EAP to authenticate the mobile device or user after 802.11 association with an AP has completed. In the terms defined in 802.1X, an AP is an authenticator and a mobile device is a supplicant. The AP, as authenticator, proxies the supplicant's communications to an authentication system. An example authentication system is a RADIUS server. The AP communicates with the RADIUS server, as a RADIUS proxy for the client, using EAP. The AP is responsible for blocking the (logical) controlled port for the associated device until the successful completion of the 802.1X authentication. At the conclusion of the 802.1X authentication, keying material is available to both the mobile device and AP that can be used for frame security.

4.3.1.2 Frame Security (802.11i)

802.11i Frame Security offers Encryption, Message Integrity and Replay Protection services.

To meet the requirements of improving security for both existing devices and new devices, 802.11i specifies two security mechanisms,

the Transition Security Network (TSN) and the Robust Security Network (RSN). Both TSN and RSN utilize keying material derived from an 802.1X-based authentication exchange to deliver a pair-wise master key (PMK) to both the mobile device and AP. TSN can also use a preshared key (PSK) to derive a PMK without the use of an authentication exchange. From the PMK is derived a pair-wise transient key (PTK). The PTK is used to create a pair-wise temporal key (TK), an EAPOL key encryption key (KEK), and an EAPOL key confirmation key (KCK). An 802.1X exchange is also used to deliver the keying material to derive a group master key (GMK). From the GMK is derived a group transient key (GTK). The GTK is used to create a group temporal key (TK) used by the AP to encrypt multicast frames and by the mobile devices to decrypt multicast frames.

TSN specifies a means to improve the security of equipment built with the original RC4-based wired equivalent privacy (WEP) cipher. TSN requires that the encryption key used with WEP be rotated on every packet. TSN specifies the algorithm for this key rotation, based on the pair-wise TK and the frame sequence counter. In addition, TSN specifies an algorithm for a keyed message integrity code (MIC) (more often called message authentication code (MAC), but that acronym is already utilized in the 802.11 standard), called Michael. Michael is a compromise between strength and computational requirements, because this must operate on legacy equipment with fixed computational capabilities. As a result, TSN also specifies some rather severe countermeasures to be implemented when an attack against the MIC is suspected.

RSN specifies an encapsulation and algorithm for new equipment that is significantly stronger than either WEP or TSN. The algorithm is an AES mode called Counter Mode with CBC-MAC (CCM)[3]. This AES mode provides data privacy, data integrity, and source integrity with low additional computational requirements beyond data privacy, alone.

4.3.2 Mutual Authentication of AP and AR

As detailed in [Section 4.3](#), the need to enforce secure communication requires a mutual strong authentication protocol and an associated Key Management protocol that derives from the trust established the authentication phase. The resulting key material is used to derive session keys and subsequent key agreement for setting up secure encapsulation of AP-AR meta-communications.

The Key Management protocol choices are governed by the worst-case specification of Lightweight AP (LAP) capabilities.

4.3.3 Path Security of AP and AR

The secure communications MUST support confidentiality, message authentication and replay protection. The choice of ciphers should consider the required strength and threat model as well as the compute capabilities, real-time nature and relative bandwidth of such traffic.

4.4 AP Provisioning

In order to create a trust model between the AP subsystem and AC subsystem for secure communications enabling automatic discovery, configuration and adaptive resource management the AP's need to be set up securely in the AC(AR)'s domain.

4.4.1 AP Identity

Identity of the AP is established reliably by cryptographically secure binding of an AP's unique identity such one of its wireline MAC addresses to a cryptographic key.

4.4.2 AP Configuration

Configuration of an AP includes providing the parameters necessary for the AP to advertise and provide service for one or more WLANs. These parameters are both physical and logical.

Physical parameters are related to the operation of the AP's radio interface. These include the channel on which the AP is to operate, the maximum power at which the AP is to transmit, antenna selections, the supported data rates, and the timing for the periodic announcements of the WLANs provisioned on the AP.

Logical parameters are related to the individual WLANs that are provisioned on the AP. These include the SSID of the WLANs, the allowed authentication methods, the allowed privacy methods, values for the contention-free period and DTIM, VLAN associations, IP addresses and netmasks, authentication server addresses, any pre-shared keys for WLANs or authentication servers, regulatory (country) information, and other 802.11-specific capabilities to be advertised for the WLANs.

4.4.3 Access Router Availability

Also discussed later in [Section 4.6](#) in discovery context, as part of provisioning an AP one may configure the ability to offer redundancy of ACs or based on negotiated architecture. Constrained architectures with limited AP-AR topologies may be unable to offer flexible

redundancies and may require hardware supported alternatives.

4.5 Access Point Service Management

In a large WLAN system with many APs, continuous management of those APs is necessary to enable quick reaction to changes in service capabilities caused by internal or external factors such as dynamically varying hotspot loads and time-variant fluctuations in RF interferences due to extraneous neighborhood devices. An adaptive RF management based on dynamic systemic monitoring and power and frequency management is needed to be driven from ACs (or at times a hierarchy of ACs).

4.5.1 Monitoring

Each AP in a WLAN must be monitored for a number of variables. This is needed to be able to assess the ability of the individual AP to meet the service demands placed upon it. Among the variables that need to be monitored are:

- a) instantaneous data load
- b) peak and average load over a configurable monitoring period
- c) Measurements of interference from neighboring BSS's
- d) number of mobile devices associated
- e) statistics for each associated mobile device
- f) Signal Strength of Received Frames
- g) RADAR detection

4.5.2 Control

Maintaining the operation of a large WLAN system at or near its peak capability requires that the individual APs that comprise that system must be controlled to adapt to changes in the internal or external factors that affect the performance of the system as a whole. In particular, the aspects of an AP that require control are the following:

- a) Access Controller to which the AP is connected

- b) enabling and disabling the operation of the AP
- c) Enabling and Disabling operation of individual radios at an AP
- d) establishment and update of session keys for protection of AC/AP communication
- e) Radio channel for transmission
- f) Transmit Power

4.6 Access Router Discovery

When a AP comes alive on a network it may authenticate and register with one or more ARs it detects on the network it is connected to. In some architectures today the ARs are the bridges they directly connect to. It performs a AR discovery procedure in its network neighborhood. Based on the Network Topology and Layering it MAY attempt a L2 or IP discovery of ACs. This will also depend partly on the architectural capabilities of the AP and of available ARs. The type of discovery protocol is also dependent on prior one-time Provisioning of AP (configuration). The identification of ARs is only dependent on the L2 or IP protocol used but is expected to be architecture-agnostic. It is the Capability Negotiation Phase ([Section 4.6.2](#)) that follows which resolves the mutual capabilities of AP and AC which lets them decide to AP register with one or more AC.

4.6.1 Access Router Availability

CAPWAP discovery entails the ability of an AP to failover to another AR in the same domain (ESS) in the event of the failure of the current AR.

Failure detection and failover may use existing IP protocols such as VRRP or extensions thereof.

4.6.2 Capabilities Negotiation

An AP performs AR discovery in its network neighborhood. Upon having discovered available ARs the AP enters into a capabilities exchange phase with the candidate ACs. If the architectural types match during the exchange - the AP registers with the AC and configures itself based on the policies it derives from the AC after mutually authenticating with the AC. The capabilities negotiated by architectural type match will decide the applicable API's between AP and AC.

4.7 Summary

The CAPWAP allows for a set of flexible architectures as described in [Section 4.1.4](#). The architecture proposes the following set of CAPWAP services to achieve the Security, Ease of Management, Enhanced QoS and Mobility objectives across the WLAN domain:

- o AC Discovery
- o Capability Negotiation
- o Mutual Authentication of AP and AC
- o Secure Encapsulation Protocol based on Secure Key Management
- o Secure AP Configuration from AC
- o Secure Encapsulation of Control and/or Data between AP & AC

5. Prior Work

Related work on such problems have been dealt with in Academia ([Section 5.1](#)) and standards ([Section 5.2](#)). The former is more directly related to the proposed CAPWAP architecture. The latter is a generic solution attempted to address distributed data-forwarding front-ends with highly available control-plane backends.

5.1 DIRAC

DIRAC is a DIstributed Router ArChitecture for wireless networks, independently developed by a research group in UCLA. DIRAC [\[16\]](#) adopts a very similar distributed architecture to what is proposed here that is composed of a generic Router Core (RC) shared by the wireless subnets and a lightweight and network specific Router Agent (RA) at each access point/base station. The Router Core in DIRAC corresponds to AR in CAPWAP while the Router Agents are the APs.

While the architecture and end goals of DIRAC are very similar to CAPWAP, there are several difference that are worth pointing out:

- o The Router Core at DIRAC is intended to be generic and agnostic to the L2 radio technology being used between the Router Agent and the client terminals. This is achieved by terminating the radio specific L2 connection at the Router Agent while the statistics/actions(i.e.,control)/events messages that are exchanged between the Router Core and the Router Agent are abstracted into a different and generic packet format. CAPWAP, on the other hand, simply encapsulates the 802.11 management frames from APs to ARs so that ARs have to fully understand 802.11 frame format.
- o No security is being considered in the DIRAC work which is probably ok for academic research but not ok for IETF standard.
- o The DIRAC paper focuses less on the protocol between the RC and RA but a lot more on the architecture and implementation issues in this work. The protocol consists of three kinds of messages: statistics, actions (i.e., controls) and (asynchronous) events. DIRAC does not consider the issue of discovery and firmware image downloading etc.
- o The DIRAC paper provides three prototype wireless services that are implemented within the DIRAC framework to demonstrate not only the potential performance gain but also the viability of these new wireless services being enabled by such a framework. These examples provide some nice academic data points for CAPWAP.

5.2 ForCES

The IETF ForCES (Forwarding and Control Element Separation) group was chartered to "define a framework and associated mechanisms for standardizing the exchange of information between the logically separate functionality of the control plane, including entities such as routing protocols, admission control, and signaling, and the forwarding plane, where per-packet activities such as packet forwarding, queuing, and header editing occur. By defining a set of standard mechanisms for control and forwarding separation, ForCES will enable rapid innovation in both the control and forwarding planes. A standard separation mechanism allows the control and forwarding planes to innovate in parallel while maintaining interoperability."

5.2.1 Similarities in Objectives and Architectural Considerations

While ForCES aims to provide interoperability between CEs and FEs from different vendors, CAPWAP has a very similar goal in mind -- to allow APs and ARs from different vendors interoperable when mixed and matched in the wireless access networks. Even though ForCES originally was heavily focused on routers to achieve interoperability between Forwarding Elements (FEs) and Control Elements (CEs) inside a router (i.e., Network Element -- NE), many similarities or analogies can be found between ForCES architecture and CAPWAP architecture:

- o "The APs can be considered as remote RF interfaces, being controlled by the AR" [LWAPP spec] -- it is clear that APs in the CAPWAP architecture can be viewed as FEs in the ForCES architecture, or more precisely, APs can be viewed as a specific wireless port function (Logical Functional Block, LFB, using ForCES's terminology) that is part of the FEs.
- o The LWAPP-related functionality of AR in the wireless access network is mostly control plane related and hence the AR can be considered a CE from the ForCES point of view. It should be noted that the AR also performs forwarding functions, and as such, could also be internally viewed as a CE/FE combination, although usage of ForCES to control APs by the AR would not necessitate usage of ForCES within the AR.
- o "AR + multiple lite-weight APs" as a whole then can be considered as a distributed router with some parts of the FEs (APs) physically separated from the CEs.

5.2.2 Overlap in Topology Considerations

While it is possible to construct a NE out of CEs and FEs which are physically separated by a routed (L3) cloud, ForCES constraint itself to focus on very close localities consisting of CE and FEs that are either components in the same physical box, or are separated at most by one local network (L3) hop. This topology overlaps with the three topologies -- directly connected, switched (L2), or routed (L3) -- considered by CAPWAP as well. But if CAPWAP support arbitrary routed cloud (with multiple L3 hops) between AP and AR, we need to carefully examine ForCES and see if it can accommodate such topology while still satisfying all the requirements including security.

5.2.3 Differences in Design Approach

The general design behind ForCES is to separate the base protocol from the actual information elements that carry the control/configuration/monitoring/events messages between the CE and FE, due to the diversity of FE functions among data plane vendors. The information elements that are specific to any particular FE (e.g., IPv4 forwarding, or DiffServ, or MPLS) are represented in FE model. Such design allows ForCES to be very flexible and extensible to accommodate wide spectrum of data plane functions, possibly including IEEE 802.11 wireless AP functions. The current LWAPP protocol is taking a very different design approach. LWAPP is a very domain specific protocol. While the general domain for LWAPP can potentially include any wireless radio technologies, the current spec of LWAPP is very much IEEE 802.11 specific and many of the 802.11 functions are assumed and built into the protocol directly.

5.2.4 Differences in the Functionality Controlled

The FE functions being controlled by CE via ForCES are mostly L3 and L4, but sometimes L2 (e.g., ARP). On the other hand, the AP functions that are being controlled by ARs are mostly L2 (IEEE 802.11MAC), but sometimes higher layer as well (if those functions reside on APs).

5.2.5 Similarities in Security Requirements

The security requirements in both the CAPWAP and ForCES appear to overlap significantly, in terms of secure association, authentication, confidentiality, integrity, anti-replay, etc. Even though ForCES has not finalized on its protocol selection (among three proposals) yet, ForCES framework document recommends that ForCES adopt one of the standard security mechanisms (IPsec or TLS). More close examination of security requirements and mechanisms employed in ForCES and CAPWAP is needed here.

5.2.6 Difference in Operation Scope

Even though no specific discovery mechanism is specified in the current LWAPP spec, CAPWAP does consider AR discovery in scope; on the other hand, ForCES considers the process of CE and FE discovering each other out of scope. ForCES assumes CEs and FEs enter the post-association phase with knowledge of which corresponding entities they are authorized to communicate with, but ForCES itself does not address how pre-association is done.

5.2.7 Comparision in Protocols

ForCES currently has three protocol proposals and the WG has just started the protocol evaluation and selection process. Therefore, it is difficult to compare LWAPP with ForCES at the moment from the protocol view-point, unless one compares LWAPP with all the three proposals first.

But ForCES requirement document captures all the important requirements that ForCES protocol is supposed to support. Merely comparing LWAPP with this set of requirements can already provide some insight.

The most obvious difference in the two protocols may very well be due to fundamentally different design philosophies behind the two as pointed out in [Section 5.2.4](#). LWAPP is a domain specific protocol with some messages assuming 802.11 semantics, while the base ForCES protocol only supports the general procedures involved for setting up association between the CE and FE, CE querying FE its capability and configuration state (if any), CE provisioning FE according to the basic capability learned in the querying stage, and FE reporting statistics and asynchronous events to CE, etc. In the context of ForCES, the messages with 802.11 specific semantics would not appear in the base ForCES protocol. Instead, an 802.11 FE (or LFB) model would have to be specified to support all the 802.11 specific configuration, statistics, and events.

Another major difference is on reliability requirement. The ForCES protocol is required to support strict reliability for mission critical payloads. On the other hand, LWAPP does not assume any reliability between the AR and AP, because it is built on top of L2 or IP directly.

One thing that LWAPP supports but none of the ForCES protocol proposals directly address is firmware image downloading.

6. Security Considerations

One of the major goals of the CAPWAP architecture is to ensure strong authentication of AP to the registered AR and secure communications between them as described in the preceding sections.

AR-AP traffic can be classified into: data traffic (e.g. from or to an end user), and control traffic which is strictly between the AR and AP. Since data traffic may be secured end-to-end security mechanisms outside the scope of this work, we confine our solution to control traffic. The resulting security consists of two components: an authenticated key exchange, and control traffic security encapsulation. The security encapsulation may be accomplished using relatively lightweight mechanisms such as CCM, described in [2]. This encapsulation provides for strong AES-based message authentication and encryption. Detailed discussions of such possible security protocol alternatives is out of scope in this document.

7. Acknowledgements

The authors wish to thank the timely inputs and discussions provided by Pat Calhoun towards completion of this document in a very short time. In no less measure our thanks go to Scott Kelly et al in kindly consenting to let us adapt from their topological and architectural analysis in [5] that helped us shorten the time to draft. The authors also wish to thank IESG & IAB for their feedback, particularly Randy Bush, James Kempf and Bernard Aboba for the discussions that has helped focus this draft objective. Thanks are also due to Dorothy Stanley in this regard for the IEEE perspective.

References

- [1] "IEEE WLAN MAC and PHY Layer Specifications", August 1999, <IEEE 802.11-99>.
- [2] "Advanced Encryption Standard (AES)", November 2001, <FIPS PUB 197>.
- [3] "Counter with CBC-MAC (CCM)", September 2003, <[RFC 3610](#)>.
- [4] "Light Weight Access Point Protocol (LWAPP)", June 2003, <<http://www.ietf.org/internet-drafts/draft-calhoun-seamoby-lwapp-03.txt>>.
- [5] "Security Requirements for a Light Weight Access Point Protocol", August 2003, <<http://www.ietf.org/internet-drafts/draft-kelly-ietf-lwapp-sec-00.txt>>.
- [6] "The Internet Standards Process Revision 3", October 1996, <<ftp://ftp.isi.edu/in-notes/rfc2026>>.
- [7] "Key words for use in RFCs to Indicate Requirement Levels", March 1997, <<ftp://ftp.isi.edu/in-notes/rfc2119>>.
- [8] "Mobility Related Terminology", April 2003, <<ftp://ftp.isi.edu/internet-drafts/draft-ietf-seamoby-terminology-04.txt>>.
- [9] "Extensible Authentication Protocol (EAP)", September 2003, <<http://www.ietf.org/internet-drafts/draft-ietf-eap-rfc2284bis-06.txt>>.
- [10] "WiFi Protected Access (WPA) ver 2.0", April 2003.
- [11] "IEEE Std 802.11i/6.0: Specification for Enhanced Security", September 2003.
- [12] "IEEE Std 802.11F: Recommended Practice for Multi-Vendor Access Point Interoperability via an Inter-Access Point Protocol across Distribution Systems Support 802.11 Operation", July 2003.
- [13] "IEEE Std 802.11h: Spectrum and Transmit Power Management Extensions in the 5 GHz Band in Europe", October 2003.
- [14] "IEEE Std 802.1X: Port-based Network Access Control", June 2001.
- [15] "IEEE Std 802.1Q: Virtual Bridged Local Area Networks", May

2003.

[16] "DIRAC: A Software-based Wireless Router System", 2003.

Authors' Addresses

Mahalingam Mani
Avaya Inc.
1001 Murphy Ranch Rd
Milpitas, CA 95035

Phone: +1 408-321-4840
EMail: mmani@avaya.com

Bob O'Hara
Airespace
110 Nortech Parkway
San Jose, CA 95134

Phone: +1 408-635-2025
EMail: bob@airespace.com

L. Lily Yang
Intel Corp.
MS JF3 206, 2111 NE 25th Avenue
Hillsboro, OR 97124

Phone: +1 503-264-8813
EMail: lily.l.yang@intel.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION

HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF
MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgment

Funding for the RFC Editor function is currently provided by the
Internet Society.