                      Architecture Taxonomy for
       Control and Provisioning of Wireless Access Points (CAPWAP)

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

Abstract

   This document provides a taxonomy of the architectures employed in
   the existing IEEE 802.11 products in the market, by analyzing
   Wireless LAN (WLAN) functions and services and describing the
   different variants in distributing these functions and services among
   the architectural entities.

Table of Contents

1. Introduction

   As IEEE 802.11 Wireless LAN (WLAN) technology matures, large scale
   deployment of WLAN networks is highlighting certain technical
   challenges. As outlined in [2], management, monitoring, and control
   of large number of Access Points (APs) in the network may prove to be
   a significant burden for network administration. Distributing and
   maintaining a consistent configuration throughout the entire set of
   APs in the WLAN is a difficult task. The shared and dynamic nature
   of the wireless medium also demands effective coordination among the
   APs to minimize radio interference and maximize network performance.
   Network security issues, which have always been a concern in WLANs,
   present even more challenges in large deployments and new
   architectures.

   Recently many vendors have begun offering partially proprietary
   solutions to address some or all of the above mentioned problems.
   Since interoperable systems allow for a broader choice of solutions,
   a standardized interoperable solution addressing the aforementioned
   problems is desirable. As the first step toward establishing
   interoperability in the market place, this document provides a
   taxonomy of the architectures employed in existing WLAN products. We
   hope to provide a cohesive understanding of the market practices for
   the standard bodies involved (including the IETF and IEEE 802.11).
   This document may be reviewed and utilized by the IEEE 802.11 Working
   Group as input in defining the functional architecture of an AP.

1.1. IEEE 802.11 WLAN Functions

   The IEEE 802.11 specifications are wireless standards that specify an
   "over-the-air" interface between a wireless client Station (STA) and
   an Access Point (AP), and also among wireless clients. 802.11 also
   describes how mobile devices can associate into a basic service set
   (BSS). A BSS is identified by a basic service set identifier (BSSID)
   or name. The WLAN architecture can be considered as a type of 'cell'
   architecture, in which each cell is the Basic Service Set (BSS), and
   each BSS is controlled by the AP. When two or more APs are connected
   via a broadcast layer 2 network and all are using the same SSID, an
   extended service set (ESS) is created.

   The architectural component used to interconnect BSSs is the
   distribution system (DS). An AP is an STA that provides access to
   the DS by providing DS services, as well as acting as an STA.
   Another logical architectural component, portal, is introduced to
   integrate the IEEE 802.11 architecture with a traditional wired LAN.
   It is possible for one device to offer both the functions of an AP
   and a portal.

   IEEE 802.11 does not specify the details of DS implementations
   explicitly. Instead, the 802.11 standard defines services that
   provide functions that the LLC layer requires for sending MAC Service
   Data Units (MSDUs) between two entities on the network. These
   services can be classified into two categories: the station service
   (SS) and the distribution system service (DSS). Both categories of
   service are used by the IEEE 802.11 MAC sublayer. Station services
   consist of the following four services:

   o Authentication: Establishes the identity of one station as a
      member of the set of stations that are authorized to associate
      with one another.

   o De-authentication: Voids an existing authentication relationship.

   o Confidentiality: Prevents the content of messages from being read
      by others than the intended recipients.

   o MSDU Delivery: Delivers the MAC service data unit (MSDU) for the
      stations.

      Distribution system services consist of the following five
      services:

   o Association: Establishes Access Point/Station (AP/STA) mapping and
      enables STA invocation of the distribution system services.

o Disassociation: Removes an existing association.

o Reassociation: Enables an established association (between AP and
   STA) to be transferred from one AP to another or the same AP.

o Distribution: Provides MSDU forwarding by APs for the STAs
   associated with them. MSDUs can be either forwarded to the
   wireless destination or to the wired (Ethernet) destination (or
   both) using the "Distribution System" concept of 802.11.

o Integration: Translates the MSDU received from the Distribution
   System to a non-802.11 format and vice versa. Any MSDU that is
   received from the DS invokes the 'Integration' services of the DSS
   before the 'Distribution' services are invoked. The point of
   connection of the DS to the wired LAN is termed as 'portal'.

Apart from these services, the IEEE 802.11 also defines additional
MAC services that must be implemented by the APs in the WLAN. For
example:

o Beacon Generation

o Probe Response/Transmission

o Processing of Control Frames: RTS/CTS/ACK/PS-Poll/CF-End/CF-ACK

o Synchronization

o Retransmissions

o Transmission Rate Adaptation

o Privacy: 802.11 Encryption/Decryption

In addition to the services offered by the 802.11, the IEEE 802.11 WG
is also developing technologies to support Quality of Service
(802.11e), Security Algorithms (802.11i), Inter-AP Protocol (IAPP, or
802.11F -- recommended practice) to update APs when a STA roams from
one BSS to another, Radio Resource Measurement Enhancements
(802.11k), etc.

IEEE 802.11 does not specify exactly how these functions are
implemented, nor does it specify that they be implemented in one
physical device. It only requires that the APs and the rest of the
DS together implement all these services. Typically, vendors
implement not only the services defined in the IEEE 802.11 standard,
but also a variety of value-added services or functions, such as load
balancing support, QoS, station mobility support, and rogue AP

detection. What becomes clear from this document is that vendors
take advantage of the flexibility in the 802.11 architecture, and
have come up with many different flavors of architectures and
implementations of the WLAN services.

Because many vendors choose to implement these WLAN services across
multiple network elements, we want to make a clear distinction
between the logical WLAN access network functions and the individual
physical devices by adopting different terminology. We use "AP" to
refer to the logical entity that provides access to the distribution
services, and "WTP" (Wireless Termination Point) to the physical
device that allows the RF antenna and 802.11 PHY to transmit and
receive station traffic in the BSS network. In the Centralized
Architecture (see section 5), the combination of WTPs with Access
Controller (AC) implements all the logical functions. Each of these
physical devices (WTP or AC) may implement only part of the logical
functions. But the DS, including all the physical devices as a
whole, implements all or most of the functions.

1.2. CAPWAP Functions

To address the four problems identified in [2] (management,
consistent configuration, RF control, security) additional functions,
especially in the control and management plane, are typically offered
by vendors to assist in better coordination and control across the
entire ESS network. Such functions are especially important when the
IEEE 802.11 WLAN functions are implemented over multiple entities in
a large scale network, instead of within a single entity. Such
functions include:

o RF monitoring, such as Radar detection, noise and interference
   detection, and measurement.

o RF configuration, e.g., for retransmission, channel selection,
   transmission power adjustment.

o WTP configuration, e.g., for SSID.

o WTP firmware loading, e.g., automatic loading and upgrading of WTP
   firmware for network wide consistency.

o Network-wide STA state information database, including the
   information needed to support value-added services, such as
   mobility and load balancing.

o Mutual authentication between network entities, e.g., for AC and
   WTP authentication in a Centralized WLAN Architecture.

The services listed are concerned with the configuration and control
of the radio resource ('RF Monitoring' and 'RF Configuration'),
management and configuration of the WTP device ('WTP Configuration',
'WTP Firmware upgrade'), and also security regarding the registration
of the WTP to an AC ('AC/WTP mutual authentication'). Moreover, the
device from which other services, such as mobility management across
subnets and load balancing, can obtain state information regarding
the STA(s) associated with the wireless network, is also reported as
a service ('STA state info database').

The above list of CAPWAP functions is not an exhaustive enumeration
of all additional services offered by vendors. We included only
those functions that are commonly represented in the survey data, and
are pertinent to understanding the central problem of
interoperability.

Most of these functions are not explicitly specified by IEEE 802.11,
but some of the functions are. For example, the control and
management of the radio-related functions of an AP are described
implicitly in the MIB, such as:

o Channel Assignment

o Transmit Power Control

o Radio Resource Measurement (work is currently under way in IEEE
   802.11k)

The 802.11h [5] 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 the following areas:

o RADAR detection

o Transmit Power Control

o Dynamic Channel Selection

1.3. WLAN Architecture Proliferation

   This document provides a taxonomy of the WLAN network architectures
   developed by the vendor community in an attempt to address some or
   all of the problems outlined in [2]. As the IEEE 802.11 standard
   purposely avoids specifying the details of DS implementations,
   different architectures have proliferated in the market. While all
   these different architectures conform to the IEEE 802.11 standard as
   a whole, their individual functional components are not standardized.

Interfaces between the network architecture components are mostly
proprietary, and there is no guarantee of cross-vendor
interoperability of products, even within the same family of
architectures.

To achieve interoperability in the market place, the IETF CAPWAP
working group is first documenting both the functions and the network
architectures currently offered by the existing WLAN vendors. The
end result is this taxonomy document.

After analyzing more than a dozen different vendors' architectures,
we believe that the existing 802.11 WLAN access network architectures
can be broadly categorized into three distinct families, based on the
characteristics of the Distribution Systems that are employed to
provide the 802.11 functions.

o Autonomous WLAN Architecture: The first architecture family is the
   traditional autonomous WLAN architecture, in which each WTP is a
   single physical device that implements all the 802.11 services,
   including both the distribution and integration services, and the
   portal function. Such an AP architecture is called Autonomous
   WLAN Architecture because each WTP is autonomous in its
   functionality, and no explicit 802.11 support is needed from
   devices other than the WTP. In such architecture, the WTP is
   typically configured and controlled individually, and can be
   monitored and managed via typical network management protocols
   like SNMP. The WTPs are the traditional APs with which most
   people are familiar. Such WTPs are sometimes referred to as "Fat
   APs" or "Standalone APs".

o Centralized WLAN Architecture: The second WLAN architecture family
   is an emerging hierarchical architecture utilizing one or more
   centralized controllers for managing a large number of WTP
   devices. The centralized controller is commonly referred to as an
   Access Controller (AC), whose main function is to manage, control,
   and configure the WTP devices that are present in the network. In
   addition to being a centralized entity for the control and
   management plane, it may also become a natural aggregation point
   for the data plane since it is typically situated in a centralized
   location in the wireless access network. The AC is often co-
   located with an L2 bridge, a switch, or an L3 router, and may be
   referred to as Access Bridge or Access Router in those particular
   cases. Therefore, an Access Controller could be either an L3 or
   L2 device, and is the generic term we use throughout this
   document. It is also possible that multiple ACs are present in a
   network for purposes of redundancy, load balancing, etc. This
   architecture family has several distinct characteristics that are
   worth noting. First, the hierarchical architecture and the

   centralized AC affords much better manageability for large scale
   networks. Second, since the IEEE 802.11 functions and the CAPWAP
   control functions are provided by the WTP devices and the AC
   together, the WTP devices themselves may no longer fully implement
   the 802.11 functions as defined in the standards. Therefore, it
   can be said that the full 802.11 functions are implemented across
   multiple physical network devices, namely, the WTPs and ACs.
   Since the WTP devices only implement a portion of the functions
   that standalone APs implement, WTP devices in this architecture
   are sometimes referred to as light weight or thin APs.

   o Distributed WLAN Architecture: The third emerging WLAN
      architecture family is the distributed architecture in which the
      participating wireless nodes are capable of forming a distributed
      network among themselves, via wired or wireless media. A wireless
      mesh network is one example within the distributed architecture
      family, where the nodes themselves form a mesh network and connect
      with neighboring mesh nodes via 802.11 wireless links. Some of
      these nodes also have wired Ethernet connections acting as
      gateways to the external network.

1.4. Taxonomy Methodology and Document Organization

   Before the IETF CAPWAP working group started documenting the various
   WLAN architectures, we conducted an open survey soliciting WLAN
   architecture descriptions via the IETF CAPWAP mailing list. We
   provided the interested parties with a common template that included
   a number of questions about their WLAN architectures. We received 16
   contributions in the form of short text descriptions answering those
   questions. 15 of them are from WLAN vendors (AireSpace, Aruba,
   Avaya, Chantry Networks, Cisco, Cranite Systems, Extreme Networks,
   Intoto, Janusys Networks, Nortel, Panasonic, Trapeze, Instant802,
   Strix Systems, Symbol) and one from the academic research community
   (UCLA). Out of the 16 contributions, one describes an Autonomous
   WLAN Architecture, three are Distributed Mesh Architectures, and the
   remaining twelve entries represent architectures in the family of the
   Centralized WLAN Architecture.

   The main objective of this survey was to identify the general
   categories and trends in WLAN architecture evolution, discover their
   common characteristics, and determine what is performed differently
   among them and why. In order to represent the survey data in a
   compact format, a "Functional Distribution Matrix" is used in this
   document, (mostly in the Centralized WLAN architecture section), to
   tabulate the various services and functions in the vendors'
   offerings. These services and functions are classified into three
   main categories:

   o Architecture Considerations: The choice of the connectivity
      between the AC and the WTP. The design choices regarding the
      physical device on which processing of management, control, and
      data frames of the 802.11 takes place.

   o 802.11 Functions: As described in Section 1.1.

   o CAPWAP Functions: As described in Section 1.2.

   For each one of these categories, the mapping of each individual
   function to network entities implemented by each vendor is shown in
   tabular form. The rows in the Functional Distribution Matrix
   represent individual functions that are organized into the above
   mentioned three categories. Each column of the Matrix represents one
   vendor's architecture offering in the survey data. See Figure 7 as
   an example of the Matrix.

   This Functional Distribution Matrix is intended for the sole purpose
   of organizing the architecture taxonomy data, and represents the
   contributors' views of their architectures from an engineering
   perspective. It does not necessarily imply that a product exists or
   will be shipped, nor an intent by the vendor to build such a product.

   The next section provides a list of definitions used in this
   document. The rest of this document is organized around the three
   broad WLAN architecture families that were introduced in Section 1.3.
   Each architecture family is discussed in a separate section. The
   section on Centralized Architecture contains more in-depth details
   than the other two families, largely due to the large number of the
   survey data (twelve out of sixteen) collected that fall into the
   Centralized Architecture category. Summary and conclusions are
   provided at the end to highlight the basic findings from this
   taxonomy exercise.

2. Conventions

   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 [3].

3. Definitions

3.1. IEEE 802.11 Definitions

   Station (STA): A device that contains an IEEE 802.11 conformant
   medium access control (MAC) and physical layer (PHY) interface to the
   wireless medium (WM).

   Access Point (AP): An entity that has station functionality and
   provides access to distribution services via the wireless medium (WM)
   for associated stations.

   Basic Service Set (BSS): A set of stations controlled by a single
   coordination function.

   Station Service (SS): The set of services that support transport of
   medium access control (MAC) service data units (MSDUs) between
   stations within a basic service set (BSS).

   Distribution System (DS): A system used to interconnect a set of
   basic service sets (BSSs) and integrated local area networks (LANs)
   to create an extended service set (ESS).

   Extended Service Set (ESS): A set of one or more interconnected basic
   service sets (BSSs) with the same SSID and integrated local area
   networks (LANs), which appears as a single BSS to the logical link
   control layer at any station associated with one of those BSSs.

   Portal: The logical point at which medium access control (MAC)
   service data units (MSDUs) from a non-IEEE 802.11 local area network
   (LAN) enter the distribution system (DS) of an extended service set
   (ESS).

   Distribution System Service (DSS): The set of services provided by
   the distribution system (DS) that enable the medium access control
   (MAC) layer to transport MAC service data units (MSDUs) between
   stations that are not in direct communication with each other over a
   single instance of the wireless medium (WM). These services include
   the transport of MSDUs between the access points (APs) of basic
   service sets (BSSs) within an extended service set (ESS), transport
   of MSDUs between portals and BSSs within an ESS, and transport of
   MSDUs between stations in the same BSS in cases where the MSDU has a
   multicast or broadcast destination address, or where the destination
   is an individual address, but the station sending the MSDU chooses to
   involve DSS. DSSs are provided between pairs of IEEE 802.11 MACs.

   Integration: The service that enables delivery of medium access
   control (MAC) service data units (MSDUs) between the distribution
   system (DS) and an existing, non-IEEE 802.11 local area network (via
   a portal).

   Distribution: The service that, by using association information,
   delivers medium access control (MAC) service data units (MSDUs)
   within the distribution system (DS).

3.2. Terminology Used in This Document

   One of the motivations in defining new terminology is to clarify
   ambiguity and confusion surrounding some conventional terms. One
   such term is "Access Point (AP)". Typically, when people talk about
   "AP", they refer to the physical entity (box) that has an antenna,
   implements 802.11 PHY, and receives/transmits the station (STA)
   traffic over the air. However, the 802.11 Standard [1] describes the
   AP mostly as a logical entity that implements a set of logical
   services so that station traffic can be received and transmitted
   effectively over the air. When people refer to "AP functions", they
   usually mean the logical functions the whole WLAN access network
   supports, and not just the subset of functions supported by the
   physical entity (box) that the STAs communicate with directly. Such
   confusion can be especially acute when logical functions are
   implemented across a network instead of within a single physical
   entity. To avoid further confusion, we define the following
   terminology:

   CAPWAP: Control and Provisioning of Wireless Access Points

   IEEE 802.11 WLAN Functions: A set of logical functions defined by the
   IEEE 802.11 Working Group, including all the MAC services, Station
   Services, and Distribution Services. These logical functions are
   required to be implemented in the IEEE 802.11 Wireless LAN (WLAN)
   access networks by the IEEE 802.11 Standard [1].

   CAPWAP Functions: A set of WLAN control functions that are not
   directly defined by IEEE 802.11 Standards, but deemed essential for
   effective control, configuration, and management of 802.11 WLAN
   access networks.

   Wireless Termination Point (WTP): The physical or network entity that
   contains an RF antenna and 802.11 PHY to transmit and receive station
   traffic for the IEEE 802.11 WLAN access networks. Such physical
   entities were often called "Access Points" (AP), but "AP" can also
   refer to the logical entity that implements 802.11 services. We
   recommend "WTP" as the generic term that explicitly refers to the
   physical entity with the above property (e.g., featuring an RF
   antenna and 802.11 PHY), applicable to network entities of both
   Autonomous and Centralized WLAN Architecture (see below).

   Autonomous WLAN Architecture: The WLAN access network architecture
   family in which all the logical functions, including both IEEE 802.11
   and CAPWAP functions (wherever applicable), are implemented within
   each Wireless Termination Point (WTP) in the network. The WTPs in

such networks are also called standalone APs, or fat APs, because
these devices implement the full set of functions that enable the
devices to operate without any other support from the network.

Centralized WLAN Architecture: The WLAN access network architecture
family in which the logical functions, including both IEEE 802.11 and
CAPWAP functions (wherever applicable), are implemented across a
hierarchy of network entities. At the lower level are the WTPs,
while at the higher level are the Access Controllers (ACs), which are
responsible for controlling, configuring, and managing the entire
WLAN access network.

Distributed WLAN Architecture: The WLAN access network architecture
family in which some of the control functions (e.g., CAPWAP
functions) are implemented across a distributed network consisting of
peer entities. A wireless mesh network can be considered an example
of such an architecture.

Access Controller (AC): The network entity in the Centralized WLAN
Architecture that provides WTPs access to the centralized
hierarchical network infrastructure in the data plane, control plane,
management plane, or a combination therein.

Standalone WTP: Refers to the WTP in Autonomous WLAN Architecture.

Controlled WTP: Refers to the WTP in Centralized WLAN Architecture.

Split MAC Architecture: A subgroup of the Centralized WLAN
Architecture whereby WTPs in such WLAN access networks only implement
the delay sensitive MAC services (including all control frames and
some management frames) for IEEE 802.11, while all the remaining
management and data frames are tunnelled to the AC for centralized
processing. The IEEE 802.11 MAC, as defined by IEEE 802.11 Standards
in [1], is effectively split between the WTP and AC.

Remote MAC Architecture: A subgroup of the Centralized WLAN
Architecture, where the entire set of 802.11 MAC functions (including
delay-sensitive functions) is implemented at the AC. The WTP
terminates the 802.11 PHY functions.

Local MAC Architecture: A subgroup of the Centralized WLAN
Architecture, where the majority or entire set of 802.11 MAC
functions (including most of the 802.11 management frame processing)
are implemented at the WTP. Therefore, the 802.11 MAC stays intact
and local in the WTP, along with PHY.

3.3. Terminology Used Historically but Not Recommended

   While some terminology has been used by vendors historically to
   describe "Access Points", we recommend deferring its use, in order to
   avoid further confusion. A list of such terms and the recommended
   new terminology is provided below:

   Split WLAN Architecture: Use Centralized WLAN Architecture.

   Hierarchical WLAN Architecture: Use Centralized WLAN Architecture.

   Standalone Access Point: Use Standalone WTP.

   Fat Access Point: Use Standalone WTP.

   Thin Access Point: Use Controlled WTP.

   Light weight Access Point: Use Controlled WTP.

   Split AP Architecture: Use Local MAC Architecture.

   Antenna AP Architecture: Use Remote MAC Architecture.

4. Autonomous Architecture

4.1. Overview

   Figure 1 shows an example network of the Autonomous WLAN
   Architecture. This architecture implements all the 802.11
   functionality in a single physical device, the Wireless Termination
   Point (WTP). An embodiment of this architecture is a WTP that
   translates between 802.11 frames to/from its radio interface and
   802.3 frames to/from an Ethernet interface. An 802.3 infrastructure
   that interconnects the Ethernet interfaces of different WTPs provides
   the distribution system. It can also provide portals for integrated
   802.3 LAN segments.

```
      +---------------+       +---------------+       +---------------+
      |  802.11 BSS 1 |       |  802.11 BSS 2 |       |  802.11 BSS 3 |
      |  ...          |       |  ...          |       |  ...          |
      |     +-----+   |       |     +-----+   |       |     +-----+   |
      +----|  WTP |----+       +----|  WTP |----+       +----|  WTP |----+
           +--+--+               +--+--+               +--+--+
              |Ethernet              |                     |
           +----------------+       |   +----------------+
                            |       |   |   |
                        +---+--+--+---+
                        | Ethernet    |
 802.3 LAN  -------------+ Switch     +------------- 802.3 LAN
 segment 1              |             |             segment 2
                        +------+------+
```
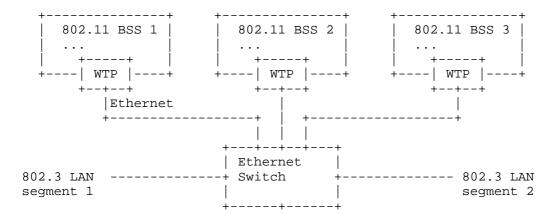
                 Figure 1: Example of Autonomous WLAN Architecture

   A single physical WTP can optionally be provisioned as multiple
   virtual WTPs by supporting multiple SSIDs to which 802.11 clients may
   associate. In some cases, this will involve putting a corresponding
   802.1Q VLAN tag on each packet forwarded to the Ethernet
   infrastructure and removing 802.1Q tags prior to forwarding the
   packets to the wireless medium.

   The scope of the ESS(s) created by interconnecting the WTPs will be
   confined by the constraints imposed by the Ethernet infrastructure.

   Authentication of 802.11 clients may be performed locally by the WTP
   or by using a centralized authentication server.

4.2. Security

   Since both the 802.11 and CAPWAP functions are tightly integrated
   into a single physical device, security issues with this architecture
   are confined to the WTP. There are no extra implications from the
   client authentication and encryption/decryption perspective, as the
   AAA interface and the key generation mechanisms required for 802.11i
   encryption/decryption are integrated into the WTP.

   One of the security needs in this architecture is for mutual
   authentication between the WTP and the Ethernet infrastructure. This
   can be ensured by existing mechanisms such as 802.1X between the WTP
   and the Ethernet switch to which it connects. Another critical
   security issue is the fact that the WTP is most likely not under lock
   and key, but contains secret information to communicate with back-end
   systems, such as AAA and SNMP. Because IT personnel uses the common
   management method of pushing a "template" to all devices, theft of
   such a device would potentially compromise the wired network.

5. Centralized WLAN Architecture

   Centralized WLAN Architecture is an emerging architecture family in
   the WLAN market. Contrary to the Autonomous WLAN Architecture, where
   the 802.11 functions and network control functions are all
   implemented within each Wireless Termination Point (WTP), the
   Centralized WLAN Architecture employs one or more centralized
   controllers, called Access Controller(s), to enable network-wide
   monitoring, improve management scalability, and facilitate dynamic
   configurability.

   The following figure schematically shows the Centralized WLAN
   Architecture network diagram, where the Access Controller (AC)
   connects to multiple Wireless Termination Points (WTPs) via an
   interconnection medium. This can be a direct connection, an L2-
   switched, or an L3-routed network as described in Section 5.1. The
   AC exchanges configuration and control information with the WTP
   devices, allowing the management of the network from a centralized
   point. Designs of the Centralized WLAN Architecture family do not
   presume (as the diagram might suggest) that the AC necessarily
   intercedes in the data plane to/from the WTP(s). More details are
   provided later in this section.

```
    +---------------+     +---------------+     +---------------+
    | 802.11 BSS 1 |     | 802.11 BSS 2 |     | 802.11 BSS 3 |
    | ...           |     | ...           |     | ...           |
    |    +-------+  |     |    +-------+  |     |    +-------+  |
    +----|  WTP  |--+     +----|  WTP  |--+     +----|  WTP  |--+
         +---+---+             +---+---+             +---+---+
             |                     |                     |
         +-----------------+   |   +----------------+
               | |...|
         +----+--+---+--------+
         | Interconnection    |
         +-------+------------+
                 |
                 |
            +-----+----+
            |    AC    |
            +----------+
```
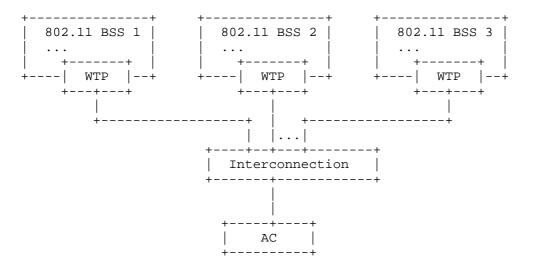
              Figure 2: Centralized WLAN Architecture Diagram

In the diagram above, the AC is shown as a single physical entity
that provides all of the CAPWAP functions listed in Section 1.2.
However, this may not always be the case. Closer examination of the
functions reveals that their different resource requirements (e.g.,
CPU, memory, storage) may be distributed across different devices.
For instance, complex radio control algorithms can be CPU intensive.
Storing and downloading images and configurations can be storage
intensive. Therefore, different CAPWAP functions might be
implemented on different physical devices due to the different nature
of their resource requirements. The network entity marked 'AC' in
the diagram above should be thought of as a multiplicity of logical
functions, and not necessarily as a single physical device. The ACs
may also choose to implement some control functions locally, and
provide interfaces to access other global network management
functions, which are typically implemented on separate boxes, such as
a SNMP Network Management Station and an AAA back-end server (e.g.,
Radius Authentication Server).

5.1. Interconnection between WTPs and ACs

   There are several connectivity options to consider between the AC(s)
   and the WTPs, including direct connection, L2 switched connection,
   and L3 routed connection, as shown in Figures 3, 4, and 5.

```
                     -------+------ LAN
                           |
                   +-------+-------+
                   |      AC       |
                   +----+-----+----+
                        |     |
                    +---+     +---+
                    |             |
                 +--+--+       +--+--+
                 | WTP |       | WTP |
                 +--+--+       +--+--+
```

                   Figure 3: Directly Connected

```
                    -------+------ LAN
                          |
                +-------+-------+
                |      AC       |
                +----+-----+----+
                     |     |
                  +--+     +---+
                  |            |
             +--+--+     +-----+-----+
             | WTP |     |   Switch  |
             +--+--+     +---+-----+-+
                             |     |
                         +-----+ +-----+
                         | WTP | | WTP |
                         +-----+ +-----+
```

Figure 4: Switched Connections

```
            +-------+-------+
            |      AC       |
            +-------+-------+
                    |
            -------+------ LAN
                    |
            +-------+-------+
            |     Router    |
            +-------+-------+
                    |
            -----+--+--+--- LAN
                 |     |
              +--+     +---+
              |            |
           +--+--+      +--+--+
           | WTP |      | WTP|
           +--+--+      +--+--+
```
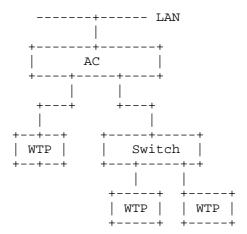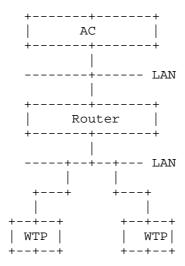
Figure 5: Routed Connections

5.2. Overview of Three Centralized WLAN Architecture Variants

   Dynamic and consistent network management is one of the primary
   motivations for the Centralized Architecture. The survey data from
   vendors also shows that different varieties of this architecture
   family have emerged to meet a complex set of different requirements
   for various possible deployment scenarios. This is also a direct
   result of the inherent flexibility in the 802.11 standard [1]
   regarding the implementation of the logical functions that are

broadly described under the term "Access Point (AP)". Because there
is no standard mapping of these AP functions to physical network
entities, several design choices have been made by vendors that offer
related products. Moreover, the increased demand for monitoring and
consistent configuration of large wireless networks has resulted in a
set of 'value-added' services provided by the various vendors, most
of which share common design properties and service goals.

In the following, we describe the three main variants observed from
the survey data within the family of Centralized WLAN Architecture,
namely the Local MAC, Split MAC, and Remote MAC approaches. For each
approach, we provide the mapping characteristics of the various
functions into the network entities from each vendor. The naming of
Local MAC, Split MAC, and Remote MAC reflects how the functions, and
especially the 802.11 MAC functions, are mapped onto the network
entities. Local MAC indicates that the MAC functions stay intact and
local to WTPs, while Remote MAC denotes that the MAC has moved away
from the WTP to a remote AC in the network. Split MAC shows the MAC
being split between the WTPs and ACs, largely along the line of
realtime sensitivity. Typically, Split MAC vendors choose to put
realtime functions on the WTPs while leaving non-realtime functions
to the ACs. 802.11 does not clearly specify what constitutes
realtime functions versus non-realtime functions, and so a clear and
definitive line does not exist. As shown in Section 5.4, each vendor
has its own interpretation on this, and there are some discrepancies
about where to draw the line between realtime and non-realtime
functions. However, vendors agree on the characterization of the
majority of MAC functions. For example, every vendor classifies the
DCF as a realtime function.

The differences among Local MAC, Split MAC and Remote MAC
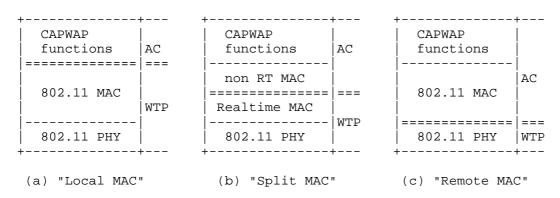architectures are shown graphically in the following figure:

```
   +--------------+---    +--------------+---    +--------------+---
   | CAPWAP       |       | CAPWAP       |       | CAPWAP       |
   | functions    |AC     | functions    |AC     | functions    |
   |==============|===    |--------------|       |--------------|
   |              |       | non RT MAC   |       |              |AC
   | 802.11 MAC   |       |==============|===    | 802.11 MAC   |
   |              |       | Realtime MAC |       |              |
   |              |WTP    |              |       |              |
   |--------------|       |--------------|WTP    |==============|===
   | 802.11 PHY   |       | 802.11 PHY   |       | 802.11 PHY   |WTP
   +--------------+---    +--------------+---    +--------------+---

   (a) "Local MAC"        (b) "Split MAC"        (c) "Remote MAC"
```

Figure 6: Three Architectural Variants within the Centralized
                 WLAN Architecture Family

5.3. Local MAC

   The main motivation of the Local MAC architecture model, as shown in
   Figure 6 (a), is to offload network access policies and management
   functions (CAPWAP functions described in Section 1.2) to the AC
   without splitting the 802.11 MAC functionality between WTPs and AC.
   The whole 802.11 MAC resides on the WTPs locally, including all the
   802.11 management and control frame processing for the STAs. On the
   other hand, information related to management and configuration of
   the WTP devices is communicated with a centralized AC to facilitate
   management of the network and maintain a consistent network-wide
   configuration for the WTP devices.

   Figure 7 shows a tabular representation of the design choices made by
   the six vendors in the survey that follow the Local MAC approach,
   with respect to the above mentioned architecture considerations.
   "WTP-AC connectivity" shows the type connectivity between the WTPs
   and AC that every vendor's architecture can support. Clearly, all
   the vendors can support L3 routed network connectivity between WTPs
   and the AC, which implies that direct connections and L2 switched
   networks are also supported by all vendors. By '802.11 mgmt
   termination', and '802.11 control termination', we denote the
   physical network device on which processing of the 802.11 management
   and control frames is done respectively. All the vendors here choose
   to terminate 802.11 management and control frames at the WTPs. The
   last row of the table, '802.11 data aggregation', refers to the
   device on which aggregation and delivery of 802.11 data frames from
   one STA to another (possibly through a DS) is performed. As shown by
   the table, vendors make different choices as to whether all the
   802.11 data traffic is aggregated and routed through the AC. The
   survey data shows that some vendors choose to tunnel or encapsulate
   all the station traffic to or from the ACs, implying that the AC also
   acts as the access router for this WLAN access network. Other
   vendors choose to separate the control and data plane by letting the
   station traffic be bridged or routed locally, while keeping the
   centralized control at the AC.

	Arch7	Arch8	Arch9	Arch10	Arch11
WTP-AC connectivity	L3	L3	L3	L3	L3
802.11 mgmt termination	WTP	WTP	WTP	WTP	WTP
802.11 control termination	WTP	WTP	WTP	WTP	WTP
802.11 data aggregation	AC	AC	WTP	AC	WTP

        Figure 7: Architecture Considerations for Local MAC Architecture

Figure 8 reveals that most of the CAPWAP functions, as described in
Section 1.2, are implemented at the AC with help from WTPs to monitor
RF channels, and collect statistics and state information from the
STAs, as the AC offers the advantages of network-wide visibility,
which is essential for many of the control, configuration, and
value-added services.

	Arch7	Arch8	Arch9	Arch10	Arch11
RF Monitoring	WTP	WTP	AC/WTP	WTP	WTP
RF Config.	AC	AC	AC	AC	AC
WTP config.	AC	AC	AC	AC	AC
WTP Firmware	AC	AC	AC	AC	AC
STA state info database	AC	AC/WTP	AC/WTP	AC/WTP	AC
AC/WTP mutual authent.	AC/WTP	AC/WTP	AC/WTP	AC/WTP	AC/WTP

     Figure 8: Mapping of CAPWAP Functions for Local MAC Architecture

The matrix in Figure 9 shows that most of the 802.11 functions are
implemented at the WTPs for Local MAC Architecture, with some minor
differences among the vendors regarding distribution service, 802.11e
scheduling, and 802.1X/EAP authentication. The difference in
distribution service is consistent with that described earlier
regarding "802.11 data aggregation" in Figure 7.

	Arch7	Arch8	Arch9	Arch10	Arch11
Distribution Service	AC	AC	WTP	AC	WTP
Integration Service	WTP	WTP	WTP	WTP	WTP
Beacon Generation	WTP	WTP	WTP	WTP	WTP
Probe Response	WTP	WTP	WTP	WTP	WTP
Power mgmt Packet Buffering	WTP	WTP	WTP	WTP	WTP
Fragmentation/ Defragment.	WTP	WTP	WTP	WTP	WTP
Association Disassoc. Reassociation	AC	WTP	WTP	WTP	WTP
WME/11e -------------					
classifying	AC				WTP
scheduling	WTP	AC/WTP	WTP	WTP	WTP
queuing	WTP		WTP	WTP	WTP

```
        Authentication
        and Privacy
        --------------
        802.1X/EAP     AC       AC       AC/WTP    AC       AC/WTP

        Keys
        Management     AC       AC       WTP       AC       AC

        802.11
        Encryption/
        Decryption     WTP      WTP      WTP       WTP      WTP
```

   Figure 9: Mapping of 802.11 Functions for Local MAC Architecture

   From Figures 7, 8, and 9, it is clear that differences among vendors
   in the Local MAC Architecture are relatively minor, and most of the
   functional mapping appears to be common across vendors.

5.4. Split MAC

   As depicted in Figure 6 (b), the main idea behind the Split MAC
   architecture is to implement part of the 802.11 MAC functionality on
   a centralized AC instead of the WTPs, in addition to providing the
   required services for managing and monitoring the WTP devices.
   Usually, the decision of which functions of the 802.11 MAC need to be
   provided by the AC is based on the time-criticality of the services
   considered.

   In the Split MAC architecture, the WTP terminates the infrastructure
   side of the wireless physical link, provides radio-related
   management, and also implements time-critical functionality of the
   802.11 MAC. In addition, the non-realtime management functions are
   handled by a centralized AC, along with higher level services, such
   as configuration, QoS, policies for load balancing, and access
   control lists. The key distinction between Local MAC and Split MAC
   relates to non-realtime functions: in Split MAC architecture, the AC
   terminates 802.11 non realtime functions, whereas in Local MAC
   architecture, the WTP terminates the 802.11 non-realtime functions
   and consequently sends appropriate messages to the AC.

   There are several motivations for taking the Split MAC approach. The
   first is to offload functionality that is specific and relevant only
   to the locality of each BSS to the WTP, in order to allow the AC to
   scale to a large number of 'light weight' WTP devices. Moreover,
   realtime functionality is subject to latency constraints and cannot
   tolerate delays due to transmission of 802.11 control frames (or
   other realtime information) over multiple-hops. The latter would
   limit the available choices for connectivity between the AC and the

WTP. Therefore, the realtime criterion is usually employed to
separate MAC services between the devices. Another consideration is
cost reduction of the WTP to make it as cheap and simple as possible.
Finally, moving functions like encryption and decryption to the AC
reduces vulnerabilities from a compromised WTP, since user encryption
keys no longer reside on the WTP. As a result, any advancements in
security protocol and algorithm designs do not necessarily obsolete
the WTPs; the ACs implement the new security schemes instead, which
simplifies the management and update task. Additionally, the network
is protected against LAN-side eavesdropping.

Since there is no clear definition in the 802.11 specification as to
which 802.11 MAC functions are considered "realtime", each vendor
interprets this in their own way. Most vendors agree that the
following services of 802.11 MAC are examples of realtime services,
and are chosen to be implemented on the WTPs.

o Beacon Generation

o Probe Response/Transmission

o Processing of Control Frames: RTS/CTS/ACK/PS-Poll/CF-End/CF-ACK

o Synchronization

o Retransmissions

o Transmission Rate Adaptation

The following list includes examples of non-realtime MAC functions as
interpreted by most vendors:

o Authentication/De-authentication

o Association/Disassociation/Reassociation/Distribution

o Integration Services: Bridging between 802.11 and 802.3

o Privacy: 802.11 Encryption/Decryption

o Fragmentation/Defragmentation

However, some vendors may choose to classify some of the above "non-
realtime" functions as realtime functions in order to support
specific applications with strict QoS requirements. For example,
Reassociation is sometimes implemented as a "realtime" function to
support VoIP applications.

The non-realtime aspects of the 802.11 MAC are handled by the AC
through the processing of raw 802.11 management frames (Split MAC).
The following matrix in Figure 10 offers a tabular representation of
the design choices made by the six vendors that follow the Split MAC
design regarding the architecture considerations. While most vendors
support L3 connectivity between WTPs and ACs, some can only support
L2 switched connections due to the tighter delay constraint resulting
from splitting MAC between two physical entities across a network.
In Figure 7, it is clear that the WTP processes the 802.11 control
frames in both the Split MAC and Local MAC. The difference between
the two lies in the termination point for 802.11 management frames.
Local MAC terminates 802.11 management frames at WTP, while at least
some of the 802.11 management frames are terminated at the AC for the
Split MAC Architecture. Since in most cases WTP devices are IP-
addressable, any of the direct connection, L2-switched, or L3-routed
connections of Section 1.2 can be used. If only Ethernet-
encapsulation is performed (e.g., as in Architecture 4), then only
direct connection and L2-switched connections are supported.

	Arch1	Arch2	Arch3	Arch4	Arch5	Arch6
WTP-AC connectivity	L3	L3	L3	L2	L3	L3
802.11 mgmt termination	AC	AC	AC	AC	AC/WTP	AC
802.11 control termination	WTP	WTP	WTP	WTP	WTP	WTP
802.11 data aggregation	AC	AC	AC	AC	AC	AC

Figure 10: Architecture Considerations for Split MAC Architecture

Similar to the Local MAC Architecture, the matrix in Figure 11 shows
that most of the CAPWAP control functions are implemented at the AC.
The exception is RF monitoring, and in some cases RF configuration,
which are performed locally at the WTPs.

	Arch1	Arch2	Arch3	Arch4	Arch5	Arch6
RF Monitoring	WTP	WTP	WTP	WTP	WTP	WTP
RF Config.	AC/WTP		AC/WTP	AC	AC	AC
WTP config.	AC		AC	AC	AC	AC
WTP Firmware	AC		AC	AC	AC	AC
STA state info database	AC		AC	AC	AC	AC
AC/WTP mutual authent.	AC/WTP	AC/WTP	AC/WTP	AC/WTP		

Figure 11: Mapping of CAPWAP Functions for Split MAC Architecture

The most interesting matrix for Split MAC Architecture is the
Functional Distribution Matrix for 802.11 functions, as shown below
in Figure 12. Vendors map the functions onto the WTPs and AC with a
certain regularity. For example, all vendors choose to implement
Distribution, Integration Service at the AC, along with 802.1X/EAP
authentication and keys management. All vendors also choose to
implement beacon generation at WTPs. On the other hand, vendors
sometimes choose to map many of the other functions differently.
Therefore, Split MAC Architectures are not consistent regarding the
exact way the MAC is split.

	Arch1	Arch2	Arch3	Arch4	Arch5	Arch6
Distribution Service	AC	AC	AC	AC	AC	AC
Integration Service	AC	AC	AC	AC	AC	AC
Beacon Generation	WTP	WTP	WTP	WTP	WTP	WTP
Probe Response	WTP	AC/WTP	WTP	WTP	WTP	WTP
Power mgmt Packet Buffering	WTP	WTP	WTP	AC	AC/WTP	WTP
Fragmentation Defragment.	WTP		WTP	AC	AC	AC
Association Disassoc. Reassociation	AC	AC	AC	AC	WTP	AC
WME/11e --------------						
classifying			AC	AC	AC	AC
scheduling	WTP/AC	AC	WTP	AC	AC	WTP/AC
queuing	WTP/AC	WTP	WTP	AC	WTP	WTP

```
Authentication
 and Privacy
 --------------

802.1X/EAP     AC        AC        AC        AC        AC        AC

Keys
Management     AC        AC        AC        AC        AC        AC

802.11
Encryption/
Decryption     WTP       AC        WTP       AC        AC        AC
```

Figure 12: Mapping of 802.11 Functions for Split MAC Architecture

5.5. Remote MAC

One of the main motivations for the Remote MAC Architecture is to
keep the WTPs as light weight as possible, by having only the radio
interfaces on the WTPs and offloading the entire set of 802.11 MAC
functions (including delay-sensitive ones) to the Access Controller.
This leaves all the complexities of the MAC and other CAPWAP control
functions to the centralized controller.

The WTP acts only as a pass-through between the Wireless LAN clients
(STA) and the AC, though they may have an additional feature to
convert the frames from one format (802.11) to the other (i.e.,
Ethernet, TR, Fiber). The centralized controller provides network
monitoring, management and control, an entire set of 802.11 AP
services, security features, resource management, channel selection
features, and guarantees Quality of Service to the users. Because
the MAC is separated from the PHY, we call this the "Remote MAC
Architecture". Typically, such architecture is deployed with special
attention to the connectivity between the WTPs and AC so that the
delay is minimized. The Radio over Fiber (RoF) from Architecture 5
is an example of Remote MAC Architecture.

5.6. Comparisons of Local MAC, Split MAC, and Remote MAC

Two commonalities across all three Centralized Architectures (Local
MAC, Split MAC, and Remote MAC) are:

o Most of the CAPWAP functions related to network control and
   configuration reside on the AC.

o IEEE 802.11 PHY resides on the WTP.

There is a clear difference between Remote MAC and the other two
Centralized Architectures (namely, Local MAC and Split MAC), as the
802.11 MAC is completely separated from the PHY in the former, while
the other two keep some portion of the MAC functions together with
PHY at the WTPs. The implication of PHY and MAC separation is that
it severely limits the kind of interconnection between WTPs and ACs,
so that the 802.11 timing constraints are satisfied. As pointed out
earlier, this usually results in tighter constraint over the
interconnection between WTP and AC for the Remote MAC Architecture.
The advantage of Remote MAC Architecture is that it offers the
lightest possible WTPs for certain deployment scenarios.

The commonalities and differences between Local MAC and Split MAC are
most clearly seen by comparing Figure 7 to Figure 10. The
commonality is that 802.11 control frames are terminated at WTPs in
both cases. The main difference between Local MAC and Split MAC is
that the WTP terminates only the 802.11 control frames in the Split
MAC, while the WTP may terminate all 802.11 frames in the Local MAC.
An interesting consequence of this difference is that the Integration
Service, which essentially refers to bridging between 802.11 and
802.3 frames, is implemented by the AC in the Split MAC and by the
WTP in the Local MAC, as shown in Figures 9 and 12, respectively.

As a second note, the Distribution Service, although usually provided
by the AC, can also be implemented at the WTP in some Local MAC
architectures. This approach is meant to increase performance in
delivering STAs data traffic by avoiding tunneling it to the AC, and
relaxing the dependency of the WTP from the AC. Therefore, it is
possible for the data and control planes to be separated in the Local
MAC Architecture.

Even though all the 802.11 traffic is aggregated at ACs in the case
of Split MAC Architecture, the data and control planes can still be
separated by employing multiple ACs. For example, one AC can
implement most of the CAPWAP functions (control plane), while other
ACs can be used for 802.11 frames bridging (data plane).

Each of the three architectural variants may be advantageous for
certain deployment scenarios. While the Local MAC retains most of
the STA's state information at the local WTPs, Remote MAC centralizes
most of the state into the back-end AC. Split MAC sits somewhat in
the middle of this spectrum, keeping some state information locally
at the WTPs, and the rest centrally at the AC. Many factors should
be taken into account to determine the exact balance desired between
the centralized and decentralized state. The impact of such balance
on network manageability is currently a matter of dispute within the
technical community.

5.7. Communication Interface between WTPs and ACs

   Before any messages can be exchanged between an AC and WTP, the WTP
   needs to discover, authenticate, and register with the AC first, then
   download the firmware and establish a control channel with the AC.
   Message exchanges between the WTP and AC for control and
   configuration can happen after that. The following list outlines the
   basic operations that are typically performed between the WTP and the
   AC in their typical order:

   1. Discovery: The WTPs discover the AC with which they will be bound
       to and controlled by. The discovery procedure can employ either
       static or dynamic configuration. In the latter case, a protocol
       is used in order for the WTP to discover candidate AC(s).

   2. Authentication: After discovery, the WTP device authenticates
       itself with the AC. However, mutual authentication, in which the
       WTP also authenticates the AC, is not always supported since some
       vendors strive for zero-configuration on the WTP side. This is
       not necessarily secure as it leaves the possible vulnerability of
       the WTP being attached to a rogue AC.

   3. WTP Association: After successful authentication, a WTP registers
       with the AC in order to start receiving management and
       configuration messages.

   4. Firmware Download: After successful association, the WTP may
       pull, or the AC may push, the WTPs firmware, which may be
       protected in some manner, such as digital signatures.

   5. Control Channel Establishment: The WTP establishes either an IP-
       tunnel or performs Ethernet encapsulation with the AC in order to
       transfer data traffic and management frames.

   6. Configuration Download: Following the control channel
       establishment process, the AC may push configuration parameters
       to the WTPs.

5.8. Security

   Given the varied distribution of functionalities for the Centralized
   Architecture, as surveyed in Section 4.3, it is obvious that an extra
   network binding is created between the WTP and the AC. This brings
   new and unique security issues and subsequent requirements.

5.8.1. Client Data Security

   The survey shows clearly that the termination point for "over the
   air" 802.11 encryption [4] can be implemented either in the WTP or in
   the AC. Furthermore, the 802.1X/EAP [6] functionality is distributed
   between the WTP and the AC where, in most cases, the AC performs the
   necessary functions as the authenticator in the 802.1X exchange.

   If the STA and AC are the parties in the 4-way handshake (defined in
   [4]), and 802.11i traffic encryption terminates at the WTP, then the
   Pairwise Transient Key (PTK) has to be transferred from the AC to the
   WTP. Since the keying material is part of the control and
   provisioning of the WTPs, a secure encrypted tunnel for control
   frames is employed to transport the keying material.

   The centralized model encourages AC implementations to use one PMK
   for many different WTPs. This practice facilitates speedy transition
   by an STA from one WTP to another that is connected to the same AC
   without establishing a separate PMK. However, this leaves the STA in
   a difficult position, as the STA cannot distinguish between a
   compromised PMK and one that is intentionally being shared. This
   issue must be resolved, but the resolution is beyond the scope of the
   CAPWAP working group. The venue for this resolution is to be
   determined by the IEEE 802 and IETF liaisons.

   When the 802.11i encryption/decryption is performed in the AC, the
   key exchange and state transitions occur between the AC and the STA.
   Therefore, there is no need to transfer any crypto material between
   the AC and the WTP.

   Regardless of where the 802.11i termination point occurs, the
   Centralized WLAN Architecture records two practices for "over the
   wire" client data security. In some cases there is an encrypted
   tunnel (IPsec or SSL) between the WTP and AC, which assumes that the
   security boundary is in the AC. In other cases, an end-to-end
   mutually authenticated secure VPN tunnel is assumed between the
   client and AC, other security gateway, or end host entity.

5.8.2. Security of Control Channel between the WTP and AC

   In order for the CAPWAP functions to be implemented in the
   Centralized WLAN Architecture, a control channel is necessary between
   the WTP and AC.

   To address potential security threats against the control channel,
   existing implementations feature one or more of the following
   security mechanisms:

   1. Secure discovery of WTP and AC.

   2. Authentication of the WTPs to the ACs (and possibly mutual
       authentication).

   3. Confidentiality, integrity, and replay protection of control
       channel frames.

   4. Secure management of WTPs and ACs, including mechanisms for
       securely setting and resetting secrets and state.

   Discovery and authentication of WTPs are addressed in the submissions
   by implementing authentication mechanisms that range from X.509
   certificates, AAA authentication to pre-shared credential
   authentication. In all cases, confidentiality, integrity, and
   protection against man-in-the-middle attacks of the control frames
   are addressed by a secure encrypted tunnel between the WTP and AC(s),
   utilizing keys derived from the authentication methods mentioned
   previously. Finally, one of the motivations for the Centralized WLAN
   Architecture is to minimize the storage of cryptographic and security
   sensitive information, in addition to operational configuration
   parameters within the WTPs. It is for that reason that the majority
   of the submissions under the Centralized Architecture category have
   employed a post WTP authenticated discovery phase of configuration
   provisioning, which in turn protects against the theft of WTPs.

5.8.3. Physical Security of WTPs and ACs

   To provide comprehensive radio coverage, WTPs are often installed in
   locations that are difficult to secure physically; it is relatively
   easier to secure the AC physically. If high-value secrets, such as a
   RADIUS shared secret, are stored in the AC instead of WTPs, then the
   physical loss of an WTP does not compromise these secrets. Hence,
   the Centralized Architecture may reduce the security consequences of
   a stolen WTP. On the other hand, concentrating all the high-value
   secrets in one place makes the AC an attractive target that requires
   strict physical, procedural, and technical controls to protect the
   secrets.

6. Distributed Mesh Architecture

   Out of the sixteen architecture survey submissions, three belong to
   the Distributed Mesh Architecture family. An example of the
   Distributed Mesh Architecture is shown in Figure 13, and reflects
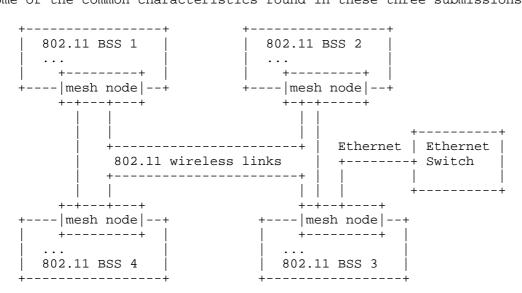   some of the common characteristics found in these three submissions.

```
    +-----------------+         +-----------------+
    | 802.11 BSS 1    |         | 802.11 BSS 2    |
    | ...             |         | ...             |
    |     +---------+ |         |     +---------+  |
    +----|mesh node|--+         +----|mesh node|--+
         +-+---+---+                 +-+-+-----+
           | |                         | |
           | |                         | |        +----------+
           | +-----------------------+ | Ethernet | Ethernet |
           |    802.11 wireless links | +--------+ | Switch   |
           |  +----------------------+ | |       |          |
           |  |                        | | |      +----------+
         +-+---+---+                 +-+--+----+
    +----|mesh node|--+         +----|mesh node|--+
    |    +---------+  |         |    +---------+  |
    | ...            |         | ...            |
    | 802.11 BSS 4    |         | 802.11 BSS 3    |
    +-----------------+         +-----------------+
```

          Figure 13: Example of Distributed Mesh Architecture

6.1. Common Characteristics

   To provide wider wireless coverage, mesh nodes in the network may act
   as APs to client stations in their respective BSS, as well as traffic
   relays to neighboring mesh nodes via 802.11 wireless links. It is
   also possible that some mesh nodes in the network may serve only as
   wireless traffic relays for other mesh nodes, but not as APs for any
   client stations. Instead of pulling Ethernet cable connections to
   every AP, wireless mesh networks provide an attractive alternative to
   relaying backhaul traffic.

   Mesh nodes can also keep track of the state of their neighboring
   nodes, or even nodes beyond their immediate neighborhood by
   exchanging information periodically amongst them; this way, mesh
   nodes can be fully aware of the dynamic network topology and RF
   conditions around them. Such peer-to-peer communication model allows
   mesh nodes to actively coordinate among themselves to achieve self-
   configuration and self-healing. This is the major distinction
   between this Distributed Architecture family and the Centralized
   Architecture -- much of the CAPWAP functions can be implemented

across the mesh nodes in a distributed fashion, without a centralized
entity making all the control decisions.

It is worthwhile to point out that mesh networks do not necessarily
preclude the use of centralized control. It is possible that a
combination of centralized and distributed control co-exists in mesh
networks. Some global configuration or policy change may be better
served in a coordinated fashion if some form of Access Controller
(AC) exists in the mesh network (even if not the full blown version
of the AC, as defined in the Centralized WLAN Architecture). For
example, a centralized management entity can be used to update every
mesh node's default configuration. It may also be more desirable to
leave certain functions, such as user authentication to a single
centralized end point (such as a RADIUS server), but mesh networks
allow each mesh AP to directly talk to the RADIUS server. This
eliminates the single point of failure and takes advantage of the
client distribution in the network.

The backhaul transport network of the mesh network can be either an
L2 or L3 networking technology. Currently, vendors are using
proprietary mesh technologies on top of standard 802.11 wireless
links to enable peer-to-peer communication between the mesh nodes.
Hence, there is no interoperability among mesh nodes from different
vendors. The IEEE 802.11 WG has recently started a new Task Group
(TGs) to define the mesh standard for 802.11.

6.2. Security

Similar security concerns for client data security, as described in
Section 5.8.1, also apply to the Distributed Mesh Architecture.
Additionally, one important security consideration for the mesh
networks is that the mesh nodes must authenticate each other within
the same administrative domain. To protect user and management data
that may not be secured at layer 3, data transmission among
neighboring nodes should be secured by a layer 2 mechanism of
confidentiality, integrity, and replay protection.

7. Summary and Conclusions

We requested existing WLAN vendors and other interested parties to
submit a short description of existing or desired WLAN access network
architectures to define a taxonomy of possible WLAN access network
architectures. The information from the 16 submissions was condensed
and summarized in this document.

New terminology has been defined wherever existing terminology was
found to be either insufficient or ambiguous in describing the WLAN
architectures and supporting functions listed in the document. For

example, the broad set of Access Point functions has been divided
into two categories: 802.11 functions, which include those that are
required by the IEEE 802.11 standards, and CAPWAP functions, which
include those that are not required by the IEEE 802.11, but are
deemed essential for control, configuration, and management of 802.11
WLAN access networks. Another term that has caused considerable
ambiguity is "Access Point", which usually reflected a physical box
that has the antennas, but did not have a uniform set of externally
consistent behavior across submissions. To remove this ambiguity, we
have redefined the AP as the set of 802.11 and CAPWAP functions,
while the physical box that terminates the 802.11 PHY is called the
Wireless Termination Point.

Based on the submissions during the architecture survey phase, we
have classified the existing WLAN architectures into three broad
classes:

1. Autonomous WLAN Architecture: Indicates a family of architectures
   in which all the 802.11 functions and, where applicable, CAPWAP
   functions are implemented in the WTPs.

2. Centralized WLAN Architecture: Indicates a family of architectures
   in which the AP functions are split between the WTPs and the AC,
   with the AC acting as a centralized control point for multiple
   WTPs.

3. Distributed WLAN Architecture: Indicates a family of architectures
   in which part of the control functions is implemented across a
   distributed network of peer entities.

Within the Centralized WLAN Architecture, there are a few visible
sub-categories that depend on how one maps the MAC functions (at a
high-level), between the WTP and the AC. Three prominent sub-
categories emerged from the information in the submissions:

1. Split MAC Architecture: The 802.11 MAC functions are split between
   the WTP and the AC. This subgroup includes all architectures that
   split the 802.11 MAC functions even though individual submissions
   differed on the specifics of the split.

2. Local MAC Architecture: The entire set of 802.11 MAC functions is
   implemented on the WTP.

3. Remote MAC Architecture: The entire set of 802.11 MAC functions is
   implemented on the AC.

The following tree diagram summarizes the architectures documented in
this taxonomy.

```
                   +----------------+
                   |Autonomous      |
      +---------->|Architecture    |
      |            |Family          |
      |            +----------------+
      |                                       +-------------+
      |                                       |Local        |
      |                              +---->|MAC          |
      |                              |        |Architecture |
      |                              |        +-------------+
      |                              |
      |            +----------------+  |     +-------------+
      |            |Centralized     |  |     |Split        |
      +---------->|Architecture    |--+---->|MAC          |
      |            |Family          |  |     |Architecture |
      |            +----------------+  |     +-------------+
      |                              |
      |                              |        +-------------+
      |                              |        |Remote       |
      |                              +---->|MAC          |
      |                              |        |Architecture |
      |                                       +-------------+
      |            +----------------+
      |            |Distributed Mesh|
      +---------->|Architecture    |
                   |Family          |
                   +----------------+
```

A majority of the submitted WLAN access network architectures (twelve
out of sixteen) followed the Centralized WLAN Architecture. All but
one of the Centralized WLAN Architecture submissions were grouped
into either a Split MAC Architecture or a Local MAC Architecture.
One submission followed the Autonomous WLAN Architecture, and three
followed the Distributed WLAN Architecture.

The WLAN access network architectures in the submissions indicated
that the connectivity assumptions were:

o Direct connection between the WTP and the AC.

o L2 switched connection between the WTP and the AC.

o L3 routed connection between the WTP and the AC.

   o Wireless connection between the mesh nodes in the distributed mesh
      architecture.

   Interoperability between equipment from different vendors is one of
   the fundamental problems in the WLAN market today. To achieve
   interoperability via open standard development, the following steps
   are suggested for IETF and IEEE 802.11.

   Using this taxonomy, a functional model of an Access Point should be
   defined by the new study group recently formed within the IEEE
   802.11. The functional model will consist of defining functional
   elements of an 802.11 Access Point that are considered atomic, i.e.,
   not subject to further splitting across multiple network elements.
   Such a functional model should serve as a common foundation to
   support the existing WLAN architectures as outlined in this taxonomy,
   and any further architecture development within or outside the IEEE
   802.11 group. It is possible, and even recommended, that work on the
   functional model definition may also include impact analysis of
   implementing each functional element on either the WTP or the AC.

   As part of the functional model definition, interfaces must be
   defined as primitives between these functional elements. If a pair
   of functional elements that have an interface defined between them is
   being implemented on two different network entities, then a protocol
   specification definition between such a pair of network elements is
   required, and should be developed by the IETF.

8. Security Considerations

   This document does not intend to provide a comprehensive threat
   analysis of all of the security issues with the different WLAN
   architectures. Nevertheless, in addition to documenting the
   architectures employed in the existing IEEE 802.11 products in the
   market, this taxonomy document also catalogues the security issues
   that arise and the manner in which vendors address these security
   threats. The WLAN architectures are broadly categorized into three
   families: Autonomous Architecture, Centralized Architecture, and
   Distributed Architecture. While Sections 4, 5, and 6 are devoted to
   each of these three architecture families, respectively, each section
   also contains a subsection to address the security issues within each
   architecture family.

   In summary, the main security concern in the Autonomous Architecture
   is the mutual authentication between the WTP and the wired (Ethernet)
   infrastructure equipment. Physical security of the WTPs is also a
   network security concern because the WTPs contain secret information
   and theft of these devices could potentially compromise even the
   wired network.

In the Centralized Architecture there are a few new security concerns
due to the new network binding between the WTP and AC. The following
security concerns are raised for this architecture family: keying
material for mobile client traffic may need to be securely
transported from the AC to WTP; secure discovery of the WTP and AC is
required, as well as mutual authentication between the WTPs and AC;
man-in-the-middle attacks to the control channel between WTP and AC,
confidentiality, integrity and replay protection of control channel
frames, and theft of WTPs for extraction of embedded secrets within.
Each of the survey results for this broad architecture category has
presented mechanisms to address these security issues.

The new security issue in the Distributed Mesh Architecture is the
need for mesh nodes to authenticate each other before forming a
secure mesh network. Encrypted communication between mesh nodes is
recommended to protect both control and user data.

9. Acknowledgements

This taxonomy is truly a collaborative effort with contributions from
a large group of people. First, we want to thank all the CAPWAP
Architecture Design Team members who have spent many hours in the
teleconference calls, over e-mails, and in writing and reviewing the
document. The full Design Team is listed here:

o Peyush Agarwal
   STMicroelectronics
   Plot# 18, Sector 16A
   Noida, U.P 201301
   India
   Phone: +91-120-2512021
   EMail: peyush.agarwal@st.com

o Dave Hetherington
   Roving Planet
   4750 Walnut St., Suite 106
   Boulder, CO 80027
   United States
   Phone: +1-303-996-7560
   EMail: Dave.Hetherington@RovingPlanet.com

o Matt Holdrege
   Strix Systems
   26610 Agoura Road
   Calabasas, CA 91302
   Phone: +1 818-251-1058
   EMail: matt@strixsystems.com

   o Victor Lin
      Extreme Networks
      3585 Monroe Street
      Santa Clara, CA 95051
      Phone: +1 408-579-3383
      EMail: vlin@extremenetworks.com

   o James M. Murphy
      Trapeze Networks
      5753 W. Las Positas Blvd.
      Pleasanton, CA 94588
      Phone: +1 925-474-2233
      EMail: jmurphy@trapezenetworks.com

   o Partha Narasimhan
      Aruba Wireless Networks
      180 Great Oaks Blvd
      San Jose, CA 95119
      Phone: +1 408-754-3018
      EMail: partha@arubanetworks.com

   o Bob O'Hara
      Airespace
      110 Nortech Parkway
      San Jose, CA 95134
      Phone: +1 408-635-2025
      EMail: bob@airespace.com

   o Emek Sadot (see Authors' Addresses)

   o Ajit Sanzgiri
      Cisco Systems
      170 W Tasman Drive
      San Jose, CA 95134
      Phone: +1 408-527-4252
      EMail: sanzgiri@cisco.com

   o Singh
      Chantry Networks
      1900 Minnesota Court
      Mississauga, Ontario L5N 3C9
      Canada
      Phone: +1 905-567-6900
      EMail: isingh@chantrynetworks.com

   o L. Lily Yang (Editor, see Authors' Addresses)

   o Petros Zerfos (see Authors' Addresses)

In addition, we would also like to acknowledge contributions from the
following individuals who participated in the architecture survey and
provided detailed input data in preparation of the taxonomy: Parviz
Yegani, Cheng Hong, Saravanan Govindan, Bob Beach, Dennis Volpano,
Shankar Narayanaswamy, Simon Barber, Srinivasa Rao Addepalli,
Subhashini A. Venkataramanan, Kue Wong, Kevin Dick, Ted Kuo, and
Tyan-shu Jou. It is simply impossible to write this taxonomy without
the large set of representative data points that they provided to us.
We would also like to thank our CAPWAP WG co-chairs, Mahalingam Mani
and Dorothy Gellert, and our Area Director, Bert Wijnen, for their
unfailing support.

10. Normative References

   [1] "IEEE WLAN MAC and PHY Layer Specifications", August 1999, <IEEE
        802.11-99>.

   [2] O'Hara, B., Calhoun, P., and J. Kempf, "Configuration and
        Provisioning for Wireless Access Points (CAPWAP) Problem
        Statement", RFC 3990, February 2005.

   [3] Bradner, S., "Key words for use in RFCs to Indicate Requirement
        Levels", BCP 14, RFC 2119, March 1997.

   [4] "IEEE Std 802.11i: Medium Access Control (MAC) Security
        Enhancements", April 2004.

   [5] "IEEE Std 802.11h: Spectrum and Transmit Power Management
        Extensions in the 5 GHz Band in Europe", October 2003.

   [6] "IEEE Std 802.1X: Port-based Network Access Control", June 2001.

Authors' Addresses

   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

   Petros Zerfos
   UCLA - Computer Science Department
   4403 Boelter Hall
   Los Angeles, CA 90095

   Phone: +1 310-206-3091
   EMail: pzerfos@cs.ucla.edu

   Emek Sadot
   Avaya
   Atidim Technology Park, Building #3
   Tel-Aviv 61131
   Israel

   Phone: +972-3-645-7591
   EMail: esadot@avaya.com