Analyzing Wireless LAN Security Overhead

Harold Lars McCarter

Thesis submitted to the Faculty of the Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Master of Science
In
Electrical Engineering

Scott F. Midkiff, Chair
Luiz A. DaSilva
Amir Zaghloul

April 17, 2006
Falls Church, Virginia

Keywords: Wireless Network Security, Encryption Overhead, Authentication Delay,
CCMP, TKIP, WEP, PEAP, LEAP

Analyzing Wireless LAN Security Overhead

Harold L. McCarter

(Abstract)

Wireless local area networks (WLAN) are beginning to play a much larger role in corporate network environments and are already very popular for home networking applications. This increase in accessibility has created large security holes for hackers and thieves to abuse, which is finally being addressed by stronger security methods such as advanced encryption algorithms and efficient authentication processes. However, these security methods often hamper network performance unbeknownst to engineers and users.

This research examines the effects of Wired Equivalent Privacy (WEP), Temporal Key Integrity Protocol (TKIP), and Counter Mode/CBC-MAC Protocol (CCMP) encryption algorithms on throughput rates for IEEE 802.11 networks as well as the authentication times for Lightweight Extensible Authentication Protocol (LEAP) and Protected Extensible Authentication Protocol (PEAP). The research shows that today's wireless hardware is capable of reducing overhead of even the most advanced encryption schemes to less than five percent of the total bandwidth.

Table of Contents

List of Tables

List of Figures

Chapter 1 Introduction

1.0 Motivation

Wireless networks have exhibited significant growth within the last few years in both home and corporate environments due in part to low cost and increased hardware quality. This growth has fueled new applications for wireless networks ranging from advanced warehouse inventory systems to wireless voice over internet protocol (VoIP) phones. The ease of use and vast distribution of these systems has created a security nightmare for home users and network administrators, which has become widely publicized in the media. Numerous articles now detail wireless vulnerabilities that have shut down corporate networks and allowed unknowing home users to become launching points for illegal activity.

As security managers begin to shore up defensive efforts on their wireless networks they decrease the chance that their network will fall victim to a malicious hacker or corporate thief. However, there is a cost to this increase of security with regards to money, time, and even network performance. Because wireless networking is still a relatively new technology many engineers are unaware of these consequences when applying them to their networks. A large goal of this research is to determine if there is any reason to forgo implementation of wireless security mechanisms because of potential degradation in performance.

1.1 Problem Statement

Security on wireless local area networks (WLANs) is a requirement in today's rapid deployment of this technology. Unfortunately, few people understand the methods, yet alone the consequences, of implementing these security measures. Compounding the problem is that there is little information available today on the cost of encryption overhead or the affects of long authentication delays on network performance. The sources that are available often fail to incorporate the latest encryption schemes and authentication methods into their analysis, such as those specified in the new Institute of Electrical and Electronics Engineers (IEEE) 802.11i standard that was ratified in 2004.

To address this problem, measurements were conducted on a test bed network to determine some of the overhead associated with wireless network security. Much of this research focuses on the throughput reduction of networks caused by advanced encryption schemes prevalent in many of today's corporate networks in addition to the authentication methods those networks implement.

Similar studies have produced mixed results with often inconclusive data, particularly with regards to encryption overhead. This research seeks to not only present new data points, but to tie together previous work in this area.

1.2 Thesis Outline

This thesis is divided into seven chapters. Each chapter builds upon the previous to develop a knowledge basis necessary for interpreting the results of Chapter 7. Those individuals who are familiar with networking technologies and security procedures may find it acceptable to skip directly to Chapter 4.

Chapter 2 provides an overview of wireless technologies and the security methods currently available today to protect those networks. This chapter was constructed to provide a basic overview of wireless terminology to help readers understand the data presented at the end of the thesis.

Chapter 3 details previous work done in the area of security analysis and presents the author's opinions on these studies. Many of techniques utilized in this research were drawn from these studies.

Chapter 4 outlines experimental background information that was important in developing the testing procedures of Chapter 5.

Chapter 5 summaries the experimental procedures that are utilized for the measurement campaign. It also provides an overview of the methods used to process the data in Chapter 6.

Chapter 6 contains the results from the various trials.

Chapter 7 presents a discussion of the results from Chapter 6 and recommends areas for future work.

1.3 Intended Contributions

The main focus of this thesis is the measured overhead associated with various forms of wireless network encryption and authentication in corporate networks. The intended contribution includes numerical averages for the percent overhead that is generated by each of these encryption mechanisms and average times for authentication, with the hope of providing engineers usable numbers for network deployment consideration. This could be extremely useful where wireless might be used in distribution or core roles on a network such as remote bridging of wired networks or where wireless access points may comprise the network backbone, which may be the case in many small office environments.

Chapter 2 Wireless Technology

2.0 Overview

This chapter provides an overview of current wireless technologies and security schemes that are part of the IEEE 802.11 standard.

Because this research focuses on the potential effects of enhanced wireless security on network performance readers should be familiar with various topics including the physical layer of IEEE 802.11, how authentication and encryption work on a secured wireless network, and how to observe these processes on the network.

This chapter begins with an in-depth look at the IEEE 802.11 protocol in order to note differences between an unsecured wireless network versus one that is protected by various layers of encryption and authentication. The chapter then provides a brief overview of the IEEE 802.11b and 802.11g standards. Finally, the chapter finishes with a complete overview of various encryption and authentication methods that are present on secured wireless networks and how they play into the IEEE 802.11i standard.

The vast majority of the IEEE 802.11 background was drawn from [O'Hara05] and the majority of all security background information was drawn from [Sankar04], [CWSP03], and [CWNA05]. Together these texts provided virtually every piece of information presented in this chapter.

2.1 IEEE 802.11 Standard

IEEE 802.11 was the first widely-used wireless local area networking standard and was selected for use in 1997. The standard consists of a medium access control (MAC) sublayer, MAC management protocols and services, and three physical layers (PHYs). The three PHYs were an infrared PHY, a frequency hopping spread spectrum (FHSS) radio PHY, and a direct sequence spread spectrum (DSSS) radio PHY. These original PHYs provided data transfer rates of 1 Mbps and 2 Mbps [O'Hara05].

The 1999 revision included two more PHYs, IEEE 802.11a and 802.11b, which would become standards in the industry with data transfer rates of 54 Mbps and 11 Mbps, respectively. The difference between the two new PHYs was that IEEE 802.11a operated with an orthogonal frequency division multiplexing (OFDM) signal at Unlicensed National Information Infrastructure (U-NII) bands verses the DSSS signal used at 2.4 GHz for IEEE 802.11b. In 2002 the widely used IEEE 802.11g standard was developed as an extension of IEEE 802.11b, providing backwards compatibility [O'Hara05].

2.1.1 MAC Layer

The MAC sublayer provides reliable data transmission for the IEEE 802.11 standard similar to a wired network. To this extent, the MAC sublayer provides three functions: a reliable method to transmit data for users, shared access to the medium among users, and the protection of transmitted data accomplished through encryption.

Because the transmission of IEEE 802.11 signals occurs wirelessly these functions must be conducted differently in the MAC sublayer because signals that are transmitted cannot simply be assumed to have been received on a wireless system.

2.1.1.1 Reliable Data Delivery

The first function, reliable delivery, is completed with a series of two frames, as shown in Figure 2.1. One is sent by the wireless client to the access point and the second is an acknowledgement frame sent from the access point to the client indicating that the frame was received. If there was no acknowledgement frame received at the client then that station can assume the access point did not receive the first frame and the client can retransmit it after a certain wait time.

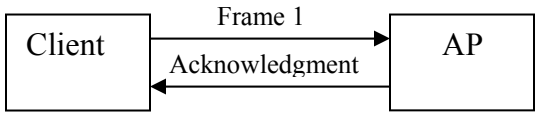

Figure 2-1. 802.11 Delivery

There is a conflict with this process that is often referred to as the "hidden node problem" [O'Hara05]. The problem occurs when one client is not in a position where it can communicate direct with another client but both clients are in a position to communicate with a common station. One station may not know that another station is transmitting and, therefore, causes a collision by transmitting a frame of its own.

To address this problem the protocol provides an optional solution with two additional frames called the request-to-send (RTS) and the clear-to-send (CTS) frames. Before a client transmits it first sends a RTS indicating its intention to send a frame; however, it will not transmit information until it receives a CTS from the destination. Because the use of these two additional frames can reduce the data throughput rate of the network it is not enabled in all situations [O'Hara05].

2.1.1.2 Shared Access

The second task of ensuring shared access to all clients is accomplished through two access mechanisms: the basic access mechanism which utilizes the distributed coordination function (DCF) and the centrally controlled access mechanism which utilizes the point coordination function (PCF).

The basic access mechanism of IEEE 802.11 utilizes carrier sense multiple access with collision avoidance (CSMA/CA) and binary exponential backoff. This access mechanism uses a "listen before you talk" approach and ensures that if the destination is already handling traffic another client will not attempt to transmit as well, avoiding a collision. If a client detects another transmission in progress it will wait a set amount of time, called the contention window (CW), before it attempts its own transmission. This value increases each time that a client detects a transmission in progress to increase the chance that the medium is available for the next transmission. This value is standard for each PHY [O'Hara05].

The DCF, which is the functional unit of the basic access mechanism, operates by checking both the physical and virtual carrier sensing mechanisms. In the event both of these mechanisms indicate that there is no transmission for a set period, based on timing intervals, then the MAC may begin a transmission. These timing intervals provide a station with a set time to wait before beginning transmission in order to help prevent collisions [O'Hara05].

The PHY determines two intervals: the short interframe space (SIFS) and the slot time. From these, three additional intervals are derived: the priority interframe space (PIFS), the distributed interframe space (DIFS), and the extended interframe space (EIFS). Each of these timing intervals changes depending on the number of times that a transmission is detected while a station is attempting to transmit [O'Hara05]. Timing intervals are illustrated in Figure 2-2.

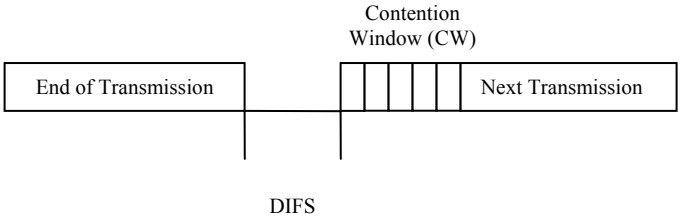

Figure 2-2. 802.11 Timing Intervals

The centrally controlled access mechanism, which utilizes the PCF, uses a poll and response protocol for the medium. This is an optional protocol that is housed within the access point and operates over the DCF, providing another method of preventing collisions. It operates by requiring stations to be added to a polling list within the access point providing traffic information to the stations [O'Hara05].

2.1.2 PHY

The PHY of IEEE 802.11 provides three levels of functionality: the coordination of frame exchanges between the MAC and the PHY under the control of the physical layer convergence procedure (PLCP) sublayer, the use of signal carrier and spread spectrum modulation to transmit frames over the radio frequency medium under the control of the

physical medium dependent (PMD) sublayer, and providing carrier sense indication back to the MAC to verify activity on the media [O'Hara05].

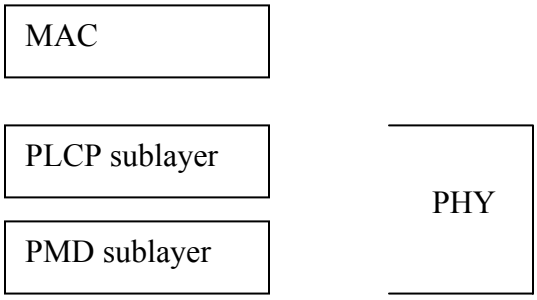

Figure 2-3. IEEE 802.11 Layers

The PHYs exist below the MAC Layer and handle data transfers in one of the three ways previously mentioned. The most widely used is direct sequence spread spectrum (DSSS), which provides eleven channels in the US in the 2.4 GHz band for IEEE 802.11. Each of these channels occupies 22 MHz of bandwidth with a 5 MHz channel separation between them, allowing for three non-overlapping channels—1,6, and 11 [O'Hara05].

2.2 IEEE 802.11b & 802.11g

Subtle differences exist between IEEE 802.11b and 802.11g that are important as they can affect network performance. IEEE 802.11b operates with complimentary code keying (CCK), so IEEE 802.11g incorporated this into its standard to maintain backwards compatibility with IEEE 802.11b. However, the backwards compatibility comes at a cost as overall throughput is reduced on IEEE 802.11g devices when the network is operating in this legacy mode. Additionally, IEEE 802.11g extends the orthogonal frequency division multiplexing (OFDM) from IEEE 802.11a to allow for 54 Mbps throughput [O'Hara05].

2.3 Authentication and Encryption Protocols

The IEEE 802.11 standard has several methods of encryption and authentication that provide varying levels of security for wireless networks [Sankar04]. This section provides an overview of those methods.

Authentication provides a method for wireless networks to verify the identity of a user and ensure they are authorized access to the network before being connected. This process allows an organization to restrict access of its wireless network to certain individuals just as it would restrict access to its wired network. Without proper authentication a wireless client will not be able to associate with a wireless access point and therefore will be unable to gain access to network resources.

Encryption is a process of shielding transmitted data by changing the structure of the data with a known process by one of the following two methods: the use of a symmetric key paradigm or an asymmetric key paradigm. Encryption helps prevent interception of transmitted data for potential malicious use.

2.3.1 Authentication

As previously mentioned, before a wireless client can gain access to network resources, such as an internet connection, it must first associate with a wireless access point. Once this is completed the access point will forward all network information to that client, much like a wired network. Because of this the process of association must ensure that only legitimate clients gain access to the network. This is where authentication comes in.

There are several authentication methods and protocols that can be implemented within a wireless network; however, only certain ones are of interest in this research. The authentication protocol used for this research is Remote Address Dial-In User Service (RADIUS) which was developed in 1996. Additionally, the methods of authentication that the research focuses on are: 802.1x, Extensible Authentication Protocol (EAP), Protected EAP (PEAP), and the Lightweight EAP (LEAP) [Sankar04].

Figure 2-4 shows the interaction between the various authentication methods and how they fall into the process. Authentication begins at the client which passes identity information through the access point to the authentication server where credentials are verified. Once verification is complete the access point will grant access to the wireless client.

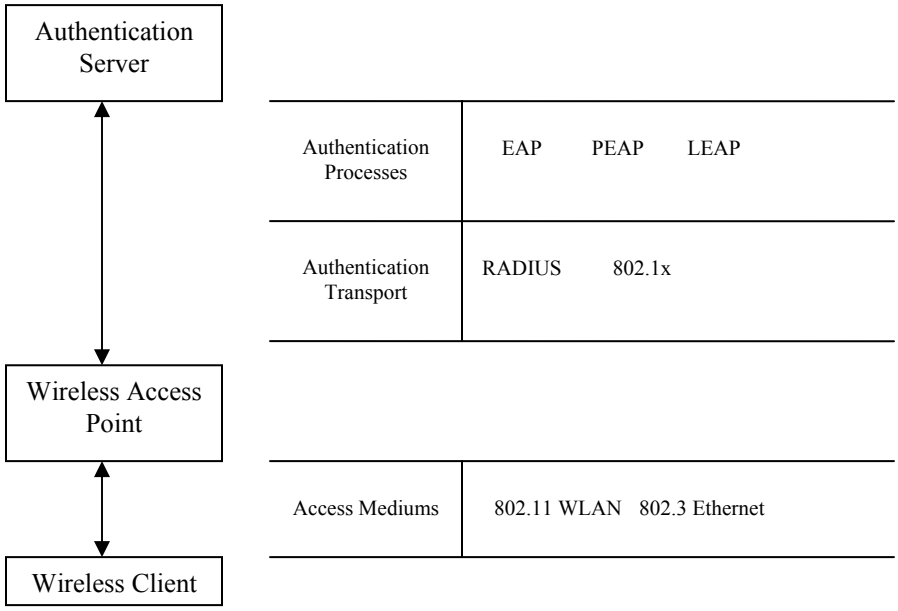

Figure 2-4. Authentication Process

The following information provides an overview of these authentication methods and protocols so that it can be established how they introduce overhead to IEEE 802.11 networks.

2.3.1.1 RADIUS

The RADIUS service acts as a mechanism to cross check authentication information against stored credentials to determine if a client should be allowed access to the network. RADIUS servers can be located anywhere within a network and can either rely on stored user credentials or a database of existing user credentials such as the usernames and passwords stored within the Active Directory Service on Microsoft Windows Server. RADIUS can interact with those databases via the Lightweight Directory Access Protocol (LDAP), as shown in Figure 2.5.

RADIUS servers have been used successfully for many years on wired networks to provide an authentication mechanism so it is only natural that they continue to be used extensively for wireless authentication. This allows organizations to continue utilizing current hardware and software during a wireless migration. In the wireless environment the mutual authentication ability of RADIUS can be used to authenticate not only the client but the wireless access point, helping to prevent man-in-the-middle attacks [Sankar04].

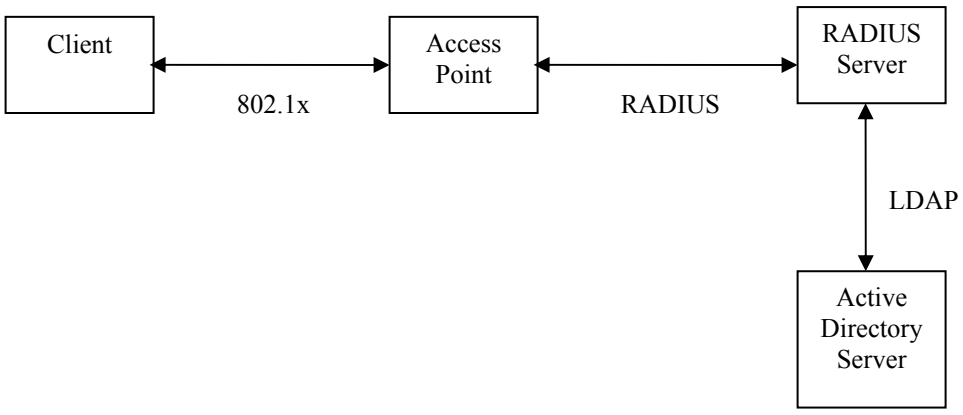

Figure 2-5. Example RADIUS Configuration

RADIUS is a preferred authentication method for several reasons. The protocol is scalable and easily interacts with other authentication mechanisms such as IEEE 802.1x and EAP. The scalability of the protocol and ability to cluster servers is significant among large organizations that may rely on multiple authentication servers and backup layers to ensure continued operation.

2.3.1.2 Overview of IEEE 802.1x

The IEEE 802.1x standard provides port-based security to both wired and wireless networks. It accomplishes this by utilizing the EAP protocol to pass authentication information over the network to the authentication system, which is usually a RADIUS server. IEEE 802.1x authentication requires three stations in place: the supplicant which is the user or client that wants to be authenticated, the authentication server, and a device that will pass information between the supplicant and the authentication server called the authenticator.

Port-based authentication begins when the supplicant connects to a closed port or the authenticator recognizes a supplicant client on a closed port. The supplicant then sends an EAP-Response/Identity packet to the authenticator which passes it on to the authentication server. The authentication server responds with a challenge requiring the supplicant to supply identity information such as a password. If the returned information is correct then the authentication server instructs the authenticator to allow access of all network traffic to the supplicant [Snyder02].

There are a few important points with regards to IEEE 802.1x. The first is that it provides a method to supply user or client-based authentication to a network with individual user names and passwords, tokens, certificates, or other methods. This is important because it can be effectively utilized in large scale networks as an authentication medium for both an organization's wired and wireless networks with almost no overhead on those networks. A second point is that the use of various types of EAP over IEEE 802.1x significantly improves the security of these individual authentication processes by generating keys which are recycled after certain intervals. Finally, in a high security wireless network IEEE 802.1x allows the credentials for authentication to be stored outside of the access point, increasing physical security.

2.3.1.3 EAP

EAP was first used in the Point-to-Point Protocol (PPP) as a method of establishing connections over dial-up. Since then EAP has been adapted for use in the wireless domain as a method to pass logon credentials between a wireless user and an authentication server. EAP and IEEE 802.1x work together to pass this logon information between the client and the authentication server [CWSP03].

As previously discussed IEEE 802.1x is a transport medium for EAP frames. When a client connects to a closed port, IEEE 802.1x opens that port for the transportation of EAP credential frames between the supplicant and the authentication server through the authenticator. The general process is illustrated in Figure 2-6.

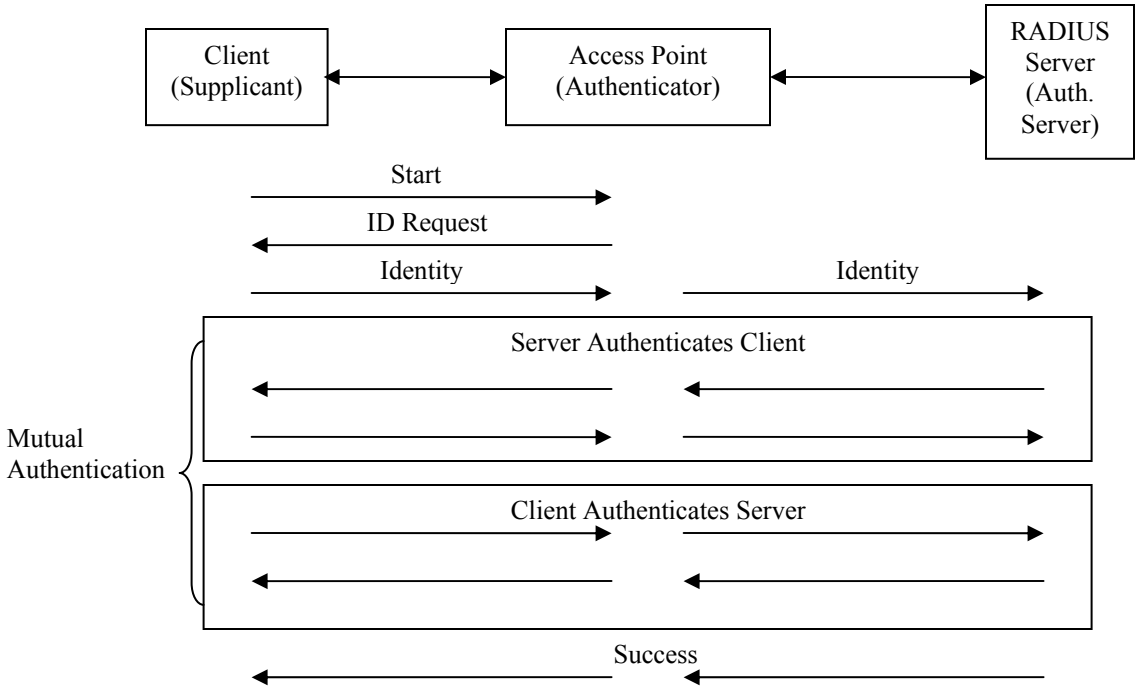

Figure 2-6. EAP Authentication

Since EAP and its various subsets support a variety of authentication methods (certificates, tokens, biometrics, etc.) the information can be passed on through the network without requiring any intermediary steps or settings. This is important on a network that may have varying levels of security between clients.

An additional security aspect of EAP is mutual authentication, which is presented in Figure 2-6. Certain subsets of EAP require the client to authenticate the server in addition to the server authenticating the client. This helps to greatly reduce the possibility of man-in-the-middle techniques as mentioned in the RADIUS section.

2.3.1.4 PEAP

PEAP is essentially EAP messages encapsulated in a Transport Layer Security (TLS) tunnel. It also provides enhanced security with the use of server side certificates and optional client side certificates. The protocol was developed to help address several security concerns with EAP: unprotected user information during the EAP negotiation, lack of support for fast connections when roaming, and lack of support for fragmentation and reassembly [CWSP03].

PEAP operates in two phase: the first is the creation of a secure tunnel using TLS and the second is client authentication using standard EAP methods. The differences from EAP can be noted in Figure 2-7. As indicated, once server side authentication is accomplished through the use of the server certificate, an encrypted tunnel is established for the transfer of EAP authentication information. Once authentication is successful the keys are

exchanged to initiate the encryption scheme of choice for the network, which is discussed in greater detail later in the chapter.

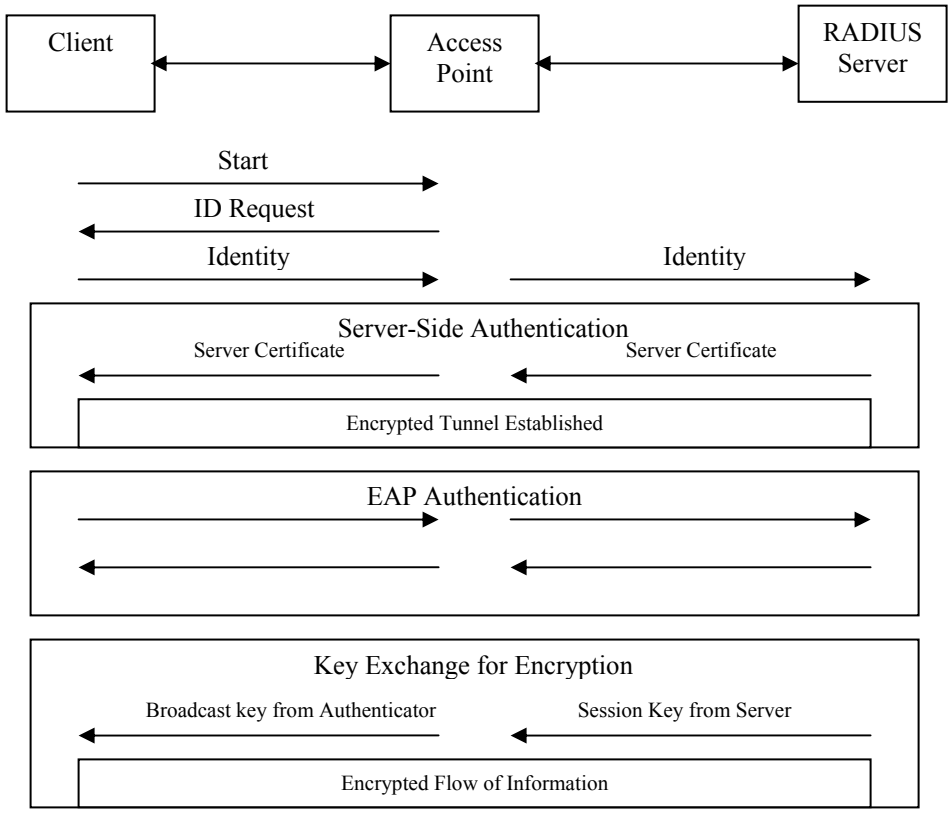

Figure 2-7. **PEAP Authentication**

2.3.1.5 LEAP

LEAP is a proprietary protocol developed by Cisco Systems to address Wired Equivalent Privacy (WEP) vulnerabilities. LEAP utilizes the passwords from a user database such as Microsoft Active Directory to create a hash of the actual password which is sent for authentication. This prevents passwords from being sent in the clear. Although LEAP is an effective security protocol it is not as robust as the certificate requirements of PEAP as it is susceptible to dictionary attacks. However, in an organization with no ability to stand up a certificate infrastructure it is the next best option [Sankar04].

2.3.2 Encryption

Encryption provides a method for wireless networks to provide end-to-end security on data streams. IEEE 802.11 networks currently have three encryption protocols available for use today: WEP, Temporal Key Integrity Protocol (TKIP), and Counter Mode/CBC-MAC Protocol (CCMP). Although WEP does not provide the security required by most networks, and TKIP and CCMP are quickly becoming the minimum standards to use for

data encryption on wireless networks, it is still in wide use and is examined in this research.

These protocols each rely on different methods to encrypt data with some form of key. This keying process typically introduces a certain amount of overhead into network communications, which is a critical part of this research. As such the manner in which these various protocols encrypt data will be covered.

2.3.2.1 WEP

The WEP protocol was originally developed to provide the same level of security as a wired network with three goals in mind: prevent disclosure of packets in transit, prevent modification of those packets, and to provide access control to the network. However, after the delivery of the WEP algorithm several vulnerabilities were discovered that severely hamper its ability to perform these functions.

WEP keys are created with two lengths: 40 and 104. However, because each WEP key includes a 24-bit initialization vector the total key lengths are 64- and 128-bits, which are the commonly used terms in the industry. The initialization vector (IV) provides added security to data as it changes with each packet. Figure 2-8 shows an IEEE 802.11 frame with WEP. The italics portions of the frame are those associated with WEP [Sankar04].

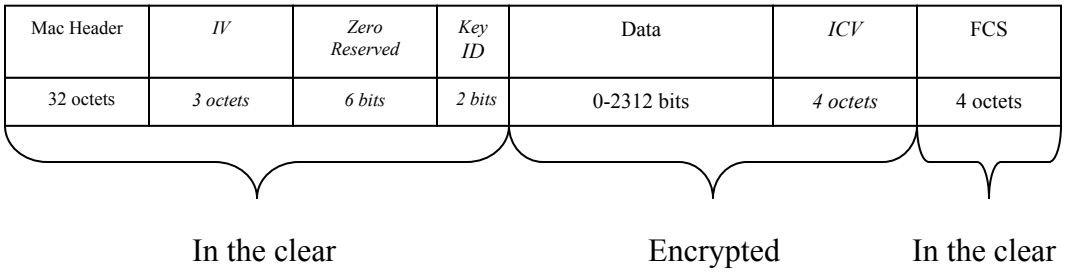

Figure 2-8. WEP Frame

The algorithm used to construct WEP keys is based on the RC4 algorithm developed by RSA Security. This is a priority stream cipher that was intended to be recycled after each key. However, WEP was designed to use the same pre-shared key (up to four different keys) for each packet which creates a huge security concern. To address the problem, the IV was developed to be attached to each WEP key, creating a WEP seed that would be different for every packet [Sankar04]. Unfortunately, the IV was not set to be unique and nonrepeating for each packet, which left further vulnerability in the algorithm. The integrity check vector (ICV) at the end of the WEP frame is a four-octet linear checksum intended to alert a station when a packet has been modified. This is commonly referred to as CRC-32. If something has been changed within a packet then the checksum will not match.

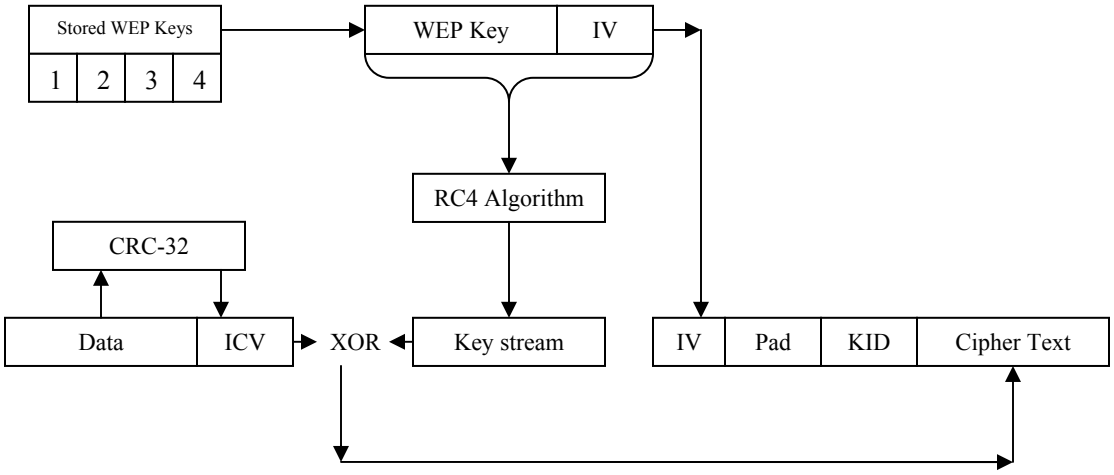

Figure 2-9. WEP Encapsulation

A WEP packet encapsulation begins with the choice of an IV. Different vendors increase the IV in different manners but since it is passed in the clear there is no problem associated with these varying methods. Next the IV is combined with the WEP key to generate the WEP seed. This seed is then run through the RC4 algorithm to generate the final key stream that will be combined with the data to encrypt it. As shown in Figure 2-9, the ICV is calculated with the CRC-32 algorithm and the data to be encapsulated. This data stream is combined with the key stream to produce cipher text which is then added to the IV and the key ID octet (labeled Pad and KID in the figure) [Sankar04].

When the WEP packet needs to be decapsulated at the receiving end the process is essentially reversed with no additional requirements, since the station at the receiving end already has the key stored and the IV vector used is sent in the clear. Because the RC4 algorithm is not computationally intensive the WEP frames can be generated quickly. This combined with the fact that WEP appends only an additional 8 octets to the frame ensures relatively low overhead for the network.

2.3.2.2 TKIP

TKIP was developed to address the vulnerabilities associated with WEP and is used in Wi-Fi Protected Access (WPA). TKIP was developed to provide backwards compatibility with WEP to prevent the need to replace all hardware that only supported WEP at the time. TKIP can be constructed by utilizing the IV, RC4 algorithm, and the ICV that WEP already used. [CWSP03]

TKIP essentially consists of three algorithms: a cryptographic message integrity algorithm, an enhancement to the IV, and a key mixing algorithm [Sankar04]. The largest change with TKIP is that it creates a new key for each packet significantly reducing the possibility of guessing a key. Additionally, the encryption scheme provides

a hash to prevent packet modification and replaces the flawed IV system by enforcing a longer packet counter and replay protection.

The message integrity check was designed with the intent of overcoming WEP's vulnerable ICV by creating a non-linear hash to prevent packet modification. Designers eventually settled on the Michael algorithm because of its relatively low computational overhead. Additionally, TKIP provides for logging, disabling, and deauthentication when a station incorrectly attempts to guess the Michael algorithm in too short of time. This ensures that new keys are generated in these instances making it very difficult to guess the ICV [CWSP03].

As indicated by the italic text in Figure 2-10, there are a total of 20 octets associated with TKIP in an IEEE 802.11 frame. This is more than twice the amount of overhead associated with a WEP frame, and is due to the extended IVs and the MIC that are used within the TKIP protocol [Sankar04].

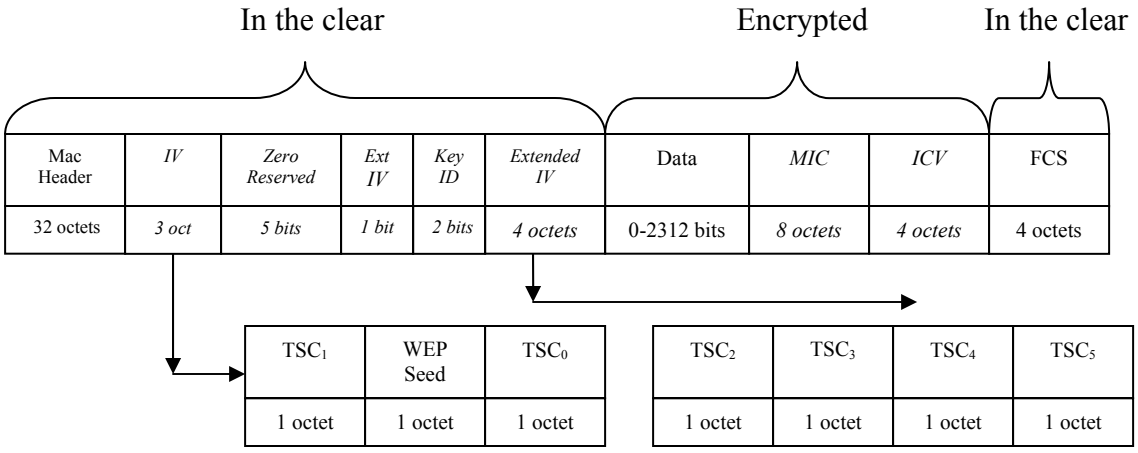

Figure 2-10. TKIP Frame

As mentioned earlier, TKIP also addresses replay attacks by adding a TKIP Sequence Counter (TSC) which prevents reuse of an IV. This is the second algorithm associated with TKIP in which a 48-bit counter is employed to ensure a unique IV for each packet. The TSC is broken down into six octets (TSC_0 through TSC_5) that are sent unencrypted as seen in Figure 2-10. This algorithm also helps prevent denial of service (DoS) attacks by ensuring that the receiver does not update the TSC until the MIC has been verified after each packet.

The final key mixing algorithm protects the Temporal Encryption Key (TEK). This is the key that is used by management algorithms to exchange keys for each authentication. This key rotation helps increase security. The TEK is used by the key mixing algorithm to combine the TEK, TSC, and transmitter address (TA) into a 128-bit WEP seed that is unique for each packet. The algorithm further avoids known keys in the RC4 process that are known to be weak. As shown in Figure 2-11, the key mixing algorithm is broken

down into two parts. During phase 1 the TKIP mixed Transmit Address and Key (TTAK) are generated by the combination of the TSC, TA, and TEK components. The process was intended to have a low computational overhead; however, it still takes some time to complete because of the multiple processes occurring simultaneously. During phase 2 the TTAK is combined with a full TEK and TSC to generate the 128-bit WEP seed [Sankar04].

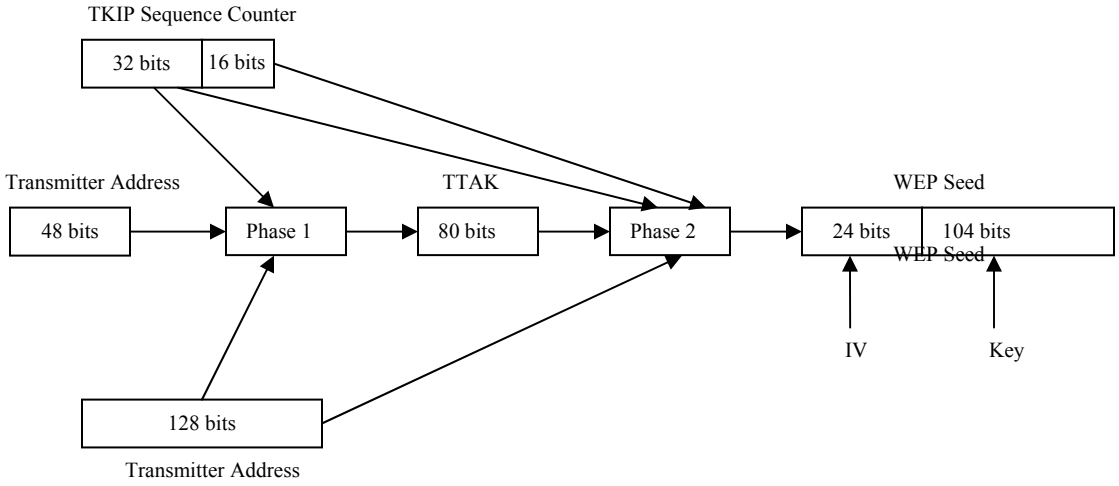

Figure 2-11. TKIP Key Mixing

With the generation of the WEP seed complete, the TKIP packet can be encapsulated with a process that is similar to the WEP encapsulation process with a few additional steps. Figure 2-12 illustrates those steps. Once the WEP seed is generated the data to be sent is run through the Michael algorithm to create the MIC key. This combination of data is combined with the CRC-32 algorithm to create the ICV which is appended to the data. As the WEP seed is run through the RC4 algorithm and the key stream is generated it is combined with the data to create the final packet [Sankar04].

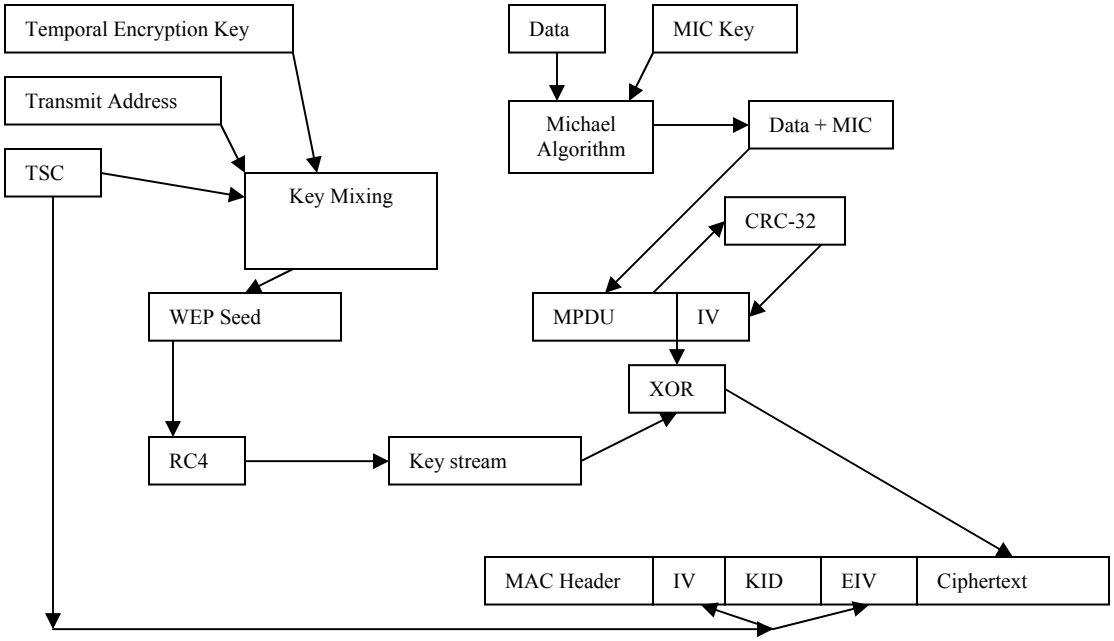

Figure 2-12. TKIP Encapsulation

The process of decapsulating a TKIP packet is the opposite of the above encapsulating process with the addition of the integrity checks. These checks are done with the ICV and the MIC to make sure that they match. If a mismatch occurs appropriate 802.11 protocols will be followed, such as packet retransmission.

The TKIP encapsulation mechanism is composed of several sub processes all working continuously to provide the encrypted packets. Although TKIP was designed not to be computationally intensive, clearly there is more work associated with TKIP packet creation than with WEP.

2.3.2.3 CCMP

CCMP is the most advanced encryption available for wireless networks today and is central to the Robust Security Network (RSN) portion of 802.11i, which will be discussed later. CCMP is based on the Advanced Encryption Standard (AES) that has currently been accepted by the US Government as a standard encryption tool. Unlike TKIP or WEP which rely on the proprietary RC4 algorithm, AES is not patented and can be used by anyone. The protocol is capable of confidentiality with its Counter Mode and ensures data integrity based on the Cipher Block Chaining Message Authentication Codes (CBC-MAC). Although AES allows different values for key and block length, the CCMP protocol utilizes 128 bits for each of these [Sankar04].

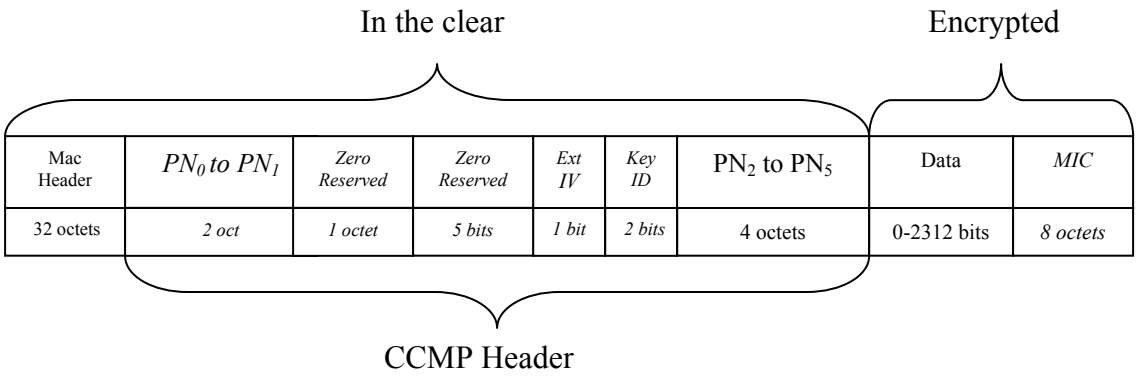

Figure 2-13. CCMP Frame

As shown in Figure 2-13, the total CCMP header length is 8 octets followed by 8 octets in the MIC. The CCMP portions are in italics. The process of CCMP encapsulation provides three types of protection: confidentiality, message integrity, and replay prevention. Confidentiality is provided by the CCM encryption within the block cipher process. Integrity is guaranteed by the development of a MIC, which is calculated differently from TKIP using AES CBC-MAC mode. Finally replay protection is ensured by the used of a packet counter, called the PN, which helps ensure packet uniqueness in CCM encryption [Sankar04].

The process of encapsulation is illustrated in Figure 2-14 below. From this it is evident how the MIC is created differently than TKIP and how the PN ensures that no two packets are the same. Although the algorithms associated with CCMP are rather complex, which is why they are not being addressed further, the actual process of encapsulation is relatively simple allowing for the associated computation to be completed quickly.

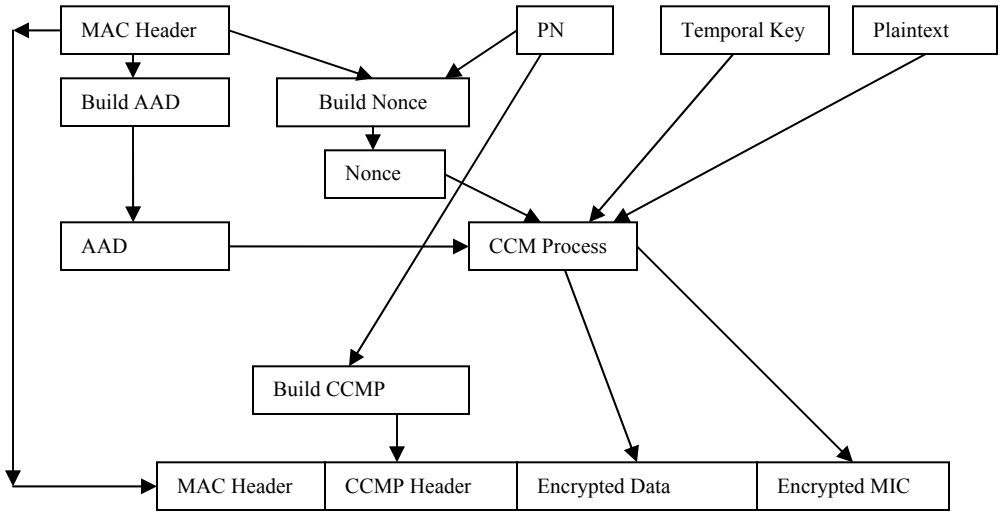

Figure 2-14. CCMP Encapsulation

2.4 IEEE 802.11i Terminology Overview

There are several terms present in the industry today describing various levels of security on IEEE 802.11 networks such as WPA, WPA2, robust security network (RSN), and transition security network (TSN) that may or may not be related depending on the source. It is important to understand that these are all subsets of the IEEE 802.11i standard, but each has different specifications.

- WPA – This specification was developed to address the security vulnerabilities of WEP and is protected by TKIP, and in a later version AES CCMP. WPA operates in two modes: WPA-Personal, which allows for the use of a pre-shared key for the TKIP algorithm and does not require authentication, or WPA-Enterprise, which provides a method of authentication via IEEE 802.1x/EAP for RADIUS connections and the use of session keys for TKIP. WPA was created when IEEE 802.11i was still in development and attempted to capture as many of its security settings as possible. [CWNA05]
- WPA2 – This is essentially the certified name for IEEE 802.11i by the Wi-Fi Alliance, and can be thought of as synonymous with IEEE 802.11i [CWNA05]. The main difference between WPA and WPA2 is the requirement of CCMP encryption with WPA2. Authentication methods between the two are similar but need to be implemented for a network to be IEEE 802.11i compliant; however, some products differentiate WPA2 into Personal and Enterprise, similarly to WPA, which allows for only encryption with the use of pre-shared keys on the WPA2-Personal setting.
- RSN – Robust Security Networks are the central theme to the IEEE 802.11i standard and require CCMP or TKIP (with the Michael Algorithm) to be classified as such [CWNA05].
- TSN – Transition Security Networks are networks that might allow older WEP encryption in addition to those security schemes of the RSN. These networks are significantly weaker in terms of security than an RSN [CWNA05].

The IEEE 802.11i standard was developed by the IEEE with the intent of shoring up security vulnerabilities in both standard WEP and WPA. It was ratified in June 2004, and addresses major security concerns associated with encryption and authentication; however, it does not address all security concerns, such as the manipulation of management frames. [CWNA05]

2.5 Summary

Now that a thorough review of IEEE 802.11 is complete readers can better understand the procedures of experimentation, results, and final analysis. The most important parts of this chapter included the overview of encryption methods as well as the network performance metrics used in the experimentation process.

Chapter 3 Previous Work

3.0 Overview

As wireless networking has grown in the market place within the past few years there has been an increasing amount of research compiled on them. However, little examination into the impact of security on the performance of those networks has been completed, particularly with the new TKIP and CCMP encryption processes that are becoming the standard for enterprise wireless solutions. In addition to encryption, some of the studies below also examined the effects that various authentication methods have on throughput and transmission delays. These studies examined both a form of authentication and encryption simultaneously while analyzing throughput and response time performance.

Several of the techniques used in prior research are utilized for the experiments described in Chapter 5. Specifically this research analyzes various combinations of encryption and authentication in various physical environments to develop conclusions on the effects these circumstances have on network performance. The testing techniques developed and utilized in many of these past research experiments are critical to the success of this research as they have already addressed many of the issues associated with analyzing wireless network performance.

3.1 Previous Lessons Learned

Although there has been little work done on analyzing the impact of security protocols on performance there has been ample research conducted on general performance issues related to wireless networks. As mentioned above, previous research has provided several experimentation patterns to follow in the analysis of the security protocols. Issues such as IEEE 802.11g backwards compatibility, Transmission Control Protocol (TCP) verses User Datagram Protocol (UDP), saturation analysis, propagation concerns, and signal power have all been thoroughly examined. This research relies heavily on these previous findings to develop the most efficient testing methods. Below is a summary of some of these issues.

3.1.1 IEEE 802.11g Backwards Compatibility

The issue of backwards compatibility between IEEE 802.11g and 802.11b networks was examined by Wang in "Performance Evaluations for Hyrbrid IEEE 802.11b and 802.11g Wireless Networks". The authors discussed the issue of allowing backwards compatibility with IEEE 802.11b devices when operating an IEEE 802.11g network. Because the timing intervals within each MAC are different between the two protocols there is an overall reduction in throughput of roughly 15% for IEEE 802.11g devices where backwards compatibility with legacy IEEE 802.11b devices is allowed. This is important to understand for any throughput measurement because if backwards compatibility is enabled the results of those tests can only be compared to similar

network environments. This prevents maximum throughput from being achieved in IEEE 802.11g devices and also adds additional variables to any studies including both types of devices.

3.1.2 UDP vs. TCP Traffic

Another area of concern to experimentation techniques used on wireless LANs is UDP verses TCP traffic. UDP traffic closely represents voice and video data that may be transmitted over a network, whereas TCP traffic would more closely represent standard network traffic such as Hyper Text Transfer Protocol (HTTP). The difference between the two types of traffic is important in understanding experimental results. UDP will help determine maximum network throughput as the overhead and reliable deliver of TCP is negated; however, since most network traffic utilizes TCP, that throughput might not resemble "real-world" data.

3.1.3 Saturation Conditions

To fully measure the maximum throughput values for a wireless link, that link must be completely loaded with packet transmissions. This condition is usually referred to as maximum saturation throughput. In their article, "Performance measurements of the saturation throughput", Pelletta and Veloys addressed the concern of ensuring the packet generation software was throttled above the maximum predicted throughput for the wireless link. For IEEE 802.11g the theoretical maximum for TCP transmission is 24.4 Mbps and 30.5 Mbps for UDP [Atheros03]. A method to ensure that this condition is reached is to reduce the throughput for the wireless link, thereby allowing the traffic generator to more easily saturate the network.

3.1.4 Propagation Concerns

Because wireless LANs rely on the transmission of RF energy, problems that arise in an RF environment compound the problem of measuring throughput on a wireless link. Issues such as multipath, signal attenuation, scattering, and fading are difficult to diagnose and correct. Further complicating the problem is that these issues often change rapidly over time if a factor in the environment changes, such as the movement of a large piece of metal near the access point. To address these concerns Atheros recommends certain techniques during the analysis of wireless throughput.

Atheros [Atheros03] suggests first checking the airwaves for potential conflicts by scanning for other IEEE 802.11 signals or noise within the area. This can be done with free software available on the Internet and can help monitor the testing area over the course of the experimentation. This helps reduce the possibility of interference from noise generated by wireless phones, microwaves, and lighting. Additionally, Atheros recommends that whenever a trial is being completed it should be done in as little time as possible as environmental factors tend to remain constant over shorter time periods.

Pelletta and Veloys [Pelleta03] and Atheros [Atheros03] also suggest paying careful attention to access point location. The location for the access point should remain constant throughout all trials and be at least one to two wavelengths away from the transmitting and receiving station. This helps to ensure that any radio frequency (RF) propagation issues such as multipath are present at equal levels throughout all trials. Additionally, the authors recommend rotating access points as signal strength and ultimately throughput can be significantly impacted at different angles despite the presence of omni-directional antennas. During initial experimentation it would be prudent to examine these conditions to determine which placement options provide the most accurate as well as repeatable data.

3.2 Previous Security Analysis

There has been some work done in the area of analyzing the security overhead associated with wireless networks, but much of it remains focused on WEP rather than on new encryption techniques available today. In fact with regards to this specific area, only a few studies were located that could be used as a starting point for this research.

3.2.1 "An Experimental Study on Wireless Security Protocols over Mobile IP Networks"

Agarwal and Wong [Agarwal] examined the security overhead and authentication delays associated with the use of WEP, EAP, and the Internet Protocol Security (IPSec) on a WLAN. They analyzed the time delays necessary to authenticate over IEEE 802.1x with varying types of EAP such Message Digest 5 Algorithm (MD5) and Transport Layer Security (TLS), and the effect on throughput that various security types can cause. As expected, they determined that more secure levels required more packet transfers and ultimately more time to complete, with EAP-TLS needing roughly double the packets and time requirement than EAP-MD5. Although authentication occurs only once in a given session it is important to understand the initial overhead associated with the connection. In their studies surrounding encryption methods they found that WEP had much less performance overhead than implementing a virtual private networks (VPNs) with IPSec.

This study was perhaps the first analysis of security overhead published and helped discover a few important experimental techniques. First, the authors found that using larger chunks of data in the traffic generators was preferred because it helped cause the encryption mechanisms to reach a fully loaded state, thereby providing measurable data. They found that using small data amounts resulted in no visible differences between encrypted and unencrypted stream throughput. Secondly, the paper addressed how different encryption techniques could be more computationally intensive than others. In their paper they observed how the 3DES encryption in IPSec required more computation power then the RC4 algorithm in WEP. Its stands to reason that by utilizing different hardware and software that these encryption times could be reduced or increased depending on processing power so it is important to utilize similar hardware and software configurations in any experimentation.

3.2.2 "IEEE 802.11 Wireless LAN Security Performance Using Multiple Clients"

The second study that was used extensively to provide information for this research was Baghaei's thesis [Baghaei03]. Baghaei completed a series of experiments on a wireless network with single and multiple client stations, comparing various levels of encryption and authentication. He used IP Traffic to generate TCP and UDP packet streams to a server from various transmitting stations. Additionally, he employed Ethereal to monitor packet arrivals at the server and to help calculate latency and authentication times.

The author utilized eight levels of security to analyze overhead:
1. No security
2. MAC address authentication
3. WEP Authentication
4. WEP Authentication with 40-bit WEP encryption
5. WEP Authentication with 128-bit WEP encryption
6. EAP-TLS authentication
7. EAP-TLS authentication with 40-bit WEP encryption
8. EAP-TLS authentication with 128-bit WEP encryption

To ensure that the network was fully saturated the author chose a traffic bandwidth of 12 Mbps which was sufficiently large to saturate the IEEE 802.11b network. Four packet sizes were chosen for the experiments, 100, 500, 1000, and 1500 bytes to prevent fragmentation of packets during transmission. The author also completed experiments on an uncongested network with transmission rates lowered to 500 kbps.

The author's results showed staggering overhead associated with these security protocols. He found that in an uncongested network (traffic rates of 500 kbps) that the level 8 security definition resulted in an approximate 35% reduction in throughput for both UDP and TCP traffic. In this experiment there was a general downward trend of throughput from security levels 4 through 8, which seemed reasonable as more complex security mechanisms were put in place. However, when the author increased the traffic rate to 12 Mbps he found that the throughput was reduced by roughly 86% for TCP and 54% for UDP from security levels 4 through 8. Additionally, he found that response times for the network increased significantly when various types of EAP authentication were implemented.

When more clients were added to the network the author found that overall throughput decreased by roughly 50% with an additional client and 67% with two additional clients. Furthermore, when testing various packet sizes he found that throughput averages varied only slightly. Out of all the studies examined, this thesis presented the most pessimistic throughput values related to encryption.

3.2.3 "Performance Investigation of Secure 802.11 Wireless LANS: Raising the Security Bar to Which Level?"

Wong [Wong03] provided perhaps the most well prepared study that was found on the subject matter. In addition to using standard TCP and UDP traffic, he examined specific types of traffic such as HTTP and file transfer protocol (FTP). This author appeared to conduct similar work to Baghaei [Baghaei03] as they both completed work at the same institution in the same year. In addition to the security levels utilized above, Wong [Wong03] implemented ten VPN levels of security. This provided for a more in-depth look into overhead associated with layer 3 security mechanisms such as VPNs, but did not expand on previously conducted studies with WEP. The security levels that were utilized with respect to VPNs were:

1. No Security
2. PPTP tunneling with CHAP
3. IPSec tunneling with CHAP
4. Firewall with PPTP and CHAP
5. Firewall with IPSec and CHAP
6. Firewall with IPSec and EAP-TLS
7. IPSec with CHAP and DES
8. IPSec with EAP-TLS and DES
9. IPSec with CHAP and 3DES
10. IPSec with EAP-TLS and 3DES

The results of his study were fairly predictable as FTP traffic provided higher throughput than HTTP over both VPNs and standard security models. One would certainly expect HTTP, with the reply requirements, to be transmitted more slowly than the more constant streams of FTP. When the results of the throughput and response times were examined from the IEEE 802.1x security model, which included the eight levels in Section 3.2.2, a direct correlation with [Baghaei03] research was visible, in that response time increased and throughput decreased dramatically as the levels were cycled higher. Wong found a 73% drop in throughput for FTP traffic and a 62% drop for HTTP traffic at level 8.

With regards to the VPN model, Wong discovered some perplexing outcomes. He found that throughput levels increase between 17% and 30% when a firewall was present with the scenario. He also compared the IEEE 802.1x model side by side with the VPN model, which generally showed VPN security had a greater effect on throughput and response time than his IEEE 802.1x security levels. The author came to the following general conclusions.

- MAC and WEP authentication created no overhead.
- Various levels of authentication created different levels of overhead with respect to response times with EAP-TLS having the longest response time.
- WEP encryption impact varied and key length only affected response times.
- Tunneling with IPSec and PPTP generated large throughput overhead.

3.2.4 "Communication, Network, and Information Security"

A paper by Jamshaid [Jamshaid03] was the only one to incorporate an analysis of TKIP into its study. The research utilized NetIQ's Qcheck software to measure throughput of a WLAN with three security settings: open system authentication, 128 bit pre-shared key WEP encryption, and WPA with TKIP pre-shared key encryption. The results showed that by implementing WEP average throughput was reduced from 21.32 Mbps to 20.69 Mbps and with WPA implemented that average throughput was further reduced to 19.10 Mbps, representing a 3% and 11% decrease in total throughput, respectively.

3.2.5 Discussion of Previous Work

After carefully examining the studies described above, several conclusions can be drawn that could potentially affect this and future research. It was noticeable that the studies' results failed to correlate with one another with the exception of the Baghaei [Baghaie03] and Wong [Wong03] work, which can be attributed to the extreme similarities in their design and the fact that both of them studied together at the same institution and appeared to be using identical measurement techniques. It was also found that several intuitive results were challenged by their data. The significant amount of overhead that the last two studies associated to WEP seemed extraordinary as it would seriously hamper network capabilities. The first study by Agarwal and Wong [Agarwal] resulted in much more reasonable overhead. It seems feasible that much of this can be associated with the fact that the hardware they used was limited in processing capabilities and unable to rapidly encrypt and decrypt packets.

The authors' studies of authentication times were found to be relevant as the idea of fast roaming among access points becomes more and more an issue in current network topologies. Although the process of measuring and analyzing authentication times between APs is a daunting task it will likely provide key information for network engineers when considering issues such as hand-over times, which could be critical to applications such as VoIP.

Jamshaid's work [Jamshaid03] seemed to be conducted with equipment that could adequately handle the encryption algorithms and returned the most reasonable results. It intuitively seems that TKIP would hamper throughput more than WEP and even AES because of the intricate encapsulation process. The use of a simple throughput measurement tool ensures that there is little computational overhead in the measurement process, which must be followed in any future work.

3.3 Technology Advances

Perhaps the largest issue to overcome in this research, rapid technology advancement, also plays a part in developing a need for the project. Most of the work done in this area focused on the most basic security mechanisms employed inIEEE 802.11 networks. These included analysis of various WEP keys and MAC address authentication. The few studies that did incorporate more advanced encryption and authentication made little effort to correlate the two with respect to performance. In addition to drawing conclusions on observed data, this research aims to tie together previous work and examine similarities and differences among it.

3.4 Summary

Previous work in this area often led to inconclusive results that provided little usable data for network engineers. Additionally, the work focused on older encryption techniques that are unlikely to be used in current network distributions. However, previous work has laid an important foundation for experimental techniques that were utilized in this research and presented in the next chapter. By building on previous work this thesis was able to expand into new experimental areas.

Chapter 4 Experimental Background

4.0 Objectives

The goal of this research is to determine the overhead associated with IEEE 802.11 security protocols with respect to both authentication and encryption. This data is important as more security is required on wireless networks to ensure reliability and data integrity. Applications associated with the use of wireless networks are continually expanding, and they could be impacted by slow response times or reduced throughput.

Some examples of these applications are the roaming of VoIP users over multiple access points, remote wireless bridging, and transmissions with low power handheld devices. Users planning on roaming with VoIP phones between access points will experience problems similar to a roaming cellular telephone customer, which could be significantly impacted by long authentication times in addition to other hand-off concerns. Wireless bridges that connect campus buildings are currently capped at IEEE 802.11g speeds of 54 Mbps which is already much lower than typical gigabit wired solutions. The addition of encryption can only further hamper throughput on these links, independent of the fact that many users could be authenticating over the links as well.

Although not all of these issues are directly addressed in this research it should help develop the need for a thorough understanding of the effects that security could cause on various types of network performance. As such the research intends to provide general overviews of the current security protocols in use today and how they compare to one another with respect to response time, latency, and throughput.

4.1 Test Bed Overview

To conduct these experiments a test bed network was developed at the United States Coast Guard's Telecommunication and Information System Command (TISCOM) in Alexandria, VA. The service is currently funding a project to develop an enterprise wireless solution that could be utilized in areas where wireless technology would be feasible and cost effective for the organization. As such this project is part of the authors' professional as well as academic responsibilities. Although the goals of these two independent areas are slightly different, work on either of them helps contribute to the other.

The Coast Guard project team chose the Cisco wireless product line because of the remote management and monitoring capabilities of the suite. Cisco already owns a wireless product line and with the acquisition of Aironet they gained an additional one. The company is currently working to bring both product lines together into a single wireless solution. The team chose to utilize products from across their lines because they were interested in different aspects from all of them. This decision was driven by management capabilities, security concerns, and scalability. Fortunately these products

have proven to work very well together and provided a solid platform for security analysis.

In addition to the network deployed at TISCOM, the team is currently working with the United States Coast Guard Academy to develop a similar network that can be remotely managed and monitored from Alexandria.

4.1.1 Hardware Selection

As mentioned above the hardware chosen is sold by Cisco Systems. The primary technology being utilized is the Lightweight Wireless Access Point Protocol (LWAPP) which involves the use of wireless controllers to manage access points and push security protocols to them. This allows network engineers to manage multiple access points from a single device and ensure that access points are essentially dummy devices with no security configurations stored within them, providing the organization with an additional layer of physical security should one of the access points be lost or stolen.

Another compounding problem with the hardware selection is that there are currently few product lines available that are fully 802.11i compliant and offer CCMP encryption options. This certainly eliminated most small office/home office (SOHO) equipment.

The hardware devices currently in use are:

Cisco 2000 Series Wireless Controller
(http://www.cisco.com/en/US/products/ps6308/index.html)

Cisco 1000 Series Access Point
(http://www.cisco.com/en/US/products/ps6306/index.html)

Cisco 1200 Series Access Point
(http://www.cisco.com/en/US/products/hw/wireless/ps430/index.html)

Cisco 1300 Series Outdoor Bridge
(http://www.cisco.com/en/US/products/ps5861/index.html)

Cisco Wireless LAN Client Adapters
(http://www.cisco.com/en/US/products/hw/wireless/ps4555/index.html)

Cisco 3500 XL Series Switch
(http://www.cisco.com/en/US/products/hw/switches/ps637/index.html)

Only three of these devices will be used in the vast majority of security testing utilizing network configurations shown in Figure 4-1.

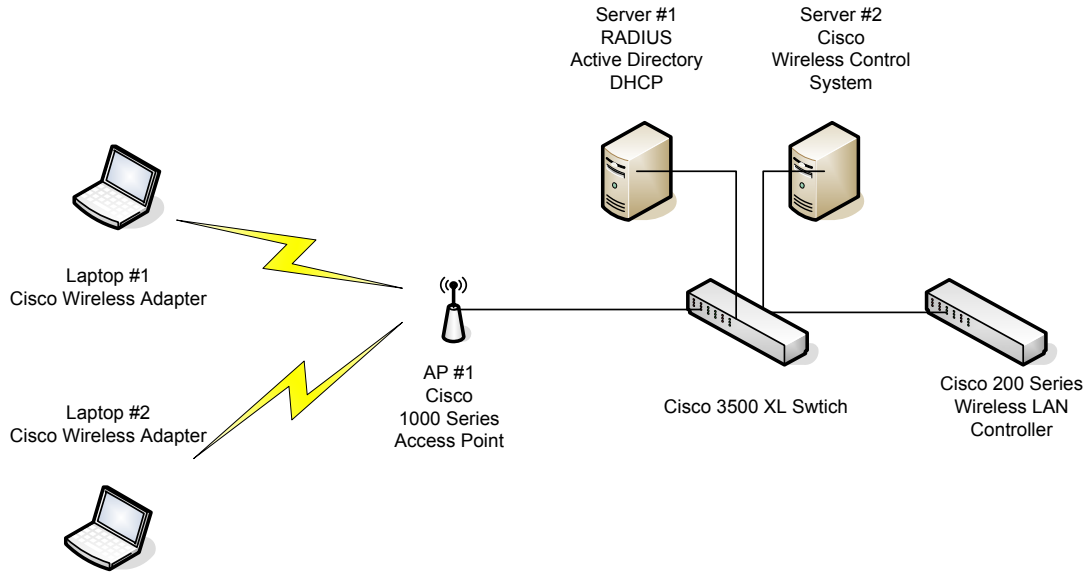

Figure 4-1. Test Bed Configuration

As shown in the diagram, the test bed is comprised of a central Cisco 3500 XL switch that makes up the Ethernet backbone. The two servers, the WLAN controller, and the access point are connected directly to the switch. The two laptops have physically constant locations during the testing process to minimize RF propagation concerns. Because this network is also part of a developing enterprise wireless solution for the Coast Guard, the switch is typically connected to a host of other devices when it is not being used for this research.

The servers and laptops being used have the hardware specifications indicated in Table 4-1.

Table 4-1. Server/Client Specifications

Machine	Laptop #1	Laptop #2	Server #1	Server #2
Brand/Model	HP Pavilion 7000	Dell Inspiron 5150	Dell Client Pro	Dell Power Edge Server
Processor	Pent M 1.6 GHz	P4 3.2 GHz	P4 2.4 GHz	Dual P3 1.26 GHz
Memory	1 GB	1 GB	512K	1 GB
Network	Cisco A/B/G Client Adapter	Cisco A/B/G Client Adapter	Integrated	Integrated

The server hardware is not top-of-the-line equipment, but should provide more than sufficient processing power to ensure there is no throughput reduction due to CPU overload. Additionally, the Cisco WLAN controller and the AP are enterprise level hardware which should provide sufficient processing power for the encryption schemes

in use. Hopefully this should help ensure prompt response times and low encryption overhead that may not be possible with less expensive devices.

4.1.2 Software Selection

There were several software issues that needed to be addressed in deciding which operating systems and throughput testing software to use during the experimentation processes. Most of these problems surrounded the fact that the newer IEEE 802.11i standard was supported in few hardware and operating system combinations. However, after careful examination it does not appear that any of the experiments could have been completed more thoroughly in different software environments, as all of the testing was completed with minimal discrepancies.

All of the experimentation is conducted within a Microsoft Windows operating system for a variety of reasons. Most importantly the Cisco a/b/g PCMCIA cards that are being used with the laptop computers do not have Linux-compatible drivers that can carry out CCMP encryption. This limits several powerful Linux networking tools freely available; however, this was not expected to affect measurement outcomes. Each server is loaded with Windows 2003 Server Edition with Service Pack 1 and both laptops are loaded with Windows XP Professional with Service Pack 2. Furthermore, instead of utilizing the Windows wireless management tools the Coast Guard team selected the Odyssey Funk Client because it holds a Federal Information Processing Standard (FIPS) compliance credential necessary for government use. This software provides an authentication solution for multiple EAP protocols and replaces the Windows wireless management software.

For throughput testing software a variety of choices were tested to determine which one was most convenient and delivered the best results for the network configuration. It was originally planned to utilize AirMagnet to complete the majority of testing, but it was discovered that the software installed independent drivers for the PCMCIA cards that were not capable of CCMP encryption. This would have caused a significant problem because a large portion of the research focused on that particular encryption scheme. Testing was initially started with Iperf; however, it appeared that UDP throughput was not being measured correctly. Because of this Qcheck, free software recently acquired by Ixia, was chosen for use during testing. This program can measure both throughput and response time quickly and efficiently. Also, because the program has little computational overhead it can be assumed that none of the throughput or response time results are negatively affected by poor CPU performance.

To measure authentication times the Ethereal protocol analyzer was chosen to record the amount of time before an authentication request begins and when access is granted. This tool also helps to verify that there is no network traffic present that could negatively impact experimental results.

A thorough discussion of these measurement tools is presented in this chapter.

4.1.3 Physical Location

The test bed's physical location at TISCOM is in the Northeast Lab building. This is a single story brick building with offices surrounding a central lab deck. Figure 4-2 is a drawing that is not to scale, but provides an adequate representation of the floor area.

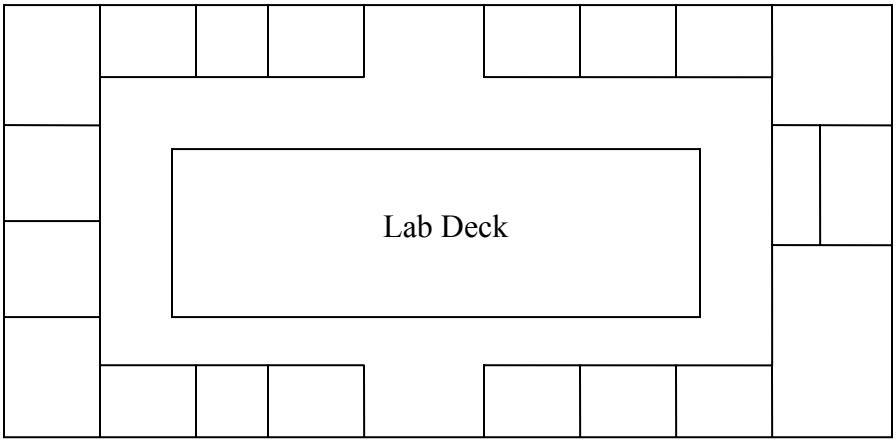

Figure 4-2. TISCOM Northeast Lab Floor Area

The outer walls of the building are composed of brick with an internal layer of dry wall. All of the other walls in the building are 4" thick with dry wall and insulation. The hardware used for the network is located close to the north lab deck wall but the access point in use is almost at the exact center of the building. The outer brick walls prevent interference from outside the building and also help prevent wireless signals from emanating outside of the building.

4.2 Network Performance Metrics

As mentioned above the research focuses on various metrics related to network performance. Because these performance metrics will comprise the bulk of the data presented in the research, it is important to understand them and how they can be measured accurately. Specifically describes the tools and methods that can be applied to ensure that the data being recorded is representative of the system.

4.2.1 Metric Descriptions

There are a large number of performance metrics that network engineers utilize to help analyze network configurations and troubleshoot problems. The ones most commonly referred to are throughput and latency, but there are more depending on the media that the network exists within. The following is a brief overview of some performance metrics commonly used today:

- Throughput – "A measurement of the data-transfer rate through a complex communications or networking scheme." [Dyson99] It is important to note that

throughput is a measurement of a certain amount of data through a link over a given time and can be affected by processor and disk performance, operating system capabilities, network hardware limitations, and the amount of data being transmitted.

- Latency – "The time delay involved in moving data traffic through a network". [Dyson99] The three sources of latency are propagation delay, which is caused by the time necessary for data to travel the length of the link; transmission delay, which is the actual time necessary for data to be moved across the network; and processing delay, which is the time needed for data encapsulation and route establishment.
- Response Time – "The time lag between sending a request and receiving the data." [Dyson99] This could be the time required to receive a response from an authentication server after a request has been past from the supplicant in an IEEE 802.1x exchange.
- Bandwidth – With regards to networking devices, bandwidth refers to the transmission capacity of a communications channel. [Dyson99] This is slightly different then throughput as bandwidth is most often a theoretical maximum where as measured throughput tends to show "real world" results.

These are just a few of the many terms used to analyze the performance of networks, but will be the only ones referred to in this research.

4.2.2 Measurement Software

In the research it was necessary to measure throughput and response times to compare the various security levels applied to the wireless test bed. As such Iperf, Qcheck, and Ethereal were used to help in this process. Each tool has its own capabilities to measure variables such as how throughput was affected by different encryption levels and the length of time required for authentication with different protocols. As mentioned above Iperf was only initially used until it was determined that Qcheck provided more accurate results; however, a brief discussion of both programs isprovided below.

4.2.2.1 Iperf

Iperf is a simple yet powerful tool for measuring throughput with both TCP and UDP traffic. Iperf allows the user to manipulate various aspects of the data that is being transmitted to most closely monitor the type of data usually present on a network. The most commonly adjusted settings within Iperf are the window size and the bandwidth associated with UDP connections. The window size that should be used is determined by the bandwidth-delay product. This value is calculated by multiplying the maximum throughput of the medium by the round-trip time of the client to the server as determined by ping. The Iperf documentation specifies an example with a bottleneck bandwidth of 45 Mbps and a round trip time of 42 ms. The result is then 45,000,000 * 0.043 = 1,890,000 bits or 236 Kbytes. The software documentation recommends setting this

value as the minimum window size to use when measuring the throughput of the link, but suggests using different values to compare results.[1]

When using Iperf it must be run on both a server and the client. Once the options are set on the server, the client can initiate a transfer for a set time interval. Upon completion both the server and client report their measurements in the client and server window.

4.2.2.2 Qcheck

Qcheck is another throughput measurement tool that is freely available on the web through Ixia. Fortunately Qcheck provides all the functionality required for this research and runs with minimal CPU overhead. This software provides the ability to measure both throughput and response time of TCP as well as UDP traffic. The program runs similarly to Iperf in that it utilizes a client/sever relationship.[2]

The only drawbacks to Qcheck are that there are not as many manipulations available with the data as there are with Iperf. Qcheck limits the amount of data to be sent from 1 to 1000 Kbytes which can often occur rather quickly over a 54-Mbps link leaving the transfer vulnerable to unforeseen shifts in the RF environment. However, this can be overcome by conducting multiple trials and computing averages, as well as by reducing the maximum bandwidth of the link.

4.2.2.3 Ethereal

Ethereal is a freely distributed packet capturing application that allows users to view all packet transfers in or out of an Ethernet interface or similar network. The captures can be filtered so that only relevant packet transfers are displayed. This is important for the authentication portion of the research as it allows interaction from the supplicant and the authentication server to be filtered allowing precise measurements of authentication times.[3]

4.3 Experimental Concerns

Before the actual experimental design is discussed, it is necessary to address potential problems or concerns that may arise during the experimentation process. It is important that these problems be mitigated by proper experimental techniques in order to present data that is as close to the actual system being tested as possible. In any experimental environment there will be variables present that can affect the outcomes and this is particularly true in a communications network. Because these experiments must mesh the inherent problems associated with studying both wireless communications and

[1] Iperf documentation is located at http://dast.nlanr.net/Projects/Iperf.
[2] Qcheck documentation is located at http://www.ixiacom.com/products/performance_applications/pa_display.php?skey=pa_q_check.
[3] Ethereal documentation is located at http://www.ethereal.com.

computer networking environments there will be a host of these variables. These concerns can be isolated to the RF and networking environment.

4.3.1 RF Environment

The RF environment in which the testing occurs is likely the largest variable associated with the experimentation. As previously mentioned, the 2.4-GHz wireless signals that are used are susceptible to reflections, refraction, diffraction, scattering, and absorption in various types of physical environments. Understanding these basic RF principles is essential when conducting testing in a varying RF environment.

- Reflection – Reflection occurs when a traveling electromagnetic wave strikes an object with physical dimensions larger than the wavelength of the signal and causes the wave to reflect towards the receiver at a different angle from the direct wave. This is how multipath interference is formed and can have a significant impact on wireless signals. Metal surfaces such as window blinds tend to have a serious reflective capacity because of there elemental makeup and physical design [CWNA05].
- Refraction – Refraction occurs when a wave travels from one medium to another with a different density and causes the wave to bend in a different direction once it passes into the new medium [CWNA05].
- Diffraction – Diffraction occurs when traveling waves bend around an object and collectively move in a different direction. Imagine water ripples bending around a fixed object in a tidal pool from a central location where a stone is dropped in the water [CWNA05].
- Scattering – Scattering occurs when a large wavelength signal hits an obstruction with smaller sections and causes the wave to scatter into multiple waves traveling in various directions. This type of behavior tends to occur outdoors around foliage [CWNA05].
- Absorption – Absorption occurs when a wave hits a material that is porous or has multiple sections. 2.4-GHz waves tend to get absorbed concrete or brick making it difficult to pass IEEE 802.11b/g signals through these materials [CWNA05].

Although most of these conditions cannot be easily observed or corrected there are ways to ensure that they are minimized in the RF environment. Unfortunately, the only measurement tool available to determine how the above conditions are affecting the network is the signal strength readouts on the clients and access points. For example, series multipath interference will significantly reduce the Receive Signal Strength Indicator (RSSI) which measures the relative signal strength of the access point or the client device.

Unfortunately, the lab deck at TISCOM is not the most wireless friendly environment as there are multiple objects within the room causing various types of interference. The lab deck has several server racks filled with equipment that are a likely source of multipath and potential electromagnetic interference. Figure 4-3 and 4-4 are from a report

generated by the Cisco Wireless Controller on noise and interference by IEEE 802.11 channels in the lab deck where the access point and notebooks are located.

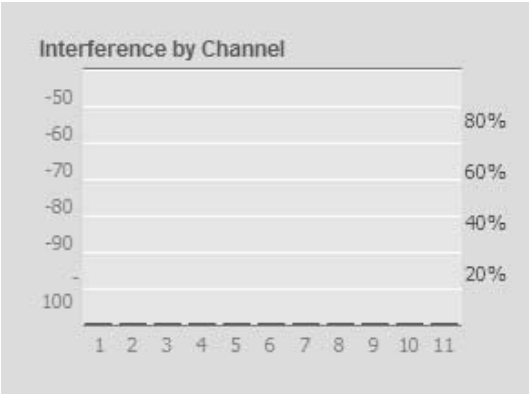

Figure 4-3. Lab Interference

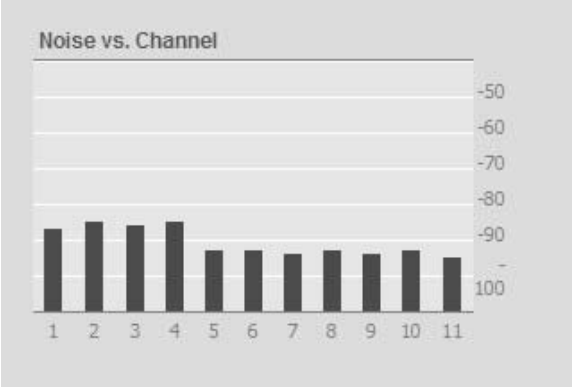

Figure 4-4. Noise in Lab Environment

As shown there is some noise in the 2.4-GHz spectrum which encompasses the wireless network with channel 11 having the least. Because of these results the tests were conducted on this channel.

This noise can be generated from a variety of sources that tend to bleed over into the 802.11 spectrum such as other access points, wireless bridges, mobile phones, and microwave ovens. There are none of these devices operating within interference range of the test bed network, but it is important to ensure that pattern continues throughout the testing.

Perhaps the most important aspect of conducting throughput and response time testing under these RF conditions is to ensure the tests are completed over as small a time interval as possible to prevent changes in the environment, and to ensure that each test is conducted against a measured maximum, which in the throughput tests involves comparison testing with no encryption and authentication. Although this may not

provide for theoretical maximums of IEEE 802.11b/g signals, it shows the results to a measured standard which could then be extrapolated to the theoretical maximums in Table 4-2 [Atheros03].

Table 4-2. Theoretical Maximums for 802.11

	# of Channels	Modulation Scheme	Max. Link Rate	Max. TCP Rate	Max. UDP Rate
802.11b	**3**	**CCK**	**11 Mbps**	**5.9 Mbps**	**71. Mbps**
802.11b/g	**3**	**OFDM/CCK**	**54 Mbps**	**14.4 Mbps**	**19.5 Mbps**
802.11g	**3**	**OFDM/CCK**	**54 Mbps**	**24.4 Mbps**	**30.5 Mbps**
802.11a	**19**	**OFDM**	**54 Mbps**	**24.4 Mbps**	**30.5 Mbps**

4.3.1.1 2.4-GHz vs. 5.0-GHz Bands

An often overlooked area in IEEE 802.11 studies is the IEEE 802.11a standard, which as mentioned previously, operates in the 5.0-GHz spectrum. Many enterprise-level wireless distributions rely on IEEE 802.11a as a backhaul option for the more standard IEEE 802.11b/g networks. Because of this, the differences between the two spectrums should be well understood by wireless engineers. Likewise those engineers should understand the advantages and disadvantages of both spectrum bands. Although IEEE 802.11a is essentially ignored in this research because it is usually only used for backhaul a brief listing of these differences from [Atheros03] is presented below.

- Antennas tend to become more directional at higher frequencies allowing better use of equipment at the 5-GHz range.
- Absorption increases at higher frequencies allowing 2.4 GHz signals to penetrate objects more effectively.
- Nearby electronics, Bluetooth devices, other IEEE 802.11 b/g devices, cordless phones, and microwaves have a large affect on 2.4-GHz signals but very little on 5-GHz signals.
- Attenuation in antenna cable runs increases with frequency allowing 2.4-GHz devices to reside on longer connections with less signal loss over the same length of cable as compared to 5-GHz devices. This is important when external antennas are utilized.

4.3.2 Networking Environment

In addition to the RF environment introducing complexities to the experimentation process so does the networking environment. Things such as system performance, link bandwidth settings, software settings, and various types of network traffic can affect the outcome of the experiments. However, just as RF conditions can be mitigated so too can the network conditions assuming there is a proper understanding of those variations likely to occur.

4.3.2.1 Bandwidth Setting

During preliminary testing it was determined that setting the maximum link bandwidth to 2 Mbps would help provide optimal results. This is because greater throughput speeds on the link resulted in larger variation of the measured throughput. It was difficult to obtain repeatable throughput numbers at higher link speeds. By reducing the throughput levels the test results became much more repeatable. This is probably due to the fact that at 2 Mbps the link was able to handle traffic with little computation strain on the system, which in turn would provide for the most easily measured encryption overhead.

4.3.2.2 System Performance

The performance of various hardware devices in the network can significantly degrade network throughput. For example, systems running at slower speeds are unable to process the software encryption as quickly as high-speed access points or other systems that may be located in the network. Slow system performance could be caused by traffic generation and monitoring software and lead to the aforementioned results. The best method for overcoming these problems is to use fast system components which can process traffic and software most efficiently. Unfortunately this is often cost prohibitive as was the case in these trials. The option implemented here was to chose software that required little computational overhead and throttle link bandwidth to reduce load.

4.4 Summary

The test bed chosen for this research was designed with the experimental procedures in mind. The hardware and software selections complemented each other during the testing processes and provided both predictable and repeatable data. The experimental RF and network concerns were developed from previous work and carried over into the processes of Chapter 5. Some concessions where made during the length of the research, but these had little impact on the final results.

Chapter 5 Measurement Campaign

5.0 Measurement Overview

The experiments were conducted on the test bed to determine several metrics when different security protocols and network configurations were applied. These metrics included end-to-end throughput, response time, and client authentication time. Several trials were conducted at various security levels to most accurately represent the system under different network configurations. These network configurations included data transfers between a single client and a wired server and between two wireless clients within various physical locations.

5.1 Overview of Security

To present the most accurate and concise information about various forms of authentication and encryption the experiments were conducted with a minimum number of security choices, and are outlined in Section 5.3. This limitation allowed the experiments to focus on those primary security mechanisms that are most likely to be implemented in an enterprise-level environment, while providing sufficient information for small office and home office users to benefit from as well.

5.1.1 Encryption Overhead

Encryption protecting a wireless link is the most persistent overhead associated with a network secure configuration. Once a wireless client authenticates with the RADIUS server and keys are exchanged the encryption remains present for the entire session affecting all data transmitted over the link. Encryption has little effect on latency associated with the link but does have an impact on throughput in a saturated network condition, where bandwidth is allocated for larger encrypted packets.

To measure this overhead, the three types of encryption were used on the test bed in the prescribed physical network configurations of Section 5.2. Throughput was measured with Qcheck during each experiment, with a minimum of ten trials per experiment, in which 1000 Kbytes of data was passed from the client to the server. Throughput values obtained with encryption were compared to throughput with no encryption present to determine the percent overhead that was associated with each encryption scheme.

To provide the most accurate representation of "real world" data the mean overhead percentages of both TCP and UDP traffic are averaged for each physical configuration.

5.1.2 Authentication Overhead

Authentication overhead was measured by noting the differences between when the first EAP message was sent from the client and when the final accept message was sent from the RADIUS server. This was observed on the RADIUS server with the help of Ethereal. This information provides general information about the time necessary for clients to authenticate with a RADIUS server using different authentication methods.

5.2 Physical Configurations

The basic network configuration presented in Chapter 4 was the primary layout for testing throughput, response, and authentication times. However, to examine how various physical obstacles or different network configurations affect overhead the layout was adjusted. To maintain consistency of data reporting there was a number associated to each physical network configuration that was used to present the data. Below is a list of those configurations and their associated number.

- Configuration 1 – The test bed was located entirely within the lab deck and only one laptop was used for testing, as seen in Figure 5-1. Server #2 acted as the server and Laptop #2 acted as the client. This was the environment used to test maximum throughput for the link because adding additional stations takes an obvious toll on throughput.

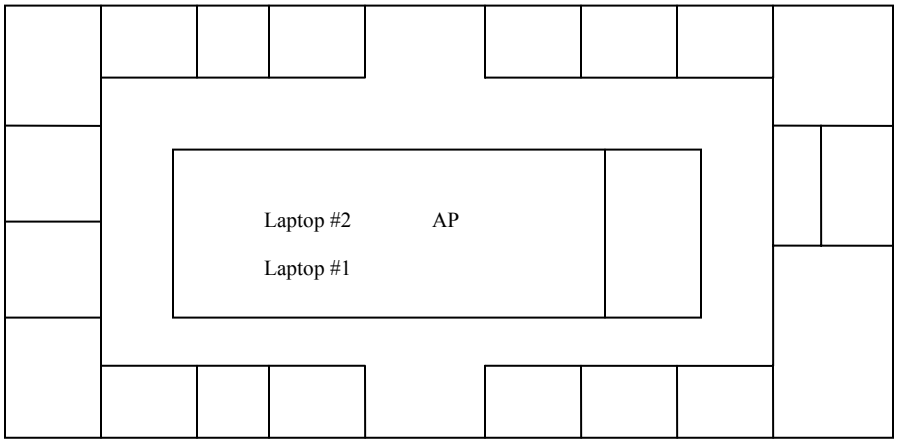

Figure 5-1. Configuration 1

- Configuration 2 – The test bed was located entirely within the lab deck but two laptops are utilized, as seen in Figure 5-1. Laptop #1 acts as the server and laptop #2 acts as the client in all scenarios.

- Configuration 3 – The test bed was located within the lab deck and only one laptop was used for testing. The laptop was moved behind one 4" wall and tested similarly to Configuration 1. Server #2 acted as the server and Laptop #2 acted as the client.

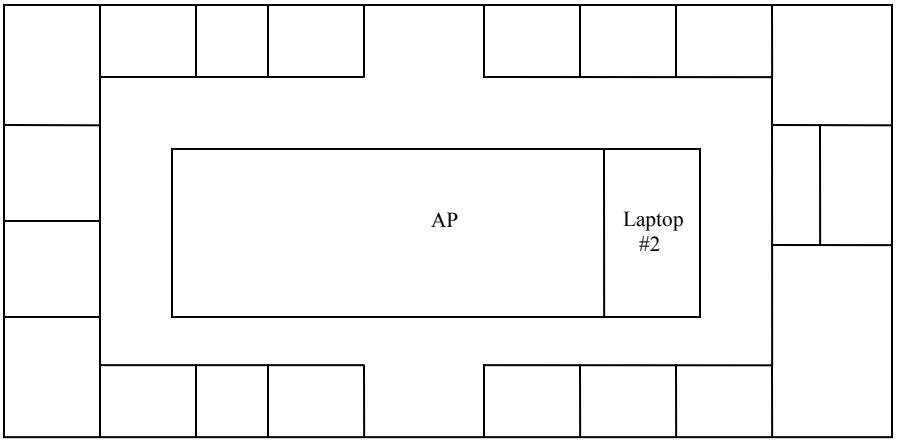

Figure 5-2. Configuration 3

- Configuration 4 – The test bed was within the lab deck and Laptop #2 remained in a similar location to Configuration 3. Laptop #1 was moved to the opposite side of the building still within range of the access point, but out of range of Laptop #2, and behind two 4" walls.

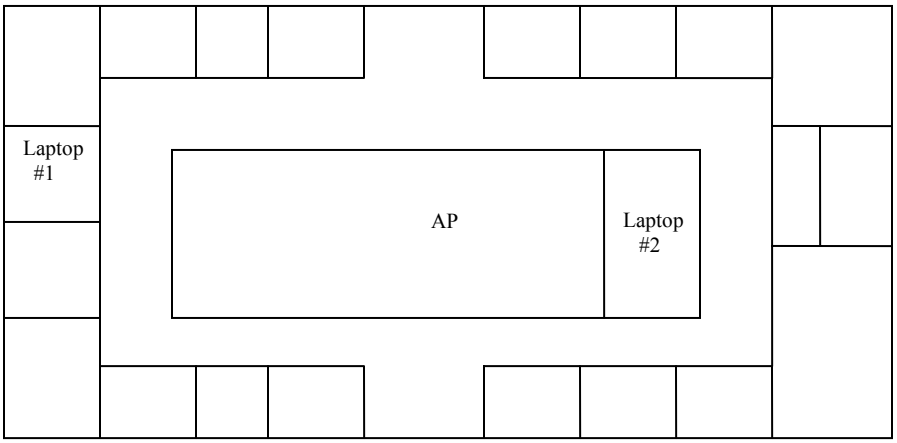

Figure 5-3. Configuration 4

Whenever data iss reported in Chapter 6 it refers to the network and security configurations used during the testing. This should help prevent confusion when the data is presente..

5.3 Security Configurations

In addition to the physical network configurations used to present data there are various security configurations utilized for the testing of response time and throughput values. Although there are a number of combinations that could be chosen, this research focuses on those most likely to be present in a corporate network environment. Because of this, the only authentication protocols that are examined are LEAP and PEAP. However, all

levels of encryption, WEP, TKIP, and CCMP, are included at some point in the trials. Below is an overview of the security combinations selected for study.

- Security Level 1 – This entails open association with no encryption on the data flow. This was the base line security scheme used as the starting point for all data comparisons with encryption and authentication.
- Security Level 2 - Open association with a 64-bit WEP key assigned for encryption. WEP utilized a pre-shared key for ease of use between the stations and should function similarly to a scenario where keys are distributed at the authentication process.
- Security Level 3 – Open association with a 128-bit WEP key for encryption. This is similar to SL2 with the exception of the longer key length.
- Security Level 4 – Authentication was completed with PEAP and encryption was handled by TKIP. Keys were generated by the handshake procedure each time the client reassociated with the access point.
- Security Level 5 – Authentication was completed with PEAP and encryption was handled by CCMP. Keys were generated by the handshake procedure each time the client reassociated with the access point.
- Security Level 6 – LEAP authentication and CCMP encryption. This was used for authentication times in coordination with SL5.

In addition to the widespread use of these security settings, they also provide ease of use in the measurement campaign. The combination of physical and security settings provided a broad look into overhead associated with encryption and authentication, and allows one to draw accurate conclusions on the effects that encryption and authentication have on network performance.

5.4 Measurement Procedures

To understand the testing process, a detailed explanation is provided of the procedures used when the actual measurements were completed. Since each physical network and security configuration changed slightly there were changes in the procedures used to complete testing in those settings.

5.4.1 Preliminary Testing

Chapter 3 detailed the methods and results of previous security measurements and was the starting point for choosing software, network layouts, and testing procedures. In addition to research already available, it was necessary to run tests similar to those conducted in the actual experimentation to determine which methods to use for each trial. Issues such as cycling security settings, the number of trials to run, access point placement, software limitations, errant values, and anticipated results were completed at this time.

5.4.1.1 Cycling Security Settings

It was necessary to develop methods ensuring security settings on the network are cycled smoothly to rotate through the various levels of security that the experiments were intended to cover. Because the Cisco products are easy to manage from a web-based interface anywhere on the network, it did not appear to pose any barriers to the testing. In fact, all settings for changing encryption schemes and turning authentication on and off were located in the same place and were adjusted in a matter of seconds. This ensured that trials were run efficiently over a small amount of time.

5.4.1.2 Number of Trials

An important part of any scientific experiment is determining the number of trials to utilize. This must balance time feasibility and ensure data accurately represents the system. Initial tests were completed with Iperf and Qcheck to determine system behaviors to that respect. It was discovered that consistent data was obtained for throughput and response times from both trials with Iperf and Qcheck, as well as strong correlation between the results of the different software programs. Ultimately, ten trials was the number chosen to be completed in each physical and security configuration to obtain suitable means for the final report. When the results are presented in Chapter 6 these averages are the focus of discussion for each physical and security configuration.

5.4.1.3 Errant Data values

Another issue to tackle was handling data points that may be vastly different from the mean values of the data. These data points occur for a variety of reasons including software “hiccups”, random changes in the RF environment, and network collisions. During preliminary trials it was found that roughly one out of every twenty trials had throughput measurements anywhere from 20 to 50 percent lower than the mean. These errant data points would significantly impact overall overhead calculations and needed to be dealt with. Fortunately the mean throughput values initially measured were consistent among all trials with the exception of these few errant outlier values. To combat the problem all trials that occurred with throughput values of more than five standard deviations below the mean were excluded. Since the standard deviation and mean were calculated in a spreadsheet after each trial it was easy to exclude a value and rerun the trial.

For example, in one of the first experiments conducted measuring throughput with various types of encryption, ten trials were run with an average throughput value of 1559.4 Kbps. The standard deviation was calculated as 26.10 Kbps, so if an eleventh trial were run any value less than approximately 1428.9 Kbps would be rerun. Fortunately, there were relatively low standard deviations throughout the entire experimentation process, which helped easily identify outliers.

5.4.1.4 Access Point Placement

Preliminary trials determined an ideal position for the access point with respect to the initially obtained throughput values. Despite that the access points had omni-directional antennas there was as much as a thirty percent reduction in throughput by rotating the access point 90 degrees and leaving the client in the same location. Once maximum throughput was reached by rotating the access point at several different angles the orientation was fixed for the entire experimentation process.

The Cisco 1000 Series access point utilized for the testing comes equipped with an internal antenna. As previously stated, obvious throughput reductions were present by rotating the device at short ranges. However, this effect was greatly reduced as the client moved further from the access point. This helped ensure tests where the clients moved around the building during the trials, were not dramatically influenced by this short range throughput anomaly.

5.4.1.5 Software Limitations

Preliminary trials also identified some software limitations that could compromise accurate reporting of information. The most noticeable was the Funk Client software's ability to report signal strength. The software's lag time when reporting signal strength could not always be relied on as accurate. To address the problem each laptop computer employed Network Stumbler to provide a real-time signal strength report on the access point.

5.4.2 Methods for Measuring Encryption Overhead

To measure the effect encryption has on wireless performance several steps were taken during each trial to ensure accuracy.

- Before tests were completed, the laptops taking part in the trials were authenticated on the network with the Funk software and manual user intervention. This ensured the software was not conducting machine authentication without the user and avoided the presence of extra clients on the network.
- Each time a security setting was changed on the network and a client was adjusted to meet that security setting the "Reconnect" button was used in the Funk software to ensure the wireless device went through the entire process of reconnecting to the network and authenticating.
- Each time a pre-shared WEP key was employed the same key was used. These were a series of '9's ranging from 5 to 16 characters depending on the size of the WEP key used. The main purpose of this was for ease of use when changing keys, but it also ensured identical streams were produced from the RC4 algorithm.

- Before trials were run an initial test was completed by Qcheck ensuring the link was established and stable. The initial test eliminated almost all errant data points discussed above.

5.4.3 Methods for Measuring Authentication Overhead

The authentication overhead is essentially being measured as the time it takes a client device to authenticate to the RADIUS server once association occurs. As previously mentioned, this time was measured within Ethereal from the moment that the first EAP packet was sent to the moment the access packet was sent from the RADIUS server. The following techniques were utilized to ensure accuracy.

- It was determined that the best way to ensure the client was disconnected and reconnected during each trial was to remove the PCMCIA card from the laptop and reinsert it. This forced the devices to reconnect to the wireless network and complete the process of authentication.
- A single capture file is utilized in Ethereal to examine all ten trials to allow the tests to be completed in a short time frame. During each test a minimum of twelve trials were conducted to ensure that if any of the first ten are errant there is additional data.
- Both LEAP and PEAP are enabled as acceptable authentication methods in the RADIUS server to ensure only client-side settings need to be adjusted between configuration changes.

5.5 Summary

The measurement campaign was designed to obtain the broadest look at encryption and authentication overhead. The physical and security configurations developed for the experiments were drawn from previous work as well as the desire to obtain more relevant data on the subject. The measurement procedures were necessary to ensure the obtained data closely represented the system, despite several inconsistent environments. The results detailed in the next chapter are an example of how good experimental design can yield positive results.

Chapter 6 Results

6.0 Overview of Results

This chapter is divided into two main sections covering encryption and authentication experiments, with the encryption section consuming the bulk of the chapter. Each of these main sections has sub-sections covering the physical network configurations outlined in Chapter 5. In general each encryption result is averaged and overhead percentages are presented in a graph. These values are those most important when discussing the results in Chapter 7.

6.1 Encryption Overhead Results

A total of four experiments were conducted corresponding to each of the physical network configurations utilized. Each experiment had trials covering five different security levels for a total of fifty trials per experiment. The results of each configuration are presented below.

6.1 Configuration 1

Configuration 1 was the baseline configuration used to determine further testing requirements with different security configurations. Configuration 1 was the only testing carried out with both Iperf and Qcheck results before it was decided all future trials would be completed with Qcheck. The data between the two programs correlated nicely with the exception of the results of UDP traffic. This configuration was tested three times to ensure the methods outlined in Chapter 5 help up during the testing process. It is the only configuration in which multiple tests were completed.

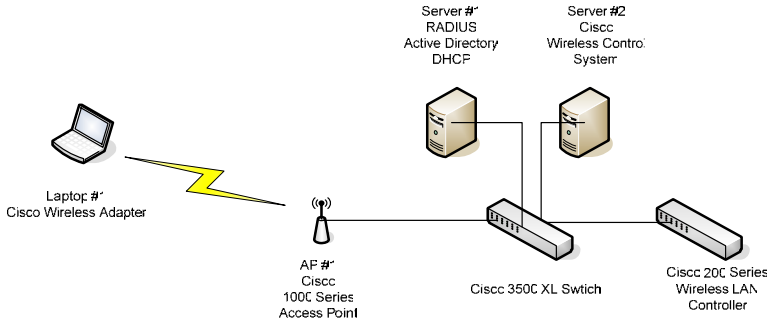

Figure 6-1. Configuration 1

This test was completed twice with Iperf and once with Qcheck. When Iperf was utilized the TCP window size was set to 128 Kbytes and the UDP bandwidth was set to 1 Mbps.

The Iperf default setting of ten seconds was used. When Qcheck was used 1000 Kbytes of data was transmitted for each trial.

6.1.1.1 TCP Results

Table 6-1 and Figure 6-2 present the numerical and graphical results for throughput from the first TCP test using Iperf.

Table 6-1. Results for Throughput Configuration 1 / TCP - Test 1

Test 1 (Iperf): TCP, 128K Window, -43 dBm

Trial	CCMP/PEAP - kbps	TKIP/PEAP - kbps	WEP(128 bits) - kbps	WEP(40 bits) - kbps	Open - kbps
1	1553	1502	1435	1547	1575
2	1546	1496	1572	1544	1590
3	1527	1511	1555	1577	1565
4	1551	1469	1575	1575	1535
5	1539	1545	1581	1573	1586
6	1553	1551	1544	1570	1563
7	1565	1537	1570	1541	1562
8	1570	1562	1557	1559	1580
9	1553	1540	1584	1560	1574
10	1555	1529	1583	1547	1573
Avg.	**1551.2**	**1524.2**	**1555.6**	**1559.3**	**1570.3**
Std. Dev.	**12.15**	**28.92**	**44.39**	**13.88**	**15.52**

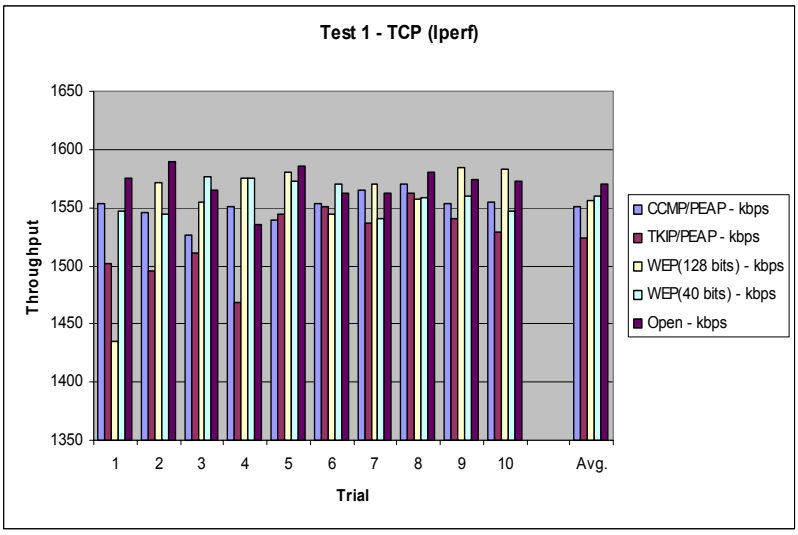

Figure 6-2. Results for Throughput Configuration 1 / TCP - Test 1

These results matched with preliminary trials and seemed correct on an intuitive level. The encryption effects are clearly represented by the averages column, with TKIP leading to the largest decrease in throughput. The open configuration generated an average throughput of 1570.3 Kbps over the 2-Mbps link. When a 40-bit WEP key was

implemented the average throughput dropped 11 Kbps or 0.7%. A 128-bit WEP key reduced throughput by 14.7 Kbps or 0.9%. TKIP degraded the link by 42.8 Kbps or 2.7%. CCMP was similar to a 128-bit WEP key with a 19.1 Kbps or a 1.21% reduction.

As previously mentioned, these results correlate to the TCP throughput results from the second test, again using Iperf, and the third test using Qcheck, which helped confirm the accuracy of the software. Results from the second and third tests are presented in Table 6-2, Figure 6-3, Table 6-3, and Figure 6-4 in numerical and graphical form.

Table 6-2. Results for Throughput Configuration 1 / TCP - Test 2

Test 2 (Iperf): TCP, 128K Window, -44 dBm

Trial	CCMP/PEAP - kbps	TKIP/PEAP - kbps	WEP(128 bits) - kbps	WEP(40 bits) - kbps	Open - kbps
1	1561	1527	1575	1530	1580
2	1552	1490	1572	1574	1592
3	1567	1550	1573	1572	1590
4	1563	1538	1558	1585	1581
5	1550	1544	1573	1562	1577
6	1561	1572	1567	1561	1595
7	1555	1559	1549	1559	1554
8	1536	1560	1551	1571	1570
9	1555	1538	1540	1570	1586
10	1554	1542	1532	1572	1577
Avg.	**1555.4**	**1542**	**1559**	**1565.6**	**1580.2**
Std. Dev.	**8.66**	**22.46**	**15.41**	**14.62**	**12.02**

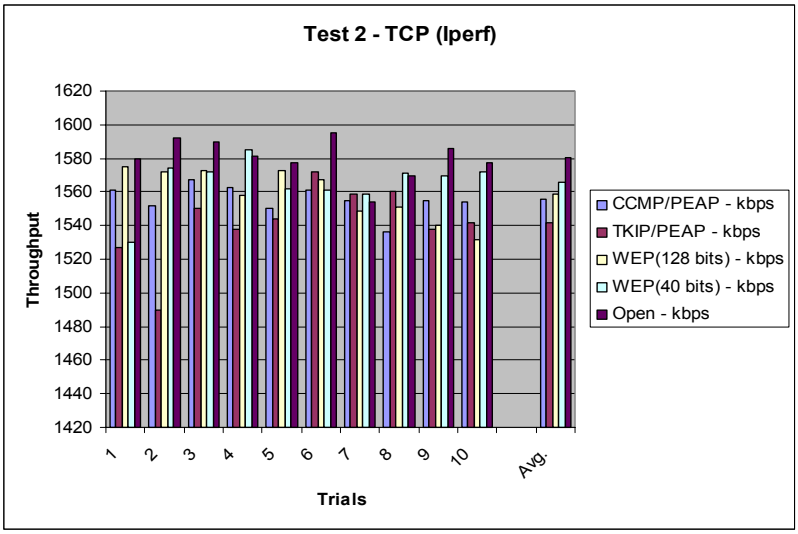

Figure 6-3. Results for Throughput Configuration 1 / TCP - Test 2

Table 6-3. Results for Throughput Configuration 1 / TCP - Test 3

Test 3 (Qcheck): TCP, 1000 kbytes, -43 dBm

Trial	CCMP/PEAP - kbps	TKIP/PEAP - kbps	WEP(128 bits) - kbps	WEP(40 bits) - kbps	Open - kbps
1	1586	1560	1594	1573	1563
2	1523	1532	1573	1547	1588
3	1539	1568	1571	1579	1585
4	1589	1560	1572	1597	1599
5	1566	1511	1587	1587	1604
6	1517	1568	1566	1580	1597
7	1572	1569	1584	1583	1568
8	1569	1552	1568	1565	1559
9	1549	1546	1599	1577	1573
10	1584	1547	1591	1606	1593
Avg.	**1559.4**	**1551.3**	**1580.5**	**1579.4**	**1582.9**
Std. Dev.	**26.10**	**18.46**	**11.90**	**16.30**	**16.07**

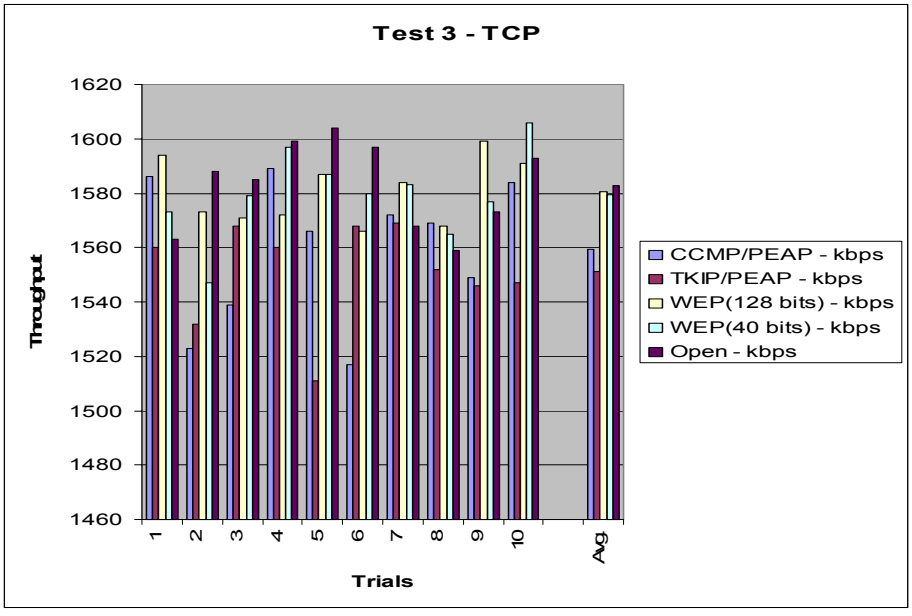

Figure 6-4. Results for Throughput Configuration 1 / TCP - Test 3

The results from these three tests were comparable. There was an obvious degradation in throughput due to TKIP encryption, but only a minor throughput loss for CCMP and WEP. The averages of these three tests are presented in Table 6-4 and Figure 6-5.

Table 6-4. Results for Throughput Configuration 1 / TCP Averages

Test	CCMP/PEAP - kbps	TKIP/PEAP - kbps	WEP(128 bits) - kbps	WEP(40 bits) - kbps	Open - kbps
Test 1	1551.2	1524.2	1555.6	1559.3	1570.3
Test 2	1555.4	1542	1559	1565.6	1580.2
Test 3	1559.4	1551.3	1580.5	1579.4	1582.9
Avg.	**1555.33**	**1539.17**	**1565.03**	**1568.10**	**1577.80**

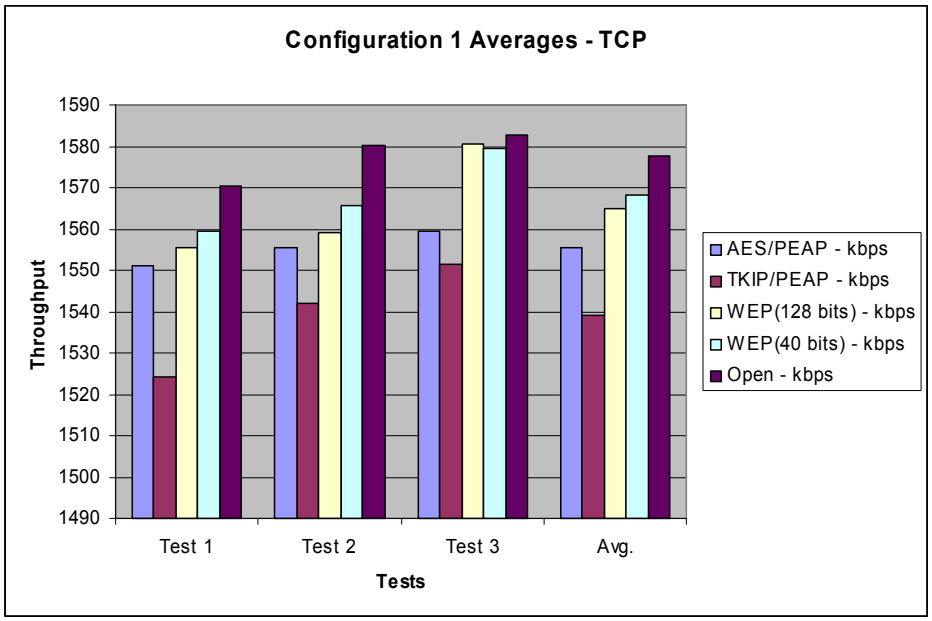

Figure 6-5. Results for Throughput Configuration 1 / TCP Averages

When examining the averages from these three tests a clear pattern can be observed between the various types of encryption and how they affect throughput on the link. In terms of overhead percentages associated with each of these encryption schemes: 64-bit WEP is 0.61%, 128-bit WEP is 0.819%, TKIP is 2.45%, and CCMP is 1.42%. The margins between these encryptions schemes are small. However, they could be detected through multiple trials at the lower 2-Mbps link rate.

6.1.1.2 UDP Results

As previously mentioned, UDP throughput data appeared to be incorrect when measured from Iperf. This is evident by the fact that UDP throughput values measured with Iperf were roughly thirty percent lower than TCP values. An engineer from Cisco Systems attributed this to some type of overloading on the server with the transmission bandwidth set to 1 Mbps over the 2 Mbps link rate. When the values were measured with Qcheck in Test 3, more reasonable values for UDP traffic were observed, where UDP throughput slightly outpaced TCP throughput. Results are given in Table 6-5 through 6-7 and Figures 6-6 through 6-8.

Table 6-5. Results for Throughput Configuration 1 / UDP – Test 1

Test 1 (Iperf): UDP, 128K Window, -43 dBm

Trial	CCMP/PEAP - kbps	TKIP/PEAP - kbps	WEP(128 bits) - kbps	WEP(40 bits) - kbps	Open - kbps
1	1050	1049	1049	1047	1041
2	1049	1045	1049	1047	1050
3	1050	1047	1041	1049	1047
4	1049	1049	1043	1024	1046
5	1043	1047	1049	1036	1049
6	1045	1049	991	1045	1043
7	1047	1044	1048	1047	1048
8	1049	1043	1047	1048	1049
9	1048	1049	1041	1019	1049
10	1050	1049	1047	1049	1034
Avg.	**1048**	**1047.1**	**1040.5**	**1041.1**	**1045.6**
Std. Dev.	**2.36**	**2.33**	**17.68**	**11.05**	**4.99**

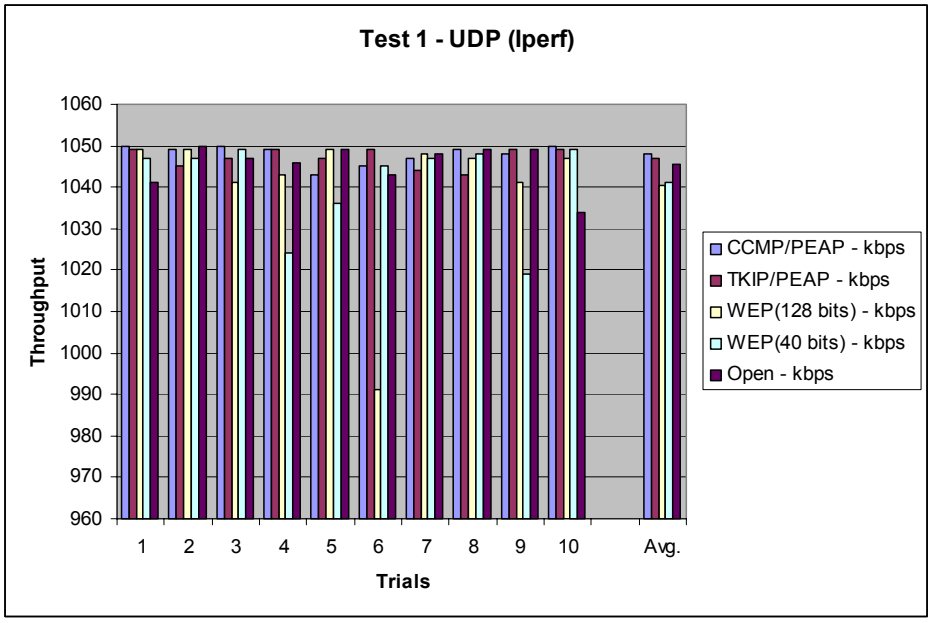

Figure 6-6. Results for Throughput Configuration 1 / UDP – Test 1

Table 6-6. Results for Throughput Configuration 1 / UDP – Test 2

Test 2 (Iperf): UDP, 128K Window, -44 dBm

Trial	CCMP/PEAP - kbps	TKIP/PEAP - kbps	WEP(128 bits) - kbps	WEP(40 bits) - kbps	Open - kbps
1	1049	1049	1045	1047	1047
2	1049	1049	1049	1049	1049
3	1049	1047	1049	1049	1047
4	1046	1049	1045	1049	1043
5	1045	1047	1050	1040	1050
6	1049	1044	1050	1047	1047
7	1047	1047	1050	1049	1047
8	1049	1049	1048	1040	1049
9	1043	1047	1042	1049	1042
10	1050	1043	1048	1049	1049
Avg.	**1047.6**	**1047.1**	**1047.6**	**1046.8**	**1047**
Std. Dev.	**2.27**	**2.13**	**2.72**	**3.68**	**2.62**

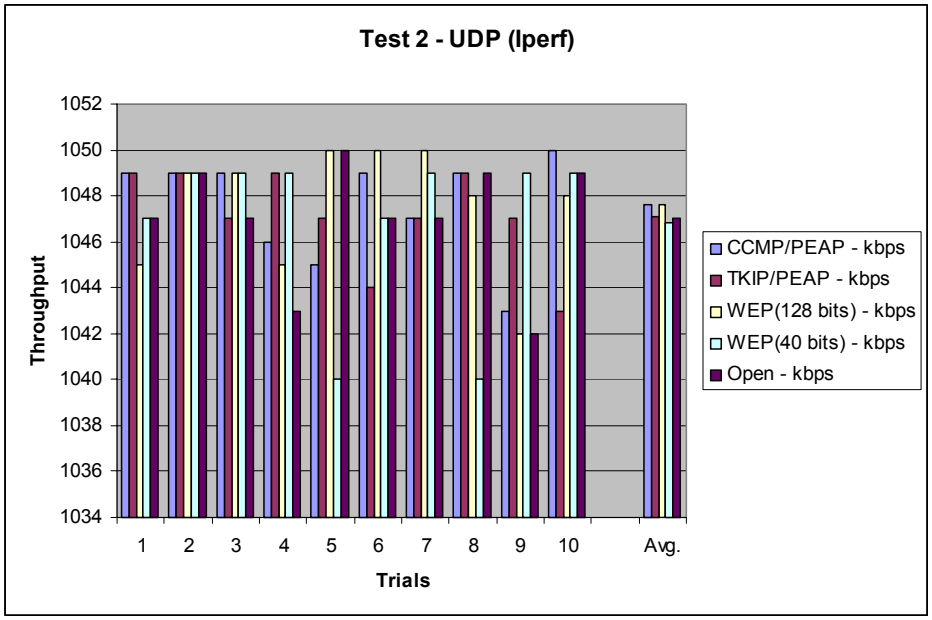

Figure 6-7. Results for Throughput Configuration 1 / UDP – Test 2

Table 6-7. Results for Throughput Configuration 1 / UDP – Test 3

Test 3 (Qcheck): UDP, 1000 kbytes, -43 dBm

Trial	CCMP/PEAP - kbps	TKIP/PEAP - kbps	WEP(128 bits) - kbps	WEP(40 bits) - kbps	Open - kbps
1	1630	1599	1616	1625	1658
2	1635	1570	1589	1594	1646
3	1574	1607	1621	1618	1599
4	1587	1622	1586	1632	1638
5	1608	1594	1634	1561	1600
6	1619	1591	1635	1579	1654
7	1629	1588	1609	1628	1644
8	1606	1596	1647	1626	1637
9	1616	1625	1618	1624	1630
10	1596	1614	1581	1634	1599
Avg.	**1610**	**1600.6**	**1613.6**	**1612.1**	**1630.5**
Std. Dev.	**19.84**	**16.76**	**22.41**	**25.16**	**22.96**

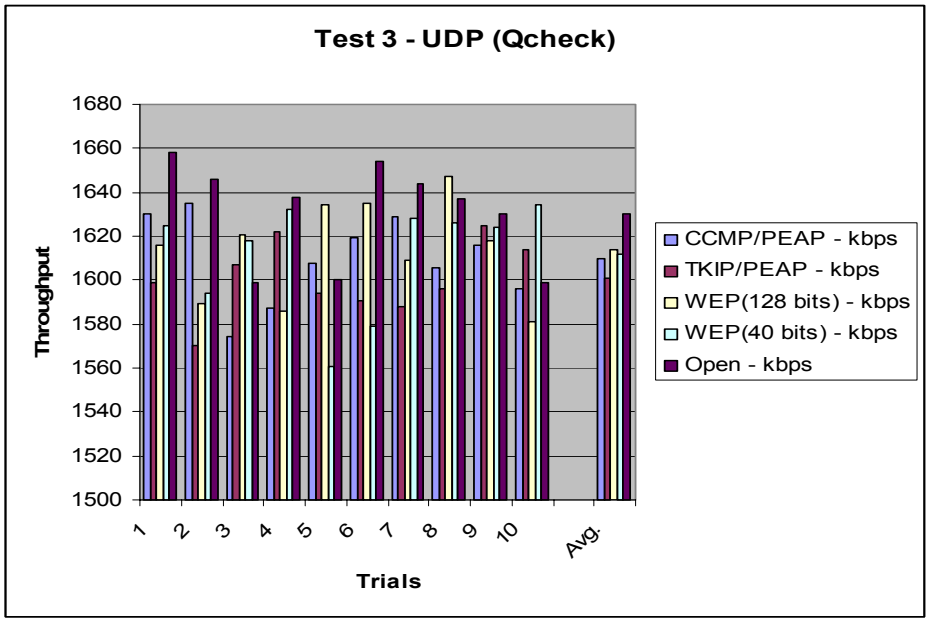

Figure 6-8. Results for Throughput Configuration 1 / UDP – Test 3

6.1.1.3 Averages

The UDP throughput values measured with Qcheck (the values from Iperf were essentially ignored) outpaced TCP values by roughly 2 to 3 percent, which was still significantly lower than the predicted 20 to 30 percent increase that UDP typically has

over TCP. Despite discrepancies between TCP and UDP traffic the effects of various encryption types appeared to be constant in the UDP values as well.

For example, when overhead percentages of TCP are compared to UDP (See Table 6-8 and Figure 6-9) there are only minor discrepancies associated with the WEP algorithms. By further averaging the overhead percentages of TCP traffic with UDP traffic a "real world" traffic environment can be partially represented and the net effect of encryption can be observed.

Table 6-8. Results for Overhead Configuration 1 Averages

	CCMP/PEAP	TKIP/PEAP	WEP(128 bits)	WEP
TCP	1.42%	2.45%	0.81%	0.61%
UDP	1.26%	1.83%	1.04%	1.13%
Avg.	1.34%	2.14%	0.92%	0.87%

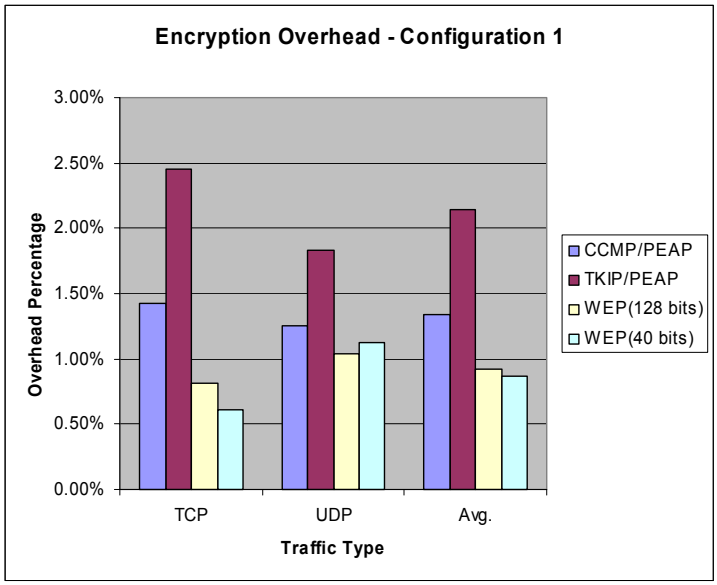

Figure 6-9. Results for Overhead Configuration 1 Averages

6.1.2 Configuration 2

This configuration was tested with both laptop computers in close proximity of one another in the lab deck area. It is obvious that the addition of another wireless client into the system reduces the overall throughput of the network, especially if testing is conducted between the two clients, which is the case in this configuration. The next reasonable question to pose is whether or not the encryption overhead from testing Configuration 1 would be doubled because the encapsulation/decapsulation process associated with the various algorithms is being conducted twice as much in Configuration 2.

It was also predictable that the number of errant values present would be roughly doubled from the first series of experiments. This could only further complicate the measurement process in an already precise environment, because widely skewed data could compromise the final averages of the system. Because of this it was very critical to conduct the testing over a small period of time to avoid even the slightest environmental changes.

6.1.2.1 TCP

Results for TCP are given in Table 6-9 and Figure 6-10.

Table 6-9. Results for Throughput Configuration 2 / TCP

TCP - 1000 Kbytes (-44dBm)

Trial	CCMP/PEAP - kbps	TKIP/PEAP - kbps	WEP(128 bits) - kbps	WEP(40 bits) - kbps	Open - kbps
1	780.6	783.9	770.1	776.6	786.9
2	765.6	780.2	770.9	786	778.6
3	792.8	750.3	773.3	792	794.1
4	784.8	783.6	775.5	770	776.2
5	778.3	782.3	770.1	788.5	786.3
6	778.4	771.1	769.5	782	793.5
7	761.9	746.1	784.5	792.7	798
8	781.5	761.1	789.3	776.6	797.9
9	777.2	773.5	790.1	797.8	781.8
10	775.3	772.1	737.9	775.6	781.7
Avg.	**777.64**	**770.42**	**773.12**	**783.78**	**787.5**
Std. Dev.	**8.85**	**13.70**	**14.76**	**9.03**	**7.98**

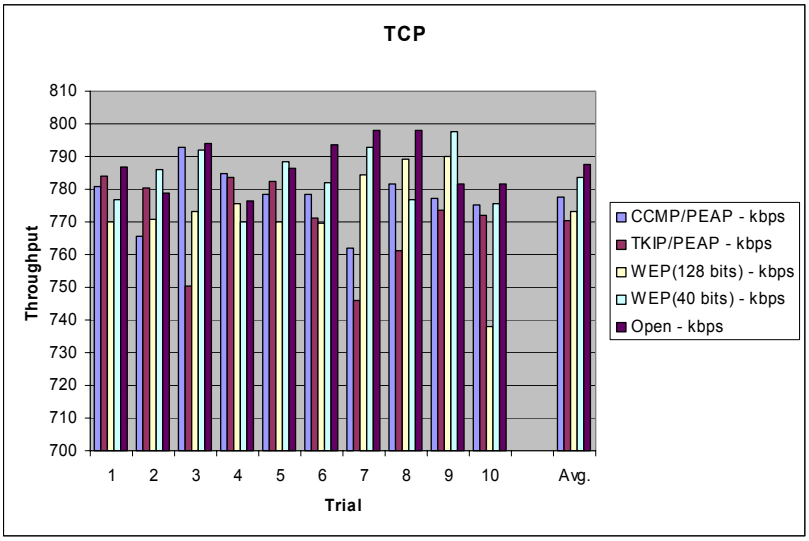

Figure 6-10. Results for Throughput Configuration 2 / TCP

The values obtained for TCP throughput were almost exactly 50 percent less than those measured in Configuration 1, which is the expected result of adding another wireless client into the system.

6.1.2.2 UDP

Results for TCP are given in Table 6-10 and Figure 6-11

Table 6-10. Results: Configuration 2 UDP

UDP - 1000 kbytes (-44 dBm)

Trial	CCMP/PEAP - kbps	TKIP/PEAP - kbps	WEP(128 bits) - kbps	WEP(40 bits) - kbps	Open - kbps
1	829.9	817.9	757.1	751.4	815.6
2	821.9	781.7	809	742.5	816.6
3	773.4	753.8	764.6	802.5	802.7
4	800.5	790.9	813.8	811.1	829.7
5	808.7	804	832	814.5	817.1
6	799.7	795.2	791.8	840	789.8
7	802.5	794.6	778.1	803.5	787.9
8	802.7	811.42	803.4	811.4	840.2
9	790.8	812.2	813.3	822.7	832.7
10	825.4	767	789.5	803.5	831.7
Avg.	**805.55**	**792.872**	**795.26**	**800.31**	**816.4**
Std. Dev.	**16.97**	**20.55**	**23.59**	**30.33**	**18.11**

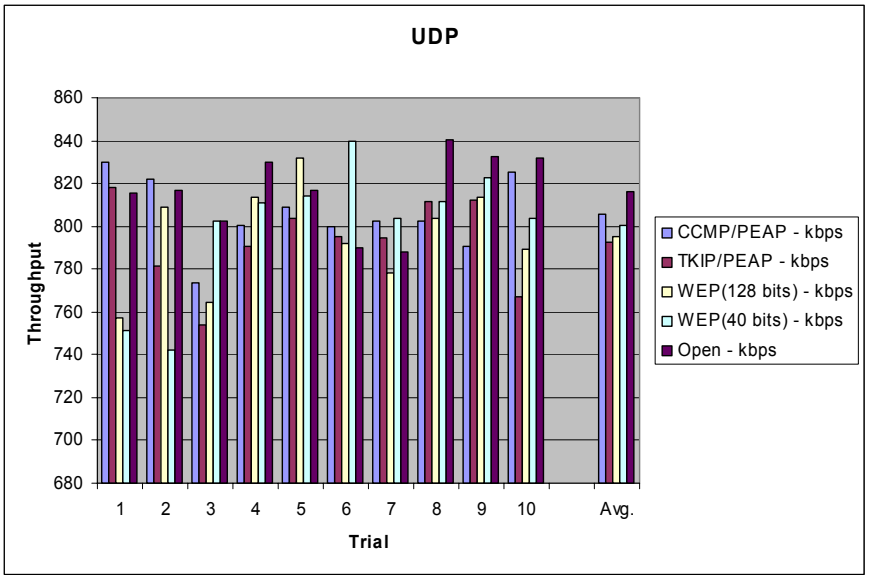

Figure 6-11. Results for Throughput Configuration 2 / UDP

6.1.2.3 Averages

The UDP results are similar in pattern to those for TCP, and were similar to the results from Configuration 1. The percent overhead results from Configuration 2 are in Table 6-11 and Figure 6-12.

Table 6-11. Results for Overhead Configuration 2 Averages

	CCMP/PEAP	TKIP/PEAP	WEP(128 bits)	WEP(40 bits)
TCP	1.25%	2.17%	1.83%	0.47%
UDP	1.33%	2.88%	2.59%	1.97%
Avg.	1.29%	2.52%	2.21%	1.22%

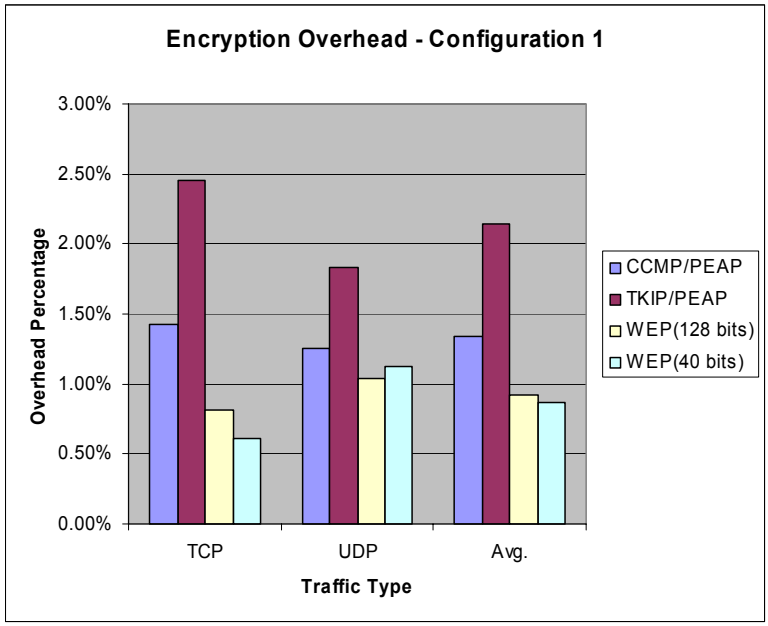

Figure 6-12. Results for Overhead Configuration 2 Averages

These results are interesting because, as expected, the WEP encryption overhead is roughly double what it is from the testing conducted in Configuration 1. However, the CCMP and TKIP overhead is almost exactly the same as in Configuration 1. Whether or not this is an effect of multiple encapsulation/decapsulation processes or just the result of errant data is hard to determine.

6.1.3 Configuration 3

Configuration 3 moved Laptop #2 behind one four-inch wall, thus reducing the signal strength approximately 6 dBm. This test was intended to determine if a small reduction in received signal strength would increase the amount of overhead generated by

encryption. The Cisco 1000 Series access point is rated for full 54 Mbps throughput up to -73 dBm signal strength. The 2 Mbps rate, which is used for, is available until signal strength drops below -94 dBm.

6.1.3.1 TCP

Results for TCP are given in Table 6-12 and Figure 6-13.

Table 6-12. Results for Throughput Configuration 3 / TCP

TCP - 1000 Kbytes (-50dBm)

Trial	CCMP/PEAP - kbps	TKIP/PEAP - kbps	WEP(128 bits) - kbps	WEP(40 bits) - kbps	Open - kbps
1	1570	1511	1553	1583	1592
2	1533	1511	1549	1587	1566
3	1576	1548	1526	1575	1604
4	1556	1517	1566	1572	1596
5	1560	1530	1545	1577	1574
6	1573	1531	1502	1589	1565
7	1536	1537	1524	1565	1587
8	1570	1523	1554	1598	1594
9	1566	1540	1593	1566	1588
10	1544	1518	1587	1581	1594
Avg.	**1558.4**	**1526.6**	**1549.9**	**1579.3**	**1586**
Std. Dev.	**15.68**	**12.66**	**28.05**	**10.42**	**13.24**

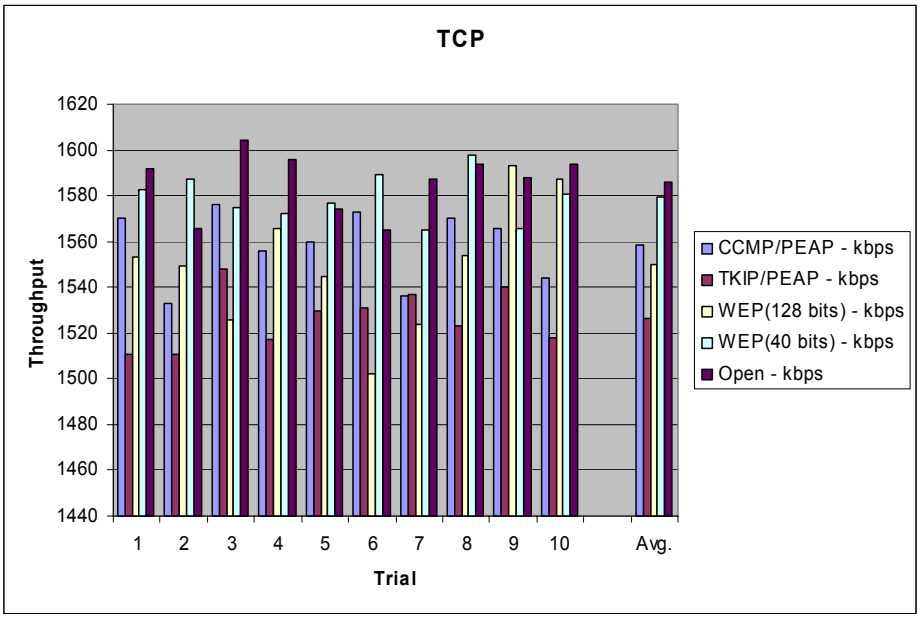

Figure 6-13. Results for Throughput Configuration 3 / TCP

6.1.3.2 UDP

Results for UDP are given in Table 6-13 and Figure 6-14.

Table 6-13. Results for Throughput Configuration 3 / UDP

UDP - 1000 kbytes (-51 dBm)

Trial	CCMP/PEAP - kbps	TKIP/PEAP - kbps	WEP(128 bits) - kbps	WEP(40 bits) - kbps	Open - kbps
1	829.9	817.9	757.1	751.4	815.6
2	821.9	781.7	809	742.5	816.6
3	773.4	653.8	764.6	802.5	802.7
4	800.5	790.9	813.8	811.1	829.7
5	808.7	804	832	814.5	817.1
6	799.7	795.2	791.8	840	789.8
7	802.5	794.6	778.1	803.5	787.9
8	802.7	811.42	803.4	811.4	840.2
9	790.8	812.2	813.3	822.7	832.7
10	825.4	767	789.5	803.5	831.7
Avg.	**805.55**	**782.872**	**795.26**	**800.31**	**816.4**
Std. Dev.	**16.97**	**47.86**	**23.59**	**30.33**	**18.11**

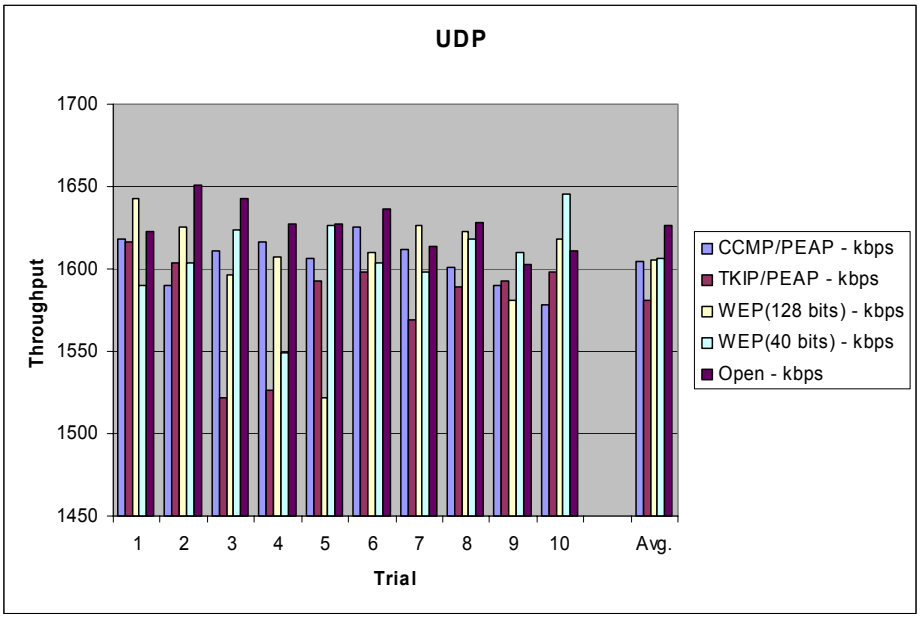

Figure 6-14. Results for Throughput Configuration 3 / UDP

6.1.3.3 Averages

These results indicate that the 6 dBm attenuation had virtually no effect on the throughput of the link for any of the encryption schemes utilized. The values were similar to those found from testing Configuration 1.

The averages presented in Table 6-14 and Figure 6-15 show the similarities to results for Configuration 1.

Table 6-14. Results for Overhead Configuration 3 Averages

	CCMP/PEAP	TKIP/PEAP	WEP(128 bits)	WEP(40 bits)
TCP	1.74%	3.75%	2.28%	0.42%
UDP	1.33%	2.80%	1.30%	1.20%
Avg	1.53%	3.27%	1.79%	0.81%

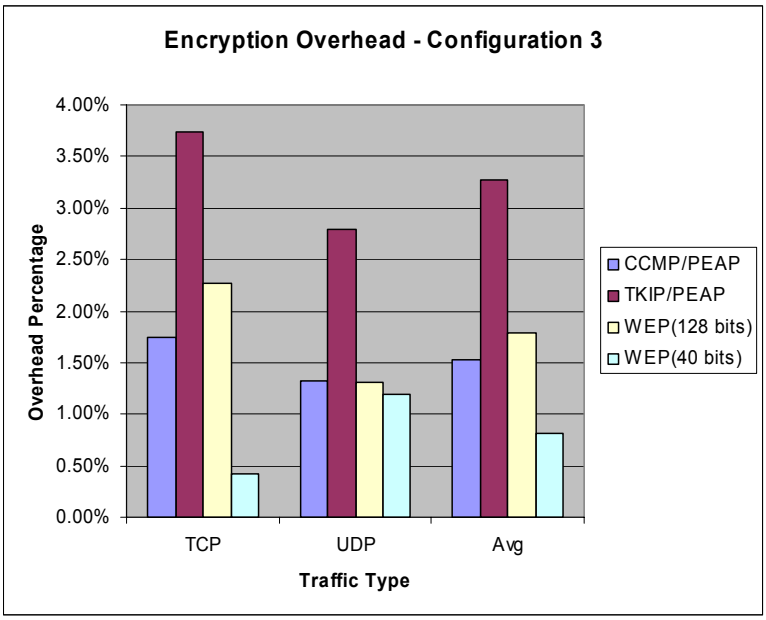

Figure 6-15. Results for Overhead Configuration 3 Averages

6.1.4 Configuration 4

Configuration 4 extended the results of Configuration 3 in that it showed that slightly reduced signal strength has little effect on throughput at 2 Mbps. The two laptop computers were each moved into locations in the building with 8 to 15 dBm attenuation levels due to physical obstructions. The throughput tests were then run as they were in Configuration 2 with the walls being the only physical change.

6.1.4.1 TCP

The results for TCP are shown in Table 6-15 and Figure 6-16.

Table 6-15. Results for Throughput Configuration 4 / TCP

TCP – 1000 Kbytes (-52dBm, -60 dBm)

Trial	CCMP/PEAP - kbps	TKIP/PEAP - kbps	WEP(128 bits) - kbps	WEP(40 bits) - kbps	Open - kbps
1	771	783	768	786	794
2	783	779	798	775	798
3	782	772	797	794	796
4	785	789	796	792	786
5	783	776	754	795	785
6	780	786	788	779	780
7	779	770	800	770	743
8	780	766	776	778	794
9	791	782	800	800	785
10	789	776	750	779	798
Avg.	**782.3**	**777.9**	**782.7**	**784.8**	**785.9**
Std. Dev.	**5.56**	**7.29**	**19.43**	**10.01**	**16.34**

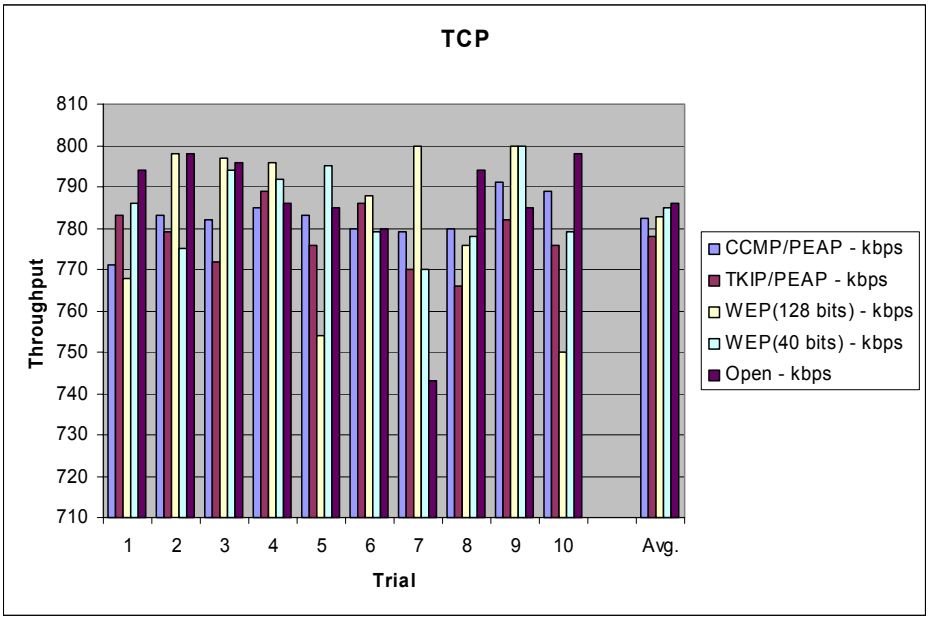

Figure 6-16. Results for Throughput Configuration 4 / TCP

6.1.4.2 UDP

Results for UDP are shown in Table 6-16 and Figure 6-17.

Table 6-16. Results for Throughput Configuration 4 / UDP

UDP - 1000 kbytes (-52 dBm, -60 dBm)

Trial	CCMP/PEAP - kbps	TKIP/PEAP - kbps	WEP(128 bits) - kbps	WEP(40 bits) - kbps	Open - kbps
1	806	808	813	800	824
2	793	761	807	813	818
3	808	814	812	819	810
4	795	810	835	796	804
5	764	789	820	812	808
6	817	809	801	804	802
7	817	817	785	788	804
8	817	750	827	758	815
9	808	745	802	810	814
10	814	750	750	823	781
Avg.	**803.9**	**785.3**	**805.2**	**802.3**	**808**
Std. Dev.	**16.48**	**30.25**	**23.98**	**18.83**	**11.75**

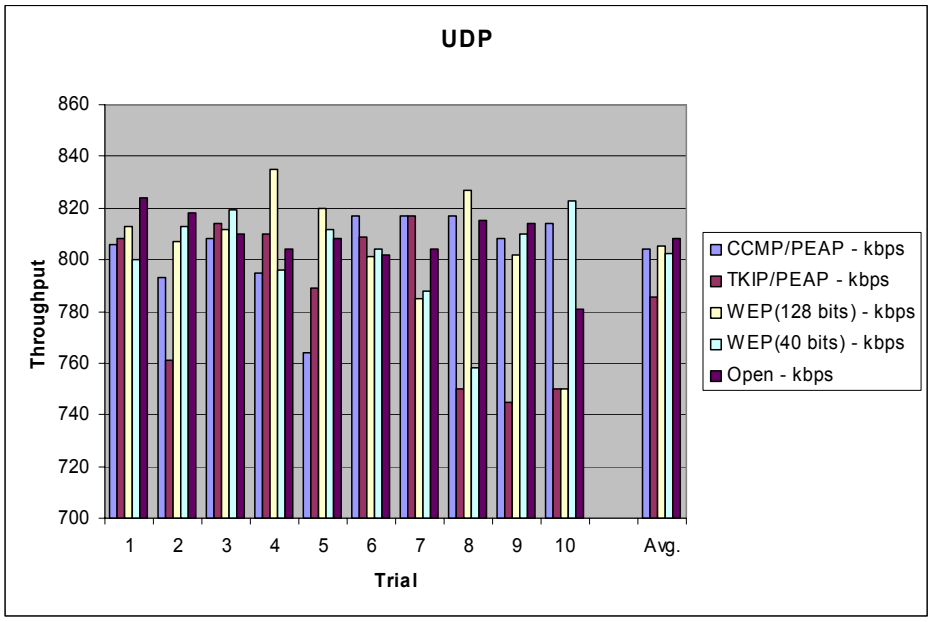

Figure 6-17. Results for Throughput Configuration 4 / UDP

6.1.4.3 Averages

The values from this test were somewhat unexpected as there was not a discernable difference between the 64- and 128-bit WEP keys. This could have been caused by slight network inconsistencies when the measurements were completed.

Table 6-17. Results for Throughput Configuration 4 Averages

	CCMP/PEAP	TKIP/PEAP	WEP(128 bits)	WEP(40 bits)
TCP	0.46%	1.02%	0.41%	0.14%
UDP	0.51%	2.81%	0.35%	0.71%
Avg	0.48%	1.91%	0.38%	0.42%

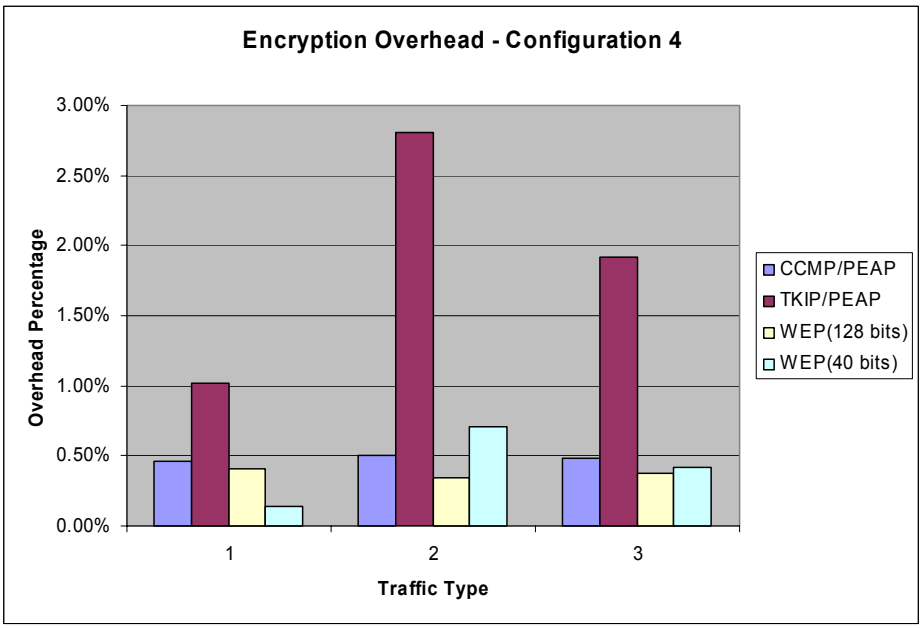

Figure 6-18. Results for Throughput Configuration 4 Averages

6.2 Authentication Overhead Results

The authentication testing was much less extensive as the only measurements that were required were response times from the client to the server and the total time authentication required over the link. These response time tests were conducted during the encryption tests and provide baseline information for the time required for packets to be acknowledged at the client by the server. These times were measured in Configurations 1 and 3, which were also used for the total authentication times with LEAP and PEAP.

6.2.1 Configuration 1

The response and authentication times in this configuration can help illustrate authentication overhead in a network that requires secure authentication. The data shows that authentication can take between fifteen to twenty times longer than a normal response time. In an environment where fast roaming is not enabled this could mean a significant delay in data transfers when roaming between access points.

6.2.1.1 Response Times

Table 6-18 and Figure 6-19 show results for response times for

Table 6-18. Results for Configuration 1 Response Times

Response Times

Trial	CCMP/PEAP - ms	TKIP/PEAP - ms	WEP(128 bits) - ms	WEP(40 bits) - ms	Open - ms
1	4	4	4	4	6
2	4	4	4	4	4
3	4	5	5	5	4
4	4	4	4	5	4
5	4	4	5	5	4
6	4	5	4	4	5
7	4	6	6	4	4
8	4	5	5	5	4
9	4	4	4	5	4
10	4	4	5	4	4
Avg.	**4**	**4.5**	**4.6**	**4.5**	**4.3**
Sd Dev.	**0.00**	**0.71**	**0.70**	**0.53**	**0.67**

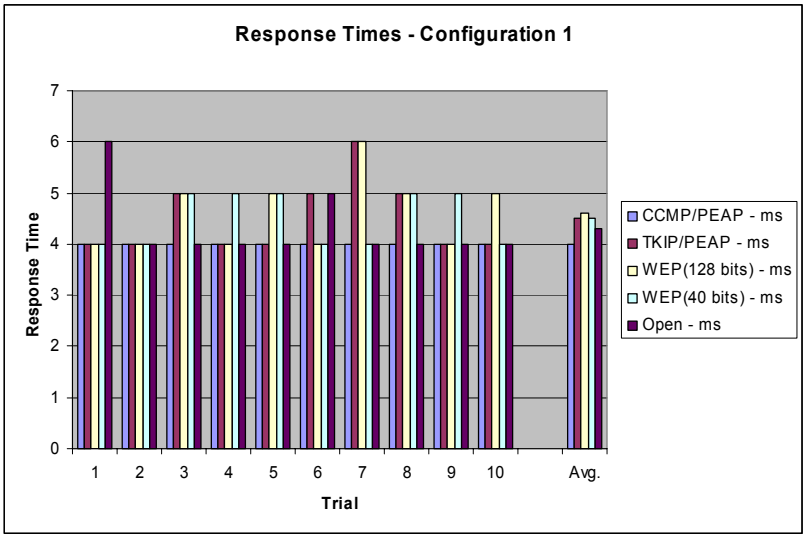

Figure 6-19. Results for Configuration 1 Response Times

6.2.1.2 Authentication Times

Authentication times are shown in Table 6-19 and Figure 6-20.

Table 6-19. Results for Configuration 1 Authentication Times

Configuration 1 (-43 dBm)		
Trials	**CCMP/PEAP - msec.**	**CCMP/LEAP - msec.**
1	95	85
2	63	49
3	77	81
4	74	43
5	67	47
6	113	71
7	91	83
8	75	43
9	68	77
10	71	75
Avg.	**79.4**	**65.4**
Std. Dev.	**15.58**	**17.66**

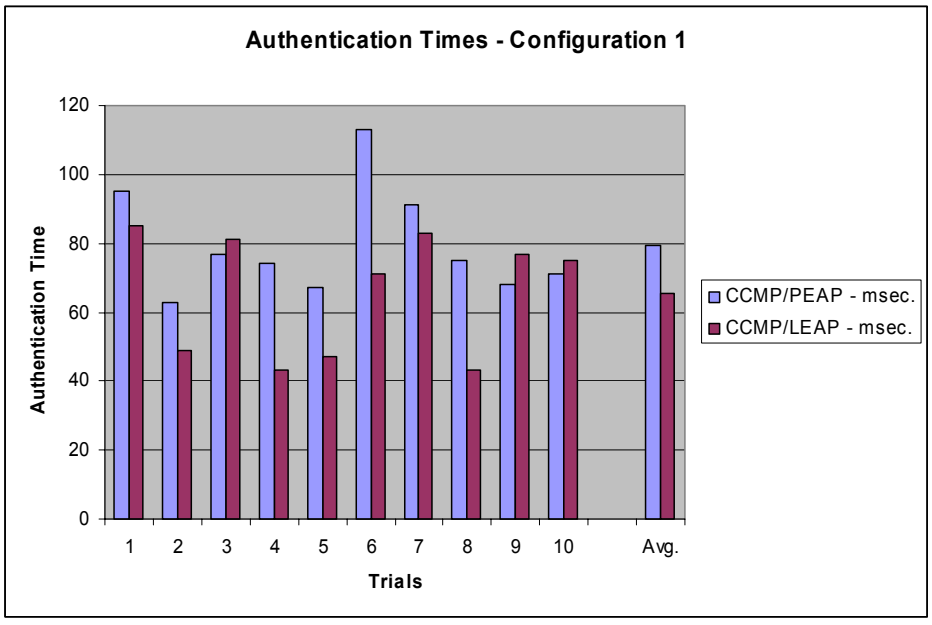

Figure 6-20. Results for Configuration 1 Authentication Times

As indicated by the results, there is a larger authentication time requirement for PEAP. This is due to the fact that PEAP requires several more packet exchanges than LEAP and therefore takes more time to accomplish. It is also apparent from this data that response time was affected little by the various encryption schemes.

6.2.2 Configuration 3

The changes between Configuration 1 and 3 were identical to those in the encryption trials. The intent was to reduce the received signal strength at the client.

6.2.2.1 Response Times

Response times are shown in Table 6-20 and Figure 6-21.

Table 6-20. Results for Configuration 3 Response Times

Response Times

Trial	CCMP/PEAP – ms	TKIP/PEAP - ms	WEP(128 bits) - ms	WEP(40 bits) - ms	Open - ms
1	7	7	6	6	7
2	7	7	6	7	7
3	7	7	6	7	6
4	7	6	6	6	7
5	6	7	7	7	7
6	7	7	7	7	8
7	7	7	6	6	7
8	7	7	7	7	7
9	7	7	6	6	7
10	7	8	6	6	7
Avg.	**6.9**	**7**	**6.3**	**6.5**	**7**
Std. Dev.	**0.32**	**0.47**	**0.48**	**0.53**	**0.47**

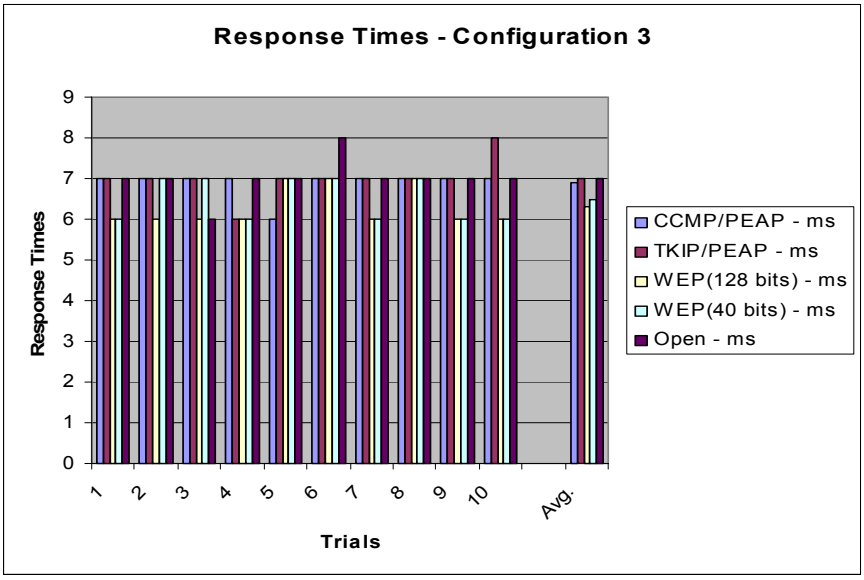

Figure 6-21. Results for Configuration 3 Response Times

6.2.2.2 Authentication Times

Authentication times are shown in Table 6-21 and Figure 6-22.

Table 6-21. Results for Configuration 3 Authentication Times

Configuration 3 (-50 dBm)		
Trial	**CCMP/PEAP - msec.**	**CCMP/LEAP - msec.**
1	110	84
2	106	50
3	92	48
4	93	44
5	85	79
6	82	48
7	103	84
8	87	118
9	91	46
10	90	66
Avg.	**93.9**	**66.7**
Std. Dev.	**9.34**	**24.28**

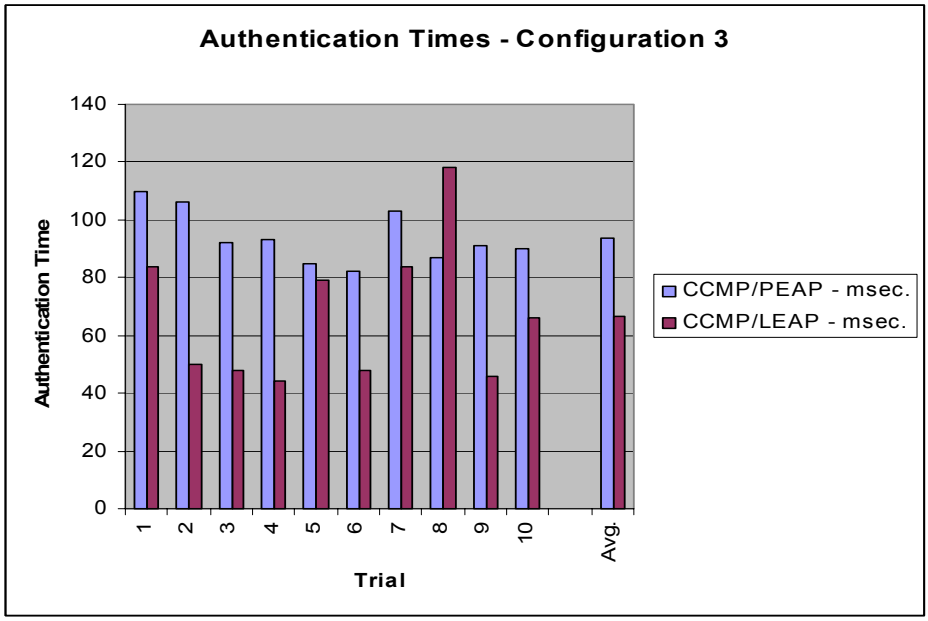

Figure 6-22. Results for Throughput Configuration 3 Authentication Times

Similar results were seen in Configuration 3. The time difference between LEAP and PEAP was roughly the same as Configuration 1. From this it can be concluded that reduced signal strength also has little effect on authentication time.

6.3 Summary

The results obtained from the experimentation seemed to meet the expectations of the research. There was a clear pattern in which TKIP generated the most overhead, followed by CCMP, and then WEP. This could be due to a number of reasons and is analyzed more closely in Chapter 7. The authentication time results reinforced data that a reduction in signal strength has little effect on throughput or response times until a defined threshold is met for that signal strength. In other words small changes in a strong IEEE 802.11 signal is likely to have little effect on performance.

Chapter 7 Analysis of Results and Future Work

7.0 Overview

The results presented in Chapter 6 represent a clear pattern that was present in throughput results associated with the various types of encryption schemes examined. The data noted a clear tendency of TKIP performance to be worse than other encryption schemes tested. More importantly, the data shows CCMP encryption scheme is only slightly less efficient than a basic WEP key, which is good news for network engineers who are hoping to implement and gain the security benefits of IEEE 802.11i. An important question to ask is why is TKIP performing worse than CCMP, and how could this affect not only enterprise level LANs, but SOHO networks.

The authentication data was intended to help show how increased authentication times could hamper network performance, particularly in roaming scenarios. However, a discussion is presented below to show how the use of fast roaming protocols should all but eliminate any potential overhead associated with wireless networks, allowing access points to rapidly hand-off clients to one another. The authentication times were helpful in reinforcing the idea that small reductions in signal strengths of a relatively strong wireless signal, has little effect on network performance.

Finally, this research is a continuation of other similar studies in the topic of wireless security overhead but there is much more to be learned. Several suggestions for future work are presented in this chapter.

7.1 Encryption Analysis

As noted in Chapter 6, there is a clear pattern present in the amount of overhead associated with WEP, TKIP, and CCMP encryption. TKIP is clearly the worst performer followed by CCMP and then WEP. However, exactly why this is and how it can hamper network resources needs to be addressed.

These encryption effects were measured on enterprise-level hardware with powerful processors capable of performing the encryption/decryption process rapidly. However, it is possible that using these encryption schemes on less powerful access points could lead to larger overhead and greater reduction in network throughput. Because the strongest level of encryption for most SOHO routers and access points is TKIP, it is quickly becoming the preferred security mechanisms at home and in small offices. If these effects are amplified by less powerful equipment it is foreseeable that TKIP could consume as much as 10 to 20 percent of the available wireless bandwidth as indicated by some of the previous studies described in Chapter 3. This would be important information for the consumer to understand before purchasing such equipment, as it may hamper their ability to conduct bandwidth intensive operations over their hardware.

Another issue to address is that only a small number of SOHO routers and access points have CCMP encryption available. Because of this TKIP will likely remain the primary encryption scheme in use for several years to come as few individuals will likely be willing to trade in their hardware for, in most cases, unnecessary security upgrades. Over time hardware vendors will likely incorporate CCMP into their equipment and the market will slowly transition to a more secure medium.

7.1.1 TKIP Performance

Because TKIP is simply a modification of the RC4 algorithm utilized in WEP, it seems intuitive that it would provide similar performance. As Hideki [Hideki06] mentions, stream algorithms, such as the RC4 algorithm, are typically fast, much faster that block ciphers. Additionally, there are only an additional four octets of data appended to a TKIP encapsulated frame than there are with a CCMP one; however, TKIP tends to introduce twice as much overhead as CCMP.

This could be based on a number of issues and is often very hard to diagnose. WEP, TKIP, and CCMP append 8, 20, and 16 octets of data, respectively, to each data frame as examined in Chapter 2, so it is feasible that the addition of 32 extra bits could be part of the problem. CCMP encryption is typically implemented within the access point or client hardware while TKIP is often done within software, which could be a major cause of throughput reduction. This is because TKIP was developed to be a software upgrade for WEP compatible hardware devices, preventing the need to replace legacy hardware. This is further compounded by the fact that TKIP has several mixing processes operating at the same time to generate the data stream. The generation of the WEP seeds itself is more complex in TKIP than standard WEP, and the added complexity of the Michael algorithm adds another layer of overhead.

In short, until legacy Wi-Fi devices that utilize WEP are completely phased out of the market TKIP will remain the primary method of encryption for home, small office, and even a small number of enterprise wireless networks. The above information should provide adequate throughput data for those individuals who intend to completely utilize their wireless link. It is also an additional reason to purchase platforms capable of CCMP encryption.

7.1.2 SOHO Comparison

To compare the Cisco 1000 Series access point to a small office/home office wireless access point, a comparison test for Configuration 1 was completed with a Linksys WRT54G access point. The intent of this alternative test was to draw comparisons between Cisco's Enterprise equipment with its home office equipment, specifically with regards to encryption overhead and to determine how much lower price models are affected by this overhead.

Similarly to Chapter 6, there are two sections that follow presenting the TCP and UDP results of this SOHO comparison along with the associated averages. However, this testing was done at 11 Mbps instead 2 Mbps due to software conflicts that prevented users from utilizing the access points at slower speeds.

7.1.2.1 TCP Results

Results for TCP are given in Table 7-1 and Figure 7-1.

Table 7-1. Results for Throughput SOHO Comparison TCP

(Qcheck): TCP, 1000 kbytes, -43 dBm

Trial	CCMP/PEAP - kbps	TKIP/PEAP - kbps	WEP(128 bits) - kbps	WEP(40 bits) - kbps	Open - kbps
1	4251	4292	4315	4333	4387
2	4267	4206	4303	4329	4384
3	4262	4267	4322	4237	4379
4	4303	4255	4245	4301	4410
5	4304	4197	4372	4310	4430
6	4269	4211	4250	4299	4386
7	4292	4233	4345	4345	4410
8	4191	4199	4348	4327	4398
9	4211	4237	4336	4378	4430
10	4290	4262	4306	4360	4422
Avg.	**4264**	**4235.9**	**4314.2**	**4321.9**	**4403.6**
Std. Dev.	**37.90**	**32.59**	**40.92**	**38.93**	**19.49**

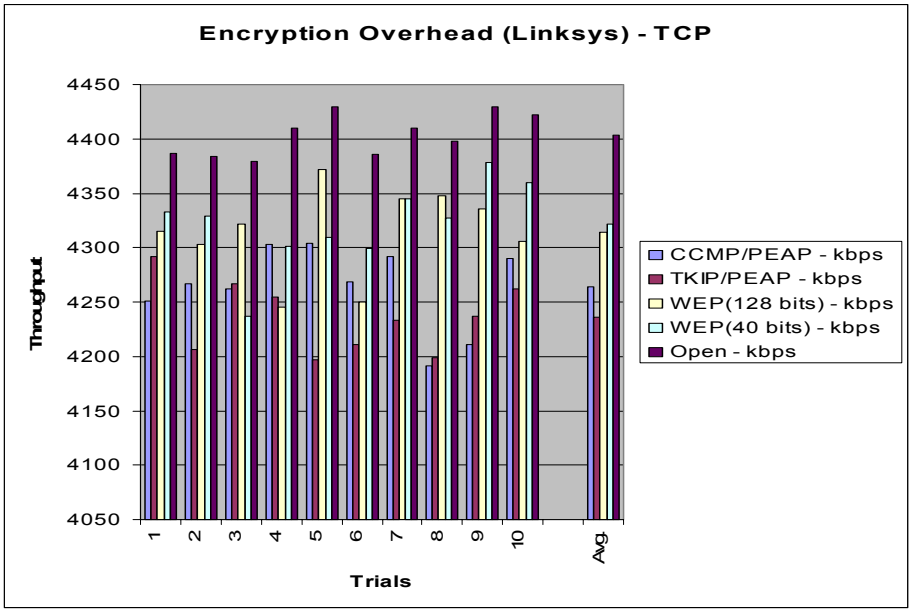

Figure 7-1. Results for Throughput SOHO Comparison TCP

7.1.2.1 UDP Results

UDP results are show in Table 7-2 and Figure 7-2.

Table 7-2. Results for Throughput SOHO Comparison UDP

Test 3 (Qcheck): UDP, 1000 kbytes, -43 dBm

Trial	CCMP/PEAP - kbps	TKIP/PEAP - kbps	WEP(128 bits) - kbps	WEP(40 bits) - kbps	Open - kbps
1	5102	5009	5122	5140	5168
2	5031	4984	5061	5109	5205
3	5079	5003	5060	5086	5215
4	5076	4994	5096	5145	5125
5	5000	5041	5128	5092	5212
6	5019	5006	5096	5151	5161
7	5047	5016	5161	5118	5205
8	5125	5017	5141	5083	5191
9	5041	5031	5125	5051	5200
10	5016	4950	5125	5101	5205
Avg.	**5053.6**	**5005.1**	**5111.5**	**5107.6**	**5188.7**
Std. Dev.	**40.53**	**25.48**	**32.91**	**31.63**	**28.71**

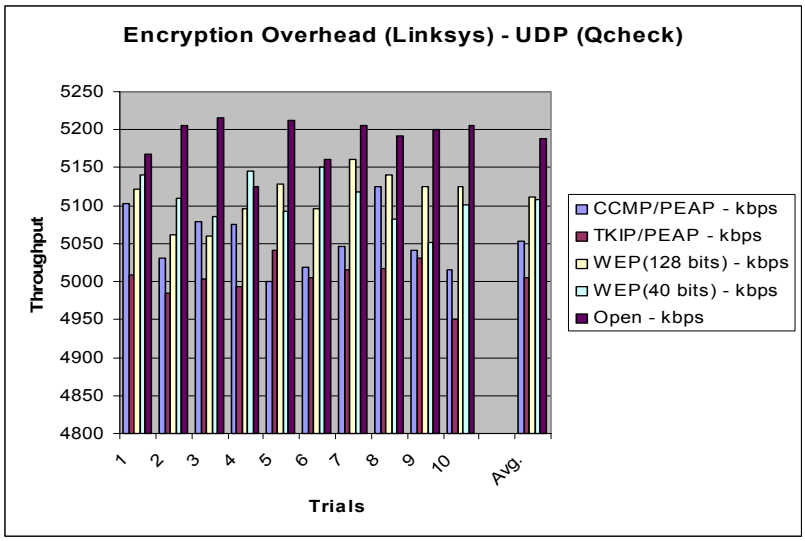

Figure 7-2. Results for Throughput SOHO Comparison UDP

7.1.2.1 Averages

Average results are show in Table 7-3 and Figure 7-3.

Table 7-3. Results for Overhead SOHO Comparison Averages

	CCMP/PEAP	TKIP/PEAP	WEP(128 bits)	WEP(40 bits)
TCP	3.17%	3.81%	2.03%	1.86%
UDP	2.60%	3.54%	1.49%	1.56%
Avg	2.89%	3.67%	1.76%	1.71%

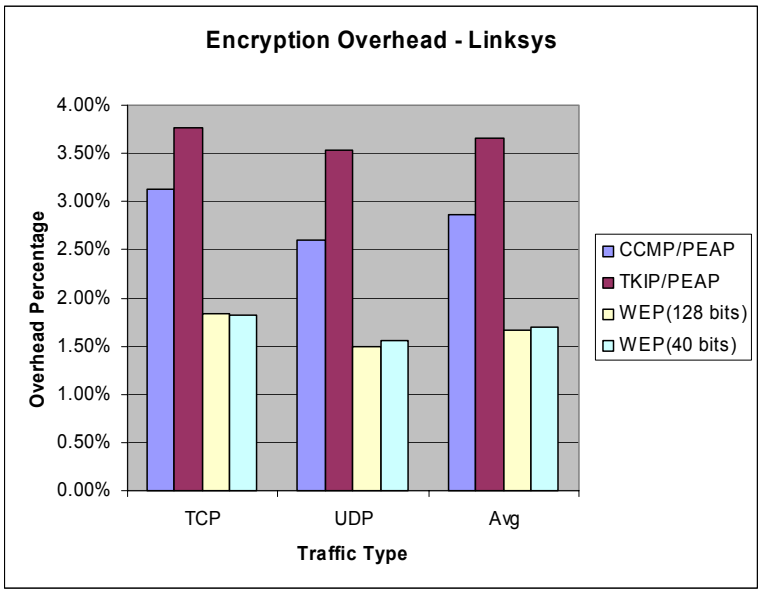

Figure 7-3. Results for Overhead SOHO Comparison Averages

The Linksys model introduced roughly twice the amount of encryption overhead as the enterprise model, Cisco 1000 Series, in Configuration 1. Although there was a decrease in performance in the SOHO model it is small enough that the added encryption would be undetectable to most users and certainly insufficient reason to avoid its use.

7.2 Authentication Analysis

The results from the authentication testing were similar to the expected results and provided little additional information. Because most enterprise wireless solutions today utilize some form of fast roaming or pre-authentication, which prevents clients from having to complete the entire authentication process during a roaming scenario, there are few problems caused by slow authentication times. However, in many networks these types of settings are not present, and upon roaming a user must reauthenticate with the RADIUS server. This can be troublesome if the RADIUS server is located several hops away from the client and can compound the authentication delay. For these scenarios it is

important that network engineers test the authentication delay in their networks to determine adequate placement for the server and other network devices. It would also seem prudent that software engineers design programs to allow for this type of delay in future product development, such as the next generation of VoIP phones.

As previously mentioned, the authentication analysis helped confirm results with regards to the signal strength. Both the encryption and authentication results showed that a small reduction in signal strength did not hamper network performance. For example, there was no reduction in throughput between Configurations 1 and 3, despite a 6 dBm signal loss. However, if the signal strength is reduced below the thresholds required for various link speeds then throughput and authentication times will likely be hampered.

7.3 Contributions

This work provided an in-depth look into the effects that encryption and authentication may have on network performance. From the results it seems that encryption on today's networks can be implemented efficiently, greatly reducing the amount of bandwidth allocated to encryption processes. Because of this there is no valid argument to forgo Layer 2 security mechanisms to prevent degradation in performance.

The work also shows the potential benefits of block ciphers such as CCMP, as a viable alternative to the stream ciphers of WEP and TKIP. Not only is CCMP more secure than previous security mechanisms, it operates efficiently. Because of this, SOHO equipment should move away from WPA solutions for wireless routers and access points and use WPA2 more prevalently.

7.4 Future Work

There are several areas of potential future work in this area that could be explored. This study attempted to test as many types of common enterprise configurations as possible but left out several that are in use or will continue to grow in the future. For example, EAP-TLS was ignored because of the requirements for client certificates within that particular authentication method. More importantly, the interaction of these other types of authentication with the current encryption schemes could be examined more thoroughly.

Although this study attempted to record data as accurately as possible it could be done even more accurately if an automated process was developed to track throughput over a period of time and report the results. This would essentially eliminate erroneous data and problems within the RF environment as the averaging power would overcome those inconsistencies.

Bibliography

[Agarwal] Agarwal, Avesh, K., Wang, Wenye. *Measuring Performance Impact of Security Protocols in Wireless Local Area Networks*. Department of Electrical and Computer Engineering, North Carolina State University, Raleigh, NC.

[Atheros03] Methodology for Testing Wireless LAN Performance. http://www.atheros.com/pt/whitepapers/Methodology_Testing_WLAN_Chariot.pdf.

[Baghaei03] Baghaei, Nilufar. *IEEE 802.11 Wireless LAN Security Performance Using Multiple Clients*. Department of Computer Science and Software Engineering. University of Canterbury, Christchurch, New Zealand, 2003 http://www.medialab.co.nz/assets/downloads/IEEE%20802.11%20Wireless%20LAN%20Security%20Performance.pdf

[CWNA05] *CWNA – Certified Wireless Network Administrator (3rd ed.)*. McGraw-Hill/Osbourne, Emeryville, CA, 2005.

[CWSP03] *CWSP – Certified Wireless Security Professional (1st ed.)*. McGraw-Hill/Osbourne, Emeryville, CA, 2003.

[Dyson99] Dyson, Peter. *Dictionary of Networking: Third Edition*. SYBEC, Alameda, CA, 1999.

[Hideki06] Hideki, Imai. *Wireless Communications Security*. Norwood, MA, Artech House, Inc., 2006.

[Jamshaid03] Jamshaid, Kamran, You, Liyu, Hamza, M. H. (2003). Performance Evaluation of Technologies for Security 802.11 Enterprise Wireless Networks in *Proceedings of the IASTED International Conference Communications, Network, and Information Security*. New York, NY, December 10-12, 2003.

[Kbar05] Kbar, Ghassan, Mansoor, Wathiq. Testing the Performance of Wireless LAN in *Asia-Pacific Conference on Communications,* Perth, Australia, October *3-5, 2003,* 492-496.

[O'Hara05] O'Hara, Bob & Petrick, Al. *IEEE 802.11 Handbook: A Designer's Companion (2nd ed.)*. IEEE Press, New York, NY, 2005.

[Pelletta03] Pelletta, Enrico, Velayos, Hector. Performance measurements of the saturation throughput. In *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks: Third International Symposium.* Riva del Garda, Italy, 3-7 April 2005, 129-138.

[Sankar05] Sankar, Krishna, Sundaralingam, Sri, Balinsky, Andrew, Miller, Darrin. *Cisco Wireless LAN Security: Expert Guidance for Securing Your 802.11 Networks*. Cisco Press, Indianapolis, IN 2005.

[Synder02] Snyder, Joel. What is 802.1x". Network World Fusion, 2002. Retrieved April 4, 2006, from Network World: http://www.networkworld.com/research/2002/0506whatisit.html.

[Wang05] Wang, Shao-Cheng, Chen, Yi-Ming, Lee, Tsern-Huei, Helmy, Ahmed. Performance Evaluations for Hybrid IEEE 802.11b and 802.11g Wireless Networks. In *Performance, Computing, and Communications Conference: 24th IEEE International.* 7-9 April 2005, 111-118.

[Wijesinha05] Wijesinha, Alexander L., Song, Yeong-tae, Krishnan, Mahesh, Marthur, Vijita, Ahn, Jin, Shyamasundar, Vijay. Throughput Measurement for UDP Traffic in 802.11g WLAN. In *Software Engineering, Artificial Intelligence, Networking and Parallel/Distribution Computing, 2005 First ACIS International Workshop on Self-Assembling Wireless Networks. SNPD/SAWN. Towson, MD,* 23-25 May 2005, 220-225.

[Wong03] Wong, Jenne. (2003). *Performance Investigation of Secure 802.11 Wireless LANs: Raising the Security Bar to Which Level?* University of Canterbury, Christchurch, New Zealand. http://www.cosc.canterbury.ac.nz/research/reports/MastTheses/2003/mast_0301.pdf