Design, Implementation, and Performance Analysis of DiscoSec – Service Pack for Securing WLANs

Ivan Martinovic, Paul Pichota, Matthias Wilhelm,
Frank A. Zdarsky, and Jens B. Schmitt

disco | Distributed Computer Systems Lab
University of Kaiserslautern, Germany
{martinovic, p_pichot, m_wilhel, zdarsky, jschmitt}@informatik.uni-kl.de

Abstract

To improve the already tarnished reputation of WLAN security, the new IEEE 802.11i security standard provides means for an enhanced user authentication and strong data confidentiality. However, the standard focuses on securing higher-layer data, i.e., protecting IEEE 802.11 data frames. Management frames used for connection administration are left unprotected and a wide spectrum of known attacks is still applicable and even extended against the IEEE 802.11i/IEEE 802.1X protocol execution.

This work describes DiscoSec, a service pack for "patching" WLANs against the most prominent vulnerabilities resulting in resource-depletion and impersonation attacks. DiscoSec provides DoS-resilient key exchange, an efficient frame authentication, and a performance-oriented implementation. By means of extensive real-world measurements the performance of DiscoSec is evaluated showing that even on very resource-limited devices the throughput is decreased by only 22 % compared to the throughput without any authentication, and by 6 % on more powerful hardware. To demonstrate its effectiveness, DiscoSec is available as an open-source WLAN device driver.

1. Introduction

The various confidentiality and integrity vulnerabilities of Wired Equivalent Privacy (WEP) and the simplicity of mounting impersonation attacks by manipulating the sender's MAC address caused the bad reputation of IEEE 802.11 security. To regain back trust in this widespread technology the IEEE Task Group *i* successfully finalized the new security standard *802.11i* [1]. The new standard provides a security framework composed of several known and approved protocols to ensure robust protection of wireless communication. An enhanced user authentication, a new underlying cipher, and a reliable integrity verification finally enabled the protection of data equivalent to the security in wired networks.

However, IEEE 802.11i focuses only on securing the user's data, i.e., it provides security for the *data frames* used to transport higher layer protocol data, leaving the *management frames* used for channel and connection administration without any protection. The reason seems to be twofold. First, the tragic end of WEP left wireless clients without standardized protection giving rise to dispersion of proprietary solutions, hence the interoperability certification program (e.g., Wi-Fi Protected Access (WPA)) was impatiently awaited by both the industry and the users. Secondly, attacks on management frames impact the *availability* of the IEEE 802.11 network, which especially in wireless networks is the most vulnerable among all security goals. Due to the frequency jamming vulnerability being an indigenous property of wireless communication, the importance of providing availability protection at the link-layer is often downgraded. Nevertheless, there is a significant difference between physical layer attacks aiming at the channel capacity, thus denying *any* communication, and link-layer attacks affecting the *services* provided by an access point and the connection states of wireless stations.

In its infrastructure mode an access point (AP) is controlling the wireless channel and providing authentication and association procedures to wireless clients. Its flawless operation, availability to manage client associations, and the administration of the wireless traffic have a direct impact on the users' security. For example, the execution of the IEEE 802.11i security standard is only possible if a wireless client reaches the final *authenticated and associated* connection state. State transitions within the IEEE 802.11 state machine are achieved by management frames, and by manipulating them even the sophisticated protection given by IEEE 802.11i is easily obviated.

Furthermore, the typical Man-In-The-Middle (MITM)

978-1-4244-2100-8/08/$25.00

attacks in wireless networks are based on abusing unauthenticated management frames. After installing a rogue AP with a stronger signal, an adversary can simply change its MAC address to any of the already associated clients (or AP) and, by sending an impersonated deauthentication or disassociation frame, it can reset a wireless client to its initial state (for more details on such attacks see [3]). Consequently, a wireless client is not able to transmit any data frames and must re-initiate the network discovery procedure which in most implementations chooses the AP with the strongest signal, hence associating with the rogue AP. Although well known, these attacks are still applicable and various tools are available to facilitate their execution (e.g., [8] demonstrates wireless phishing attacks within public hotspot scenarios).

To sustain the fast deployment of IEEE 802.11 technology there is a need for protection against such low-cost, yet very effective attacks. In this work, we describe and evaluate DiscoSec, a solution against the most prominent vulnerabilities within present WLANs. A similar goal was also set within the IEEE 802.11 Task Group *w* which is still in proposal stages, and therefore, we implement the concept of DiscoSec as an open-source IEEE 802.11 device driver to serve as a benchmark and prototype for future developments.[1]

The rest of the paper is structured as follows. In Section 2, we discuss various security objectives influencing the design of DiscoSec. The key exchange and implementation-related decisions are presented in Section 3, while Section 4 illustrates the impact of resource-depletion attacks and introduces a protection to mitigate them. In Section 5, we evaluate DiscoSec using three different hardware platforms and conduct real-world measurements to analyze the key exchange, the frame authentication, and the impact of DiscoSec on overall network throughput. Various design and implementation decisions were based on measurements using modern equipment, hence this work describes not only the final results but also different lessons learned during our research.

2. Design Goals of DiscoSec

In contrast to wired networks where end-devices have comparable hardware capabilities and executing expensive computations does not present a performance problem, introducing cryptography-based protection in wireless networks opens various performance-related issues. Especially critical are protections using stateful protocol execution and complex message exchanges, which by abusing the broadcast nature of wireless communication, often result in new resource-depletion attacks. Such examples can also be found in IEEE 802.11i where a resource demanding protocol and an unauthenticated exchange of key material are prone to various protocol-blocking attacks (e.g., [5, 6]). To mitigate such attacks and to allow interoperability between stations not supporting DiscoSec, we identified the three most important requirements on which the solution should be based:

1. *Simple and lightweight authentication protocol*
2. *DoS resilient protocol execution*
3. *No alterations to the current IEEE 802.11 state machine*

Simple and lightweight authentication is necessary for several reasons. Simplicity is a property affecting not only protocol design but also its implementation. In wireless communication where no assumption on a reliable channel should be made, protocols consisting of many round-trips (e.g., many message exchanges) often create deadlock vulnerabilities. The simplicity of authentication also assists us in reusing well-established cryptographic primitives available within the standard Linux (kernel) Crypto API and the OpenSSL library, thereby minimizing the potential for faulty implementation.

The *lightweight* property of an authentication protocol focuses on the key exchange phase where public key cryptography is used. To avoid many message round-trips we abandon the negotiation of security properties and rather utilize an anonymous Diffie-Hellman (DH) key exchange. The idea behind this decision is to shift the key exchange phase to the very beginning of the communication so that no resource reservation is made before the key exchange is finalized. At this stage no user identities are known but only their link-layer addresses, and therefore DiscoSec *binds* the sender's and receiver's link-layer address for the remainder of the session. Being protected from frame injection, devices can utilize identity authentication within the later stages of communication relying upon more heavy-weight protocols. As a result, the key exchange is executed within only *one round-trip* (i.e. two messages) whilst supporting the second important design property — DoS resilient protocol execution.

DoS resilience of DiscoSec concerns both computation- and memory-depletion attacks. The key exchange is the most vulnerable part of the authentication protocol and its arbitrary initiation should be avoided. For this reason we implement a rate limitation of key exchange requests which takes advantage of broadcast communication to provide fairness in the association process. The DoS protection is provided as a configuration parameter and its dimensioning can be adapted to comply with performance characteristics of the dedicated AP.

[1] DiscoSec's source code for using it as a wireless device driver is available at disco.informatik.uni-kl.de/downloads/

	Management Frames	Data Frames
State 1	Beacon, Probe Req. /Resp., Traffic Indication Message, **Authentication Req./Resp., Deauthentication.**	None (infrastructure BSS)
State 2	**Association Req./Resp., Reassociation Req./Resp., Deauthentication.**	None
State 3	**Deauthentication Disassociation**	**All frames**

Table 1. IEEE 802.11 frame types and connection states within they are allowed to be transmitted. Bold frames are authenticated by DiscoSec.

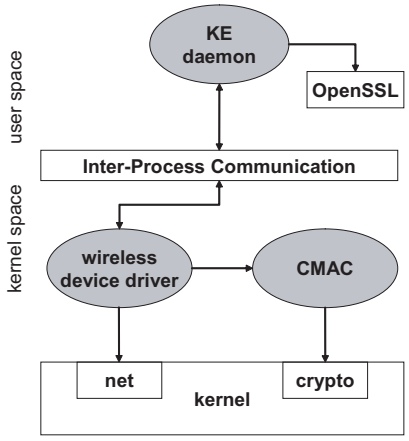

Figure 1. Architecture of DiscoSec

To support legacy and DiscoSec protected stations within the same basic service set (BSS), we design and implement DiscoSec *without changing the current IEEE 802.11 state machine*. All required information is embedded within the existing frames as *Information Elements* (IEs), i.e., the key-value data structure reserved by the IEEE 802.11 standard for transmission of custom data. The authentication data is therefore only processed if DiscoSec is implemented, otherwise it is simply discarded by the legacy driver. The frame structure remains unchanged, and a DiscoSec enhanced AP has no impact on the association procedure for legacy stations.

Various other properties identified as performance vs. security tradeoffs are offered as configuration parameters of DiscoSec's implementation and can be adapted to the network requirements.

2.1. Contribution

The security goals which DiscoSec fulfills are the *authentication* and *integrity* of management and (optionally) data frames exchanged between a wireless station and an AP. This eliminates the most prominent attacks based on injecting fake or impersonated frames, such as Deauthentication and Disassociation attacks.

Table 1 provides an overview of all management and data frames used within the three connection states of wireless clients (DiscoSec authenticated frames are depicted as bold). Using DiscoSec, both station and AP hold a shared secret before entering the second connection state, which demands reservation of the AP's memory resources. Both participants are able to prove that all subsequent unicast frames are sent from the devices that participated in the network association procedure. Furthermore, due to its high performance DiscoSec offers authentication of all *data frames* transmitted during the session and thus protects the execution of more heavy-weight protocols transmitted within them.

As a proof of concept DiscoSec is available as a WLAN device driver for all Atheros-based chip-sets, using Linux kernel (ver. 2.6.14+), and OpenSSL (ver. 0.9.8+).

3. DiscoSec's Design and Implementation

The design of DiscoSec followed the requirements discussed in the previous section. The cryptographic primitives used as building blocks for implementation of DiscoSec were chosen based on their performance properties. Decisions such as the choice of the underlying cipher, the use of a block cipher-based message authentication code (CMAC) and elliptic curve cryptography (ECC) for the key exchange were based on extensive measurements on dedicated APs (a more detailed discussion is given in Section 5).

3.1. Architecture of DiscoSec

The architecture of DiscoSec is depicted in Figure 1. The functionality is split into modules and logically divided into three functional units: *(i)* the wireless LAN device driver that controls the WLAN hardware and contains the 802.11 network stack, *(ii)* the Key Exchange (KE) daemon which provides public key cryptography features utilizing primitives offered by the OpenSSL library. It processes the key exchange requests issued by the wireless device driver via Inter-Process Communication, and *(iii)* the CMAC kernel module which provides functions for calculating Message Authentication Codes (MACs) using the kernel's standard Crypto API.

The calculation of the MACs is a time-critical operation and is thus implemented in *kernel space*. Contrary to CMAC, the Key Exchange daemon runs in *user space*

AP	Link-layer address of access point
STA	Link-layer address of wireless station
EC_{Param}	Elliptic curve parameters
PK_{AP}	Access Point's public-key
PK_{STA}	Station's public-key
MK	Master key computed from ECDH
SK	Session key used for authentication
$MAC_{SK}(m)$	Message Authentication Code
AT	Association Token

Table 2. Notation used.

with a lower priority. In case of a high CPU load, the Key Exchange daemon is scheduled less often, so the already associated stations are not influenced by expensive computations and their data throughput remains stable as shown later in the performance section of this work.

3.2. Terminology and Cryptographic Primitives Used

Table 2 shows the notation used in DiscoSec's frame authentication. All exchanged variables are defined as custom Information Elements appended to existing IEEE 802.11 frames.

The key exchange utilizes Elliptic Curve Diffie-Hellman (ECDH) as it enjoys the advantage of much smaller key sizes compared to Diffie-Hellman based on the discrete logarithm problem. The public keys PK_{AP} and PK_{STA} are available in 128 bit length (expandable to 256 bit). EC_{Param} defines an elliptic curve over a finite field supported by the OpenSSL library and changeable through DiscoSec's configuration parameters.

The shared secret MK is computed from the ECDH key exchange using PK_{AP}, PK_{STA} and the station's and AP's private keys. It serves as a *master key* to derive key material for authentication.

The association token AT is used as a *nonce* for computing a fresh session key SK for frame authentication, and additionally for providing DoS protection to control the rate of association requests (more information on tokens and the DoS protection is given in Subsection 4.1).

The MAC_{SK} is a cipher-based message authentication code utilizing AES. We selected AES as the underlying cipher due to its inclusion in the IEEE 802.11i standard and its good performance characteristics. It is provided within the Linux Crypto API (ver. 2.4+). An implementation of secure CMAC based on AES for authentication of messages with variable length was not available during the development of DiscoSec. Therefore, we implemented RFC 4493 [14], which defines AES-CMAC and serves as a NIST recommendation for CMAC message authentication using the AES block cipher [12].

3.3. Association Procedure - Key Derivation

We omit the detailed description of a public/private key initiation and the ECDH shared secret computation. Both methods are standardized and their implementations are given by OpenSSL ver. 0.9.8+ [7, 13]. In the following we describe DiscoSec's specific parameter exchange and the derivation of an authentication key.

The key exchange is accomplished within a single round-trip:

$$\text{AP} \rightarrow \text{STA} \ : \{PK_{AP}, EC_{Param}, AT\}$$

$$\text{AP} \leftarrow \text{STA} \ : \{PK_{STA}, AT\}$$

During start-up the AP initializes its key pair based on the elliptic curve parameters EC_{Param} (the AP's key pair can also be precomputed and loaded during start-up). The resulting PK_{AP} and EC_{Param} are sent to the stations via periodically emitted Beacon frames or a triggered Probe Response frame depending on either active or passive network discovery.

The wireless station extracts the supplied values, generates its key pair based on EC_{Param}, and computes the master key MK using the $ECDH$ method. The session key SK is created by applying the cryptographic hash function *SHA-256* on the MK and a previously received association token AT. The 128 least-significant bits of the hash are selected to provide the authentication key.

The reason for deriving the authentication key from the master key is to support the *key-caching* technique similar to the IEEE 802.11i standard. If the same wireless station decides to associate with the same AP and uses the static key pair, the computationally expensive $ECDH$ key exchange can be omitted. The fresh SK is then derived by applying a single hash computation on the new association token and the cached MK.

The PK_{STA} and the AT are returned within the Authentication Request to the AP, which computes the MK and SK analogously to the station side. This finalizes the key exchange and both participants use their session key SK as the secret key for the AES frame authentication ($MAC_{SK}(m)$).

This is also the most critical part of the association procedure. While the AP's key pair can be calculated offline and loaded during start-up, the session key derivation is triggered by an Authentication Request and its computation depends on PK_{STA} and AT. This opens a new vulnerability because the Authentication Request can easily be faked, and validation is only possible after the AP has derived MK by performing complex modular computations. If Authentication Request frames are received faster than the AP's transmission queue is processed, the AP can suffer from high frame loss and various operational anomalies. To prevent

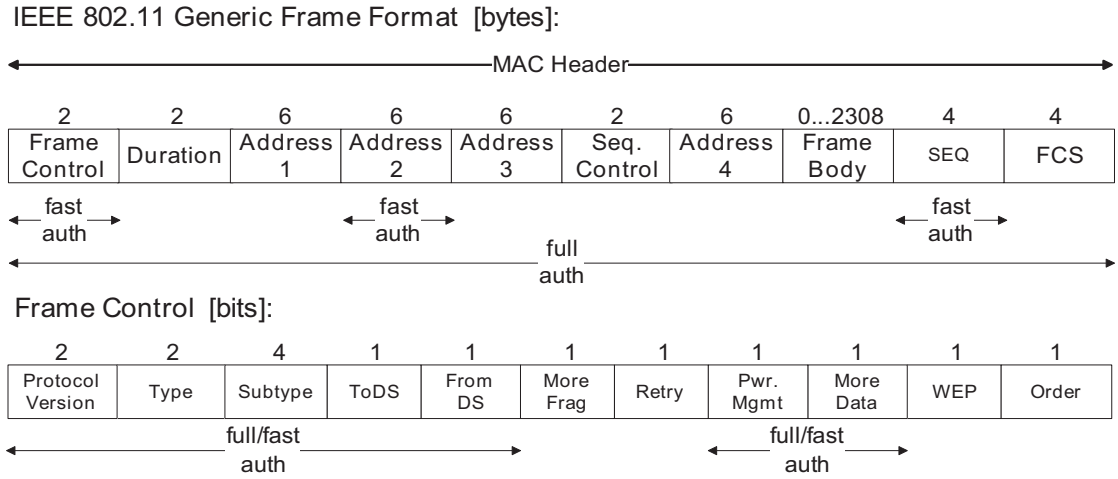

Figure 2. Authentication Modes: authenticated frame fields in *full_auth* and *fast_auth* modes.

this kind of resource-depletion attack, DiscoSec provides a countermeasure based on an association rate control which is described in Section 4.

3.4. Frame Authentication - Variable vs. Fixed Frame Length

The challenge of frame authentication lies in its performance. While the low number of transmitted management frames only marginally increases computational load, the authentication of every data frame significantly stresses performance-limited APs. Consequently, data frame authentication directly impacts the throughput of the wireless connection.

For efficient frame authentication the most influential parameter is the frame's *size*. The Maximum Transmission Unit (MTU) of expected 1500 bytes is overrun in IEEE 802.11 networks and currently, frames up to 3000 byte are transmitted over the wireless channel. The reason lies in the proprietary features of various WLAN cards whose purpose is to increase throughput by frame aggregation. For example the *SuperG* [2] extensions of Atheros utilize the so-called *FastFrame* and *Bursting* techniques. FastFrame exploits the wireless channel more efficiently by increasing the amount of data contained within a single frame, i.e., it minimizes the frame overhead, while Bursting increases a throughput by sending multiple frames within a single transmission opportunity. Although not standardized, these properties are common to various IEEE 802.11 vendors (under different names such as *108G Technology*, *Xtreme G*, *Plus*) and their standardization should be finalized within the *802.11n* standard.

The authentication of such frames can often present a computational burden for performance-limited APs. To be able to support these extensions DiscoSec provides two modes of authentication - the *full_auth* mode for APs with sufficient computational capabilities and the *fast_auth* mode for APs where full authentication of frames would result in a new performance bottleneck. Both modes are based on CMAC-AES authentication.

The *full_auth* mode supports authentication of frames with variable lengths, while the *fast_auth* mode limits the data included in the MAC to a fixed amount of 128 bits (both modes and authenticated frame fields are depicted in Figure 2). In the *fast_auth* mode only certain header fields are authenticated and the frame's payload is omitted from computation. The authenticated fraction of the *fast_auth* mode matches the block size of the AES cipher and can be authenticated within a single AES block computation (the *AT* of 4 byte is included into the calculation, although not depicted in the figure). Accordingly, the authentication is more lightweight and may be performed much faster, for a quantitative comparison see Section 5. On the other side, this presents a tradeoff between security and performance, i.e., a complete authentication vs. higher throughput. While in our opinion a meaningful on-the-fly manipulation of the single bits during wireless transmission is hard to achieve and therefore *fast_auth* is sufficient for wireless communication, the decision on which mode to use is left to the user as a configuration parameter.

3.5. Replay Protection

Frames transmitted over the wireless channel can easily be intercepted and used for replay attacks. To detect the resending of such frames, DiscoSec implements a replay protection by authenticating the frame's sequence numbers. The IEEE 802.11 generic frame format contains a sequence number field which is 16 bits long out of which 4 bits are used for fragment count and only 12 bits as frame sequence.

In order to provide a frame counter sufficient for long sessions and to avoid resynchronization problems, DiscoSec implements an independent 32 bit sequence number (SEQ) field. It is included as an IE in each frame and used for MAC computation within both authentication modes (as shown in Figure 2). The semantics of the legacy sequence number field (Sequence Control) remains therefore unchanged.

The verification is based on accepting a received sequence number within a *window:*

$$seq_{previous} < seq_{current} < seq_{previous} + window + 1$$

This way, the false positive rejections of frames that are retransmitted due to loss or corruption are minimized. Frames containing a value less than the current sequence number are rejected. The magic number for the window length is usually selected around 10 which we also verified by real-world measurements. Since it obviously depends on the wireless environment, it can be changed in DiscoSec's configuration.

4. Resource-depletion Protection

The association procedure in the infrastructure mode of an IEEE 802.11 network utilizes a stateful protocol execution which is prone to DoS attacks, especially to a memory-depletion attack of wireless access points (APs). After receiving an authentication request, an AP reserves memory for a client's connection state. Flooding an AP with a high number of fake authentication requests exhausts the AP's memory which consequently results in faulty operation or even a full device crash. As an example, the results of flooding a commonly used AP are shown in Figure 3. We flooded the AP with 100 authentication requests per second. Shortly after, the AP started freezing for longer periods of time, i.e., it would not respond to *any* frames. This example demonstrates how a simple attack heavily impacts the operation of an AP and motivates the need for resource protection (more details on abusing such performance bottlenecks to prepare more sophisticated attacks is given in [8]).

When using cryptographic primitives on resource-limited devices, such an attack can even be extended to computational resources. In case of DiscoSec, flooding with authentication requests results in initiating many expensive key-exchange computations and thus exhausts AP's computational power. To avoid such attacks, DiscoSec provides simple but effective protection based on a rate limitation.

4.1. DiscoSec's DoS Protection

The protection is based on using *association tokens* which allow only a certain number of authentication requests per second.

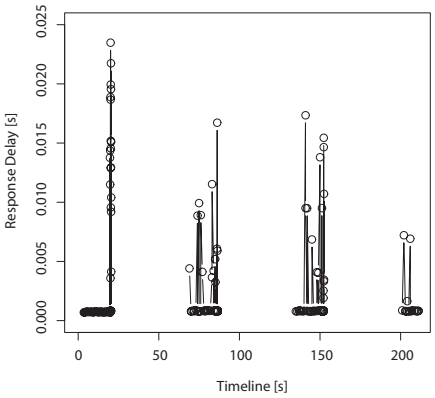

Figure 3. The outcome of a resource-depletion attack - the unprotected AP periodically freezes.

The AP generates a set Θ_t containing association tokens which are 32 bit randomly chosen numbers. They are published at a discrete time t and applied using the following concurrent phases:

- *publication phase* P_t - the AP broadcasts Θ_t to surrounding stations,
- *acceptance phase* A_t - the AP allows associations containing unused tokens from Θ_{t-1}.

Set Θ_t is *repeatedly* sent within the Beacon frames for the duration of the time interval $[t, t+1[$. The Beacon frame additionally contains the *Counter Field* which reports the number of Beacon frames until $t+1$.

The association proceeds as follows: after receiving a Beacon frame, the wireless station *randomly* chooses one token from Θ_t. It waits until the Beacon frame signals the beginning of $t+1$ and then sends the authentication frame containing the chosen token. If the token is unused the AP accepts the station's request and initiates the key exchange.

Using this mechanism, a successful authentication is independent of the time at which a station discovers the tokens. Every station knows the beginning of A_{t+1} and possesses Θ_t, therefore the authentication success probability equals to roughly fair medium access mechanism of IEEE 802.11. Choosing a random token helps legitimate STAs to increase their chance of successful authentication and increases the cost of a successful attack. To ensure that no legitimate station can authenticate, an attacker would have to send all the tokens before the legitimate station sends its requests, and even then, the attacker must succeed for each published Θ_t.

The implementation of this protection is simple as it only requires Θ_{t-1} and Θ_t to be saved at the AP. The length of the

Device	CPU [MHz]	RAM [MB]	Kernel
Cube	MIPS, 324	64	2.6.14
Routerboard	Geode, 266	256	2.6.17
Laptop	Pentium 3, 1400	1024	2.6.17

Table 3. Platforms Used.

Counter Field is 1 byte and Beacons are per default broadcast every 100 ms. The number of tokens within Θ_t depends on the performance characteristics of the dedicated AP.

During our measurements the performance-weakest AP could afford 10 authentications per second, i.e, every second the AP publishes 10 new tokens and accepts 10 tokens. It is important to mention that the tokens are only verified by the AP if the DoS Protection is enabled. On the other hand, when running with DoS Protection, only stations supporting the token mechanism can associate. This tradeoff is the unavoidable consequence of extending the AP's protection functionality.

While the primary objective to protect an AP's resources and assure its operational stability is fulfilled within this version of DiscoSec, more sophisticated techniques to differentiate between legitimate stations and attacker's requests are part of our current research [11].

5. Performance Analysis of DiscoSec

5.1. Evaluated Platforms and Methodology

The selection of platforms for testing DiscoSec's performance focused on hardware discrepancies in order to represent the computational capabilities of broadly available devices. Their hardware characteristics are shown in Table 3.

The performance-weakest device is a 4G AccessCube[2]. The device is from the year 2004 and runs on an architecture other than x86, thus making cross compiling necessary. The other two devices are a Routerboard[3] 230 using Voyage Linux[4] 0.3 and a medium-class Laptop operating in the master mode of the wireless device driver, i.e., offering authentication and association procedures.

To provide insight into all authentication-related delays, DiscoSec was configured to protect both management and data frames. For throughput measurements we generated a continuous stream of UDP packets at various bit rates and under various AP utilizations. For measurements of the key exchange and MAC computation we set the AP utilization to levels of 0 %, 15 %, 30 % and 50 % while monitoring delay as a response variable. The utilization was increased

[2]www.meshcube.org
[3]www.routerboard.com
[4]linux.voyage.hk/

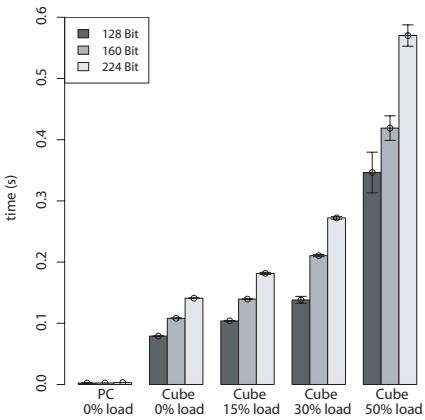

Figure 4. Cost of key exchange using ECDH under various AP loads and different key sizes.

either by using already associated clients sending with the maximal throughput or artificially by additional CPU computations (if frame transmission did not result in a high AP utilization). The measurements were repeated 10 times and depicted results represent the mean with 0.95 confidence intervals.

In the remainder of the paper, the analysis of the key exchange and frame authentication is given for the performance-weakest device AP_{Cube}, while the final network throughput results of DiscoSec are provided for all three devices.

5.2. Key Exchange and Frame Authentication

Before going into details of the delays introduced by the key exchange, we briefly mention the state-of-the art delays imposed by the IEEE 802.11i security standard. For mutual identity authentication the security standard requires an Authentication Server that undertakes the shared secret computation instead of the AP. From [9] measurements show that the IEEE 802.11i delay imposed by the key exchange using mutual authentication (e.g., EAP-TLS) varies between $\approx$ 300 ms and 4 s, depending on different platforms and various implementations of the standard. Concerning DiscoSec's key exchange, Figure 4 depicts delays using different key sizes. The 128 bit elliptic curve key takes only $\approx$ 79 ms on the performance weakest device. The varying AP utilization does not influence the key exchange much and at 50% utilization the 128 bit key exchange remains under 400 ms. Clearly, longer keys increase the computational time, nevertheless even the key exchange using 224 bit keys (equivalent to 2048 bit RSA public key) remains under 600 ms.

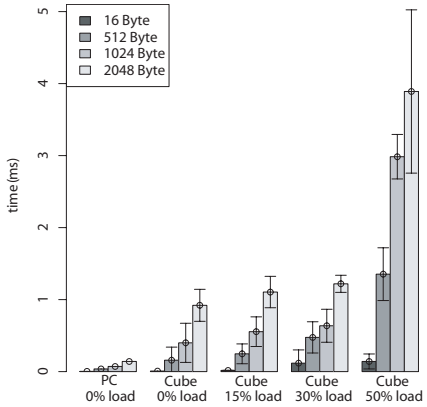

Figure 5. Computation of MACs for various message sizes.

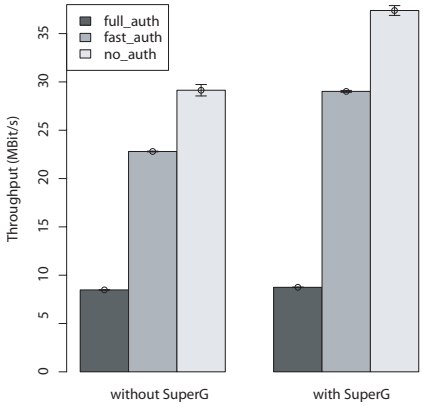

Figure 6. Throughput using standard transmission vs. SuperG throughput enhancing features.

In contrast to the shared secret which is generated only once at the beginning of the session, a MAC is calculated for each transmitted management and data frame, implicitly influencing the connection throughput.In order to evaluate the measurement results given in Figure 5, it is important to consider the impact of the MAC computation on a frame transmission. The maximum throughput of the AP_{cube} is around 29 Mbit/s using the plain driver without any extensions. This means that every $\approx 433\,\mu s$ a packet is transmitted. As a back-of-the-envelope calculation, if the MAC generation takes just as long, which includes the overhead imposed from the driver, then the data rate will halve. On the tested hardware, the processing of 1024 bytes of data already takes $\approx 400\,\mu s$, hence not leaving much space. Nevertheless, the same figure shows that the computation time remains stable and varies less (given by the interval length of the confidence intervals) for all keysizes if load is under 50 %, otherwise delay and its variance dramatically increase exhibiting the device's computational limits which may result in unpredictable behavior. For this reason, the *fast_auth* mode becomes inevitable. When using *fast_auth* the computation of authenticated data equals 16 bytes which requires an AES key length of 128 bit and does not exceed $\approx 150\,\mu s$ even at 50 % CPU load. This significantly relieves the computational burden resulting in much higher throughput as shown in the next subsection.

5.3. DiscoSec Featuring SuperG

The throughput comparison between plain and *SuperG* enhanced transmission is depicted in Figure 6 for three different configurations: *no_auth*, *fast_auth*, *full_auth*. Considering the frame transmission without SuperG extensions, the *no_auth* bar denotes the maximum possible unauthenticated throughput of 29 Mbit/s equal to transmission without DiscoSec which serves as reference. Using *full_auth*, the complete IEEE 802.11 frame is authenticated. This security feature is the most computationally demanding and AP_{Cube} offers only 8 Mbit/s throughput.

Using the *fast_auth* mode the AP relaxes the computational requirements and achieves data rates of $\approx$ 23 Mbit/s (78% of unauthenticated throughput). This scenario shows the importance of providing *fast_auth* as a tradeoff parameter for performance-limited APs. Enabling the SuperG extensions (FastFrame and Bursting) leads to shorter transmission delays and larger frames, increasing the *no_auth* throughput to 37 Mbit/s. But more importantly, since SuperG does not impact the IEEE 802.11 header, the computational effort of *fast_auth* mode is equal to a transmission without SuperG features although more data is being transmitted. Therefore, even on a very performance-limited device like AP_{Cube}, using SuperG with *fast_auth* authenticated transmission is at $\approx$ 29 Mbit/s (78%).

5.4. Overall Throughput

Until now, the presented analysis focused only on the weakest measured device. By using the *fast_auth* mode, performance degradation can be mitigated, though not eliminated. The trend of modern APs aims at offloading computations of cryptographic primitives, especially symmetric ciphers to specialized hardware. For example, the new generation of Geode CPUs features a hardware implementation of the 128 bit AES cipher and a true random number generator.

In our measurements we analyzed an older version of the Geode CPU within a *Routerboard RB230.* It is a multifunctional device running at only 266 MHz (less than AP_{Cube}), without any special-purpose hardware for faster computations. It uses a Geode x86 SC1100 processor, equivalent to

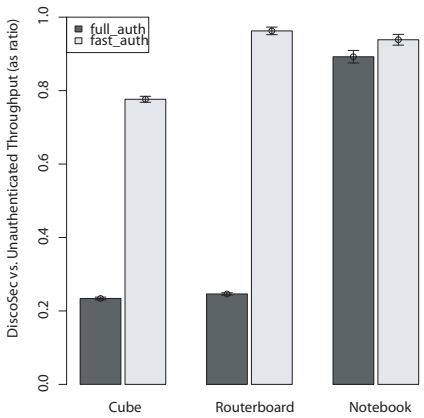

Figure 7. Throughput analysis of all tested platforms.

the Intel Pentium MMX architecture. Although its low CPU clock frequency does not allow for much faster computation, in *fast_auth* it achieves 97 % of the possible throughput (see Figure 7). The modest-looking 2-3% throughput increase of *full_auth* mode compared with AP_{Cube} implies that Routerboard succeeds in authenticating $\approx$ 1 Mbit/s more data. Using hardware capabilities of an older Notebook running as an AP, it shows an exemplary authentication performance. The throughput of both, *full_auth* and *fast_auth* mode is at 89 % and 94 %, respectively.

To summarize, this section provided an overview of what to expect of the network throughput using computationally-limited devices. It demonstrates the importance of identifying security vs. performance tradeoffs which in turn may smoothen throughput differences among heterogeneous hardware platforms.

6. Related Work

Concerning attacks based on unauthenticated management and data frames, [3] demonstrates their devastating effect on IEEE 802.11 networks. Based on the same vulnerabilities, in [6] various attacks are successfully mounted even against the new security standard IEEE 802.11i. While the empirical demonstration is a frequently used method to illustrate the problem of link-layer security, protection against such attacks prevalently remains conceptual. Only [4] discusses the implementation issues of a proposed solution. The authors employ two protocols, SIAP and SLAP, to establish a secure association utilizing public key infrastructure. While their solution offers encryption, it also modifies the IEEE 802.11 state machine and requires a SIAP server.

In commercial products, Cisco offers a feature called Management Frame Protection (MFP), but there is regrettably no detailed information other than white papers [15]. Interestingly, MFP does not seem to be a client-side supported feature, and thus only protects APs, while clients remain vulnerable to management frame attacks.

The initial idea of the authentication mechanism used in DiscoSec was described in [10]. In this paper we significantly improved and implemented the concept, and analyzed its performance. To the best of our knowledge, DiscoSec is the first solution with a design supporting the IEEE 802.11 state machine, extensively tested on performance-limited hardware and available for use on present devices.

7. Conclusion

This work described DiscoSec, a lightweight authentication protocol designed to protect WLANs against the most prominent attacks based on resource-depletion and injection of impersonated management and data frames. DiscoSec followed the idea of "patching", i.e., providing a small, effective and easily applicable solution to a variety of devices. During the development of DiscoSec, we came across various design and implementation decisions such as providing a DoS-resilient key exchange, efficient authentication, support for throughput-increasing features like SuperG, and the usage of widely accepted cryptographic primitives. Most of these decisions are offered as configuration parameters to facilitate the balance between security and performance tradeoffs.

Using real-world measurements, we demonstrated that even a performance-limited device achieves 78 % of the maximum throughput, while using a more powerful device the price paid is only a 9-11 % throughput decrease for full and fast authentication, respectively.

8. Acknowledgement

We gratefully acknowledge the *madwifi.org* project whose device driver implementation was used as basis for DiscoSec implementation, and the anonymous reviewers for their valuable comments.

References

[1] IEEE 802.11i/D10.0. Security Enhancements, Amendment 6 to IEEE Standard for Information Technology. IEEE Standard, April 2004.

[2] Atheros. SuperG – Maximizing Wireless Performance . Available at www.super-g.com, (last accessed 12.10.2007).

[3] J. Bellardo and S. Savage. 802.11 Denial-of-Service attacks: Real Vulnerabilities and Practical Solutions.

In *Proceedings of the USENIX Security Symposium*, pages 15–28, August 2003.

[4] D. Faria and D. Cheriton. DoS and Authentication in Wireless Public Access Networks. In *Proceedings of the 2004 ACM Workshop on Wireless Security*, pages 47–56, September 2002.

[5] C. He and J. C. Mitchell. Analysis of the 802.11i 4-way handshake. In *Proceedings of the 2004 ACM Workshop on Wireless Security*, pages 43–50, October 2004.

[6] C. He and J. C. Mitchell. Security Analysis and Improvements for IEEE 802.11i. In *Proceedings of the 12th Annual Network and Distributed System Security Symposium (NDSS'05)*, pages 90–110, February 2005.

[7] IETF. Diffie-Hellman Key Agreement Method. RFC 2631, 1999.

[8] I. Martinovic, F. A. Zdarsky, A. Bachorek, C. Jung, and J. B. Schmitt. Phishing in the Wireless: Implementation and Analysis. In *Proceedings of the 22nd IFIP International Information Security Conference (SEC 2007)*, pages 145–156, May 2007.

[9] I. Martinovic, F. A. Zdarsky, A. Bachorek, and J. B. Schmitt. Introduction of IEEE 802.11i and Measuring its Security vs. Performance Tradeoff. In *Proceedings of the 13th European Wireless Conference*, April 2007.

[10] I. Martinovic, F. A. Zdarsky, and J. B. Schmitt. On the Way to IEEE 802.11 DoS Resilience. In *Proceedings of the Workshop on Security and Privacy in Mobile and Wireless Networking*, May 2006.

[11] I. Martinovic, F. A. Zdarsky, M. Wilhelm, C. Wegmann, and J. B. Schmitt. Wireless Client Puzzles in IEEE 802.11 Networks: Security by Wireless. In *Proc. ACM Conference on Wireless Network Security (WiSec 2008)*, pages 36–45, March 2008.

[12] NIST Special Publication 800-38B. Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication, May 2005.

[13] SECG. Elliptic Curve Cryptography, Standards for Efficient Cryptographic Group. Available at www.secg.org/collateral/sec2.pdf, (last accessed 08.10.2007).

[14] J. Song, R. Poovendran, J. Lee, and T. Iwata. The AES-CMAC Algorithm. RFC 4493 (Informational), June 2006.

[15] www.cisco.com; Document ID: 82196. Infrastructure Management Frame Protection (MFP) with WLC and LAP Configuration Example, (last accessed 28.05.2007).