

Also available as: [PDF](#) and [Postscript](#).

Using the Fluhrer, Mantin, and Shamir Attack to Break WEP

August 6, 2001

[Adam Stubblefield](#)¹ John Ioannidis [Aviel D. Rubin](#)
Rice University AT&T Labs - Research, Florham Park, NJ
astubble@cs.rice.edu {ji,rubin}@research.att.com

Abstract:

We implemented an attack against WEP, the link-layer security protocol for 802.11 networks. The attack was described in a recent paper by Fluhrer, Mantin, and Shamir. With our implementation, and permission of the network administrator, we were able to recover the 128 bit secret key used in a production network, with a passive attack. The WEP standard uses RC4 IVs improperly, and the attack exploits this design failure. This paper describes the attack, how we implemented it, and some optimizations to make the attack more efficient. We conclude that 802.11 WEP is totally insecure, and we provide some recommendations.

1. Introduction

Wireless networking has taken off, due in large part to the availability of the 802.11 standard. While another standard, Bluetooth, is also gaining in popularity, the longer range and higher speeds achieved by 802.11 make it the protocol of choice for wireless LANs. Office buildings, conferences, and even many residences now offer 802.11 connectivity. The PC cards that are most often used in these networks provide a security protocol called Wired Equivalent Privacy (WEP).

WEP is easy to administer. The device using the 802.11 card is configured with a key, that in practice usually consists of a password or a key derived from a password. The same key is deployed on all devices, including the access points. The idea is to protect the wireless communication from devices that do not know the key.

Borisov, Goldberg and Wagner demonstrated some security flaws in WEP [1]. They explained that WEP fails to specify how IVs for RC4 are specified. Several PC cards reset IVs to zero every time they are initialized, and then increment them by one for every use. This results in high likelihood that keystreams will be reused, leading to simple cryptanalytic attacks against the cipher, and decryption of message traffic. The authors verified this experimentally and describe other weaknesses as well. For example, the space from which IVs are chosen is too small, virtually guaranteeing reuse, and leading to the same cryptanalytic attacks just described. The paper also shows that message authentication in WEP is broken.

Fluhrer, Mantin, and Shamir describe a passive ciphertext-only attack against RC4 as used in WEP [2]. The attack exploits the method in which the standard describes using IVs for the RC4 stream cipher. In their paper, the authors state, *Note that we have not attempted to attack an actual WEP connection, and hence do not claim that WEP is actually vulnerable to this attack.* Based on the description in their paper, we successfully implemented the attack, proving that WEP is in fact completely vulnerable. The purpose of this paper is to describe our implementation, along with some enhancements to improve the performance of the attack.

2. Overview of the WEP attack

In this section we present an overview of the WEP protocol and review briefly how the attack of Fluhrer, Mantin, and Shamir can be applied to WEP. For a detailed description of WEP we refer the reader to the official 802.11 standard [4].

Encryption in WEP uses a secret key, k , shared between an access point and a mobile node. To compute a WEP frame, the plaintext frame data, M , is first concatenated with its checksum $c(M)$, to produce $M \cdot c(M)$ (where \cdot denotes concatenation). Then, a per packet initialization vector (IV) is prepended to the secret key, k , to create the packet key, $IV \cdot k$. The RC4 stream cipher is then initialized using this packet key, and the output bytes of the cipher are exclusive-ored (denoted $+$) with the checksummed plaintext to generate the ciphertext:

$$C = (M \cdot c(M)) + RC4(IV \cdot k)$$

Since the IV is transmitted in the clear, if we are able to correctly guess the first byte of M , we can apply the known IV attack of section 7.1 and appendix A of Fluhrer, Mantin, and Shamir [2]. We give a brief description of the attack here, but refer the reader to that paper for the details. We defer the discussion of how to guess the first byte of the plaintext to Section 3.

To begin, we must describe the structure of the RC4 stream cipher (a full description can be found in [6]). RC4 consists of two parts, a key scheduling algorithm and an output generator. In WEP, the key scheduling algorithm uses either a 64-bit packet key (40-bit secret key plus 24-bit IV) or a 128-bit key (104-bit secret key plus 24-bit IV) to set up the RC4 state array, S , which is a permutation of $\{0, \dots, 255\}$. The output generator uses the state array S as well as two counters, i and j , to create a pseudorandom sequence. The Fluhrer, Mantin, and Shamir known IV attack utilizes the fact that, in some cases, knowledge of the IV and the first output byte leaks information about the key bytes. They refer to these key-leaking cases as *resolved*. It is simple to test whether a particular packet provides an IV and output byte that result in a resolved condition, though we refer the reader to the Fluhrer *et. al.* paper for the full explanation. By looking at a number of these resolved cases, we can see a bias toward the true key bytes.

3. Implementation

In implementing this attack, we had three goals. First and foremost, we wanted to verify that the attack could work in the real world. Second, we were interested in how cheaply and easily the attack could be launched. Lastly, we wanted to see what sort of improvements could be made to both the general RC4 attack and the WEP attack in particular. In this section we report on our success at the first two goals,

while reserving discussion about attack optimizations to Section [4](#).

3.1 Simulating the Attack

Before trying to break WEP, we created a simulation of the RC4 attack to both verify our understanding of the weakness and to gather information about how many resolved packets we could expect would be required when mounting the actual attack. The coding of the simulated attack took under two hours, including a few optimizations. The simulation showed that the attack was always able to recover the full key when given 256 probable resolved cases.² We also observed that although 60 resolved cases (the number recommended in the Fluhrer *et. al.* paper) was usually enough to determine a key byte, there were instances in which more cases were required. Because at this point we had not thoroughly investigated how accurately we would be able to determine the first output byte of the RC4 pseudorandom sequence, we also simulated the effect that sometimes guessing wrong would have on the attack. We were pleased to see that as long as the number of incorrect guess was kept small, the correct key byte would still be returned correctly, though sometimes more resolved cases were needed.

3.2 Capturing the Packets

Surprisingly, capturing WEP encrypted packets off of our wireless network proved to be the most difficult part of the attack. There are a number of commercial software programs that are able to both capture and decode 802.11 packets, such as NAI's "Sniffer" and Wildpacket's "AiroPeek," though both products cost thousands of dollars. Because we wanted to show that the attack could be done by an adversary with limited resources, we ordered a \$100 Linksys wireless card, based on the Intersil Prism II chipset. We made this choice because the Prism II allows much of its computation to be completed in software and because there was a Linux driver available that claimed to be able to grab raw WEP encrypted packets. Though we did not know it at the time, this chipset has been used by others to mount dictionary and brute force attacks against WEP.³

After many hardware headaches, we were able to make the card work in Linux by using both the `linux-wlan-ng prism2` driver⁴ and a modified version of Tim Newsham's patch to re-enable raw packet monitoring.⁵ We were then able to use a modified version of the packet sniffer `ethereal`⁶ to capture raw WEP encrypted packets and to decode the data necessary for our attack tool.

Even with the hardware and software problems we ran into, from the time that we first decided to look at this problem, it took less than a week for the the card to be ordered and shipped, the test-bed to be set up, the problems to be debugged, and a full key to be recovered.

3.3 Mounting the Attack

The last piece in actually mounting the attack was determining the true value of the first plaintext byte of each packet, so that we could infer the first byte of the pseudorandom sequence from the first ciphertext byte. We originally looked at `tcpdump` output of decrypted traffic (using a correctly keyed card⁷), and were planning on using packet length to differentiate between ARP and IP traffic (both of which have well known first bytes in their headers) as these were by far the two most common types of traffic on our network. After implementing this, however, we discovered that the attack didn't seem to work. We then tried hand decrypting packets to determine whether `tcpdump` was working correctly and discovered that an additional 802.2 encapsulation header is added for both ARP and IP traffic.⁸ This discovery actually made the attack even easier, as all packets would now have the same first plaintext byte (0xAA, the SNAP designation). Note that a correctly keyed card is not needed for the attack; we

simply used one to design the attack.

Although our actual attack used the improvements discussed in the next section, we present an outline of how a naive attack could work here. It is interesting to note that even this baseline version of the attack would still be successful in a short period of time (a day or two at most) and with an even smaller amount of computation when compared to the improved implementation, assuming that the wireless network in question had a reasonable amount of traffic.

To begin, we collected a large number of packets from our wireless network. To speed the process up for some of our experiments late at night when network volume was low, we artificially increased the load on the wireless network by ping flooding a wireless node. (We could have waited until more traffic was created; this is not an active attack.) Because we are able to predict the value of the first byte of any plaintext, the fact that we changed the makeup of the network traffic did not affect these experiments. In looking at the IVs of these collected packets, we discovered that the wireless cards use a simple counter to compute the IV, wherein the first byte is incremented first. In section A.1 of Fluhrer *et. al.*, the authors postulate that 4,000,000 packets would be sufficient with this type of counter, we found the number to be between 5,000,000 and 6,000,000 for our key. This number is still not unreasonable, as we were able to collect that many packets in a few hours on a partially loaded network.

4. Improving the attack

In this section we discuss several modifications that can be made to improve the performance of the key recovery attack on WEP. While not necessary for the compromise to be effective, they can decrease both time and space requirements for an attacker.

4.1 Choosing IVs

In the baseline attack (the one described in Fluhrer *et. al.*), only IVs of a particular form are considered. However, we found that there are other IVs that can result in a resolved state, and that testing all IVs instead of only the subset suggested by the Fluhrer *et. al.* paper can be done in parallel with receiving packets. This conclusion was verified by Adi Shamir [7], who also noted that these packets appear more often for higher key bytes.

4.2 Guessing Early Key Bytes

As the Fluhrer, Mantin, and Shamir attack works by building on previously discovered key bytes, recovering early key bytes is critical. There are two approaches that we tried both separately and together. The first utilized the way that the IVs were generated, namely that we would receive packets that resolved for lots of different key bytes before necessarily receiving enough resolving packets to predict the early key bytes.⁹ We would therefore use the resolving cases that we had received to narrow down the possibilities for the early key bytes. We were then able to test candidate keys by determining if the WEP checksum on a decrypted packet turned out correctly.

The second approach exploited the poor key management available in WEP implementations. Since WEP keys have to be entered manually, we assumed that instead of giving clients a long string of hex digits, a user memorable passphrase would be used. After examining the test wireless cards at our disposal, we determined that the user-memorable passphrase is simply used raw as the key (i.e. the

ASCII is used; no hashing is done). Although hashing does not protect against a dictionary attack, it would have helped in this circumstance, as we were able to determine directly whether each key byte was likely to be part of a user memorable passphrase by checking whether the byte value corresponded to an ASCII letter, number, or punctuation symbol.

4.3 Special Resolved Cases

As Shamir pointed out to us, there are cases when a resolved case can provide an even better indication as to a particular key byte. The following paragraph assumes that the reader is familiar with the Fluhrer *et. al.* attack and uses the same notation as that paper.

If there is a duplication among the three values at positions 1 , x , and $x+y$ (i.e. there are only two distinct values), then the probability that these positions in the S permutation remain unchanged jumps from e^{-3} (5%) to e^{-2} (13%). We can thus treat the evidence from these cases as about three times more convincing as a standard resolved case.

5. Conclusions

We were able to implement the attack described by Fluhrer *et. al.* in several hours. It then took a few days to figure out which tools to use and what equipment to buy to successfully read keys off of 802.11 wireless networks. Our attack used off of the shelf hardware and software, and the only piece we provided was the implementation of the RC4 attack, along with some optimizations. We believe that we have demonstrated the ultimate break of WEP, which is the recovery of the secret key by observation of traffic.

Given this attack, we believe that 802.11 networks should be viewed as insecure. We recommend the following for people using such wireless networks.

- Assume that the link layer offers no security.
- Use higher-level security mechanisms such as IPsec [3] and SSH [8] for security, instead of relying on WEP.
- Treat all systems that are connected via 802.11 as external. Place all access points outside the firewall.
- Assume that anyone within physical range can communicate on the network as a valid user. Keep in mind that an adversary may utilize a sophisticated antenna with much longer range than found on a typical 802.11 PC card.

The experience with WEP shows that it is difficult to get security right. Flaws at every level, including protocol design, implementation, and deployment, can render a system completely vulnerable. Once a flawed system is popular enough to become a target, it is usually a short time before the system is defeated in the field.

Acknowledgments

We thank Bill Aiello, Steve Bellovin, Bob Miller, Adi Shamir, Dave Wagner, and Dan Wallach for helpful discussions.

We informed Stuart Kerry, the 802.11 Working Group Chair, that we successfully implemented the Fluhrer, et al. attack. Stuart replied that the 802.11 Working Group is in the process of revising the security, among other aspects, of the standard and appreciates this line of work as valuable input for developing robust technical specifications.

Bibliography

- 1 BORISOV, N., GOLDBERG, I., AND WAGNER, D.
Intercepting mobile communications: The insecurity of 802.11.
MOBICOM 2001 (2001).
- 2 FLUHRER, S., MANTIN, I., AND SHAMIR, A.
Weaknesses in the key scheduling algorithm of RC4.
Eighth Annual Workshop on Selected Areas in Cryptography (August 2001).
- 3 KENT, S., AND ATKINSON, R.
Security architecture for the Internet protocol.
Request for Comments 2401, Internet Engineering Task Force, November 1998.
- 4 L. M. S. C. OF THE IEEE COMPUTER SOCIETY.
Wireless LAN medium access control (MAC) and physical layer (PHY) specifications.
IEEE Standard 802.11, 1999 Edition (1999).
- 5 POSTEL, J., AND REYNOLDS, J. K.
Standard for the transmission of IP datagrams over IEEE 802 networks.
Request for Comments 1042, Internet Engineering Task Force, Feb. 1988.
- 6 SCHNEIER, B.
Applied Cryptography - Protocols, Algorithms, and Source Code in C.
John Wiley & Sons, Inc., 1994.
- 7 SHAMIR, A.
Personal communications, 2001.
- 8 YLONEN, T.
SSH - secure login connections over the Internet.
USENIX Security Conference VI (1996), 37-42.

Footnotes

... Stubblefield¹

Research done while a summer intern at AT&T Labs

... cases.²

Cases corresponding to IVs of the form (B+3, 255, N) as in the Fluhrer *et al.* paper.

... WEP.³

See Blackhat `01 presentation at

http://www.lava.net/~newsham/wlan/WEP_password_cracker.ppt

... driver⁴

Available from <http://www.linux-wlan.com/>

... monitoring⁵

Available from <http://www.lava.net/~newsham/wlan/>

... ethereal⁶

Available from <http://www.ethereal.com/>

... card⁷

Note that a correctly keyed card is not needed; we simply used one to design the attack.

... traffic⁸

We eventually traced this back to RFC 1042 [[5](#)].

... bytes⁹

See Figure 6 of Fluhrer *et. al.*; resolved cases are much more likely to occur for later key bytes.

© AT&T Labs, 2001.