

A DOS ATTACK AGAINST THE INTEGRITY-LESS ESP (IPSEC)

Ventzislav Nikov
Philips TASS and AppTech
Leuven, Belgium

Keywords: IPsec (ESP) Standard, Denial of Service Attack.

Abstract: This paper describes a new practical DoS attack that can be mounted against the “encryption-only” configuration (i.e. without authenticated integrity) of ESP as allowed by IPsec. This finding can serve as a strong argument to convince those in charge of the IPsec standardization to improve it by banning the “encryption-only” configuration from the standard.

1 INTRODUCTION

As illustrated in (M. Bellare, 2004; Bellare, 1996; N. Borisov, 2001; Canvel et al., 2003; C. McCubbin, 2000; T. Yu, 2004), it is, today, common knowledge in the cryptographic community that encryption without authenticated integrity opens the door to various attacks. More recently, the authors of (K. Paterson, 2005) illustrate this argument on the encryption-only configuration allowed by IPsec for the secure payload encapsulation (ESP) (Kent, 2005).

IPsec encryption-only configuration was shown to be exposed to variety of “initialization vector attacks” (C. McCubbin, 2000). The authors of (C. McCubbin, 2000) analyze the security risks posed on encryption-only IPsec when it is used as intermediate layer of the protocol stack. But their analysis has a limited scope - just the implication on the next (possible) upper-layer protocol (e.g. TCP, UDP, IP). We develop further their idea of properly modifying certain fields in the payloads of the whole protocol stack (e.g. IP/UDP/RTP or IP/TCP). In this way we build a more powerful attack.

The authors of (K. Paterson, 2005) have exposed even more serious weakness of the encryption-only configuration on IPsec. It has been shown that the “encryption-only” configuration is subject to a variety of ciphertext-only attacks, i.e. revealing the encryption key. The attack consists of two phases: in the first phase the attacker modifies certain fields in the first few blocks of the ciphertext such that the upper-layer protocol submits either an error message or the

decryption of the modified datagram directly to the attacker’s machine. The second phase proceeds with further recovery of the plaintext. The attacks presented in (K. Paterson, 2005) are efficient and have been proven practical by the authors, i.e. against an implementation of the IPsec stack in the Linux kernel. Our attack is completely different, since it is not a ciphertext-only attack, the attacker goal is to mount with minimum efforts a Denial of Service Attack (DoS) against the IPsec gateway.

One of the purposes of (K. Paterson, 2005) was to help closing the gap between the currently known cryptography theory and the current practices of standard bodies, implementers and users. Among others, providing practical efficient attacks against existing (or in development) standards or products helps the various communities (theoreticians and practitioners) to work together to establish cryptographically sound but workable standards and products.

The main contribution of this paper is to present yet another attack against the “encryption-only” configuration for ESP. This attack has been unearthed years ago but hasn’t lead to a publication because at that time it was perceived as of no theoretical interest. However, following the example of (K. Paterson, 2005), we hope to contribute with the publication of our finding to help improving the IPsec standard (and the corresponding implementations).

Our attack is a very practical Denial of Service (DoS) attack against the IPsec gateway providers. It only requires a minimal effort from the attacker while having a devastating potential impact on the

gateways performance. We think that this attack is an even stronger and convincing argument than those already published to support the banishment of the “encryption-only” configuration of ESP.

Recall that the ESP configuration is not integrity protected, thus the option provided in the ESP for replay protection makes no sense since it requires obviously data integrity. We stress that our attack is not just a replay attack since the attacker can deviate/multiply the flow to different targets. We also assume, as in (C. McCubbin, 2000; K. Paterson, 2005), that the datagrams are not checked after IPsec processing to see if the correct IPsec policies were applied; this is the case in the Linux kernel implementation (as pointed out in (K. Paterson, 2005)).

Although it is commonly known that encryption without authenticated integrity is dangerous and more specifically that ESP without data authentication leads to all kind of problems there are still security IPsec specialists that claim “when IPsec gateway is protected in addition with a stateful firewall which does filter the packets, most of the described above attacks will not work.” We think that such an approach is just an attempt to shift the problem, instead of eradicate it.

The rest of this paper is organized as follows: After an introduction to the IPsec ESP in Sect. 2, we present the DoS attack itself in Sect. 3.

2 PRELIMINARIES

2.1 Ipvsec

IPsec provides security at the level of the IP layer. The secure encapsulation of the payload of IP datagrams is part of the IPsec standard and is provided by ESP (Encapsulating Security Payload, (Kent, 2005)). ESP provides integrity, authentication and encryption of the IP datagrams (as an option, it is possible to add a replay protection).

The encapsulation is connectionless, i.e. it is performed on a per-datagram basis. One common use of IPsec is to build virtual private networks where IPsec is configured to use ESP in tunnel mode. The way ESP modifies the datagrams in tunnel mode is shown in the Appendix Fig. 4.

Fig. 4 shows that given a datagram to protect, ESP creates a new datagram made of, in sequence, a new IP header, an ESP header, the original datagram (IP header and Inner Payload), an ESP trailer and an ESP authentication payload. The detailed description of the ESP header, trailer and authentication payload is not provided since they are outside of the scope of our attack.

IPsec supports different encryption algorithms: AES (J. Daemen, 2002; NIST, 2002) is the most commonly used but other block ciphers as DES (NIST, 1977), 3DES (NIST, 1999), CAST128 (Adams, 1997b; Adams, 1997a), RC5 (Rivest, 1994), IDEA (X. Lai, 1990), and Blowfish (Schneier, 1993) are also allowed. All block ciphers are used in CBC mode (R. Pereira, 1998; S. Frankel, 2003). Additionally, AES in CTR mode is a valid alternative (Housley, 2004) but in this case the standard states that “*AES-CTR implementations MUST employ a non-NUL ESP authentication method, since it is trivial to forge an AES-CTR ciphertext*”. As a consequence, the hereafter described attack targets only the block ciphers in CBC mode. The attack depends only on the block size n of the cipher, where $n = 64$ for all block ciphers except for AES where $n = 128, 192, 256$.

2.2 Cbc Mode

If the to be transmitted data after ESP padding is made of q blocks, if the plaintext blocks are denoted by P_1, P_2, \dots, P_q , and if $e_K(\cdot)$ ($d_K(\cdot)$) denotes the encryption (decryption) of blocks using an n -bit key K then the CBC mode (NIST, 1980) works as follows. After a random n -bit IV (Initialization Vector) is generated, the ciphertext blocks are computed according to the equations:

$$C_0 = IV, \quad C_i = e_K(C_{i-1} \oplus P_i), \quad (1 \leq i \leq q).$$

At the receiver side the plaintext is recovered according to the equations:

$$P_i = C_{i-1} \oplus d_K(C_i), \quad (1 \leq i \leq q).$$

As pointed out in (A. Menezes, 1996; K. Paterson, 2005) a well known weakness of the CBC mode is the “*bit flipping attack*”. An attacker can flip (invert) a specific bit in the ciphertext block C_{i-1} , then this specific bit in the recovered plaintext block P_i is also flipped (since $P_i = C_{i-1} \oplus d_K(C_i)$). This allows an attacker to introduce controlled changes into the recovered plaintext block P_i , but the previous block P_{i-1} is randomized. Hence the integrity of the IV in CBC should be protected otherwise uncontrolled change on the first recovered plaintext block P_1 is possible.

2.3 Ip, TCP, Udp and RTP Datagrams

The hereafter presented attack depends on the way the IP stack is structured. The presentation is limited to IPv4 headers as specified in (Tanenbaum, 2002).

The layout of the IP header is illustrated in the Appendix Fig. 5. The potential targets of the attack are the source and the destination IP addresses (32 bits

each) and the header checksum (16 bits). The first 80 bits are disregarded and the last optional fields are assumed to be absent. Note that, in the absence of the optional fields, there is no padding.

There are two commonly used protocols in the transport layer, a connectionless one - the User Datagram Protocol (UDP) (Postel, 1980) and a connection-oriented one - the Transmission Control Protocol (TCP) (Postel, 1981). Thus an IP datagram may contain either UDP or TCP embedded datagrams. This gives additional targets for the attack.

The TCP header (Postel, 1981) is illustrated in the Appendix Fig. 6. The source and the destination ports (16 bits each), the sequence and acknowledgement number (32 bits each), which together form the first 96 bits of the TCP header as well as the checksum located between 128-th and 144-th bits are additional important targets of the attack.

The UDP header (Postel, 1980) is illustrated in Fig. 1. In this case the source and destination ports as well as the checksum (16 bits each) are additional targets of the attack.

Source Port	Destination Port
Length	Checksum
Data...	

Figure 1: UDP Segment Format.

Sometimes, another protocol is encapsulated in the UDP datagrams, e.g. the Real-Time Transport Protocol (RTP) (Schulzrinne et al., 2003). A description of the RTP header is given in the Appendix Fig. 7, giving yet another target for the attack: the sequence number field located between 16-th and 32-th bits.

This list of protocols and targets is not exhaustive, the attack is applicable to other protocols. The rationale behind the choice of those targets is that these are the fields that are consulted or manipulated by the IP gateways in an IP network.

2.4 Internet Checksum

The (internet) checksum (R. Braden, 1988) is used to discover errors while datagrams are transmitted. The targets chosen in this paper to illustrate the attack: the IP header, the TCP and UDP datagrams, all use checksums to protect their content. The purpose of the internet checksum is to provide an efficient protection against transmission errors but not to provide cryptographic integrity protection of the content.

The outline of the internet checksum is very simple: the data is presented as a sequence of 16 bit integers, then the 1's complement sum of these integers is computed and the checksum is the 1's complement of this sum.

The way this checksum is computed allows the fast incremental update of the checksum when the data over which it is computed is changed. For example, to update the checksum C in the IP header m to the new IP header m' , but without recomputing it over the entire new IP header, the updated checksum C' can be easily computed as $C' = C + (m - m')$, as shown in (T. Mallory, 1990). Note that, to compute the new checksum, one needs to know $m - m'$, but neither m , m' nor the old checksum C are to be known.

3 A DOS ATTACK

3.1 Fields Manipulation

The choice made by some implementers or users for the "encryption-only" ESP configuration (i.e. without the "costly" authenticated integrity) is based on the false belief that confidentiality protection together with the structure of IP, TCP (or IP, UDP, RTP) datagrams is enough to detect any data change (malicious or not). This amounts at making the hypothesis that any modification to the encrypted datagrams will be detected during the parsing of the decrypted datagrams (i.e. they are supposedly malformed) or during the verification of the checksums.

This is not true; an attacker can, in fact very easily, modify the encrypted (in CBC mode) datagrams in such a way that they are still acceptable for the embedded protocols such as IP, TCP, UDP, RTP. In the attack presented here, the IV is changed in such a way that the first decrypted block of plaintext P'_1 is changed in a predictable way (see Sect. 2.2). More precisely, the IV is replaced with a new one which differs only in the positions to be altered in the first block of plaintext.

Consider a confidentiality protection with a block cipher with length $n = 128$ (see Fig. 2). In this case, P_1 , the first block of data, contains the IP header checksum and the source IP address (see Sect. 2.3). Now an attacker can change the source IP address to whatever she/he wants, then using the difference between the old and the new source IP addresses she/he can compute the new correct checksum (see Sect. 2.4) and finally she/he can compute the new IV corresponding to the altered ESP datagram (the rest of the data remains the same). When a gateway receives such a datagram, it accepts it as a valid datagram (i.e. the IP header is valid) and forwards it further to the corresponding destination address.




Figure 2: $n=128$.

In the case where $n = 192$ (see the Appendix Fig. 8), the attacker has more choices. Assuming that TCP/IP (or UDP/IP) datagrams are sent, the attacker can also change the destination IP address, the source and destination ports. But since the TCP (UDP) checksum is in block P_2 , the modifications brought by the attacker must be such that the TCP (UDP) checksum does not change, i.e. it is still a valid TCP (UDP) datagram.

In the case where $n = 256$, the attacker has even more choices (see the Appendix Fig. 9). Assuming that TCP/IP datagrams are sent, the attacker can also change the source and destination ports from the TCP datagram as well as the sequence number and the acknowledgement number but again she/he should do that in such a way that the TCP checksum is not changed since it is still in block P_2 . If RTP/UDP/IP datagrams are sent, the attacker can change the source and destination ports from the UDP datagram without being restricted anymore to keep unchanged their UDP checksum, since the checksum is in P_1 . In addition the sequence number in the RTP datagram can also be changed.

The attack cannot be applied when $n = 64$ because the first 64 bits of the IP header do not contain any interesting target fields (see Fig. 3).




Figure 3: $n=64$.

3.2 Mounting the Attack

Assume a source device communicates with a destination device through IPSec gateways A and B (see the Appendix Fig. 10), using “encryption-only” IPSec

(ESP). Consider the $n = 128$ case. As shown in Sect. 3.1, an attacker can intercept a legitimate ESP datagram, modify the source address, such that the altered ESP datagram and the embedded IP header are still valid and re-inject the new ESP datagram in the network. Gateway B will accept this modified ESP datagram, extract the IP header and embedded payload (see the Appendix Fig. 4) and forward them to the destination device. Notice that, for any legitimate ESP datagram, the attacker can generate up to 2^{32} false source addresses and hence so many false ESP datagrams, which will have the same content as the legitimate one but will claim to come from different sources. All of them will be accepted by the gateway B and forwarded to the destination device. In this way, the gateway as well as the destination device are overloaded with junk datagrams. The attacker can mount this attack on the fly in real-time.

If we assume (as in (K. Paterson, 2005)) that the attacker knows several source IP addresses of machines which legitimately communicate through the IPSec gateway, then even stateful firewall could not defeat the attack. Moreover it is enough that the attacker knows a significant portion of this IP addresses, which might be more realistic since it may be that the host is on a network with network prefix known to the attacker, or because it is widely used address prefix.

Consider now the $n = 192$ and $n = 256$ cases. In addition to the source address, the attacker can also alter the destination address along with the source and destination ports. The rest of the attack is similar to the $n = 128$ case, except that the junk datagrams can be targeted to any destination address. So the attacker can mount a DoS attack on any selected server in the Internet.

4 CONCLUSIONS

This paper presents a very strong and practical DoS attack against the “encryption-only” configuration for ESP (IPSec). It also shows that, contrary to the intuition, increasing the cipher block size jeopardizes the security of the protocol even more, i.e. the attack becomes stronger. We sincerely hope that this attack will serve as a strong argument to improve the existing standard by banning the “encryption-only” configuration and will be a compelling example for the implementors/users of the standard convincing them not to use such a configuration.

ACKNOWLEDGEMENTS

This work has been prepared in conjunction with the design and development of the IPSec communication

suite for the Philips Semiconductors Nexperia Mobile platform. The author would like to thank Marijke De Soete, Xavier Serret, Marc Vauclair and Bruno Keymolen for the valuable discussions.

REFERENCES

- A. Menezes, P. van Oorschot, S. V. (1996). Handbook of applied cryptography. In *CRC Press*.
- Adams, C. (1997a). Constructing symmetric ciphers using the CAST design procedure. In *Designs, Codes, and Cryptography 12(3)* pp. 283–316.
- Adams, C. (1997b). The CAST-128 encryption algorithm. In *RFC 2144*.
- Bellovin, S. (1996). Problem areas for the IP security protocols. In *Usenix Unix Security Symposium*, pp. 1–16.
- C. McCubbin, A. Selcuk, D. S. (2000). Initialization vector attacks on the IPSec protocol suite. In *WETICE'00, IEEE Computer Society*, pp. 171–175.
- Canvel, B., Hiltgen, A., Vaudenay, S., and Vuagnoux, M. (2003). Password interception in a SSL/TLS channel. In *CRYPTO'03, LNCS 2729*, pp. 583–599.
- Housley, R. (Jan. 2004). Using advanced encryption standard (AES) counter mode with IPSec encapsulating security payload (esp). In *RFC 3686*.
- J. Daemen, V. R. (2002). The design of rijndael: AES - the advanced encryption standard. In *LNCS*.
- K. Paterson, A. (2005). Cryptography in theory and practice: The case of encryption in IPSec. In *Cryptology ePrint Archive, 2005/416, to appear at EUROCRYPT'06*.
- Kent, S. (March 2005). IP encapsulating security payload (ESP). In *Internet-Draft*.
- M. Bellare, T. Kohno, C. N. (2004). Breaking and provably repairing the SSH authenticated encryption scheme: A case study of the encode-then-encrypt-and-MAC paradigm. In *ACM Transactions on Information and System Security (TISSEC) 7(2)*, pp. 206–241.
- N. Borisov, I. Goldberg, D. W. (2001). Intercepting mobile communications: The insecurity of 802.11. In *MOBI-COM'01*, pp. 180–189.
- NIST (1977). Encryption standard (DES). In *FIPS PUB 46*.
- NIST (1980). DES modes of operation. In *FIPS PUB 81*.
- NIST (1999). Triple-DES (3DES). In *FIPS PUB 46-3*.
- NIST (2002). Advanced encryption standard (AES). In *FIPS PUB 197*.
- Postel, J. (Aug. 1980). User datagram protocol. In *RFC 768*.
- Postel, J. (Sept. 1981). Transmission control protocol. In *RFC 793*.
- R. Braden, D. Borman, C. P. (Sept. 1988). Computing the internet checksum. In *RFC 1071*.
- R. Pereira, R. A. (Nov. 1998). The ESP CBC-mode cipher algorithms. In *RFC 2451*.
- Rivest, R. (1994). The RC5 encryption algorithm. In *FSE'94*, pp. 86–96.
- S. Frankel, R. Glenn, S. K. (Sept. 2003). The AES-CBC cipher algorithm and its use with IPSec. In *RFC 3602*.
- Schneier, B. (1993). Description of a new variable-length key, 64-bit block cipher (blowfish). In *FSE'93*, pp. 191–204.
- Schulzrinne, H., Casner, S., Frederick, R., and Jacobson, V. (July 2003). RTP: A transport protocol for real-time applications. In *RFC 3550*.
- T. Mallory, A. K. (Jan. 1990). Incremental updating of the internet checksum. In *RFC 1141*.
- T. Yu, S. Hartman, K. R. (2004). The perils of unauthenticated encryption: Kerberos version 4. In *NDSS'04*.
- Tanenbaum, A. (2002). Computer networks. In *Prentice Hall*.
- X. Lai, J. M. (1990). A proposal for a new block encryption standard. In *EUROCRYPT'90*, pp. 389–404.

APPENDIX

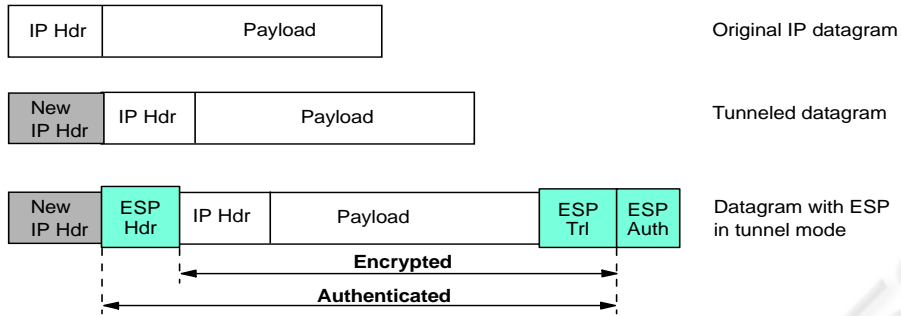


Figure 4: ESP in Tunnel Mode.

0	4	8	16	17	18	20	24
VERS	HLEN	Service Type	Total Length				
ID			FLG	Fragment Offset			
TTL		Protocol	Header Checksum				
Source IP Address							
Destination IP Address							
IP Options						Padding	
Data							

Figure 5: Format of an IP Datagram Header.

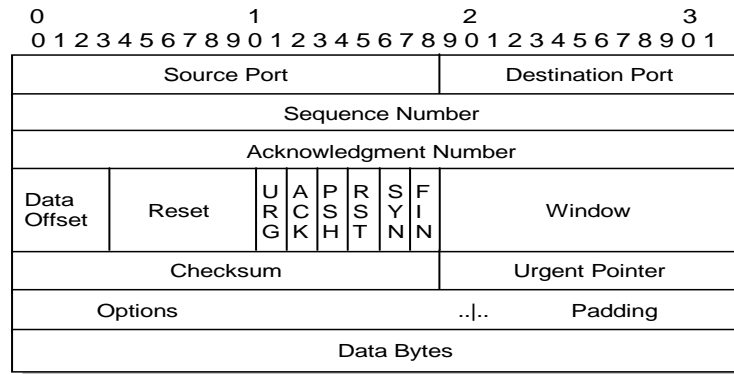


Figure 6: TCP Segment Format.

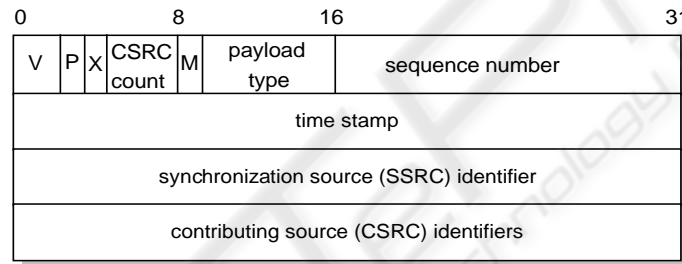


Figure 7: RTP Header Format.




Figure 8: n=192.

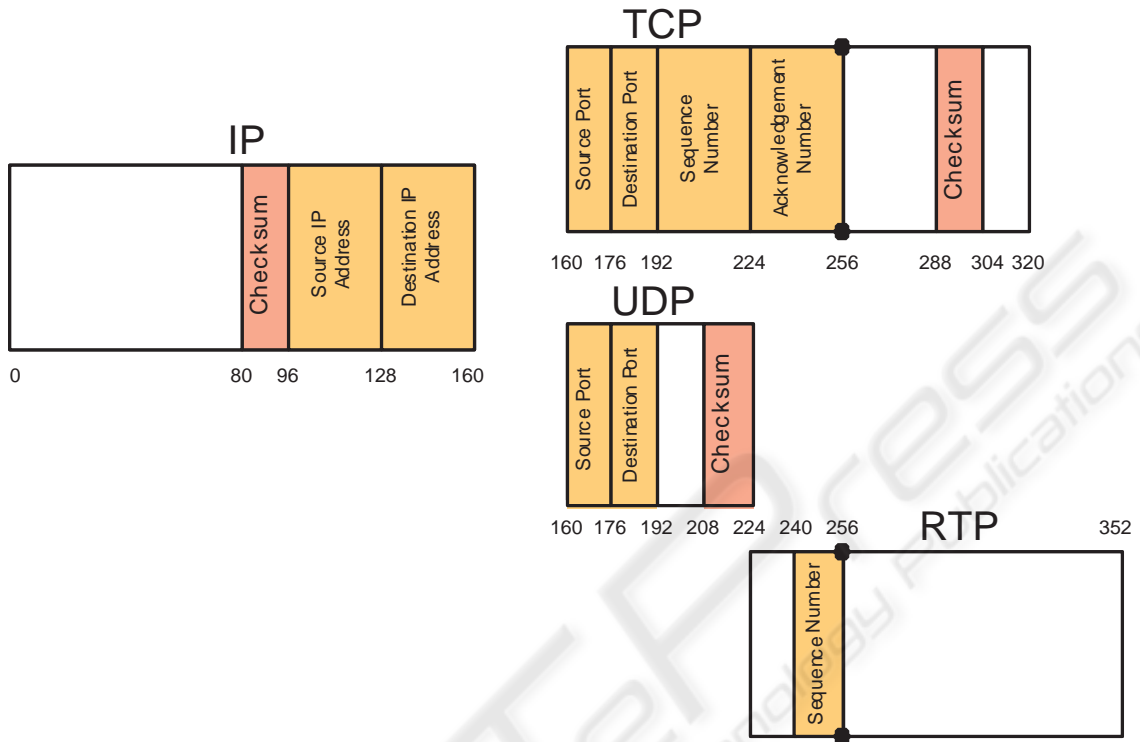


Figure 9: n=256.




Figure 10: IPsec.