CS276 Cryptography	**Spring 2004**

Lecture 10 : 2.25.04

Lecturer: David Wagner *Scribe: Cheng Tien Ee*

Disclaimer: *These notes have not been subjected to the usual scrutiny reserved for formal publications. They may be distributed outside this class only with the permission of the Instructor.*

10.1 Security of Old IPSec

IPSec is the network layer security protocol used to encrypt IP packets between a pair of end hosts. IPSec requires one key per pair of hosts (say A and B). In host A's network stack, each packet with destination IP address matching that of B's will be encrypted. Since this encryption takes place at the network layer, it is transparent to the application, and thus we might run Telnet with improved security.

The previous version of IPSec used Cipher-Block Chaining (CBC) mode (shown in Figure 10.1) to encrypt IP packets, where each block cipher implements Advanced Encryption Standard (AES). Neither Message Authentication Codes (MAC) nor Cyclic Redundancy Checks (CRC) were used in that version. MACs are optional in the new version of IPSec. Although the AES blocks are chained, the randomizing effect of each block implies that slight changes to the input, either the Initialization Vector (IV) or one of the message blocks, will cause only limited unpredictable change to the rest of the ciphertext. Also, from Figure 10.1 we can see that the first message block `M[1]` is unprotected if the adversary is able to change the IV.

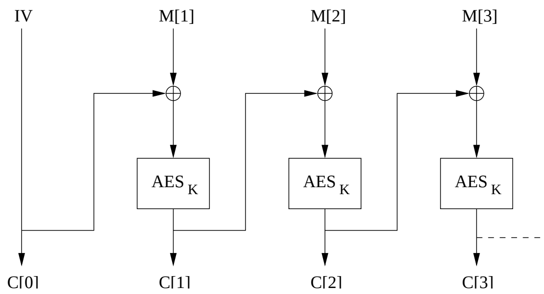

Figure 10.1: CBC mode with AES

10.1.1 Possible Attacks on Old IPSec

In this section we list 3 possible attacks on the old version of IPSec. The attacks use the facts that (1) machines may have multiple hosts who will then share the same network stack and keys and (2) IPSec uses a symmetric key for each pair of hosts.

1. Suppose hosts A and B have just completed a session and both close their individual TCP ports. An adversary can now log into B's machine, open the same TCP port, and replay the traffic.

2. Session hijacking can take place in the following manner: if the port numbers used for communication between A and B are known, the adversary can first open a connection from A's to B's machine, close the connection, then replay traffic when B logs in.

3. The third attack allows the adversary to send encrypted messages to a reserved port (ie. with number less than 1024). With reference to Figure 10.2, suppose the adversary intends to send IP packets to port 23 at B's machine. She proceeds as follows:

   (a) log into A's machine, record and send encrypted message $E_K(m')$ to port 4000 on B's machine

   (b) next send message $E_K(m)$ to B on port 23, prepended with the ciphertext $E_K(m')$ to give $E_K(m')||E_K(m)$.

   (c) At B's machine, the entire ciphertext is decrypted with CBC at the IP layer, giving us $m'||\cdots||m$, where $\cdots$ is some string obtained from the decryption of $E_K(m)$'s IV.

   (d) The entire string $m'||\cdots||m$ is then passed up to the transport layer, which then extracts m' and m, and discards $\cdots$. Demultiplexing is performed, where m' is sent to port 4000, and m to port 23.

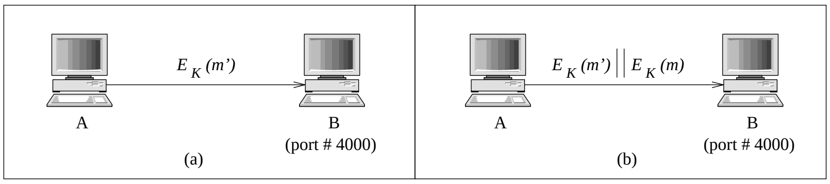

Figure 10.2: Adversary sending data to reserved port

This attack does not work for TCP because of its 1s'-complement checksum, but is possible for UDP packets as their checksum is optional.

10.1.2 Important Lessons

We should never encrypt messages without Message Authentication Codes (MAC). Even if only confidentiality is required, the system may unintentionally reveal the plaintext, which may then be subject to alteration by the adversary. Also, the overall system must be designed to be secure against combined chosen plaintext and ciphertext attacks.

10.2 IND-CCA2

Definition 10.1 *The advantage of an adversary A in indistinguishability under chosen-ciphertext attack 2, or IND-CCA2, in the real-or-random sense is given by:*

$$\mathbf{Adv\ A} = Pr[A^{E_K, D_K} = 1] - Pr[A^{E_K \circ \$, D_K} = 1]$$

where the oracle $E_K \circ \$$ returns the encryption of a random message of the same size as the initial one.

Note that the above definition is too easily broken since on the right side we will almost certainly not get back the initial message if we decrypt the encryption of a random message, whilst we will definitely obtain the initial message if we decrypt its encryption. Thus, a restriction is required:

Restriction: The adversary is not allowed to query D_K on any ciphertext c returned by the first oracle (either E_K or $E_K \circ \$$).

We have the following relationship

$$(t, q, \epsilon) - RoR - CCA2 \Leftrightarrow \mathbf{Adv\ A} \leq \epsilon \qquad \forall\ \mathbf{A}, \quad time \leq t, \quad queries \leq q.$$

Also, if the encryption algorithm is secure in CCA2, then it is secure against lunchtime attack (or CCA1). Recall that the difference between CCA1 and CCA2 is that the adversary in CCA1 has access to the oracle only before she gets the ciphertext.