

Notes on Security Protocols
15-440, Fall 2011
Carnegie Mellon University
Randal E. Bryant

References:

Tannenbaum: 9.1, 9.2 (skip 9.2.3)

BASICS

Desired attributes of security:

Confidentiality: Do not reveal information to untrusted parties

Integrity: Ensure that information has not been corrupted

Authentication: Ensure identity of source of information

Availability: Ensure that desired resource can be used.

Attacking strategies (Depict with box diagrams)

Passive

Code breaking: Determine sensitive data or keys

Traffic analysis: Learn from overall communication patterns

Active

Masquerade: Pretend to be someone else

Replay: Reuse old information

Alteration: Modify message

Create spurious messages: (For denial of service)

CRYPTO BUILDING BLOCKS:

1. Private key crypto:

Have key K that must be kept secret between parties

Operations

$K(M)$: Encrypt message M using key K to create cyphertext C

$K^{-1}(C)$: Decrypt cyphertext using key K . Should give M .

(Desired) Attributes:

Can be made fast and safe

Requires key distribution & preservation

Given C , hard to determine M or K

Given C & P , hard to determine K

Attack strategies:

Ciphertext. Given multiple samples C_1, C_2, \dots

Deduce message(s) or K

Known plaintext. Given multiple pairs $(M_1, C_1), (M_2, C_2), \dots$

Deduce K

Chosen plaintext. Given ability to compute $K()$

Generate pairs $(M_1, C_1), \dots$ for chosen messages

Deduce K

Examples:

DES: 56-bit keys, encrypt 64-bit blocks

Security: No longer good enough.

Can buy hardware DES crackers that do brute-force attack in 1 day

Triple DES: $K = (K_1, K_2, K_3)$ (168 bits total)

$K(M) = K_1(K_2^{-1}(K_3(M)))$

$K^{-1}(M) = K_3^{-1}(K_2(K_1^{-1}(C)))$

Why the middle inversion?

If let $K_1=K_2=K_3$, then get regular DES.
If let $K_1=K_3$, then get double-DES.

AES (a.k.a. Rijndael): 128-bit blocks, key = 128, 192, or 256

2. Public Key Crypto Systems

Two keys:

K^+ : Public key. Can be widely disseminated
 K^- : Private key. Known only to key owner.

Work both ways as encryption/decryption pair:

$K^+(K^-(M)) = M$
 $K^-(K^+(M)) = M$

Stream encryption

Standard encryption schemes are block-based. Work on fixed size blocks.

What if have message that is longer than single block.

Obvious (Electronic Codebook [ECB]):

Break M into blocks M_1, M_2, \dots, M_n (Pad final if necessary)
Encrypt each separately $K(M) = [K(M_1), \dots, K(M_n)]$

Better (Cypher block chaining [CBC]):

(Assume block size and cyphertext size are same)

Generate random block R

$K(M) = C_0, C_1, C_2, \dots, C_n$
where $C_0 = R, C_i = K(C_{i-1} \text{ XOR } M_i)$

Decoding

Recover M_1, M_2, \dots
 $M_i = K^{-1}(C_i) \text{ XOR } C_{i-1}$

SECURITY PROTOCOLS

Shared Key authentication

Suppose parties Alice (A) and Bob (B) have shared secret key K_{ab} . A wants to initiate session with B where each party is sure of who is at the other end. (or at least, that party at each end knows K_{ab}).

Requires 4 rounds:

1. A→B: A

2. B→A: R_b

R_b is a "nonce" short for "number used once". Should be randomly generated so that no one will be able to guess next nonce to be used by B. See go function `crypto/rand.Int`

3. A→B: $K_{ab}(R_b)$, R_a

First part is A's way of proving to B that she knows K_{ab} .

4. B→A: $K_{ab}(R_a)$

This is B's way of proving that he knows K_{ab} .

Interesting idea: Can we reduce this to 3 rounds:

1. A→B: A, R_a

2. B→A: R_b , $K_{ab}(R_a)$

A sure that B knows K_{ab}

3. $K_{ab}(R_b)$

B sure that A knows K_{ab}

Vulnerability:

B is providing a free encryption service that can be used by attacker C:

1. C→B: A, R_a

2. B→A (intercepted by C): R_b , $K_{ab}(R_a)$

1'. C→B: A, R_b

2'. B→A (intercepted by C): R_b' , $K_{ab}(R_b)$

3. C→B: $K_{ab}(R_b)$

Bob now thinks that C knows K_{ab} , and so C can masquerade as A.

General principle. Can label each agent by what it knows.

Say $\text{Kappa}(A, X)$ means "A knows X"

Examples:

If $\text{Kappa}(A, K)$, and B sends A message $K(M)$, then A knows M.

If A knows K, A sends nonce R_a to B, and B sends A $K(R_a)$, then
A knows that B knows K. Write $\text{Kappa}(A, \text{Kappa}(B, K))$

Revisiting protocol:

$\text{Kappa}(A, K_{ab})$
 $\text{Kappa}(B, K_{ab})$
 $\text{Kappa}(C, \{\})$ [What does a potential attacker know?]

1. $A \rightarrow B: A$

2. $B \rightarrow A: R_b$

3. $A \rightarrow B: K_{ab}(R_b), R_a$

$\text{Kappa}(B, \text{Kappa}(A, K_{ab}))$
 $\text{Kappa}(C, [R_b, K_{ab}(R_b)])$

4. $B \rightarrow A: K_{ab}(R_a)$

$\text{Kappa}(A, \text{Kappa}(B, K_{ab}))$
 $\text{Kappa}(C, [R_a, K_{ab}(R_a)])$

All an attacker learns is two plain/ciphertext pairs. Can do even better by modifying steps to be:

3'. $A \rightarrow B: K_{ab}(R_b+1), R_a$.

4'. $B \rightarrow A: K_{ab}(R_a+1)$

These are good enough for authentication but yield even less information about K_{ab} .

Key Distribution

Suppose have many clients: A, B, C, ...

Impractical to require that each pair has secret key K_{xy} .

Instead, have centralized "key distribution center" or KDC. Call it S. Assume KDC is trustworthy. Keeps its secrets. Does not attempt malicious behavior.

KDC maintains keys with each client. K_{as} , K_{bs} , ...

When A & B want to communicate, get S to create "session key" K_{ab} with which they can conduct a conversation. As name implies, generally used for bounded duration.

Version 1:

1. A→S: A,B

KDC generate K_{ab}

2. S→A: $K_{as}(K_{ab})$

$Kappa(A, K_{ab})$

3. S→B: $K_{bs}(K_{ab})$

$Kappa(B, K_{ab})$

Observe: A & B aren't really sure of each others identity, but could go through above authentication protocol & work things out.

Version 2:

Same idea, but want to eliminate communication from S to B.

1. A→S: A,B

KDC generate K_{ab}

2. S→A: $K_{as}(K_{ab}), K_{bs}(K_{ab})$

$Kappa(A, K_{ab})$

Second part is a "ticket". A cannot use this information directly. It's like a sealed envelope that it can present to B.

3. A→B: A, $K_{bs}(K_{ab})$

$Kappa(B, K_{ab})$

Again, could follow earlier protocol to have A & B authenticate to each other.

VULNERABILITY:

Both of these protocols are vulnerable to replay attacks. E.g., suppose that C had earlier communicated with A via session key K_{ac} . C could spoof the server:

1'. A→S: A,B. Intercepted by C.

2'. C→A: $K_{as}(K_{ac}), K_{cs}(K_{ac})$ [C saw earlier message going to A]

Now when A & C engage in authentication protocol, A will think it's talking to B.

Needham-Schroeder Protocol

Original protocol proposed in 1978. We show slightly enhanced version here, and describe a known weakness.

Purpose:

Set up secure channel between A & B with session key K_{ab} .
Authenticate A & B to each other.
Make sure all information is fresh.

Use three nonces: R_{a1} , R_{a2} , R_b

1. A→S: R_{a1} , A, B

2. S→A, $K_{as}(R_{a1}, B, K_{ab}, K_{bs}(A, K_{ab}))$

$K_{appa}(A, K_{ab})$.
 $K_{appa}(A, K_{appa}(S, K_{as}))$

A knows this message was in direct response to #1. So, this is a fresh session key from S. Include B to differentiate from other session keys A might be creating.

Ticket $K_{bs}(A, K_{ab})$ includes identity of A. Prevents ticket from some other channel from being reused.

Encrypt everything under K_{as} , so that attacker cannot learn anything about session, or any plain/ciphertext pairs.

3. A→B: $K_{ab}(R_{a2}), K_{bs}(A, K_{ab})$

$K_{appa}(B, K_{ab})$

(Note that sending ticket in step #2 encrypted was useless, since we've revealed it here.)

4. B→A: $K_{ab}(R_{a2}-1, R_b)$

$K_{appa}(A, K_{appa}(B, K_{ab}))$

This is a slight variation on previous nonce technique:

If A knows K, A sends $K(R_a)$ to B, and B sends A $K(R_a-1)$,
then A knows that B knows K.

5. A→B: $K_{ab}(R_b-1)$

$K_{appa}(B, K_{appa}(A, K_{ab}))$

Note that the standard N-S protocol does not include R_{a2} , and therefore there is no way for A to be sure that it's really talking to B (more specifically, that it is talking to a party that was able to get the value of K_{ab} .)

This protocol is still considered to be a bit weak, in that there is nothing that can assure B that the ticket it receives on step #3 is fresh, not an old ticket that is being reused. This weakness was first described publicly in 1981 by Denning & Sacco, who suggested fixing it by adding a time stamp to the ticket. Later (1987) Needham & Schroeder published a fix that adds two more steps to the beginning, as well as an additional nonce R_{b0} . Here are the two steps (numbered -1 and 0, to indicate their relationship to the other 5 steps):

-1. A→B: A

A tells B that it wants to set up a session

0. B→A: $K_{bs}(R_{b0})$

B provides a nonce that A cannot decode.

We incorporate this nonce into the next 3 steps of the protocol:

1'. A->S: Ra1, A, B, Kbs(Rb0)

2'. S->A: Kas(Ra1, B, Kab, Kbs(A, Kab, Rb0))

3'. A->B: Kab(Ra2), Kbs(A, Kab, Rb0)

We see that nonce Rb0 gets incorporated into the ticket that A provides B. B can now be sure that this is a newly generated ticket.

Kerberos

Uses a version of secret key N-S to set up secure channels between users & services. Uses time stamp rather than nonce's to avoid reuse of stale keys. Based on following:

If A knows K, B sends $K(t)$ to A, and current time $\tilde{t} = t$,
then A knows that B knows K.

Requires A & B to maintain synchronized clocks. More precise ==>
less vulnerable to replay attacks.

Public key N-S:

1. A→B: $K_b + (A, R_a)$

B generates secret session key K_{ab} .

$Kappa(B, K_{ab})$

2. B→A: $K_a + (R_a, R_b, K_{ab})$

A certain that message came from B.

$Kappa(A, K_{ab})$

$Kappa(A, Kappa(B, K_{ab}))$

3. A→B: $K_{ab}(R_b)$

$Kappa(B, Kappa(A, K_{ab}))$

B certain that message came from A.