Network Security

Chapter 3

Cryptographic Protocols
for Encryption, Authentication
and Key Establishment

Overview

□ Part I: Introduction

□ Part II: The Secure Channel

□ Part III: Authentication and Key Establishment Protocols

- Key Distribution Centers (KDC)

- Public Key Infrastructures (PKI)

- Building Blocks of key exchange protocols

Cryptographic Protocols

❑ Definition:

A *cryptographic protocol* is defined as a series of steps and message exchanges between multiple entities in order to achieve a specific security objective.

❑ Properties of a protocol (in general):

- Everyone involved in the protocol must know the protocol and all of the steps to follow in advance.

- Everyone involved in the protocol must agree to follow it.

- The protocol must be unambiguous, that is every step is well defined and there is no chance of misunderstanding.

- The protocol must be complete, i.e. there is a specified action for every possible situation.

❑ Additional property of a cryptographic protocol:

- It should not be possible to do or learn more than what is specified in the protocol.

Applications of Cryptographic Protocols

- Key establishment
- Authentication
  - Data origin authentication
  - Entity authentication
- Authenticated key establishment
- Data integrity
- Confidentiality

treated in this course

- Secret sharing
- Key escrow (ensuring that only an authorized entity can recover keys)
- Zero-Knowledge proofs (proof of knowledge of an information without revealing the information)
- Blind signatures (useful for privacy-preserving time-stamping services)
- Secure elections
- Electronic money

Data Origin Authentication - Data Integrity

- Definitions:
    - *Data integrity* is the security service that enables entities to verify that a message has not been altered by unauthorized entities.
    - *Data origin authentication* is the security service that enables entities to verify that a message has been originated by a particular entity and that it has not been altered afterwards. Therefore, in contrast to the data integrity service, data origin authentication necessarily involves identifying the *source* of a message.

- *Data origin authentication* implies *data integrity.*
- Although it is possible to achieve data integrity without data origin authentication, they are normally achieved by the same mechanisms
    - By the application of a cryptographic hash function with a digital signature and appending the signed hash value to the message.
    - By the use of MAC that is appended to the message .

- Data integrity of messages exchanged is a fundamental building block to cryptographic protocols.
- Without data integrity, hardly any security goal can be achieved.

Entity Authentication

❑ Definition:

Entity authentication is the security service that enables communication partners to verify the identity of their peer entities.

❑ Entity authentication is the most fundamental security service, as all other security services build upon it.

❑ In general it can be accomplished by various means:

▪ *Knowledge:* e.g. passwords

▪ *Possession:* e.g. physical keys or cards

▪ *Immutable characteristic:* e.g. biometric properties like fingerprint, etc.

▪ *Location:* evidence is presented that an entity is at a specific place (example: people check rarely the authenticity of agents in a bank).

▪ *Delegation of authenticity:* the verifying entity accepts that somebody who is trusted has already established authentication.

❑ In communication networks, direct verification of the above means is difficult or insecure which motivates the need for cryptographic protocols.

Overview

□   Part I: Introduction

□   Part II: The Secure Channel

□   Part III: Authentication and Key Establishment Protocols

- Key Distribution Centers (KDC)

- Public Key Infrastructures (PKI)

- Building Blocks of key exchange protocols

Problem Statement (1)

(c.f. Niels Ferguson, Bruce Schneier: Practical Cryptography, Ch 8, pp. 111ff)

❑ Goal

▪ This chapter illustrates the functionality of a "secure channel" between two parties Alice and Bob, provided by a simple security protocol, the so-called secure channel algorithm.

▪ The functionality of this secure channel is a good start for understanding the functionality of common security protocols such as IPSec and SSL.

Alice   Eve   Bob

Problem Statement (2)

❑ Assumptions

- The channel is bi-directional, i.e. Alice sends messages to Bob and Bob sends messages to Alice (almost all communications are bi-directional).

- Eve tries to attack the secure channel in any possible way.

  - Eve can read all of the communication between Alice and Bob and arbitrarily manipulate exchanged messages.

  - Particularly, Eve can delete, insert, or modify exchanged messages.

❑ Requirement

- Alice and Bob share a secret session key K that is known only to both of them.

- The way how this key is established is discussed in the next Chapter.

Security Properties

❑ Processing

- ▪ Alice needs to send a sequence of messages (Service Data Units: SDU) $m_1, m_2, \ldots$
- ▪ These messages are processed by the secure channel algorithm (i.e. the security protocol), which generates PDUs (Protocol Data Units) and sends them to Bob.
- ▪ Bob processes the received PDUs using the corresponding secure channel algorithm and ends up with a sequence of messages $m_1', m_2', \ldots$
- ▪ In the ideal case $\{m_1', m_2', \ldots\} = \{m_1, m_2, \ldots\}$

Security Properties

- Security properties of the secure channel algorithm
  - Eve does not learn anything about the messages m_i except for their timing and size.
  - $\{m_1{'}, m_2{'}, \ldots \} \leq \{m_1, m_2, \ldots \}$
  - Bob can not prevent Eve from deleting a message in transit
    (and Bob can not prevent message loss neither).
  - The messages received are in the correct order.
  - There are no duplicate messages, no modified messages and no bogus messages sent by someone else other than Alice.
  - Bob knows exactly which messages he has missed.

General Remark

❑ Some protocols have some acknowledgement mechanisms for recovering from message loss.

❑ However, this is not handled by the secure channel, since it would make it more complicated.

Outline

❑ The message processing in the secure channel consists of
  ▪ Message numbering
  ▪ Authentication
  ▪ Encryption

❑ There are two approaches for the order of applying the authentication and the encryption to a message

  (1) One may either encrypt the message first, and then authenticate the obtained cipher text

  (2) Or one might authenticate first and then encrypt the message with the MAC value together
  (this approach is used subsequently)

❑ Both approaches have advantages and disadvantages
  ▪ If encryption is applied first (1), Bob can discard bogus messages before spending CPU resources on decrypting them
  ▪ If authentication is applied first (2), the MAC value will be also protected.
  ▪ Also, the Horton principle: "Authenticate what you mean, not what you say" ➔ (2)
  ▪ See [Fer03] for further details on this discussion

Message Numbering

❑ Message numbers allow Bob to reject replayed messages.

❑ They tell Bob which messages got lost in transit.

❑ They ensure that Bob receives the messages in their correct order.

❑ Messages numbers increase monotonically,

i.e. later messages have a greater message number.

❑ Message numbers have to be unique;

i.e. no two messages may have the same message number.

❑ A simple message numbering scheme functions as follow:

- Alice numbers the first message as 1, the second message as 2, etc.
- Bob keeps track of the last message number he has received.
- Any new message must have a message number that is larger than the message number of the previous message.
- If the message number overflows,

  e.g. message number is an 32-bit integer and the current message is

  2^{32} -1, then Alice needs to stop using the current session key K before it can wrap back to 0.

Authentication/ Encryption

- For the data authentication, we need a MAC function,
    - e.g. HMAC-SHA-256
    - With a hash-value of 256 bits the collision rate is extremely low.
- The input to the MAC consists of
    - the message number i
    - the message m_i
    - extra authentication data x_i that is required by Bob to interpret m_i,
      e.g., protocol version number, negotiated field size, etc.
    - Note: the length of x_i must be fix
- Let $a_i := $ MAC $(i \,||\, x_i \,||\, m_i)$
    - The way how x_i is interpreted is out-of-scope and not a part of the functionality of the secure channel algorithm. The secure channel algorithm just considers it as a string.
    - However, the secure channel assures the integrity of x_i
- For encryption, we need an encryption algorithm,
    - e.g. AES in CTR mode with 256 bits (since it is pretty fast and secure)
- Frame Format
    - the message that Alice finally sends to Bob consists of message number i, followed by $E(m_i \,||\, a_i)$.

Initialization of the Secure Channel (1)

❑ The initialization procedure of the secure channel generates 4 different keys from the existing shared session key *K:*

- An encryption key and an authentication key to send messages from Alice to Bob.

- An encryption key and an authentication key to send messages from Bob to Alice.

- It is strongly recommended not to reuse the same key for different purposes.

❑ Moreover, the initialization procedure sets the initial message numbers.

Possible Attacks on Secure Channels

❑ If the same key K is used for different purposes, different attacks become possible.

❑ E.g. If K is used for encryption in both directions, then *known-plain-text* attacks become possible:

- Since encryption is done with AES in CTR mode

- and assuming that Alice and Bob initialize the counter for the generation of the key streams k_i in the same way (starting with 0 and incrementing by 1 for each message)

- The key stream k_i generated for each message m_i depends only on K and the sequence number i

- Therefore, for each sequence number i the same key stream k_i will be generated on both sides.

- If an attacker can guess a plain text m_i

  then, it can decrypt m'_i

  $m'_i \| a'_i = c_i \oplus c'_i \oplus (m_i \| a_i)$

- Note: Eve does not need to guess a_i for this attack, since it can perform the xor operation only on the first bits that include m_i.

Alice   Eve   Bob

$$c_i = (m_i \| a_i) \oplus k_i \longrightarrow$$

$$\longleftarrow c'_i = (m'_i \| a'_i) \oplus k_i$$

Initialization of the Secure Channel (2)

Function InitializeSecureChannel
Input: K   Key of the channel
            R   Role: Specified if this party is Alice or Bob
Output   S   State for the secure channel

// First compute the 4 keys that are needed
KeySendEnc ←   SHA-256 (K || „Enc Alice to Bob")
KeyRecEnc   ←   SHA-256 (K || „Enc Bob to Alice")
KeySendAuth ←   SHA-256 (K || „Auth Alice to Bob")
KeyRecAuth ←   SHA-256 (K || „Auth Bob to Alice")
// The strings „Enc Alice to Bob", „Enc Bob to Alice", etc. are simply used to generate different
(uncorrelated) keys.They can be also substitued by other strings, e.g. „A", „B", „C" and „D".

// Swap the encryption and decryption keys if this party is Bob
If R = „Bob" then {   SWAP (KeySendEnc, KeyRecEnc)
                       SWAP (KeySendAuth, KeyRecAuth)
                       }
// Set the send and receive counters to zero. The send counter is the number of the last sent
message. The receive counter is the number of the last received message
(MsgCNTSend, MsgCNTRec) ← (0,0)

// Package the state
S ← (KeySendEnc, KeyRecEnc, KeySendAuth, KeyRecAuth, MsgCNTSend, MsgCNTRec)
return S

Sending a Message

Function SendMessage

Input: S   Secure session state

             m   message to be sent

             x   addtional data to be authenticated

Output   t   data to be transmitted to the receiver

// First check the number and update it

if (MsgCNTSend >= MAX_MSG_NUMBER){

     print „MsgCNTSend overflow; re-keying is required"

     exit

}

MsgCNTSend ← MsgCNTSend + 1

i ← MsgCNTSend

// Compute the authentication

a ← HMAC-SHA-256 (KeySendAuth , i || x || m)

// Generate key stream k

k ← E(KeySendEnc, nonce || 0) || E(KeySendEnc, nonce || 1) || ...

// the message number *i* can be used as nonce as this would garantee

// that (nonce || block-number) will be unique each time E is applied

// Form the final text (i is an integer of 4 bytes length)

m_a ← m||a

t ← i || ($m_a \oplus$ first-lenght(m_a)-bytes(k))

return t

Receiving a Message

Function ReceiveMessage
Input: S   Secure session state
               t   text received from transmitter
               x   addtional data to be authenticated
Output   m   message that was sent
// Split t into i and the encrypted message plus authenticator.
// This split is unambiguous since i is an integer of 4 bytes length
i || t' ← t
// Generate the key stream, just as the sender did
k ← E(KeyRecEnc, nonce || 0) || E(KeyRecEnc, nonce || 1) || ...
// Decrypt the message and MAC field, and split.
// This split is also unambiguous since length of MAC value a is known (in this case 256 bits)
m || a ← t' ⊕ first-lenght(t')-bytes(k)
// Recompute the authentication
a' ← HMAC-SHA-256 (KeyRecAuth, i || x || m)
if (a' ≠ a) {
        destroy k, m
        return MsgAuthenticationFailure }
else if (i <= MsgCNTRec) {
         destroy k, m
         return MessageOrderError }
MsgCNTRec ← i

return m

Message Reordering during Transmission

❏ The presented algorithm for processing received messages guarantees that no message is received twice.

❏ However, messages that are re-ordered during transmission, otherwise perfectly valid, will be lost.

❏ In some situations this can be inefficient, e.g. with IP packets, since they can be reordered during transport.

❏ We will see that IPSec, the IP security protocol that encrypts and authenticate IP packets, deals with this problem by maintaining a replay protection window instead of a single counter.

Further Design Criteria

- Negotiation of cryptographic algorithms
  - AES-256 and SHA-256 are just examples in this secure channel
  - Most security protocols support the negotiation of the cryptographic algorithms to be used.
  - This is the case, e.g., for IPSec and SSL
  - The Kerberos authentication protocol version 4 (which will be explained later in this lecture) was based only on DES because DES was considered to be secure by the time Kerberos v4 was developed.
  - This flaw was fixed for Kerberos v5.

Further Design Criteria

❑ Multiple communication partners simultaneously

- In many cases, Alice wishes to communicate with several partners simultaneously, e.g. with Bob and Carol.

- Bob and Carol may use different cryptographic algorithms.

- Therefore, Alice needs to know how to handle a message received from Bob or from Carol.

- Each message needs to include a unique identifier for the connection in order to facilitate this task.

- E.g.
    - In IPSec, this identifier is called the Security Parameter Index (SPI)
    - In SSL/TLS, there is a so-called Session Identifier (SessionID)
    - (Both IPSec and SSL/TLS will be explained in subsequent chapters of this lecture)

Secure Channels - Conclusions

❑ The secure channel is one of the most useful application of cryptography.

❑ Given good encryption and authentication primitives, it is possible to construct a secure channel.

❑ However, there are a lot of small details to pay attention to.

❑ Some applications require encryption.

❑ However, in most cases, authentication is more important.

❑ In fact Eve can cause a lot more damage if she manipulates messages and sends bogus messages, than by just listening to the message sent by Alice.

❑ In practice, the main difficulty frequently is not the secure channel itself.

❑ It can be instead

  ▪ Establishing the shared secret

  ▪ Knowing to whom you are talking to ➔ Entity authentication

Overview

Part I: Introduction

Part II: The Secure Channel

Part III: Authentication and Key Establishment Protocols

- **Introduction**
- Key Distribution Centers (KDC)
- Public Key Infrastructures (PKI)
- Building Blocks of key exchange protocols

Problem Statement

❑ Goal

▪ Run a key exchange protocol such that at the end of the protocol:

- Alice and Bob have agreed on a shared „session key" for a secure channel
- Alice and Bob have agreed on the cryptographic algorithms to be used for the secure channel
- Alice (Bob) must be able to verify that Bob (Alice) knows K and that he (she) is "alive"
- Alice and Bob must know that K is newly generated

Entity Authentication or Key Establishment? (1)

❑ Many authentication protocols – as a side effect of the authentication exchange - do establish a secret session key for securing the session (to be used only for the current session).

❑ Some opinions about the relationship between authentication and key establishment:

- „It is accepted that these topics should be considered jointly rather separately" [Diff92]

- „… authentication is rarely useful in the absence of an associated key distribution" [Bell95]

- „In our view there are situations when entity authentication by itself may useful, such as when using a physically secured communication channel." [Boyd03]

Entity Authentication or Key Establishment? (2)

❑ Example

- Alice wants to use the online banking service provided by her bank

- Alice can perform an online banking session from any terminal using a (secure) Internet Browser

- The Internet browser authenticates the web server based on the certificate (see below) which includes the public key of the web server.

- Authentication of the web server:
  - as a consequence of this authentication mechanism, a shared session key $K_{A,B}$ is generated, which can be used for this session (it is important that this session key is correctly destroyed when the session is over)

- Authentication of the client:
  - the web server authenticates Alice based on her PIN number. (As a consequence of the successful authentication of Alice, no additional secret key is established.)

- This example shows that both cases are common:
  - Entity authentication with key establishment
  - Entity authentication without key establishment

- The goals of a protocol have to be carefully set up for each application scenario

- Entity authentication   - Mutual entity authentication

- Entity auth. with key establishment   - Mutual entity auth. with key establishment

First Try: Key Establishment with Diffie-Hellman

❑ Assume Alice and Bob want to establish a secure channel with a shared secret $K_{A,B}$

❑ The Diffie-Hellman protocol introduced in Chapter 4 is our first example of a cryptographic protocol for key exchange. So what's wrong with it?

❑ The problem with a "simple DH exchange" is that a man-in-the-middle attack is possible.

❑ Neither Alice nor Bob know after a protocol run with whom they actually have exchanged a key

Second Try: Static Approach

❑ Static Approach for the negotiation of "session keys" and cryptographic algorithms

  ▪ Keys are manually exchanged. Cryptographic algorithms are agreed on personally

❑ Pro's

  ▪ Simple,

  ▪ session keys are automatically authenticated

❑ Con's

  ▪ Manual process is required (either by a direct meeting or by phone call)

  ▪ Does not scale for a large set of hosts

    $\dfrac{n*(n-1)}{2}$ symmetric keys would be needed for n entities

  ▪ Renewing of keys or cryptographic algorithms require another manual process

  ▪ If the key is compromised, all sessions can be compromised (also previous recorded sessions!)

  ▪ Keys are not changed frequently

Example: Static Approach in GSM/UMTS Networks

❑ The user mobile phone share a long-term secret key with the home network.

❑ The secret key is stored in the SIM card that the user received from his provider.

❑ Note: in GSM/UMTS networks, the scalability issue is not severe

  ▪ A mobile device does not communicate directly with other mobile devices.

  ▪ Communication takes place between the mobile device and the network instead.

  ▪ Only n symmetric keys are required (instead of $\frac{n*(n-1)}{2}$ keys).

❑ More details on GSM/UMTS security will be provided in a subsequent chapter in this lecture.

Trusted Third Parties (TTP)

❑ Boyd's Theorem [Boyd03]

▪ *„Assuming the absence of a secure channel, two entities cannot establish an authenticated session without the existence of an entity that can mediate between the two and which both parties trust and have a secure channel with".*

❑ A TTP is a special entity which has to be trusted by its users

❑ A TTP can significantly reduce the key management complexity

❑ "Trusted" means that it is expected to always behave honestly.

❑ The TTP is assumed to always respond exactly according to the protocol specification, and, therefore, will never deliberately compromise the security of its clients.

Key Distribution Centers (KDC)

❑ A KDC is an option for providing authentication and key establishment.

❑ A KDC is a TTP that shares secrets with all entities (an entity may be a user or a host).

❑ Alice asks KDC for a secret to (securely) talk to Bob.

❑ KDC generates a secret $K_{A,B}$

❑ Example of KDCs:

  ▪ The Kerberos protocol is based on a KDC.

  ▪ In fact, a Kerberos server is often called a KDC.

❑ Drawbacks:

  ▪ KDC can monitor all authentication and key establishment activities.

  ▪ KDC knows the session key.

  ▪ KDC needs to be online during the authentication and key establishment procedure.

  ▪ KDC is a potential single-point-of-failure/ bottleneck.

Public Key Infrastructures (PKI)

❑ A Certificate Authority (CA) asserts the correctness of the certificate by signing it with her private key.

❑ CA is a trusted third party (TTP) that is trusted by all the entities.

❑ All entities know the public key of the CA.

❑ Since Alice knows CA's public key, she can verify the signature of Bob's certificate that was generated by CA.

❑ See later in this chapter for more details on PKIs.

Trusted Third Parties (TTP) – General Remarks

❑ The TTP is a very powerful entity in this topology. If an attacker manages to compromise TTP, he will be in control of the whole network!

❑ The TTP may directly be involved in the authentication procedure, which is the case for KDCs.

  ➔ Online TTP

❑ TTP may not be required for the authentication.

❑ In case a CA signs the public key of Alice, and Bob knows the public key of the CA, he will be able to verify the validity of Alice's certificate that is signed by CA without talking to CA.

  ➔ Offline TTP (provides more scalability)
     However, Certificate Revocation Lists (CRLs) are still required.

Some Notation...

Notation of Cryptographic Protocols (1)	
Notation	Meaning
A	Name of A, analogous for B, E, TTP, CA
CA_A	Certification Authority of A
r_A	Random value chosen by A
t_A	Timestamp generated by A
$(m_1, ..., m_n)$	Concatenation of messages $m_1, ..., m_n$
$A \rightarrow B: m$	A sends message m to B
$K_{A, B}$	Secret key, only known to A and B

Some Notation...

Notation of Cryptographic Protocols (2)	
Notation	**Meaning**
$K_{A\text{-}pub}$	Public key of A
$K_{A\text{-}priv}$	Private key of A
$\{m\}_K$	Message m encrypted with key K, synonym for $E(K, m)$
	(also integrity protection in case of shared key protocols)
$H(m)$	Cryptographic hash value over message m, computed with function H
$A[m]$	Shorthand notation for $(m, \{H(m)\}_{K_{A\text{-}priv}})$
$Cert_{CK_{CA\text{-}priv}}(K_{A\text{-}pub})$	Certificate of CA for public key $K_{A\text{-}pub}$ of A, signed with private certification key $CK_{CA\text{-}priv}$
CA<<A>>	Shorthand notation for $Cert_{CK_{CA\text{-}priv}}(K_{A\text{-}pub})$

Overview

❑ Invented in 1978 by Roger Needham and Michael Schroeder [Nee78]

Roger Needham

Michael Schroeder

The Needham-Schroeder Protocol (2)

❑ The Needham-Schroeder Protocol is a protocol for mutual authentication and key establishment

❑ It aims to establish a session key between two users (or a user and an application server, e.g. email server) over an insecure network

❑ The protocol has 2 versions:

- The *Needham Schroeder Symmetric Key Protocol*:
  - based on symmetric encryption
  - Forms the basis for the *Kerberos* protocol
- The *Needham Schroeder Public Key Protocol*:
  - *Uses public key cryptography*
  - A flaw in this protocol was published by Gavin Lowe [Lowe95] 17 years later!
 - Lowe proposes also a way to fix the flaw in [Lowe95]

Gavin Lowe

Overview

- Part I: Introduction

- Part II: The Secure Channel

- Part III: Authentication and Key Establishment Protocols

  - Introduction

  - Key Distribution Centers (KDC)

    - Needham-Schroeder Protocol

      - Symmetric Version

      - Asymmetric Version

    - Kerberos

  - Public Key Infrastructures (PKI)

  - Building Blocks of a key exchange protocol

Authentication
Server AS

(user data)

Long term symmetric key: $K_{AS, A}$

Long term symmetric key $K_{AS, B}$

1. Request Ticket, Challenge$_1$

2. Response$_1$, Session Key, Ticket

User Alice: *A*

User Bob: *B*

3. Ticket, Challenge$_2$

4. Response$_2$, Challenge$_3$

5. Response$_3$

The Needham Schroeder Symmetric Key Protocol - Overview

❑   *AS* shares symmetric keys with all users, in particular with Alice ($K_{AS,A}$) and Bob ($K_{AS,B}$)

1.) $A \rightarrow AS: (A, B, r_1)$

- Alice sends a message to *AS* with her name und Bob's name, telling the server she wants to communicate with Bob.

- In other words, Alice asks the *KDC* to supply a session key and a "ticket" for secure communication with Bob.

- The freshly generated random number r_1 is used to authenticate *AS* and avoid that a man-in-the-middle is pretending to be *AS*.

2.) $AS \rightarrow A: \{r_1, K_{A,B}, B, Ticket_{A,B}\}_{K_{AS,A}}$ where $Ticket_{A,B} = \{K_{A,B}, A\}_{K_{AS,B}}$

- *AS* generates the session key $K_{A,B}$ and sends it to Alice encrypted with $K_{AS,A}$

- *AS* includes r_1 in the encrypted message, so Alice can confirms that r_1 is identical to the number generated by her in the first step, thus she knows the reply is a fresh reply from *AS*.

- Furthermore, *AS* includes a copy of the session key $K_{A,B}$ for Bob included in *Ticket*$_{A,B}$

- Note here that during this protocol run, *AS* does not communicate directly with Bob

- Since Alice may be requesting keys for several different people, the inclusion of Bob's name tells Alice who she is to share this key with.

❑ Needham-Schroeder protocol definition (continued):

3.) A → B: (Ticket$_{A,B}$)

- Alice forwards the ticket to Bob.
- Bob can decrypt the ticket with $K_{AS,B}$ and get the session key $K_{A,B}$.
- Since Alice's name A is included in the ticket, Bob knows that this ticket was granted by AS for Alice.

4.) B → A: {r$_2$ }$_{K_{A,B}}$

- After decrypting message (3), Bob generates the new random number r_2 and includes it in message (4) which is encrypted with the freshly generated session key $K_{A,B}$.
- However, Bob still also needs to verify that Alice knows the session key $K_{A,B}$ and that she is alive (otherwise, an attacker could send an "old" ticket pretending to be Alice). Therefore, Bob challenges Alice with this new random number r_2

5.) A → B: {r$_2$ - 1}$_{K_{A,B}}$

- Alice checks if message 4 was encrypted with the freshly generated session key $K_{A,B}$. Since Alice does not know r_2, she has to check the integrity of the message (or detect by similar means that Bob used key $K_{A,B}$).
- After decrypting Bob's message, Alice computes r_2 - 1 and answers with message (5)
- Bob decrypts the message and verifies that it contains $r_2 - 1$.

❑ Needham-Schroeder also proposed a protocol variant where Alice reuses the Ticket from the server. Key $K_{A,B}$ is therefore not new anymore and it cannot be used to authenticate Bob. As a consequence Alice needs to include a challenge in message (3).

❑ Protocol variant with reuse of ticket and shared key:

1.)+ 2.) Not necessary, Alice reuses the ticket.

3.) $A \to B$: ($Ticket_{A,B}$, $\{r_2\}_{K_{A,B}}$)

- Alice sends the ticket again to Bob.
- Bob either still knows the ticket or he can decrypt the ticket again with $K_{AS,B}$ and get the session key $K_{A,B}$.
- Since Alice's name A is included in the ticket, Bob knows that this ticket was granted by AS for Alice.
- As the session key is not fresh anymore, Alice cannot authenticate Bob with $K_{A,B}$. In order to verify that Bob is alive, receiving Alice's messages and still has the correct session key, Alice includes a challenge in message (3) which consists of a nonce random number r_2

4.) $B \to A$: $\{r_3, r_2 - 1\}_{K_{A,B}}$

- After decrypting message (3), Bob calculates ($r_2 - 1$) and includes it in message (4) which is encrypted with the freshly generated session key $K_{A,B}$
- However, Bob still also needs to verify that Alice really knows the session key $K_{A,B}$ and that she is alive (otherwise, an attacker could send an "old" ticket pretending to be Alice).
- Therefore, Bob must challenge Alice with a new random number r_3

5.) $A \to B$: $\{r_3 - 1\}_{K_{A,B}}$

- After decrypting Bob's message, Alice computes $r_3 - 1$ and answers with message (5)
- Bob decrypts the message and verifies that it contains $r_3 - 1$.

❑ Discussion:

- The *Needham-Schroeder Symmetric Key Protocol* can be considered as secure (no known attacks so far) if the session key $K_{A,B}$ can not be "brute-forced" or discovered by an attacker.

- However, if an attacker, Eve, can manage to get to know a session key $K_{A,B}$, she can later use this to impersonate as Alice by *replaying* the message 3:

  3') $E \rightarrow B: (Ticket_{A,B}, r_2)$

  4') $B \rightarrow A: \{r_3, r_2 - 1\}_{K_{A,B}}$, Eve has to intercept this message

  Since Eve knows $K_{A,B}$ knows she will be able to decrypt Bob's reply 4') and answers with

  5') $E \rightarrow B: \{r_3 - 1\}_{K_{A,B}}$

  So, if an attacker Eve is able to compromise **one** session key $K_{A,B}$, she will be able to impersonate Alice in the future even though she doesn't know $K_{A,TTP}$

- This problem is solved in the Kerberos protocol with *timestamps*.

The Needham-Schroeder Symmetric Key Protocol (6)

❑ Note:

- The term „ticket" was not used in the original description of the Needham-Schroeder Protocol. [Nee78]
- However, it is used here to provide an analogy with the Kerberos protocol.
- In the Kerberos protocol, the ticket includes more data than $K_{A,B}$ and A.

Overview

- ❑ Part I: Introduction

- ❑ Part II: The Secure Channel

- ❑ Part III: Authentication and Key Establishment Protocols

  - ▪ Introduction

  - ▪ Key Distribution Centers (KDC)

    - • Needham-Schroeder Protocol

      - – Symmetric Version

      - – Asymmetric Version

    - • Kerberos

  - ▪ Public Key Infrastructures (PKI)

  - ▪ Building Blocks of a key exchange protocol

The Needham-Schroeder Public Key Protocol (1)

❑ The Needham-Schroeder Public Key Protocol

- Protocol description

- Attack published by Gavin Lowe in 1995

❑ Assumptions

- ▪ *AS* is a trusted server.
- ▪ *AS* knows the public keys of all users
- ▪ All users know *AS*'s public key

❑ Protocol run

1.) $A \rightarrow AS: (A, B)$

- ▪ Alice requests Bob's public key from *AS*.

2.) $AS \rightarrow A: \{ K_{B\text{-}pub}, B \}_{K_{AS\text{-}priv}}$

- ▪ AS asserts that Bob's public key is $K_{B\text{-}pub}$

3.) $A \rightarrow B: \{ r_A, A \}_{K_{B\text{-}pub}}$

- ▪ Alice generates a random number r_A and sends it to Bob together with her name, encrypted with Bob's public key $K_{B\text{-}pub}$

4.) $B \rightarrow AS: (B, A)$

- ▪ Bob requests Alice's public key from *AS*.

❑ Needham-Schroeder public key protocol definition (continued):

5.) $AS \rightarrow B: \{ K_{A\text{-}pub}, A \}_{K_{AS\text{-}priv}}$

▪ AS asserts that Alice's public key is $K_{A\text{-}pub}$

6.) $B \rightarrow A: \{ r_A, r_B \}_{K_{A\text{-}pub}}$

▪ Bob generates a random number r_B and sends it to Alice together with r_A encrypted with $K_{A\text{-}pub}$. Thus, Bob proves to Alice that he was able to decrypt message (3) successfully and therefore proving his identity to Alice. Here in message (6), Bob challenges also, whether she can decrypt the message and extracts r_B.

7.) $A \rightarrow B: \{ r_B \}_{K_{B\text{-}pub}}$

▪ Alice decrypts message (6) with her private key, extracts r_B and encrypts it with Bob's public key.
▪ Upon receipt, Bob can verify that r_B is correct and thus verify that he is talking to Alice.

❑ At the end of the protocol run, Alice and Bob know each other's identities, know both r_A, r_B but r_A, r_B are not known to eavesdroppers. Therefore, a symmetric session key $K_{A,B}$ can be now easily derived on both sides: e.g. $K_{A,B} = H(r_A, r_B)$, where H is cryptographic hash function that has been agreed on a priori.

The Needham-Schroeder Public Key Protocol (4)

❑ Attack:

▪ The *Needham-Schroeder Public Key Protocol* is vulnerable to a *man-in-the-middle attack.*

▪ If an attacker *M* can persuade *A* to initiate a session with him, he can relay the messages to *B* and convince *B* that he is communicating with *A*.

▪ For simplicity, we don't illustrate the communication with *AS* here, which remains unchanged.

3') $A \rightarrow M$: $\{ r_A , A \}_{K_{M\text{-}pub}}$

  • A sends r_A to M, who decrypts the message with $K_{M\text{-}priv}$

3'') $M \rightarrow B$: $\{ r_A , A \}_{K_{B\text{-}pub}}$

  • *M relays the message to B, pretending that A is communicating*

6') $B \rightarrow M$: $\{ r_A , r_B \}_{K_{A\text{-}pub}}$

  • *B* sends r_B

6'') $M \rightarrow A$: $\{ r_A , r_B \}_{K_{A\text{-}pub}}$

  • *M* relays it to *A*

❑ Attack on the Needham-Schroeder public key protocol (continued):

7') $A \rightarrow M: \{ r_B \}_{K_{M\text{-}pub}}$

- ▪ A decrypts r_B and confirms it to M, who learns it

7'') $M \rightarrow B: \{ r_B \}_{K_{B\text{-}pub}}$

- ▪ M re-encrypts r_B and convinces B that he has decrypted it.

❑ At the end of the attack, B falsely believes that A is communicating with him, and that r_A and r_B are known only to A and B.

❑ The attack was first described in 1995 by Gavin Lowe [Lowe95].

❑ The paper also describes a fixed version of the protocol, referred to as the *Needham-Schroeder-Lowe* protocol. The fix involves the modification of message (6)

6.) $\qquad B \rightarrow A: \{ r_A , r_B \}_{K_{A\text{-}pub}}$

which is replaced with the fixed version

6.) $\qquad B \rightarrow A: \{ r_A , r_B , B \}_{K_{A\text{-}pub}}$

Overview

Kerberos (1)

- Kerberos is an authentication and access control service for work-station clusters that was designed at the MIT during the late 1980s

- Design goals:
  - *Security:* eavesdroppers or active attackers should not be able to obtain the necessary information to impersonate a user when accessing a service
  - *Reliability:* as every use of a service requires prior authentication, Kerberos should be highly reliable and available
  - *Transparency:* the authentication process should be transparent to the user beyond the requirement to enter a password
  - *Scalability:* the system should be able to support a large number of clients and servers

- The underlying cryptographic primitive of Kerberos is symmetric encryption (Kerberos V. 4 uses DES, V. 5 allows other algorithms)

- A good tutorial on the reasoning beyond the Kerberos design is given in [Bry88a], that develops the protocol in a series of fictive dialogues

Kerberos (2)

- The basic usage scenario of Kerberos is a user, *Alice,* who wants to access one or more different services (e.g. AFS, printing server, email server), that are provided by different servers S1, S2, ... connected over an insecure network
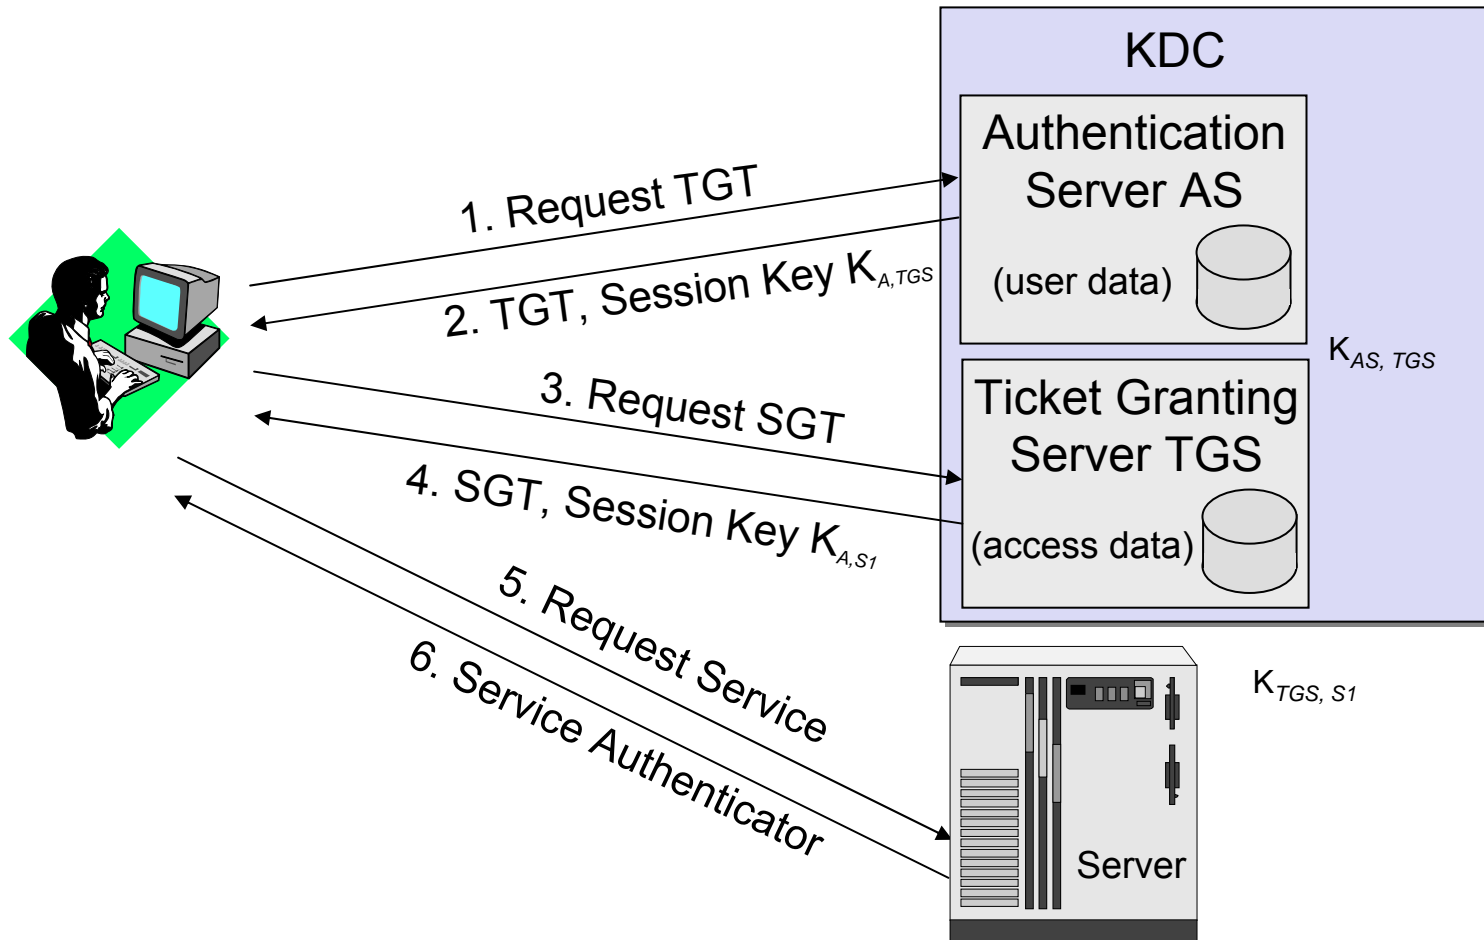
- Kerberos deals with the following security aspects of this scenario:

  - *Authentication:* Alice will authenticate to an *authentication server (AS)* who will provide a *temporal permit* to demand access for services. This permit is called *ticket-granting ticket (TGT, also called Ticket$_{TGS}$)* and is comparable to a temporal passport.

  - *Access control:* by presenting her *ticket-granting ticket (Ticket$_{TGS}$)* Alice can demand a *ticket granting server (TGS)* to obtain *access to a service* provided by a specific server *S1*. The TGS decides if the access will be permitted and answers with a *service granting ticket* (SGS, also called *Ticket$_{S1}$*) for server *S1*.

  - *Key exchange:* the authentication server provides a *session key $K_{A,TGS}$* for communication between Alice and TGS and the TGS provides a *session key $K_{A,S1}$* for communication between Alice and *S1*. The use of these session keys also serves for authentication purposes.

KDC

Authentication
Server AS

(user data)

$K_{AS, TGS}$

Ticket Granting
Server TGS

(access data)

$K_{TGS, S1}$

Server

1. Request TGT

2. TGT, Session Key $K_{A,TGS}$

3. Request SGT

4. SGT, Session Key $K_{A,S1}$

5. Request Service

6. Service Authenticator

Accessing a Service with Kerberos Version 4 - Protocol Overview

Kerberos (4)

- At the beginning, the user *A* logs on at his workstation and requests to access a service:
  - The workstation represents him in the Kerberos protocol and sends the first message to the authentication server *AS*, containing his *name A*, a *timestamp t_A* , the name of an appropriate *ticket granting server TGS* and the requested ticket lifetime:

    1.) $A \rightarrow AS: (A, t_A, TGS, RequestedTicketLifetime_{TGS})$

- *AS* looks up *A* and his password in the user's database, generates the master key $K_{A,AS}$ out of A's password ($K_{A,AS}$ = MD5(Password$_A$)), extracts the workstation *IP address Addr$_A$*, creates a *ticket granting ticket Ticket$_{TGS}$* and a *session key $K_{A,TGS}$*, and sends the following message to A:

    2.) $AS \rightarrow A: \{K_{A,TGS}, TGS, t_{AS}, LifetimeTicket_{TGS}, Ticket_{TGS}\}_{K_{A,AS}}$

    with $Ticket_{TGS} = \{K_{A,TGS}, A, Addr_A, TGS, t_{AS}, LifetimeTicket_{TGS}\}_{K_{AS,TGS}}$

- Upon receipt of this message, the workstation asks user A to type in her password, computes the key $K_{A,AS}$ from it, and uses this key to decrypt the message. If Alice does not provide her "authentic" password, message (2) can not be decrypted correctly and the protocol run will fail.

Kerberos (5)

❏ Alice creates a so-called *authenticator* and sends it together with the ticket-granting ticket and the server name *S1* to TGS:

> 3.) $A \rightarrow TGS: (S1, Ticket_{TGS}, Authenticator_{A,TGS})$
>
> with $Authenticator_{A,TGS} = \{A, Addr_A, t'_A\}_{K_{A,TGS}}$

- With the Authenticator, *A* can prove to TGS that she knows the secret $K_{A,TGS}$

- In order to counter reply attacks, a fresh timestamp t'_A is included in the *Authenticator.*

- An authenticator must be used only once.

❏ Upon receipt, *TGS* decrypts $Ticket_{TGS}$, extracts the session key $K_{A,TGS}$ and uses this key to decrypt $Authenticator_{A,TGS}$.

❏ If the name and address in the authenticator and in the ticket are matching and the timestamp t'_A is still fresh (not older than 5 minutes), it checks if *A* may access the service S1 based on the access policies database and creates the following message:

> 4.) $TGS \rightarrow A: \{K_{A,S1}, S1, t_{TGS}, Ticket_{S1}\}_{K_{A,TGS}}$
>
> with $Ticket_{S1} = \{K_{A,S1}, A, Addr_A, S1, t_{TGS}, LifetimeTicket_{S1}\}_{K_{TGS,S1}}$

Kerberos (6)

❑ Alice decrypts the message and does now hold a session key for secure communication with *S1*. She now sends a message to *S1* to show him her ticket and a new authenticator:

  5.) $A \rightarrow S1$: (Ticket$_{S1}$, Authenticator$_{A,S1}$)

  with Authenticator$_{A,S1}$ = {A, Addr$_A$, t''$_A$}$_{K_{A,S1}}$

❑ Here, the *Authenticator* is used to counter *replay attacks.*

❑ Upon receipt, *S1* decrypts the ticket with the key $K_{TGS,S1}$ he shares with *TGS* and obtains the session key $K_{A,S1}$ for secure communication with *A*. Using this key he checks the authenticator and responds to *A*:

  6.) $S1 \rightarrow A$: {t''$_A$ + 1}$_{K_{A,S1}}$

❑ By decrypting this message and checking the contained value, Alice can verify that she is really communicating with *S1*, as only he (besides *TGS*) knows the key $K_{TGS,S1}$ required to decrypt *Ticket$_{S1}$* which contains the session key $K_{A,S1}$, and so only he is able to decrypt *Authenticator$_{A,S1}$* and to answer with *t''$_A$* + 1 encrypted with $K_{A,S1}$
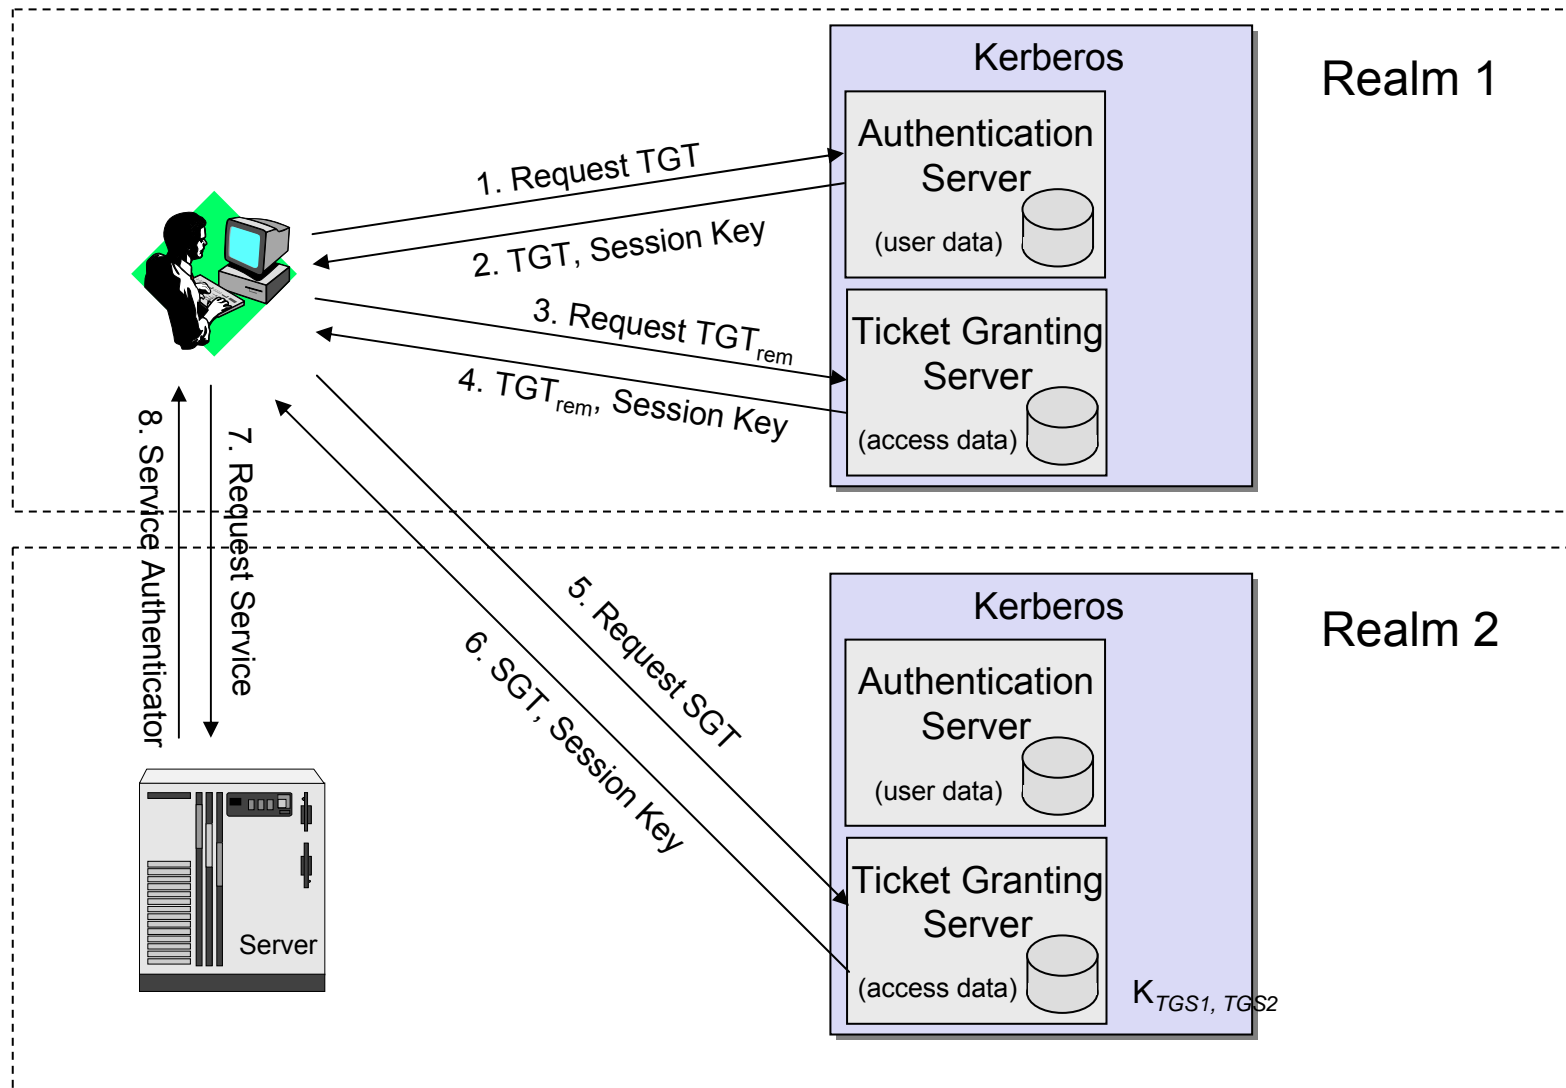
Multiple Domain Kerberos (1)

❑ Consider an organization with workstation clusters on two different sites, and imagine that user *A* of site 1 wants to use a server of site 2:

- ▪ If both sites do use their own Kerberos servers and user databases (containing passwords) then there are in fact two different *domains*, also called *realms* in Kerberos terminology.

- ▪ In order to avoid that user A has to be registered in both realms, Kerberos allows to perform an inter-realm authentication.

❑ Inter-realm authentication requires, that the ticket granting servers of both domains share a secret key $K_{TGS1,TGS2}$

- ▪ The basic idea is that the *TGS of another realm* is viewed as a *normal server* for which the TGS of the local realm can hand out a ticket.

- ▪ After obtaining the ticket for the remote realm, Alice requests a service granting ticket from the remote TGS

- ▪ However, this implies that remote realm has to trust the Kerberos authentication service of the home domain of a "visiting" user!

- ▪ Scalability problem: *n* realms require $n \times (n-1)/2$ secret keys!

Multiple Domain Kerberos (3)

❑ Messages exchanged during a multiple domain protocol run:

1.) $A \rightarrow AS1: (A, t_A, TGS1, RequestedTicketLifetime_{TGS})$

2.) $AS1 \rightarrow A: \{K_{A,TGS1}, TGS1, t_{AS}, LifetimeTicket_{TGS1}, Ticket_{TGS1}\}_{K_A}$

with $Ticket_{TGS1} = \{K_{A,TGS1}, A, Addr_A, TGS1, t_{AS}, LifetimeTicket_{TGS1}\}_{K_{AS1,TGS1}}$

3.) $A \rightarrow TGS1: (TGS2, Ticket_{TGS1}, Authenticator_{A,TGS1})$

with $Authenticator_{A,TGS1} = \{A, Addr_A, t'_A\}_{K_{A,TGS1}}$

4.) $TGS1 \rightarrow A: \{K_{A,TGS2}, TGS2, t_{TGS1}, Ticket_{TGS2}\}_{K_{A,TGS1}}$

with $Ticket_{TGS2} = \{K_{A,TGS2}, A, Addr_A, TGS2, t_{TGS1}, LifetimeTicket_{TGS2}\}_{K_{TGS1,TGS2}}$

5.) $A \rightarrow TGS2: (S2, Ticket_{TGS2}, Authenticator_{A,TGS2})$

with $Authenticator_{A,TGS2} = \{A, Addr_A, t''_A\}_{K_{A,TGS2}}$

6.) $TGS2 \rightarrow A: \{K_{A,S2}, S2, t_{TGS2}, Ticket_{S2}\}_{K_{A,TGS2}}$

with $Ticket_{S2} = \{K_{A,S2}, A, Addr_A, S2, t_{TGS2}, LifetimeTicket_{S2}\}_{K_{TGS2,S1}}$

7.) $A \rightarrow S2: (Ticket_{S2}, Authenticator_{A,S2})$

with $Authenticator_{A,S2} = \{A, Addr_A, t'''_A\}_{K_{A,S2}}$

8.) $S1 \rightarrow A: \{t'''_A + 1\}_{K_{A,S1}}$

Kerberos V.4 Pro's and Con's (1)

❑ Advantages of Kerberos V.4

- Simpler than V.5 and provides high performance due to hard coding of the parameters

- But: ticket lifetimes are encoded in 1 byte (i.e. 256 different possible values) with 5 minutes as the smallest unit

  ➔ The maximal ticket lifetime is 21 hours and 20 minutes

❑ Disadvantages of Kerberos V.4

- Hard encoding of the parameters can also cause many limitations.

- Offline dictionary attacks on user passwords can easily be mounted by simply requesting a ticket granting ticket for some user from the AS, which is then encrypted with this user's guessed master key.

- Only DES is available for encryption.

- Only IP (and only IPv4) is supported

Kerberos V.4 Pro's and Con's (2)

❑ Maximum ticket lifetime is about 21 hours. This is insufficient for some long running applications.

❑ Tickets can only be used from one host (delegation of rights is not possible).

Kerberos V.5 (1)

❑ Encoding

- Kerberos V.5 uses ASN.1 syntax which is more flexible than binary hard-coded type length

- e.g.

```
HostAddress ::=   SEQUENCE {
                           addr-type[0]      INTEGER,
                           address[1]        OCTET STRING
        }
```

❑ Ticket lifetimes

- Kerberos V.5 allows for much longer ticket lifetimes, since time encoding in ASN.1 allows for times until Dec 31, 9999.

- Since it would useful to invalidate Kerberos tickets that have a long lifetime, an additional management of the tickets is required.
    - E.g. in case employee X has left the company and had root access.

- Kerberos V5 offers the option that tickets can be re-validated by the KDC with a fresh timestamp before they can be re-used.

- The new ticket lifetime features render the management of master key versions at the KDC more complicated than in V.4

Kerberos V.5 (2)

❑ Delegation of rights

- In contrast to Kerberos V.4, in V.5 Alice can request that multiple network addresses should be included in the ticket or no address at all, which means that it can be used from any user's address

- This is useful, e.g. if the user wants to execute some batch scripts even if he is not logged in, e.g. for periodic backups

- It is a policy decision if the KDC issues such tickets, and if a service accepts tickets from specific addresses or not.

- The KDC can log all delegation events and can provide an audit trail in case of a security compromise of a service.

➔ This access control model provides for a lot of flexibility but is also inherently dangerous if not configured correctly.

Kerberos V.5 (3)

❑ Some improvements of cryptographic primitives

- ▪ The master key is a hash function of Alice's password <u>and</u> the realm name.

- ▪ Kerberos V.5 also allows DES and permits new modes which should provide confidentiality and integrity protection.

Kerberos - Evading password guessing attacks

❑ In Kerberos v4, any user, or attacker Mallory, can request a ticket for Alice.

❑ Mallory can not immediately decrypt the message (2) received from *AS*, since she does not know Alice's master key K_A

❑ However, since the algorithm how K_A is derived from Alice's password is known (a hash of the Alice's password), Mallory can perform a password guessing attack (also called dictionary attack) using message (2)

❑ Kerberos v5 has the *pre-authentication* option to prevent this attack

  ▪ Alice needs to include a fresh timestamp encrypted with her master key $\{t_A\}_{K_A}$ when sending the authentication request, i.e. message (1)

❑ This measure is effective against active attackers.

❑ However, passive attacks are still possible

  ▪ Mallory can record authentication exchanges of other users and perform a password guessing attack.

❑ Therefore, it is important to choose good passwords for the security of Kerberos.

❑ Kerberos V4 uses DES in a special mode called *Propagating Cipher Block Chaining* (*PCBC*).

❑ A modification in a cipher text encrypted in CBC mode results into a damage in the next two blocks in the decrypted plain text.

❑ A modification in a cipher text encrypted in PCBC mode results into a damage in all the remaining blocks in the decrypted plain text.
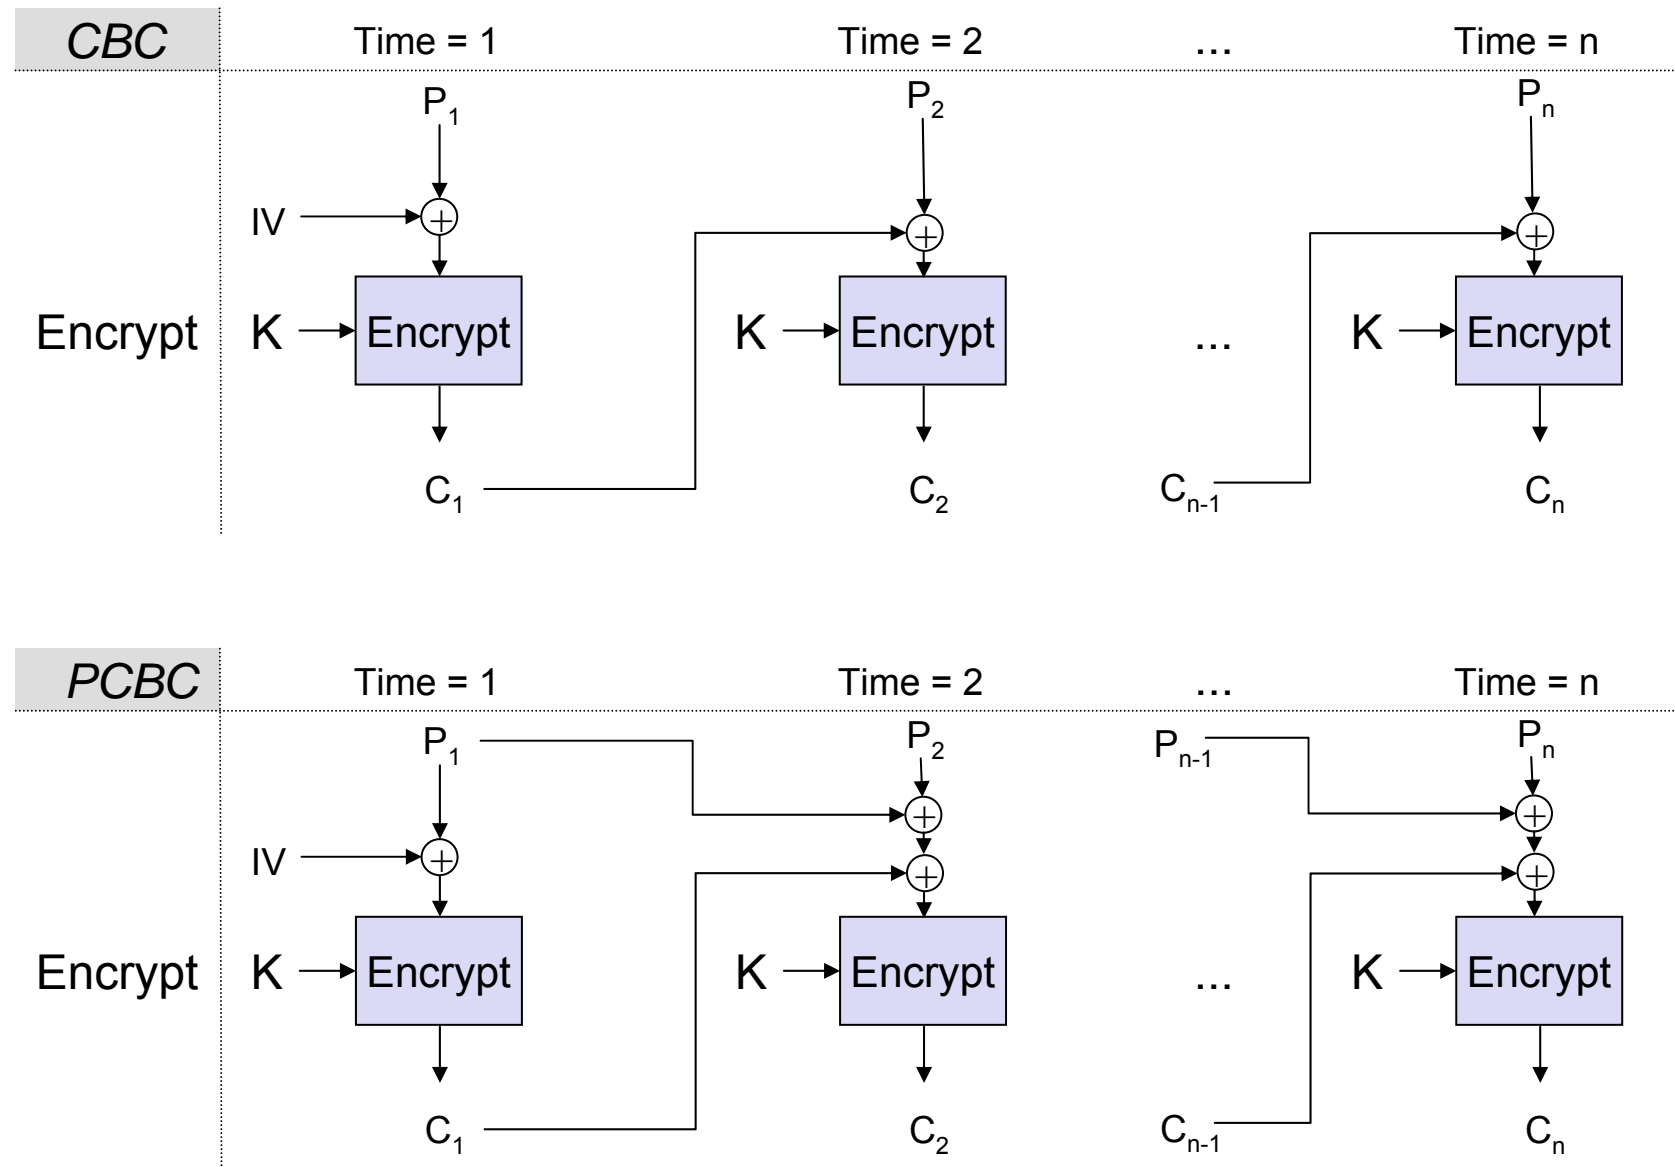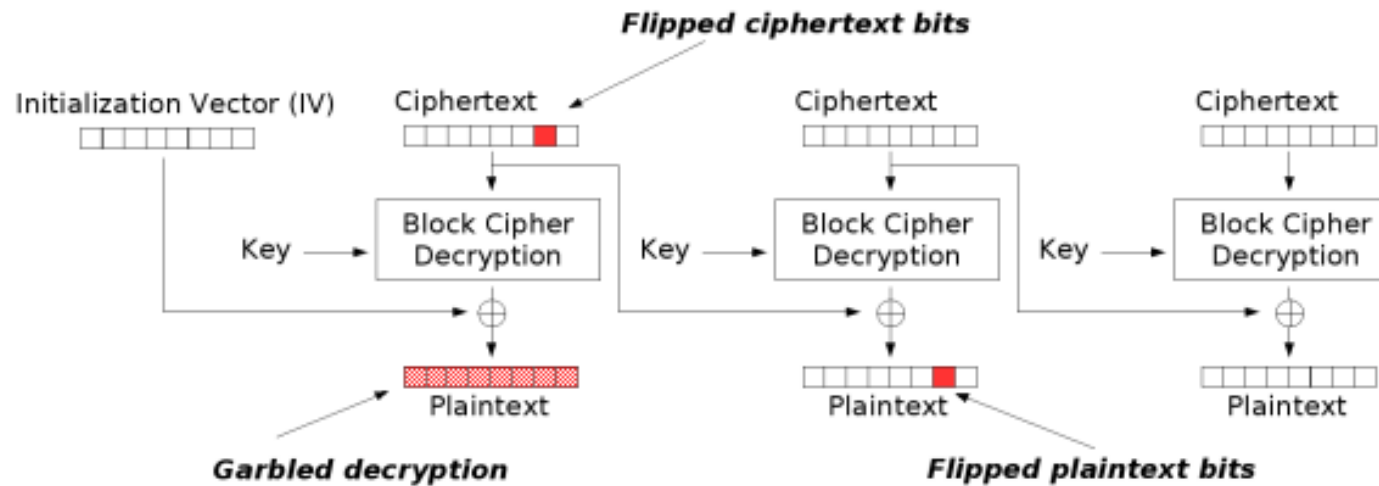
Reminder: CBC Error Propagation

❑ A distorted cipher text block results in two distorted plaintext blocks, as $p_i´$ is computed using c_{i-1} and c_i

Flipped ciphertext bits

Initialization Vector (IV)   Ciphertext   Ciphertext   Ciphertext

Key → Block Cipher Decryption   Key → Block Cipher Decryption   Key → Block Cipher Decryption

Plaintext   Plaintext   Plaintext

Garbled decryption   **Flipped plaintext bits**

Modification attack or transmission error for CBC

Source: http://www.wikipedia.org/

- In Kerberos V4, the PCBC mode was supposed to allow for providing message encryption and integrity by processing the message only once:
  - If a block in the cipher text is manipulated by an attacker, all the remaining blocks in the decrypted plain text at the receiver's side will be damaged.
  - If order to be able to verify the latter case (i.e. whether the message is damaged) Kerberos V4 uses a checksum computed with the key ($K_{AS,A}$, $K_{A,TGS}$ or $K_{A,S1}$) and the message.
  - The encrypted message is transmitted together with the checksum.
  - The algorithm for the checksum is not documented, only implemented.

- Although not broken. Not believed to be strong.

- Potential problems:
  - The checksum is not a cryptographic hash function.
  - Therefore, it might reveal some information about the key (the one-way property of cryptographic hash functions is not necessarily satisfied)
  - It might be also easier to find two messages with the same check sum (the 2nd pre-image resistance of cryptographic hash function is not necessarily satisfied)

- Not used in V5.

Kerberos V5: are things getting better? (1)

❑ Algorithms for encryption and message integrity for Kerberos V5 are specified in [RFC3961]

❑ [RFC3961] allows for *unkeyed* checksums for verifying the data integrity,

  ▪ e.g. CRC, MD5, SHA-1

  ▪ The checksum is encrypted together with the message to be transmitted.

  ▪ *"An unkeyed checksum mechanism can be used with any encryption type, as the key is ignored (a key is not needed for the computing of the checksum), but its use must be limited to cases where the checksum itself is protected, to avoid trivial attacks."* [RFC3961] (Section 4)

  ▪ *"These (unkeyed) checksum types use no encryption keys and thus can be used in combination with any encryption type, but they may only be used with caution, in limited circumstances where the lack of a key does not provide a window for an attack, preferably as part of an encrypted message. Keyed checksum algorithms are recommended."* [RFC3961] (Section 6.1)

Kerberos V5: are things getting better? (2)

❑ Fortunately, Kerberos V5 allows also for keyed checksums
- ▪ e.g. HMAC with MD5 or SHA-1

❑ In this case, different keys can be used for encryption and for data integrity.
- ▪ *"Due to advances in cryptography, some cryptographers consider using the same key for multiple purposes unwise. Since keys are used in performing a number of different functions in Kerberos, it is desirable to use different keys for each of these purposes, even though we start with a single long-term or session key."*
  [RFC3961] (Section 2)

❑ More information on the algorithms used in Kerberos for encryption and data integrity can be found in
- ▪ [RFC3961] Encryption and Checksum algorithms for Kerberos V5
- ▪ [RFC3962] Adds AES cipher suites for Kerberos V5
- ▪ [RFC4757] Microsoft implementation of Kerberos (with RC4 and HMAC)

❑ If you want to use Kerberos, you should use the latest version of it.

❑ It has been around for a while and many competent people have looked at it.

Kerberos – Reality check (1)

❑ In many environements, the application of the user is not „kerberized".

user name & password

auth. success/failure

Server

Kerberos protocol exchange

KDC

❑ In this case, the application servers needs to perform the Kerberos exchange on behalf on the user, and get a ticket for itself, if authentication is successful.

❑ Since the application server requires the user name and password in order to perform the Kerberos authentication exchange successfully, it is very important that the user name and password are protected by another mean, e.g. a TLS tunnel between the user's application and the application server.

❑ Note: the user's password will be revealed to the application server.

Kerberos – Reality check (2)

❑ Even worse: most of the application servers do not support Kerberos by themselves.

❑ However, they can be enabled to use Kerberos with so-called *Pluggable Authentication Modules (PAM)*

Services and applications

| FTP Server | SSH Server | IMAP Server | Apache Web Server |

| PAM Configuration file | → | PAM API |

PAM authentication service modules

| Kerberos | RADIUS | LDAP |

addressed in this lecture

❑ PAM can authenticate users based on different sources of authentication databases and potentially with different protocols.

❑ The native Kerberos protocol as described by MIT is rarely used today.

Kerberos - Discussion

❑ General properties

- Kerberos provides authentication of users and authorization to access services over an insecure network.

- The KDC is logically separated into an *Authentication Server* and a *Ticket Granting Server*

- The user needs to enter her password only once (*Single-Sign-On*). Authenticated access to services with the service tickets occurs transparently to the user.

❑ Reliability

- The KDC is involved in each authentication process.

  ➔ The KDC must be highly reliable

- The design of the Kerberos protocol does not foresee a solution for reliability

- Reliability is implemented with backup KDCs.

Kerberos - Discussion

❑ Security and complexity

- ▪ Although V5 fixes some flaws of V4, it introduces higher complexity, e.g. due to the complex management of tickets.

- ▪ The requirement that all hosts need to have synchronized clocks could be fixed if random numbers are used for the *Authenticators* instead of timestamps. (will be treated at assignments)

- ▪ However, such a fix is currently not foreseen.

- ▪ Kerberos V4 tries to protect the password by using it only once in an authentication protocol run.

- ▪ However, dictionary attacks are very easy to perform

- ▪ An „attacker" Mallory can request a ticket for any user „Alice" and performs a dictionary attacks based on the received response from the KDC.

Kerberos - Discussion

- Kerberos V5 tries to fix this problem with the pre-authentication

- However, dictionary attacks are still possible.

- Mallory needs to wait until Alice requests an ticket.

- Mallory can also listen to Kerberos authentication protocol runs in the network and collects the messages required for a dictionary attack.

- Dictionary attacks on Kerberos could be avoided with an additional DH exchange (will be treated at assignments)

- However, such a fix is currently not foreseen neither.

Overview

PKI - Overview

❑ Each entity has a public key/private key pair,

e.g. RSA or ECC public/private keys

❑ Each entity has a „certificate" that binds its „name" to its public key

❑ Note: in a networking environment "names" could be

▪ a user name (optionally with an email address)

▪ But it could be also e.g. IP addresses, the DNS name of the node, etc.

❑ A Certificate Authority (CA) asserts the correctness of the certificate by signing it with her private key.

❑ CA is a trusted third party (TTP) that is trusted by all the entities.

❑ Furthermore, each entity knows the public key of CA

❑ When Alice whishes to communicate with Bob, she can receives Bob's certificate

▪ E.g. from a directory service or from Bob himself at the beginning of the authentication procedure

❑ Since Alice knows CA's public key, she can verify the signature of Bob's certificate that was generated by CA
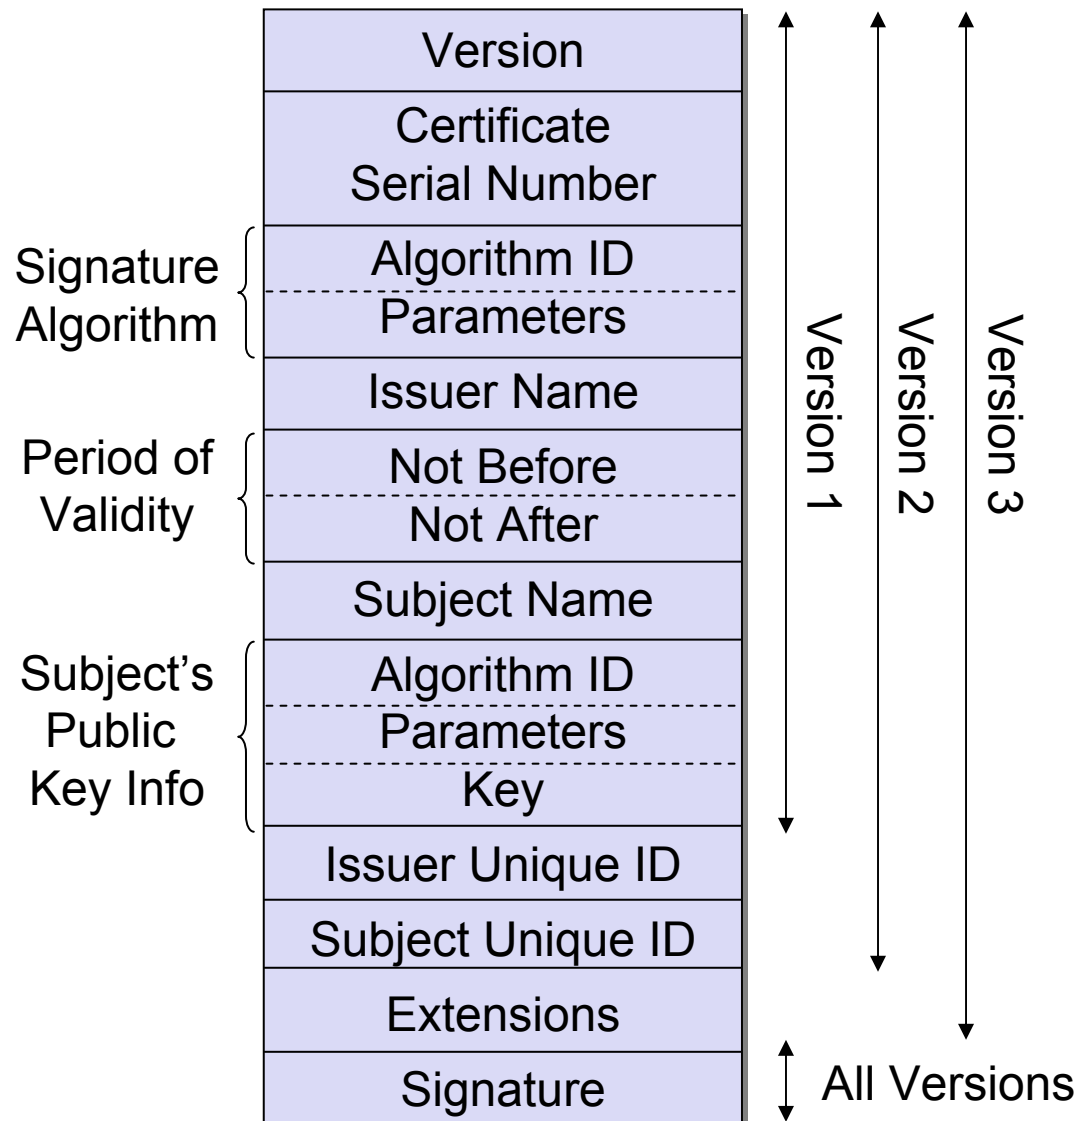
X.509 PKI Authentication Services – Introduction

❑ X.509 is an international recommendation of ITU-T and is part of the X.500-series defining directory services:

- The first version of X.509 was standardized in 1988
- A second version standardized 1993 resolved some security concerns
- A third version was drafted in 1995

❑ X.509 defines a framework for provision of authentication services, comprising:

- Certification of public keys and certificate handling:
  - Certificate format
  - Certificate hierarchy
  - Certificate revocation lists

X.509 – Public Key Certificates (1)

Version
Certificate Serial Number
Algorithm ID
Parameters
Issuer Name
Not Before
Not After
Subject Name
Algorithm ID
Parameters
Key
Issuer Unique ID
Subject Unique ID
Extensions
Signature

Signature Algorithm

Period of Validity

Subject's Public Key Info

Version 1

Version 2

Version 3

All Versions

- ❑ A *public key certificate* is some sort of passport, certifying that a public key belongs to a specific name
- ❑ Certificates are issued by *certification authorities (CA)*
- ❑ If all users know for sure the public key of the CA, every user can check every certificate issued by this CA
- ❑ Certificates can avoid online-participation of a TTP
- ❑ The security of the private key of the CA is crucial to the security of all users!

X.509 – Public Key Certificates (2)

❑ Notation of a certificate binding a public key K_{A-puc} to user *A* issued by certification authority *CA* using its private key $K_{CA-priv}$:

- $Cert_{KCA-priv}(K_{A-pub}) = CA[V, SN, AI, CA, T_{CA}, A, K_{A-pub}]$
  with: *V*   = version number

  SN   = serial number

  AI   = algorithm identifier of signature algorithm used

  CA   = name of certification authority

  T_{CA} = period of validity of this certificate

  A   = name to which the public key in this certificate is bound

  K_{A-pub}= public key to be bound to a name

- The shorthand notation *CA[m]* stands for $(m, \{H(m)\}_{KCA-priv})$
- Another shorthand notation for $Cert_{KCA-priv}(K_{A-pub})$ is CA<<A>>

❑ Consider now two users Alice and Bob, living in different countries, who want to communicate securely:

- Chances are quite high that their public keys are certified by different CAs
- Let's call Alice's certification authority *CA* and Bob's *CB*
- If Alice does not trust or even know *CB* then Bob's certificate *CB<<B>>* is useless to her, and the same applies in the other direction

❑ A solution to this problem is to construct *certificate chains*:

- Imagine for a moment that *CA* and *CB* know and trust each other
  - A real world example of this concept is the mutual trust between countries considering their passport issuing authorities
- If *CA* certifies *CB*'s public key with a certificate *CA<<CB>>* and *CB* certifies *CA*'s public key with a certificate *CB<<CA>>*, then *A* and *B* can check their certificates by checking a certificate chain:
  - Upon being presented *CB<<B>>* Alice tries to look up if there is a certificate *CA<<CB>>*
  - She then checks the chain: *CA<<CB>>, CB<<B>>*

❑ Certificate chains need not to be limited to a length of two certificates:

- CA<<CC>>, CC<<CD>>, CD<<CE>>, CE<<CG>>, CG<<G>
  would permit Alice to check the certificate of user G issued by CG even if
  she just knows and trusts her own certification authority CA

- In fact, A's trust in the key $K_{G\text{-}priv}$ is established by a *chain of trust* between
  certification authorities

- However, if Alice is presented *CG<<G>>,* it is not obvious which
  certificates she needs for checking it

❑ X.509 therefore suggests that authorities are arranged in a *certification hierarchy,* so that navigation is straightforward:

X.509 – Certificate Revocation (1)

❑ When a certificate is issued, it is expected to be in use for its entire validity period.

❑ However, various circumstances may cause a certificate to become invalid prior to the expiration of the validity period.

❑ Reasons for revoking a certificate:

- The information in the certificate is not valid anymore.

- The private key can not be used anymore, e.g. because
  - the physical medium where the private key was stored becomes defect, e.g. the hard disk, the USB stick or the smart card.
  - the physical medium where the private key is stored has been stolen.
  - the private is protected with a password and the password can not be recovered.

- The private key is (partially) revealed or at least assumed to be revealed, e.g. a Trojan horse or a key logger has been discovered on the computer.

- The parameters of the certificate become inadequate, e.g.
  - The cryptographic algorithm is broken.
  - The key length is considered as inappropriate.

X.509 – Certificate Revocation (2)

❑ An even worse situation occurs if the private key of a certification authority is compromised:

  ▪ This implies that all certificates signed with this key have to be revoked.

❑ Certificate revocation is realized by maintaining *certificate revocation lists (CRL)*:

❑ CRLs are stored in the X.500 directory

❑ Each CA issues a signed data structure periodically called a certificate revocation list (CRL).

→ Certificate revocation is a relatively slow and expensive operation

Online Certificate Status Protocol (OCSP)

❑ The CRL can be accessed with the *Online Certificate Status Protocol (OCSP)*

❑ An OCSP client issues a status request to an OCSP server and suspends acceptance of the certificate in question until the responder provides a response.

❑ CAs that support an OCSP service, either hosted locally or provided by an Authorized Responder, provide the necessary information for the online validation of the status of the certificate.

Simple Certificate Validation Protocol (SCVP)

❑ OCSP makes no claim about the validity of the certificate other than its revocation status.

❑ However, the certificate consumer needs to build the certificate validation path and needs to inquire the revocation status of each of the certificates in the certificates' chain.

❑ The certificate validation process is rather resource-consuming.

❑ Therefore, in some environments, e.g. with cell phones, it would be desirable to fully off-load the certificate validation process to an external trusted entity.

❑ The *Simple Certificate Validation Protocol (SCVP)* [RFC5055] offers this functionality.

PKI - Discussion

❑ PKIs assume a relationship between the CA and the entities, which is not always available:

  ▪ There is no „global" PKI

  ▪ There is no worldwide CA. (But CAs might "cross-certify" each others)

❑ It remains questionable whether a CA executes its task faithfully, i.e., whether a CA verify the identity of the users thoroughly.

❑ In particular, if the CA certifies millions of users.

❑ Nevertheless, PKIs are very commonly used

  ▪ They are integrated, e.g. in each Internet browser

  ▪ Every Internet-Browser has a list of „root CAs" that are considered as trusted.

Overview

- ❑ Part I: Introduction

- ❑ Part II: The Secure Channel

- ❑ Part III: Authentication and Key Establishment Protocols

  - ▪ Introduction

  - ▪ Key Distribution Centers (KDC)

  - ▪ Public Key Infrastructures (PKI)

  - ▪ Building Blocks of key exchange protocols

Problem Statement

(c.f. Niels Ferguson, Bruce Schneier: Practical Cryptography, Ch. 15, pp. 261ff)

❑ Assumption

- Alice and Bob are able to authenticate messages to each other, e.g.
  - Using RSA signatures, if Alice and Bob know each other's public keys or using a PKI
  - Using a long term pre-shared secret key and a MAC function

❑ Goal

- Run a key exchange protocol such as at the end of the protocol:
  1. Alice and Bob have agreed on a shared „session key" for a secure channel
  2. Alice and Bob have agreed on the cryptographic algorithms to be used for the secure channel
  3. Alice (Bob) must be able to verify that Bob (Alice) knows K and that he (she) is "alive"
  4. Alice and Bob must know that K is newly generated

❑ Note: even if Alice and Bob possess a long term pre-shared secret key, it is recommended to perform a key exchange in order to derive a separate session key

❑ Why do we need a session key if we already have a (long term) key?

❑ De-couple the session key from the long-term key
  1. If the session key is compromised, e.g. because of a flawed implementation of the secure channel, then the long-term shared secret should remain safe.
  2. If the long-term key is compromised *after* the key negotiation has been run, the attacker who learns the shared secret key still does not learn the session key negotiated by the protocol, i.e. yesterday's data is still protected if the long-term key is compromised today.
  ▪ These properties are important and make the entire system more robust
  ▪ The 2nd property is called „*Forward Secrecy*"

❑ Definition: Forward Secrecy [Boyd03]
  ▪ A key establishment protocol provides *forward secrecy* if compromise of the long-term keys of a set of entities (private keys or symmetric keys) does not compromise the session keys established in previous protocol runs involving these entities.

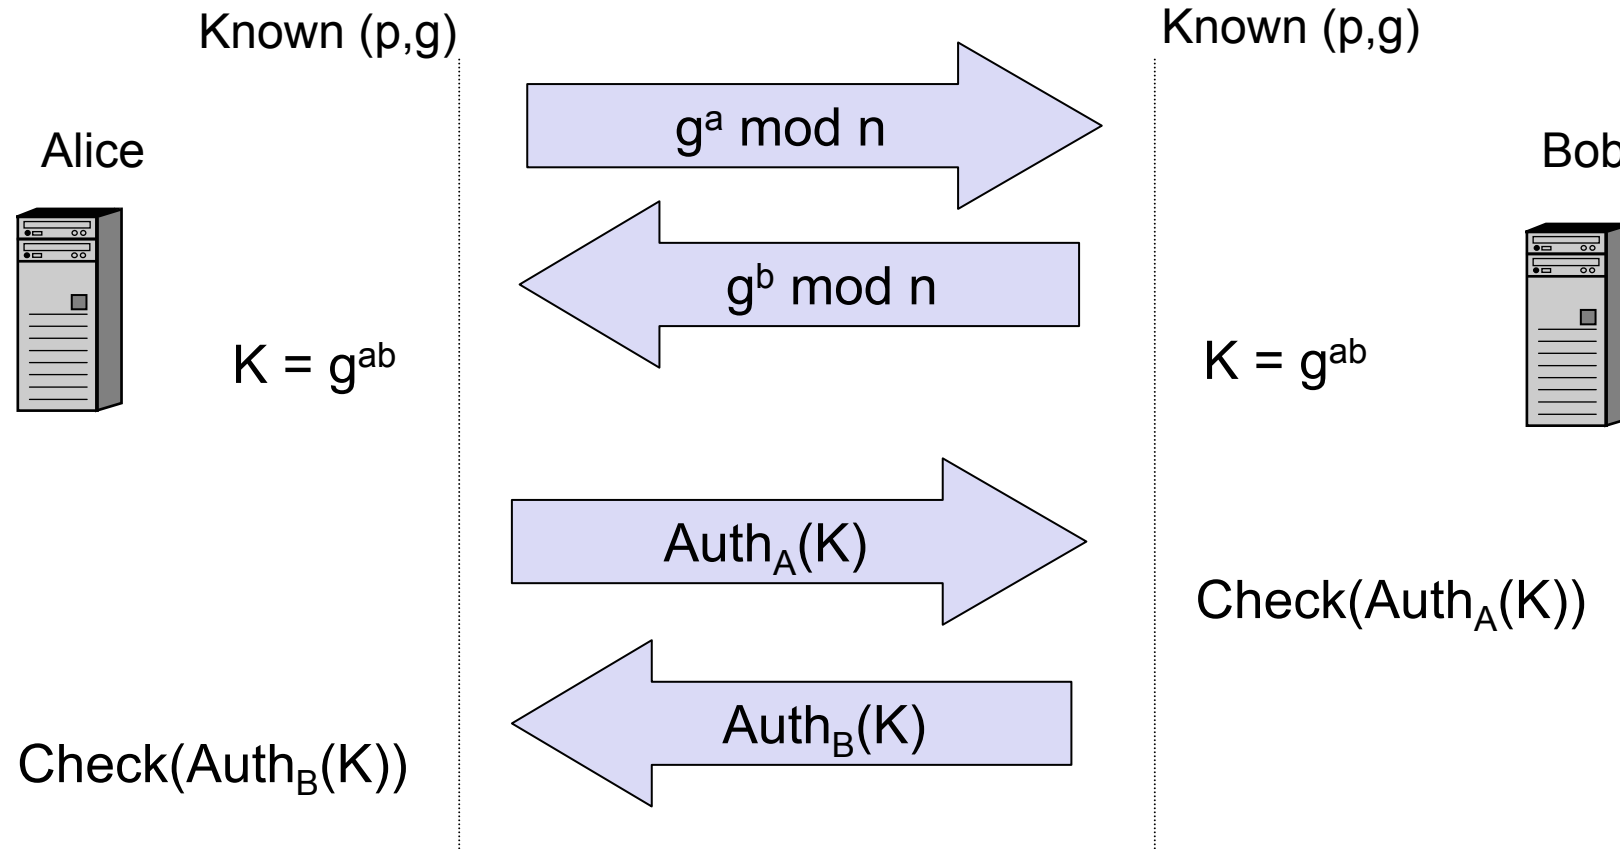**Other Reasons for Separating Session Keys and Long-Term Keys**

❑ Sometimes the long-term key is weak, e.g. passwords

  ▪ Users do not want to memorize a 30-letters password

  ▪ They tend to choose much simpler ones

❑ In some cases, the session key needs to be changed before the session is over (re-keying)

  ▪ This is, e.g., the case if the message sequence numbers overflow and need to be reset

  ▪ This would be problematic if the session key is equal to the long-term key

First Try

- Alice and Bob perform a Diffie-Hellman key exchange and then authenticate the obtained key *k*

Known (p,g)   Known (p,g)

Alice   Bob

$$g^a \bmod n \rightarrow$$

$$\leftarrow g^b \bmod n$$

$$K = g^{ab}$$ $$K = g^{ab}$$

$$\text{Auth}_A(K) \rightarrow$$

$$\text{Check}(\text{Auth}_A(K))$$

$$\leftarrow \text{Auth}_B(K)$$

$$\text{Check}(\text{Auth}_B(K))$$

Problems with "First Try"

❑ Alice and Bob use constant DH parameters *p* and *g*

- This is a bad design, since

  • *p* and *g* might be considered as insecure after a while

  • Protocols live for a long time. Using the same constants raises interoperability issues

❑ The exchange uses 4 messages, whereas it is possible to achieve the goal using 3 messages

❑ *K* is used as input for the authentication function *Auth*

- This would be fine, if *Auth* is a strong function

- But if *Auth*(*K*) leaks some knowledge about *K* this would require a new analysis of the entire protocol

- A rule of thumb: "*Secrets should be used only for a single purpose*".

❑ The authentication messages are too similar

- If *Auth* is a MAC function, then $Auth_B(K) = Auth_A(K)$

➔ Bob can just send the authentication value that he received from Alice.

➔ At the end of the protocol run, Alice can not be sure that Bob has the correct key

Second Attempt

❑ Alice chooses the DH parameters *p* and *g*

  ▪ Bob verifies that he supports *p* and *g*

❑ The protocol exchange is reduced to 2 messages

Alice

Bob

$(p, g, g^a, \text{Auth}_A(p, g, g^a))$

- Check$(p, g, g^a, \text{Auth}_A(p, g, g^a))$
- $k = g^{ab}$

$g^b, \text{Auth}_B(g^b)$

- Check$(g^b, \text{Auth}_B(g^b))$
- $k = g^{ab}$

Second Attempt, Evaluation

❑ DH parameters are chosen dynamically
  ▪ If p is not large enough, Bob can send an error message to Alice with the minimal supported length for p and abort the protocol run

❑ The protocol run requires only 2 messages

❑ The key g^{ab} is not used anymore for the authentication of messages

❑ String that are being authenticated are not the same

❑ However, a replay attack possible
  ▪ Bob can not be sure that he is actually talking to Alice
  ▪ Anybody can record the first message that Alice sends and then later send it to Bob
  ▪ Bob verifies $Auth_A$ and finishes the protocol thinking that he has just shared a session key k with Alice

❑ This problem is called *the lack of liveliness*
  ▪ Bob can not be sure that Alice is „alive", and he is not talking to a replaying attacker
  ▪ The typical way to solve this problem is to make sure that Alice's authenticator $Auth_A$ covers a random value that has been chosen by Bob

Alice

Bob

$N_a, (p,g), g^a$

$N_b, g^b, Auth_B(g^b, N_a)$

$k = g^{ab}$

$k = g^{ab}$

$Auth_A(p,g,g^a, N_b)$

Check($Auth_A$)

Check($Auth_B$)

❑ Alice and Bob need to agree on the cryptographic algorithms to be used for encryption and data integrity

- Facilitates the support of new stronger cryptographic algorithms
- Deprecated cryptographic algorithms can be removed easily
- Upgrades do not require an additional standardization process

Alice

Bob

List of supported crypto algorithms

Chosen crypto algorithms
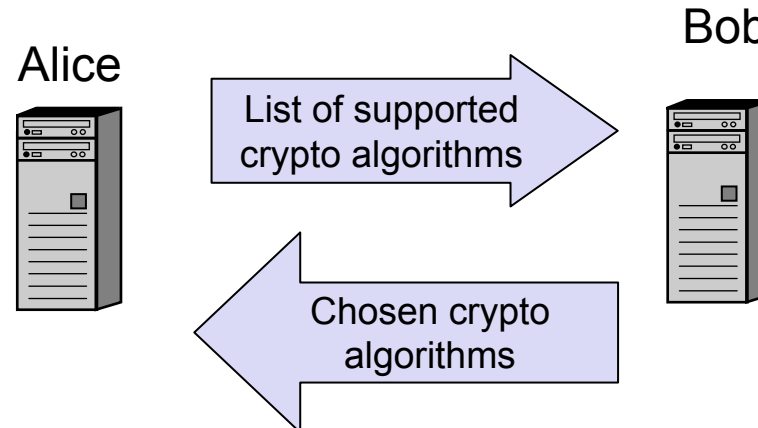
Further Design Issues: Denial-of-Service Protection (1)

❑ Bob may be flooded with a large number of requests for establishing a secure channel from a large number of attackers

❑ This phenomena is called *Distributed Denial-of-Service attacks (DDoS)*

❑ Since Bob needs to store state and perform computation for each request, a DoS attack would exhaust Bob's resources, such as CPU and memory

❑ Possible Countermeasures:

  ▪ Before processing a new request, verify if the "initiator" can receive messages sent to the claimed source of the request (see next slide)

Denial-of-Service Protection with Cookies (1)

- Upon receiving a request from Alice, Bob calculates a Cookie and sends it to Bob.
- Alice will receive the Cookie and resend the request with the Cookie together.
- Bob verifies that the Cookie is correct and then starts to process Alice's request.
- An attacker that is sending requests with a spoofed (i.e. falsified) source address will not be able to send the Cookie.

Denial-of-Service Protection with Cookies (2)

❑ Requirements:

- An attacker that is not on the path between Alice and Bob must not be able to guess the correct value of the Cookie

- Bob must be able to generate the Cookie after receiving message 1 with minimal processing (CPU friendly)

- Bob must be able to verify that the Cookie is correct upon receipt of message 3, without necessarily storing any information after message 1 (memory friendly)

  ➔ Bob must be able to re-calculate the Cookie sent in message 2 and verify that the received Cookie from Alice in message 3 is correct

❑ One possible way to compute the cookie could be as follow:

$Cookie$ = Hash(N_a | $Address_{Alice}$ | <$secret$>)

where

- N_a is the nonce sent by Alice (as above)
- <$secret$> is randomly generated secret known only to Bob
- Hash is a cryptographic hash function.

❑ Only a legitimate initiator (Alice) or a host on the path can read the "cookie" and can send the cookie back to the responder (Bob)

Denial-of-Service Protection with Cookies (3)

❑ Additional requirement:
- *<secret>* needs to be changed regularly. Otherwise, it can be brute-forced successfully after a while

➔ Another possible way to compute the cookie could be as follow:

Cookie = *<Version ID of Secret>* | Hash(N_a | IP_a | *<secret>*)

where
- *<Version ID of Secret>* is changed whenever *<secret>* is regenerated.

❑ Cookies discussion:
- Advantage: allows to counter simple address spoofing attacks
- Drawbacks:
  - requires one additional message roundtrip.

Further Design Issues: Reuse of the DH Values

❑ The calculation if the DH values g^a and g^b is computationally expensive

❑ Alice and Bob may re-use the values g^a and g^b

❑ However, Alice and Bob must ensure that the key has been freshly generated

➔ The random numbers N_a and N_b can be included in the computation of the shared key

➔ One possible way to compute the session key:

   ▪ $K = H(\,N_a\,|\,N_b\,|\,g^{ab}\,)$ where H is a cryptographic hash function

❑ However, the re-use of the DH values affects the property of (perfect) forward secrecy (see next slide).

Forward Secrecy (1)

❑ The DH exchange is not only used to gain a shared secret g^{ab} (that needs to be authenticated).

❑ The DH exchange offers also the property of "forward secrecy"
  ▪ If any long term keys,
    • the long-term pre-shared secret key between Alice and Bob
    • or Alice/Bob private key

  is compromised, an attacker that has recorded previous protocol runs, would need to compromise the DH exchange as well in order to gain the session keys for these previous sessions.

❑ Forward Secrecy was originally called "Perfect Forward Secrecy" (PFS)
  ▪ Many cryptographers did not agree with the term "perfect", so it is usually skipped.

❑ PFS requires that when a session is closed, each endpoint forgets
  ▪ all the keying material used for this session
  ▪ any information that could be used to recompute those keys
  ▪ In particular, it needs to forget the secrets used in the DH calculation and the state of a pseudo-random number generator that could be used to re-compute the DH secrets.

Forward Secrecy (2)

❑ Note

- By running a key exchange protocol, PFS is usually provided with the DH exchange

- Many protocols do not provide PFS, since DH is computationally intensive

- Examples
  - IPSec IKEv (Version 1 and Version 2): yes,
  - TLS: PFS optionally provided with ephemeral (temporary) DH
  - WLAN: WEP, WPA: no
  - GSM/UMTS Authentication and Key Exchange (AKA): no

    (although some commercial products do already support PFS for GMS networks. But both mobile phones need to support it)

Further Design Issues: Simplicity

❑ „A more complex system loses on all fronts. It contains more weaknesses to start with, it is much harder to analyze, and it is much harder to implement without introducing security-critical errors in the implementation." [Fer00]

❑ An important design criterion for a new protocol is that the protocol state machine should be as simple as possible.

❑ Especially for security protocols, the simpler the state machine is the easier the security analysis of the protocol can be.

❑ Remember that an attacker can send any type of message at any time to any participant in the protocol.

❑ One way to reduce the complexity is to design the protocol such as it consists of pairs of messages:

  ▪ a request

  ▪ and a response.

❑ Every request requires a response.

Alice

Bob

$N_a, (p,g)$, g^a, proposed crypto algs

N_b, g^b, chosen crypto algs

$k = g^{ab}$

$k = g^{ab}$

$Auth_A(p,g,g^a$, proposed crypto algs, $N_b)$

Check($Auth_A$)

$Auth_B(g^b$, chosen crypto algs, N_a)

Check($Auth_B$)

References

[Bell95]   M. Bellare and P. Rogaway, Provably Secure Session Key Distribution - The Three Party Case, Proc. 27th STOC, 1995, pp 57--64

[Boyd03] Colin Boyd, Anish Mathuria, "Protocols for Authentication and Key Establishment", Springer, 2003

[Bry88a] R. Bryant. *Designing an Authentication System: A Dialogue in Four Scenes.* Project Athena, Massachusetts Institute of Technology, Cambridge, USA, 1988.

[Diff92]   W. Diffie, P. C. van Oorschot, and M. J. Wiener. Authentication and authenticated key exchanges. Designs, Codes, and Cryptography, 1992

[Dol81a] D. Dolev, A.C. Yao. *On the security of public key protocols.* Proceedings of IEEE 22nd Annual Symposium on Foundations of Computer Science, pp. 350-357, 1981.

[Fer00]   Niels Ferguson, Bruce Schneier, "A Cryptographic Evaluation of IPsec". http://www.counterpane.com/ipsec.pdf 2000

[Fer03]   Niels Ferguson, Bruce Schneier, „Practical Cryptography", John Wiley & Sons, 2003

[Gar03]   Jason Garman, "Kerberos. The Definitive Guide", O'Reilly Media, 1st Edition, 2003

References

[Kau02a]   C. Kaufman, R. Perlman, M. Speciner. *Network Security*. Prentice Hall, 2nd edition, 2002.

[Koh94a]   J. Kohl, C. Neuman, T. T'so, *The Evolution of the Kerberos Authentication System.* In Distributed Open Systems, pages 78-94. IEEE Computer Society Press, 1994.

[Mao04a]   W. Mao. *Modern Cryptography: Theory & Practice*. Hewlett-Packard Books, 2004.

[Nee78]   R. Needham, M. Schroeder. *Using Encryption for Authentication in Large Networks of Computers.* Communications of the ACM, Vol. 21, No. 12, 1978.

[Woo92a]   T.Y.C Woo, S.S. Lam. *Authentication for distributed systems.* Computer, 25(1):39-52, 1992.

[Lowe95]   G. Lowe, „An Attack on the Needham-Schroeder Public-Key Authentication Protocol", *Information Processing Letters*, volume 56, number 3, pages 131-133, 1995.

Additional references from the IETF

[RFC2560]   M. Myers, et al., "X.509 Internet Public Key Infrastructure Online Certificate Status Protocol – OCSP", June 1999

[RFC3961]   K. Raeburn, "Encryption and Checksum Specifications for Kerberos 5", February 2005

[RFC3962]   K. Raeburn, "Advanced Encryption Standard (AES) Encryption for Kerberos 5", February 2005

[RFC4757]   K. Jaganathan, et al., "The RC4-HMAC Kerberos Encryption Types Used by Microsoft Windows ", December 2006

[RFC4120]   C. Neuman, et al., "The Kerberos Network Authentication Service (V5)", July 2005

[RFC4537]   L. Zhu, et al, "Kerberos Cryptosystem Negotiation Extension", June 2006

[RFC5055]   T. Freeman, et al, "Server-Based Certificate Validation Protocol (SCVP)", December 2007