



EEC-484/584 Computer Networks

Lecture 24

Wenbing Zhao

wenbing@ieee.org

(Lecture notes are based on materials supplied by
Dr. Louise Moser at UCSB and Prentice-Hall)



Outline

- Review of last lecture
 - Symmetric-key algorithms
 - Public-key algorithms
 - Digital signatures
- Today's topics
 - Management of public keys
 - Communication security



Symmetric-Key Algorithms

- DES – The Data Encryption Standard
- AES – The Advanced Encryption Standard
- Cipher Modes
- Other Ciphers
- Cryptanalysis



Cipher Modes

- Despite all the complexity, AES and DES (or any block cipher) is basically a monoalphabetic substitution cipher using big characters
- Electronic Code Book Mode
- Cipher Block Chaining Mode
- Cipher Feedback Mode
- Stream Cipher Mode
- Counter Mode



RSA

- Rivest, Shamir, Adleman, 1978: a good method for public-key cryptography
- RSA method:
 - Choose two large primes, p and q (typically 1024 bits)
 - Compute $n = p \times q$ and $z = (p-1) \times (q-1)$
 - Choose a number relatively prime to z and call it d
 - Find e such that $e \times d = 1 \pmod{z}$
- To encrypt a message, P , Compute $C = P^e \pmod{n}$
- To decrypt C , compute $P = C^d \pmod{n}$
- The public key consists of the pair (e, n)
- The private key consists of the pair (d, n)

Digital Signatures

- Requirement on **digital signatures**: one party can send a signed message to another party in such a way that the following conditions hold:
 - The receiver can verify the claimed identity of the sender.
 - The sender cannot later repudiate the contents of the message
 - The receiver cannot possibly have concocted the message himself

Digital Signatures



Message Digests

- Often, authentication is needed but secrecy is not
- **Message digest (MD)**: using a one-way hash function that takes an arbitrarily long piece of plaintext and from it computes a fixed-length bit string
 - Given P , it is easy to compute $MD(P)$
 - Given $MD(P)$, it is effectively impossible to find P
 - Given P no one can find P' such that $MD(P') = MD(P)$
 - A change to the input of even 1 bit produces a very different output

Message Digests

- Digital signatures using message digests




MD5

- One of the most widely used hash functions
- MD5 is the fifth in a series of message digests designed by Ronald Rivest (1992)
- It operates by mangling bits in a sufficiently complicated way that every output bit is affected by every input bit
- MD5 generates a 128-bit fixed value

SHA-1

- **SHA-1: Secure Hash Algorithm 1**, developed by National Security Agency (NSA) and blessed by NIST. It generates 160-bit message digest
- Use of SHA-1 and RSA for signing nonsecret messages



Need for Better Hash Functions

- Both MD5 and SHA-1 have found to have weakness
 - It takes much less time to find two plaintexts that hash to the same value (i.e., collision) than expected
 - Research was carried out by a research group lead by a Chinese professor, [Xiaoyun Wang, Shandong University](http://www.infosec.sdu.edu.cn/people/wangxiaoyun.htm)
 - Interested students can visit the following URL for details: <http://www.infosec.sdu.edu.cn/people/wangxiaoyun.htm>

The Birthday Attack

- How many operations it takes to find a collision for an m -bit message digest?
 - Only $2^{m/2}$, **NOT** 2^m , using the **birthday attack**
 - For MD5, need 2^{64} operations
 - For SHA-1, need 2^{80} operations => **actually** $\leq 2^{69}$ ops with latest discovery
- **Birthday attack.** How many students do you need in a class before the probability of having two people with the same birthday exceeds $1/2$? => 23
 - Intuitively, with 23 people, we can form $(23 \times 22) / 2 = 253$ different pairs, each of which has a probability of $1/365$ of being a hit

The Birthday Attack

- Generally, if there is some mapping between inputs and outputs with n inputs and k possible outputs,
 - There are $n(n-1)/2$ input pairs
 - If $n(n-1)/2 > k$, the chance of having at least one match is pretty good
 - Thus, approximately, a match is likely for $n > \text{sqr}(k)$

The Birthday Attack

- Example: construct two letters with same hash value
 - I have [*known* / *worked with*] Prof. Wilson for [*about* / *almost*] six years. He is an [*outstanding* / *excellent*] researcher of great [*talent* / *ability*] known [*worldwide* / *internationally*] for his [*brilliant* / *creative*] insights into [*many* / *a wide variety of*] [*difficult* / *challenging*] problems.
 - I have [*known* / *worked with*] Tom for [*about* / *almost*] six years. He is a [*poor* / *weak*] researcher not well known in his [*field* / *area*]. His research [*hardly ever* / *rarely*] shows [*insight in* / *understanding of*] the [*key* / *major*] problems of [*the* / *our*] day.


Management of Public Keys

- Certificates
- X.509
- Public key infrastructure

Problems with Public-Key Management

17

- If Alice and Bob do not know each other, how do they get each other's public keys to start the communication process ?
 - It is essential Alice gets Bob's public key, not someone else's
- A way for Trudy to subvert public-key encryption



9 December 2005

EEC484/584

Wenbing Zhao

Certificates

18

- **Certification Authority (CA):** an organization that certifies public keys
 - It certifies the public keys belonging to people, companies, or even attributes
 - CA does not need to be on-line all the time
- A possible certificate and its signed hash

```

    I hereby certify that the public key
    19836A8B03030CF83737E3837837FC3s87092827262643FFA82710382828282A
    belongs to
    Robert John Smith
    12345 University Avenue
    Berkeley, CA 94702
    Birthday: July 4, 1958
    Email: bob@superdupernet.com

    -----
    SHA-1 hash of the above certificate signed with the CA's private key
    
```

9 December 2005

EEC484/584

Wenbing Zhao

X.509

19

- Devised and approved by ITU
- The basic fields of an X.509 certificate

Field	Meaning
Version	Which version of X.509
Serial number	This number plus the CA's name uniquely identifies the certificate
Signature algorithm	The algorithm used to sign the certificate
Issuer	X.509 name of the CA
Validity period	The starting and ending times of the validity period
Subject name	The entity whose key is being certified
Public key	The subject's public key and the ID of the algorithm using it
Issuer ID	An optional ID uniquely identifying the certificate's issuer
Subject ID	An optional ID uniquely identifying the certificate's subject
Extensions	Many extensions have been defined
Signature	The certificate's signature (signed by the CA's private key)

9 December 2005

EEC484/584

Wenbing Zhao

Public-Key Infrastructures

20

- A **Public-Key Infrastructure (PKI)** is needed for reasons of
 - Availability, Scalability, Ease of management
- A PKI has multiple components
 - Users, CAs, Certificates, Directories
- A PKI provides a way of structuring these components and define standards for the various documents and protocols
 - A simple form of PKI is hierarchical CAs

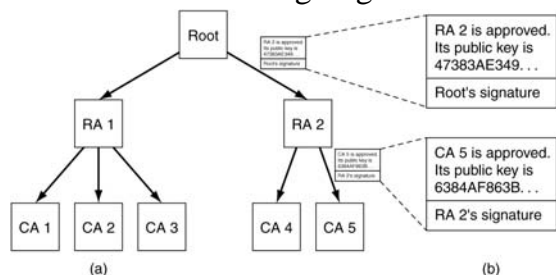
9 December 2005

EEC484/584

Wenbing Zhao

Public-Key Infrastructures

- Hierarchical PKI
- A **chain of trust/certification path**:
A chain of certificates going back to the root



9 December 2005

EEC484/584

Wenbing Zhao

Public-Key Infrastructures

- **Directories**: where certificates (and their chains back to some known trust anchor) are stored
 - Each user store his or her own certificates. While doing this is safe (i.e., there is no way for users to tamper with signed certificates without detection), it is also inconvenient
 - Use DNS as a certificate directory
 - Dedicated directory servers whose only job is managing X.509 certificates

9 December 2005

EEC484/584

Wenbing Zhao

Public-Key Infrastructures

- **Revocation**: sometimes certificates can be revoked, due to a number of reasons
 - Person or organization holding it has abused it in some way
 - The subject's private key has been exposed
 - The CA's private key has been compromised
- **Reinstatement**: a revoked certificate could conceivably be reinstated
 - E.g., if it was revoked for nonpayment of some fee that has since been paid

9 December 2005

EEC484/584

Wenbing Zhao

Public-Key Infrastructures

- Each CA periodically issue a **CRL (Certificate Revocation List)** giving the serial numbers of all certificates that it has revoked
 - A user who is about to use a certificate must now acquire the CRL to see if the certificate has been revoked
- Having to deal with revocation (and possibly reinstatement) eliminates one of the best properties of certificates, namely, that they can be used without having to contact a CA
- Where should CRLs be stored?
 - A good place would be the same place the certificates themselves are stored

9 December 2005

EEC484/584

Wenbing Zhao

Communication Security

- IPsec
- Firewalls
- Virtual private networks
- Wireless security

IPsec

- **IPsec (IP security):** a solution for Internet security
 - Described in RFCs 2401, 2402, and 2406, among others
 - Not all users want encryption (because it is computationally expensive). Rather than make it optional, it was decided to require encryption all the time but permit the use of a **null algorithm**

IPsec

- The complete IPsec design is a framework for multiple services, algorithms and granularities
 - The major services are **secrecy, data integrity, and protection from replay attacks** (intruder replays a conversation)
 - All of these are based on **symmetric-key cryptography** because high performance is crucial

IPsec

- Why multiple services ?
 - Not everyone wants to pay the price for having all services all the time
- Why multiple algorithms ?
 - An algorithm now thought to be secure may be broken in the future
 - By making IPsec algorithm-independent, the framework can survive even if some particular algorithm is later broken

IPsec

- Why multiple granularities ? To make it possible to protect
 - A single TCP connection,
 - All traffic between a pair of hosts, or,
 - All traffic between a pair of secure routers, among other possibilities

IPsec

- IPsec is connection oriented
- **Security Association (SA):** A simplex connection between two end points and has a security identifier associated with it
 - If secure traffic is needed in both directions, two security associations are required
 - **Security identifiers** are carried in packets traveling on these secure connections and are used to look up keys and other relevant
- IPsec can be used in either of two modes
 - Transport mode and Tunnel mode

IPsec – Transport Mode

- In **transport mode**, the IPsec header is inserted just after the IP header
 - The *Protocol* field in the IP header is changed to indicate that an IPsec header follows the normal IP header (before the TCP header)
 - The IPsec header contains security information, primarily the SA identifier, a new sequence number, and possibly an integrity check of the payload

IPsec – Tunnel Mode


- In **tunnel mode**, the entire IP packet, header and all, is encapsulated in the body of a new IP packet with a completely new IP header
 - Typically the whole IP packet is encrypted
- Tunnel mode is useful when the tunnel ends at a location other than the final destination, e.g., a firewall
 - In this mode, the firewall encapsulates and decapsulates packets as they pass through the firewall
 - By terminating the tunnel at a secure machine, the machines on the company LAN do not have to be aware of IPsec

IPsec – Tunnel Mode

- Tunnel mode is also useful when a bundle of TCP connections is aggregated and handled as one encrypted stream
 - It prevents an intruder from seeing who is sending how many packets to whom
 - Sometimes just knowing how much traffic is going where is valuable information

IPsec – Authentication Header

- **AH (Authentication Header):** It provides integrity checking and antireplay security, but not secrecy (i.e., no data encryption)



IPsec - AH header

- The *Next header field* is used to store the previous value that the IP *Protocol* field had before it was replaced with 51 to indicate that an AH header follows.
- The *Payload length* is the number of 32-bit words in the AH header minus 2.
- The *Security parameters index* is the connection identifier. It is inserted by the sender to indicate a particular record in the receiver's database. This record contains the shared key used on this connection and other information about the connection.

IPsec - AH header

- The *Sequence number* field is used to number all the packets sent on an SA. **Every packet gets a unique number, even retransmissions**
 - The purpose of this field is to detect **replay attacks**
 - These sequence numbers may not wrap around. If all 2^{32} are exhausted, a new SA must be established to continue communication.

IPsec - AH header


- The **Authentication data**, which is a variable-length field that contains the payload's digital signature
 - When the SA is established, the two sides negotiate which signature algorithm they are going to use and the shared key to use
 - The shared key is also used in the signature computation
 - Compute the hash over the packet plus the shared key. The shared key is not transmitted
 - A scheme like this is called an **HMAC (Hashed Message Authentication Code)**
 - It is much faster to compute than first running SHA-1 and then running RSA on the result

IPsec - AH header

- The **Authentication data** also provides integrity check covers some of the fields in the IP header, namely, those that do not change as the packet moves from router to router
 - The *Time to live* field changes on each hop, for example, so it cannot be included in the integrity check
 - However, the IP source address is included in the check, making it impossible for an intruder to falsify the origin of a packet

IPsec – ESP Header

- **ESP (Encapsulating Security Payload)**: provides both authentication and confidentiality guarantee for packets. Can be used for both transport mode and tunnel mode




IPsec – ESP Header

- The ESP header consists of two 32-bit words
 - They are the *Security parameters index* and *Sequence number* fields that we saw in AH
 - A third word that generally follows them (but is technically not part of the header) is the *Initialization vector* used for the data encryption, unless null encryption is used, in which case it is omitted

Firewalls

- A firewall consisting of two packet filters and an application gateway



42

Firewalls

- Each **packet filter** is a standard router equipped with some extra functionality
- The extra functionality allows every incoming or outgoing packet to be inspected
 - Packets meeting some criterion are forwarded normally
 - Those that fail the test are dropped

9 December 2005

EEC484/584

Wenbing Zhao

Firewalls

- Packet filters are typically driven by tables configured by the system administrator
 - These tables list sources and destinations that are acceptable
 - Sources and destinations that are blocked, and
 - Default rules about what to do with packets coming from or going to other machines

43

9 December 2005

EEC484/584

Wenbing Zhao

Firewalls

- Firewalls cannot solve all the security problems
 - If a firewall is configured to allow in packets from only specific networks (e.g., the company's other plants), an intruder outside the firewall can put in false source addresses to bypass this check
 - If an insider wants to ship out secret documents, he can encrypt them or even photograph them and ship the photos as JPEG files, which bypasses any word filters
 - 70% of all attacks come from inside the firewall, for example, from disgruntled employees
 - Firewalls cannot cope with **DoS** attacks

44

9 December 2005

EEC484/584

Wenbing Zhao

Denial of Service Attack


- **DoS (Denial of Service)** attacks: Attacks in which the intruder's goal is to shut down the target rather than steal data
 - Usually, the request packets have false source addresses so the intruder cannot be traced easily
 - An intruder can send a TCP *SYN* packet to establish a connection
 - The site will then allocate a table slot for the connection and send a *SYN + ACK* packet in reply
 - If the intruder does not respond, the table slot will be tied up for a few seconds until it times out

DDoS

- **DDoS (Distributed Denial of Service)** attack
 - Intruder has already broken into hundreds of computers elsewhere in the world
 - Intruder then commands all of them to attack the same target at the same time
 - Not only does this approach increase the intruder's firepower, it also reduces his chance of detection, since the packets are coming from a large number of machines belonging to unsuspecting users

Virtual Private Networks

- **VPN (Virtual Private Network):** overlay network on top of public networks but with most of the properties of private networks



A leased-line private network

A virtual private network:
IPsec is typically used

802.11 Security


- **WEP (Wired Equivalent Privacy):** a data link-level security protocol prescribed by 802.11 standard
 - It is designed to make the security of a wireless LAN as good as that of a wired LAN
 - Many base stations come with security disabled by default

802.11 Security

- When 802.11 security is enabled, each station has a secret key shared with the base station
 - WEP encryption uses a stream cipher based on the RC4 algorithm
 - In WEP, RC4 generates a keystream that is XORed with the plaintext to form the ciphertext

Packet Encryption Using WEP

- Payload is checksummed using CRC-32 polynomial
- Checksum is appended to the payload to form the plaintext
- This plaintext is XORed with a chunk of keystream its own size. The result is the ciphertext
- The IV used to start RC4 is sent along with the ciphertext



WEP Weakness

- Many installations use the same shared key for all users
- Many 802.11 cards reset IV to 0 when the card is inserted into computer, and increment it by one on each packet sent
 - Packets with low IV values are common => subject to the keystream reuse attack

WEP Weakness

- IVs are only 24 bits, so after 2^{24} packets have been sent, they have to be reused
 - With randomly chosen IVs, the expected number of packets that have to be sent before the same one is used twice is about 5000, due to birthday attack
- Since IVs are used at random, once an intruder has determined a valid (IV, keystream) pair, he/she can use it to generate traffic that interfere with normal communication.
- RC4 is weak

Bluetooth Security

- Bluetooth provides security in multiple layers
- Physical layer, frequency hopping provides a tiny bit of security
- To establish a channel, the slave and master each check to see if the other one knows the **passkey** (i.e., the shared key between the slave and master)
 - If so, they negotiate whether that channel will be encrypted, integrity controlled, or both
 - Then they select a random 128-bit session key, some of whose bits may be public

Bluetooth Security

- Encryption uses a stream cipher called **E₀**; integrity control uses **SAFER+**. Both are traditional symmetric-key block ciphers
 - E₀ may have fatal weaknesses
 - SAFER+ slower than Rijndael, eliminated in the first round of the NIST contest