

Security and secure systems



COS 518: *Advanced Computer Systems*
Lecture 18

Michael Freedman

Intro to crypto in 10 minutes

2

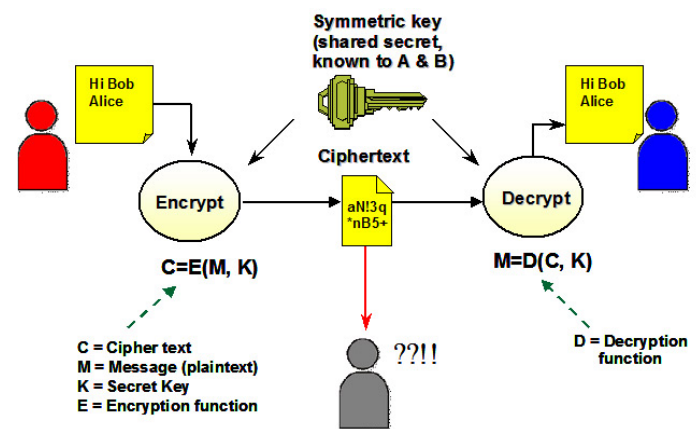
What is Cryptography?

- From Greek, meaning “secret writing”
- Confidentiality: encrypt data to hide content
- Include “signature” or “message authentication code”
 - Integrity: Message has not been modified
 - Authentication: Identify source of message

plaintext $\xrightarrow{\text{encryption}}$ ciphertext $\xrightarrow{\text{decryption}}$ plaintext

- Modern encryption:
 - *Algorithm* public, *key* secret and provides security
 - Symmetric (shared secret) or asymmetric (public-private key)

Symmetric Cipher Model



4

Symmetric (Secret Key) Crypto

- Sender and recipient share common key
 - **Main challenge: How to distribute the key?**
- Provides dual use:
 - Confidentiality (encryption)
 - Message authentication + integrity (MAC)
- 1000x more computationally efficient than asymmetric

5

Public-Key Cryptography

- **Each party has (public key, private key)**
- **Alice's public key PK**
 - Bob uses PK to encrypt messages *to* Alice
 - ciphertext = encrypt (message, PK)
 - Bob uses PK to verify signatures *from* Alice
 - isValid = verify (signature, message, PK)
- **Alice's private/secret key: sk**
 - Alice uses sk to decrypt ciphertexts sent to her
 - message = decrypt (ciphertext, sk)
 - Alice uses sk to generate new signatures on messages
 - signature = sign (message, sk)

6

(Simple) RSA Algorithm

- **Generating a key:**
 - Generate composite $n = p * q$, where p and q are secret primes
 - Pick public exponent **e**
 - Solve for secret exponent **d** in $d * e \equiv 1 \pmod{(p-1)(q-1)}$
 - Public key = (e, n), private key = d
- Encrypting message m: $c = m^e \pmod n$
- Decrypting ciphertext c: $m = c^d \pmod n$
- **Security** due to cost of factoring large numbers
 - Finding (p,q) given n takes $O(e^{\log n \log \log n})$ operations
 - n chosen to be 2048 or 4096 bits long

7

Cryptographic hash function

(and using them in systems)

8

Cryptography Hash Functions I

- Take message m of arbitrary length and produces fixed-size (short) number $H(m)$
- **One-way function**
 - Efficient: Easy to compute $H(m)$
 - **Hiding property:** Hard to find an m , given $H(m)$
 - Assumes “ m ” has sufficient entropy, not just {“heads”, “tails”}
 - **Random:** Often assumes for output to “look” random


9

Cryptography Hash Functions II

- Collisions exist: | possible inputs | \gg | possible outputs |
... but hard to find
- **Collision resistance:**
 - Strong resistance: Find any $m \neq m'$ such that $H(m) == H(m')$
 - Weak resistance: Given m , find m' such that $H(m) == H(m')$
 - For 160-bit hash (SHA-1)
 - Finding any collision is birthday paradox: $2^{\{160/2\}} = 2^{80}$
 - Finding specific collision requires 2^{160}


10

Hash Pointers



11


Self-certifying names



- P2P file sharing software (e.g., Limewire)
 - File named by $F_{name} = H(\text{data})$
 - Participants verify that $H(\text{downloaded}) == F_{name}$

12


Self-certifying names



- BitTorrent
 - Large file split into smaller chunks (~256KB each)
 - Torrent file specifies the name/hash of each chunk
 - Participants verify that $H(\text{downloaded}) == C_{\text{name}}$
 - Security relies on getting torrent file from trustworthy source

13


Hash chains



Creates a “tamper-evident” log of data

14

Hash chains



If data changes, all subsequent hash pointers change
Otherwise, found a hash collision!

15

Examples:
HTTP and DNS security


16

“Securing” HTTP

- Threat model
 - Eavesdropper listening on conversation (confidentiality)
 - Man-in-the-middle modifying content (integrity)
 - Adversary impersonating desired website (authentication, and confidentiality)
- Enter HTTP-S
 - HTTP sits on top of secure channel (SSL/TLS)
 - All (HTTP) bytes written to secure channel are encrypted and authenticated
 - **Problem:** What is actually authenticated to prevent impersonation? Which keys used for crypto protocols?

17

Learning a valid public key



- What is that lock?
 - Securely binds domain name to public key (PK)
 - Believable only if you trust the attesting body
 - Bootstrapping problem: Who to trust, and how to tell if this message is actually from them?
 - If PK is authenticated, then any message signed by PK cannot be forged by non-authorized party

18

How to authenticate PK

General Details

This certificate has been verified for the following uses:
SSL Server Certificate

Issued To

Common Name (CN)	www.wellsfargo.com
Organization (O)	Wells Fargo and Company
Organizational Unit (OU)	ISG
Serial Number	41:C5:CD:90:95:3C:A1:4B:C1:8A:

Issued By

Common Name (CN)	<Not Part Of Certificate>
Organization (O)	VeriSign Trust Network
Organizational Unit (OU)	VeriSign, Inc.

Validity

Issued On	5/12/10
Expires On	5/13/11

Fingerprints

SHA1 Fingerprint	C5:EC:18:24:50:9D:90:93:96:69:
MD5 Fingerprint	1C:51:99:C9:EA:7B:FB:64:3F:92:F

Certificate Hierarchy

- Builtin Object Token:Verisign Class 3 Public Primary Certific
- VeriSign, Inc.
- www.wellsfargo.com

Certificate Fields

- Not After
- Subject
- Subject Public Key Info
 - Subject Public Key Algorithm
 - Subject's Public Key

Extensions

- Certificate Basic Constraints
- Certificate Key Usage
- CRL Distribution Points

Field Value

Modulus (1024 bits):
c9 b3 e9 c0 4a 42 be 1a c4 0a a0 b5 e0 9c 79 89
52 82 b1 89 b3 82 dc 2d 03 2b 1e 77 c3 4c 7d 97
37 e2 c6 7b 31 b5 6b 25 d3 9e 7e 7e 07 95 7e e6
ab 6a 5c 88 ec 27 9d 72 3e a0 80 0c a5 ea d4 ef

22

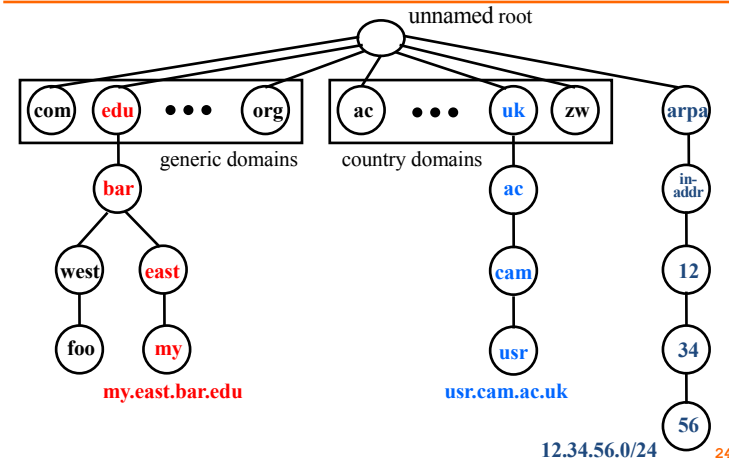
The trouble with CAs

- Browse/OS vendors pick which CAs to trust
 - Sometimes they revoke this trust – e.g. DigiNotar.
- No notion of CAs having authority over only given TLD
- Trust the {Iranian, Chinese, US} national authorities?
- What standards does Apple use to pick root certs? Google? MSFT?
 - There's a restraint-of-trade issue here. Can't enter the CA business without vendor support...

DNS Security

23

Hierarchical naming in DNS



12.34.56.0/24 24

DNS Integrity: Trust the TLD operators?

- If domain name doesn't exist, DNS should return NXDOMAIN (non-existent domain) msg
- Verisign instead creates wildcard DNS record for all [.com](#) and [.net](#) domain names not yet registered
 - September 15 – October 4, 2003
- Redirection for these domain names to Verisign web portal: “to help you search”
 - and serve you ads...and get “sponsored” search
 - Verisign and online advertising companies make money..

25

DNS Integrity:

Answer from authoritative server?

- DNS cache poisoning
 - Client asks for `www.evil.com`
 - Nameserver authoritative for `www.evil.com` returns additional section for (`www.cnn.com`, `1.2.3.4`, A)
 - Thanks! I won't bother check what I asked for

26

DNS Integrity: Answer from authoritative server?

- To prevent cache poisoning, client remembers domain and 16-bit request ID (used to demux UDP response)
- But...DNS hijacking attack:
 - 16 bits: 65K possible IDs
 - What rate to enumerate all in 1 sec? 64B/packet
 - $64 * 65536 * 8 / 1024 / 1024 = 32$ Mbps
 - Prevention: Also randomize the DNS source port
 - Windows DNS alloc's 2500 DNS ports: ~164M possible IDs
 - Would require 80 Gbps
 - Kaminsky attack: this source port...wasn't random after all

27

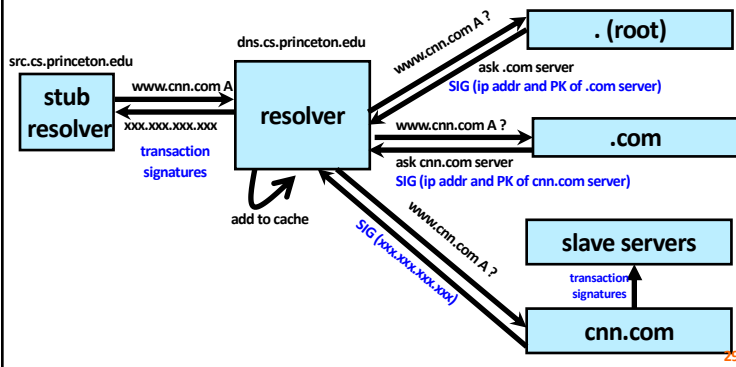
Let's strongly believe the answer! Enter DNSSEC

- The DNS servers sign the hash of resource record set with its private (signature) keys
- Public keys can be used to verify the SIGs
- Leverages hierarchy:
 - Authenticity of nameserver's public keys is established by a signature over the keys by the parent's private key
 - In ideal case, only roots' public keys need to be distributed out-of-band

28

Verifying the tree

Question: **www.cnn.com** ?



29