

Security Engineering (4)

Email: christian.urban at kcl.ac.uk

Office: N7.07 (North Wing, Bush House)

Slides: KEATS (also home work is there)



last week: buffer overflow attacks

- this required some cheating on a modern OS
- but the main point: no cheating needed in practice (remember the quote about toasters)

Case-In-Point: Android

- a list of common Android vulnerabilities (5 BOAs out of 35 vulnerabilities; all from 2013 and later):

<http://androidvulnerabilities.org/>

- a paper that attempts to measure the security of Android phones:

“We find that on average 87.7% of Android devices are exposed to at least one of 11 known critical vulnerabilities...”

<https://www.cl.cam.ac.uk/~drt24/papers/spsm-scoring.pdf>

Survey at KEATS


Thanks!

Two General Counter Measures against BOAs etc

Both try to reduce the attack surface (trusted computing base):

- **unikernels** – the idea is to not have an operating system at all
- all functionality of the server is implemented in a single, stand-alone program
- all functionality an operating system would normally provide (network stack, file system) is available through libraries
- the best known unikernel is MirageOS using Ocaml (<https://mirage.io>)

Network Applications: Privilege Separation



- the idea is make the attack surface smaller and mitigate the consequences of an attack

Access Control in Unix

- access control provided by the OS
- authenticate principals
- mediate access to files, ports, processes etc according to **roles** (user ids)
- roles get attached with privileges (some special roles: root)

principle of least privilege:

users and programs should only have as much privilege as they need to accomplish a task

Access Control in Unix (2)

- privileges are specified by file access permissions (“everything is a file”)
- there are 9 (plus 2) bits that specify the permissions of a file

- r-- rw- rwx bob staff file
directory user group other

Unix-Style Access Control

- Q: “I am using Windows. Why should I care?”

A: In Windows you have similar AC:

administrators group

(has complete control over the machine)

authenticated users

server operators

power users

network configuration operators

- Modern versions of Windows have more fine-grained AC than Unix; they do not have a setuid bit, but have runas (asks for a password).

Weaknesses of Unix AC

Not just restricted to Unix:

- if you have too many roles (i.e. too finegrained AC), then hierarchy is too complex
you invite situations like...let's be root
- you can still abuse the system...

A “Cron”-Attack

The idea is to trick a privileged person to do something on your behalf:

- root:

```
rm /tmp/*/*
```

A “Cron”-Attack

The idea is to trick a privileged person to do something on your behalf:

- root:

```
rm /tmp/*/*
```

the shell behind the scenes:

```
rm /tmp/dir1/file1 /tmp/dir1/file2 /tmp/dir2/file1 ...
```

this takes time

A “Cron”-Attack

- 1 attacker (creates a fake passwd file)

```
mkdir /tmp/a; cat > /tmp/a/passwd
```

- 2 root (does the daily cleaning)

```
rm /tmp/*/*
```

records that /tmp/a/passwd

should be deleted, but does not do it yet

- 3 attacker (meanwhile deletes the fake passwd file, and establishes a link to the real passwd file)

```
rm /tmp/a/passwd; rmdir /tmp/a;
```

```
ln -s /etc /tmp/a
```

- 4 root now deletes the real passwd file

A “Cron”-Attack

- 1 attacker (creates a fake passwd file)
`mkdir /tmp/a; cat > /tmp/a/passwd`

- 2 root
rm
To prevent this kind of attack, you need additional policies (for example don't do such operations as root).

should be deleted, but does not do it yet

- 3 attacker (meanwhile deletes the fake passwd file, and establishes a link to the real passwd file)
`rm /tmp/a/passwd; rmdir /tmp/a;`
`ln -s /etc /tmp/a`
- 4 root now deletes the real passwd file

Subtleties

- Can Bob write file?

`- r--rw-rwx bob staff file`

directory user group other

Subtleties

- Can Bob write file?
- What if Bob is member of staff?

`- r--rw-rwx bob staff file`
directory user group other

Login Processes

- login processes run under $\text{UID} = 0$

```
ps -axl | grep login
```

- after login, shells run under $\text{UID} = \text{user}$ (e.g. 501)

```
id cu
```

Login Processes

- login processes run under $\text{UID} = 0$

```
ps -axl | grep login
```

- after login, shells run under $\text{UID} = \text{user}$ (e.g. 501)

```
id cu
```

- non-root users are not allowed to change the UID — would break access control
- but needed for example for accessing passwd

Setuid and Setgid

The solution is that Unix file permissions are 9 + 2 Bits: **Setuid** and **Setgid** bits

- When a file with setuid is executed, the resulting process will assume the UID given to the owner of the file.
- This enables users to create processes as root (or another user).
- Essential for changing passwords, for example.

```
chmod 4755 fobar_file
```

Discretionary Access Control

- Access to objects (files, directories, devices, etc.) is permitted based on user identity. Each object is owned by a user. Owners can specify freely (at their discretion) how they want to share their objects with other users, by specifying which other users can have which form of access to their objects.
- Discretionary access control is implemented on any modern multi-user OS (Unix, Windows NT, etc.).

Mandatory Access Control

- Access to objects is controlled by a system-wide policy, for example to prevent certain flows of information. In some forms, the system maintains security labels for both objects and subjects (processes, users) based on which access is granted or denied. Labels can change as the result of an access. Security policies are enforced without the cooperation of users or programs.
- This is implemented in banking or military operating system versions (SELinux).

Mandatory Access Control

- Access to objects is controlled by a system-wide policy, for example to prevent certain flows of information. In some forms, the system maintains security labels for both objects and subjects (processes, users) based on which access is granted or denied. Labels can change as the result of an access. Security policies are enforced without the cooperation of users or programs.
- This is implemented in banking or military operating system versions (SELinux).
- A simple example: Air Gap Security. Uses a completely separate network and computer hardware for different application classes (Bin Laden, Bruce Schneier had airgaps).

Mandatory Access Control

- Access to objects is controlled by a system-wide policy, for example to prevent certain flows of information. In some forms, the system maintains security labels for both objects and subjects (processes, users) based on which access is granted or denied. Labels can change as the result of an access. Security policies are enforced without the cooperation of users or programs.
- This is implemented in banking or military operating system versions (SELinux).
- A simple example: Air Gap Security. Uses a completely separate network and computer hardware for different application classes (Bin Laden, Bruce Schneier had airgaps).
- What do we want to protect: Secrecy or Integrity?

The Bell-LaPadula Model

- Formal policy model for mandatory access control in a military multi-level security environment. All subjects (processes, users, terminals, files, windows, connections) are labeled with a confidentiality level, e.g.

unclassified < confidential < secret < top secret

- The system policy automatically prevents the flow of information from high-level objects to lower levels. A process that reads top secret data becomes tagged as top secret by the operating system, as will be all files into which it writes afterwards.

Bell-LaPadula

- **Read Rule:** A principal P can read an object O if and only if P 's security level is at least as high as O 's.
- **Write Rule:** A principal P can write an object O if and only if O 's security level is at least as high as P 's.

This restricts information flow \Rightarrow military

Bell-LaPadula

- **Read Rule:** A principal P can read an object O if and only if P 's security level is at least as high as O 's.
- **Write Rule:** A principal P can write an object O if and only if O 's security level is at least as high as P 's.

This restricts information flow \Rightarrow military

Bell-LaPadula: **'no read up'** - **'no write down'**

Principle of Least Privilege

A principal should have as few privileges as possible to access a resource.

- Bob (T_S) and Alice (S) want to communicate
⇒ Bob should lower his security level

Biba Policy

Data Integrity (rather than data secrecy)

- Biba: **'no read down'** - **'no write up'**
- **Read Rule:** A principal P can read an object O if and only if P 's security level is lower or equal than O 's.
- **Write Rule:** A principal P can write an object O if and only if O 's security level is lower or equal than P 's.

Biba Policy

Data Integrity (rather than data secrecy)

- Biba: **'no read down'** - **'no write up'**
- **Read Rule:** A principal P can read an object O if and only if P 's security level is lower or equal than O 's.
- **Write Rule:** A principal P can write an object O if and only if O 's security level is lower or equal than P 's.

E.g. Firewalls: you can read from inside the firewall, but not from outside

Phishing: you can look at an approved PDF, but not one from a random email


Security Levels (2)

- Bell-La Padula preserves data secrecy, but not data integrity

Security Levels (2)

- Bell-La Padula preserves data secrecy, but not data integrity
- Biba model is for data integrity
 - read: your own level and above
 - write: your own level and below

Shared Access Control



To take an action you need at least either:


- 1 CEO
- 2 MDs, or
- 3 Ds

Lessons from Access Control

Not just restricted to Unix:


- if you have too many roles (i.e. too finegrained AC), then hierarchy is too complex
you invite situations like...lets be root
- you can still abuse the system...

Protocols




- Other examples: Wifi, Http-request, TCP-request, card readers, RFID (passports)...

Protocols



- Other examples: Wifi, Http-request, TCP-request, card readers, RFID (passports)...
- The point is that we cannot control the network: An attacker can install a packet sniffer, inject packets, modify packets, replay messages...fake pretty much everything.

Keyless Car Transponders



- There are two security mechanisms: one remote central locking system and one passive RFID tag (engine immobiliser).
- How can I get in? How can thieves be kept out? How to avoid MITM attacks?

Papers: Gone in 360 Seconds: Hijacking with Hitag2,
Dismantling Megamos Crypto: Wirelessly Lockpicking
a Vehicle Immobilizer

Problems with Key Fobs

Circumventing the ignition protection:

- either dismantling Megamos crypto,
- or use the diagnostic port to program blank keys

MONDAY 27 OCTOBER 2014 EVENING STANDARD

ebook.com/eveningstandard
us on Twitter @standardnews



Insurers refuse to cover Range Rovers due to security flaw

Kiran Randhawa

INSURANCE companies are refusing to cover new Range Rovers in London after thieves found a way of bypassing the vehicles' keyless ignition systems.

Criminals use hand-held electronic devices, available on eBay, to get around the feature. Unless owners have secure parking, underwriters are now said to be refusing to insure them.

Insurers have asked to meet Jaguar Range Rover to discuss the growing problem. Thatcham Research, the motor insurers' automotive research centre, said that 294 Range Rover Evoque and Sport vehicles were stolen in London between January and July. In the same period, 63 BMW X5s, a rival to the Range Rover, were taken.

James Wasdell, co-founder of Quantum Underwriting, said: "If you are an owner of a street-parked Range Rover, nine out of 10 insurers will now say no. However, we have found a solution by combining the use of physical protection [for the car] and advising clients to insure all assets with one insurer." Jaguar Land Rover said: "Our line-up continues to meet the insurance industry requirements. Nevertheless, we are taking this issue very seriously."

Dismantling Megamos Crypto: Wirelessly Lockpicking a Vehicle Immobilizer

Paul Yorkel¹, Florin D. Garcia², and Rory Ege³

¹ Institute for Computing and Information Science, York University, Toronto, ON, Canada

² School of Computer Science, University of Birmingham, United Kingdom

³ e-mail: rory.ege@utoronto.ca

1. Introduction

This is a preprint submitted to the High Court of London on Tuesday 26th June 2013. The authors are restricted from publishing the technical content of the article until the proceedings of the 13th USENIX Security Symposium, Washington DC, August 2013.

2. Historical clues

Figure 1 contains the cryptographic hash (SHA-256) of the original final paper which was scheduled to appear in the proceedings of the 13th USENIX Security Symposium, Washington DC, August 2013.

```
7d050a897404979ec0ca3d8609174b444  
43643da139c70b79366954c0c065d8e  
64031881348df0c23ba6f62a6a80c04df  
bbb629a1d8ffc60fa31880b4d5b4aca
```

Figure 1: SHA-256 hash of the final paper

References

1. Paul Yorkel, Florin D. Garcia, and Rory Ege. Dismantling megamos crypto: wirelessly lockpicking a vehicle immobilizer. In *13th USENIX Security Symposium (USENIX Security 2013)*. USENIX Association, 2013.

HTTPS / GSM



- I am sitting at Starbucks. How can I be sure I am really visiting Barclays? I have no control of the access point.
- How can I achieve that a secret key is established in order to encrypt my mobile conversation? I have no control over the access points.


G20 Summit in 2009



- Snowden documents reveal “that during the G20 meetings...GCHQ used ‘ground-breaking intelligence capabilities’ to intercept the communications of visiting delegations. This included setting up internet cafes where they used an email interception program and key-logging software to spy on delegates’ use of computers...”
- “The G20 spying appears to have been organised for the more mundane purpose of securing an advantage in meetings.”

Handshakes

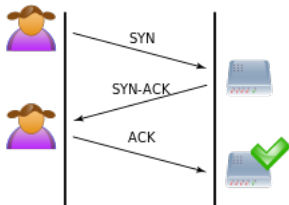
- starting a TCP connection between a client and a server initiates the following three-way handshake protocol:



Alice: Hello server!
Server: I heard you
Alice: Thanks

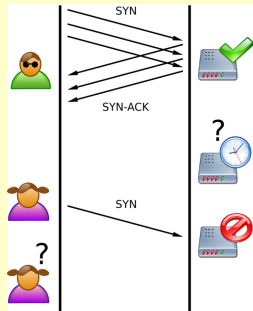
Handshakes

- starting a TCP connection between a client and a server initiates the following three-way handshake protocol:



Alice:
Server:
Alice:

SYNflood attacks:



Protocols

$A \rightarrow B : \dots$

- by convention A , B are named principals *Alice...* but most likely they are programs, which just follow some instructions (they are more like roles)


Protocols

$$\begin{array}{l} A \rightarrow B : \dots \\ B \rightarrow A : \dots \\ \vdots \end{array}$$

- by convention A , B are named principals *Alice...* but most likely they are programs, which just follow some instructions (they are more like roles)
- indicates one “protocol run”, or session, which specifies some order in the communication
- there can be several sessions in parallel (think of wifi routers)

Handshakes

- starting a TCP connection between a client and a server initiates the following three-way handshake protocol:



Alice: Hello server!
Server: I heard you
Alice: Thanks

$A \rightarrow S: \text{ SYN}$

$S \rightarrow A: \text{ SYN-ACK}$

$A \rightarrow S: \text{ ACK}$

Cryptographic Protocol Failures

Ross Anderson and Roger Needham wrote:

A lot of the recorded frauds were the result of this kind of blunder, or from management negligence pure and simple. However, there have been a significant number of cases where the designers protected the right things, used cryptographic algorithms which were not broken, and yet found that their systems were still successfully attacked.

Oyster Cards



- good example of a bad protocol (security by obscurity)

Wirelessly Pickpocketing a Mifare Classic Card

The Mifare Classic is the most widely used contactless smartcard on the market. The stream cipher CRYPTO1 used by the Classic has recently been reverse engineered and serious attacks have been proposed. The most serious of them retrieves a secret key in under a second. In order to clone a card, previously proposed attacks require that the adversary either has access to an eavesdropped communication session or executes a message-by-message man-in-the-middle attack between the victim and a legitimate reader. Although this is already disastrous from a cryptographic point of view, system integrators maintain that these attacks cannot be performed undetected.

This paper proposes four attacks that can be executed by an adversary having only wireless access to just a card (and not to a legitimate reader). The most serious of them recovers a secret key in less than a second on ordinary hardware. Besides the cryptographic weaknesses, we exploit other weaknesses in the protocol stack. A vulnerability in the computation of parity bits allows an adversary to establish a side channel. Another vulnerability regarding nested authentications provides enough plaintext for a speedy known-plaintext attack. (a paper from 2009)

Oyster Cards



- good example of a bad protocol (security by obscurity)
- *“Breaching security on Oyster cards should not allow unauthorised use for more than a day, as TfL promises to turn off any cloned cards within 24 hours...”*

Authentication Protocols

Alice (A) and Bob (B) share a secret key K_{AB}

Passwords:

$$B \rightarrow A : K_{AB}$$

Authentication Protocols

Alice (A) and Bob (B) share a secret key K_{AB}

Passwords:

$$B \rightarrow A : K_{AB}$$

Problem: Eavesdropper can capture the secret and replay it; A cannot confirm the identity of B

Authentication?



"On the Internet, nobody knows you're a dog."

Authentication Protocols

Alice (A) and Bob (B) share a secret key K_{AB}

Simple Challenge Response:

$$\begin{aligned} A &\rightarrow B : N \\ B &\rightarrow A : \{N\}_{K_{AB}} \end{aligned}$$

Authentication Protocols

Alice (A) and Bob (B) share a secret key K_{AB}

Mutual Challenge Response:

$$A \rightarrow B : N_A$$

$$B \rightarrow A : \{N_A, N_B\}_{K_{AB}}$$

$$A \rightarrow B : N_B$$

Nonces

- 1 I generate a nonce (random number) and send it to you encrypted with a key we share
- 2 you increase it by one, encrypt it under a key I know and send it back to me

I can infer:

- you must have received my message
- you could only have generated your answer after I send you my initial message
- if only you and me know the key, the message must have come from you

$A \rightarrow B: N_A$
 $B \rightarrow A: \{N_A, N_B\}_{K_{AB}}$
 $A \rightarrow B: N_B$

The attack (let A decrypt her own messages):

$A \rightarrow E: N_A$
 $E \rightarrow A: N_A$
 $A \rightarrow E: \{N_A, N'_A\}_{K_{AB}}$
 $E \rightarrow A: \{N_A, N'_A\}_{K_{AB}}$
 $A \rightarrow E: N'_A (= N_B)$

$A \rightarrow B: N_A$
 $B \rightarrow A: \{N_A, N_B\}_{K_{AB}}$
 $A \rightarrow B: N_B$

The attack (let A decrypt her own messages):

$A \rightarrow E: N_A$
 $E \rightarrow A: N_A$
 $A \rightarrow E: \{N_A, N'_A\}_{K_{AB}}$
 $E \rightarrow A: \{N_A, N'_A\}_{K_{AB}}$
 $A \rightarrow E: N'_A (= N_B)$

Solutions: $K_{AB} \neq K_{BA}$ or include an id in the second message

Encryption to the Rescue?

- $A \rightarrow B : \{A, N_A\}_{K_{AB}}$ encrypted
- $B \rightarrow A : \{N_A, K'_{AB}\}_{K_{AB}}$
- $A \rightarrow B : \{N_A\}_{K'_{AB}}$

Encryption to the Rescue?

- $A \rightarrow B : \{A, N_A\}_{K_{AB}}$ encrypted
- $B \rightarrow A : \{N_A, K'_{AB}\}_{K_{AB}}$
- $A \rightarrow B : \{N_A\}_{K'_{AB}}$

means you need to send separate “Hello” signals (bad), or worse share a single key between many entities

Protocol Attacks

- replay attacks
- reflection attacks
- man-in-the-middle attacks
- timing attacks
- parallel session attacks
- binding attacks (public key protocols)
- changing environment / changing assumptions

- (social engineering attacks)

Public-Key Infrastructure

- the idea is to have a certificate authority (CA)
- you go to the CA to identify yourself
- CA: “I, the CA, have verified that public key P_{Bob}^{pub} belongs to Bob”
- CA must be trusted by everybody
- What happens if CA issues a false certificate?
Who pays in case of loss? (VeriSign explicitly limits liability to \$100.)

Man-in-the-Middle

“Normal” protocol run:

- A sends public key to B
- B sends public key to A
- A sends message encrypted with B 's public key, B decrypts it with its private key
- B sends message encrypted with A 's public key, A decrypts it with its private key

Man-in-the-Middle

Attack:

- A sends public key to B — C intercepts this message and send his own public key
- B sends public key to A — C intercepts this message and send his own public key
- A sends message encrypted with C 's public key, C decrypts it with its private key, re-encrypts with B 's public key
- similar for other direction

Man-in-the-Middle

Potential Prevention?

- A sends public key to B
- B sends public key to A
- A encrypts message with B 's public key, send's **half** of the message
- B encrypts message with A 's public key, send's **half** of the message
- A sends other half, B can now decrypt entire message
- B sends other half, A can now decrypt entire message

Man-in-the-Middle

Potential Prevention?

- A sends public key to B
- B sends public key to A
- A encrypts message with B 's public key, send's **half** of the message
- B encrypts message with A 's public key, send's **half** of the message
- A sends other half, B can now decrypt entire message
- B sends other half, A can now decrypt entire message

Under which circumstances does this protocol prevent MiM-attacks, or does it?

Splitting Messages

0 X 1 p e U V T G J K + H 7 0 m M j A M 8 p

$\{A, m\}_{K_B^{pub}}$

0 X 1 p e U V T G J K

H_1

+ H 7 0 m M j A M 8 p

H_2

- you can also use the even and odd bytes
- the point is you cannot decrypt the halves, even if you have the key

$$A \rightarrow C : K_A^{pub}$$

$$C \rightarrow B : K_C^{pub}$$

$$B \rightarrow C : K_B^{pub}$$

$$C \rightarrow A : K_C^{pub}$$

$$\{A, m\}_{K_C^{pub}} \mapsto H_1, H_2$$

$$\{B, m'\}_{K_C^{pub}} \mapsto M_1, M_2$$

$$\{C, a\}_{K_B^{pub}} \mapsto C_1, C_2$$

$$\{C, b\}_{K_A^{pub}} \mapsto D_1, D_2$$

$$A \rightarrow C : H_1$$

$$C \rightarrow B : C_1$$

$$B \rightarrow C : \{C_1, M_1\}_{K_C^{pub}}$$

$$C \rightarrow A : \{H_1, D_1\}_{K_A^{pub}}$$

$$A \rightarrow C : \{H_2, D_1\}_{K_C^{pub}}$$

$$C \rightarrow B : \{C_2, M_1\}_{K_B^{pub}}$$

$$B \rightarrow C : M_2$$

$$C \rightarrow A : D_2$$

$$A \rightarrow C : K_A^{pub}$$

$$C \rightarrow B : K_C^{pub}$$

$$B \rightarrow C : K_B^{pub}$$

$$C \rightarrow A : K_C^{pub}$$

$$\{A, m\}_{K_C^{pub}} \mapsto H_1, H_2$$

$$\{B, m'\}_{K_C^{pub}} \mapsto M_1, M_2$$

$$\{C, a\}_{K_B^{pub}} \mapsto C_1, C_2$$

$$\{C, b\}_{K_A^{pub}} \mapsto D_1, D_2$$

$$A \rightarrow C : H_1$$

$$C \rightarrow B : C_1$$

$$B \rightarrow C : \{C_1, M_1\}_{K_C^{pub}}$$

$$C \rightarrow A : \{H_1, D_1\}_{K_A^{pub}}$$

$$A \rightarrow C : \{H_2, D_1\}_{K_C^{pub}}$$

$$C \rightarrow B : \{C_2, M_1\}_{K_B^{pub}}$$

$$B \rightarrow C : M_2$$

$$C \rightarrow A : D_2$$

m = How is your grandmother? m' = How is the weather today in London?

- you have to ask something that cannot be imitated (requires A and B know each other)
- what happens if m and m' are voice messages?

- you have to ask something that cannot be imitated (requires A and B know each other)
- what happens if m and m' are voice messages?
- So C can either leave the communication unchanged, or invent a complete new conversation

Car Transponder (HiTag2)

- 1 C generates a random number N
- 2 C calculates $(F, G) = \{N\}_K$
- 3 $C \rightarrow T: N, F$
- 4 T calculates $(F', G') = \{N\}_K$
- 5 T checks that $F = F'$
- 6 $T \rightarrow C: N, G'$
- 7 C checks that $G = G'$

Car Transponder (HiTag2)


- 1 C generates a random number N
- 2 C calculates $(F, G) = \{N\}_K$
- 3 $C \rightarrow T: N, F$
- 4 T calculates $(F', G') = \{N\}_K$
- 5 T checks that $F = F'$
- 6 $T \rightarrow C: N, G'$
- 7 C checks that $G = G'$

This process means that the transponder believes the car knows the key K , and the car believes the transponder knows the key K . They have authenticated themselves to each other, or have they?


A Man-in-the-middle attack in real life:

- the card only says yes to the terminal if the PIN is correct
- trick the card in thinking transaction is verified by signature
- trick the terminal in thinking the transaction was verified by PIN

Normal PIN check



Fraudulent PIN check



- the moral: establishing a secure connection from “zero” is almost impossible—you need to rely on some established trust
- that is why PKI relies on certificates, which however are badly, badly realised

Trusted Third Parties

Simple protocol for establishing a secure connection via a mutually trusted 3rd party (server):

$$A \rightarrow S : A, B$$

$$S \rightarrow A : \{K_{AB}, \{K_{AB}\}_{K_{BS}}\}_{K_{AS}}$$


$$A \rightarrow B : \{K_{AB}\}_{K_{BS}}$$

$$A \rightarrow B : \{m\}_{K_{AB}}$$

PKI: The Main Idea

- the idea is to have a certificate authority (CA)
- you go to the CA to identify yourself
- CA: “I, the CA, have verified that public key P_{Bob}^{pub} belongs to Bob”
- CA must be trusted by everybody
- certificates are time limited, and can be revoked
- What happens if CA issues a false certificate?
Who pays in case of loss? (VeriSign explicitly limits liability to \$100.)

PKI: Chains of Trust



- CAs make almost no money anymore, because of stiff competition
- browser companies are not really interested in security; only in market share

PKI: Weaknesses

CAs just cannot win (make any profit):

- there are hundreds of CAs, which issue millions of certificates and the error rate is small
- users (servers) do not want to pay or pay as little as possible
- a CA can issue a certificate for any domain not needing any permission (CAs are meant to undergo audits, but...DigiNotar)
- if a CA has issued many certificates, it “becomes too big to fail”
- Can we be sure CAs are not just frontends of some government organisation?

PKI: Weaknesses

- many certificates are issued via Whois, whether you own the domain...if you hijacked a domain, it is easy to obtain certificates
- the revocation mechanism does not work (Chrome has given up on general revocation lists)
- lax approach to validation of certificates (Have you ever bypassed certification warnings?)
- sometimes you want to actually install invalid certificates (self-signed)

PKI: Attacks

- Go directly after root certificates
 - governments can demand private keys
 - 10 years ago it was estimated that breaking a 1024 bit key takes one year and costs 10 - 30 Mio \$; this is now reduced to 1 Mio \$
- Go after buggy implementations of certificate validation
- Social Engineering
 - in 2001 somebody pretended to be from Microsoft and asked for two code-signing certificates

The eco-system is completely broken (it relies on thousands of entities to do the right thing). Maybe DNSSEC where keys can be attached to domain names is a way out.

Real Attacks

- In 2011, DigiNotar (Dutch company) was the first CA that got compromised comprehensively, and where many fraudulent certificates were issued to the wild. It included approximately 300,000 IP addresses, mostly located in Iran. The attackers (in Iran?) were likely interested “only” in collecting gmail passwords.
- The Flame malware piggy-bagged on this attack by advertising malicious Windows updates to some targeted systems (mostly in Iran, Israel, Sudan).

PKI is Broken

- PKI and certificates are meant to protect you against MITM attacks, but if the attack occurs you are presented with a warning and you need to decide whether you are under attack.
- Webcontent gets often loaded from 3rd-party servers, which might not be secured
- Misaligned incentives: browser vendors are not interested in breaking webpages with invalid certificates

Why are there so many invalid certificates?

- insufficient name coverage (www.example.com should include example.com)
- IoT: many appliances have web-based admin interfaces; the manufacturer cannot know under which IP and domain name the appliances are run (so cannot install a valid certificate)
- expired certificates, or incomplete chains of trust (servers are supposed to supply them)

Protocols are Difficult

- even the systems designed by experts regularly fail
- the one who can fix a system should also be liable for the losses
- cryptography is often not the problem