

CSE 451: Operating Systems Spring 2006

Module 20 Distributed Systems Authentication / Authorization

John Zahorjan
zahorjan@cs.washington.edu
Allen Center 534

Some Fundamental Concepts

- **Authentication** (who are you)
 - identifying principals (users / programs)
- **Authorization** (what are you allowed to do)
 - determining what access users and programs have to things
- **Auditing** (what happened)
 - record what users and programs are doing for later analysis / prosecution

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

2

The Local Case: Unix Password File

- Encrypt passwords with passwords

$K=[\text{alison}]_{\text{alison}}$

Bob: 14: S6Uu0cYDVdTAK
David: 15: J2ZI4ndBL6X.M
Mary: 16: VW2bqvTalBJKg

- David's password, "alison," is encrypted using itself as the key and stored in that form.
- Password supplied by user is encrypted with itself as key, and result compared to stored result.
- "No problem if someone steals the file"

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

3

The Dictionary Attack

- Encrypt many (all) possible password strings offline, and store results in a dictionary
 - I may not be able to invert any particular password, but the odds are very high I can invert one or more
- 26 letters used, 7 letters long
 - 8 billion passwords (33 bits)
 - Generating 100,000/second requires 22 hours
- But most people's passwords are not random sequences of letters!
 - girlfriend's/boyfriend's/spouse's/dog's name/words in the dictionary
- Dictionary attacks have traditionally been incredibly easy

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

4

Making it harder

- Using symbols and numbers and longer passwords
 - 95 characters, 14 characters long
 - 10^{27} passwords = 91 bits
 - Checking 100,000/second breaks in 10^{14} years
- Require frequent changing of passwords
 - guards against loaning it out, writing it down, etc.

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

5

Do longer passwords work?

- People can't remember 14-character strings of random characters
- People write down difficult passwords
- People give out passwords to strangers
- Passwords can show up on disk
- If you are forced to change your password periodically, you probably choose an even dumber one
 - "feb04" "mar04" "apr04"
- How do we handle this in CSE?

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

6

Countermeasure to the Dictionary Attack: Salt

- Unix (1979): **salted** passwords
 - The salt is just a random number from a large space

$K=[alison392]_{alison392}$

Bob: 14: T7Vs1dZEWeRcL: 45
David: 15: K3AJ5ocCM4ZMS: 392
Mary: 16: WX3crwUbmCKLf: 152

Encryption is computed after affixing a number to the password. Thwarts pre-computed dictionary attacks

Okay, are we done? Problem solved?

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

7

Attack Models

- Besides the problems already mentioned that obviously remain (people give out their passwords / write them down / key loggers / ...), there may be other clever attacks that we haven't thought of
- **Attack Model:** when reasoning about the security of a mechanism, we need typically need to carefully describe what kinds of attacks we're thinking of
 - helps us reason about what vulnerabilities still remain

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

8

Example 1: Login spoofers

- Login spoofers are a specialized class of Trojan horses
 - Attacker runs a program that presents a screen identical to the login screen and walks away from the machine
 - Victim types password and gets a message saying "password incorrect, try again"
- Can be circumvented by requiring an operation that unprivileged programs cannot perform
 - E.g., start login sequence with a key combination user programs cannot catch, CTRL+ALT+DEL on Windows

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

9

Example 2: Cool password attack

- VMS (early 80's) password checking flaw
 - password checking algorithm:



```
for (I=0; I<password.length( ); I++) {
    if password[I] == supplied_password[I]
        return false;
}
return true;
```
 - can you see the problem?
 - hint: think about virtual memory...
 - another hint: think about page faults...
 - final hint: who controls where in memory supplied_password lives?

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

10

Distributed Authentication (Single Domain)




5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

11

Aside on Encryption

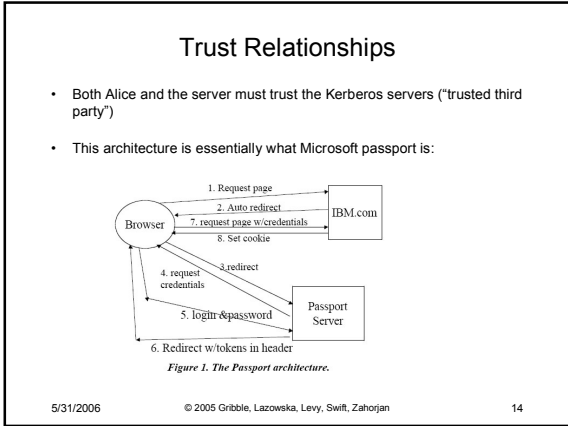
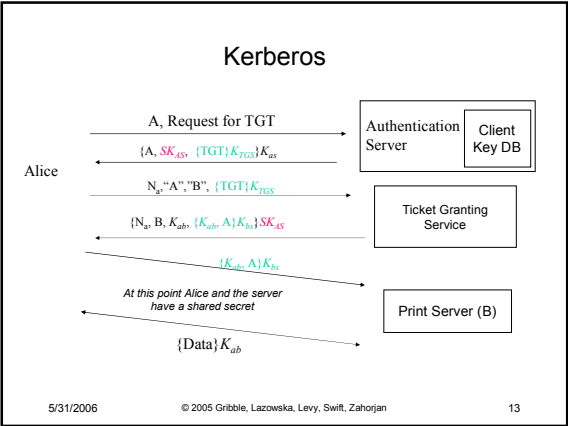


- **Encryption:** takes a **key** and **plaintext** and creates **ciphertext**: $\{M\}_{k1} = C$
- **Decryption:** takes ciphertext and a key and recovers plaintext: $\{C\}_{k2} = M$
- **Symmetric algorithms** (aka secret-key aka shared secret algorithms):
 - $k1 = k2$ (or can get $k2$ from $k1$)
- **Public-Key Algorithms**
 - decryption key ($k2$) cannot be calculated from encryption key ($k1$)
 - encryption key can be made public!
 - encryption key = "public key", decryption key = "private key"
- **Computational requirements:**
 - Deducing M from $\{M\}_k$ is "really hard"
 - Computing $\{M\}_{k1}$ and $\{C\}_{k2}$ is efficient

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

12



CBS NEWS

Microsoft's Passport Flaw Fixed
WASHINGTON, May 8, 2003

(AP) A computer researcher in Pakistan discovered how to breach Microsoft Corp.'s security procedures for its popular Internet Passport service, designed to protect customers visiting some retail Web sites, sending e-mails and in some cases making credit-card purchases.

Microsoft acknowledged the flaw affected all its 200 million Passport accounts but said it fixed the problem early Thursday, after details were published on the Internet. Product Manager Adam Sohn said the company was unaware of hackers actually hijacking anyone's Passport account, but several experts said they successfully tested the procedure overnight.

In theory, Microsoft could face a staggering fine by U.S. regulators of up to \$2.2 billion. Under a settlement with the Federal Trade Commission last year over lapsed Passport security, Microsoft pledged to take reasonable safeguards to protect personal consumer information during the next two decades or risk fines up to \$11,000 per violation.

The FTC said it was investigating this latest lapse. The agency's assistant director for financial practices, Jessica Rich, said Thursday that each vulnerable account could constitute a separate violation - raising the maximum fine that could be assessed against Microsoft to \$2.2 billion.

"If we were to find that they didn't take reasonable safeguards to protect the information, that could be an order violation," Rich said.

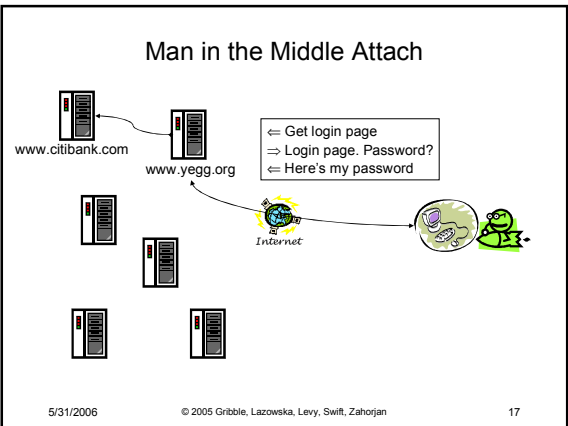
The researcher, Muhammad Faisal Rauf Danka, determined that by typing a specific Web address that included the phrase "email@wrednet", he could seize any person's Passport account and change the password associated with it.

5/31/2006 © 2005 Gribble, Lazowska, Levy, Swift, Zahorjan 15

Distributed Authentication at World Scale

- Bill Gates wants to login to his Citibank account to move \$10 from savings to checking
- Both Bill and Citibank are worried:
 - Citibank:
 - How do I know that I'm talking with Bill?
 - Does Bill have \$10 in his savings account?
 - ...
 - Bill:
 - How do I know that I'm talking with Citibank?

5/31/2006 © 2005 Gribble, Lazowska, Levy, Swift, Zahorjan 16

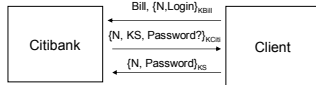


Authentication Solutions

- Citibank authenticating Bill
 - This is just a client accessing a server. Citibank can use shared secrets.
 - Bill has to use some secret communicated out-of-band (e.g., ATM PIN number) to create a shared secret for online access.
- Bill authenticating Citibank
 - Could shared secret work for the bank to authenticate itself to the client?
 - ...
 - In the end, we rely on a trusted third party (just like Kerberos, but implemented differently)

5/31/2006 © 2005 Gribble, Lazowska, Levy, Swift, Zahorjan 18

Why not this?



5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

19

Public Key Encryption

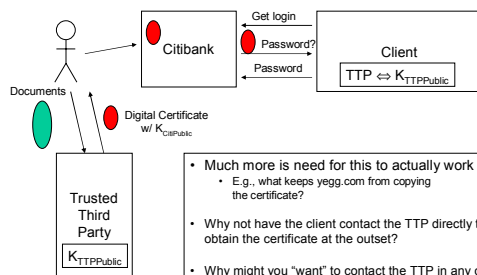
- Key pairs, K_{Public} / $K_{Private}$
 - $\{(M)_{K_{Public}}\}_{K_{Private}} = \{(M)_{K_{Private}}\}_{K_{Public}} = M$
 - Each key is the decryption key for the other used as an encryption key
 - It is computationally infeasible to deduce $K_{Private}$ from K_{Public}
 - You can distribute K_{Public} freely
- $\{(M)_{K_{Public}}\}$ can be decrypted only by the holder of the private key
- $\{(M)_{K_{Private}}\}$ can be created only by the holder of the private key
 - "Signing"

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

20

Authentication by Certificate: Basic Idea




- Much more is need for this to actually work
 - E.g., what keeps yegg.com from copying the certificate?
- Why not have the client contact the TTP directly to obtain the certificate at the outset?
- Why might you "want" to contact the TTP in any case?

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

21

Client/Server Communication: ssl (tls)



Notes:

1. Master/session key determined independently by both client and server as:
 $F(N_{Client}, N_{Server}, \text{Pre-master})$
2. I've taken some liberties to simplify the explanation... (cf. CSE 461)

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

22

The Larger Security Problem

- **Integrity**
My data should be protected against modification by malicious parties
 - "Modification" includes deletion
- **Privacy**
My data should not be disclosed without my consent
- Both issues have become much more complicated in the last decade
 - Attackers exploit bugs/weaknesses accessible through the net
 - We all run third-party code

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

23

Spyware

- Software that is installed that collects information and reports it to third party
 - key logger, adware, browser hijacker, ...
- Installed one of two ways
 - piggybacked on software you choose to download
 - "drive-by" download
 - your web browser has vulnerabilities
 - web server can exploit by sending you bad web content
- Estimates
 - majority (50-90%) of Internet-connected PCs have it
 - 1 in 20 executables on the Web have it
 - about 0.5% of Web pages attack you with drive-by-downloads

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

24

kingsofchaos.com

- A benign web site for an online game
 - earns revenue from ad networks by showing banners
 - but, it relinquishes control of the ad content

banner ad from
adworldnetwork.com
(a legitimate ad network)

inline javascript loads
HTML from ad provider



5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

26

Incident

- kingsofchaos.com was given this "ad content"

```
<script type="text/javascript">document.write('
\u003c\u0062\u006f\u0064\u0079\u0020\u006f\u006e\u0055\u006f\u0077\u0050\u006f\u0070\u0075\u0028\u0029\u003b\u0073\u0068\u006f\u006f\u0077\u0048\u0069 ...etc.
```

- This "ad" ultimately:
 - bombarded the user with pop-up ads
 - hijacked the user's homepage
 - exploited an IE vulnerability to install spyware

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

26

What's going on?

- The advertiser was an ex-email-spammer
- His goal:
 - **force** users to see ads from his servers
 - **draw revenue** from ad "affiliate programs"
 - Apparently earned several millions of dollars
- Why did he use spyware?
 - control PC and show ads even when not on the Web

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

27

Principle of Least Privilege

- Figure out exactly which capabilities a program needs to run, and grant it only those
 - start out by granting none
 - run program, and see where it breaks
 - add new privileges as needed.
- Unix: concept of root is **not** a good example of this
 - some programs need root just to get a small privilege
 - e.g., FTP daemon requires root:
 - to listen on network port < 1024
 - to change between user identities after authentication
 - but root also lets you read any file in filesystem

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

28

Principle of Complete Mediation

- Check **every** access to **every** object
 - in rare cases, can get away with less (caching)
 - but only if sure nothing relevant in environment has changed...and there is a lot that's relevant!
- A TLB caches access control information
 - page table entry protection bits
 - is this a violation of the principle?

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

29

Modern security problems

- Confinement
 - How do I run code that I don't trust?
 - e.g., RealPlayer, Flash
 - How do I restrict the data it can communicate?
 - What if trusted code has bugs?
 - e.g., Internet Explorer
- Solutions
 - Restricted contexts – let the user divide their identity
 - ActiveX – make code writer identify self
 - Java – use a virtual machine that intercepts all calls
 - Binary rewriting – modify the program to force it to be safe

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

30

Restricted contexts

- Role-based access control (RBAC)
 - Add extra identity information to a process
 - e.g., both username and program name (mikesw.navigator)
 - Use both identities for access checks
 - add extra security checks at system calls that use program name
 - add extra ACLs on objects that grant/deny access to the program
 - Allows users to **sub-class** themselves for less-trusted programs
- `chroot`
- Browse in a VMWare machine


5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

31

ActiveX

- All code comes with a public-key signature
- Code indicates what privileges it needs
- Web browser verifies certificate
- Once verified, code is completely trusted



5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

32

Java / C#

- All problems are solved by a layer of indirection
 - All code runs on a virtual machine
 - Virtual machine tracks security permissions
 - Allows fancier access control models - allows stack walking
- Interposition using language VM doesn't work for other languages
- Virtual machines can be used with all languages
 - Run virtual machine for hardware
 - Inspect stack to determine *subject* for access checks

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

33

Binary rewriting

- Goal: enforce code safety by *embedding* checks in the code
- Solution:
 - Compute a mask of accessible addresses
 - Replace system calls with calls to special code

Original Code:

```
lw $a0, 14($s4)
jal ($s5)
move $a0, $v0
jal $printf
```

Rewritten Code:

```
and $t6,$s4,0x001fff0
lw $a0, 14($t6)
and $t6,$s5, 0x001fff0
jal ($t6)
move $a0, $v0
jal $sfi_printf
```

5/31/2006

© 2005 Gribble, Lazowska, Levy, Swift, Zahorjan

34