

CS162  
Operating Systems and  
Systems Programming  
Lecture 25

Protection and Security  
in Distributed Systems

May 5, 2008  
Prof. Anthony D. Joseph  
<http://inst.eecs.berkeley.edu/~cs162>

Goals for Today

- Security Mechanisms
  - Authentication
  - Authorization
  - Enforcement
- Cryptographic Mechanisms

Note: Some slides and/or pictures in the following are adapted from slides ©2005 Silberschatz, Galvin, and Gagne. Many slides generated from my lecture notes by Kubiawicz.

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.2

Protection vs Security

- **Protection:** one or more mechanisms for controlling the access of programs, processes, or users to resources
  - Page Table Mechanism
  - File Access Mechanism
- **Security:** use of protection mechanisms to prevent misuse of resources
  - Misuse defined with respect to policy
    - » E.g.: prevent exposure of certain sensitive information
    - » E.g.: prevent unauthorized modification/deletion of data
  - Requires consideration of the external environment within which the system operates
    - » Most well-constructed system cannot protect information if user accidentally reveals password
- What we hope to gain today and next time
  - Conceptual understanding of how to make systems secure
  - Some examples, to illustrate why providing security is really hard in practice

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.3

Preventing Misuse

- Types of Misuse:
  - Accidental:
    - » If I delete shell, can't log in to fix it!
    - » Could make it more difficult by asking: "do you really want to delete the shell?"
  - Intentional:
    - » Some high school brat who can't get a date, so instead he transfers \$3 billion from B to A.
    - » Doesn't help to ask if they want to do it (of course!)
- Three Pieces to Security
  - **Authentication:** who the user actually is
  - **Authorization:** who is allowed to do what
  - **Enforcement:** make sure people do only what they are supposed to do
- Loopholes in any carefully constructed system:
  - Log in as superuser and you've circumvented authentication
  - Log in as self and can do anything with your resources; for instance: run program that erases all of your files
  - Can you trust software to correctly enforce Authentication and Authorization?????

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.4

## Authentication: Identifying Users

### • How to identify users to the system?

#### - Passwords

- » Shared secret between two parties
- » Since only user knows password, someone types correct password  $\Rightarrow$  must be user typing it
- » Very common technique

#### - Smart Cards

- » Electronics embedded in card capable of providing long passwords or satisfying challenge  $\rightarrow$  response queries
- » May have display to allow reading of password
- » Or can be plugged in directly; several credit cards now in this category

#### - Biometrics

- » Use of one or more intrinsic physical or behavioral traits to identify someone
- » Examples: fingerprint reader, palm reader, retinal scan
- » Becoming quite a bit more common



5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.5

## Passwords: Secrecy

### • System must keep copy of secret to check against passwords

- What if malicious user gains access to list of passwords?

- » Need to obscure information somehow

- Mechanism: utilize a transformation that is difficult to reverse without the right key (e.g. encryption)

### • Example: UNIX /etc/passwd file

- passwd  $\rightarrow$  one way transform(hash)  $\rightarrow$  encrypted passwd

- System stores only encrypted version, so OK even if someone reads the file!

- When you type in your password, system compares encrypted version

### • Problem: Can you trust encryption algorithm?

- Example: one algorithm thought safe had back door

- » Governments want back door so they can snoop

- Also, security through obscurity doesn't work

- » GSM encryption algorithm was secret; accidentally released; Berkeley grad students cracked in a few hours



5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.6

## Passwords: How easy to guess?

### • Ways of Compromising Passwords

#### - Password Guessing:

- » Often people use obvious information like birthday, favorite color, girlfriend's name, etc...

#### - Dictionary Attack:

- » Work way through dictionary and compare encrypted version of dictionary words with entries in /etc/passwd

#### - Dumpster Diving:

- » Find pieces of paper with passwords written on them
- » (Also used to get social-security numbers, etc)

### • Paradox:

- Short passwords are easy to crack

- Long ones, people write down!

### • Technology means we have to use longer passwords

- UNIX initially required lowercase, 5-letter passwords: total of  $26^5=10$ million passwords

- » In 1975, 10ms to check a password  $\rightarrow$  1 day to crack

- » In 2005, .01 $\mu$ s to check a password  $\rightarrow$  0.1 seconds to crack

- Takes less time to check for all words in the dictionary!

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.7

## Passwords: Making harder to crack

### • How can we make passwords harder to crack?

- Can't make it impossible, but can help

### • Technique 1: Extend everyone's password with a unique number (stored in password file)

- Called "salt". UNIX uses 12-bit "salt", making dictionary attacks 4096 times harder

- Without salt, would be possible to pre-compute all the words in the dictionary hashed with the UNIX algorithm: would make comparing with /etc/passwd easy!

- Also, way that salt is combined with password designed to frustrate use of off-the-shelf DES hardware

### • Technique 2: Require more complex passwords

- Make people use at least 8-character passwords with upper-case, lower-case, and numbers

- »  $70^8=6 \times 10^{14}=6$ million seconds=69 days@0.01 $\mu$ s/check

- Unfortunately, people still pick common patterns

- » e.g. Capitalize first letter of common word, add one digit

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.8

### Passwords: Making harder to crack (con't)

- **Technique 3: Delay checking of passwords**
  - If attacker doesn't have access to /etc/passwd, delay every remote login attempt by 1 second
  - Makes it infeasible for rapid-fire dictionary attack
- **Technique 4: Assign very long passwords**
  - Long passwords or pass-phrases can have more entropy (randomness→harder to crack)
  - Give everyone a smart card (or ATM card) to carry around to remember password
    - » Requires physical theft to steal password
    - » Can require PIN from user before authenticates self
  - Better: have smartcard generate pseudorandom number
    - » Client and server share initial seed
    - » Each second/login attempt advances to next random number
- **Technique 5: "Zero-Knowledge Proof"**
  - Require a series of challenge-response questions
    - » Distribute secret algorithm to user
    - » Server presents a number, say "5"; user computes something from the number and returns answer to server
    - » Server never asks same "question" twice
  - Often performed by smartcard plugged into system

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.9

### Administrivia

- Project #4 code deadline is Wed 5/14 at 11:59pm
- Final Exam
  - May 21<sup>st</sup>, 12:30-3:30pm
- Final Topics: Any suggestions?
  - Please send them to me...

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.10

### Authentication in Distributed Systems

- What if identity must be established across network?



- Need way to prevent exposure of information while still proving identity to remote system
- Many of the original UNIX tools sent passwords over the wire "in clear text"
  - » E.g.: telnet, ftp, yp (yellow pages, for distributed login)
  - » Result: Snooping programs widespread
- What do we need? Cannot rely on physical security!
  - Encryption: Privacy, restrict receivers
  - Authentication: Remote Authenticity, restrict senders

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.11

### Private Key Cryptography

- **Private Key (Symmetric) Encryption:**
    - Single key used for both encryption and decryption
  - **Plaintext:** Unencrypted Version of message
  - **Ciphertext:** Encrypted Version of message
- 
- Important properties
    - Can't derive plain text from ciphertext (decode) without access to key
    - Can't derive key from plain text and ciphertext
    - As long as password stays secret, get both secrecy and authentication
  - Symmetric Key Algorithms: DES, Triple-DES, AES

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.12

## Key Distribution


- How do you get shared secret to both places?
  - For instance: how do you send authenticated, secret mail to someone who you have never met?
  - Must negotiate key over private channel
    - » Exchange code book
    - » Key cards/memory stick/others
- Third Party: Authentication Server (like Kerberos)
  - Notation:
    - »  $K_{xy}$  is key for talking between x and y
    - »  $(...)^K$  means encrypt message (...) with the key K
    - » Clients: A and B, Authentication server S
  - A asks server for key:
    - » A→S: [Hi! I'd like a key for talking between A and B]
    - » Not encrypted. Others can find out if A and B are talking
  - Server returns *session* key encrypted using B's key
    - » S→A: **Message** [ Use  $K_{ab}$  (This is A! Use  $K_{ab}$ )<sup>Ksb</sup> ]<sup>Ksa</sup>
    - » This allows A to know, "S said use this key"
  - Whenever A wants to talk with B
    - » A→B: **Ticket** [ This is A! Use  $K_{ab}$  ]<sup>Ksb</sup>
    - » Now, B knows that  $K_{ab}$  is sanctioned by S

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.13

## Authentication Server Continued



- Details
  - Both A and B use passwords (shared with key server) to decrypt return from key servers
  - Add in timestamps to limit how long tickets will be used to prevent attacker from replaying messages later
  - Also have to include encrypted checksums (hashed version of message) to prevent malicious user from inserting things into messages/changing messages
  - Want to minimize # times A types in password
    - » A→S (Give me temporary secret)
    - » S→A (Use  $K_{temp-sa}$  for next 8 hours)<sup>Ksa</sup>
    - » Can now use  $K_{temp-sa}$  in place of  $K_{sa}$  in protocol

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.14

## Public Key Encryption

- Can we perform key distribution without an authentication server?
  - Yes. Use a Public-Key Cryptosystem.
- Public Key Details
  - Don't have one key, have two:  $K_{public}$ ,  $K_{private}$ 
    - » Two keys are mathematically related to one another
    - » Really hard to derive  $K_{public}$  from  $K_{private}$  and vice versa
  - Forward encryption:
    - » Encrypt:  $(cleartext)^{K_{public}} = ciphertext_1$
    - » Decrypt:  $(ciphertext_1)^{K_{private}} = cleartext$
  - Reverse encryption:
    - » Encrypt:  $(cleartext)^{K_{private}} = ciphertext_2$
    - » Decrypt:  $(ciphertext_2)^{K_{public}} = cleartext$
  - Note that  $ciphertext_1 \neq ciphertext_2$ 
    - » Can't derive one from the other!
- Public Key Examples:
  - RSA: Rivest, Shamir, and Adleman
    - »  $K_{public}$  of form  $(k_{public}, N)$ ,  $K_{private}$  of form  $(k_{private}, N)$
    - »  $N = pq$ . Can break code if know p and q
  - ECC: Elliptic Curve Cryptography


5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.15

## Public Key Encryption Details

- Idea:  $K_{public}$  can be made public, keep  $K_{private}$  private




- Gives message privacy (restricted receiver):
  - Public keys (secure destination points) can be acquired by anyone/used by anyone
  - Only person with private key can decrypt message
- What about authentication?
  - Use combination of private and public key
  - Alice→Bob: [(I'm Alice)<sup>A\_private</sup> Rest of message]<sup>B\_public</sup>
  - Provides restricted sender and receiver
- But: how does Alice know that it was Bob who sent her  $B_{public}$ ? And vice versa...

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.16

## Secure Hash Function



- Hash Function: Short summary of data (message)
  - For instance,  $h_1 = H(M_1)$  is the hash of message  $M_1$ 
    - »  $h_1$  fixed length, despite size of message  $M_1$ .
    - » Often,  $h_1$  is called the "digest" of  $M_1$ .
- Hash function  $H$  is considered secure if
  - It is infeasible to find  $M_2$  with  $h_1 = H(M_2)$ ; i.e. can't easily find other message with same digest as given message.
  - It is infeasible to locate two messages,  $m_1$  and  $m_2$ , which "collide", i.e. for which  $H(m_1) = H(m_2)$
  - A small change in a message changes many bits of digest/can't tell anything about message given its hash

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.17


## Use of Hash Functions

- Several Standard Hash Functions:

- MD5: 128-bit output
- **Obsolete: Use SHA-256 or SHA 512**

- Can we use hashing to securely reduce load on server?

- Yes. Use a series of insecure mirror servers (caches)
  - » Use secure channel with server
- First, ask server for digest of desired file
  - » Can be insecure channel
- Then ask mirror server for file
  - » Check digest of result and catch faulty or malicious mirrors



5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.18

BREAK

## Signatures/Certificate Authorities

- Can use  $X_{\text{public}}$  for person  $X$  to define their identity
  - Presumably they are the only ones who know  $X_{\text{private}}$ .
  - Often, we think of  $X_{\text{public}}$  as a "principle" (user)
- Suppose we want  $X$  to sign message  $M$ ?
  - Use private key to encrypt the digest, i.e.  $H(M)^{X_{\text{private}}}$
  - Send both  $M$  and its signature:
    - » Signed message =  $[M, H(M)^{X_{\text{private}}}]$
  - Now, anyone can verify that  $M$  was signed by  $X$ 
    - » Simply decrypt the digest with  $X_{\text{public}}$
    - » Verify that result matches  $H(M)$
- Now: How do we know that the version of  $X_{\text{public}}$  that we have is really from  $X$ ???
- Answer: Certificate Authority
  - » Examples: Verisign, Entrust, Etc.
- $X$  goes to organization, presents identifying papers
  - » Organization signs  $X$ 's key:  $[X_{\text{public}}, H(X_{\text{public}})^{C_{\text{private}}}]$
  - » Called a "Certificate"
- Before we use  $X_{\text{public}}$ , ask  $X$  for certificate verifying key
  - » Check that signature over  $X_{\text{public}}$  produced by trusted authority
- How do we get keys of certificate authority?
  - Compiled into your browser, for instance!


5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.20

## Security through SSL

- **SSL Web Protocol**
  - Port 443: secure http
  - Use public-key encryption for key-distribution
- Server has a **certificate** signed by certificate authority
  - Contains server info (organization, IP address, etc)
  - Also contains server's public key and expiration date
- Establishment of Shared, 48-byte "master secret"
  - Client sends 28-byte random value  $n_c$  to server
  - Server returns its own 28-byte random value  $n_s$ , plus its certificate  $cert_s$
  - Client verifies certificate by checking with public key of certificate authority compiled into browser
    - » Also check expiration date
  - Client picks 46-byte "premaster" secret (pms), encrypts it with public key of server, and sends to server
  - Now, both server and client have  $n_c$ ,  $n_s$ , and pms
    - » Each can compute 48-byte master secret using one-way and collision-resistant function on three values
    - » Random "nonces"  $n_c$  and  $n_s$  make sure master secret fresh



5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.21

## SSL Pitfalls

- Netscape claimed to provide secure comm. (SSL)
  - So you could send a credit card # over the Internet
- Three problems (reported in NYT):
  - Algorithm for picking session keys was predictable (used time of day) - brute force key in a few hours
  - Made new version of Netscape to fix #1, available to users over Internet (unencrypted!)
    - » Four byte patch to Netscape executable makes it always use a specific session key
    - » Could insert backdoor by mangling packets containing executable as they fly by on the Internet.
    - » Many mirror sites (including Berkeley) to redistribute new version - anyone with root access to any machine on LAN at mirror site could insert the backdoor
  - Buggy helper applications - can exploit *any* bug in either Netscape, or its helper applications

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.22

## Authorization: Who Can Do What?

- How do we decide who is authorized to do actions in the system?
- **Access Control Matrix:** contains all permissions in the system
  - Resources across top
    - » Files, Devices, etc...
  - Domains in columns
    - » A domain might be a user or a group of permissions
    - » E.g. above: User D3 can read F2 or execute F3
  - In practice, table would be huge and sparse!



object	F <sub>1</sub>	F <sub>2</sub>	F <sub>3</sub>	printer
domain				
D <sub>1</sub>	read		read	
D <sub>2</sub>				print
D <sub>3</sub>		read	execute	
D <sub>4</sub>	read write		read write	

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.23

## Authorization: Two Implementation Choices

- **Access Control Lists:** store permissions with object
  - Still might be lots of users!
  - UNIX limits each file to: r,w,x for owner, group, world
  - More recent systems allow definition of groups of users and permissions for each group
  - ACLs allow easy changing of an object's permissions
    - » Example: add Users C, D, and F with rw permissions
- **Capability List:** each process tracks which objects has permission to touch
  - Popular in the past, idea out of favor today
  - Consider page table: Each process has list of pages it has access to, not each page has list of processes ...
  - Capability lists allow easy changing of a domain's permissions
    - » Example: you are promoted to system administrator and should be given access to all system files

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.24

## Authorization: Combination Approach



- Users have capabilities, called "groups" or "roles"
  - Everyone with particular group access is "equivalent" when accessing group resource
  - Like passport (which gives access to country of origin)
- Objects have ACLs
  - ACLs can refer to users or groups
  - Change object permissions by modifying ACL
  - Change broad user permissions via changes in group membership
  - Possessors of proper credentials get access

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.25

## Authorization: How to Revoke?

- How does one revoke someone's access rights to a particular object?
  - Easy with ACLs: just remove entry from the list
  - Takes effect immediately since the ACL is checked on each object access
- Harder to do with capabilities since they aren't stored with the object being controlled:
  - Not so bad in a single machine: could keep all capability lists in a well-known place (e.g., the OS capability table).
  - Very hard in distributed system, where remote hosts may have crashed or may not cooperate (more in a future lecture)

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.26

## Revoking Capabilities

- Various approaches to revoking capabilities:
  - Put expiration dates on capabilities and force reacquisition
  - Put epoch numbers on capabilities and revoke all capabilities by bumping the epoch number (which gets checked on each access attempt)
  - Maintain back pointers to all capabilities that have been handed out (Tough if capabilities can be copied)
  - Maintain a revocation list that gets checked on every access attempt

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.27

## How fine-grained should access control be?

- Example of the problem:
  - Suppose you buy a copy of a new game from "Joe's Game World" and then run it.
  - It's running with your userid
    - » It removes all the files you own, including the project due the next day..
- How can you prevent this?
  - Have to run the program under *some* userid.
    - » Could create a second *games* userid for the user, which has no write privileges.
    - » Like the "nobody" userid in UNIX - can't do much
  - But what if the game needs to write out a file recording scores?
    - » Would need to give write privileges to one particular file (or directory) to your *games* userid.
  - But what about non-game programs you want to use, such as Quicken?
    - » Now you need to create your own private *quicken* userid, if you want to make sure tha the copy of Quicken you bought can't corrupt non-quicken-related files

- But - how to get this right??? Pretty complex...

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.28

### Authorization Continued

- **Principle of least privilege:** programs, users, and systems should get only enough privileges to perform their tasks
  - Very hard to do in practice
    - » How do you figure out what the minimum set of privileges is needed to run your programs?
  - People often run at higher privilege than necessary
    - » Such as the "administrator" privilege under windows
- **One solution: Signed Software**
  - Only use software from sources that you trust, thereby dealing with the problem by means of authentication
  - Fine for big, established firms such as Microsoft, since they can make their signing keys well known and people trust them
    - » Actually, not always fine: recently, one of Microsoft's signing keys was compromised, leading to malicious software that looked valid
  - What about new startups?
    - » Who "validates" them?
    - » How easy is it to fool them?

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.29

### Summary

- **User Identification**
  - Passwords/Smart Cards/Biometrics
- **Passwords**
  - Encrypt them to help hid them
  - Force them to be longer/not amenable to dictionary attack
  - Use zero-knowledge request-response techniques
- **Distributed identity**
  - Use cryptography
- **Private Key Encryption (also Symmetric Key)**
  - Single Key used to encode and decode
  - Pros: Very Fast
    - » can encrypt at network speed (even without hardware)
  - Cons: Need to distribute *secret* key to both parties

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.30

### Summary (cont'd)

- **Secure Hash Function**
  - Fixed length summary of data
  - Hard to find another block of data with same hash
- **Public Key Encryption (also Asymmetric Key)**
  - Two keys: a public key and a private key
    - » Not derivable from one another
  - Pros: Can distribute keys in public
    - » Need certificate authority (Public Key Infrastructure)
  - Cons: Very Slow
    - » 100–1000 times slower than private key encryption
- **Session Key**
  - Randomly generated private key used for single session
  - Often distributed via public key encryption

5/5/08

Joseph CS162 ©UCB Spring 2008

Lec 25.31