

CS162 Operating Systems and Systems Programming Lecture 25

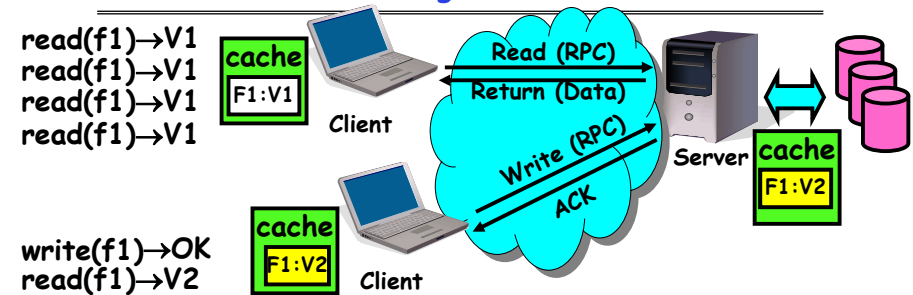
Protection and Security in Distributed Systems

November 28, 2007

Prof. John Kubiatowicz

<http://inst.eecs.berkeley.edu/~cs162>

Review: Use of caching to reduce network load



- Idea: Use caching to reduce network load
 - In practice: use buffer cache at source and destination
- Advantage: if open/read/write/close can be done locally, don't need to do any network traffic...fast!
- Problems:
 - Failure:
 - » Client caches have data not committed at server
 - Cache consistency!
 - » Client caches not consistent with server/each other

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.2

Goals for Today

- Finish discussing distributed file systems/Caching
- 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 Kubiatowicz.

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.3

Network File System (NFS)

- Three Layers for NFS system
 - **UNIX file-system interface**: open, read, write, close calls + file descriptors
 - **VFS layer**: distinguishes local from remote files
 - » Calls the NFS protocol procedures for remote requests
 - **NFS service layer**: bottom layer of the architecture
 - » Implements the NFS protocol
- NFS Protocol: RPC for file operations on server
 - Reading/searching a directory
 - manipulating links and directories
 - accessing file attributes/reading and writing files
- **Write-through caching**: Modified data committed to server's disk before results are returned to the client
 - lose some of the advantages of caching
 - time to perform write() can be long
 - Need some mechanism for readers to eventually notice changes! (more on this later)

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.4

NFS Continued

- NFS servers are **stateless**; each request provides all arguments require for execution
 - E.g. reads include information for entire operation, such as `ReadAt(inumber, position)`, not `Read(openfile)`
 - No need to perform network `open()` or `close()` on file - each operation stands on its own
- **Idempotent**: Performing requests multiple times has same effect as performing it exactly once
 - Example: Server crashes between disk I/O and message send, client resend read, server does operation again
 - Example: Read and write file blocks: just re-read or re-write file block - no side effects
 - Example: What about "remove"? NFS does operation twice and second time returns an advisory error
- **Failure Model**: Transparent to client system
 - Is this a good idea? What if you are in the middle of reading a file and server crashes?
 - Options (NFS Provides both):
 - » Hang until server comes back up (next week?)
 - » Return an error. (Of course, most applications don't know they are talking over network)

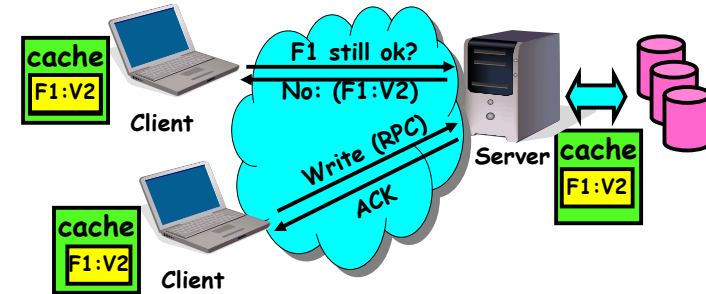
11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.5

NFS Cache consistency

- NFS protocol: weak consistency
 - Client polls server periodically to check for changes
 - » Polls server if data hasn't been checked in last 3-30 seconds (exact timeout is tunable parameter).
 - » Thus, when file is changed on one client, server is notified, but other clients use old version of file until timeout.



- What if multiple clients write to same file?
 - » In NFS, can get either version (or parts of both)
 - » Completely arbitrary!


11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.6

Sequential Ordering Constraints

- What sort of cache coherence might we expect?
 - i.e. what if one CPU changes file, and before it's done, another CPU reads file?
- Example: Start with file contents = "A"



- What would we actually want?
 - Assume we want distributed system to behave exactly the same as if all processes are running on single system
 - » If read finishes before write starts, get old copy
 - » If read starts after write finishes, get new copy
 - » Otherwise, get either new or old copy
 - For NFS:
 - » If read starts more than 30 seconds after write, get new copy; otherwise, could get partial update

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.7

NFS Pros and Cons

- NFS Pros:
 - Simple, Highly portable
- NFS Cons:
 - Sometimes inconsistent!
 - Doesn't scale to large # clients
 - » Must keep checking to see if caches out of date
 - » Server becomes bottleneck due to polling traffic

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.8

Andrew File System

- Andrew File System (AFS, late 80's) → DCE DFS (commercial product)
- **Callbacks:** Server records who has copy of file
 - On changes, server immediately tells all with old copy
 - No polling bandwidth (continuous checking) needed
- Write through on close
 - Changes not propagated to server until close()
 - Session semantics: updates visible to other clients only after the file is closed
 - » As a result, do not get partial writes: all or nothing!
 - » Although, for processes on local machine, updates visible immediately to other programs who have file open
- In AFS, everyone who has file open sees old version
 - Don't get newer versions until reopen file

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.9

Andrew File System (con't)

- Data cached on local disk of client as well as memory
 - On open with a cache miss (file not on local disk):
 - » Get file from server, set up callback with server
 - On write followed by close:
 - » Send copy to server; tells all clients with copies to fetch new version from server on next open (using callbacks)
- What if server crashes? Lose all callback state!
 - Reconstruct callback information from client: go ask everyone "who has which files cached?"
- AFS Pro: Relative to NFS, less server load:
 - Disk as cache ⇒ more files can be cached locally
 - Callbacks ⇒ server not involved if file is read-only
- For both AFS and NFS: central server is bottleneck!
 - Performance: all writes→server, cache misses→server
 - Availability: Server is single point of failure
 - Cost: server machine's high cost relative to workstation

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.10

World Wide Web

- Use client-side caching to reduce number of interactions between clients and servers and/or reduce the size of the interactions:
 - Time-to-Live (TTL) fields - HTTP "Expires" header from server
 - Client polling - HTTP "If-Modified-Since" request headers from clients
 - Server refresh - HTML "META Refresh tag" causes periodic client poll
- What is the polling frequency for clients and servers?
 - Could be adaptive based upon a page's age and its rate of change
- Server load is still significant!

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.11

WWW Proxy Caches

- Place caches in the network to reduce server load
 - But, increases latency in lightly loaded case
 - Caches near servers called "reverse proxy caches"
 - » Offloads busy server machines
 - Caches at the "edges" of the network called "content distribution networks"
 - » Offloads servers and reduce client latency
- Challenges:
 - Caching static traffic easy, but only ~40% of traffic
 - Dynamic and multimedia is harder
 - » Multimedia is a big win: Megabytes versus Kilobytes
 - Same cache consistency problems as before
- Caching is changing the Internet architecture
 - Places functionality at higher levels of comm. protocols

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.12

Administrivia

- **MIDTERM II: Monday December 3rd**
 - Next Monday
 - 6:00-9:00pm, 2050 Valley LSB
 - All material from last midterm and up to today (lectures 12-25)
 - Includes virtual memory
 - One page of handwritten notes, both sides
- **Review Session: Sunday, Dec 2nd**
 - 7:00-9:00, 306 Soda (Hopefully this time!)
- **Final Exam**
 - December 17th, 5:00-8:00pm, 10 Evans
 - Covers whole course (except last lecture)
 - Two pages of handwritten notes, both sides
- **Final Topics: Any suggestions?**

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.13

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

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.14

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?????

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.15

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 ⇒ must be user typing it
 - » Very common technique
 - **Smart Cards**
 - » Electronics embedded in card capable of providing long passwords or satisfying challenge → 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



11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.16

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 → one way transform(hash) → 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

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.17

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 → 1 day to crack
 - » In 2005, .01μs to check a password → 0.1 seconds to crack
 - Takes less time to check for all words in the dictionary!

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.18

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 @ 0.01μs/check
 - Unfortunately, people still pick common patterns
 - » e.g. Capitalize first letter of common word, add one digit

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.19

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


11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.20

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


11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.21

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

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.22

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:
 - » $A \rightarrow S$: [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
 - » $S \rightarrow A$: **Message** [Use K_{ab} (This is A ! Use K_{ab}) ^{K_{sb}}] ^{K_{sa}}
 - » This allows A to know, " S said use this key"
 - Whenever A wants to talk with B
 - » $A \rightarrow B$: **Ticket** [This is A ! Use K_{ab}] ^{K_{sb}}
 - » Now, B knows that K_{ab} is sanctioned by S

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.23

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
 - » $A \rightarrow S$ (Give me temporary secret)
 - » $S \rightarrow A$ (Use $K_{temp-sa}$ for next 8 hours) ^{K_{sa}}
 - » Can now use $K_{temp-sa}$ in place of K_{sa} in protocol

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.24

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: $(\text{cleartext})^{K_{\text{public}}} = \text{ciphertext}_1$
 - » Decrypt: $(\text{ciphertext}_1)^{K_{\text{private}}} = \text{cleartext}$
 - Reverse encryption:
 - » Encrypt: $(\text{cleartext})^{K_{\text{private}}} = \text{ciphertext}_2$
 - » Decrypt: $(\text{ciphertext}_2)^{K_{\text{public}}} = \text{cleartext}$
 - Note that $\text{ciphertext}_1 \neq \text{ciphertext}_2$
 - » 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


11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.25

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: $[(I'm Alice)^{A_{\text{private}}} \text{ Rest of message}]^{B_{\text{public}}}$
 - Provides restricted sender and receiver
- But: how does Alice know that it was Bob who sent her B_{public} ? And vice versa...

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.26

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

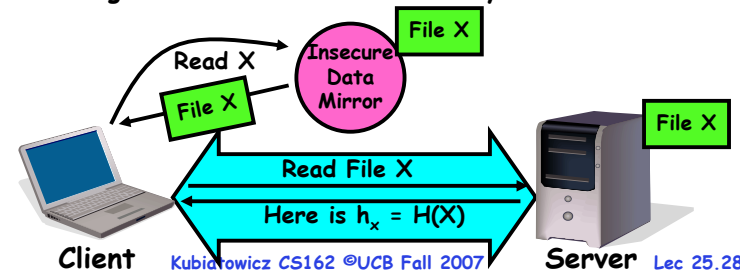
11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.27

Use of Hash Functions

- Several Standard Hash Functions:
 - MD5: 128-bit output
 - SHA-1: 160-bit output
- Can we use hashing to securely reduce load on server?
 - Yes. Use a series of insecure mirror servers (caches)
 - First, ask server for digest of desired file
 - » Use secure channel with server
 - Then ask mirror server for file
 - » Can be insecure channel
 - » Check digest of result and catch faulty or malicious mirrors



11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.28

Signatures/Certificate Authorities

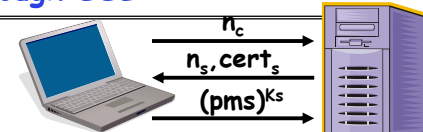
- Can use X_{public} for person X to define their identity
 - Presumably they are the only ones who know X_{private} .
 - Often, we think of X_{public} as a "principle" (user)
- Suppose we want X to sign message M?
 - Use private key to encrypt the digest, i.e. $H(M)^{X_{\text{private}}}$
 - Send both 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_{public}
 - » Verify that result matches $H(M)$
- Now: How do we know that the version of X_{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_{public} , ask X for certificate verifying key
 - » Check that signature over X_{public} produced by trusted authority
- How do we get keys of certificate authority?
 - Complied into your browser, for instance!

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.29

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

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.30

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

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.31

Conclusion

- 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
- Symmetrical (or Private Key) Encryption
 - Single Key used to encode and decode
 - Introduces key-distribution problem
- Public-Key Encryption
 - Two keys: a public key and a private key
 - » Not derivable from one another
- Secure Hash Function
 - Used to summarize data
 - Hard to find another block of data with same hash

11/28/07

Kubiatowicz CS162 ©UCB Fall 2007

Lec 25.32