

## Secure Communication with an Insecure Internet Infrastructure

## But first: some spam!

- If you rocked 15-441 (or are doing so), you might consider:
  - 15-610 next semester
  - Conviva (Very CMU CSD startup)'s looking for summer interns
  - Dave's looking for a few students for projects

2

## 15-610: Engineering Complex Large-scale Computer Systems

<http://www.cs.cmu.edu/~15-610>

**M. Satyanarayanan & Jan Harkes**  
**School of Computer Science**  
**Carnegie Mellon University**

## Vision of this Course

This is a master's level course to prepare students for technical leadership roles in creating and evolving the complex, large-scale computer systems that society will increasingly depend on in the future.

The course will teach the organizing principles of such systems, identifying a core set of versatile techniques that are applicable across many system layers.

Students will acquire the knowledge base, intellectual tools, hands-on skills and modes of thought needed to build well-engineered computer systems that withstand the test of time, growth in scale, and stresses of live use.

Strong design and implementation skills are expected of all students. The course assumes a high level of proficiency in all aspects of operating system design and implementation.

# Course Overview

## Target audience

- already possess strong hands-on systems skills
- desire careers as creators of major computer systems
- seek mastery of system design and implementation skills

15-410++

## Approach

- small but versatile *conceptual toolkit* of systems techniques
- immersive *hands-on experience* in applying this toolkit
- *case studies* to learn hard-won experience of others



© 2006-2008 M. Satyanarayanan

Preview of 15-610 for Spring 2009

preview -

# Conceptual Toolkit

*Caching* for performance and availability

*Prefetching* for performance and availability

*Content-Addressable Storage* for performance

*Damage containment & replication* for reliability and availability

## Challenges of size and longevity

- *Scale reduction* for performance and usability
- *Reducing fragmentation* for performance and manageability
- *Hints* for performance and scaling

*Coping with human foibles* for robustness

- limitations of individual users
- limitations of large groups of users

© 2006-2008 M. Satyanarayanan

Preview of 15-610 for Spring 2009

preview -

# Hands-on Projects

## Series of 4 projects

### Based on a single open-source base (Coda File System)

- embodies many of concepts discussed in class
- almost entirely user-level implementation
- local expertise

## Individual projects

### Hardware donated by Intel for course

- loaner laptop for each student

© 2006-2008 M. Satyanarayanan

Preview of 15-610 for Spring 2009

preview -

# Conviva Internship

## Live Internet media streaming

- Directly from CMU/Berkeley research (Hui Zhang et al.) on overlay multicast
- (How do you stream media to 100,000 people on the Internet, with high quality, without a huge fixed infrastructure??)
- Highly-available, scalable back-end services
- Large-scale data analytics and visualization
- Distributed software testing and automation
- If parts of this sound similar to a 441 project, don't be surprised. :) Overlay & p2p multicast is becoming important in the real world.
- Std. qualifications - network programming, C/C++/Java, Python, etc. [university@conviva.com](mailto:university@conviva.com) for more info.

© 2006-2008 M. Satyanarayanan

Preview of 15-610 for Spring 2009

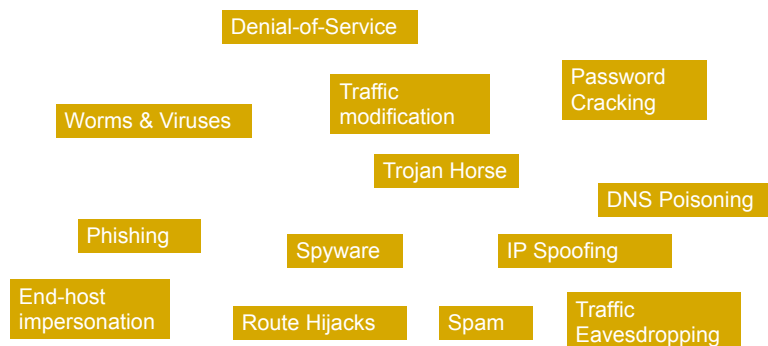
preview -

## dga summer projects

- Building systems for improving Web security
  - <http://www.cs.cmu.edu/~dwendlan/perspectives/>
- Prototyping novel Internet architecture features:
  - <http://www.cs.cmu.edu/~dga/papers/aip-hotnets2007-abstract.html>
- Building large-scale data analysis techniques on mid-sized clusters

## Back to our schedule...

## What is “Internet Security” ?



Many things to many people!

- 1) Attacks and vulnerabilities at all layers of the stack
- 2) Attackers will attack the most vulnerable / profitable components

## Internet Design Decisions: (ie: how did we get here? )

- Origin as a small and cooperative network  
(=> largely trusted infrastructure)
- Global Addressing  
(=> every sociopath is your next-door neighbor\*)
- Connection-less datagram service  
(=> can't verify source, hard to protect bandwidth)

\* Dan Geer

## Internet Design Decisions: (ie: how did we get here?)

- Anyone can connect
  - ANYONE can connect...
- Millions of hosts run nearly identical software
  - single exploit can create epidemic
- Most Internet users know about as much as Senator Stevens (aka “the tubes guy”)
  - God help us all...

## Our “Narrow” Focus


Yes:

- 1) Creating a “secure channel” for communication (today)
- 2) Protecting network resources and limiting connectivity (last time)


No:

- 1) Preventing software vulnerabilities & malware, or “social engineering”.


## Secure Communication with an Untrusted Infrastructure



## Secure Communication with an Untrusted Infrastructure



## Secure Communication with an Untrusted Infrastructure



## What do we need for a secure communication channel?

- Authentication (Who am I talking to?)
- Confidentiality (Is my data hidden?)
- Integrity (Has my data been modified?)
- Availability (Can I reach the destination?)

## What is cryptography?

"cryptography is about communication in the presence of adversaries."

- Ron Rivest

"cryptography is using math and other crazy tricks to approximate magic"

- Unknown 441 TA

## What is cryptography?

Tools to help us build secure communication channels that provide:

- 1) Authentication
- 2) Integrity
- 3) Confidentiality

## Cryptography As a Tool

- Using cryptography securely is not simple
- Designing cryptographic schemes correctly is near impossible.

Today we want to give you an idea of what can be done with cryptography.

Take a security course if you think you may use it in the future (e.g. 18-487)

## The Great Divide

Symmetric Crypto:  
(Commonly (mis)-  
called Private  
key)

Asymmetric Crypto:  
(Public key)  
Example: RSA

Requires a pre-  
shared secret  
between  
communicating  
parties?

Yes

No

Overall speed of  
cryptographic  
operations

Fast

Slow

## Symmetric Key: Confidentiality

### Motivating Example:

You and a friend share a key  $K$  of  $L$  random bits, and a message  $M$  also  $L$  bits long.

### Scheme:

You send her the  $xor(M, K)$  and then they “decrypt” using  $xor(M, K)$  again.

- 1) Do you get the right message to your friend?
- 2) Can an adversary recover the message  $M$ ?

## Symmetric Key: Confidentiality

- One-time Pad (OTP) is secure but usually impractical
  - Key is as long as the message
  - Keys cannot be reused (why?)

In practice, two types of ciphers are used  
that require only constant key length:

### **Stream Ciphers:**


Ex: RC4, A5

### **Block Ciphers:**

Ex: DES, AES, Blowfish

## Symmetric Key: Confidentiality

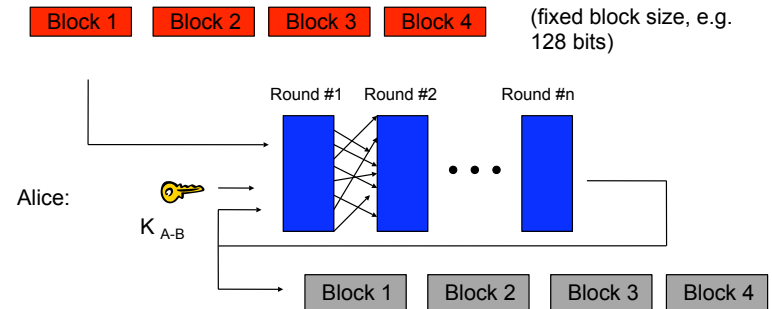
### Stream Ciphers (ex: RC4)



Bob uses  $K_{A-B}$  as PRNG seed, and XORs encrypted text to get the message back (just like OTP).

## Symmetric Key: Confidentiality

### Block Ciphers (ex: AES)



Bob breaks the ciphertext into blocks, feeds it through decryption engine using  $K_{A-B}$  to recover the message.

## Symmetric Key: Integrity


### Background: Hash Function Properties

- Consistent  
hash(X) always yields same result
- One-way  
X, can't find Y s.t. hash(Y) = X
- Collision resistant

given

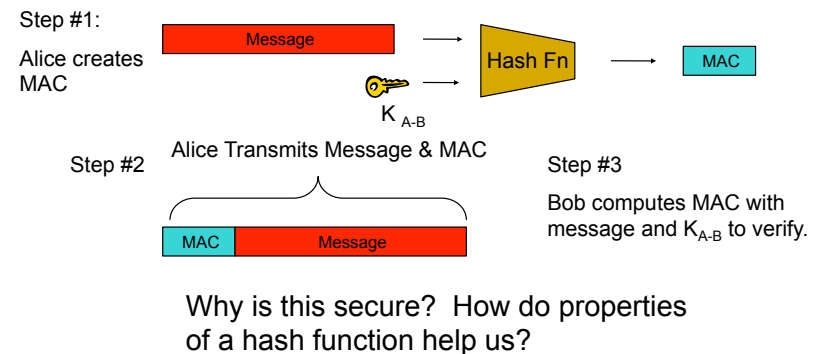
hash(W) = Z, can't find X such that hash(X) = Z

given




## Symmetric Key: Integrity

### Hash Message Authentication Code (HMAC)




## Symmetric Key: Authentication

- You already know how to do this!  
(hint: think about how we showed integrity)




## Symmetric Key: Authentication

What if Mallory overhears the hash sent by Bob, and then "replays" it later?



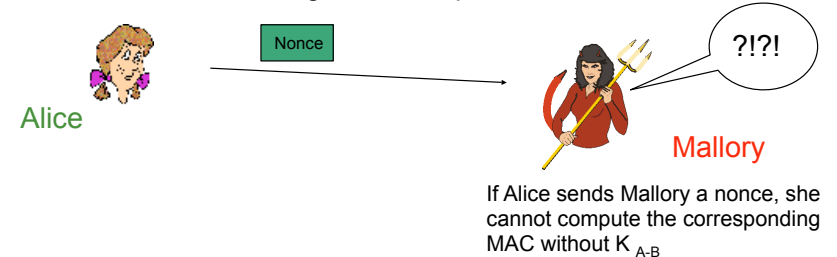
## Symmetric Key: Authentication

- A "Nonce"
  - A random bitstring used only once. Alice sends nonce to Bob as a "challenge". Bob Replies with "fresh" MAC result.



## Symmetric Key: Authentication

- A "Nonce"
  - A random bitstring used only once. Alice sends nonce to Bob as a "challenge". Bob Replies with "fresh" MAC result.





## Symmetric Key Crypto Review

- Confidentiality: Stream & Block Ciphers
- Integrity: HMAC
- Authentication: HMAC and Nonce


Questions??

Are we done? Not Really:

- 1) Number of keys scales as  $O(n^2)$
- 2) How to securely share keys in the first place?

## Asymmetric Key Crypto:

- Instead of shared keys, each person has a “key pair”



- The keys are inverses, so:  $K_B^{-1}(K_B(m)) = m$

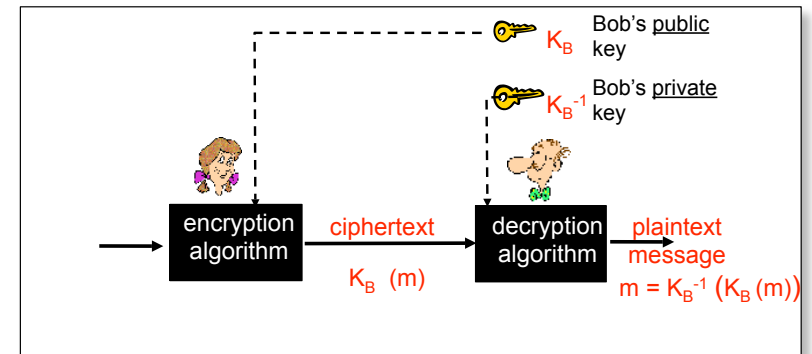
## Asymmetric Key Crypto:

- It is believed to be computationally unfeasible to derive  $K_B^{-1}$  from  $K_B$  or to find any way to get  $M$  from  $K_B(M)$  other than using  $K_B^{-1}$ .

=>  $K_B$  can safely be made public.

Note: We will not detail the computation that  $K_B(m)$  entails, but rather treat these functions as black boxes with the desired properties.

## Asymmetric Key: Confidentiality



## Asymmetric Key: Sign & Verify

- If we are given a message  $M$ , and a value  $S$  such that  $K_B(S) = M$ , what can we conclude?
- The message must be from Bob, because it must be the case that  $S = K_B^{-1}(M)$ , and only Bob has  $K_B^{-1}$ !
- This gives us two primitives:
  - $\text{Sign}(M) = K_B^{-1}(M) = \text{Signature } S$
  - $\text{Verify}(S, M) = \text{test}(K_B(S) == M)$

## Asymmetric Key: Integrity & Authentication


- We can use  $\text{Sign}()$  and  $\text{Verify}()$  in a similar manner as our HMAC in symmetric schemes.

Integrity:



Receiver must only check  $\text{Verify}(M, S)$

Authentication:



## Asymmetric Key Review:

- Confidentiality: Encrypt with Public Key of Receiver
- Integrity: Sign message with private key of the sender
- Authentication: Entity being authenticated signs a nonce with private key, signature is then verified with the public key

But, these operations are computationally expensive\*

## One last “little detail”...

How do I get these keys in the first place??

Remember:

- Symmetric key primitives assumed Alice and Bob had already shared a key.
- Asymmetric key primitives assumed Alice knew Bob's public key.

This may work with friends, but when was the last time you saw Amazon.com walking down the street?


## Symmetric Key Distribution

- How does Andrew do this?

Andrew Uses Kerberos, which relies on a Key Distribution Center (KDC) to establish shared symmetric keys.

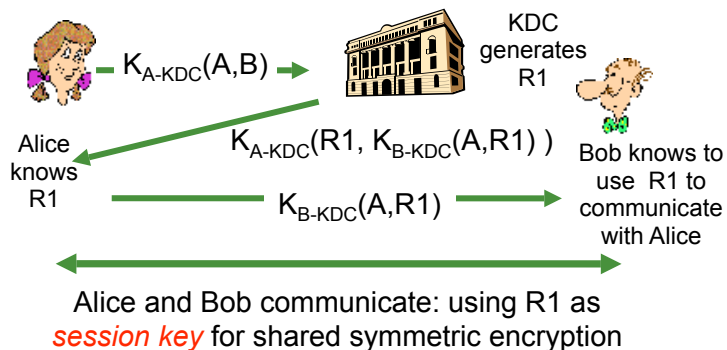
## Key Distribution Center (KDC)

- Alice, Bob need shared symmetric key.
- KDC**: server shares different secret key with *each* registered user (many users)
- Alice, Bob know own symmetric keys,  $K_{A-KDC}$   $K_{B-KDC}$ , for communicating with KDC.



## Key Distribution Center (KDC)

**Q:** How does KDC allow Bob, Alice to determine shared symmetric secret key to communicate with each other?




## How Useful is a KDC?

- Must always be online to support secure communication
- KDC can expose our session keys to others!
- Centralized trust and point of failure.

In practice, the KDC model is mostly used within single organizations (e.g. Kerberos) but not more widely.

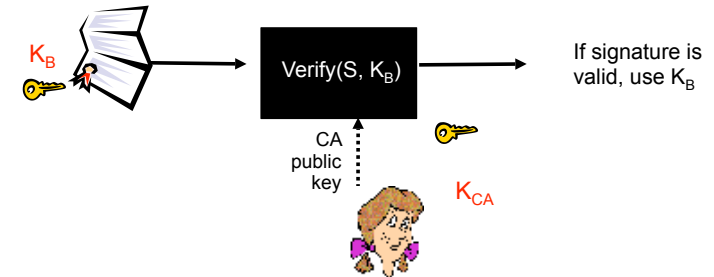
## Certification Authorities

- **Certification authority (CA):** binds public key to particular entity, E.
- An entity E registers its public key with CA.
  - E provides “proof of identity” to CA.
  - CA creates certificate binding E to its public key.
  - Certificate contains E’s public key AND the CA’s signature of E’s public key.



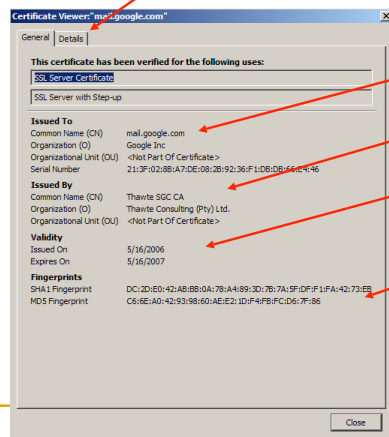
## Certification Authorities

- When Alice wants Bob's public key:
  - Gets Bob's certificate (Bob or elsewhere).
  - Use CA's public key to verify the signature within Bob's certificate, then accepts public key



## Certificate Contents

- info algorithm and key value itself (not shown)



- Cert owner
- Cert issuer
- Valid dates
- Fingerprint of signature

## Which Authority Should You Trust?

- Today: many authorities
- What about a shared Public Key Infrastructure (PKI)?
  - A system in which “roots of trust” authoritatively bind public keys to real-world identities
  - So far it has not been very successful

## Transport Layer Security (TLS) aka Secure Socket Layer (SSL)

- Used for protocols like HTTPS
- Special TLS socket layer between application and TCP (small changes to application).
- Handles confidentiality, integrity, and authentication.
- Uses “hybrid” cryptography.
  - e.g., encryption: encrypt with symmetric key;  
encrypt symmetric key w/public key (smaller!)


## What to take home?

- Internet design and growth => security challenges
- Symmetric (pre-shared key, fast) and asymmetric (key pairs, slow) primitives provide:
  - Confidentiality
  - Integrity
  - Authentication
- “Hybrid Encryption” leverages strengths of both.
- Great complexity exists in securely acquiring keys.
- Crypto is hard to get right, so use tools from others, don’t design your own (e.g. TLS).

## Resources

- Textbook: 8.1 – 8.3
- Wikipedia for overview of Symmetric/Asymmetric primitives and Hash functions.
- OpenSSL ([www.openssl.org](http://www.openssl.org)): top-rate open source code for SSL and primitive functions.
- “Handbook of Applied Cryptography” available free online: [www.cacr.math.uwaterloo.ca/hac/](http://www.cacr.math.uwaterloo.ca/hac/)

## Setup Channel with TLS “Handshake”



### Handshake Steps:

- 1) Clients and servers negotiate exact cryptographic protocols
- 2) Client's validate public key certificate with CA public key.
- 3) Client encrypt secret random value with server's key, and send it as a challenge.
- 4) Server decrypts, proving it has the corresponding private key.
- 5) This value is used to derive symmetric session keys for encryption & MACs.

## How TLS Handles Data

1) Data arrives as a stream from the application via the TLS Socket



2) The data is segmented by TLS into chunks



3) A session key is used to encrypt and MAC each chunk to form a TLS "record", which includes a short header and data that is encrypted, as well as a MAC.



4) Records form a byte stream that is fed to a TCP socket for transmission.

