


CS 417 – DISTRIBUTED SYSTEMS

# Week 12: Security in Distributed Systems

## Part 1: Cryptography Intro



Paul Krzyzanowski

© 2021 Paul Krzyzanowski. No part of this content, may be reproduced or reposted in whole or in part in any manner without the permission of the copyright owner.

# Goals of Security

Keep systems, programs, and data secure

The **CIA\* Triad**:

**1. Confidentiality**

**2. Integrity**

**3. Availability**

*\*No relationship to the Central Intelligence Agency*

# Confidentiality

- Keep data & resources hidden
  - Data will only be shared with authorized individuals
  - Sometimes – conceal the existence of data or communication
- Traditional focus of computer security

## Data confidentiality

“The property that information is not made available or disclosed to unauthorized individuals, entities, or processes [i.e., to any unauthorized system entity].”

– *RFC 4949, Internet Security Glossary*

# Integrity

The trustworthiness of the data or resources

- *Preventing unauthorized changes to the data or resources*

- **Data integrity**

- Data integrity: property that data has not been modified or destroyed in an unauthorized or accidental manner

- **Origin integrity**

- Authentication

- **System integrity**

- The ability of a system to perform its intended function, free from deliberate or inadvertent manipulation

Often more important than confidentiality!

# Availability

- Being able to use the data or resources
- Property of a system being accessible and capable of working to required performance specifications

*Turning off a computer provides confidentiality & integrity but hurts availability*

***Denial of Service (DoS) attacks target availability***

# Thinking about security

Security is not

just adding encryption

... or using a 512-bit key instead of a 64-bit key

... or changing passwords

... or setting up a firewall

It is a systems issue

= Hardware + firmware + OS + app software + networking + people

= Processes & procedures, policies, detection, forensics

*“Security is a chain: it’s only as secure as the weakest link”*

*– Bruce Schneier*

# The Internet Introduces Risks

“The internet was designed to be open, transparent, and interoperable. Security and identity management were secondary objectives in system design. This lower emphasis on security in the internet’s initial design not only gives attackers a built-in advantage. It can also make intrusions difficult to attribute, especially in real time. This structural property of the current architecture of cyberspace means that we cannot rely on the threat of retaliation alone to deter potential attackers. Some adversaries might gamble that they could attack us and escape detection.”

– *William J. Lynn III, Deputy Defense Secretary, 2010*

<http://archive.defense.gov/speeches/speech.aspx?speechid=1593>

# The Internet Makes It Easier To Attack

- Security was not a design consideration
- Intelligence is at the edges of the network – distributed among many players
- Access and routing not centrally managed
  - Routing decisions distributed
  - No access control: any system can be added to the Internet
- Bad actors can hide!



# How the Internet Creates Vulnerabilities

<b>Action at a distance</b>	<ul style="list-style-type: none"><li>• People can be beyond our control or visibility</li></ul>
<b>Asymmetric medium</b>	<ul style="list-style-type: none"><li>• Actors can project or harness greater force. Low barriers to entry. Offense can be more effective than defense. A small number of actors can have a large effect.</li><li>• E.g., Anonymous, fraud spam email, or Facebook requests for money.</li><li>• Sending millions of messages costs almost nothing</li><li>• Small counties can hurt countries like the US or China.</li></ul>
<b>Actors can be anonymous</b>	<ul style="list-style-type: none"><li>• Identifying a source can be difficult.</li><li>• Attack with impunity. Trust becomes a challenge. Are you really communicating with your bank? We don't know who fired the missile.</li></ul>
<b>No borders or checkpoints</b>	<ul style="list-style-type: none"><li>• China and North Korea (maybe Russia) are the only counties that control data flow to/from their country</li></ul>
<b>No distinction in data</b>	<ul style="list-style-type: none"><li>• Hard to distinguish valid data from attacks</li><li>• Can't tell what code will be harmful until it's executed</li></ul>

# The Operating System

The operating system normally handles security

- User authentication – *passwords*
- Access control – *file permissions, system call access*
- Resource management – *memory limits, scheduling*

But it can only control resources it owns

- Other systems may have different policies

# Distributed Control

Distributed systems often use components that belong to different entities

Programs may:

- Call remote services – *are they trustworthy?*
- Store data on remote servers – *who manages them?*
- Send data over a network – *what route do the packets take?*

# Cryptography $\neq$ Security

Cryptography may be a component of a secure system

Adding cryptography may not make a system secure

# Cryptography: what is it good for?

- Confidentiality
  - Others cannot read contents of the message
- Authentication
  - Determine origin of message
- Integrity
  - Verify that message has not been modified
- Nonrepudiation
  - Sender should not be able to falsely deny that a message was sent

# Cryptography is hard

Designing (and cracking) codes is an art & a science

- Algorithm security is hard to evaluate – takes skill and time

Weak algorithms destroy effectiveness of cryptography

- Enigma cipher machine
- Every NATO and Warsaw Pact algorithm during Cold War
- DVD (DeCSS, 1999), HD DVD (Dec 2006) and Blu-Ray (Jan 2007)
- High-bandwidth Digital Content Protection (HDCP)
- RC4 cipher
- Wired Equivalent Privacy (WEP for Wi-Fi)
- All digital cellular encryption algorithms
- Firewire

# Confidentiality

# Encryption

**Plaintext** (cleartext) message  $P$

**Cipher** = cryptographic algorithm

**Encryption**  $E(P)$

Produces **Ciphertext**,  $C = E(P)$

**Decryption**,  $P = D(C)$



# Properties of a good cryptosystem

1

Ciphertext should be indistinguishable from random values

2

Given ciphertext, there should be no way to extract the original plaintext or the key short of enumerating all possible keys (i.e., a *brute force attack*)

3

The keys should be large enough that a brute force attack is not feasible

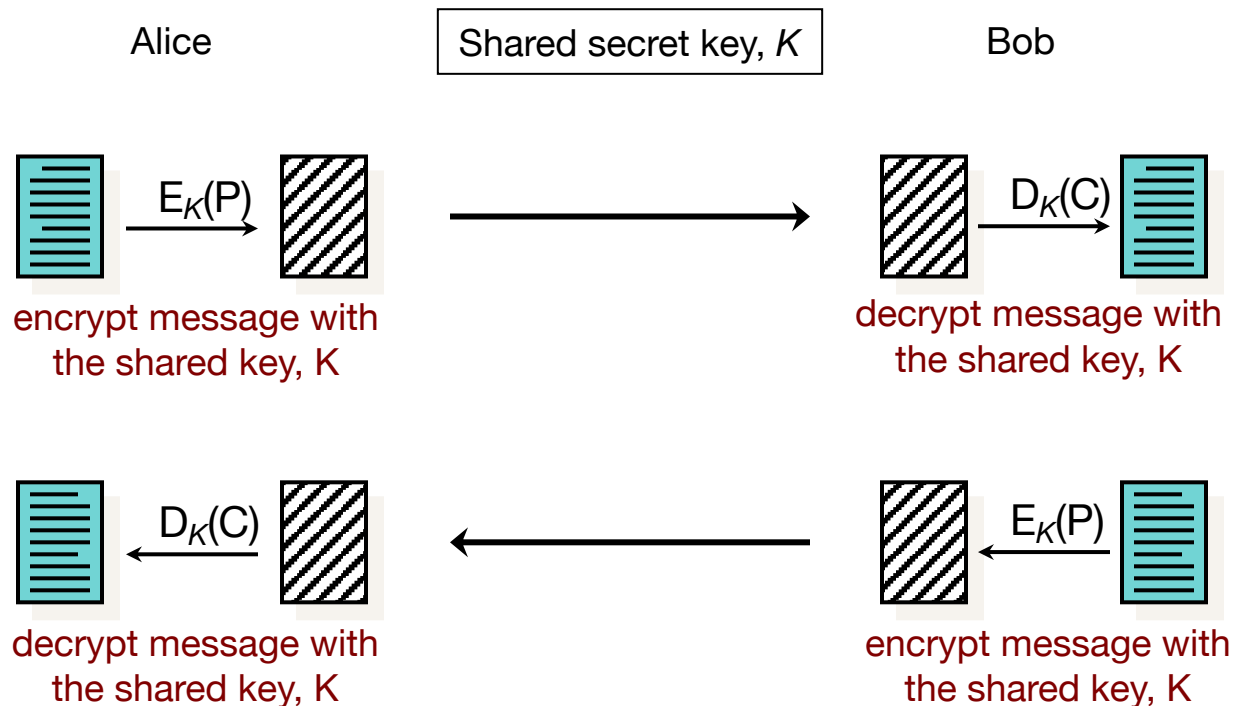
# Symmetric key ciphers

Same shared secret key,  $K$ , for encryption & decryption

$$C = E_K(P)$$

$$P = D_K(C)$$

# Communication with a symmetric cipher



AES (Advanced Encryption Standard)	<ul style="list-style-type: none"><li>• FIPS standard since 2002</li><li>• 128, 192, or 256-bit keys; operates on 128-bit blocks</li><li>• By far the most widely used symmetric encryption algorithm</li></ul>
DES, 3DES	<ul style="list-style-type: none"><li>• FIPS standard since 1976; 56-bit key; operates on 64-bit (8-byte) blocks</li><li>• Triple DES recommended since 1999 (112 or 168 bits)</li><li>• Not actively used anymore; AES is better by any measure</li></ul>
ChaCha20	128 or 256-bit keys – stream cipher
Twofish	128, 192 or 256 bits – block cipher
IDEA	128-bit keys; operates on 64-bit blocks More secure than DES but faster algorithms are available

# Public Key Cryptography

# Public-key algorithm

Two related keys ( $A, a$ )

$$\begin{array}{ll} C = E_A(P) & = D_a(C) \\ C' = E_a(P) & = D_A(C') \end{array} \left. \vphantom{\begin{array}{l} C = E_A(P) \\ C' = E_a(P) \end{array}} \right\} \begin{array}{l} A \text{ is a **public** key} \\ a \text{ is a **private** key} \end{array}$$

- Examples:
  - RSA
  - Elliptic Curve Cryptography (ECC)
- Key length
  - Unlike symmetric cryptography, not every number is a valid key
  - 3072-bit RSA  $\approx$  256-bit elliptic curve  $\approx$  128-bit symmetric cipher
  - 15360-bit RSA  $\approx$  521-bit elliptic curve  $\approx$  256-bit symmetric cipher

# Trapdoor functions

Public key cryptography relies on **trapdoor functions**

## Trapdoor function

- Easy to compute in one direction
- Inverse is difficult to compute without extra information

Example:

**125356232168767786330093** is the product of two prime #s  
What are they?

But if you're told that one of them is **139420885313**  
then it's easy to compute the other: **899120902061**


# Communication with public key algorithms

Different keys for encrypting and decrypting

- No need to worry about key distribution




# Communication with public key algorithms



# Communicating with symmetric cryptography

- Both parties must agree on a secret key,  $K$
- Message is encrypted, sent, decrypted at other side



**Key distribution must be secret**

*Otherwise, messages can be decrypted*

*Users can be impersonated*

# Key distribution

Secure key distribution is the biggest problem with symmetric cryptography

# Distributing Keys

- **Pre-shared keys**
  - Initial configuration, out of band (send via USB key, recite, ...)
- **Trusted third party**
  - Knows all keys
  - Alice creates a temporary key (**session key**)
  - Encrypts it with her key – sends to Trent
  - Trent decrypts it and sends it to Bob
  - Alternatively: Trent creates a session key – encrypts it for Alice & for Bob
- **Public key cryptography**
  - Alice encrypts a message with Bob's public key
  - Only Bob can decrypt
- **Diffie-Hellman**
- **Hybrid cryptosystems**

# Permanent vs. Ephemeral Keys

## Permanent keys

- Keys you use over and over again – e.g., your password

## Ephemeral keys

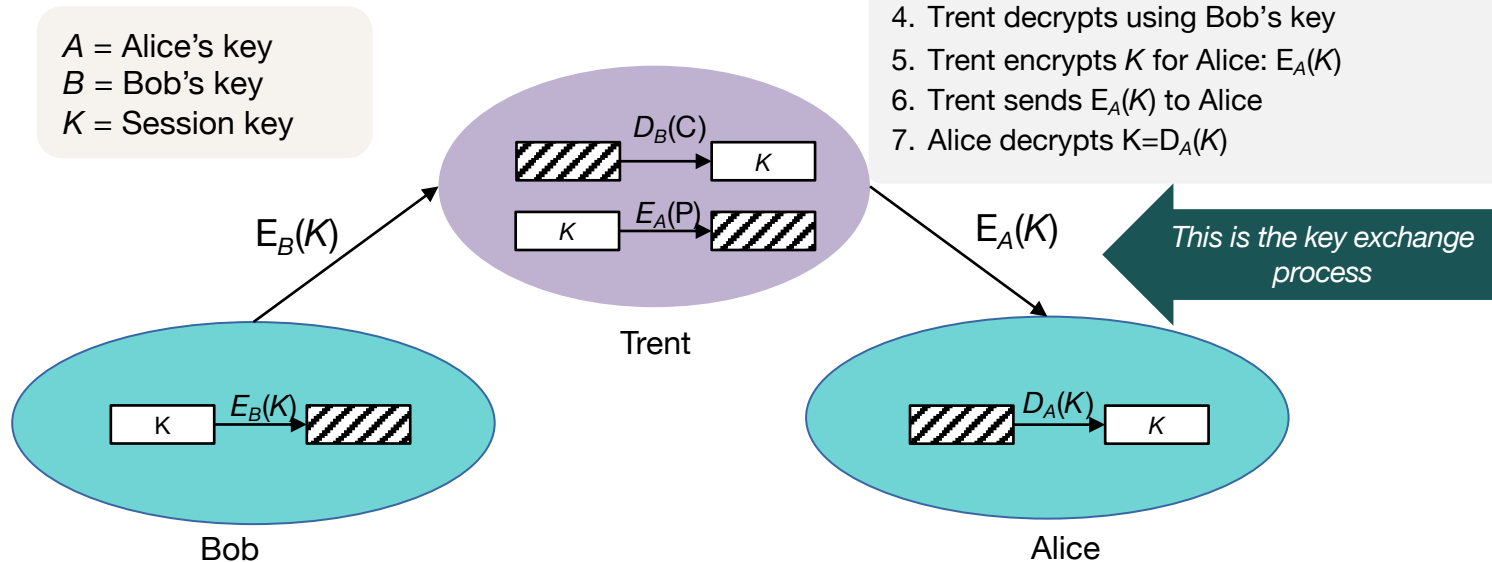
- Keys that are created spontaneously for one use, such as a communication session, and then never used again
- Session keys

## Why use ephemeral keys?

- The more data is encrypted with the same key, the easier it is for cryptanalysts to try to mount attacks
- We may have key exchange protocols that need to create a key for two parties to communicate

# Key exchange with a trusted third party

- Trusted third party, Trent, knows all the keys
- Everyone else only knows their own keys




We'd need to enhance this to avoid *replay attacks*: time stamps, sequence numbers

# Key exchange with a trusted third party

... continued

$A$  = Alice's key  
 $B$  = Bob's key  
 $K$  = Session key

1. Bob creates a random session key,  $K$
2. Bob encrypts it with his secret key:  $E_B(K)$
3. Bob sends  $E_B(K)$  to Trent
4. Trent decrypts using Bob's key
5. Trent encrypts  $K$  for Alice:  $E_A(K)$
6. Trent sends  $E_A(K)$  to Alice
7. Alice decrypts  $K = D_A(K)$
8. **Alice & Bob communicate, encrypting messages with the session key,  $K$**



# Diffie-Hellman Key Exchange

## Key distribution algorithm

- Allows two parties to share a secret key over a non-secure channel
- Not public key encryption
- Based on difficulty of computing discrete logarithms in a finite field compared with ease of calculating exponentiation

Allows us to negotiate a secret **common key** without fear of eavesdroppers:

$$\text{common key} = f(\text{your\_private\_key}, \text{their\_public\_key})$$



# Diffie-Hellman Key Exchange

- All arithmetic performed in a field of integers modulo some large number
- Both parties agree on
  - a **large prime number  $p$**
  - and a **number  $\alpha < p$**
- Each party generates a public/private key pair

Private key for user  $i$ :  $X_i$

Public key for user  $i$ :  $Y_i = \alpha^{X_i} \bmod p$

# Diffie-Hellman exponential key exchange

- Alice has secret key  $X_A$
- Alice sends Bob public key  $Y_A$
- Alice computes
- Bob has secret key  $X_B$
- Bob sends Alice public key  $Y_B$

$$K = Y_B^{X_A} \bmod p$$

**$K = (\text{Bob's public key}) (\text{Alice's private key}) \bmod p$**

# Diffie-Hellman exponential key exchange

- Alice has secret key  $X_A$
- Alice sends Bob public key  $Y_A$
- Alice computes

$$K = Y_B^{X_A} \bmod p$$

- Bob has secret key  $X_B$
- Bob sends Alice public key  $Y_B$
- Bob computes

$$K = Y_A^{X_B} \bmod p$$

$$\mathbf{K' = (Alice's\ public\ key)\ (Bob's\ private\ key)\ mod\ p}$$

# Diffie-Hellman exponential key exchange

- Alice has secret key  $X_A$
- Alice sends Bob public key  $Y_A$
- Alice computes

$$K = Y_B^{X_A} \bmod p$$

- expanding:

$$\begin{aligned} K &= Y_B^{X_A} \bmod p \\ &= (\alpha^{X_B} \bmod p)^{X_A} \bmod p \\ &= \alpha^{X_B X_A} \bmod p \end{aligned}$$

- Bob has secret key  $X_B$
- Bob sends Alice public key  $Y_B$
- Bob computes

$$K = Y_A^{X_B} \bmod p$$

- expanding:

$$\begin{aligned} K &= Y_B^{X_A} \bmod p \\ &= (\alpha^{X_A} \bmod p)^{X_B} \bmod p \\ &= \alpha^{X_A X_B} \bmod p \end{aligned}$$

$$\mathbf{K = K'}$$

$K$  is a common key, known *only* to Bob and Alice

# Hybrid Cryptosystems

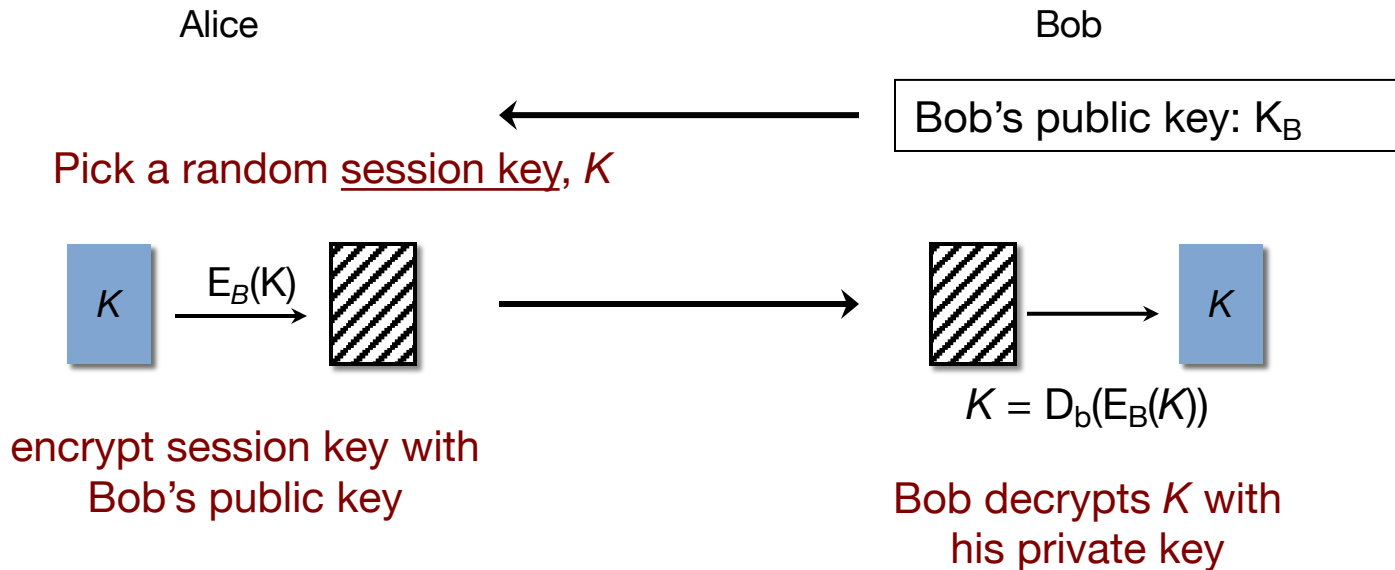
# Hybrid Cryptosystems

- **Session key**: randomly-generated key for one communication session
- Use a **public key algorithm** to send the session key
- Use a **symmetric algorithm** to encrypt data with the session key

Public key algorithms are almost never used to encrypt messages

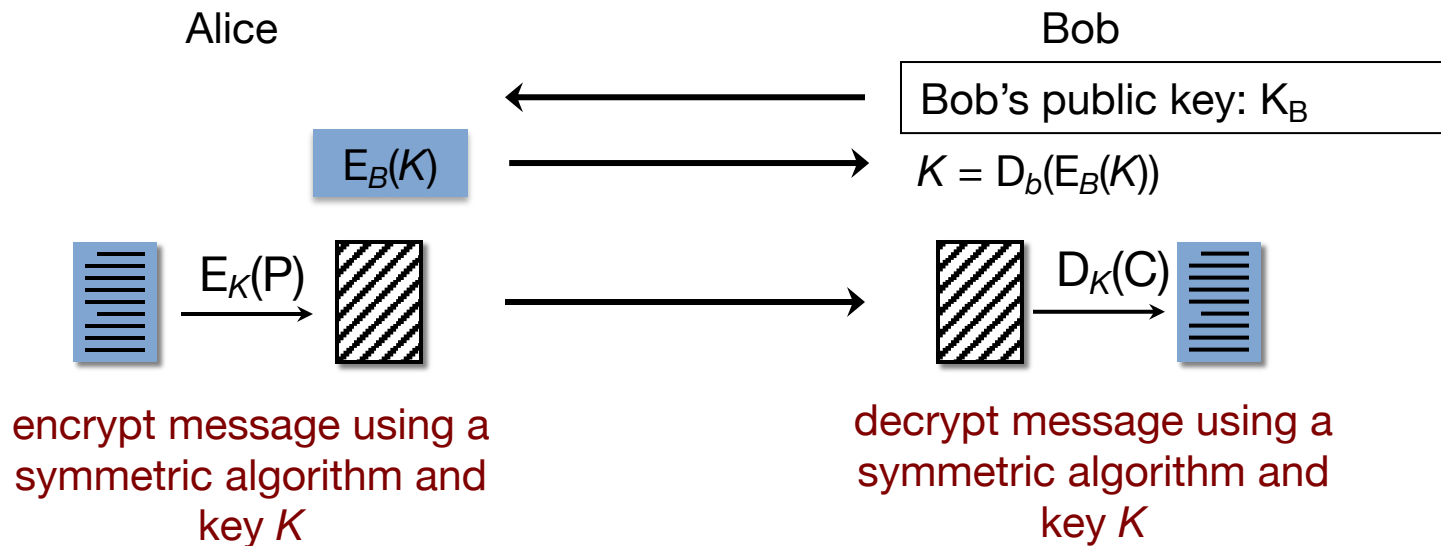
- MUCH slower; vulnerable to *chosen-plaintext attacks*
- RSA-2048 approximately 55x slower to encrypt and 2,000x slower to decrypt than AES-256

# Communication with a hybrid cryptosystem




Now Bob knows the secret session key,  $K$

# Communication with a hybrid cryptosystem





# Communication with a hybrid cryptosystem



decrypt message using a  
symmetric algorithm and  
key  $K$

encrypt message using a  
symmetric algorithm and  
key  $K$

# Cryptographic systems: summary

- Symmetric ciphers
  - Shared secret key
- Asymmetric ciphers – public key cryptosystems
  - Based on **trapdoor** functions
- Hybrid cryptosystem
  - Use a **public key algorithm** to send a randomly-chosen **session key**
  - Use a **symmetric key** algorithm to encrypt traffic back & forth
- Key exchange algorithms
  - Trusted Third Party
  - Diffie Hellman
  - Public key

} *Enables secure communication without using a 3<sup>rd</sup> party or knowledge of a shared secret*

# The End