



Security Potpourri!



CS 161: Computer Security

**Guest Lecturers: Frank Li, Rebecca Portnoff,
Grant Ho, Rishabh Poddar**


Instructor: Prof. Vern Paxson

TAs: Paul Bramsen, Apoorva Dornadula, David Fifield, Mia Gil Epner,
David Hahn, Warren He, Grant Ho, Frank Li, Nathan Malkin, Mitar
Milutinovic, Rishabh Poddar, Rebecca Portnoff, Nate Wang

<http://inst.eecs.berkeley.edu/~cs161/>

April 27, 2017

Side Channel Attacks



Side Channels

- Security systems are implemented in software or hardware on physical devices, which interact with their environment.
- Sometimes, attackers can monitor or affect these physical interactions, leaking useful “side channel” information.
- **Side channel attacks** use this information.

Note: Hard to identify due to abstractions.



Attacking Password Checker

```
/* Tenex (old OS) system call to check if submitted password is correct. */
```

```
bool CheckPassword(char* submitted_password, char* user){  
    char* real_password = GetUserPassword(user);  
    for (int i = 0; submitted_password[i] && real_password[i]; ++i) {  
        if (submitted_password[i] != real_password[i])  
            return False;  
    }  
    /* Ensures both strings are same len. */  
    return submitted_password[i] == real_password[i];  
}
```

Attacking Password Checker

```
/* Tenex (old OS) system call to check if submitted password is correct. */
```

```
bool CheckPassword(char* submitted_password, char* user){  
    char* real_password = GetUserPassword(user);  
    for (int i = 0; submitted_password[i] && real_password[i]; ++i) {  
        if (submitted_password[i] != real_password[i])  
            return False;  
    }  
    /* Ensures both strings are same len. */  
    return submitted_password[i] == real_password[i];  
}
```

Attacking Password Checker

```
/* Tenex (old OS) system call to check if submitted password is correct. */  
  
bool CheckPassword(char* submitted_password, char* user){  
    char* real_password = GetUserPassword(user);  
    for (int i = 0; submitted_password[i] && real_password[i]; ++i) {  
        if (submitted_password[i] != real_password[i])  
            return False;  
    }  
    /* Ensures both strings are same len. */  
    return submitted_password[i] == real_password[i];  
}
```

Attacking Password Checker

```
/* Tenex (old OS) system call to check if submitted password is correct. */  
bool CheckP  
char* re  
for (int i  
    if (s  
    }  
/* Ensur  
return submitted_password[i] == real_password[i];  
}
```

Say passwords are only alphanumeric.

To brute force a 10-character password, requires guessing:

$62^{10} = 8.39 \cdot 10^{17}$ possible passwords.

Better “Side Channel” Attack

Leverage memory layout of the submitted password, by spreading it out across multiple pages.



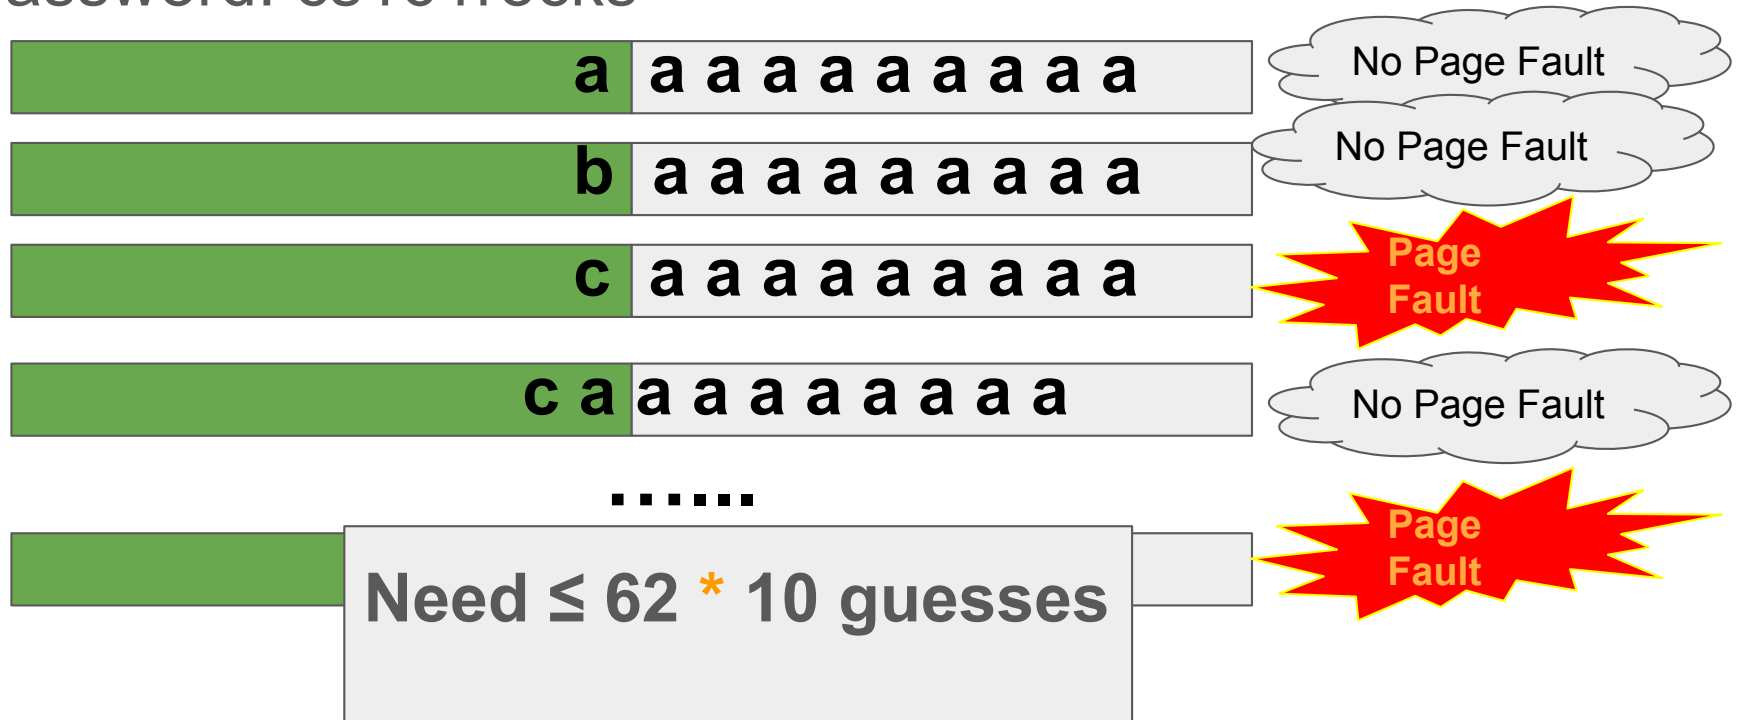
If password doesn't start with 'W', CheckPassword returns immediately (loop exits after 1 iteration).
Page-out (or unmap) this page

If password **DOES** start with 'W', CheckPassword looks for second character of submitted password, and **page faults!**

Better “Side Channel” Attack

Page faults are slow, *timing side channel!* (Seg faults also visible)

Real password: cs161rocks



Potential Fixes?

Fix 1: Always check entire password.

- Might still leak password length based on how long the check takes!

Fix 2: Assume a max length for password. Always loop that many times, even if password is shorter.

- Constant time algorithm: Eliminates timing side channel, but now caps password length and has worse performance.

Power Analysis on RSA

RSA decryption: $M = C^d \pmod N$

Common algorithm for exponentiation is “square and multiply”.

```
def exponentiate(base C, exponent d):  
    V = 1  
    For each bit b in d (most to least significant):  
        V = V^2 mod N  
        If b==1: V = V*C mod N  
    return V
```

Ex: $d=1010$ in binary = 10 in decimal.

Old $V = 1$

Power Analysis on RSA

RSA decryption: $M = C^d \bmod N$

Common algorithm for exponentiation is “square and multiply”.

```
def exponentiate(base C, exponent d):  
    V = 1  
    For each bit b in d (most to least significant):  
        V = V^2 mod N  
        If b==1: V = V*C mod N  
    return V
```

Ex: $d=1010$ in binary = 10 in decimal.

Old $V = 1$, New $V = 1^2 * C = C$ (bit is 1)

Power Analysis on RSA

RSA decryption: $M = C^d \bmod N$

Common algorithm for exponentiation is “square and multiply”.

```
def exponentiate(base C, exponent d):  
    V = 1  
    For each bit b in d (most to least significant):  
        V = V^2 mod N  
        If b==1: V = V*C mod N  
    return V
```

Ex: $d=1010$ in binary = 10 in decimal.

Old $V = C$, New $V = C^2$ (bit is 0)

Power Analysis on RSA

RSA decryption: $M = C^d \bmod N$

Common algorithm for exponentiation is “square and multiply”.

```
def exponentiate(base C, exponent d):  
    V = 1  
    For each bit b in d (most to least significant):  
        V = V^2 mod N  
        If b==1: V = V*C mod N  
    return V
```

Ex: $d=10\mathbf{1}0$ in binary = 10 in decimal.

Old $V = C^2$, New $V = C^{2*2} * C = C^5$ (bit is 1)

Power Analysis on RSA

RSA decryption: $M = C^d \bmod N$

Common algorithm for exponentiation is “square and multiply”.

```
def exponentiate(base C, exponent d):  
    V = 1  
    For each bit b in d (most to least significant):  
        V = V^2 mod N  
        If b==1: V = V*C mod N  
    return V
```

Ex: $d=1010$ in binary = 10 in decimal.


Old $V = C^5$, New $V = C^{10}$ (bit is 0), **as expected!**

Power Analysis on RSA

RSA decryption: $M = C^d \bmod N$


Common algorithm for exponentiation is “square and multiply”.

Square + multiply computation produces different power usage profile than just squaring! Can distinguish between a 0 or 1 bit in secret key based on power usage!



AES Cache Timing Attack

Did not cover
in lecture, just
for reference!



AES's computation accesses tables of values. Which indices are accessed is based on the secret key.

If these tables are stored in memory shared by the attacker and victim process (e.g. memory deduplication), attacker can load the tables into cache, *except* for one index. Later attacker can load that missing index.

- Cache miss = slower, victim didn't use the missing index
- Cache hit = faster, victim used the missing index

Can learn the secret key based on indices accessed.

Other side channels used for attacks

- Timing
- Cache hits
- Power usage
- Data remanence (“deleted” but uncleared memory)
- Row Hammer (change off-limit memory by accessing adjacent memory)
- Network side channels (recall global IP ID scanning from Homework 4)
- Electromagnetic radiation
- Acoustics
- Optical

Backpage and Bitcoin: Uncovering Human Traffickers

Bitcoin


- Bitcoin ownership is pseudonymous
 - Exchange bitcoin using pseudonyms
 - Pseudonyms are public keys, tied to private key the user owns
 - Sign out-going transactions with private key

Blockchain

- Public, distributed, peer-to-peer, hash-chained audit log of all transactions
 - Hash chain is public, broadcasted on peer-to-peer network, and append-only

Blockchain cont'd

- How do you get bitcoin?
 - Mining
 - Append block to most recent/longest version of blockchain
 - Buy it



Sex Trafficking and the Internet

- Internet has opened new ways for traffickers to advertise and find victims
- Broader goal : use computer science tools/techniques to fight sex trafficking and slavery
- Detect traffickers from advertisements they pay for and post

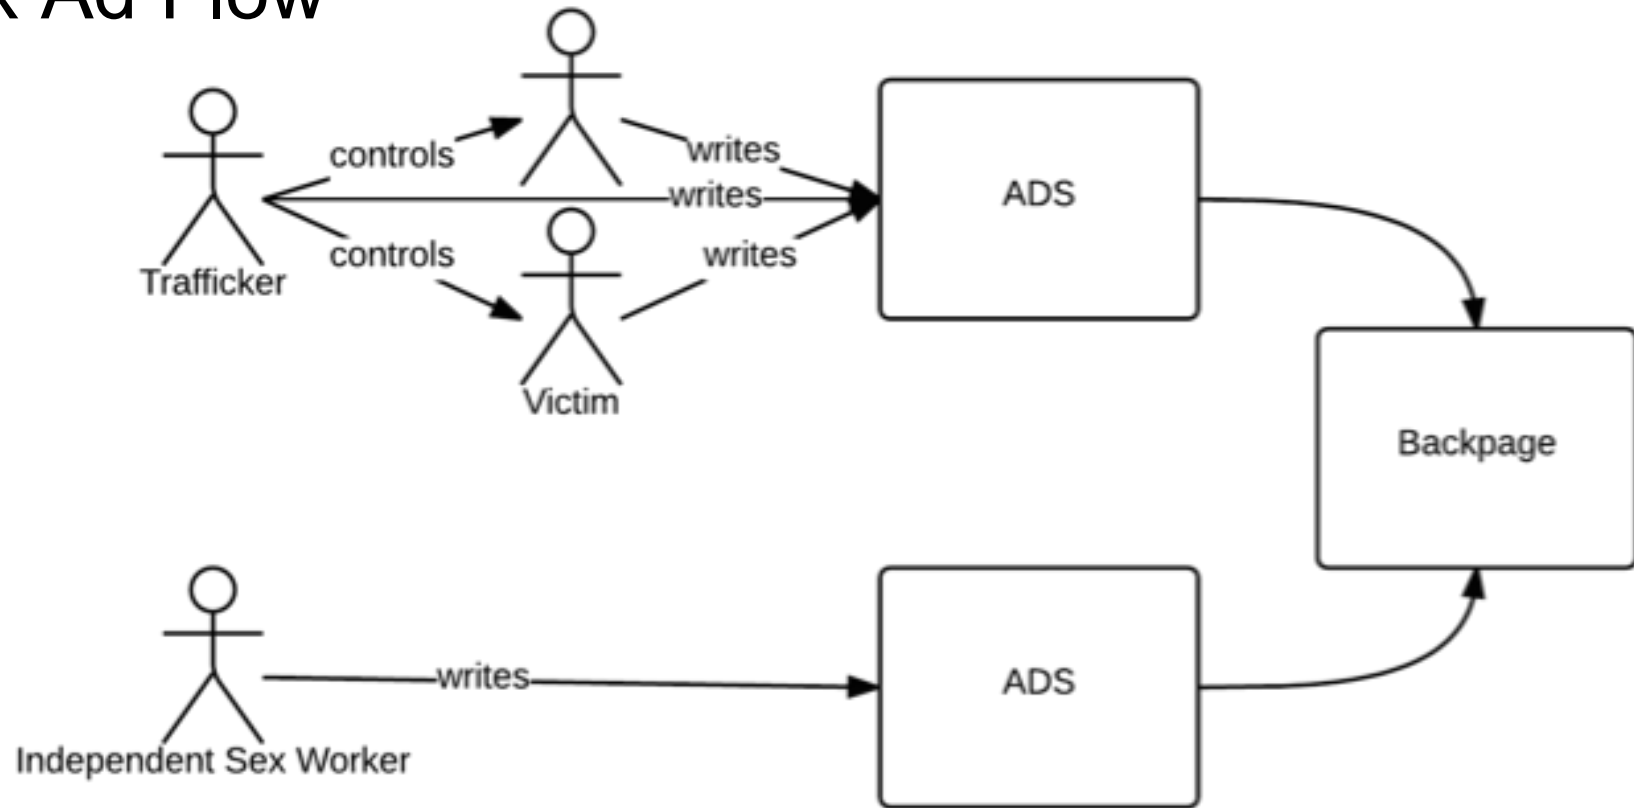
Problem Statement

- Can I distinguish traffickers from independent sex workers on classified ad sites?
 - Too much data

Backpage

- 2nd largest online classified ad site in the US
- 80% percent of the market for online sex ads in USA
- Running since 2004, listings all over the world
- Used by traffickers to advertise their victims
- Two forms of payment for adult entertainment listings
 - Bitcoin
 - Check/money order sent via regular mail

Sex Ad Flow




Goal

- Develop techniques to cluster sex ads by owner
 - Current best clustering is via hard-link; unreliable

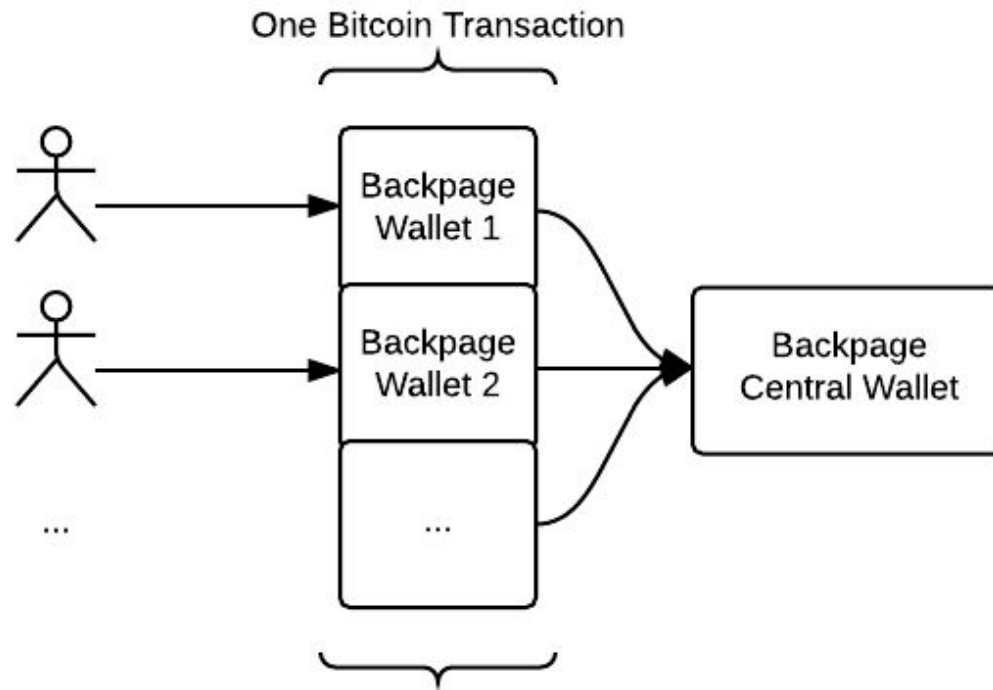
Results

- Two different methodologies that combine the classifier, linking technique and existing hard identifiers to group ads by owner
 - Stylometry classifier that distinguishes between sex ads posted by the same vs. different authors with 90% TPR and 1% FPR
 - Side channel attack that takes advantage of leakages from the Bitcoin blockchain and sex ad site to link a subset of sex ads to Bitcoin public wallets and transactions
- Analyzed 4-weeks of scraped sex ads from Backpage
 - Rebuild the price of each Backpage sex ad, and analyze the output of the two different methodologies

Stylometry model

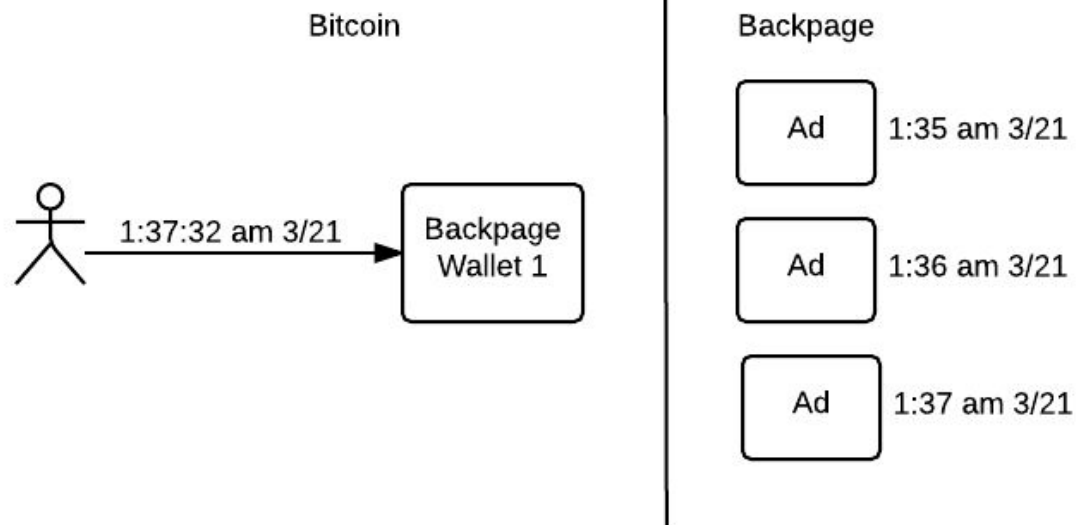


Backpage Payment Flow



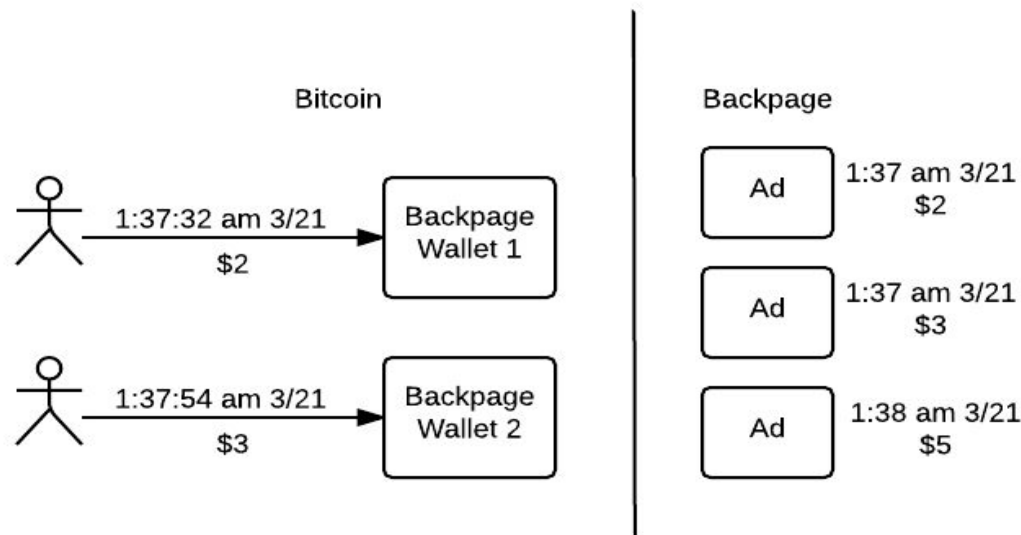
Timing Side Channel Attack

- Backpage posts ad onto its site one minute after payment appears on Bitcoin mempool




Timing & Price Side Channel Attack

- Backpage's pricing algorithm takes ad posting frequency and location as variables, and can be reverse engineered



Persistent Bitcoin Identity Methodology

- Goal: map each ad to its true owner wallet
- Persistent Bitcoin Identity: any wallet that sends the change from each of its transactions back into itself, and has one exact match
 - Use stylometry model to distinguish non-exact match
 - All ads that match to this wallet are clustered under this PBI



4-week case study

- 26 'ground truth' test ads
 - 25 required payment, 1 free
 - Placed from Dec 12th, 2016 to Dec 24th, 2016
 - Price range from \$2 to \$20
 - Posted in 27 distinct US regions
- Scraped all the sex ads in every US location every hour, for 4 weeks
- 741,443 unique ads scraped
 - 151,482 required payment
 - Placed from Dec 10th, 2016 to Jan 9th, 2017
 - Price range from \$1 to >\$100
 - Posted in 60 distinct US regions

4-week case study: PBI

- 11 ground truth ads paid using a PBI
 - 8 transactions were exact match for correct ad
 - 3 transactions matched two ads, one of which was the correct ad
- 249 PBI's total
- 90 of those PBI's had at least one exact match
- Results:
 - Links between hard identifiers
 - Evidence of networks across multiple locations
 - Owners of sex ad clusters spending a lot of money on ads

Conclusion

- Promising!
- First work to try to link specific purchases to specific transactions on the Blockchain
- Lots of work left to be done

User Authentication & Passwords

CS 161: Computer Security

Guest Lecturer: Grant Ho

Instructor: Prof. Vern Paxson

TAs: Paul Bramsen, Apoorva Dornadula, David Fifield, Mia Gil Epner, David Hahn, Warren He, Grant Ho, Frank Li, Nathan Malkin, Mitar Milutinovic, Rishabh Poddar, Rebecca Portnoff, Nate Wang

<http://inst.eecs.berkeley.edu/~cs161/>

April 27, 2017

With content from Raluca Ada Popa & Dan Boneh

Attacks & Defenses on Password Authentication

- Often worry about 3 classes of attacks (threat models)
 1. Online guessing
 2. Server compromise (“offline guessing”)
 3. Client password compromise

Attacks & Defenses on Password Authentication

- Often worry about 3 classes of attacks (threat models)
 1. Online guessing
 2. Server compromise (“offline guessing”)
 3. Client password compromise
- We’ll just focus on the last two threat models b/c of time constraints

Threat Model 1: Server Compromise

Attacker breaks into server and steals password database
(also called “offline guessing attacks”)

Threat Model #1: Server Compromised

- Attacker breaks into server and steals password database
- Happens all the time ☹️



Threat Model #1: Server Compromised

- Insecure Defense: Server stores encrypted passwords in its database
- But server needs easy access to secret key in order to verify users when they login
 - So, if Mallory breaks into the server, then she can just steal secret key too!

Encrypting passwords is ***not*** a secure solution

Secure Password Storage

- Server should store *salted + hashed* passwords (Section 6, Problem #1)

- **Setup**

1. During account registration, server generates random number (salt)
2. Server computes $h = \text{hash}(\text{salt}, \text{password})$
3. Server stores (username, salt, h) and *deletes user's password*

username	salt	$h = \text{hash}(\text{salt}, \text{password})$
Alice	235545235	Hash(Alice's pwd, 235545235)
Bob	678632523	Hash(Alice's pwd, 678632523)

- **Authentication**

- User's browser sends {username, password} to server
- Server computes $\text{hash}(\text{salt}, \text{password})$ and checks if it matches h

Secure Password Storage

- **Secure Defense:** Server should store *salted + hashed* passwords

username	salt	h = hash(salt, password)
Alice	235545235	Hash(Alice's pwd, 235545235)
Bob	678632523	Hash(Alice's pwd, 678632523)

- Attacker steals password database, but:
 - Only sees salts & h 's: salt is random & secure hash functions are **one-way**.
 - Attacker can still compute big table of guesses for Alice & check for matching h :



'123456'	Hash('123456', 235545235)
'password'	Hash('password', 235545235)
'aaaaaaaa'	Hash('aaaaaaaa', 235545235)
...	...

But **salting** forces attacker to re-compute table for each user and prevents pre-computation.

Secure Password Storage

- **Secure Defense:** Server should store *salted + securely hash* passwords
- The secure hash function should also be *slow* to compute
 - Usually we want fast crypto for performance
 - But here we want attacker to wait... and wait... and wait... for guessing to succeed.
 - Examples: Argon, bcrypt, scrypt
- Conceptually, $\text{Slow-Hash}(x) = \text{hash}(\text{hash}(\text{hash}(\text{hash}(\dots(\text{hash}(x))))))$
 - where hash is a regular secure hash (e.g., SHA-256 or HMAC)
- If Slow-Hash is 1,000 times slower, attack that previously took 1 day now takes ~3 years

Threat Model 2: Client Password Compromise

Attacker obtains Alice's password

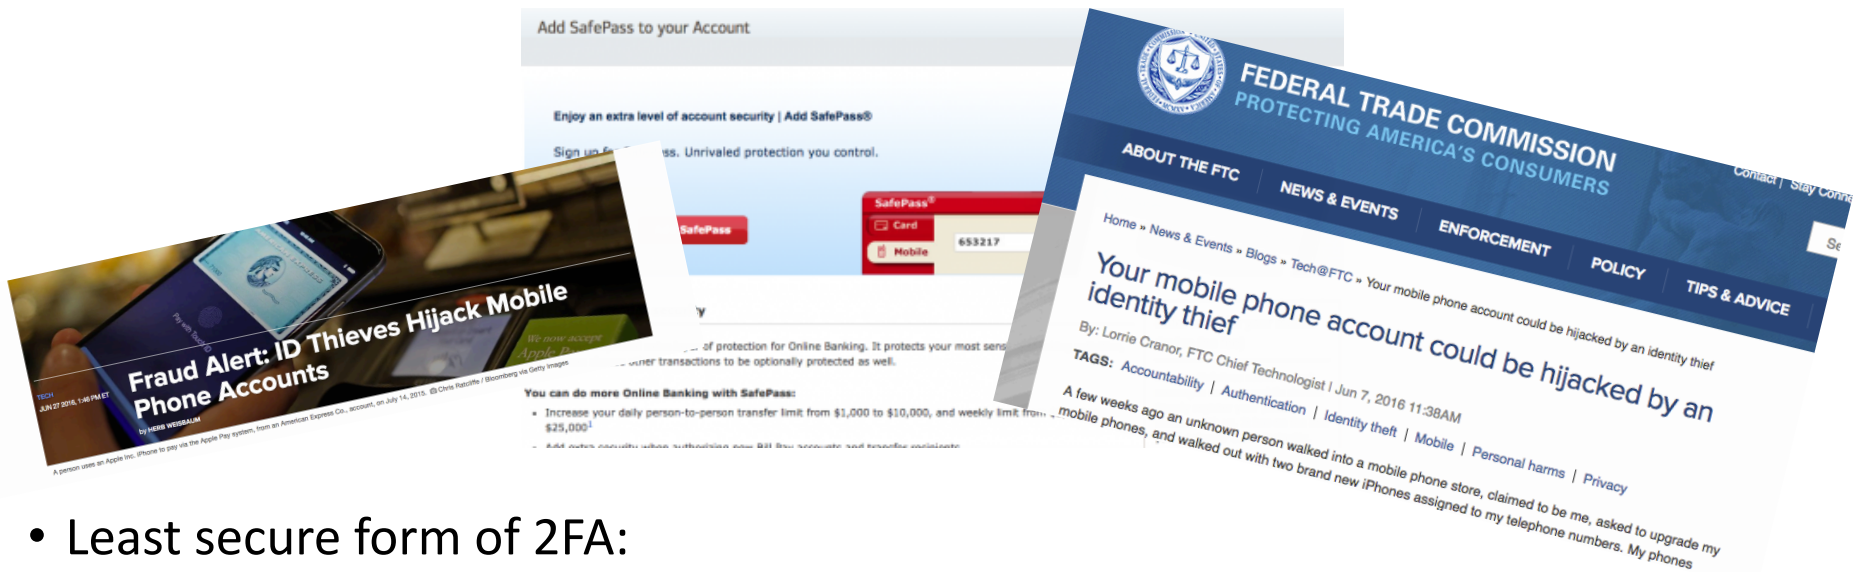
- Phishing
- Surveillance camera (airport, cafe, etc.) records Alice typing password
- ...

Threat Model #2: Password Compromise

- Defense: Two-factor authentication (2FA)
 1. Something you know (password)
 2. Something you have (smartphone/authentication device)
 3. Something you are (fingerprints/iris scanner)
 - Require 2 methods from above
- Most common 2FA: password + authentication device
 - User enters password at login
 - If password correct, user then needs to use authentication device
- Let's examine some 2FA designs for the authentication device

Common 2FA Designs

1. Text message: server generates random number & texts it to you



- Least secure form of 2FA:
 - Hijack phone number
 - Mobile malware or any app w/ text message permissions (e.g. Tinder, Uber, etc.)
 - ...



Common 2FA Designs

1. Text message: server generates random number & texts it to you

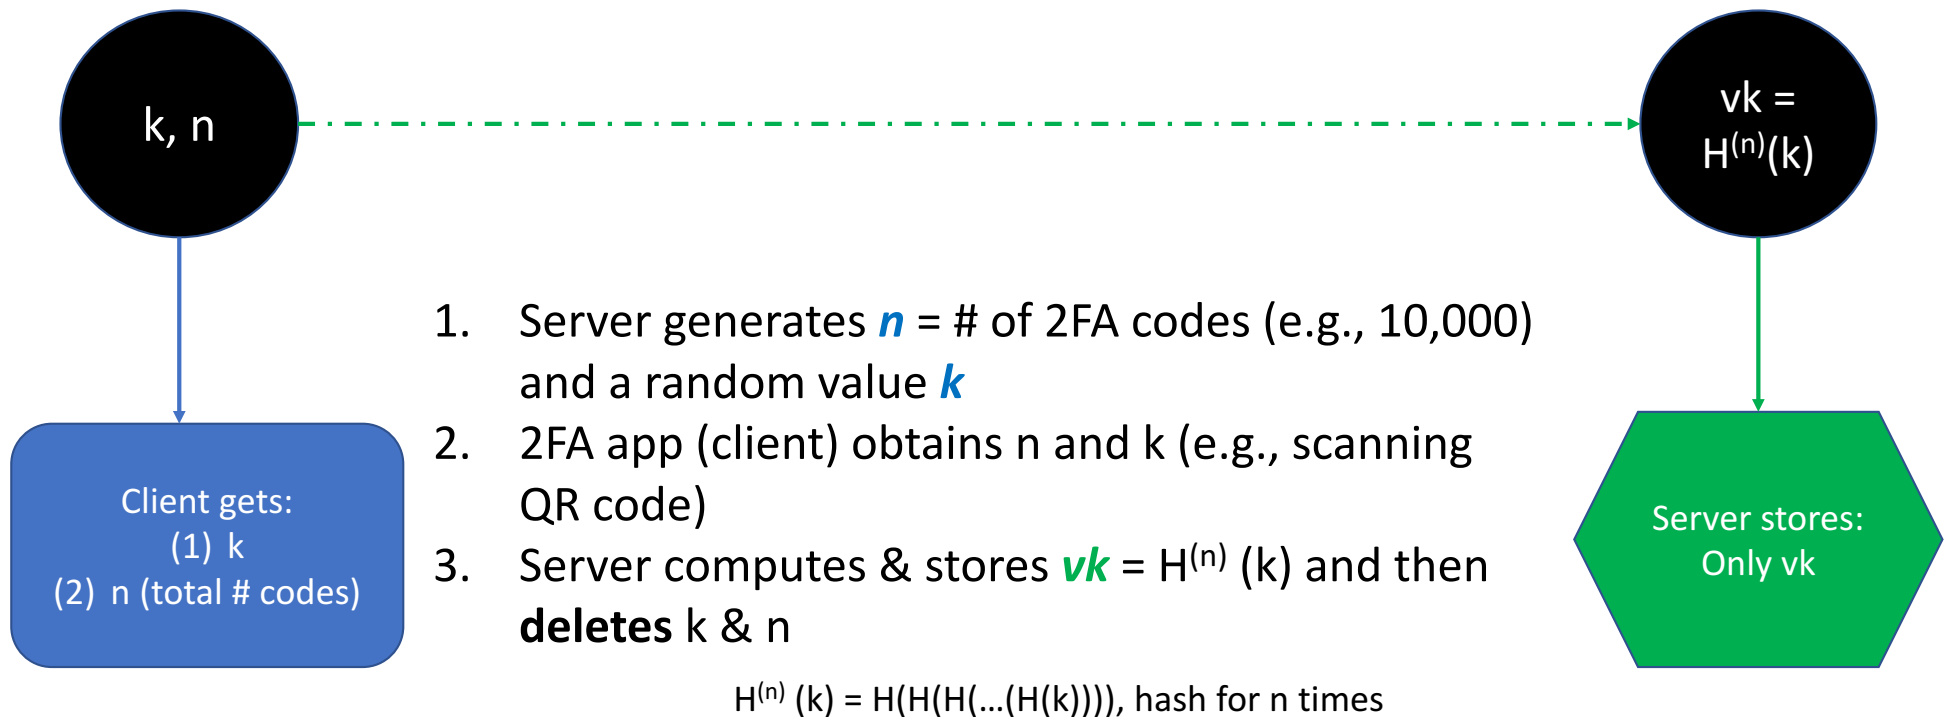
- Least secure: hijack phone number or hack telephone company (e.g., nation state)

2. Authenticator apps (more secure)


- Google Authenticator, Duo, etc.
- Protocols:
 - S/KEY
 - TOTP
 - Push notification



The S/Key Protocol: Setup



The S/Key Protocol: Authenticating to Server



Client stores:
(1) k
(2) n (total # codes)
(3) j (# total logins)

1. Client computes & sends $sk = H^{(n-j)}(k)$
2. Server checks if $H(sk) = vk$
3. Server updates $vk = sk$
4. Repeat 1-3

Secure even if attacker breaks into server and steals vk for each user!

Server stores:
Only vk

Common 2FA Designs

1. Text message: server generates random number & texts it to you

- Least secure: hijack phone number or hack telephone company (e.g., nation state)

2. Authenticator apps (more secure)

- Common protocols: S/KEY, TOTP, Push notification
- **Still vulnerable to phishing!**
 1. Phishing page asks for user's password
 2. Next, phishing page asks user to enter 2FA code
 3. Attacker then uses both to login

Common 2FA Designs

1. Text message: server generates random number & texts it to you

- Least secure: hijack phone number or hack telephone company (e.g., nation state)

2. Authenticator app (more secure)

- Push notification, S/KEY, TOTP
- Still vulnerable to phishing!

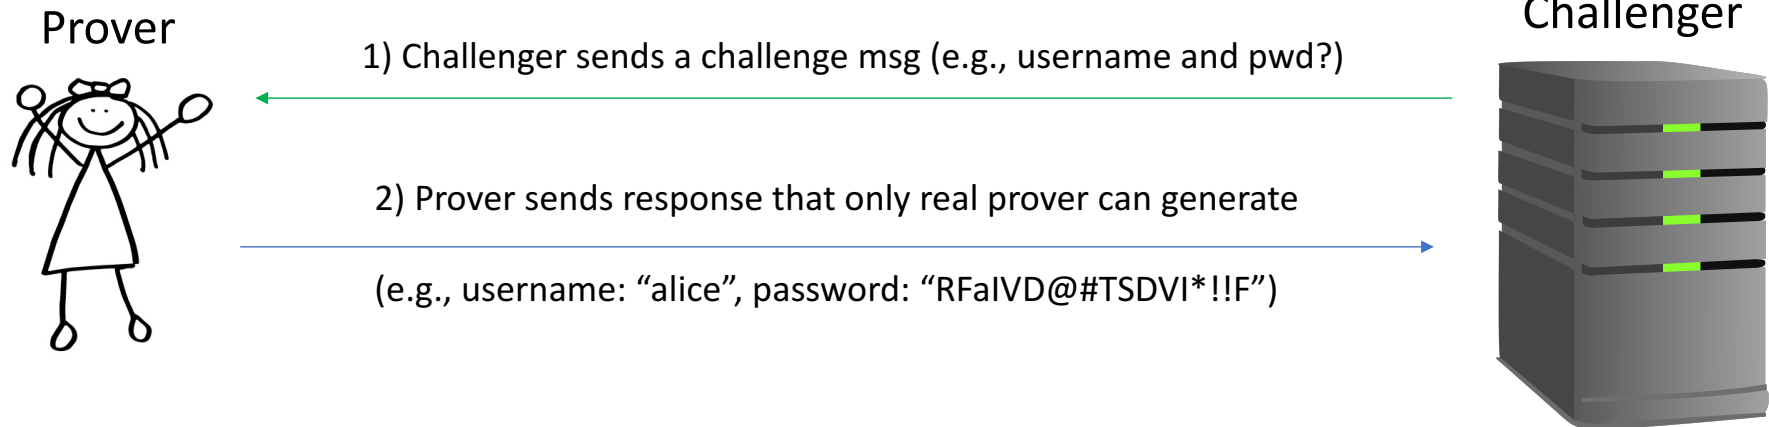
3. Hardware tokens: challenge-response (most secure)

- Hardware device that's plugged into laptop
- Can protect against phishing attacks



Challenge-Response (General)

- General protocols for authentication
- A “**prover**” wants to authenticate to a “**challenger**”
 - E.g., a user (prover) wants to login to Gmail (challenger) as Alice



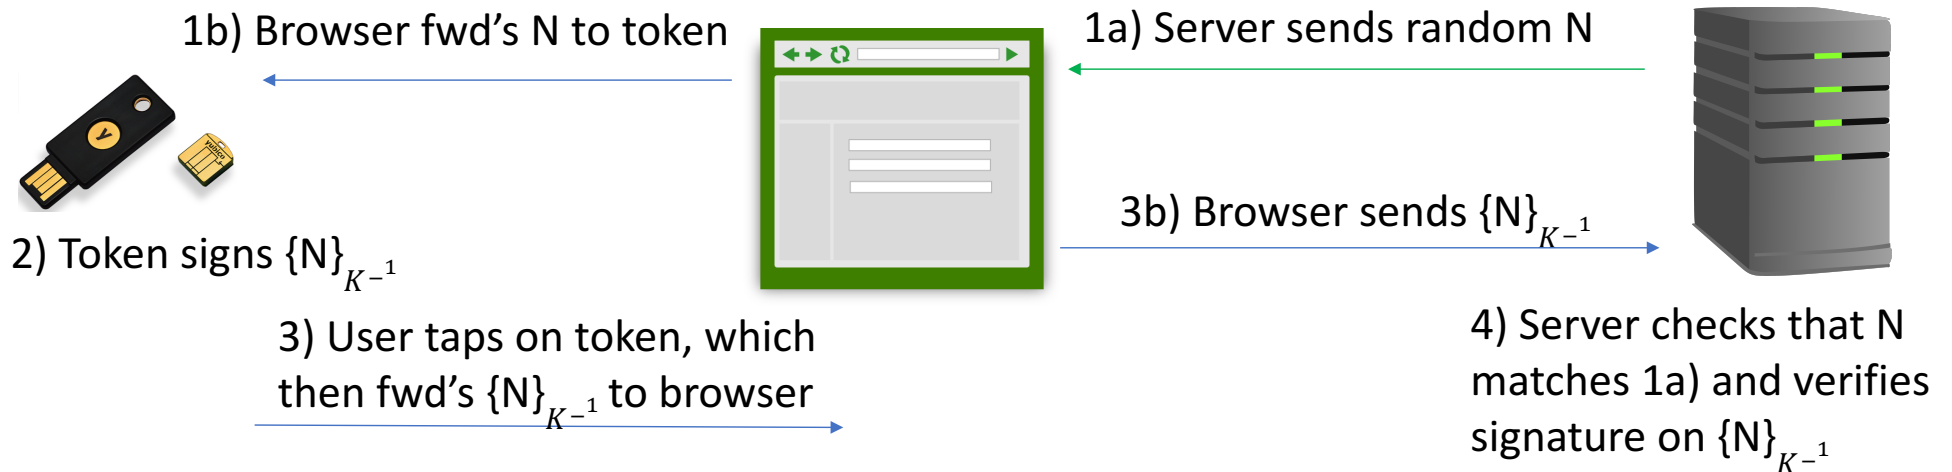
2FA Challenge-Response

- Hardware 2FA token has a public & private key pair embedded in device

A. Setup

1. Alice's browser gets K = 2FA token's public key and sends K to server
2. Server stores (username, K) in its 2FA database

B. Authentication



Adding Phishing Resistance

1b) Browser fwd's N to token

AND it includes **D = domain of actual webpage in browser**




2) Token signs $\{N, D\}_{K^{-1}}$

3) User taps on token, which then fwd's $\{N, D\}_{K^{-1}}$ to browser



1a) Server sends random number N

3b) Browser sends $\{N, D\}_{K^{-1}}$



4) Server checks:

- **D** matches its domain
- N matches what it sent
- Valid signature on $\{N, D\}_{K^{-1}}$


Phishing Attack Now Fails!

- During phishing attack, browser will be at website w/ domain $D' = \text{gmai1.com}$, instead of real domain $D = \text{gmail.com}$



Phishing Attack Now Fails!

- During phishing attack, browser will be at website w/ domain $D' = \text{gmail1.com}$, instead of real domain $D = \text{gmail.com}$



Practical Advice for Future Security Engineers

Applicable to your users and your employees:

1. Use HTTPS (prevent MITM from seeing passwords)
2. Securely store passwords (Threat Model #1)
3. Enable 2FA, ideally hardware tokens (Threat Model #2)
4. Securely check passwords & rate limit (not covered b/c of time)
5. Incorporate detection systems if you can (not covered b/c of time)
 1. Access logging
 2. Spearphishing detection
 3. Honey accounts/Tripwire

Computing on private data

CS 161: Computer Security

Instructor: Prof. Vern Paxson

TAs: Paul Bramsen, Apoorva Dornadula, David Fifield, Mia Gil Epner, David Hahn, Warren He, Grant Ho, Frank Li, Nathan Malkin, Mitar Milutinovic, Rishabh Poddar, Rebecca Portnoff, Nate Wang

<http://inst.eecs.berkeley.edu/~cs161/>

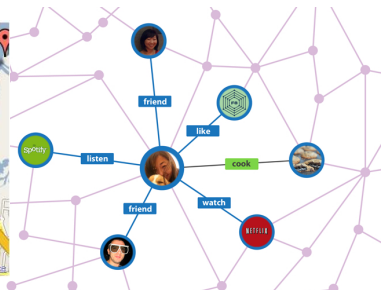

April 27, 2017

Guest Lecturer: Rishabh Poddar

*With content from
Raluca Ada Popa, Dan Boneh, and Taesoo Kim*

Many decisions are made on private data

- User data (e.g. email, social)
- Medical data
- Financial data
- Location data



Data stored unencrypted in order to allow applications to compute queries / make decisions

Defense: try to build **walls** around the data (e.g. access control, firewalls, IDS, etc.)

Attackers eventually break into systems

EVERYDAY MONEY IDENTITY THEFT

Data Breach Tracker: All the Major Companies That Have Been Hacked

Anthem Suffers the Largest Healthcare Data Breach to Date

by Fred Pennic 02/05/2015 2 Comments

[Twitter](#) [LinkedIn](#) [Facebook](#) [Email](#)

A photograph of a modern, multi-story brick and glass building, likely an Anthem headquarters or regional office, with a clear blue sky and some trees in the foreground.

- Sometimes, they even obtain root access or have admin privilege

WordPress firm Automattic suffers root-level hack

Summary: Hackers gained administrative privileges to a number of Automattic servers, WordPress founder Matt Mullenweg has said

By Tom Espiner | April 14, 2011 -- 13:31 GMT (06:31 PDT)

[Follow @tomespiner](#)



How can we prevent attackers from obtaining the data **even if they gain access to the system?**

The data should be encrypted at all times!

- Not just sometimes when the data is *at rest* (i.e. when no computations are being performed)

Problem: How does the server carry out its service (i.e. **perform computations on the data**) if the data is encrypted?

Two main approaches

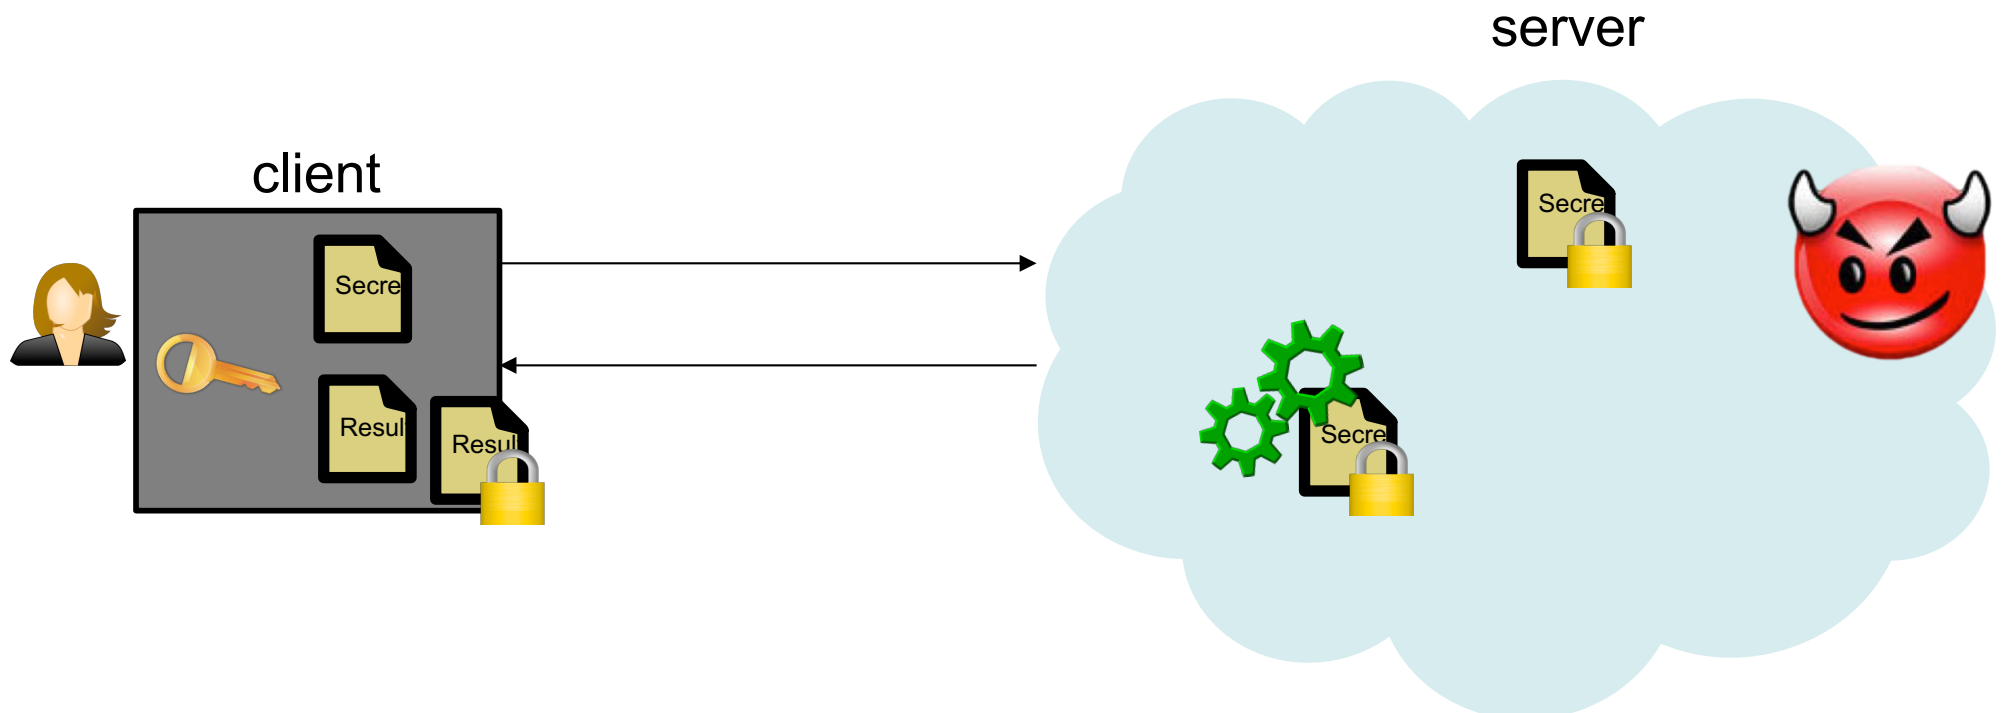
1. Compute ***directly on encrypted data*** (uses specialized cryptography)
2. ***Shielded computation*** on data (uses specialized hardware)

Which approach to use?

- **Security**: confidentiality / integrity guarantees
- **Functionality**: what computations can be supported
- **Performance**: how efficient is it to compute

Approach #1: Computation on encrypted data

Computation on encrypted data



Server performs computations on the encrypted data without ever decrypting it

Computation on encrypted data

- Option #1: **Property preserving** encryption

Computation on encrypted data

- Option #1: **Property preserving** encryption
 - **Deterministic** encryption:
If $x = y$ then $\text{Enc}(x) = \text{Enc}(y)$

Computation on encrypted data

- Option #1: **Property preserving** encryption
 - **Deterministic** encryption:
If $x = y$ then $\text{Enc}(x) = \text{Enc}(y)$

Can compute queries such as:

```
SELECT * FROM table WHERE name = 'Alice'
```

Computation on encrypted data

- Option #1: **Property preserving** encryption
 - **Deterministic** encryption:
If $x = y$ then $\text{Enc}(x) = \text{Enc}(y)$

Can compute queries such as:

```
SELECT * FROM table WHERE name = 0xfadc...
```

Computation on encrypted data

- Option #1: **Property preserving** encryption
 - Deterministic encryption:
If $x = y$ then $\text{Enc}(x) = \text{Enc}(y)$

Can compute queries such as:

```
SELECT * FROM table WHERE name = 0xfadc...
```

- **Order preserving** encryption:
If $x > y$ then $\text{Enc}(x) > \text{Enc}(y)$

Computation on encrypted data

- Option #1: **Property preserving** encryption
 - Deterministic encryption:
If $x = y$ then $\text{Enc}(x) = \text{Enc}(y)$

Can compute queries such as:

```
SELECT * FROM table WHERE name = 0xfadc...
```

- **Order preserving** encryption:
If $x > y$ then $\text{Enc}(x) > \text{Enc}(y)$

Can compute queries such as:

```
SELECT * FROM table WHERE age > 10
```

Computation on encrypted data

- Option #1: **Property preserving** encryption
 - Deterministic encryption:
If $x = y$ then $\text{Enc}(x) = \text{Enc}(y)$

Can compute queries such as:

```
SELECT * FROM table WHERE name = 0xfadc...
```

- **Order preserving** encryption:
If $x > y$ then $\text{Enc}(x) > \text{Enc}(y)$

Can compute queries such as:

```
SELECT * FROM table WHERE age > 0x1d3e...
```

Computation on encrypted data

- Option #1: **Property preserving** encryption

Performance

- Nearly as fast as computing on plaintext

Security

- Leaks some information about the plaintexts (e.g. frequency distribution of values, or order of ciphertexts)

Functionality

- Very limited: e.g. only equality for deterministic encryption, range comparison for order preserving encryption

Computation on encrypted data

- Option #2: **Partially homomorphic encryption**

Computation on encrypted data

- Option #2: **Partially homomorphic** encryption
 - **ElGamal** cryptosystem (enables multiplication over ciphertexts)

$$\text{Enc}(x) \cdot \text{Enc}(y) = \text{Enc}(x \cdot y)$$

Computation on encrypted data

- Option #2: **Partially homomorphic encryption**

- **ElGamal** cryptosystem (enables multiplication over ciphertexts)

$$\text{Enc}(x) \cdot \text{Enc}(y) = \text{Enc}(x \cdot y)$$

- **Paillier** cryptosystem (enables addition over ciphertexts)

$$\text{Enc}(x) + \text{Enc}(y) = \text{Enc}(x + y)$$

Computation on encrypted data

- Option #2: **Partially homomorphic encryption**

Performance

- Reasonably efficient, but not as fast as computing on plaintext

Security

- Similar level of confidentiality guarantees as standard AES-based encryption

Functionality

- Very limited: only specific operations can be computed (i.e. can only add, or can only multiply; can't do both)

Computation on encrypted data

- Option #3: **Fully homomorphic** encryption
 - Enables *arbitrary* functions

$$F(\text{Enc}(x), \text{Enc}(y)) = \text{Enc}(F(x, y))$$

Computation on encrypted data

- Option #3: **Fully homomorphic** encryption
 - Enables *arbitrary* functions

$$F(\text{Enc}(x), \text{Enc}(y)) = \text{Enc}(F(x, y))$$

Performance

- **Prohibitively slow** (currently 6 orders of magnitude slower)

Security

- Similar level of confidentiality guarantees as standard AES-based encryption

Functionality

- Allows arbitrary computations

Approach #2: Shielded
computation on data using
Intel Software Guard
Extensions (SGX)

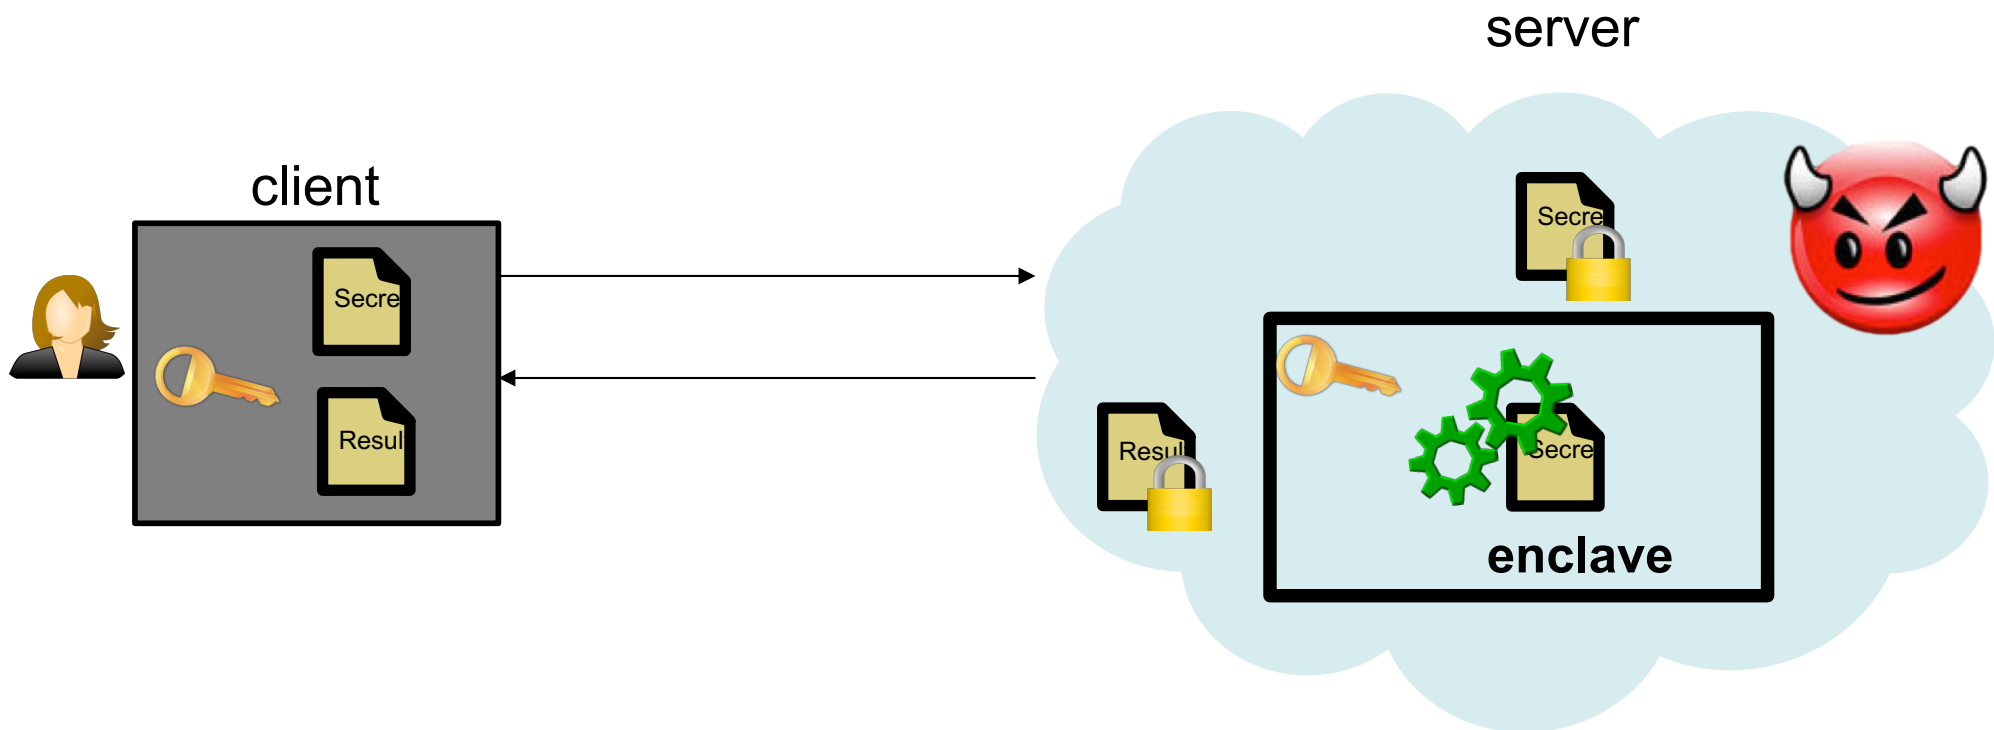
(Extensions to Intel processors)

Intel SGX

- **Feature #1:** Can run code in hardware-protected *containers* (called **enclaves**)

Intel SGX

- **Feature #1:** Can run code in hardware-protected *containers* (called **enclaves**)



Intel SGX

- **Feature #1:** Can run code in hardware-protected *containers* (called **enclaves**)
 - **Secure region of address space**, protected by the processor from all external software access (**even from the operating system**)

Intel SGX

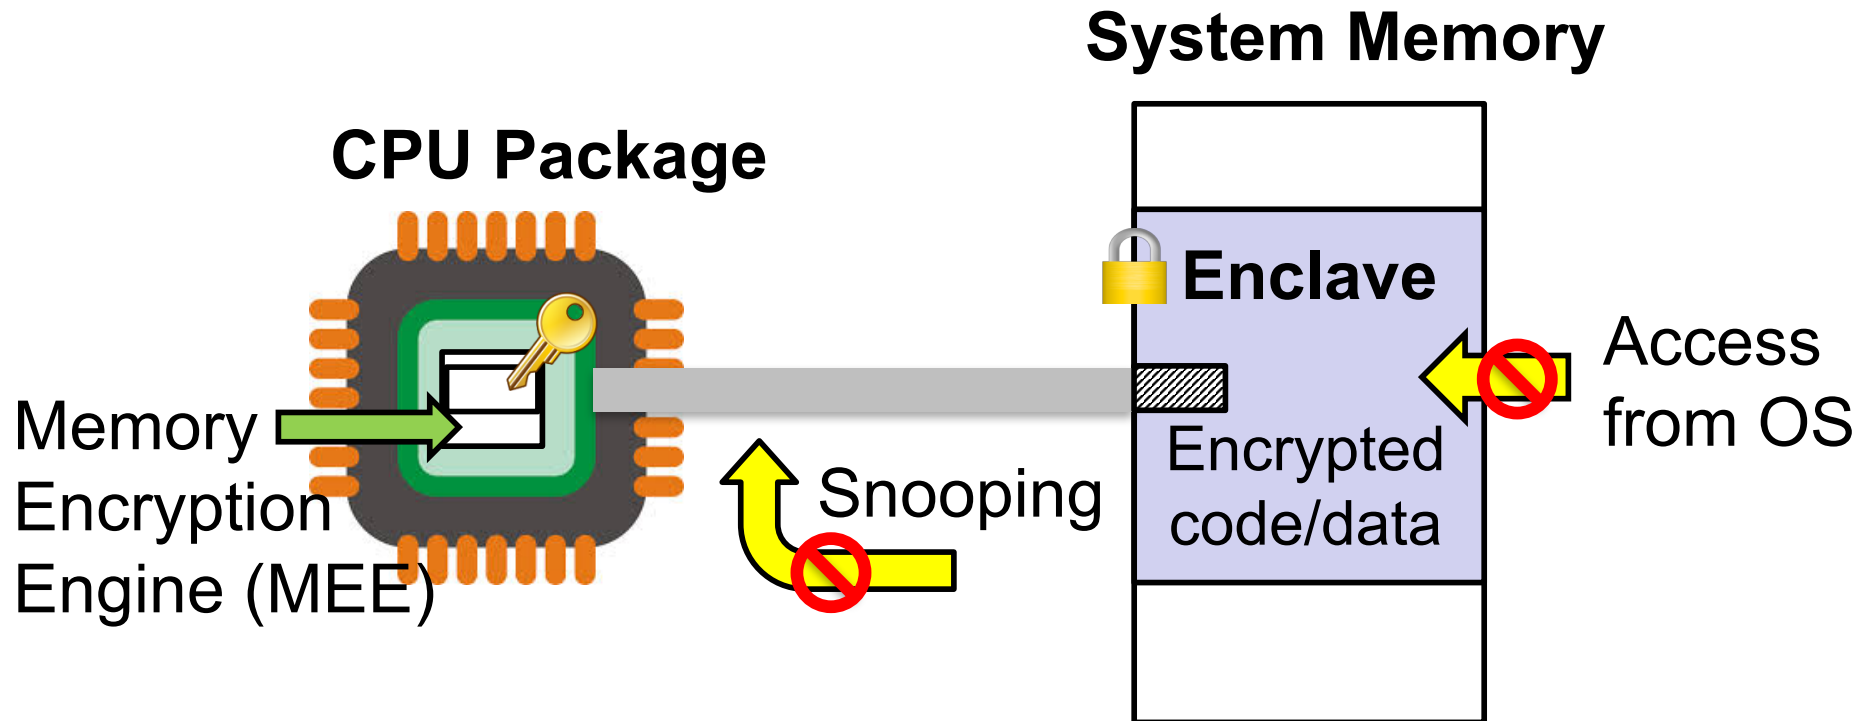
- **Feature #1:** Can run code in hardware-protected *containers* (called **enclaves**)
 - Secure region of address space, protected by the processor from all external software access (even from the operating system)
 - Code and data in enclave region of main memory **always encrypted** using **processor specific keys**
 - Decrypted only within the CPU package (i.e. when loaded into registers / cache)

Intel SGX

- **Feature #1:** Can run code in hardware-protected *containers* (called **enclaves**)
 - Secure region of address space, protected by the processor from all external software access (even from the operating system)
 - Code and data in enclave region of main memory **always encrypted** using **processor specific keys**
 - Decrypted only within the CPU package (i.e. when loaded into registers / cache)

Code and data loaded into an enclave is ***isolated*** from the rest of the system

SGX: How enclaves work



Problem: How to verify correct code has been loaded?

- Enclave code allowed to access unencrypted data
- Malicious / tampered code in enclave could *exfiltrate* data (i.e. leak it to the attacker)

Intel SGX

Extensions to Intel processors that support:

- Feature #2: **Attestation**

Intel SGX

- Feature #2: **Attestation**
 - Prove to local / remote system that the **correct** code has been loaded into the enclave (i.e. **verify the integrity of the enclave** using a **hash measurement** of the loaded code/data)

Intel SGX

- Feature #2: **Attestation**

- Prove to local / remote system that the *correct* code has been loaded into the enclave (i.e. **verify the integrity of the enclave** using a **hash measurement** of the loaded code/data)
- Verify that **measurement was generated by an enclave** running on the same platform (using a **MAC**)

Intel SGX

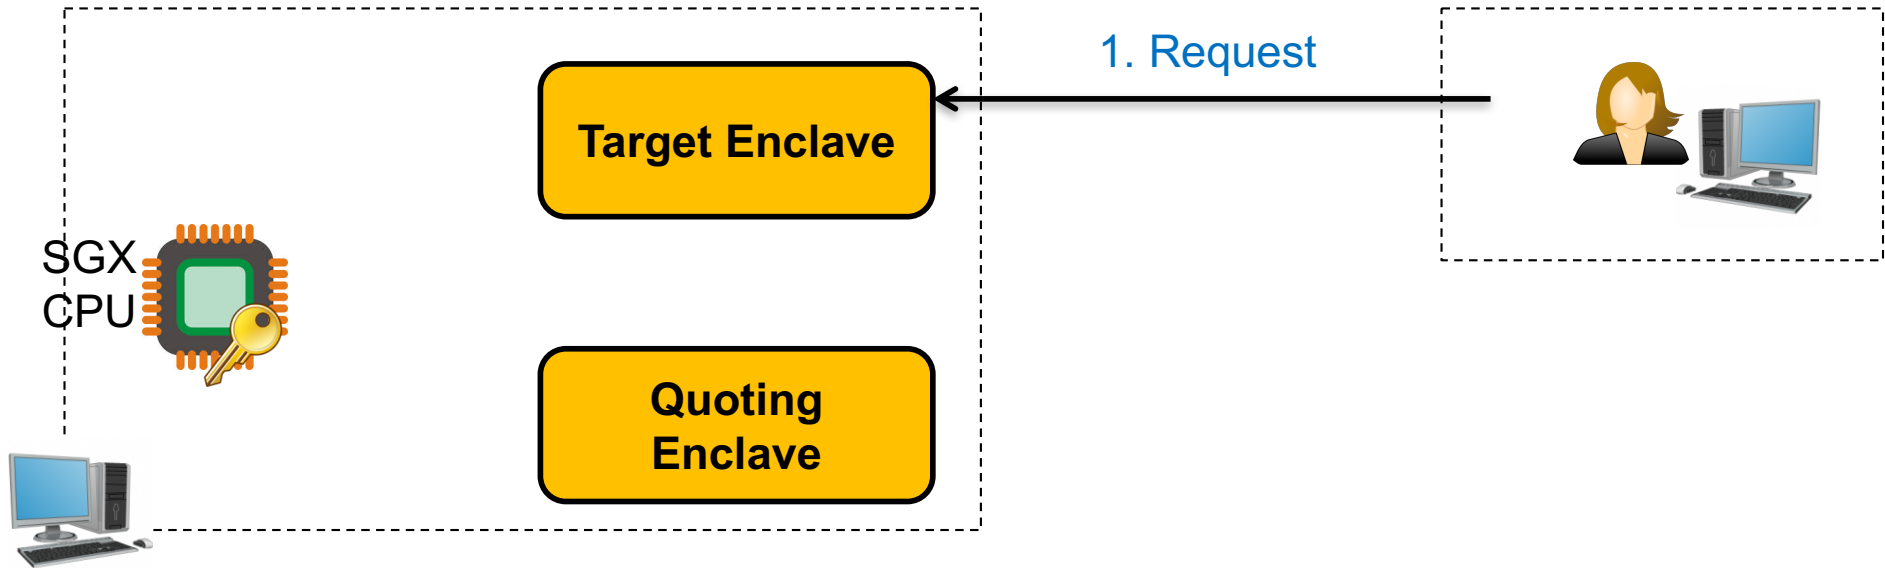
- Feature #2: **Attestation**

- Prove to local / remote system that the *correct* code has been loaded into the enclave (i.e. **verify the integrity of the enclave** using a **hash measurement** of the loaded code/data)
- Verify that **measurement generated by an enclave** running on the same platform (using a **MAC**)
- Uses a special **quoting enclave** for this purpose that signs the measurement and sends it to the client for verification

SGX: How attestation works

Server

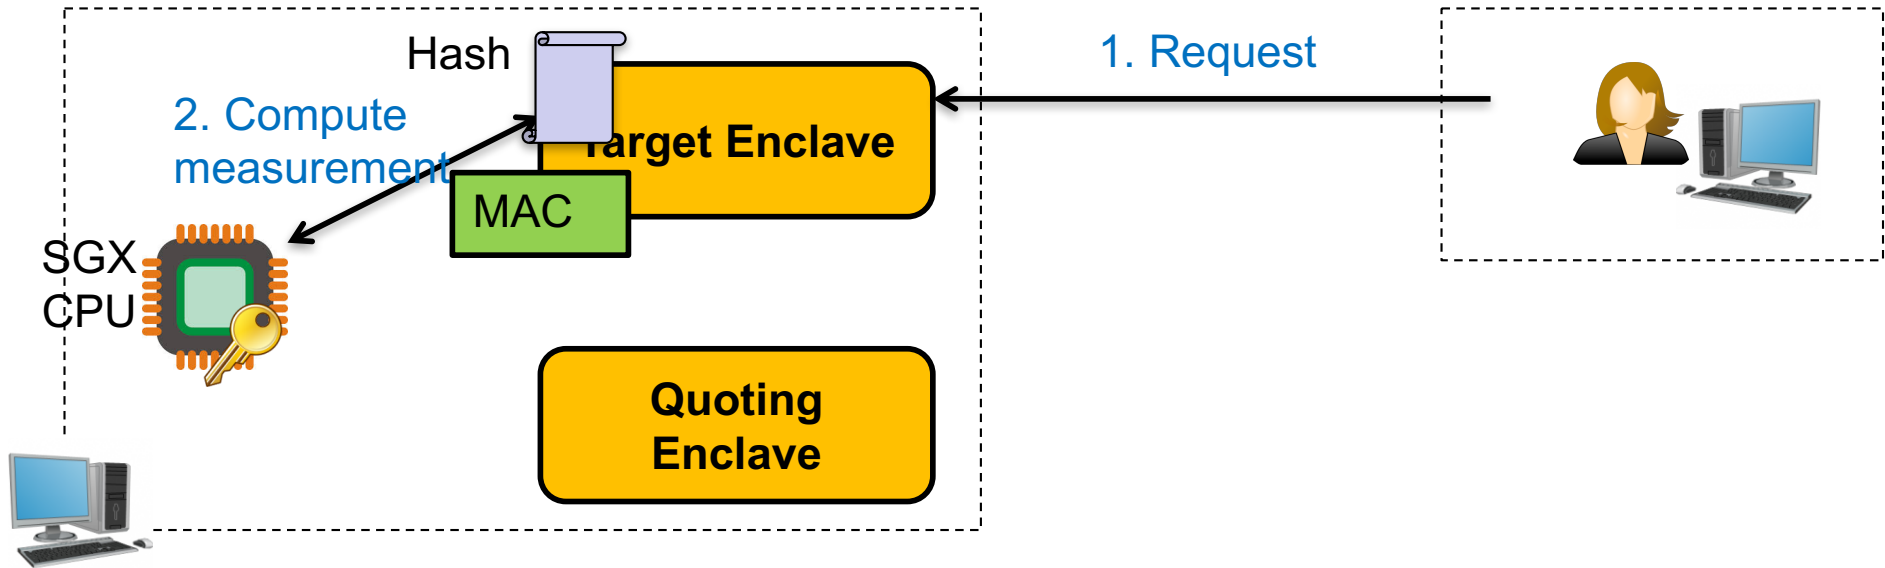
Client



SGX: How attestation works

Server

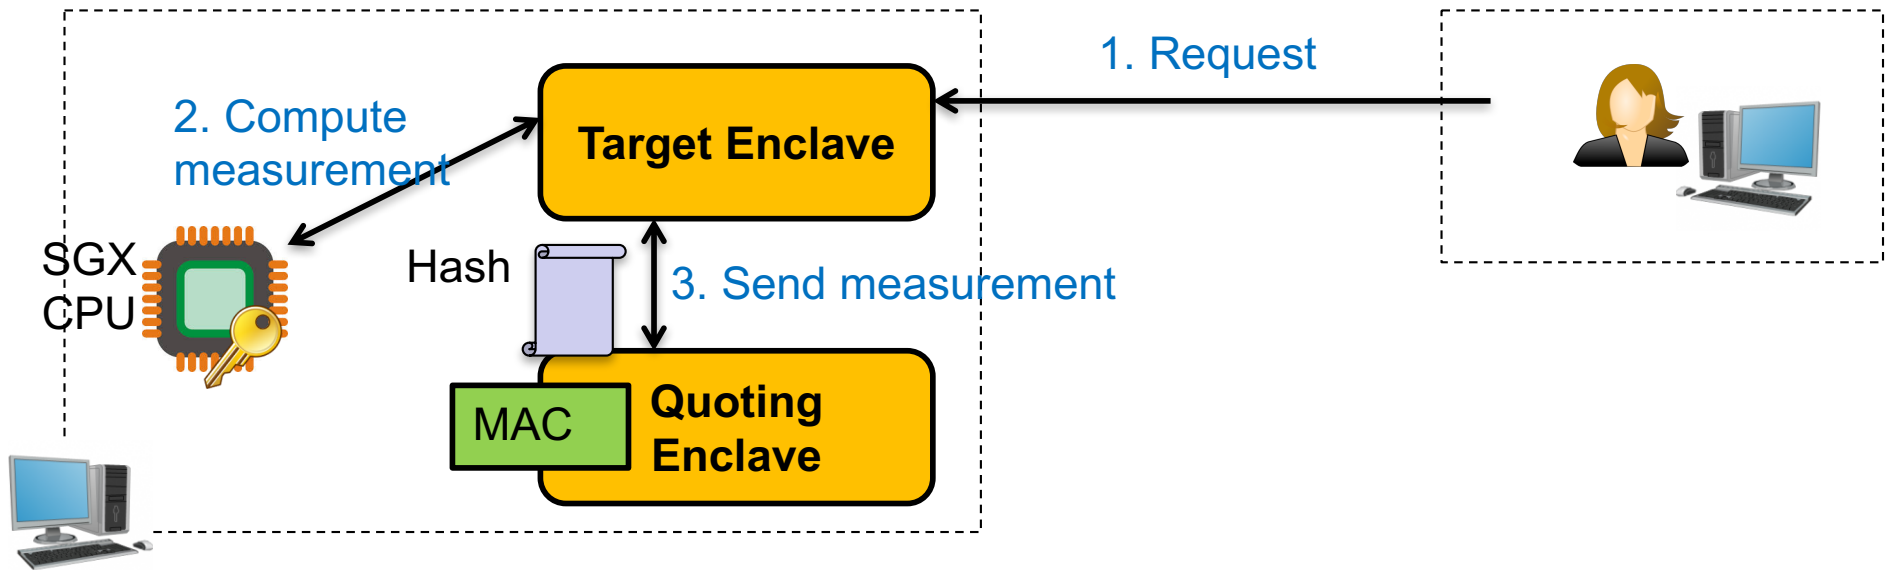
Client



SGX: How attestation works

Server

Client

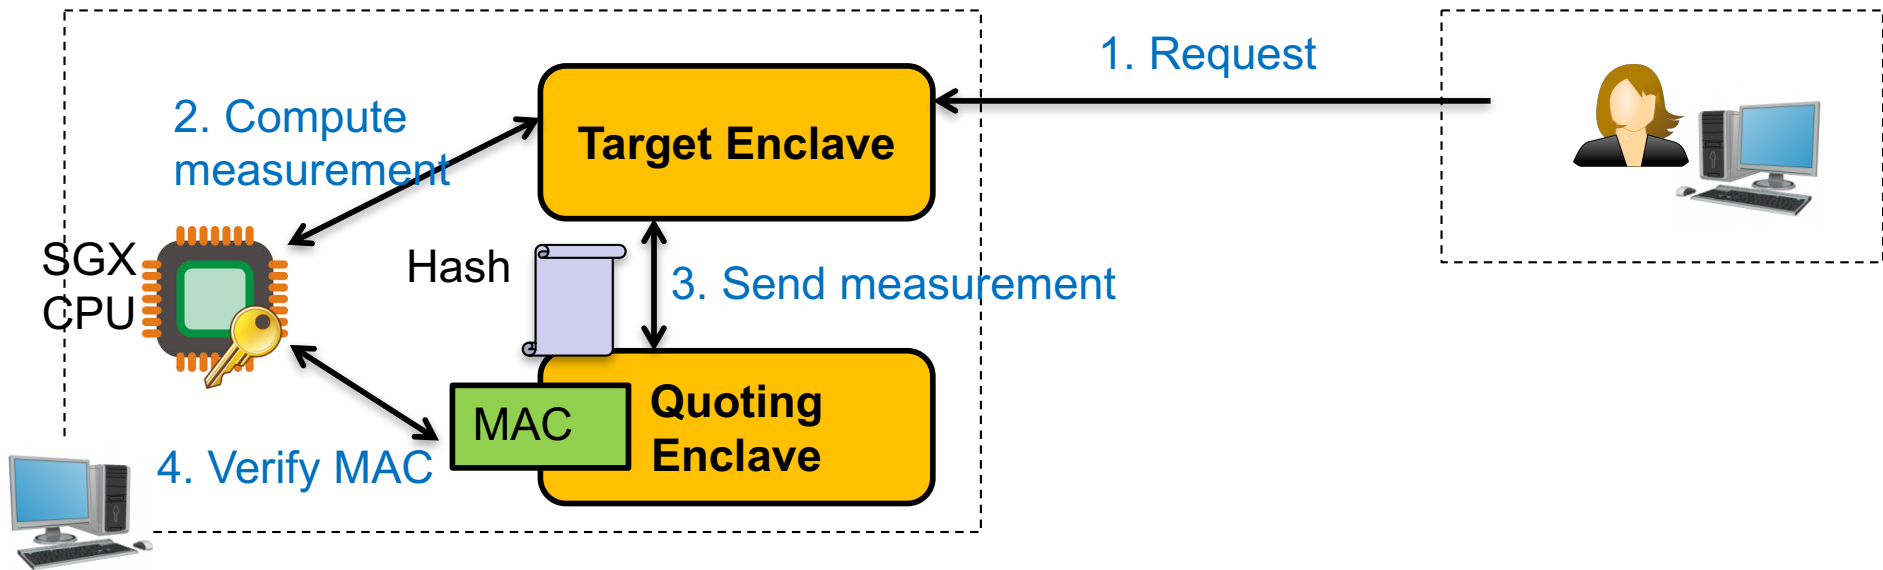


- Can also establish a **secure channel** between client and the enclave by exchanging Diffie-Hellman keys as part of the attestation process

SGX: How attestation works

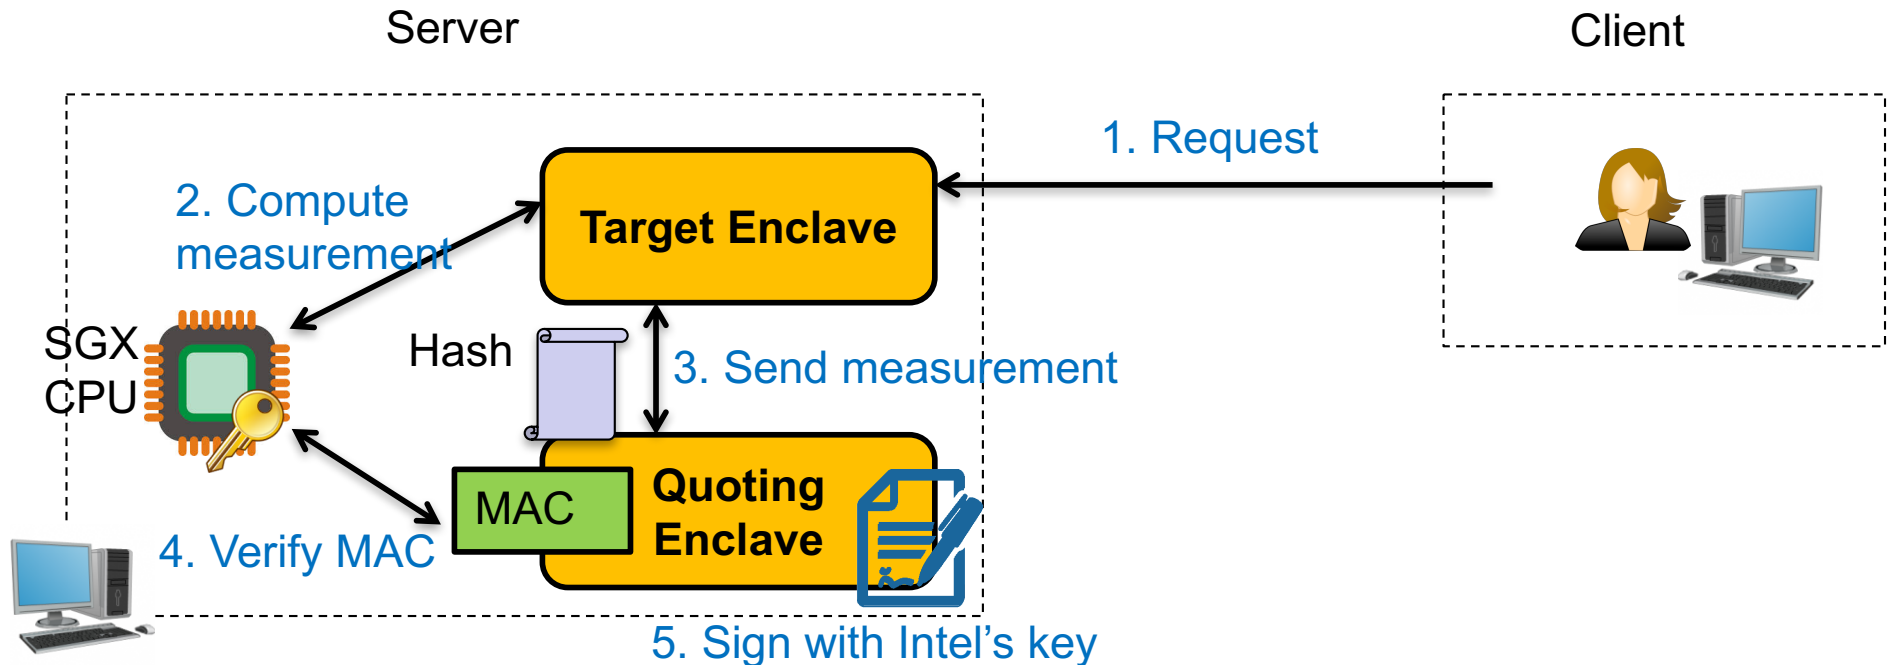
Server

Client



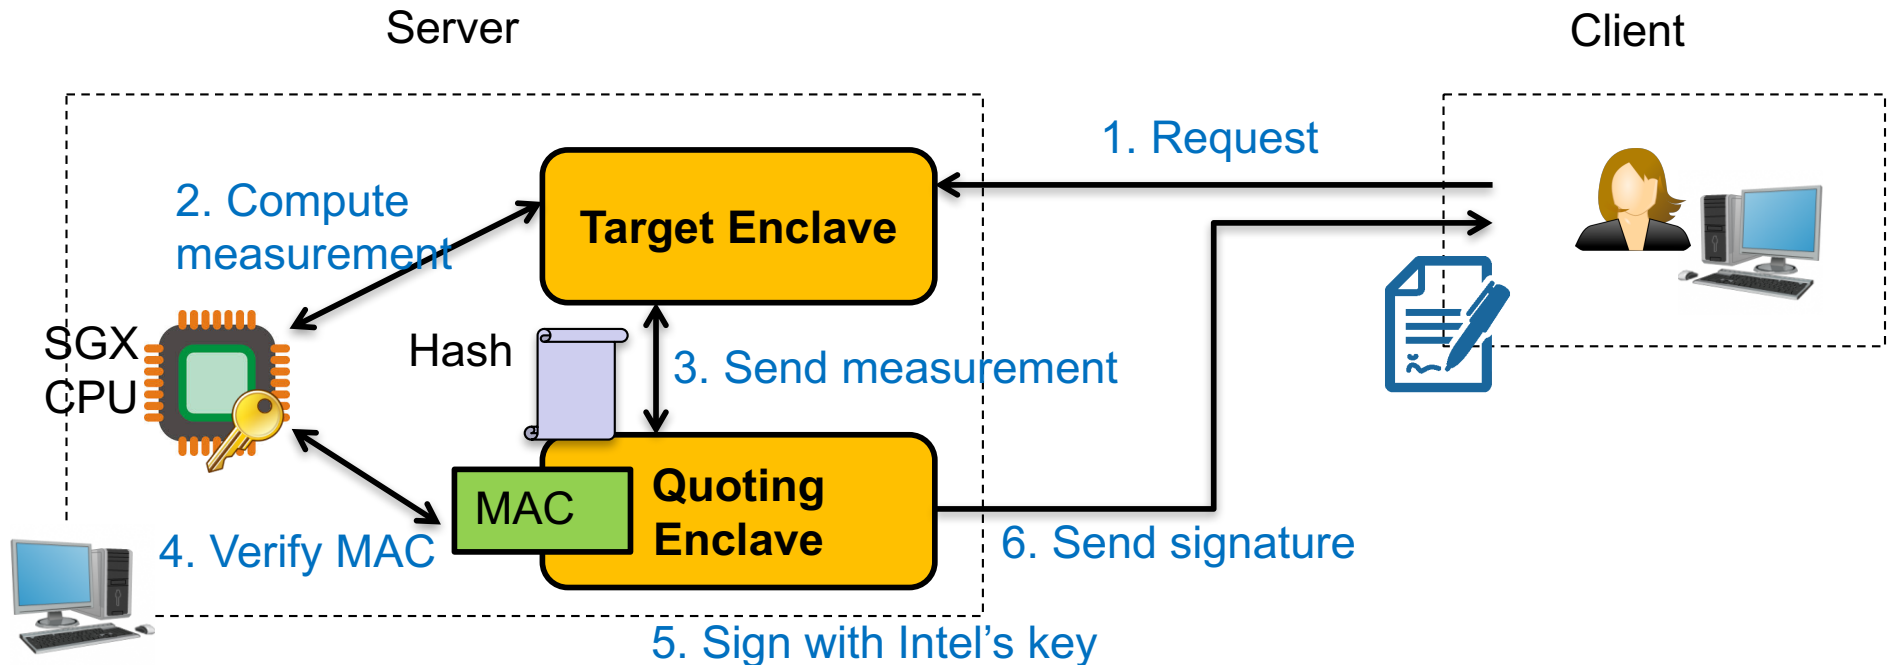
- Can also establish a **secure channel** between client and the enclave by exchanging Diffie-Hellman keys as part of the attestation process

SGX: How attestation works



- Can also establish a **secure channel** between client and the enclave by exchanging Diffie-Hellman keys as part of the attestation process

SGX: How attestation works



- Can also establish a **secure channel** between client and the enclave by exchanging Diffie-Hellman keys as part of the attestation process

Intel SGX

- **Minimal *TCB*** (trusted computing base):
 - Only the **processor + the code loaded into the enclave** need to be trusted
 - Nothing else (DRAM, peripherals, operating system, etc.) needs to be trusted

Intel SGX

- **Minimal *TCB*** (trusted computing base):
 - Only the **processor + the code loaded into the enclave** need to be trusted
 - Nothing else (DRAM, peripherals, operating system, etc.) needs to be trusted

So even if an attacker manages to gain root access on the server, won't be able to learn the data

Summary

Comparison

Computation with cryptography	Computation with SGX
No need to trust server-side hardware	Need to trust Intel's processor
No need to trust server-side software	Software running in enclave can leak unencrypted data No need to trust other privileged software (including the OS)
Can execute only a few simple functions efficiently	Runs arbitrary computation at processor speeds
Still vulnerable to side-channels	Still vulnerable to side-channels

Thank you, and good luck!