

Lec 16: User Authentication (2)


CSED415: Computer Security
Spring 2025

Seulbae Kim

POSTECH
POHANG UNIVERSITY OF SCIENCE AND TECHNOLOGY

Midterm exam

- Statistics
 - Max: 87
 - Average: 62
 - (vs 46 in 2024)
- To dispute:
 - Please meet me after today's class




Administrivia

- Lab 04 has been released!
 - About password-based authentication and entropy
 - Due on April 25

Recap

- Password-based authentication
 - Most widely used authentication method
 - Very easy to use and deployable
- Passwords are valuable, but considered weak due to
 - Human factors
 - Inevitable brute-force attacks
 - Incorrect policy

Means of authentication

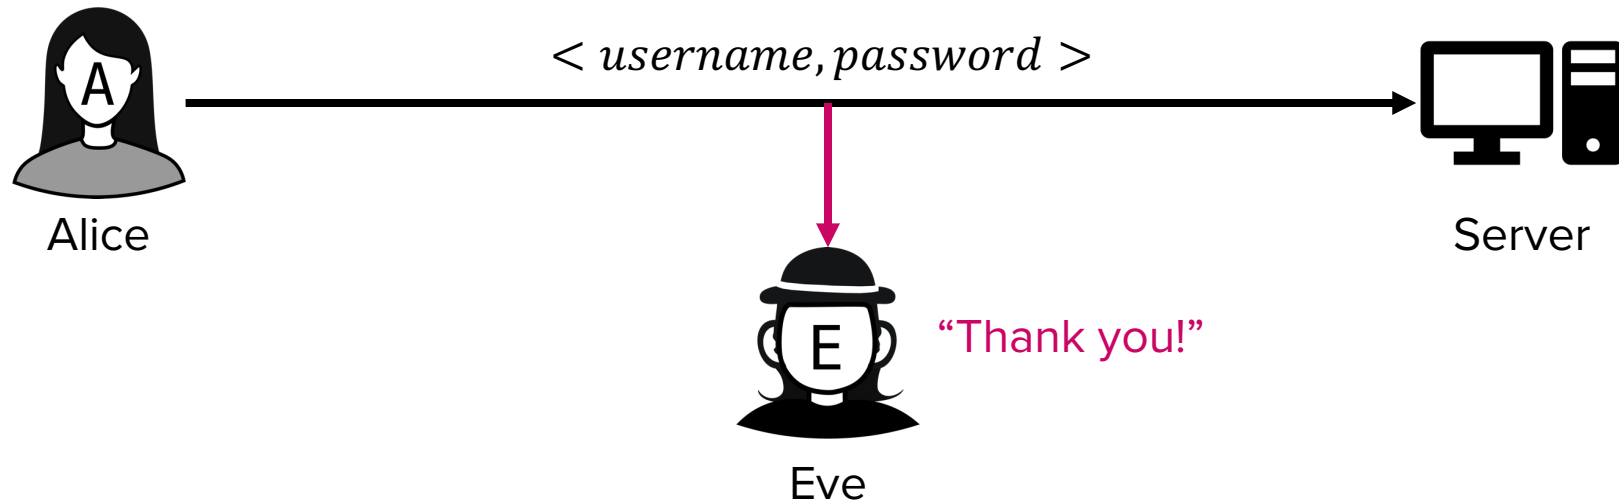
- Password-based 
- Challenge-response
- Biometric
- Zero-knowledge
- Multi-factor

Today's topic!

Challenge-Response Authentication

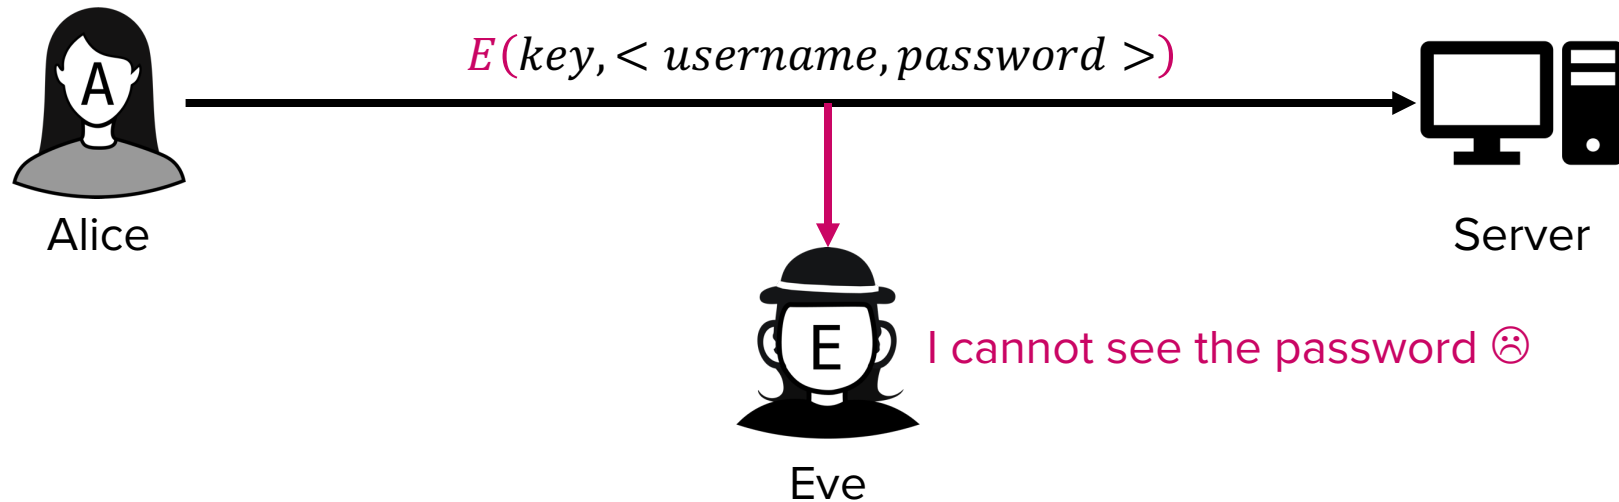
Transmitting a password

- How should a user transmit a password to a system?
 - Worst idea: Send the password in the clear (as plaintext)



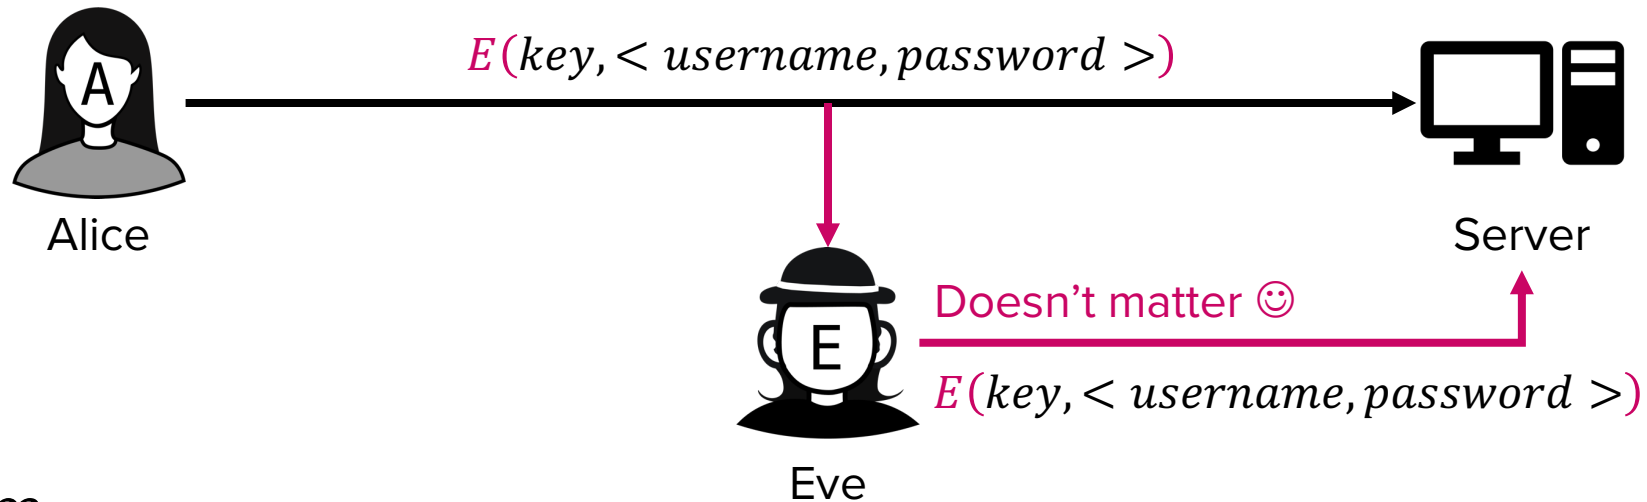
Transmitting a password

- How should a user transmit a password to a system?
 - Slightly better idea: Send the encrypted password



Transmitting a password

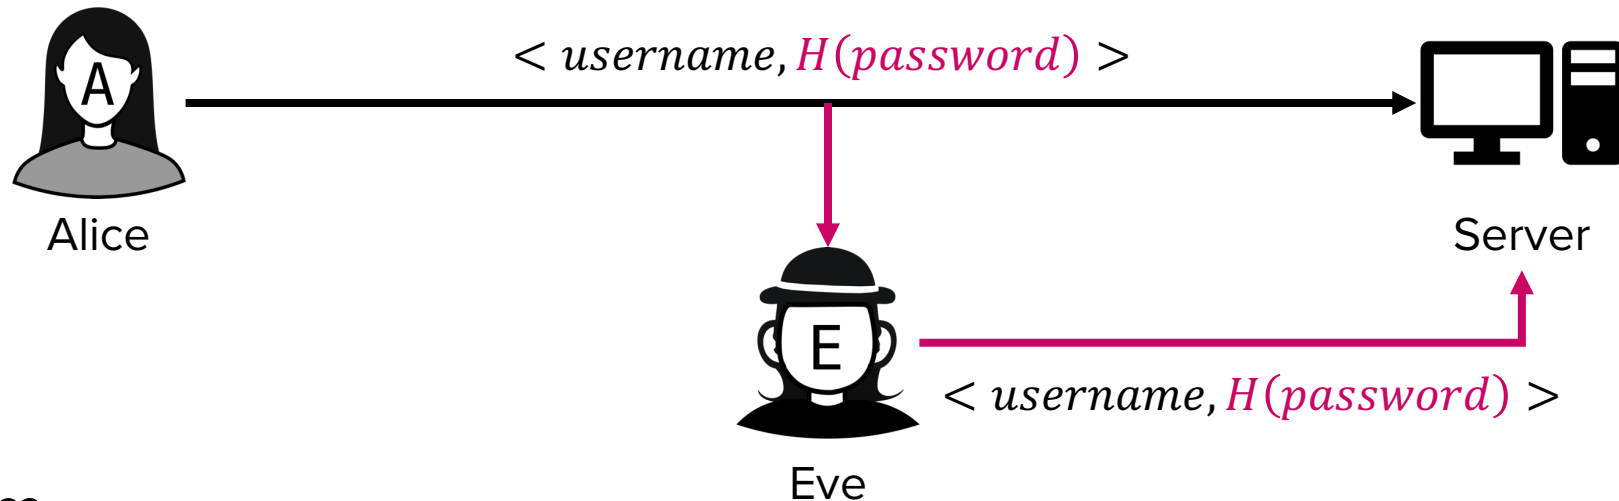
- How should a user transmit a password to a system?
 - Slightly better idea: Send the encrypted password



- Problem
 - An MitM attacker can record and replay the identification

Transmitting a password

- How should a user transmit a password to a system?
 - Another idea: Send the hashed password



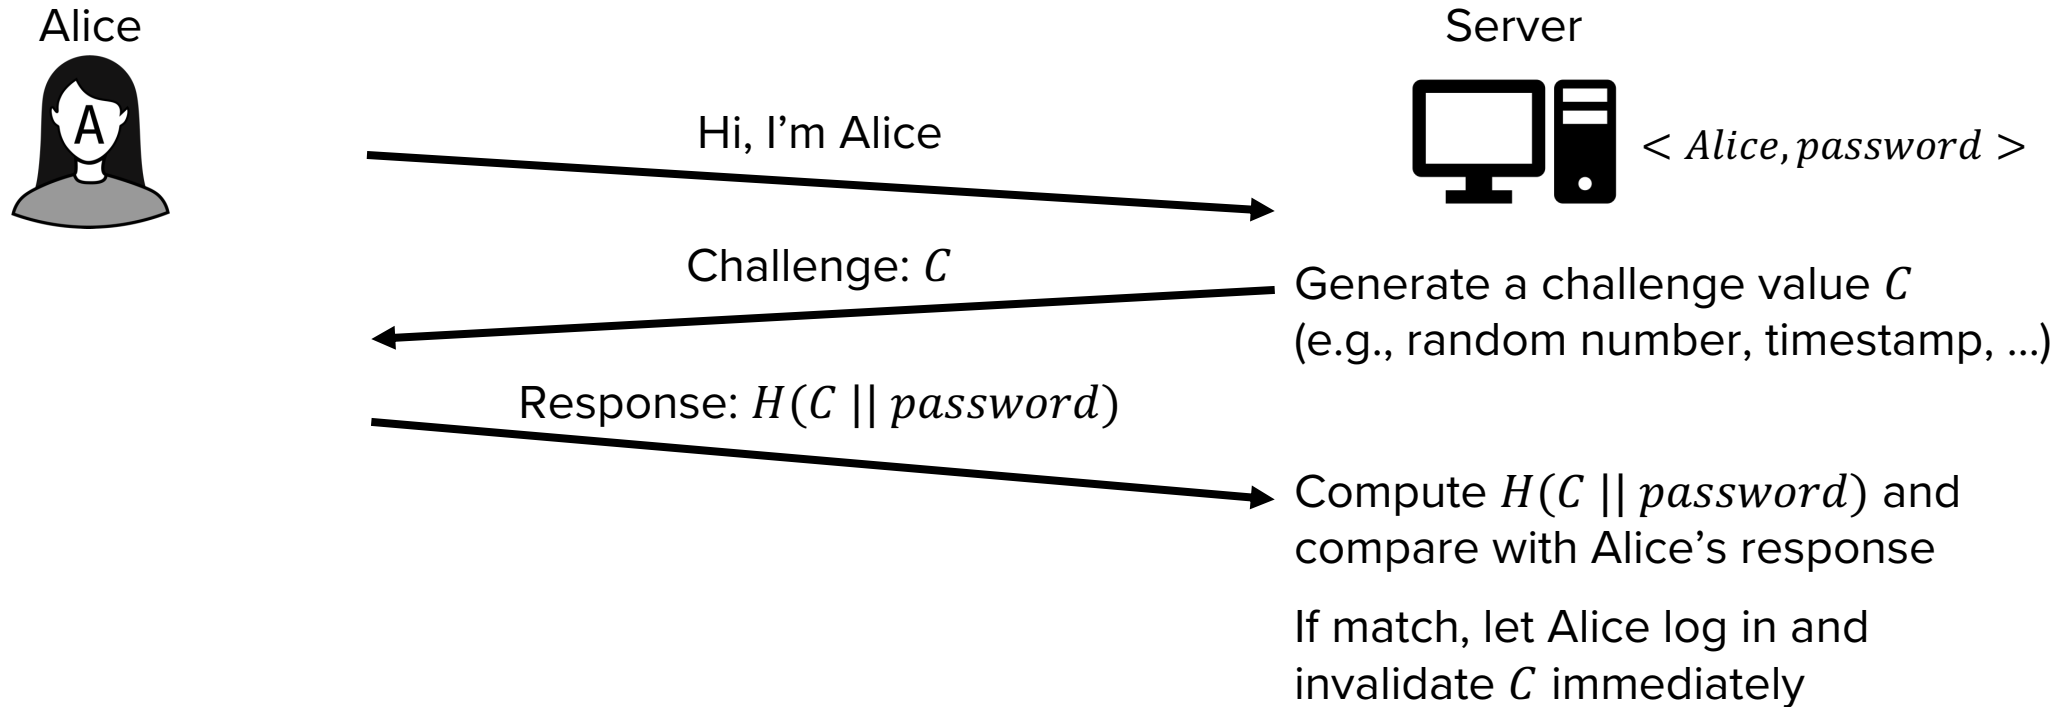
- Problem
 - Hashing does not improve security, since the hash can also be replayed

Transmitting a password

- How should a user transmit a password to a system?
 - Encryption and hashing do not automatically add security
 - A better idea: Challenge-response protocol

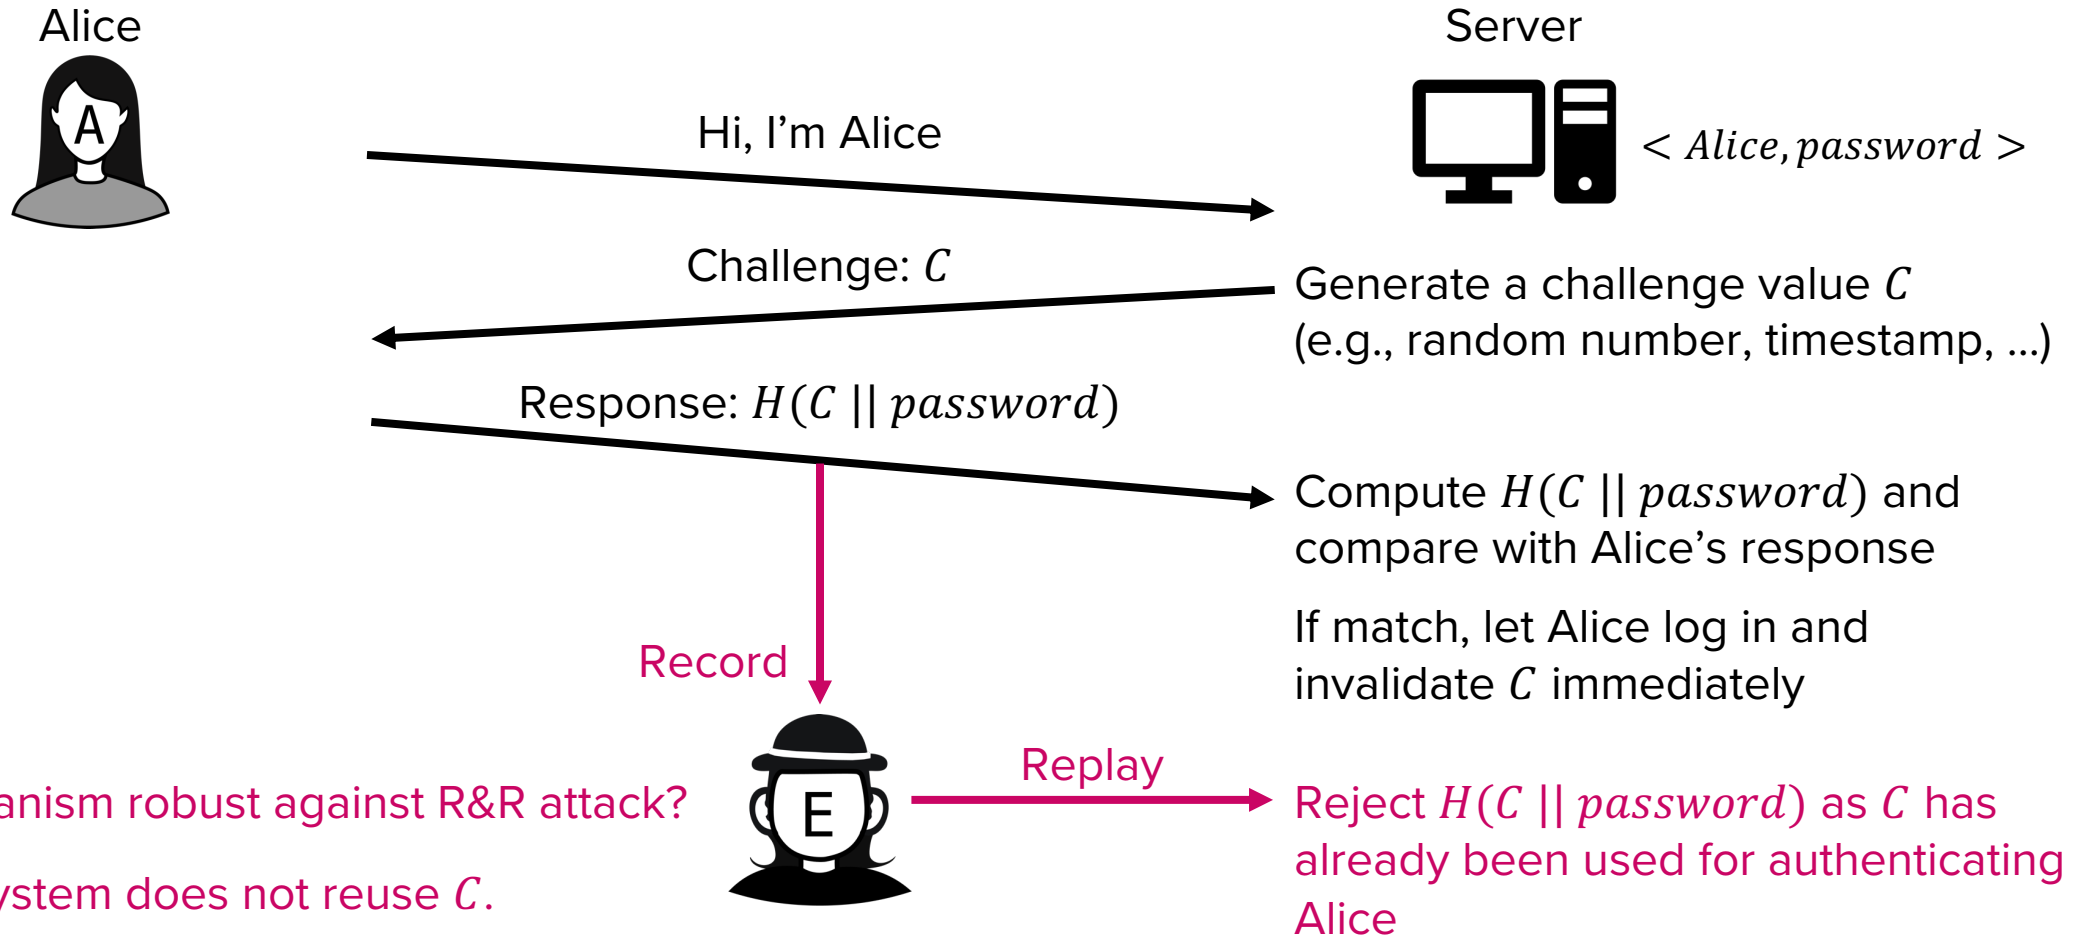
Challenge-response authentication

- Idea



Challenge-response authentication

- Idea

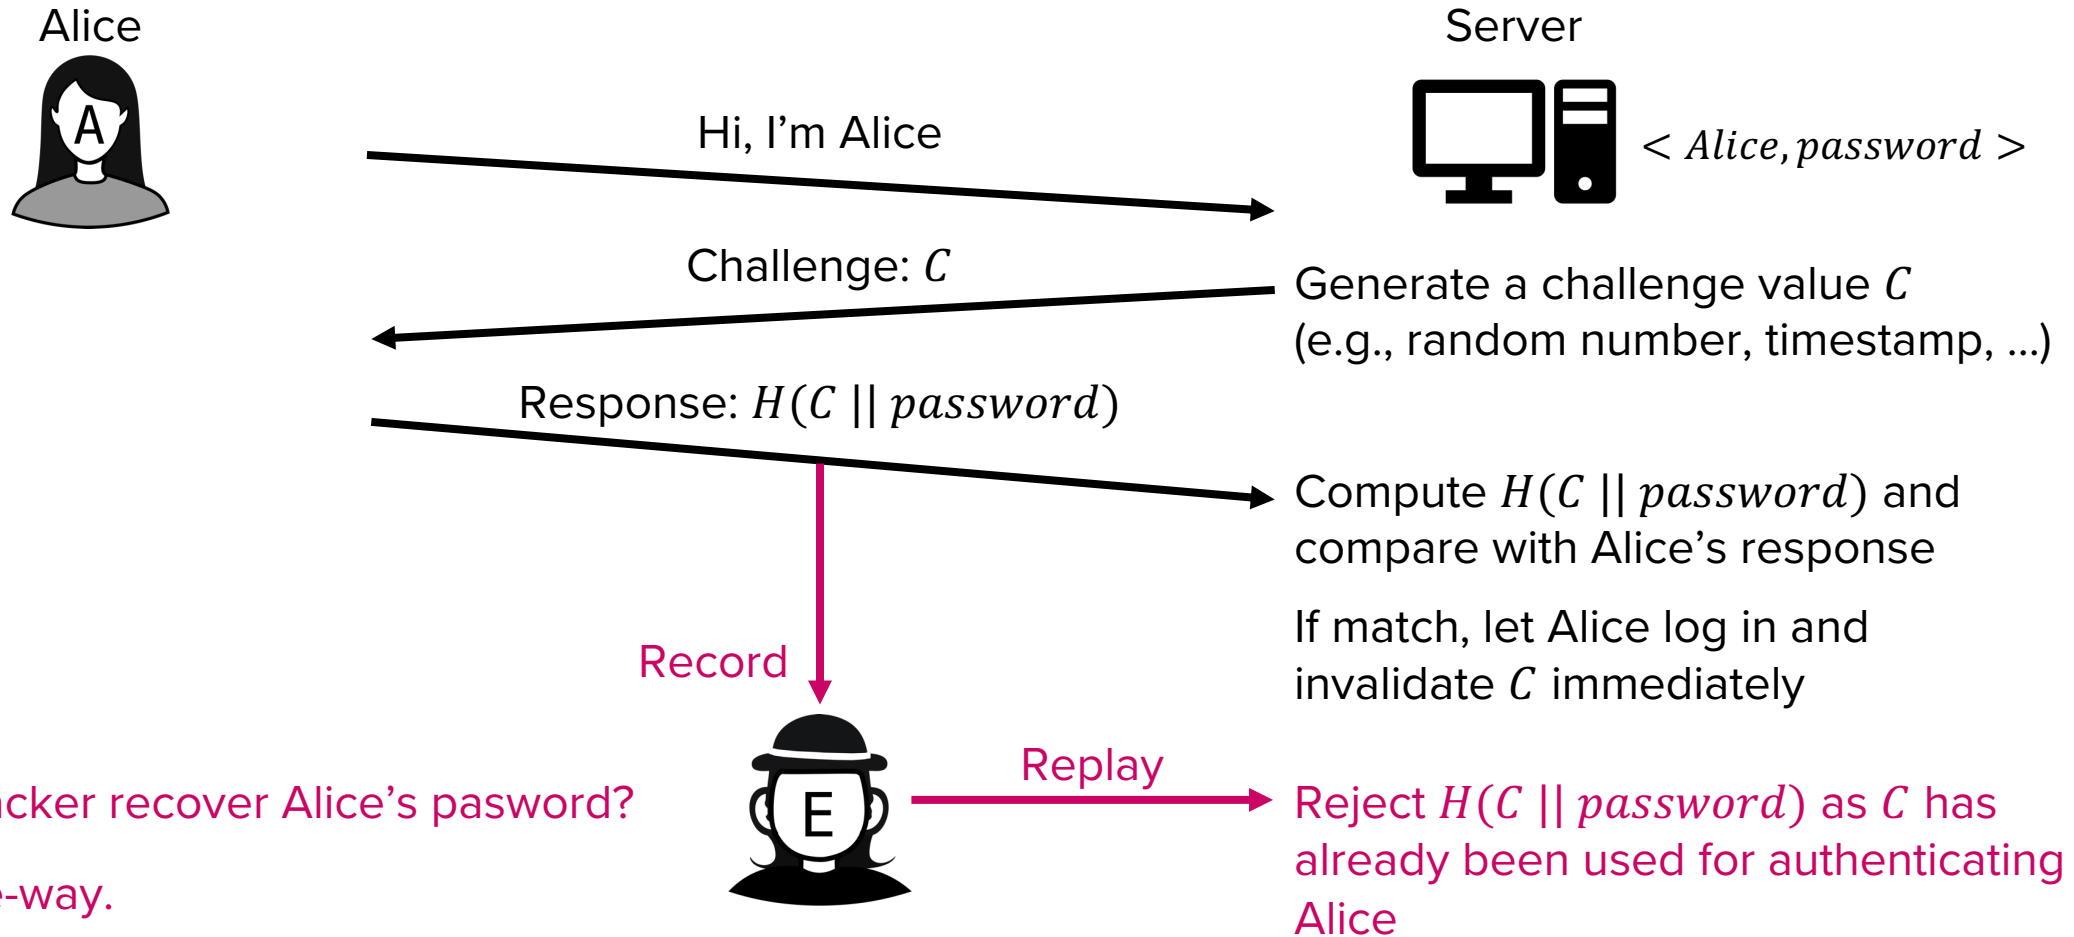


Q) Is the mechanism robust against R&R attack?

A) Yes! If the system does not reuse C .

Challenge-response authentication

- Idea




Q) Can the attacker recover Alice's password?

A) No! H is one-way.

Challenge-response in practice

- Symmetric key-based implementation (1)
 - Using shared key k and timestamp t (current time)



$E(k, t || password)$

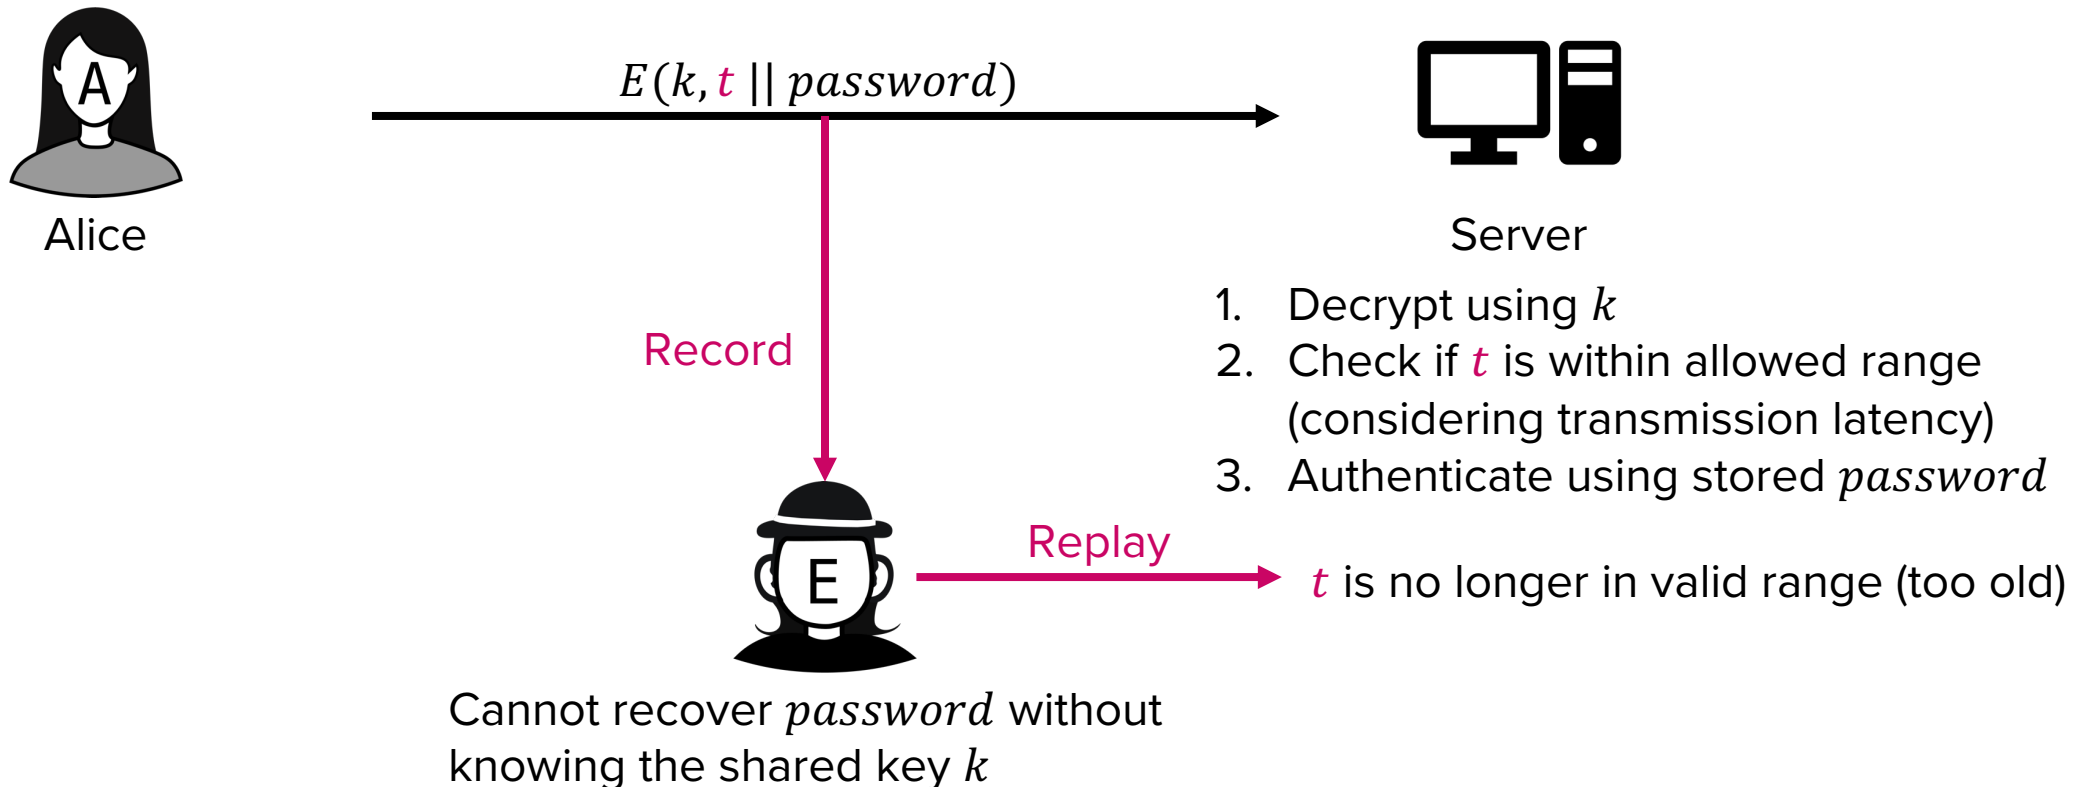


Server

1. Decrypt using k
2. Check if t is within allowed range (considering transmission latency)
3. Authenticate using stored *password*

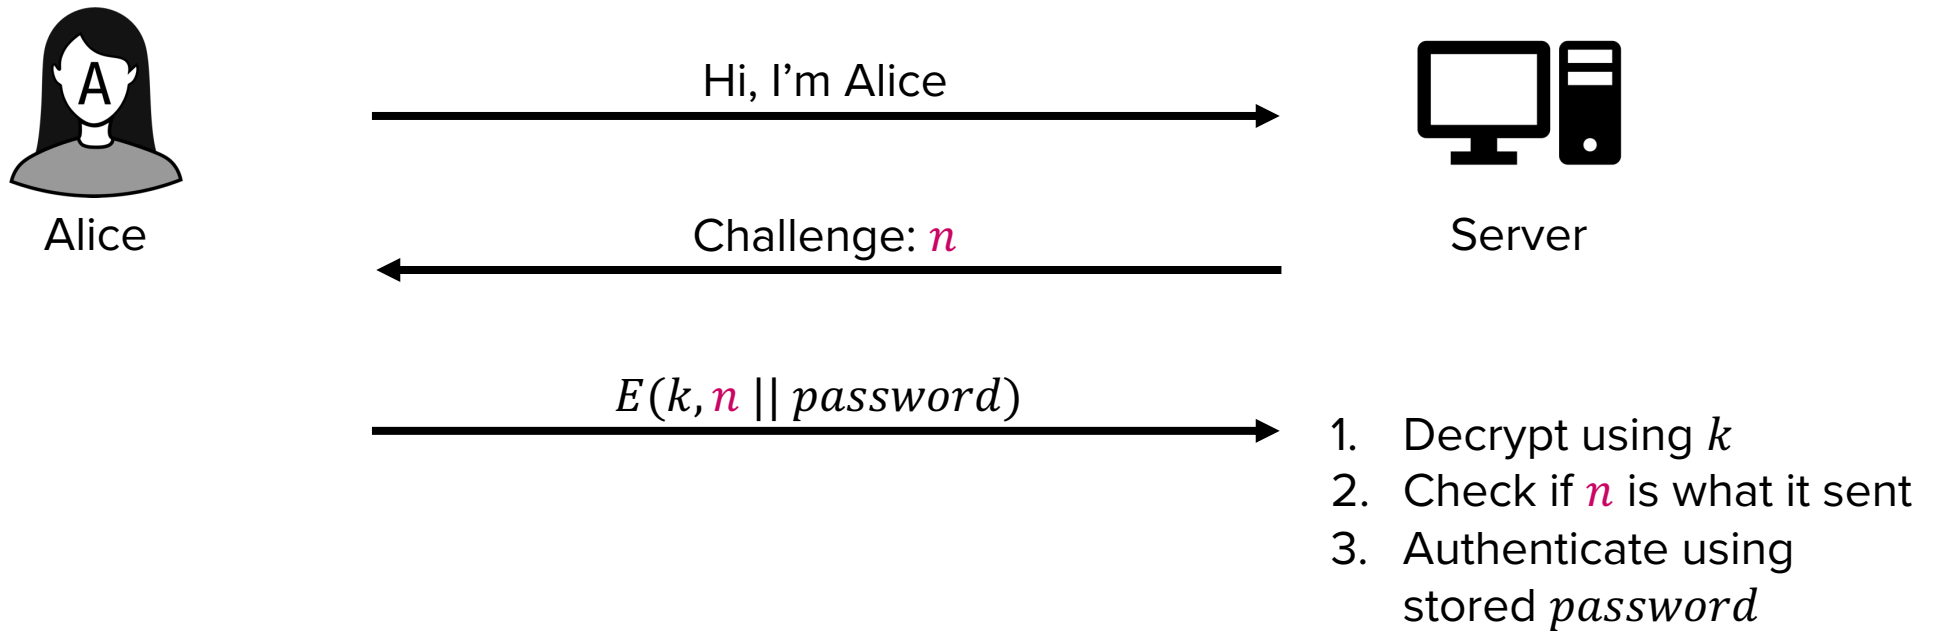
Challenge-response in practice

- Symmetric key-based implementation (1)
 - Using shared key k and timestamp t (current time)



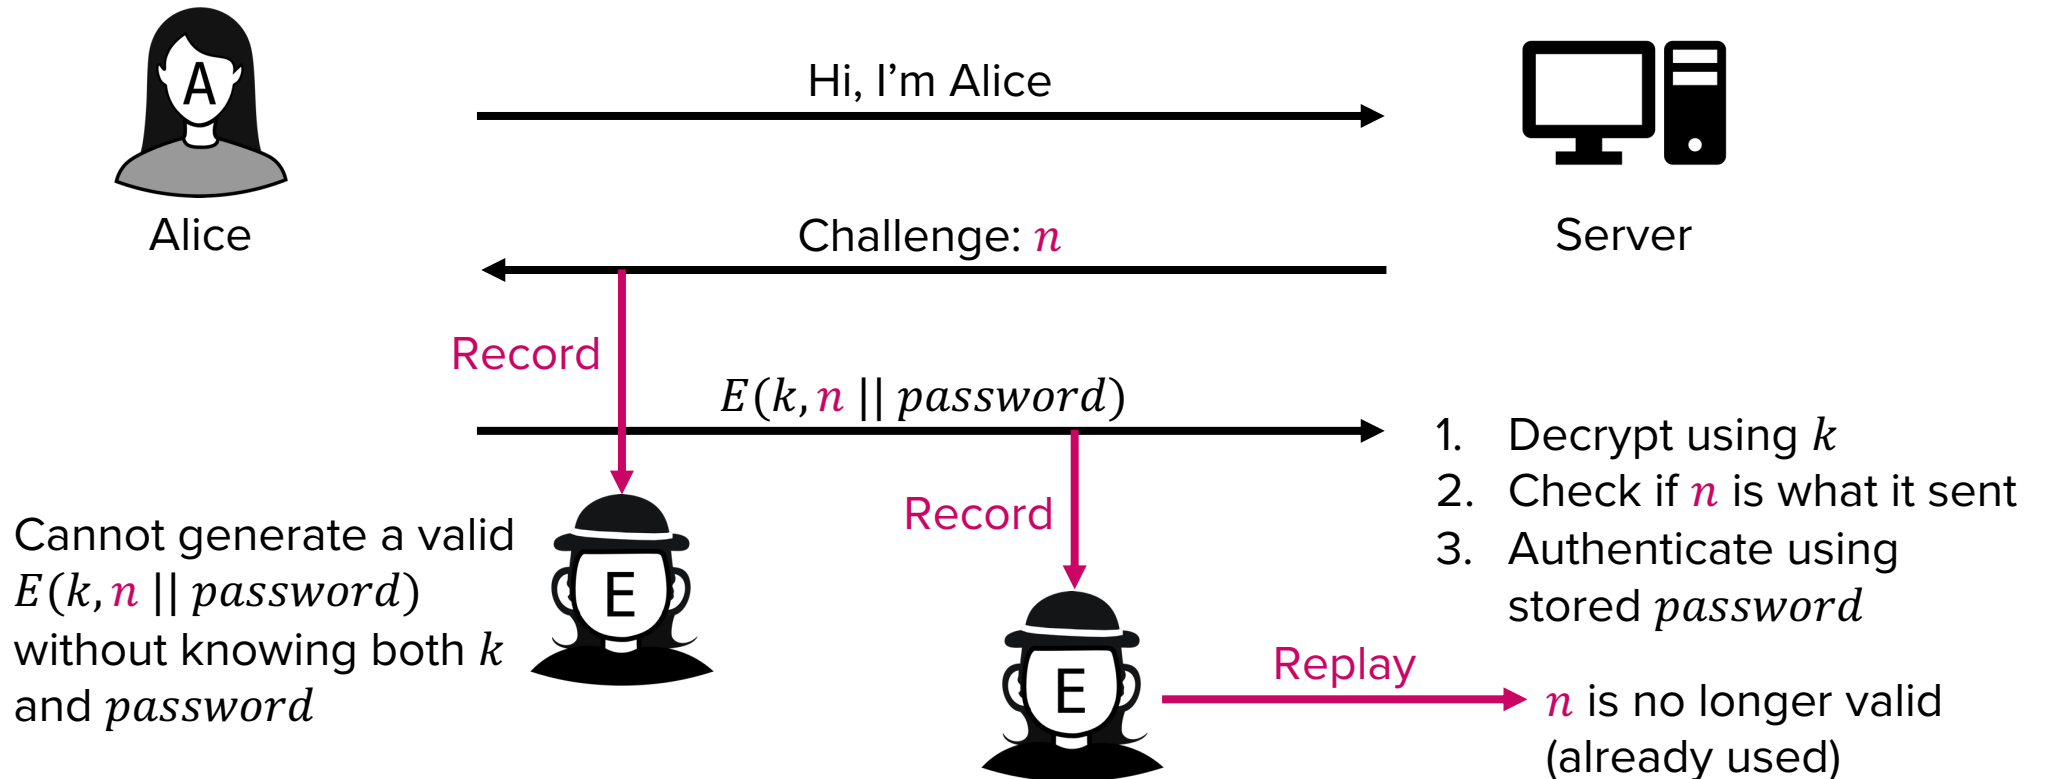
Challenge-response in practice

- Symmetric key-based implementation (2)
 - Using shared key k and a nonce n (random number)




Challenge-response in practice

- Symmetric key-based implementation (2)
 - Using shared key k and a nonce n (random number)



Challenge-response in practice

- Asymmetric key-based implementation (1)
 - Using public key k_p^S , secret key k_s^S , and timestamp t (current time)



$E(k_p^S, t || password)$




Server

1. Decrypt using k_s^S
2. Check if t is within allowed range (considering transmission latency)
3. Authenticate using stored *password*

Challenge-response in practice

- Asymmetric key-based implementation (1)
 - Using public key k_p^S , secret key k_s^S , and timestamp t (current time)



$E(k_p^S, t || \text{password})$



Server

1. Decrypt using k_s^S
2. Check if t is within allowed range (considering transmission latency)
3. Authenticate using stored *password*

Record



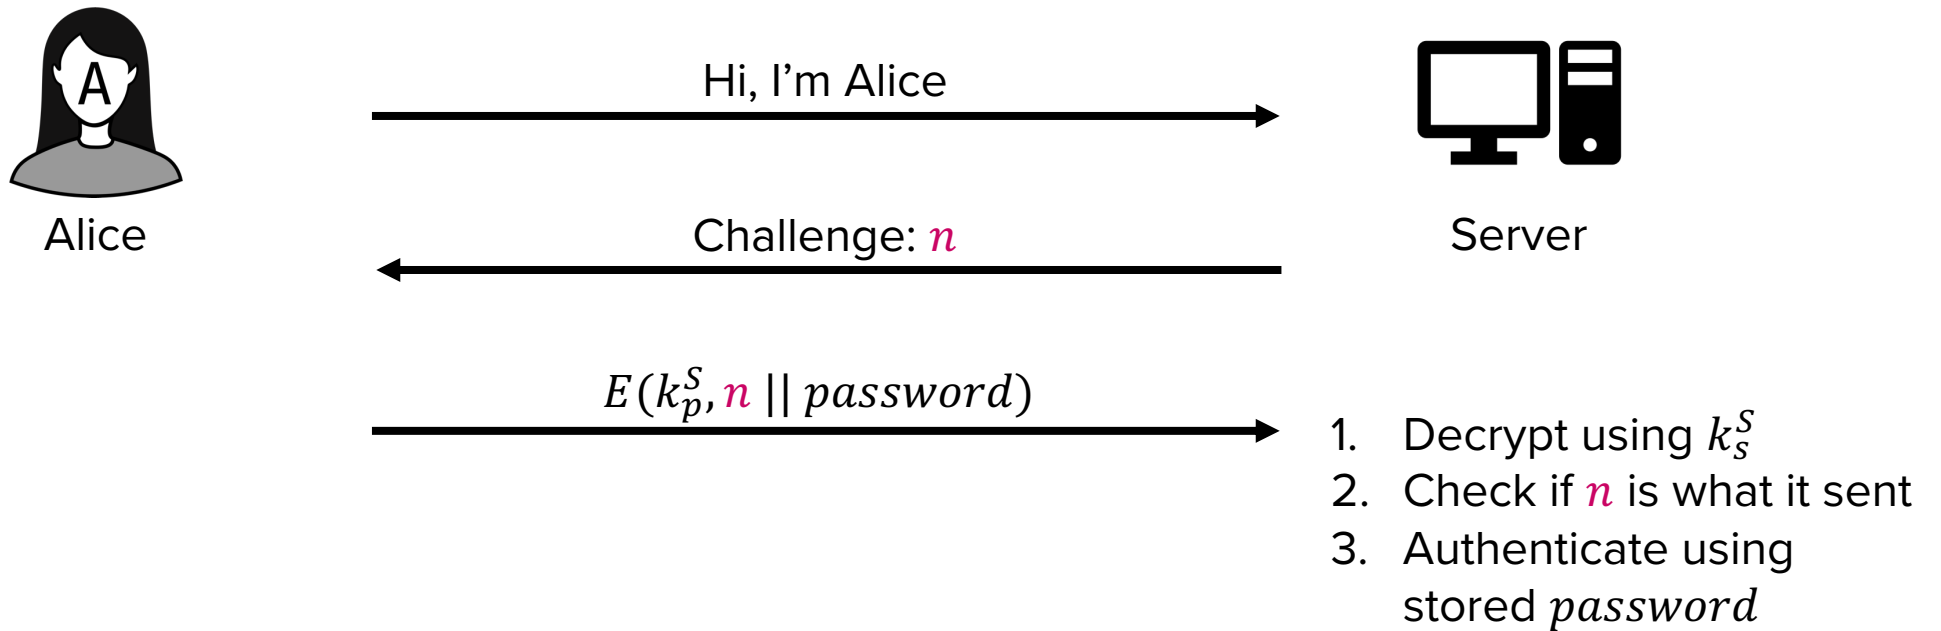
Replay

t is no longer in valid range (too old)

Cannot recover *password* without knowing server's secret key k_s^S

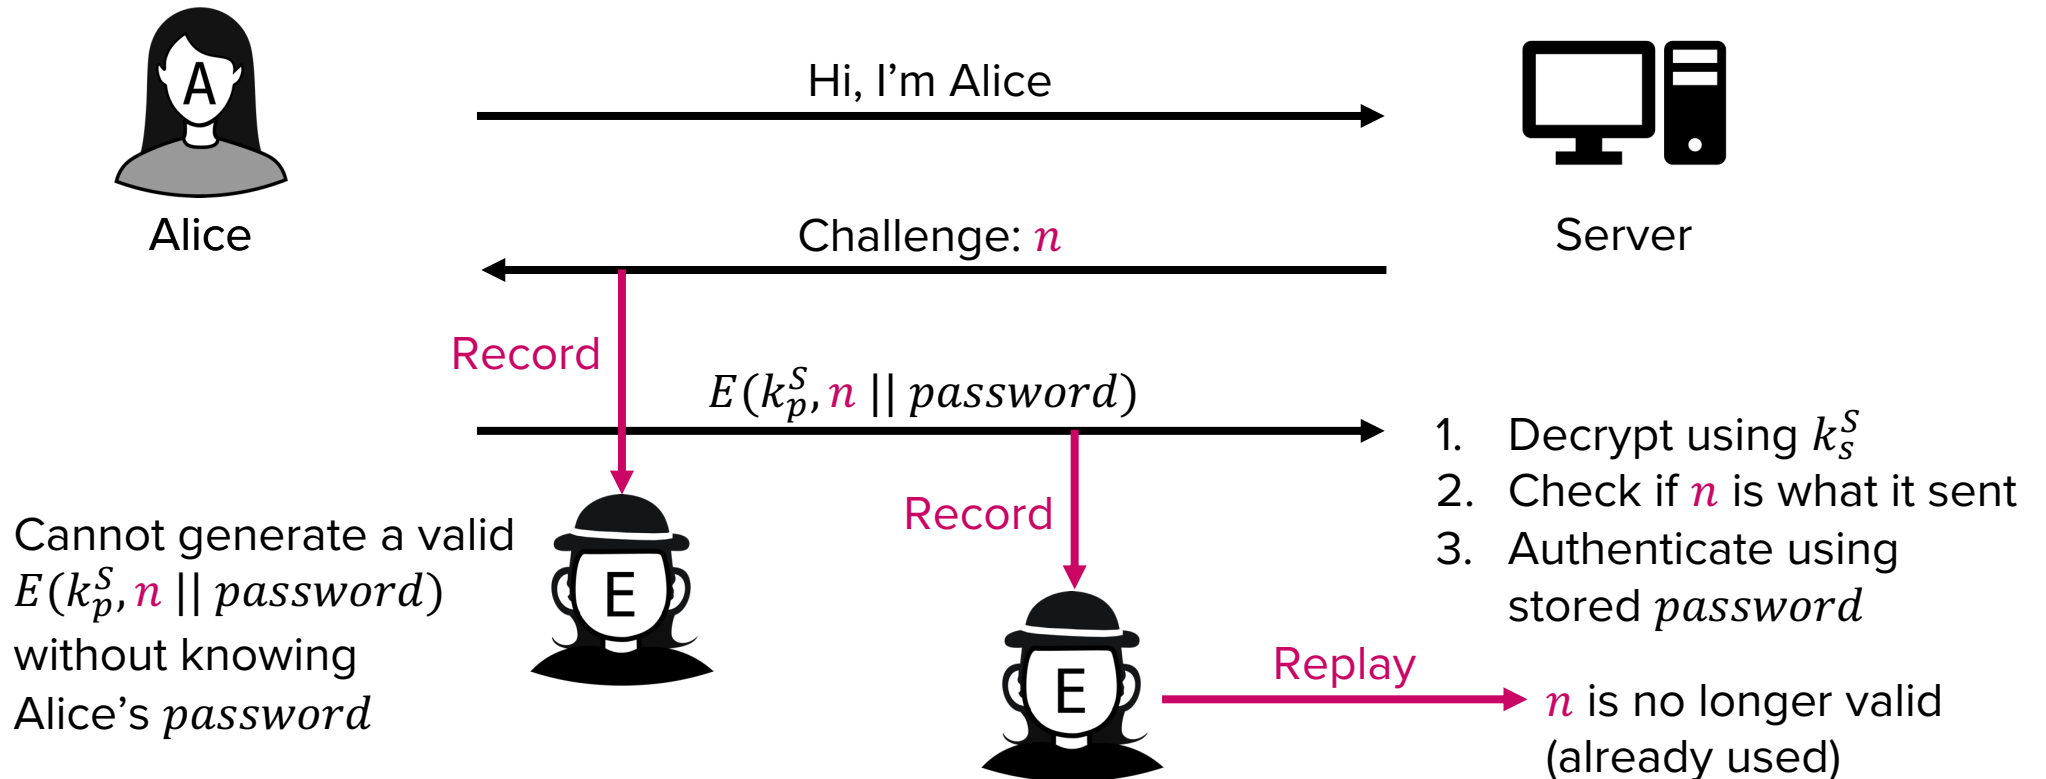
Challenge-response in practice

- Asymmetric key-based implementation (2)
 - Using public key k_p^S , secret key k_s^S , and a nonce n



Challenge-response in practice

- Asymmetric key-based implementation (2)
 - Using public key k_p^S , secret key k_s^S , and a nonce n



Biometric Authentication

Biometric authentication

- Use “something you are” for authentication
 - Authenticate users based on their unique physical characteristics
 - Characteristics include
 - Facial characteristics (e.g., Apple’s Face ID)
 - Fingerprints (e.g., Apple’s Touch ID)
 - Retina (Pattern of retinal blood vessels)
 - Iris
 - Voice




Image from All About Vision

Biometric authentication

- Advantages of using something you are for authentication
 - No need to remember anything (== can never forget the secret)
 - No need to carry anything (== can never lose the secret)
- Problems
 - Once compromised, cannot easily be changed
 - Not as accurate as digital methods (e.g., password matching)
 - Authentication is costly
 - Biometric information is considered more sensitive than a password
 - Your personal data needs to be stored on the service

Biometric authentication

- Problems
 - Accuracy: Not as accurate as digital methods, such as password matching
 - <https://www.youtube.com/watch?v=e8-yupM-6Oc>



The probability that a random person in the population could look at your iPhone or iPad Pro and unlock it using Face ID is less than 1 in 1,000,000 with a single enrolled appearance whether or not you're wearing a mask. As an additional protection, Face ID

<https://support.apple.com/en-us/102381>

Biometric authentication

- Problems
 - **Recovery:** If stolen or compromised, it is very hard to change biometric information
 - **Cost:** Authentication is slow and costly
 - Need a dedicated hardware (e.g., retina scanner, LiDAR, etc.)
 - **Privacy:** Your biometric data needs to be stored on the service
 - Biometric information is considered more sensitive than a password

Zero-knowledge Authentication


Your identity matters

- Problem of existing authentication methods
 - Your identity must be revealed during authentication
 - What you know (password / challenge-response)
 - What you have (token)
 - What you are (biometric information)

Zero-knowledge proofs (ZKP)


- Problem setting
 - Peggy is a prover and Victor is a verifier
 - Peggy wants to prove to Victor that **she knows the secret**
 - However, she does not want to reveal any other information to Victor
 - Including the secret itself
- Can Peggy authenticate without revealing her identity?

The Ali Baba cave example




The Ali Baba cave example

1. Peggy enters the cave and randomly selects a path w/o Victor seeing the path




The Ali Baba cave example

2. Victor enters and shouts the name of the randomly selected path



The Ali Baba cave example


3. If Peggy knows the password, she can return to Victor using the correct path





* If Peggy doesn't know the password, she still has a 50% chance to succeed


The Ali Baba cave example

4. Repeat multiple times until Victor is confident



Color-blind Victor example


- Victor has a “red-green color blindness”
 - He cannot tell red from green
- Setting
 - Prepare two balls  
 - One red ball, one green ball
 - All properties (weight, size, ...) are identical except for the color
 - Peggy should prove to Victor that the two balls have different colors



Ishihara Plate #9


Color-blind Victor example

1. Victor randomly selects a ball and shows it to Peggy




Color-blind Victor example

2. Victor enters a room and makes a random decision about switching the ball (switch or not switch)




Color-blind Victor example

3. Victor shows the ball to Peggy and asks if he switched the balls



Color-blind Victor example

4. Repeat steps 1-3 until Victor is confident



Color-blind Victor example







- Probability that Peggy is also color-blind but gets the answer right is 50%
 - Experiment repeated 10 times, probability that Peggy does not know the secret becomes $\frac{1}{2^{10}}$ (less than 0.1%)
- Victor learns that the two balls are distinguishable without learning the color of each ball

ZKP for user authentication

- Secure Remote Password (SRP) protocol
 - User authentication using ZKP
 - Server does not store user's password
 - Server verifies that the user knows the password w/o seeing the password (zero-knowledge!)

ZKP for user authentication


- Recap: Diffie-Hellman key exchange
 - Public information
 - Large prime number p and its generator g
 - Secret information
 - Alice's secret key a and Bob's secret key b
 - Exchange
 - Alice sends $A = g^a \bmod p$ to Bob, Bob sends $B = g^b \bmod p$ to Alice
 - Key derivation
 - Alice derives a shared key $k = B^a \bmod p = g^{ab} \bmod p$
 - Bob derives the same shared key $k = A^b \bmod p = g^{ab} \bmod p$

	Public	Secret
	$p = 11$	$a = 15$ 
	$g = 6$	$b = 8$ 
	$A = 10$	$k = 1$ 
	$B = 4$	

SRP protocol

• Step 1: Registration

Secret:
Password $pass$



Compute $x = H(pass || salt)$
Compute $v = g^x \bmod p$

- Username *Alice*
- Randomly selected $salt$
- Verifier $v = g^x \bmod p$




Store: $(Alice, salt, v)$

Public: Prime p , generator g

SRP protocol

- Step 2: Parameter sharing

Secret:
Password *pass*



Derive $x = H(\textit{pass} || \textit{salt})$

Generate random secret *a*

Done sharing

username = Alice

salt

$$A = g^a \bmod p$$

$$u, B = v + g^b \bmod p$$

Public: Prime p , generator g

salt A u B



Storage:
(*Alice, salt, v*)
($v = g^x \bmod p$)


Fetch Alice's *salt*

Generate random param *u*
and random secret *b*

SRP protocol

- Step 3: Session key derivation

Secret:
Password *pass*



Derive $x = H(\text{pass} || \text{salt})$

Generate random secret *a*

Done sharing

Derive:

- $S = (B - g^x \text{ mod } p)^{a+ux}$
- Session key $K = H(S)$

username = Alice

salt

$A = g^a \text{ mod } p$

$u, B = v + g^b \text{ mod } p$

Same key has been derived without revealing *pass*

Public: Prime p , generator g

salt A u B



Storage:
(Alice, salt, v)
($v = g^x \text{ mod } p$)

Fetch Alice's salt

Generate random param u
and random secret *b*


Derive:

- $S = (A * v^u)^b \text{ mod } p$
- Session key $K = H(S)$

SRP protocol

- Step 4: Mutual authentication

Secret:
Password $pass$



$username = Alice$

Public: Prime p , generator g
 $salt \ A \ u \ B$



Storage:
 $(Alice, salt, v)$
 $(v = g^x \bmod p)$

Derive:

- $S = (B - g^x \bmod p)^{a+ux}$
- Session key $K = H(S)$

Derive:

- $S = (A * v^u)^b \bmod p$
- Session key $K = H(S)$

Compute and send M_1

$$M_1 = H(A || B || K)$$

If valid, login success!

$$M_2 = H(A || M_1 || K)$$

If valid,
compute and send M_2

SRP protocol

- Strengths
 - Resistant to leaks
 - Server does not store any password
 - Resistant to dictionary attacks
 - $pass$ or $x = H(pass || salt)$ are never sent in public
 - Resistant to active attacks
 - Mallory cannot derive the session key K from any publicly transmitted information
- Weakness
 - Slow!

Multi-factor Authentication

Multi-factor authentication (MFA)

- User provides two or more identifications
 - What you know (password) + what you are (fingerprint)
 - What you know (password) + what you also know (PIN)
 - ...
- Fortifies inherently weak password-based authentication by providing an additional layer of security
 - Leaked passwords → Fingerprint cannot be leaked
 - Brute-forcing → Cannot brute-force fingerprint

Practical MFA implementation

- Password + One-time code sent via SMS message
 - Server stores the user's phone number
 - Advantage:
 - Easy to implement
 - Compromised server does not automatically break security unless the user's phone is also compromised
 - Disadvantage:
 - Phone network and carriers should be trusted
 - Could lead to phishing attacks

Practical MFA implementation

- Password + One-time code sent via SMS message
 - Known attacks:
 - SIM swapping
 - Attacker collects various personal information of the victim
 - The attacker impersonates the victim and convinces the victim's phone carrier to port the number to a new SIM card
 - The victim loses phone connection and the attacker's phone is activated with the victim's phone number
 - The attacker attempts to log into a service using victim's leaked credentials
 - The attacker receives the one-time login code sent to the victim and breaks 2FA
 - The victim should make phone calls for recovery, but cannot do so without a number

Practical MFA implementation

- Password + Time-based one-time passwords (TOTP)
 - Server and user device agree on a secret value
 - Google's Authenticator app allows users to scan a QR code to register secret
 - User device generates $TOTP = H(secret || cur_time)$
 - Use coarse-grained time (e.g., *cur_time* is updated every 30 seconds)
 - User enters the *TOTP* and server checks if it is valid
 - Advantages:
 - Do not need network connection, do not need to trust phone carriers
 - Disadvantages:
 - Needs extra steps for app installation and setup
 - If the server is compromised, all secret values need to be re-registered

Evaluating Authentication Method

Evaluating authentication method

- Metric for usability and security: Confusion matrix
 - True/False: Intended/Unintended
 - Positive/Negative: Allow/Disallow

System

		System	
		Allow	Disallow
User	Alice logs in as Alice	True Positive	False Negative
	Attacker logs in as Alice	False Positive	True Negative

High FN causes inconvenience
(bad usability)

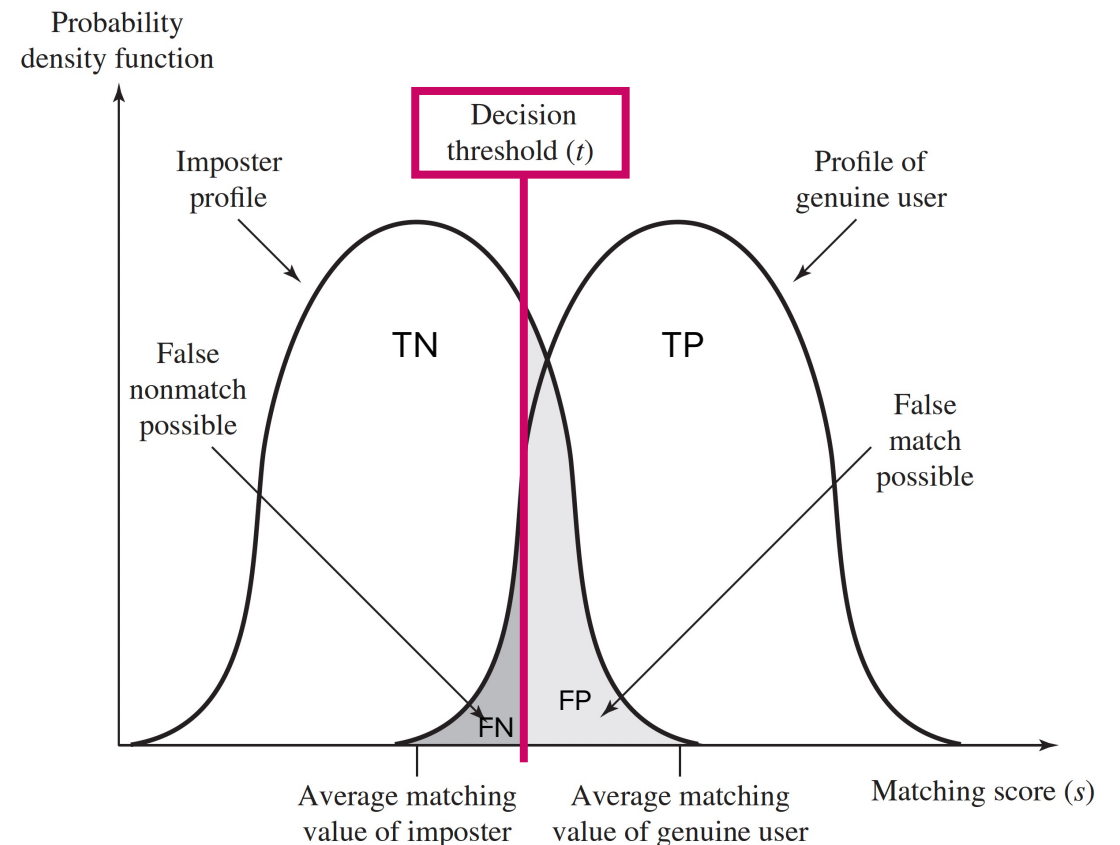
High FP means high exploitability
(bad security)

→ Goals:

- Very high TP
- Very low FN
- Zero FP

Evaluating authentication method

- A dilemma: There is no clean separation between imposter and user profiles
 - Increase the threshold to get:
 - Increased security (FP↓)
 - Decreased convenience (FN↑)
 - Decrease the threshold to get:
 - Decreased security (FP↑)
 - Increased convenience (FN↓)



Profiles of a biometric characteristic of an imposter and an authorized user

Summary

- User authentication is hard
- Password-based auth is a long-lasting solution
- Strengthen passwords with password managers and MFA

Coming up next

- Authentication: To open the front door or not
 - Coarse-grained control for the entire system accessibility
- Access control: After opening the door to a user
 - Fine-grained control for system resources

Questions?