Week 1

AP Cybersecurity

Homework bundle for this week

本周作业合集

exercises.lol

Contents 目录

1. Introduction to Security

Starter Quiz · 课前小测

AP Cybersecurity Name: _____ Score: _____

1. What is the main goal of social engineering?
   - ☐ To break the encryption
   - ☐ To trick a person into breaking security
   - ☐ To overload a server
   - ☐ To scan for open ports
2. Which is a sign of an online password attack in the logs?
   - ☐ One successful login at noon
   - ☐ Many failed logins in a short time
   - ☐ A file being downloaded
   - ☐ A printer coming online
3. An adversary sets up a fake access point copying the real SSID. This is a(n)...
   - ☐ jamming attack
   - ☐ evil twin attack
   - ☐ war driving attack
   - ☐ smurf attack
4. An AI clone of a person's voice used to impersonate them is a...
   - ☐ dictionary
   - ☐ deepfake
   - ☐ honeypot
   - ☐ keylogger
5. Why is AI useful for reading network logs?
   - ☐ It makes the logs shorter
   - ☐ It can sort millions of events humans cannot
   - ☐ It deletes the logs
   - ☐ It encrypts the logs
6. On an evil-twin network, HTTPS sites still keep your data encrypted.

   ☐ True ☐ False

7. AI security recommendations should be reviewed by a human before use.

   ☐ True ☐ False

8. Giving an adversary your pet's name and birthday could help them impersonate you.

   ☐ True ☐ False

9. Multifactor authentication (MFA) adds security beyond just a password.

   ☐ True ☐ False

10. Enabling MFA can stop an adversary who has only cloned your voice.

    ☐ True ☐ False

1. Introduction to Security

Answers · 答案

AP Cybersecurity

1. **To trick a person into breaking security**

   Social engineering targets the **human**, not the code.

2. **Many failed logins in a short time**

   A burst of **failed** logins signals guessing.

3. **evil twin attack**

   Copying the SSID to lure victims is the **evil twin**.

4. **deepfake**

   A **deepfake** clones a voice or face.

5. **It can sort millions of events humans cannot**

   AI handles the **scale** humans cannot.

6. **True**

   HTTPS encryption protects you even if traffic is intercepted.

7. **True**

   AI is probabilistic; a human must check.

8. **True**

   That info often answers security **challenge questions**.

9. **True**

   MFA asks for extra proof, like a texted code.

10. **True**

    A second factor blocks a voice-only attack.

AP CYBERSECURITY • EXERCISE SHEET

1.1 Understanding Social Engineering

Name: _____ Class: _____ Date: _____ **Total: 14 marks**

KEY WORDS urgency 紧迫感 • impact 影响 • adversary 对手 • intimidation 恐吓 • social engineering 社会工程学

Objective

Build the skills to answer exam questions on **social engineering** 社会工程学 and its psychological tactics.

You must be able to:

- name social-engineering tactics: **intimidation** 恐吓, **urgency** 紧迫感, authority, consensus
- explain how a tactic pressures a victim to act
- describe the impact when a victim reveals information or downloads **malware** 恶意软件

1 Worked examples

Study these first. Each one shows the method for a task used later.

■ Spotting the tactic

Read the message and ask what feeling it creates. A threat of a bad result = **intimidation**. A tight deadline = **urgency**. A claim that everyone else already acted = **consensus**. Borrowed power from a boss = **authority**.

■ Explaining the impact

State what the victim gives up and what the adversary can then do. Example: giving a one-time password lets the adversary log in as the victim; giving personal details lets them answer challenge questions and impersonate the victim.

2 Practice

2.1 Define **social engineering** in one sentence. [2]

2.2 Name two psychological tactics used in social engineering and describe each briefly. [4]

2.3 An email says: "Confirm your details in the next hour or your account is deleted."

(a) Identify the main tactic used. [1]

(b) Explain how it pressures the reader. [2]

3 Exam-style questions

3.1 An adversary threatens to report an employee to the police unless they send a file. Which tactic is this? [1]

- **A** Consensus
- **B** Intimidation
- **C** Familiarity
- **D** Scarcity

3.2 A victim gives an adversary their one-time password (OTP). The most likely impact is that the adversary can [1]

- **A** encrypt the network
- **B** log in to a service as the victim
- **C** jam the Wi-Fi
- **D** clone an access card

3.3 A message reads: "Over 90% of your team has already verified their account."

(a) Name the tactic being used. [1]

(b) Explain why this tactic is persuasive. [2]

Solutions

2.1 manipulating people (not code) into breaking security by revealing information, clicking a link, or running malware.

2.2 any two of: intimidation = threatens a bad result; urgency = invents a deadline; authority = borrows a boss's power; consensus = claims everyone else acted.

2.3 (a) urgency. (b) the 1-hour deadline makes the reader act quickly without stopping to check whether the request is safe.

3.1 B. a threat of a negative consequence is **intimidation**.

3.2 B. an OTP lets the adversary authenticate as the victim.

3.3 (a) consensus. (b) people tend to follow the crowd, so believing others acted pressures the victim to act too.

Questions 8 through 10 refer to the following.

Consider the following email that was identified and reported as a phishing attempt.

From: Human Resources
To: All Employees

Subject: Final Reminder — Deadline Today!

Our records indicate that your employee profile is incomplete. In order to remain in compliance with company-wide HR standards, all staff must confirm their personal information by the end of the day. Failure to respond by this deadline may result in suspension of payroll access until your profile is verified.

Please note: Over 92% of employees have already completed this verification, and you are among the last to finalize your record.

To confirm your employee file, reply to this email with the following details:

1. Personal phone number
2. Pet's name (for verification)
3. Date of birth

We appreciate your cooperation in keeping our records accurate and ensuring continued access to employee services.

Thank you,

Human Resources

8. Which of the following shows how the email uses urgency to pressure the recipient into responding?
   (A) The recipient is warned that they must confirm their personal information by the end of the day.
   (B) The recipient is warned that failure to comply could result in payroll suspension.
   (C) The recipient is encouraged to help human resources keep their records accurate.
   (D) The recipient is told that the message comes from a trusted internal department.

9. How does the email's claim that "Over 92% of employees have already completed the verification" attempt to pressure the recipient into responding to the email?
   (A) By framing the request as routine HR communication that seems safe, showing the use of familiarity
   (B) By implying that the opportunity to comply is limited and may soon be lost, showing the use of scarcity
   (C) By presenting the request as coming from those in power and intimidating the recipient, showing the use of authority
   (D) By applying social pressure and making the recipient believe most coworkers have already acted, showing the use of consensus

10. Which of the following is the most likely impact if the recipient responds to this email with the requested information?
   (A) An adversary could use the recipient's personal details to impersonate them.
   (B) An adversary could use the recipient's device as a foothold and move laterally on the network compromising other devices.
   (C) An adversary could infect the recipient's device with malicious code hidden in the reply email.
   (D) An adversary could disrupt the entire payroll system by disabling access for all employees.

AP CYBERSECURITY • EXERCISE SHEET

1.2 Suspicious Website Logins

Name: _____ Class: _____ Date: _____ **Total: 11 marks**

KEY WORDS authentication 身份验证 • password attack 密码攻击 • mfa 多因素身份验证 • password manager 密码 • dictionary 字典

Objective

Build the skills to answer exam questions on **password attacks** 密码攻击 and stronger **authentication** 身份验证.

You must be able to:

- identify signs of an online password attack in a log
- explain why weak, patterned passwords fail
- recommend stronger authentication: long unique passwords and **MFA** 多因素身份验证

1 Worked examples

Study these first. Each one shows the method for a task used later.

■ Signs in the log

An online password attack shows many failed logins in a short time, logins at unusual hours, or logins from unknown devices.

■ Making authentication stronger

Use long, random, unique passwords (a password manager helps), avoid names and dates, and enable multifactor authentication so a stolen password alone is not enough.

2 Practice

2.1 State two signs, visible in a log, of an online password attack. [2]

2.2 Explain why a password like `Summer2024!` is weak. [2]

2.3 State two ways to make authentication stronger. [2]

3 Exam-style questions

3.1 Which log pattern most strongly suggests an online password attack? [1]

- **A** one successful login
- **B** many failed logins in a short time
- **C** a file download
- **D** a printer coming online

3.2 Which is the strongest password? [1]

- **A** your dog's name
- **B** P@ssword!
- **C** a 16-character random string
- **D** your birthday

3.3 A company wants to reduce account takeovers.

(a) Name one authentication improvement they could require. [1]

(b) Explain how it helps even if a password is stolen. [2]

Solutions

2.1 any two of: many failed logins in a short time; logins at unusual hours; logins from unknown devices.

2.2 it follows a common pattern - word + year + symbol - that an adversary's dictionary of guesses expects.

2.3 any two of: use long/random/unique passwords; use a password manager; enable MFA.

3.1 B. a burst of failed logins signals guessing.

3.2 C. long + random + unique beats any pattern.

3.3 (a) multifactor authentication / MFA. (b) MFA needs a second factor, so a stolen password alone cannot log in.

AP CYBERSECURITY • EXERCISE SHEET

1.3 Best Practices for Public Networks

Name: _____ Class: _____ Date: _____ **Total: 11 marks**

KEY WORDS evil twin 双胞胎恶意热点 • access point 接入点 • wireless attacks 无线攻击 • encrypted 加密的 • jamming 干扰攻击

Objective

Build the skills to answer exam questions on adversary types and **wireless attacks** 无线攻击 on public networks.

You must be able to:

- classify adversaries by skill and **motivation** 动机
- describe wireless attacks: **evil twin** 双胞胎恶意热点, jamming, war driving
- recommend protections: verify the SSID, prefer HTTPS, use a **VPN** 虚拟专用网络

1 Worked examples

Study these first. Each one shows the method for a task used later.

■ Wireless attacks

An evil twin copies the real network name (SSID) so victims connect to the attacker's access point. Jamming floods the air with a signal to block access (a denial of service). War driving detects networks and where their signal leaks.

■ Encryption still protects you

On an evil twin, the adversary can read unencrypted traffic - but HTTPS sites stay encrypted, so your login is safe even on a fake access point.

2 Practice

2.1 Describe an **evil twin** attack in one or two sentences. [2]

2.2 State two things a user can do to stay safe on public Wi-Fi. [2]

2.3 An attacker floods an area with a strong radio signal so no one can connect.

(a) Name this attack. [1]

(b) State the broader category of attack it belongs to. [1]

3 Exam-style questions

3.1 An adversary who buys ready-made tools online and reuses known exploits is best described as [1]

- **A** a highly skilled attacker
- **B** a low-skilled attacker
- **C** a security analyst
- **D** a system administrator

3.2 On an evil-twin network, which of your activities stays protected? [1]

- **A** visiting an HTTP site
- **B** visiting an HTTPS site
- **C** sending an unencrypted email
- **D** browsing over plain DNS

3.3 A café's Wi-Fi signal reaches the car park outside.

(a) Explain why this is a security concern. [2]

(b) State one control that reduces this exposure. [1]

Solutions

2.1 the adversary sets up a fake access point with an SSID copied from the real network, so victims connect to it and their traffic can be read.

2.2 any two of: verify the exact network name; prefer HTTPS sites; use a VPN.

2.3 (a) jamming. (b) denial of service / DoS.

3.1 B. reusing others' tools indicates low skill.

3.2 B. HTTPS encryption protects data even if traffic is intercepted.

3.3 (a) an attacker outside can detect and probe the network without entering the building. (b) control the signal strength so it stops at the walls.

AP CYBERSECURITY • EXERCISE SHEET

1.4 AI-Based Cybersecurity Attacks

Name: _____ Class: _____ Date: _____ **Total: 11 marks**

KEY WORDS deepfake 深度伪造 • phishing 钓鱼 • shared secret 共享秘密 • adversary 对手 • impersonate 冒充

Objective

Build the skills to answer exam questions on how adversaries use AI, including **deepfakes** 深度伪造 and AI-written **phishing** 钓鱼.

You must be able to:

- explain how AI augments attacks (deepfakes, flawless phishing, data poisoning)
- describe defenses: a **shared secret** 共享秘密, MFA, verifying AI output

1 Worked examples

Study these first. Each one shows the method for a task used later.

■ How AI helps attackers

AI can clone a voice or face (a deepfake) to impersonate someone on a call, and can write phishing emails in perfect, native-sounding language - removing the bad grammar that once exposed scams.

■ Defending against AI attacks

Agree a shared secret word with close contacts to verify identity, enable MFA so a cloned voice cannot log in, and never trust AI output without checking a reliable non-AI source.

2 Practice

2.1 Explain how AI makes phishing emails harder to spot. [2]

2.2 State one defense against a deepfake voice call and explain how it helps. [2]

2.3 Why should you not type sensitive data into a chatbot? [2]

3 Exam-style questions

3.1 An AI clone of a person's voice, used to impersonate them, is called a [1]

- **A** honeypot
- **B** deepfake
- **C** keylogger
- **D** baseline

3.2 Which defense best stops an adversary who has only cloned an employee's voice? [1]

- **A** a longer password
- **B** multifactor authentication
- **C** a brighter screen
- **D** a faster network

3.3 An employee receives an urgent video call from the "CEO" asking to transfer money.

(a) State what kind of attack this could be. [1]

(b) Describe one way the employee could verify the caller's identity. [2]

Solutions

2.1 AI writes in perfect, native-sounding language, removing the poor grammar that used to reveal a scam.

2.2 a shared secret question only the two parties know; the deepfake cannot answer it, exposing the fake.

2.3 some AI tools feed user input back into training, and an adversary could later extract that data.

3.1 B. a deepfake clones a voice or face.

3.2 B. a second factor blocks a voice-only attack.

3.3 (a) an AI deepfake impersonation. (b) ask a prearranged shared-secret question that only the real CEO would know; or verify through a separate known channel.

AP CYBERSECURITY • EXERCISE SHEET

1.5 Leveraging AI in Cyber Defense

Name: _____ Class: _____ Date: _____ **Total: 11 marks**

Objective

Build the skills to answer exam questions on how defenders use AI to protect networks, applications, and data.

You must be able to:

- explain how AI recommends safer configurations (with human review)
- explain how AI handles the scale of network events for faster detection

1 Worked examples

Study these first. Each one shows the method for a task used later.

■ AI at scale

A network produces millions of events daily - too many for people to read. AI sorts likely-malicious events from harmless ones, alerts staff, or acts automatically, cutting response from days to seconds.

■ Human in the loop

AI advice is a recommendation, not a command. A knowledgeable human must review a suggested firewall rule or code fix before it is applied, because AI is probabilistic.

2 Practice

2.1 Explain why AI is useful for analysing network logs. [2]

2.2 Give one benefit and one limitation of AI-powered threat detection. [2]

2.3 Why must a human review an AI's security recommendation before applying it? [2]

3 Exam-style questions

3.1 The main advantage of AI for threat detection is that it can [1]

- **A** remove all false alarms
- **B** respond in seconds instead of days
- **C** replace every analyst
- **D** delete the logs

3.2 Which task is best left to a human, not the AI? [1]

- **A** sorting millions of events
- **B** flagging likely threats
- **C** making the final risk-acceptance decision
- **D** matching known signatures

3.3 A hospital's network logs 20 million events a day.

(a) Explain why human staff alone cannot review all these events. [1]

(b) Describe how AI helps and what role the human keeps. [2]

Solutions

2.1 a network logs millions of events daily that humans cannot read, but AI can sort them quickly.

2.2 benefit: responds in seconds not days; limitation: it is probabilistic and needs human review.

2.3 AI guesses are probabilistic and can be wrong; a bad rule could block real users or miss a threat.

3.1 B. speed of response is AI's key advantage.

3.2 C. the final risk decision is a human/policy call.

3.3 (a) the volume is far too large for people to read. (b) AI sorts and flags the likely-malicious events; a human reviews and decides on the flagged cases.