

Lab Guide

Four exercises: target profiling for awareness training, phishing-email forensics, lookalike-domain detection, and email-header analysis — detect and defend against social engineering with AI.

Authorized Use Only & Defensive Framing

This lesson frames all exercises from the **defensive and awareness-training perspective**. The goal is to analyze phishing campaigns that have already been sent, so defenders can detect, train, and prevent future incidents. All target profiles, emails, and headers are **synthetic**. Never send phishing messages to real people without documented authorization for a sanctioned engagement.

Prerequisites

- Completed the [Theory](#) reading
- Access to an AI assistant (Claude, ChatGPT, or local LLM)
- Downloaded all four files into a folder named `phishing-lab/`
- Familiarity with: SMTP headers, SPF/DKIM/DMARC, basic OSINT, MITRE ATT&CK Initial Access
- Optional: `dnstwist`, `urlscan.io`, `PhishTool` for cross-checking

Lab Scenario

AcmeCorp's SOC flagged a phishing wave that hit multiple departments. You have the inbound emails, the targeted employee profiles, the suspicious domains involved, and the raw email headers from the gateway. Your job: figure out how the attacker profiled targets, how sophisticated the campaign was, which controls failed, and what awareness training would prevent the next round.

Download Lab Files

Save all four files into a folder named `phishing-lab/` before starting the exercises.

↓ `target-profiles.txt` ~14 KB

↓ `phishing-emails.txt` ~10 KB

↓ `domain-analysis.txt` ~14 KB

↓ `email-headers.txt` ~13 KB

Exercise 1 — Target Profile Analysis (Defensive)

Open `target-profiles.txt`. These are the OSINT profiles the attacker assembled before the campaign.

Task 1.1 — What made each target vulnerable?

You are a security awareness trainer reviewing employee profiles used in a phishing campaign that ALREADY HAPPENED. For each profile, identify:

1. What public information the attacker used to personalize the lure (job title, department, recent activity, interests).
2. Which details SHOULD NOT have been public (and where they leaked from).
3. The pretext category this person was most vulnerable to:
 - Authority (boss / IT / vendor)
 - Urgency (deadline, limited offer, security alert)
 - Scarcity (last chance, limited seats)
 - Social proof (colleagues did this)
 - Reciprocity (favor, gift)
 - Curiosity (leaked doc, scandal)
4. What awareness-training module would best protect this person (executives / finance / engineers / new hires / etc.)?
5. Three specific things this employee should remove or lock down on their public profiles.

Output as a per-person training worksheet.

[paste target-profiles.txt]

✓ **Self-check: You should end with concrete, actionable guidance for each employee — not generic "be careful" advice.**

Exercise 2 — Phishing Email Forensics

Open `phishing-emails.txt`. These are the messages that actually reached inboxes.

Task 2.1 — Dissect the campaign

Below are phishing emails that landed in employee inboxes.

For each message, identify:

1. The pretext (what is the attacker pretending to be?).
2. The call-to-action (click link / open attachment / reply with info / call number).
3. The psychological trigger (urgency / authority / fear / reward).
4. Linguistic red flags (typos, unusual phrasing, generic greetings, inconsistent tone, machine-translation artifacts).
5. Technical red flags (display name spoofing, reply-to mismatch, suspicious link domain, encoded URLs, tracking pixels, attachment types).
6. The MITRE ATT&CK technique (T1566.001 attachment / T1566.002 link / T1566.003 via service).
7. The likely goal (credential phish / malware / BEC / recon).

Then rank the emails by sophistication and note which would most likely bypass both user skepticism and gateway filters.

[paste phishing-emails.txt]

✓ **Self-check: You should be able to classify every email with a specific MITRE technique and name at least two red flags the recipient should have caught.**

Exercise 3 — Lookalike Domain Detection

Open `domain-analysis.txt`. Typo-squatting and homograph domains are the infrastructure of phishing.

Task 3.1 — Catalog the attacker infrastructure

Below is a list of domains extracted from a phishing wave.

For each domain:

1. Identify the legitimate brand being spoofed.
2. Classify the lookalike technique:
 - Typo-squat (missing letter, swapped letters)
 - Homograph (Cyrillic / Unicode lookalikes)
 - TLD swap (.com -> .co / .net)
 - Subdomain spoof (acmecorp.com.attacker.xyz)
 - Combo-squat (acmecorp-login.com)
 - Brand-in-path (attacker.com/acmecorp)
3. Registration red flags:
 - Domain age (under 30 days = very suspicious)
 - Privacy-protected registrant
 - Cheap registrar with anonymous purchase
 - Registered with hosting on bulletproof providers
4. DNS / MX indicators suggesting it is configured to RECEIVE credentials (A record + HTTPS, MX record = BEC, TXT with SPF include of a known abuse provider).
5. A blocklist recommendation: which domains should be blackholed at the corporate DNS, and with what rule?

Output as a defender triage table with a "block / monitor / ignore" verdict per domain.

[paste domain-analysis.txt]

✓ Self-check: You should end with a list of domains to add to a DNS blocklist and a reasoning snippet for each that a SOC analyst could paste into a ticket.

Exercise 4 — Email Header Analysis

Open `email-headers.txt`. Headers tell the story the message body tries to hide.

Task 4.1 — Trace the real sender

Below are full email headers. For each message:

1. Trace the Received: chain from origin to destination.
Identify the true origin IP (first external Received:).
2. Check authentication results:
 - SPF: pass / fail / softfail / neutral / none
 - DKIM: pass / fail / none, signer domain
 - DMARC: pass / fail / policy (none / quarantine / reject)
3. Compare the From: display name vs. the return-path / envelope sender vs. the reply-to.
4. Identify any Received: header that is clearly forged (out-of-order timestamps, inconsistent IPs, fake hostnames).
5. Determine whether the message was:
 - Legitimate
 - Forged but aligned (authenticated but socially engineered)
 - Spoofed (failed auth, gateway let it through anyway)
 - Relayed through a compromised third-party server
6. Identify the upstream mail infrastructure (provider, ASN, country) for blocklist correlation.
7. Give a verdict per message: allow / quarantine / block / alert.

Finally: list 3 policy changes the AcmeCorp email admin should make to prevent this class of message from reaching users.

[paste email-headers.txt]

✓ **Self-check:** You should be able to point at a specific header line and say "this is where the attacker's real IP leaks" for at least one message.

Final Exercise — Campaign Reconstruction & Training Plan

Task 5.1 — Connect all four sources

Here are four analyses of the same phishing wave (target profiles, email bodies, lookalike domains, raw headers). Reconstruct the campaign as a defender would for an after-action review:

1. Attacker profile: capability level, likely motive, possible attribution (nation-state / criminal / hacktivist / insider).
2. Target selection logic: how did the attacker pick these employees?
3. Infrastructure map: domains, IPs, ASNs, and the relationships between them.
4. The most successful variant in the campaign and why.
5. Which technical controls failed (SPF / DKIM / DMARC / gateway rules / DNS filter) and which held.

Then produce:

- A one-page incident report for the CISO.
- A 5-point awareness-training outline for employees.
- A 5-point technical hardening checklist for the email admin.
- A detection rule (Sigma or Sublime) that would catch the next variant of this campaign.

Format as markdown, max 1500 words.

✓ Self-check: Your report should connect at least 3 of the 4 sources. The best campaign reconstructions tie target selection -> lure content -> sending infrastructure -> authentication failure into a single story.

Reflection

- Where does AI help most in phishing defense: detecting bad emails, training users, or reconstructing campaigns?
- What information in an employee's public profile was most dangerous? What is reasonable to remove vs. what is operationally required?
- How do you distinguish a sophisticated targeted phish from a mass-market one? What signals scale and which do not?
- If you were the AcmeCorp CISO, what is the single biggest change you would make after this incident?

What's Next

Lesson 10 is the **capstone** — combine everything from recon through exploitation to social engineering into a single end-to-end red team engagement walkthrough.