

Answers

Reference answers for the lab exercises and homework tasks. Try the work first — this page is for grading yourself, not shortcutting.

These answers describe what a *good* AI-assisted run should produce. Your own outputs may differ slightly because AI responses are non-deterministic. What matters is whether you hit the key findings, not whether the wording matches.

Part 1 — Lab Answers

Exercise 1 — WHOIS Triage

Expected findings

- **Primary domain:** `acmecorp.com` stands out by creation date, registrar, and unique NS records shared with `acmecorp-internal.com` and `acme-staging.net`.
- **Shared infrastructure:** at least two side-domains share the same name servers or registrant email, tying them to AcmeCorp.
- **Contact leak:** the technical contact email (format `firstinitial.lastname@acmecorp.com`) gives you the internal username convention for the entire org.
- **Privacy posture:** the primary domain has privacy protection; the staging domain does not — classic misconfiguration.

Common mistakes

Treating expired or parked domains as active, or missing the fact that two domains are related because they share NS records even though registrant info is hidden.

Exercise 2 — DNS & Subdomain Mapping

Expected findings

- **High-priority targets:** anything matching `vpn-*`, `admin-*`, `staging-*`, `dev-*`, `jenkins.*`, or `jira.*`.
- **MX leak:** the mail provider (Google Workspace / Office 365) reveals whether password-spray should target Azure AD or Google endpoints.
- **SPF / DKIM / DMARC:** a permissive SPF (`~all` soft-fail) or missing DMARC means spoofing is viable.
- **TXT leakage:** SaaS verification tokens reveal which third parties AcmeCorp uses (Salesforce, Atlassian, Mailchimp).

The top two candidates

Expect the AI to flag a dev/staging environment and an admin panel as the top two. Both are gold because they are usually under-hardened versions of production.

Exercise 3 — Certificate Transparency Mining

Expected findings

- At least one subdomain from CT logs that is NOT in DNS — usually an internal-sounding name like `internal-api.acmecorp.com` or `vault.acmecorp.com` with a recent certificate issuance.
- A wildcard cert covering `*.dev.acmecorp.com` that implies an entire dev environment exists.

- Certificates issued in the last 90 days indicating active services (and therefore worth probing).

The delta

The CT-only subdomain is the single highest-value finding in this whole exercise because it was deliberately hidden but still leaked through a CA log.

Exercise 4 — Social & People Recon

Expected findings

- **Top three high-value targets:** IT Director, DevOps Lead, and CFO (or Finance Director). Each has 2-3 plausible username formats based on the email pattern inferred from WHOIS.
- **Tech-stack leakage:** engineers' bios usually leak the framework (React, Django), the cloud (AWS / GCP), and sometimes specific internal project names.
- **Phishing angles:** the CFO is vulnerable to BEC/authority pretexts; the DevOps Lead to "urgent AWS alert" / fake vendor tickets.

Pretext quality check

A good AI answer names a specific defensive control for each pretext (e.g., "DMARC reject + user awareness training" for display-name spoofing). Generic "be careful" advice is a sign of a lazy answer.

Exercise 5 — Paste-Dump Triage

Expected findings

- At least one `username:password` pair or API key that an attacker could try against an asset you already identified in Exercises 2-3.

- Internal hostname or IP that confirms infrastructure you inferred from WHOIS/DNS.
- A source-code snippet revealing either a hard-coded secret or a connection string.
- A reference to an internal project codename that gives you pretext material for Task 4.2.

The #1 finding

A credential that maps to an identified endpoint. "A leaked password with nowhere to use it" is worth far less than "a password for a VPN portal we found in DNS" — the correlation is the whole point of paste-dump triage.

Final Exercise — The Recon Report

The #1 attack path

In a well-executed synthesis, the #1 path should combine at least three sources: a leaked credential from the paste dump + a subdomain from CT logs + the username format from WHOIS. If your top path only uses one source, you are thinking like a scanner instead of an attacker.

Report structure check

Executive summary: three sentences max. Top 3 recommendations: each must name a specific *asset*, a specific *technique*, and a specific *expected outcome*. Vague recommendations mean the report is not actionable.

Part 2 — Homework Answers

Task 1 — Reusable OSINT Prompt Library

Expected output

A markdown file with five H2 sections (one per source). Each prompt should:

- Use a `{{target_domain}}` style placeholder, not a hard-coded name.
- Specify the output format explicitly (markdown table, JSON, bullet list).
- Include a "confidence" clause asking the AI to mark low-confidence findings.
- Include a "when to use" one-liner.

Grading signal

A good library is short and reusable. If your prompts exceed 20 lines each, you are probably over-specifying.

Task 2 — AI vs. Manual Benchmark

Expected outcome

- AI will usually beat manual on: exhaustive extraction, format conversion, correlation across similar records, base64/hex decoding.
- Manual will usually beat AI on: catching subtle "this does not fit the pattern" anomalies, applying engagement-specific context (e.g., "we know the client uses Okta, so this CNAME matters"), and noticing deceptive data (planted honeypots).

Typical pitfalls

The AI sometimes hallucinates plausible-but-wrong findings (e.g., a CVE that does not apply to the version you actually have). Always verify high-value AI claims with a second source.

Task 3 — Defender's Mirror Exercise

Expected remediation ranking (example shape)

1. **Rotate leaked credential + enforce password manager** — low effort, immediate blast-radius reduction.
2. **Take down the dev subdomain or put it behind VPN** — moderate effort, removes an entire attack surface.
3. **Enforce DMARC reject** — moderate effort, blocks the phishing pretexts from Task 4.2.
4. **Remove technical contact email from WHOIS (privacy protection)** — low effort but only cosmetic once leaked.
5. **Enterprise-wide LinkedIn hygiene policy** — high effort, low ROI — belongs in the backlog, not sprint 1.

A common mistake is ranking by severity instead of ROI. "Training the whole company about LinkedIn" is high severity but has poor ROI; "rotate the one leaked password" has better ROI.