

Answers

Reference answers for Lesson 10 capstone.

Part 1 — Lab Answers

Exercise 1 — Scope Document Review

Expected briefing

Clear lists of in-scope and out-of-scope assets. A flagged ambiguity — scope documents always have one. Common ambiguities: whether subsidiaries count, whether physical tailgating is permitted, whether production DBs can be read.

Exercise 2 — Recon Data Review

Expected output

A prioritized target list where every entry references a specific technique from Lessons 01-09 and a specific MITRE ATT&CK Initial Access technique.

Exercise 3 — Vulnerability Scan Triage

Expected attack path

Primary path built on a HIGH or CRITICAL with a public PoC against an in-scope host. Operator checklist should be reproducible — anyone on the team can run it from the notes.

Exercise 4 — Attack Timeline Reconstruction

Expected chain

Initial access → foothold → privesc → persistence → lateral movement → objective. Each step mapped to MITRE ATT&CK. At least three defender-detection opportunities identified.

Final Exercise — Client Report

Expected report quality

CISO can read the exec summary in 10 minutes and understand the business impact. SecOps can implement the tactical recommendations for three months. Tone is direct and evidence-based, not fear-mongering.

Part 2 — Homework Answers

Task 1 — ROE Template

Expected template

Placeholders for target, scope boundaries, engagement window, contacts, and authorization letter. Red-flags checklist should include at least: scope boundary

ambiguity, executive phishing permission, destructive test permission, production data handling, blackout windows.

Task 2 — Blue Team Debrief

Expected gap ranking

Initial access failures are worth the most (every subsequent step is a consequence). Persistence and lateral movement detections are second. "Exfil detection" is last because by then the damage is done.

Task 3 — AI Retrospective

Typical conclusions

AI excels at: exhaustive extraction, format conversion, pattern matching on known vulnerability classes, boilerplate generation. AI struggles with: novel bugs, business-logic flaws, engagement-specific context, adversarial reasoning ("how will defenders react"). A good personal policy distinguishes "AI drafts, human validates" from "human drafts, AI reviews" based on task risk.