

Homework — Two Scenario Walkthroughs

~25 minutes · one red scenario, one blue scenario

Two scenarios. Work through both before looking at the answer key. Write your reasoning down; the point of the exercise is the reasoning, not the conclusion.

Do not open [answer.html](#) until you have written out your own analysis for both scenarios. The exercise is worthless if you peek early.

Red scenario — mid-engagement triage

You're on day 2 of a red team engagement. You've captured an NTLM hash, a partial Kerberos ticket, and a list of internal hostnames. You want AI help to identify which hash type is which and to draft the next-step plan.

- Which of these assets can you paste into an AI, and which can't? Why?
- What's the safe workflow for each asset — specifically, what do you do instead of pasting?
- How do you get AI help on the next-step plan without leaking engagement-specific context?
- If your personal account already has *any* engagement content in its history, what does that mean for the account going forward?

✓ **Self-check: answer before looking at the key.**

Blue scenario — DLP telemetry triage

You're reviewing last week's DLP telemetry. Three engineers show up in the AI-provider egress rules:

- **Engineer A** uploaded 12 files to chat.openai.com on Tuesday between 14:00 and 14:30.
- **Engineer B** pasted 8 code blocks into claude.ai the same day over 6 hours.
- **Engineer C** submitted one 40-line prompt to gemini.google.com.

Work through:

- Which is most concerning and why?
- What do you do next for each?
- Which signals in the data — volume, burst pattern, content class, time-of-day — are carrying your decision?
- At what point does a finding escalate from a soft conversation to a formal incident?

✓ **Self-check: answer before looking at the key.**