LAB 12 - WEEK 14 - 50 POINTS

Tabletop Incident Response: Guide & Worksheet

CIS 340: AI for Cybersecurity - Smart Defense for the Digital Business

Colorado State University - Computer Information Systems

What this is (and a note on scope)

A tabletop is a TALK-THROUGH crisis simulation - no real systems are touched and there are no trick questions. Working in a team, you respond to a simulated security incident involving an AI/ML system, using a scenario adapted from CISA's free Tabletop Exercise Packages (CTEP). You are graded on participation and the quality of your reasoning, not on getting a single 'right' answer. AI tools are not permitted for the write-up.

Team name: _____ Date: _____ Members: _____

1. What This Lab Is

You and your team are the incident-response group for a mid-size company. An AI/ML system starts misbehaving during a busy day, and you must work through the crisis together - detecting, deciding, containing, recovering, and communicating - then debrief on what you learned. This builds a job-relevant skill: calm, structured response under pressure, with AI tools assisting and humans deciding.

Work in teams of 4-6. Time: about 60-90 minutes (a compressed version runs in Wednesday's class). Submit THIS completed worksheet in Canvas - there is no code to hand in.

Learning objectives

- Apply the incident-response (IR) lifecycle: prepare; detect & analyze; contain, eradicate, recover; post-incident.
- Make and justify decisions under time pressure as a team.
- See how AI incidents differ (faults hide in data/models; failures can be silent).
- Connect the response to the four trustworthy-AI pillars (esp. robustness & accountability).

Setup

- Form a team of 4-6 and assign the five roles in the table below.
- Scenario: an AI-powered fraud detector begins behaving strangely (adapted from CISA CTEP).
- Your instructor (facilitator) will reveal 'injects' - new developments - each round.
- The Scribe records decisions and reasoning here as you go.

CISA Tabletop Exercise Packages (free): https://www.cisa.gov/resources-tools/services/cisa-tabletop-exercise-packages

2. Your Team & Roles

Assign each role to a person (one person may hold two on a small team).

Role	Team member	Responsibility
Incident Lead		runs the response, makes the call
Comms		customers, leadership, regulators
Tech / Analyst		investigates the model & data
Legal / Compliance		law, privacy, disclosure duties
Scribe		logs decisions & timeline

3. The Scenario

Your company uses an AI fraud-detection model to approve or block transactions. This morning, on a busy sales day, the model's behavior changed sharply. You are called into the incident bridge. Work each round below as a team, following the IR lifecycle. Record your decisions AND your reasoning - the

reasoning is where the credit is.

Ground rules

Discuss and decide together. Justify every decision. There are no trick questions and no single correct answer - responders are judged on sound reasoning and clear communication. Stay in role.

4. Round-by-Round Decision Log

Round 1 - Detect & Analyze

Inject: the model's block rate jumped 20x this morning. Is it an attack, a bug, data drift, or a false alarm? How do you tell? What severity, and who do you notify first? What do you NOT do yet?

Decision & who owns it:

Why (your reasoning / trade-offs):

Round 2 - Contain & Decide

Inject: logs show the model was retrained overnight on new data, and real fraud is now slipping through. Do you roll back, add human review, or something else? What are the trade-offs (blocking real customers vs. fraud losses)? Who approves a high-impact action?

Decision & who owns it:

Why (your reasoning / trade-offs):

Round 3 - Recover & Communicate

Inject: a reporter and customers are asking what happened. What is your safe recovery plan? What do you tell customers, leadership, and regulators, and when? What do you preserve for the post-incident review?

Decision & who owns it:

Why (your reasoning / trade-offs):

5. Debrief (Hotwash)

The most valuable part of any exercise. Answer as a team, in complete sentences.

Q1. What went well in your response, and what was confusing or too slow?

Q2. Where did AI HELP (e.g., fast detection), and where did HUMANS have to decide? Give a specific example of each.

Q3. Which trustworthy-AI pillar (transparency, explainability, robustness, accountability) would most have prevented or eased this incident, and why?

Q4. What ONE change to your PREPARE phase would most improve your response next time?

Q5. What was the likely root cause of the incident (tie it to a Week 10-13 concept), and how would you confirm it?

6. Submission & Grading

- Submit this completed worksheet (roles, all three decision logs, and the five debrief answers) in Canvas.
- You are graded on PARTICIPATION and the QUALITY of your reasoning - not on a 'correct' answer.
- See rubric.md for the 50-point breakdown. Late work: -10%/day up to 48 hours (syllabus).
- AI tools are not permitted for the write-up.

How this connects to the course

This lab turns Week 14's four pillars into practice and pulls together the whole Exam-3 unit: the root cause is usually a Week 10-13 idea (poisoning, drift, a privacy or fairness failure), and a good response depends on explainability, robustness, and accountability. AI assists; humans decide.