

Blue Team Capstone: Full Defense Scenario

The culmination of everything you have learned — bringing together monitoring, detection, response, forensics, and hardening into a single integrated defense operation, augmented by AI at every step.

What We Cover

Full lifecycle defense against a multi-stage attack using every tool and technique from Lessons 1-9

Capstone Challenge

Defend a live network, detect intrusions, respond, recover, and write the final report

Defense-in-Depth

SOC Workflow

AI-Augmented

STRIDE/DREAD

Incident Response

Forensics

Educational Use Disclaimer

Educational Use Only

This capstone combines all defensive techniques from the course. These skills are for authorized security operations and CTF competitions. Always follow your organization's security policies and procedures.

Everything in this lesson assumes you are operating within a sanctioned environment — your own lab, an authorized engagement, or a CTF competition. Unauthorized use of these techniques against systems you do not own is illegal.

Course Recap: Lessons 1-9

Lesson	Core Skill	Capstone Role
1 — SOC Fundamentals	Alert triage & SIEM basics	Initial detection layer
2 — Log Analysis	Pattern recognition in logs	Evidence collection
3 — Network Defense	Traffic analysis & IDS	Network monitoring
4 — Endpoint Security	Host-based detection	Endpoint protection
5 — Threat Intelligence	IOC feeds & enrichment	Contextual awareness
6 — Incident Response	Containment & eradication	Active response
7 — Digital Forensics	Evidence preservation & analysis	Post-incident analysis
8 — AI-Augmented Defense	LLM-powered analysis	AI at every stage

Lesson	Core Skill	Capstone Role
9 — Hardening & Prevention	Configuration & patching	Post-incident hardening


Defense-in-Depth Architecture

No single control stops a determined attacker. Defense-in-depth layers multiple controls so that failure at one layer is caught by the next.

1. **Perimeter** — Firewalls, WAF, DDoS mitigation, DNS filtering
2. **Network** — IDS/IPS, network segmentation, traffic analysis, NetFlow monitoring
3. **Endpoint** — EDR, host firewalls, application whitelisting, OS hardening
4. **Application** — Input validation, authentication, authorization, secure coding
5. **Data** — Encryption at rest/in transit, DLP, access controls, backup verification
6. **Human** — Security awareness training, phishing simulations, access reviews
7. **AI Layer** — LLM-assisted anomaly correlation, automated triage, predictive analysis

Notes: The AI layer is not a replacement for any other layer — it enhances every layer by accelerating analyst decision-making and surfacing patterns humans might miss.

Defense-in-Depth: Architecture Diagram



Each layer feeds telemetry into the centralized SIEM, where the AI correlation engine connects dots across layers in real time.

SOC Workflow with AI Augmentation

A modern Security Operations Center operates in a continuous loop. AI enhances each phase without replacing human judgment.

Traditional SOC

- Analyst manually triages alerts
- Playbook-driven response
- High alert fatigue
- Hours to correlate events
- Reactive posture

AI-Augmented SOC

- AI pre-triages and prioritizes
- Dynamic playbook suggestions
- Reduced false positive noise
- Seconds to correlate events
- Proactive threat hunting


SOC Tier Model & AI Integration

SOC Tier	Role	AI Augmentation
Tier 1 — Alert Analyst	Initial triage, ticket creation	Auto-classification, severity scoring, context enrichment
Tier 2 — Incident Responder	Deep investigation, containment	Attack chain reconstruction, IOC correlation, suggested actions
Tier 3 — Threat Hunter	Proactive hunting, tool development	Anomaly detection, hypothesis generation, query building
SOC Manager	Metrics, reporting, staffing	Automated KPI dashboards, trend analysis, executive summaries

Notes: AI does not eliminate tiers — it elevates each tier's capability. A Tier 1 analyst with AI assistance can perform at near Tier 2 level for routine incidents.

The Full Defense Lifecycle

Every defense scenario follows this cycle. The capstone challenge will test your ability to execute each phase under pressure.



1. **Monitor** — Collect and aggregate telemetry from all layers
 2. **Detect** — Identify anomalies and confirmed threats via rules + AI
 3. **Respond** — Contain, eradicate, and recover from the threat
 4. **Forensics** — Preserve evidence, analyze root cause, build timeline
 5. **Harden** — Patch, reconfigure, add detection rules, update policies
-

Scenario Introduction: "Operation Shadow Breach"

Scenario Overview

You are the lead blue team analyst for Nexus Financial Corp. At 02:47 UTC, your SIEM fires a cluster of alerts. Over the next several hours, a sophisticated threat actor executes a multi-stage attack across your infrastructure. Your mission: detect every phase, contain the damage, preserve evidence, and harden the environment — using AI assistance at every step.

Target Environment

- **DMZ** — Web servers (nginx), mail gateway, public DNS
- **Internal** — Active Directory domain, file servers, database cluster
- **Cloud** — AWS S3 buckets, Lambda functions, RDS instances

- **Endpoints** — 250 Windows workstations, 30 Linux servers
-

Phase 1: Monitoring — Baseline & Collection

Before you can detect anomalies, you need to know what "normal" looks like. Your monitoring stack must cover every layer.

Log Sources for the Scenario

```
# Monitoring stack configuration
log_sources:
  network:
    - zeek_conn_logs      # Connection metadata
    - suricata_alerts     # IDS signature matches
    - netflow_v9          # Traffic flow records
    - dns_query_logs      # DNS resolution tracking
  endpoint:
    - windows_security_log # Event IDs 4624, 4625, 4688, etc.
    - sysmon_events        # Process creation, network, file
    - linux_auditd         # Syscall auditing
    - edr_telemetry        # CrowdStrike/Defender events
  application:
    - nginx_access_log     # Web request logs
    - nginx_error_log      # Server errors
```

AI-Powered Baseline Monitoring

Use an LLM to analyze your baseline and flag deviations before traditional rules catch them.

```
import anthropic

def ai_baseline_analysis(current_metrics, baseline_stats):
    """Use AI to compare current metrics against established baselines."""
    client = anthropic.Anthropic()

    prompt = f"""Analyze these current network metrics against our baseline.
Flag anything anomalous and explain why it matters.

BASELINE (30-day average):
{baseline_stats}

CURRENT (last 15 minutes):
{current_metrics}"""
```

Notes: This approach catches "low and slow" attacks that operate within normal thresholds individually but form suspicious patterns collectively.

Phase 2: Detection — The Attack Begins

02:47 UTC — Initial Alerts Fire

SIEM Alert Cluster

- Alert 1:** Suricata — ET EXPLOIT Apache Struts RCE Attempt (10.0.1.50)
- Alert 2:** Suricata — ET EXPLOIT Apache Struts RCE Attempt (10.0.1.50)
- Alert 3:** Sysmon — Suspicious PowerShell Encoded Command (WS-042)
- Alert 4:** Zeek — Unusual DNS TXT Query Volume (internal → external)
- Alert 5:** CloudTrail — AssumeRole from unusual IP (198.51.100.47)

Five alerts in 3 minutes. A Tier 1 analyst might treat these as separate events. AI correlation reveals they are one attack with multiple fronts.

AI-Powered Alert Correlation

```
def correlate_alerts_with_ai(alerts):  
    """Feed multiple alerts to AI for cross-layer correlation."""  
    client = anthropic.Anthropic()  
  
    alert_text = "\n".join([  
        f"[{a['timestamp']}] {a['source']}: {a['description']} "  
        f"(src={a['src_ip']}, dst={a['dst_ip']})"  
        for a in alerts  
    ])  
  
    response = client.messages.create(  
        model="claude-sonnet-4-20250514",  
        max_tokens=3000,  
        messages=[{"role": "user", "content": f"""You are a SOC analyst.  
Correlate these alerts and determine if they represent:
```

The AI identifies this as a coordinated attack: web exploitation + lateral movement + data staging
+ cloud pivot. Confidence: 94%.

MITRE ATT&CK Mapping of the Attack

Tactic	Technique	Evidence
Initial Access	T1190 — Exploit Public-Facing App	Struts RCE attempts on web server
Execution	T1059.001 — PowerShell	Encoded command on WS-042
Persistence	T1053.005 — Scheduled Task	New task on compromised host
Defense Evasion	T1027 — Obfuscated Files	Base64 PowerShell, packed binaries
Lateral Movement	T1021.002 — SMB/Admin Shares	Unusual SMB connections
Collection	T1560 — Archive Collected Data	7zip activity on file server
Exfiltration	T1048.001 — Exfil Over DNS	High-volume DNS TXT queries
Cloud	T1078.004 — Cloud Accounts	AssumeRole from attacker IP

Phase 3: Response — Containment

With the attack confirmed, containment must be swift and surgical. Contain too broadly and you disrupt business; too narrowly and the attacker pivots.

Immediate Containment Actions

1. **Network isolation** — Move compromised web server (10.0.1.50) to quarantine VLAN
2. **Block attacker IPs** — 198.51.100.47 and related C2 infrastructure at perimeter firewall
3. **DNS sinkhole** — Redirect exfiltration domain to internal sinkhole
4. **Credential reset** — Force password change for accounts on WS-042 and any AD accounts accessed
5. **Cloud lockdown** — Revoke compromised IAM role sessions, rotate access keys

```
# Emergency firewall block
sudo iptables -I INPUT -s 198.51.100.47 -j DROP
sudo iptables -I OUTPUT -d 198.51.100.47 -j DROP

# DNS sinkhole entry
echo "198.51.100.47 exfil-c2.evil.example" >> /etc/hosts
```


```
# AWS session revocation
aws iam put-role-policy --role-name CompromisedRole \
  --policy-name EmergencyDeny \
  --policy-document '{"Version":"2012-10-17","Statement":[{"Effect":"Deny","Action":"*","Resource":"*"}]}'
```

AI-Assisted Incident Response Playbook

```
def generate_response_playbook(incident_summary, environment_context):  
    """Generate a tailored IR playbook using AI."""  
    client = anthropic.Anthropic()  
  
    response = client.messages.create(  
        model="claude-sonnet-4-20250514",  
        max_tokens=4000,  
        messages=[{"role": "user", "content": f"""As a senior incident responder,  
generate a step-by-step playbook for this incident.  
  
INCIDENT SUMMARY:  
{incident_summary}  
  
ENVIRONMENT:  
{environment_context}"""}]  
    )
```

Why AI Playbooks?

Static playbooks cannot anticipate every scenario. AI generates context-specific steps based on the exact attack chain, environment topology, and available tools — then updates in real time as the situation evolves.



Phase 3 (cont.): Eradication

After containment stops the bleeding, eradication removes every trace of the attacker from the environment.

Eradication Checklist

- **Malware removal** — Identify and quarantine all dropped binaries, scripts, webshells
- **Persistence removal** — Delete scheduled tasks, registry run keys, cron jobs, startup items
- **Backdoor account cleanup** — Audit all accounts created or modified during the attack window
- **Patch the vulnerability** — Apply Apache Struts patch that enabled initial access
- **Cloud cleanup** — Remove unauthorized IAM policies, Lambda functions, S3 bucket policies

```
# Find and remove attacker scheduled tasks (Windows)
Get-ScheduledTask | Where-Object {
    $_.Date -gt "2026-04-08" -and
    $_.Author -notmatch "SYSTEM|Microsoft"
} | ForEach-Object {
    Write-Host "Removing suspicious task: $($_.TaskName)"
}
```

```
Unregister-ScheduledTask -TaskName $_.TaskName -Confirm:$false  
}
```

Phase 4: Forensics — Evidence & Timeline

Forensics runs in parallel with response. The golden rule: **preserve first, analyze second**.

Evidence Collection Priority

Priority	Evidence Type	Volatility	Collection Method
1	Memory dumps	Highest	<code>winpmem</code> / <code>LiME</code>
2	Network captures	High	Full PCAP from tap/span
3	Running processes	High	<code>volatility3</code> analysis
4	Log files	Medium	SIEM export + raw copy
5	Disk images	Low	<code>dd</code> / FTK Imager
6	Cloud logs	Low	CloudTrail S3 export

Notes: Always calculate and record hashes (SHA-256) of forensic images. Maintain chain of custody

documentation for any evidence that may be used in legal proceedings.

AI-Assisted Timeline Reconstruction

```
def build_attack_timeline(log_entries, memory_artifacts, network_captures):  
    """Use AI to reconstruct a unified attack timeline."""  
    client = anthropic.Anthropic()  
  
    evidence = f"""LOG ENTRIES (sorted by time):  
{log_entries}  
  
MEMORY ARTIFACTS:  
{memory_artifacts}  
  
NETWORK CAPTURES (summary):  
{network_captures}"""  
  
    response = client.messages.create(  
        model="claude-sonnet-4-20250514",
```

Scenario Timeline: Operation Shadow Breach

Time (UTC)	Event	Layer
02:31	Reconnaissance — port scan from 198.51.100.47	Network
02:38	Vulnerability scan targeting Apache Struts	Application
02:44	Successful RCE exploit on web server 10.0.1.50	Application
02:45	Reverse shell established, C2 beacon activated	Network
02:47	First SIEM alerts fire (detection point)	SOC
02:52	Credential harvesting via Mimikatz on web server	Endpoint
03:01	Lateral movement to WS-042 via stolen credentials	Endpoint
03:08	Pivot to file server, data staging begins	Endpoint
03:15	DNS exfiltration channel established	Network

Time (UTC)	Event	Layer
03:22	Cloud pivot — AssumeRole with stolen keys	Cloud
03:30	Containment actions initiated by blue team	SOC

Dwell time before detection: 3 minutes. **Dwell time before containment:** 46 minutes.

Threat Modeling: STRIDE & DREAD

After the incident, use threat modeling to understand what went wrong and what else could go wrong.

STRIDE (Threat Categories)

- **S**poofing — Identity falsification
- **T**ampering — Data modification
- **R**epudiation — Denying actions
- **I**nformation Disclosure — Data leaks
- **D**enial of Service — Availability attacks
- **E**levation of Privilege — Unauthorized access

DREAD (Risk Scoring)

- **D**amage — How severe is the impact?
- **R**eproducibility — How easy to replicate?
- **E**xploitability — How easy to exploit?
- **A**ffected Users — How many impacted?
- **D**iscoverability — How easy to find?

AI-Assisted Threat Modeling

```
def ai_threat_model(system_description, known_vulnerabilities):  
    """Use AI to perform STRIDE/DREAD analysis."""  
    client = anthropic.Anthropic()  
  
    response = client.messages.create(  
        model="claude-sonnet-4-20250514",  
        max_tokens=4000,  
        messages=[{"role": "user", "content": f"""Perform a threat model for this system.  
  
SYSTEM:  
{system_description}  
  
KNOWN VULNERABILITIES / RECENT INCIDENTS:  
{known_vulnerabilities}"""}]  
    )
```

Notes: AI excels at threat modeling because it can consider thousands of known attack patterns and map them to your specific architecture. Always validate AI output with domain experts.

STRIDE Analysis: Operation Shadow Breach

STRIDE Category	Threat Identified	DREAD Score	Mitigation
Spoofing	Stolen credentials used for lateral movement	8.2	MFA everywhere, credential monitoring
Tampering	Attacker modified scheduled tasks	6.8	File integrity monitoring, GPO lockdown
Repudiation	Cloud audit logs could be disabled	5.4	Immutable logging, log forwarding
Info Disclosure	Data exfiltrated via DNS tunneling	9.1	DNS monitoring, DLP, encryption at rest
Denial of Service	Web server taken offline during response	4.6	Redundancy, failover, WAF

STRIDE Category	Threat Identified	DREAD Score	Mitigation
Elevation	RCE to admin via unpatched Struts	9.4	Patching SLA, WAF rules, least privilege

Building a Complete Monitoring Stack


A production-grade security monitoring stack integrates open-source and commercial tools, unified by a central SIEM with AI analytics.

```
# Complete security monitoring stack
monitoring_stack:
  siem:
    primary: "Elastic Security (ELK)"
    retention: "90 days hot, 1 year warm, 7 years cold"
    ai_integration: "LLM API for alert enrichment"

  network:
    ids: "Suricata"
    traffic_analysis: "Zeek"
    flow_collection: "ntopng / NetFlow"
    packet_capture: "Arkime (full PCAP)"
    dns_monitoring: "passivedns + Pi-hole logs"

  endpoint:
```

Monitoring Stack Data Flow



Security Metrics & KPIs

You cannot improve what you do not measure. These metrics prove the value of your SOC and guide investment decisions.

Metric	Definition	Target	Our Scenario
MTTD	Mean Time To Detect	< 5 min	3 min
MTTR	Mean Time To Respond	< 60 min	43 min
MTTC	Mean Time To Contain	< 30 min	46 min
FPR	False Positive Rate	< 10%	8%
Alert Volume	Alerts per analyst per day	< 50	35
Dwell Time	Attacker presence duration	< 24 hrs	46 min
Coverage	ATT&CK techniques detected	> 70%	82%

Measuring AI Impact on SOC Performance

Before AI Integration

- MTTD: 14 minutes
- MTTR: 2.5 hours
- False Positive Rate: 32%
- Alerts per analyst: 120/day
- ATT&CK coverage: 54%

After AI Integration

- MTTD: 3 minutes (-79%)
- MTTR: 43 minutes (-71%)
- False Positive Rate: 8% (-75%)
- Alerts per analyst: 35/day (-71%)
- ATT&CK coverage: 82% (+52%)

AI does not just make analysts faster — it fundamentally changes what a SOC team can achieve. These numbers are the language that justifies budget to leadership.

Phase 5: Hardening — Closing the Gaps

Every incident is a lesson. Hardening transforms incident findings into permanent defensive improvements.

Post-Incident Hardening Actions

1. **Patch management** — Emergency patch for Apache Struts; review patch SLAs for all internet-facing services
 2. **Network segmentation** — Isolate web tier from internal network; restrict SMB to necessary paths only
 3. **DNS security** — Deploy DNS-over-HTTPS monitoring; add anomaly detection for TXT record queries
 4. **Credential hardening** — Enforce MFA for all admin accounts; deploy credential guard on Windows endpoints
 5. **Cloud security** — Restrict AssumeRole to known IPs; enable CloudTrail log file validation
 6. **Detection rules** — Add 12 new SIEM rules based on observed TTPs
-

AI-Generated Detection Rules

```
def generate_detection_rules(incident_ttps, log_format="sigma"):
    """Generate new detection rules from incident findings."""
    client = anthropic.Anthropic()

    response = client.messages.create(
        model="claude-sonnet-4-20250514",
        max_tokens=4000,
        messages=[{"role": "user", "content": f"""Based on these TTPs observed
during an incident, generate detection rules in {log_format} format.

OBSERVED TTPs:
{incident_ttps}

For each rule provide:
1. Rule in {log_format} format
```

Example AI-Generated Sigma Rule

```
title: DNS TXT Query Volume Anomaly
status: experimental
description: Detects high volume DNS TXT queries indicating possible data exfil
logsource:
```

```
    category: dns
detection:
  selection:
    query_type: "TXT"
  timeframe: 5m
  condition: selection | count(query) by src_ip > 50
level: high
tags:
  - attack.exfiltration
  - attack.t1048.001
```

Continuous Improvement: Threat Intelligence Feeds

Static defenses lose effectiveness over time. Continuous improvement means your defenses evolve as fast as the threat landscape.

AI-Updated Detection Pipeline

```
def update_detections_from_threat_intel(new_iocs, existing_rules):  
    """Use AI to merge new threat intel into existing detection rules."""  
    client = anthropic.Anthropic()  
  
    response = client.messages.create(  
        model="claude-sonnet-4-20250514",  
        max_tokens=3000,  
        messages=[{"role": "user", "content": f"""Review these new IOCs from  
threat intelligence feeds and recommend updates to our detection rules.  
  
NEW THREAT INTELLIGENCE:  
{new_iocs}  
  
EXISTING RULES (summary):  
{existing_rules}"""}]
```

Feed Sources to Integrate

MISP (community IOCs), **AlienVault OTX** (open threat exchange), **abuse.ch** (malware/botnet tracking), **VirusTotal** (file/URL reputation), **Shodan** (exposure monitoring), and sector-specific ISACs.

Report Writing: Executive Summary

Every incident ends with a report. The executive summary is what leadership reads. The technical appendix is what your peers review.

```
def generate_executive_summary(incident_data):  
    """Generate an executive-friendly incident summary."""  
    client = anthropic.Anthropic()  
  
    response = client.messages.create(  
        model="claude-sonnet-4-20250514",  
        max_tokens=2000,  
        messages=[{"role": "user", "content": f"""Write an executive summary  
for this security incident. The audience is C-suite and board members.  
  
INCIDENT DATA:  
{incident_data}  
  
Requirements:  
- No jargon — explain technical terms in business language"""
```

Report Writing: Technical Report Structure

Full Technical Incident Report Outline

1. **Executive Summary** — Business impact, key findings, recommendations
2. **Incident Overview** — Classification, severity, scope, affected systems
3. **Timeline** — Minute-by-minute chronology with evidence sources
4. **Attack Chain Analysis** — MITRE ATT&CK mapping, full kill chain walkthrough
5. **Technical Findings** — IOCs, malware analysis, vulnerability details
6. **Response Actions** — What was done, when, by whom, and the outcome
7. **Forensic Evidence** — Chain of custody, artifacts collected, analysis results
8. **Root Cause Analysis** — Why it happened, contributing factors
9. **Lessons Learned** — What worked, what did not, what to change
10. **Recommendations** — Prioritized hardening actions with timelines and owners
11. **Appendices** — Full IOC list, raw logs, tool output, Sigma rules created

Notes: AI can draft each section from your raw notes and evidence. Always review AI-generated reports for accuracy — AI may infer causation where only correlation exists.

Combining All Tools & Techniques

The capstone challenge requires you to use every major tool and technique from the course. Here is the integrated toolkit.

Phase	Tools	AI Role
Monitor	ELK, Zeek, Suricata, Sysmon, CloudTrail	Baseline anomaly detection
Detect	SIEM rules, Sigma, YARA, Snort	Alert correlation & prioritization
Respond	TheHive, Cortex, iptables, AWS CLI	Playbook generation & guidance
Forensics	Volatility, Autopsy, Wireshark, plaso	Timeline reconstruction & analysis
Harden	CIS Benchmarks, Ansible, Terraform	Rule generation & config review
Report	Markdown, LaTeX, JIRA	Drafting & summarization

CTF Capstone Challenge: Operation Iron Shield

Mission Briefing

You are the blue team for MedTech Industries. An active, multi-stage attack is underway against your healthcare infrastructure. Patient data, research IP, and operational systems are all at risk. Detect, respond, and recover using AI at every step. You have 4 hours.

Challenge Phases

1. **Phase A: Detection (60 min)** — Analyze provided SIEM alerts, PCAP files, and log bundles. Identify all attack stages and map to MITRE ATT&CK.
2. **Phase B: Response (60 min)** — Write and execute containment commands. Generate an AI-assisted response playbook. Document every action taken.
3. **Phase C: Forensics (60 min)** — Build a complete attack timeline from evidence. Identify IOCs, determine data exposure, perform memory analysis.

4. **Phase D: Harden & Report (60 min)** — Create detection rules for observed TTPs. Write both executive and technical reports. Propose a 90-day hardening plan.
-

Capstone: Attack Scenario Details

The Attack You Will Face

- **Stage 1 — Initial Access:** Spear-phishing email with macro-enabled document targeting HR department
- **Stage 2 — Execution:** Macro drops PowerShell loader, establishes C2 via HTTPS to legitimate-looking domain
- **Stage 3 — Privilege Escalation:** PrintNightmare exploit against unpatched print server
- **Stage 4 — Lateral Movement:** Pass-the-hash across internal network using domain admin credentials
- **Stage 5 — Collection:** Access to patient database, research file shares, email archives
- **Stage 6 — Exfiltration:** Data staged and exfiltrated via encrypted channel to cloud storage
- **Stage 7 — Impact:** Ransomware deployment to 40% of endpoints as a diversion

Key Challenge

The ransomware (Stage 7) is a distraction. The real objective was data theft in Stages 5-6. Can your team identify the true threat while managing the visible crisis?

Capstone: Provided Materials

What You Receive

- **SIEM export** — 2,400 alerts from a 6-hour window (JSON format)
- **PCAP files** — Network captures from DMZ and internal segments (4 files, ~500MB)
- **Windows Event Logs** — Security, Sysmon, PowerShell logs from 8 key hosts (EVTX)
- **Memory dump** — RAM capture from compromised HR workstation (4GB, Lime format)
- **CloudTrail logs** — 48 hours of AWS API activity (JSON)
- **Email headers** — Original phishing email with full headers
- **Network diagram** — Topology of the target environment
- **Asset inventory** — All systems with OS versions, roles, and patch levels

AI Access

You have full access to the Claude API throughout the challenge. How you use it is part of the evaluation. Be strategic — document every AI interaction and how it influenced your decisions.

Scoring Rubric: Phase A — Detection

Criterion	Points	Description
Attack stages identified	20	Each of 7 stages correctly identified (partial credit for incomplete)
MITRE ATT&CK mapping	15	Accurate technique IDs for each stage
IOC extraction	10	Complete list of IPs, domains, hashes, file paths
AI-assisted correlation	10	Effective use of AI for cross-log correlation
False flag identification	5	Recognizing ransomware as diversion

Phase A Total: 60 points

Scoring Rubric: Phase B — Response

Criterion	Points	Description
Containment actions	15	Appropriate, timely, well-scoped containment

Criterion	Points	Description
Eradication completeness	10	All persistence mechanisms identified and removed
AI playbook quality	10	Context-specific, actionable, prioritized steps
Business continuity	5	Minimized operational disruption during response
Command documentation	5	Every response command logged with timestamp

Phase B Total: 45 points

Scoring Rubric: Phases C & D

Phase C — Forensics (45 points)

Criterion	Points	Description
Timeline accuracy	15	Complete, chronologically correct attack timeline
Evidence handling	10	Proper chain of custody, hash verification
Memory analysis	10	Artifacts extracted from RAM dump (processes, network, injected code)
Data exposure assessment	5	Accurate determination of what data was accessed/exfiltrated
AI-assisted analysis	5	Effective use of AI for artifact interpretation

Phase D — Harden & Report (50 points)

Criterion	Points	Description
Detection rules	10	Valid Sigma/YARA rules for observed TTPs
90-day hardening plan	10	Prioritized, realistic, addresses root causes
Executive summary	10	Clear, non-technical, business-impact focused
Technical report	10	Complete, evidence-backed, properly structured
AI documentation	10	Clear record of how AI was used and its impact on outcomes

Grand Total: 200 points

Evaluation Criteria: Grading Scale

Grade	Points	Description
Distinguished	180-200	Exceptional defense; identified all attack stages including the diversion; professional-quality reports
Proficient	140-179	Solid defense; detected most stages; effective containment and good documentation
Developing	100-139	Adequate defense; missed some attack stages or containment gaps; reports need improvement
Beginning	< 100	Significant gaps in detection or response; incomplete analysis and reporting

Bonus Points (up to 20)

- **+5** — Creative use of AI beyond suggested patterns
- **+5** — Identifying attacker infrastructure or attribution clues
- **+5** — Proposing detection rules that catch the attack at Stage 1

- **+5** — Exceptional report quality (presentation-ready for a real board meeting)
-

Capstone Tips & Strategy

Do This

- Start with the network diagram — understand the environment first
- Use AI for initial alert triage to save time
- Document everything as you go, not after
- Preserve evidence before analyzing it
- Look for what is NOT in the alerts (gaps in visibility)
- Manage your time — 60 min per phase is strict

Avoid This

- Do not chase the ransomware first — it is a distraction
 - Do not skip the memory analysis — critical evidence lives there
 - Do not blindly trust AI output — validate with evidence
 - Do not forget cloud logs — the attack spans on-prem and cloud
 - Do not write the report last — draft sections as you complete phases
 - Do not work alone if you have a team — divide by phase
-

Course Conclusion: The AI-Augmented Defender

Over 10 lessons, you have built a complete defensive skill set enhanced by AI at every stage. Here is what you now know how to do:

- **Monitor** — Build and operate a multi-layer security monitoring stack
- **Detect** — Write detection rules and use AI for cross-layer correlation
- **Analyze** — Triage alerts, hunt threats, and interpret complex evidence
- **Respond** — Execute containment, eradication, and recovery under pressure
- **Investigate** — Perform digital forensics with proper evidence handling
- **Harden** — Transform incident findings into permanent defensive improvements
- **Report** — Communicate findings to both technical and executive audiences
- **Automate** — Use AI to accelerate every phase of the defense lifecycle

What Sets You Apart

You do not just know how to use security tools — you know how to think like a defender, leverage AI as a force multiplier, and communicate your findings effectively. That combination is rare and valuable.

Notes: Encourage students to continue practicing with CTF platforms (HackTheBox, TryHackMe, CyberDefenders) and to pursue certifications like CompTIA Security+, CySA+, or GCIA.
