

Theory

Reading time: ~25 minutes · Combining every skill from Lessons 1–9 into a cohesive red team attack chain

Warning: Red team operations require **explicit written authorization** before any testing begins. Unauthorized penetration testing is illegal in virtually every jurisdiction and can result in criminal prosecution, civil liability, and career-ending consequences. Never execute any technique described in this course against systems you do not own or have documented permission to test.

1. Recap: How Lessons 1–9 Fit Together

Over the previous nine lessons, you built an arsenal of AI-assisted offensive security skills. Each lesson addressed a specific phase or capability within the broader red team workflow. Individually, these skills are useful. Combined into a coherent attack chain, they become transformative.

Lessons 1–3 focused on **reconnaissance and data analysis** — gathering intelligence through OSINT, identifying patterns in log data, and dissecting network traffic to understand the target environment. Lessons 4–7 shifted to **vulnerability discovery** — auditing source code, scanning web applications, reversing binaries, and probing infrastructure for misconfigurations. Lessons 8–9 addressed **exploitation and social engineering** — weaponizing discovered vulnerabilities into working exploits and crafting AI-enhanced phishing campaigns.

This capstone lesson ties everything together. You will learn the frameworks that professional red teams use to plan, execute, and report full attack chains, and you will see exactly where each previous lesson fits within those frameworks.

2. The Cyber Kill Chain (Lockheed Martin)

The Cyber Kill Chain, developed by Lockheed Martin in 2011, models the stages an adversary must complete to achieve their objective. Understanding this model allows red teams to simulate realistic attacks and helps defenders identify where to break the chain.

2.1 The Seven Phases

- 1. **Reconnaissance** — Harvesting email addresses, identifying technologies, mapping the attack surface. The attacker gathers as much information as possible about the target without direct interaction or through passive means.
- 2. **Weaponization** — Coupling an exploit with a payload (e.g., embedding a backdoor in a PDF, creating a malicious Office macro). This phase happens entirely on the attacker's side with no target interaction.
- 3. **Delivery** — Transmitting the weaponized payload to the target via email attachments, malicious links, USB drops, or watering-hole attacks.
- 4. **Exploitation** — Triggering the vulnerability to execute code on the target system. This might exploit a software bug, a misconfiguration, or human trust.
- 5. **Installation** — Establishing persistence by installing a backdoor, web shell, or implant that survives reboots and allows return access.
- 6. **Command & Control (C2)** — Opening a communication channel back to the attacker's infrastructure. The compromised system beacons out to receive instructions, exfiltrate data, or download additional tools.
- 7. **Actions on Objectives** — Achieving the ultimate goal: data exfiltration, lateral movement, privilege escalation, destruction, or ransomware deployment.

2.2 Mapping Lessons 1–9 to Kill Chain Phases

Kill Chain Phase	Lesson	Skill Covered
Reconnaissance	01 — AI-Powered OSINT Reconnaissance	Target profiling, attack surface mapping, data harvesting
Reconnaissance	02 — Log Analysis & Pattern Discovery	Identifying patterns, anomaly detection in operational data
Reconnaissance	03 — Network Traffic Analysis with AI	Protocol analysis, traffic fingerprinting, network mapping

Kill Chain Phase	Lesson	Skill Covered
Weaponization	04 — AI-Assisted Code Auditing	Discovering exploitable vulnerabilities in source code
Weaponization	06 — Binary & Cryptographic Analysis	Reverse engineering, identifying cryptographic weaknesses
Weaponization	08 — AI-Generated Exploit Development	Building working exploits and payloads
Delivery	09 — Social Engineering & Phishing with AI	Phishing campaigns, pretexting, payload delivery
Exploitation	05 — Web Application Vulnerability Scanning	SQLi, XSS, authentication bypass, web shell upload
Exploitation	07 — Configuration & Infrastructure Weaknesses	Misconfiguration exploitation, default credentials, exposed services
Installation / C2 / Actions	10 — Red Team Capstone (this lesson)	Persistence, lateral movement, objective completion, reporting

3. MITRE ATT&CK Framework Overview

While the Kill Chain provides a high-level linear model, the MITRE ATT&CK (Adversarial Tactics, Techniques, and Common Knowledge) framework offers a granular, matrix-based view of adversary behavior. ATT&CK catalogs real-world tactics and techniques observed in actual intrusions.

3.1 Tactics, Techniques, and Procedures

- **Tactics** — The adversary's tactical objective or the "why" behind an action. ATT&CK defines 14 tactics for Enterprise, including Initial Access, Execution, Persistence, Privilege Escalation, Defense Evasion, Credential Access, Discovery, Lateral Movement, Collection, Command and Control, Exfiltration, and Impact.
- **Techniques** — The "how" — specific methods used to achieve a tactic. For example, under Initial Access, techniques include Phishing (T1566), Exploit Public-

Facing Application (T1190), and Valid Accounts (T1078). Each technique has sub-techniques for even finer granularity.

- **Procedures** — The specific implementation of a technique by a particular threat actor or tool. For example, APT29 uses spearphishing with OAuth consent links (a procedure implementing T1566.002, Phishing: Spearphishing Link).

3.2 Kill Chain vs. ATT&CK: Complementary Models

The Kill Chain and ATT&CK are not competing frameworks — they are complementary. The Kill Chain provides a sequential narrative of an attack from start to finish, making it excellent for planning and communicating the overall flow of an engagement. ATT&CK provides depth, cataloging the specific techniques available at each stage and enabling precise mapping of red team activities to known adversary behaviors.

In practice, professional red teams use the Kill Chain to structure their engagement plan and ATT&CK to select and document the specific techniques they employ at each phase. When writing reports, Kill Chain stages organize the narrative while ATT&CK IDs provide the technical precision that defenders need to build detections.

4. Planning a Full Attack Chain

A professional red team engagement does not begin with exploitation. It begins with meticulous planning that combines every discipline from this course into a coherent strategy.

4.1 Phase 1: Intelligence Gathering

Start with the OSINT techniques from Lesson 1. Map the target's external footprint: domains, subdomains, IP ranges, employee names, technology stack, cloud providers, and third-party relationships. Use AI to correlate findings from multiple sources and identify patterns a human analyst might miss. Feed network captures (Lesson 3) through AI classifiers to fingerprint services and identify unusual traffic patterns.

4.2 Phase 2: Vulnerability Discovery

With the attack surface mapped, systematically probe for weaknesses. Use AI-assisted code auditing (Lesson 4) on any open-source components or leaked code. Scan web applications (Lesson 5) for injection points, authentication flaws, and business logic errors. Analyze any binaries or cryptographic implementations (Lesson 6) for reversible

weaknesses. Audit infrastructure configurations (Lesson 7) for default credentials, exposed management interfaces, and overly permissive access controls.

4.3 Phase 3: Exploitation and Access

Develop exploits (Lesson 8) targeting the most promising vulnerabilities. Craft social engineering campaigns (Lesson 9) as an alternative or supplementary access vector. The best red teams maintain multiple access paths so that if one is detected and blocked, others remain available.

4.4 Phase 4: Post-Exploitation

Once initial access is achieved, the real work begins. Establish persistence, escalate privileges, move laterally through the network, and work toward the engagement objectives. Document every step meticulously — the value of a red team engagement lies not in the access itself but in the detailed report of how that access was achieved and what it enabled.

5. AI as a Force Multiplier Across the Kill Chain

Throughout this course, AI has served as a force multiplier at every stage. In reconnaissance, AI processes and correlates vast amounts of OSINT data faster than any human analyst. In vulnerability discovery, AI identifies patterns in code and configurations that manual review would miss. In exploitation, AI assists with payload generation, evasion techniques, and adapting exploits to specific targets. In social engineering, AI crafts personalized phishing content that is linguistically convincing and contextually relevant.

The key insight is that AI does not replace the red team operator — it amplifies their capabilities. The operator provides judgment, creativity, and ethical boundaries. AI provides speed, pattern recognition, and the ability to process information at scale. A skilled operator with AI assistance can accomplish in hours what previously took days or weeks.

However, this amplification cuts both ways. The same AI capabilities available to red teams are available to real adversaries. Understanding how AI enhances offensive operations is therefore essential for building effective defenses.

6. Red Team Engagement Methodology

6.1 Rules of Engagement (ROE)

Every professional red team engagement begins with a detailed Rules of Engagement document. The ROE defines the boundaries of the engagement and protects both the red team and the client. A proper ROE must include:

- **Scope** — Exactly which systems, networks, applications, and personnel are in scope. Anything not explicitly listed is out of scope.
- **Authorized techniques** — Which ATT&CK techniques are permitted. Some clients may prohibit denial-of-service, data destruction, or social engineering targeting specific individuals.
- **Time windows** — When testing may occur. Some clients restrict testing to business hours; others want after-hours testing to simulate real attacks.
- **Communication protocols** — How the red team communicates with the client's point of contact. This includes emergency stop procedures if testing causes unintended disruption.
- **Data handling** — How sensitive data encountered during testing will be handled, stored, and destroyed.
- **Legal authorization** — Signed authorization from someone with legal authority to approve the testing (typically a C-level executive or general counsel).

6.2 Engagement Types

Red team engagements vary in scope and knowledge level:

- **Black box** — The red team has no prior knowledge of the target. Simulates an external attacker.
- **Grey box** — The red team has partial knowledge (e.g., network diagrams, credentials for a low-privilege account). Simulates an insider or a compromised contractor.
- **White box** — The red team has full knowledge of the environment. Maximizes coverage but does not test detection of initial access.
- **Assumed breach** — The red team starts with an established foothold. Focuses on post-exploitation, lateral movement, and objective completion.

7. Writing a Professional Red Team Report

The report is the most important deliverable of any red team engagement. A brilliant attack chain that is poorly documented provides no value to the client. The report must communicate findings to multiple audiences: executives who control budgets, managers who prioritize remediation, and engineers who implement fixes.

7.1 Report Structure

A professional red team report typically follows this structure:

1. **Executive Summary** (1–2 pages) — High-level findings, overall risk assessment, and top recommendations. Written for non-technical stakeholders.
2. **Engagement Overview** — Scope, objectives, methodology, timeline, and team composition.
3. **Attack Narrative** — A chronological walkthrough of the attack chain, from initial reconnaissance through objective completion. Include timestamps, screenshots, and ATT&CK technique IDs.
4. **Findings & Recommendations** — Each finding documented with severity rating, affected systems, evidence, business impact, and specific remediation steps.
5. **ATT&CK Mapping** — A matrix or table mapping every technique used to its ATT&CK ID, with detection recommendations for each.
6. **Appendices** — Raw data, tool output, exploit code, and detailed technical evidence.

7.2 Writing the Executive Summary

The executive summary is the most-read section of the report. It must be concise, accurate, and actionable. Here is an example structure:

Example: Executive Summary

Engagement: External Red Team Assessment for Acme Corporation

Period: March 1–21, 2026

Overall Risk Rating: Critical

Summary: The red team achieved full domain administrator access within 72 hours of engagement start, beginning from a zero-knowledge external position.

Initial access was obtained through a spearphishing campaign targeting the engineering department, which yielded credentials for a developer workstation. From this foothold, the team escalated privileges by exploiting a misconfigured Active Directory certificate template, moved laterally to the finance server, and exfiltrated simulated sensitive financial data.

Key Findings:

- Email gateway did not detect AI-crafted phishing payloads (Critical)
- Active Directory certificate services misconfiguration allowed privilege escalation (Critical)
- Network segmentation between engineering and finance was insufficient (High)
- Endpoint detection did not alert on lateral movement tooling (High)

Top Recommendations:

1. Deploy advanced email filtering with AI-based phishing detection
2. Audit and remediate Active Directory certificate templates immediately
3. Implement micro-segmentation between business units
4. Upgrade EDR to a solution with behavioral detection capabilities

8. Metrics and Measuring Red Team Effectiveness

Effective red team programs measure their impact with quantifiable metrics. Without metrics, it is impossible to demonstrate improvement over time or justify the investment in red team operations.

8.1 Offensive Metrics

- **Time to initial access** — How long from engagement start to first foothold. Shorter times indicate weaker perimeter defenses.
- **Time to objective** — How long to achieve the defined objective (e.g., domain admin, data exfiltration). This is the most important offensive metric.

- **Number of access paths discovered** — How many independent paths to the objective exist. More paths indicate a larger attack surface.
- **Techniques successfully employed** — Mapped to ATT&CK, this shows which categories of attack the organization is vulnerable to.

8.2 Defensive Metrics

- **Mean time to detect (MTTD)** — How long before the blue team noticed red team activity. Lower is better.
- **Mean time to respond (MTTR)** — How long from detection to effective containment. This measures the incident response capability.
- **Detection coverage ratio** — What percentage of red team techniques were detected. Map this against ATT&CK to identify blind spots.
- **False positive rate** — How many alerts fired that were not related to red team activity. High false positive rates indicate noisy detections that lead to alert fatigue.

8.3 Program Metrics

- **Remediation rate** — What percentage of findings from previous engagements have been resolved. This is the ultimate measure of program value.
- **Year-over-year improvement** — Compare time-to-access, time-to-objective, and detection coverage across engagements to demonstrate security posture improvement.
- **Cost per finding** — The engagement cost divided by the number of actionable findings. Helps justify budget allocation.

9. Summary

A full red team attack chain is more than the sum of its parts. The Kill Chain provides the narrative structure; ATT&CK provides the technical vocabulary; lessons 1–9 provide the skills; and professional methodology ties everything together into an engagement that delivers real value to the client.

Key takeaways:

- **Frameworks guide execution** — Use the Kill Chain for planning and ATT&CK for technique selection and documentation

- **AI amplifies every phase** — From reconnaissance through reporting, AI makes red team operators faster and more thorough
- **Authorization is non-negotiable** — Written ROE and legal authorization must precede any testing activity
- **The report is the product** — A red team engagement without a clear, actionable report delivers zero value
- **Metrics drive improvement** — Measure offensive performance, defensive response, and program-level outcomes to demonstrate and improve effectiveness
- **Ethics are foundational** — Every technique in this course must only be applied within the bounds of explicit authorization and applicable law

"The purpose of a red team is not to break things — it is to make the organization stronger by revealing how things can be broken."

10. Challenge Resources

Download the following files to use during the capstone lab. These provide the scenario data for the MegaCorp Industries red team engagement.

Download Files

Capstone scenario data for the MegaCorp Industries red team engagement. All files are plain text (.txt) and safe to open.

↓ [recon-data.txt](#) ~8 KB

↓ [vulnerability-scan.txt](#) ~9 KB

↓ [attack-timeline.txt](#) ~10 KB

↓ [scope-document.txt](#) ~12 KB