

Ethical Hacking & Penetration Testing

Course Duration ~90 Hours · 12 Weeks · Live Online · Lab-Heavy · Isolated Environments, Assessed Reports, Capstone Engagement

Curriculum reviewed: July 2026 — offensive tooling and the AI attack surface move fast; this syllabus is checked against official sources

Authorised Testing Only — read first. Every technique on this course is taught for **authorised security testing only**, performed exclusively in isolated lab environments and against deliberately vulnerable targets provided for the course. You must **never test any system you do not own or lack explicit written authorisation to test** — unauthorised access to computer systems is a criminal offence, including under Indian law, and this course covers that legal framework and its consequences. Professional engagements begin with scoping, rules of engagement and a **signed authorisation document**; the course teaches that process as the required first step of every engagement, and covers responsible disclosure as professional practice.

[Engagement Lifecycle](#)[Web & API](#)[Active Directory](#)[AI Red Teaming](#)[Reporting](#)[Capstone](#)

In short: Find it before they do — legally, methodically, and with a report a client can act on. This is the **offensive-security specialisation** of our security track, for people who want to become professional security consultants. The recommended prerequisite is our Cybersecurity course (the defensive foundation) or equivalent experience: we do not re-teach cryptography, identity or SOC operations — we assume them and go deep on methodology, exploitation, chaining, post-exploitation, and the two things that separate a practitioner from a lab-solver. First, **the report is the product**: clients buy a document that says what is broken, how badly, and what to do — the leading practical certification enforces this by scoring an inadequate report zero regardless of machines compromised, so reporting gets a full module and a real assessed deliverable. Second, **AI red teaming**: a full module on prompt injection, retrieval-pipeline exploitation, multi-agent chaining and tool-layer abuse, mapped to the LLM application vulnerability list — and our genuine advantage is that this catalogue already teaches how those agent and retrieval systems are built, which is what makes someone effective at breaking them. Across 11 modules you scope an engagement with written authorisation, map an attack surface, test web and API applications with proxy-driven manual technique, compromise a realistic Active Directory domain, run an AI red-team engagement, and deliver a full simulated engagement as a professional report and presentation. Mapped to the recognised progression: entry-level practical, the leading hands-on penetration-testing certification (OffSec OSCP/OSCP+), and emerging AI red teaming credentials (OffSec AI Red Teamer, HackTheBox Certified Offensive AI Expert). Honest note: the leading practical exam is demanding and most build foundations first. On AI as a tool for the tester: it accelerates recon, payloads and report drafting — it does not replace methodology, creativity or judgment.

Duration	12 weeks · ~90 hours · live + hands-on	Format	Live online · lab-heavy · isolated environments, assessed reports, capstone engagement
Prerequisite	Security fundamentals (our Cybersecurity course or equivalent), comfort with Linux and networking, and basic scripting. Programming experience an advantage, particularly for web and API testing.	Curriculum	11 modules, from the professional engagement to reporting and a capstone engagement
Certification	Mapped to the recognised progression: entry-level practical, the leading hands-on penetration-testing certification, and emerging AI red teaming credentials	Ideal for	SOC analysts moving to offence, developers moving into application security, and AI/ML engineers assessing what they build

The Report Is the Product

Clients do not buy exploitation; they buy a document that tells them what is broken, how badly, and what to do about it. The industry's most respected practical certification enforces this directly — **an inadequate report scores zero, regardless of how many machines were compromised**. So professional reporting, evidence capture, risk rating and client communication get a full module and a real assessed deliverable here — the single biggest gap between people who can pass a lab and people who can be billed to a client.

AI Red Teaming — Where This Course Is Different

AI red teaming has become a distinct offensive discipline, with dedicated certifications now emerging from the same organisations behind the established penetration-testing credentials. Organisations deploy language models, retrieval pipelines, copilots and autonomous agents faster than they can assess them, and traditional methodology does not cover the resulting attack surface. A full module covers direct and indirect

prompt injection, system prompt extraction, retrieval-pipeline exploitation, multi-agent chaining, tool and integration-layer abuse, insecure output handling, and mapping findings to the recognised LLM application vulnerability list. Our genuine advantage: understanding how agent and retrieval systems are built is what makes someone effective at breaking them — and this catalogue already teaches that. AI is also honestly covered as a tool for the tester: it accelerates recon, payloads and drafting; it does not replace methodology or judgment.

Course Highlights

- Every module leads with what you do, exploit, chain or report — not what you “learn about”
- The engagement lifecycle taught as a profession — scoping, written authorisation, rules of engagement, disclosure
- The report as the deliverable — a full module and a real assessed professional report
- A dedicated AI red teaming module — the discipline generic bootcamps don't cover
- Active Directory to real depth — the single most valuable enterprise skill in the field
- Web and API testing driven by proxy-based manual technique, not scanner output
- Every technique in isolated labs against deliberately vulnerable targets — authorised testing only
- A capstone engagement scoped, executed, reported and presented end to end

Hands-on Projects & Portfolio

- **Recon & Attack Surface Map** — full reconnaissance and attack-surface map of a lab target organisation
- **Web Application Assessment** — against deliberately vulnerable applications, documented to client standard
- **API Assessment** — authorisation, token handling and business logic abuse
- **Active Directory Engagement** — compromise a realistic multi-host domain and map the full attack path
- **Privilege Escalation Gauntlet** — across Linux and Windows targets
- **Cloud Misconfiguration Exercise** — identity escalation in an isolated account
- **AI Red Team Engagement** — assess a retrieval-based agent application, chain a prompt injection to a real impact, and recommend guardrails
- **Authorised Phishing Simulation** — designed, executed in lab and analysed within scope
- **Assessed Professional Report** — the deliverable that decides whether you're employable
- **Capstone Engagement** — complete simulated engagement, scoped, executed, reported and presented

Tools & Stack

Kali Linux · Nmap · Burp Suite · Metasploit · BloodHound · Impacket · Responder · CrackMapExec · Hashcat · John the Ripper · ffuf · sqlmap · Nuclei · Nessus or OpenVAS · Wireshark · PowerShell & Bash · Python · Sliver or Covenant · Gophish · Pacu or ScoutSuite · Garak or PyRIT · MITRE ATT&CK · OWASP Testing Guide · OWASP LLM Top 10 · Docker

Certification Alignment

- Mapped to the recognised progression in offensive security — an **entry-level practical credential**, the leading **hands-on penetration-testing certification (OffSec OSCP / OSCP+)**, and the **emerging AI red teaming credentials** (such as OffSec's AI Red Teamer and HackTheBox's Certified Offensive AI Expert)
- Honest note: the leading practical certification is genuinely demanding, and most candidates build foundations first — this course is designed to get you ready to attempt it, not to promise a shortcut
- Presented as preparation and alignment — no official partnership or accreditation implied, and no exam outcome guaranteed

Career Outcomes

Offensive security is a consulting-shaped profession — demand is driven by regulatory requirements, client assurance and annual testing obligations, which makes it steady work concentrated in consultancies, managed security providers and large in-house security teams. Two honest points: employers hire on demonstrated ability rather than coursework, so lab write-ups, disclosed findings and a professional report portfolio matter more than a certificate alone — which is why this course produces exactly those artefacts; and **developers are unusually well positioned to enter this field**, because understanding how applications are built is a direct advantage in web, API and AI testing. This course targets:

- Penetration Tester · Security Consultant · Application Security Engineer · Red Team Operator (with experience)
- Vulnerability Assessment Analyst · Bug Bounty Hunter · Security Researcher · AI Red Teamer · Purple Team Engineer

Who This Course Is For

- Security analysts and SOC professionals moving from defence into offensive work
- **Developers moving into application security** — a strong, under-recognised path that connects to the development courses in this catalogue
- System and network administrators with strong infrastructure knowledge
- Final-year students and fresh IT graduates who have completed security foundations
- **AI and ML engineers who need to assess the systems they build** — an emerging and underserved audience
- Consultants and auditors who need genuine technical testing capability
- Anyone preparing for a practical, hands-on penetration-testing certification

Module 1: The Professional Engagement

7 Hours

YOU PRODUCE

- A complete engagement scope with rules of engagement and a written authorisation document — the required first step of every real test, drafted before any tooling is touched

TOPICS

- What a penetration test actually is and how it differs from a vulnerability scan; engagement types (black, grey and white box; internal, external, red team, purple team)
- Scoping, rules of engagement and written authorisation; legal frameworks and liability; established testing methodologies and standards
- Ethics and responsible disclosure

Assessment: Quiz + Lab

Module 2: Reconnaissance & Enumeration

8 Hours

YOU DO

- A full reconnaissance and attack-surface map of a lab target organisation — the enumeration discipline that quietly decides most engagements

TOPICS

- Passive intelligence gathering and open-source reconnaissance; attack surface mapping; DNS and infrastructure discovery
- Active scanning and service enumeration; fingerprinting
- The discipline that decides most engagements: thorough enumeration beats clever exploitation

Assessment: Quiz + Lab

Module 3: Vulnerability Analysis & Exploitation Fundamentals

8 Hours

YOU EXPLOIT

- A validated vulnerability on a lab target — researched, safely adapted and understood well enough to explain exactly what it does and why it works

TOPICS

- Vulnerability identification and validation; separating real findings from scanner noise; public exploit research and safe adaptation
- Exploitation frameworks; manual exploitation; payloads and shells; evading basic defences
- Understanding an exploit well enough to explain what it does

Assessment: Quiz + Lab

Module 4: Web Application Testing

9 Hours

YOU DO

- A proxy-driven manual assessment of a deliberately vulnerable application — the dominant vulnerability classes found, proven and documented to client standard

TOPICS

- Injection; authentication and session flaws; access control failures; server-side request forgery
- Deserialisation; file upload abuse; business logic flaws; client-side issues
- Proxy-driven manual testing throughout

Assessment: Quiz + Lab**Module 5: API & Modern Application Attack Surface**

8 Hours

YOU DO

- An API assessment covering authorisation, token handling and business logic abuse — where more risk now lives than in the interfaces in front of it

TOPICS

- REST and GraphQL testing; authentication and token handling flaws including JWT and OAuth misuse
- Broken object and function level authorisation; mass assignment; rate limiting and abuse
- Why APIs now carry more risk than the interfaces in front of them

Assessment: Quiz + Lab**Module 6: Network & Infrastructure Exploitation**

8 Hours

YOU EXPLOIT

- A chained path across a lab network — initial foothold to privilege escalation to pivot — on both Linux and Windows targets

TOPICS

- Service-specific attacks; credential attacks and password cracking; man-in-the-middle and network-level techniques
- Pivoting and tunnelling; privilege escalation on Linux and Windows
- Persistence concepts; living-off-the-land techniques

Assessment: Quiz + Lab**Module 7: Active Directory Attacks**

9 Hours

YOU DO

- A realistic multi-host domain compromise with the full attack path mapped — the single most valuable enterprise skill in this field

TOPICS

- Domain enumeration; Kerberos attack techniques; credential theft and reuse; delegation abuse
- Access control weaknesses; lateral movement; domain privilege escalation
- Attack-path mapping across a realistic multi-host domain

Assessment: Quiz + Lab**Module 8: Cloud & Container Testing**

7 Hours

YOU DO

- An identity-escalation and misconfiguration exercise in an isolated cloud account — staying strictly within what provider testing policies permit

TOPICS

- Cloud attack surface; identity and permission escalation paths as the dominant cloud risk; metadata and instance exposure
- Storage misconfiguration; serverless considerations; container escape concepts and orchestrator misconfiguration
- What is and is not permitted under provider testing policies

Assessment: Quiz + Lab

Module 9: AI Red Teaming

9 Hours

YOU DO

- A red-team assessment of a retrieval-based agent application — chain an indirect prompt injection to a real impact, then write up the finding and recommend guardrails, mapped to the LLM application vulnerability list

TOPICS

- Assessing language model applications, retrieval pipelines, copilots and autonomous agents; direct and indirect prompt injection; system prompt extraction
- Retrieval boundary and data exfiltration failures; multi-agent chaining; tool and integration-layer abuse; insecure output handling
- Model and data poisoning concepts; guardrail evasion; mapping findings to the recognised industry vulnerability list for LLM applications

Assessment: Quiz + Lab

Module 10: Social Engineering & Physical Considerations

6 Hours

CONSENT & SCOPE

When the target is a person, the ethical and consent boundaries are stricter, not looser. Every activity here is performed only within an authorised engagement scope agreed in writing, never against unconsenting people.

YOU DESIGN

- An authorised phishing simulation — designed, executed in the lab and analysed — within a scope agreed in writing

TOPICS

- Phishing campaign design within authorised scope; pretexting; awareness testing
- AI-assisted lure generation and what it changes; physical security concepts
- The strict ethical and consent boundaries that apply when the target is a person rather than a system

Assessment: Quiz + Lab

Module 11: Reporting, Communication & Capstone

11 Hours

YOU DELIVER

- One full simulated engagement, end to end — scoped, tested, findings chained — delivered as a professional report and a stakeholder presentation: the assessed deliverable that decides employability

TOPICS

- Evidence capture and note discipline during testing; writing findings developers can act on; risk rating and prioritisation
- Executive summaries for non-technical stakeholders; remediation advice; retesting and verification
- Client debriefs and difficult conversations; the full end-to-end capstone engagement

Assessment: Capstone Engagement + Professional Report + Presentation