Faysal Aziz

Hacked at Machine Speed

Managing cybersecurity vulnerabilities in the age of AI.

FA-0836 · 1 day

Cybersecurity

Learn practical personal, Linux-server and web-application defences, including the risks of AI-assisted attacks and overprivileged agents.

Introduction

Internet-facing systems are routinely scanned for exposed services, weak credentials and known vulnerabilities. Defenders need to understand what is exposed and how to reduce avoidable risk before and after a service goes live.

Attackers use AI to assist vulnerability research, automate parts of attacks and generate convincing impersonation content. Google threat researchers have documented agentic attack frameworks and automated analysis of software patches. These developments increase the importance of timely updates, evidence-based verification and layered defences.

Automation can lower the effort needed to target websites, servers and personal accounts. Both individuals and organisations benefit from practical security habits and clear recovery plans.

Established controls remain essential: strong authentication, timely patching, least privilege and reduced exposure. These measures reduce risk but do not guarantee that an attacker will move on. This course applies them methodically and in a practical order.

The day starts from zero. Participants learn what cybersecurity protects, how an attack unfolds, and what vulnerability, exploit, red team, blue team and penetration testing mean in practice. The course then examines how AI has reshaped the threat landscape, including the newer risks created by AI assistants and agents themselves, which can be manipulated through hidden instructions or given far more access than their task requires.

The second half is practical. Participants strengthen their own accounts and devices, then work through the hardening of a Linux server in a prepared training environment: reducing what is exposed, securing remote access, enabling a firewall, banning repeated failed logins and keeping the system patched. The course then moves up to the web application, the front door that stays open after the server is locked,

covering admin login protection, two-factor sign-in, least privilege for accounts and AI agents, secrets, logging, backups and recovery. Core steps are guided hands-on work, and the remainder is demonstrated live. AI is also put to work for the defender, as an auditor and an independent reviewer.

The course is designed for beginners who build or manage websites, professionals seeking to understand organisational risks, and individuals protecting personal accounts and devices. Core steps use guided exercises in a prepared environment; other topics are demonstrated. Controls must be adapted to the system and its operational needs.

Learning outcomes

By the end of this course, participants will be able to:

- Explain core cybersecurity concepts, including assets, threats, vulnerabilities, exploits and risk.
- Describe the roles of red teams, blue teams and penetration testing, and the legal boundaries that govern them.
- Explain how attackers use AI to find vulnerabilities, automate attacks and impersonate trusted people.
- Recognise warning signs of impersonation and apply independent verification to suspicious messages, voices and video calls.
- Protect personal accounts and devices with a password manager, two-factor authentication, updates and backups.
- Identify what a server exposes to the internet and reduce that exposure.
- Harden a Linux server through secure remote access, a firewall, automatic banning and patching.
- Apply application hardening principles, including admin login protection, least privilege and secrets management.
- Limit the access given to AI assistants and agents, and recognise prompt injection risks.
- Use logs, backups, tested restores and rollback to detect, respond to and recover from an incident.
- Use AI tools to audit and review their own security.

Prerequisites

- Confident everyday use of a computer, web browser and email.
- A laptop with a current browser and an SSH client.
- A smartphone on which an authenticator app can be installed.
- Internet access permitted for the training environment.
- Permission and connectivity to use SSH with the designated training systems.

Training outline

1. Cybersecurity Foundations

1. Assets, threats, vulnerabilities and risk
2. Confidentiality, integrity and availability
3. Attack surface and the stages of an attack
4. Vulnerabilities, CVEs, patches and zero-days
5. Defence in depth and least privilege

2. Red Team, Blue Team and Penetration Testing

1. Red, blue and purple team roles
2. Vulnerability scanning versus penetration testing
3. Penetration testing stages and rules of engagement
4. Legal and ethical boundaries
5. Acting on findings

3. How AI Changed the Threat Landscape

1. AI-assisted reconnaissance and vulnerability discovery
2. From published patch to working exploit
3. Autonomous attack agents
4. AI-written phishing, voice cloning and deepfakes
5. Prompt injection and excessive agency in AI tools
6. AI coding assistants and supply-chain risk

4. Personal Cyber Defence

1. Passwords, passphrases and password managers
2. Two-factor authentication and passkeys
3. Recognising AI-generated phishing and impersonation
4. Verification habits for calls, messages and payments
5. Device, browser and home network updates
6. Backups and account recovery

5. Server Hardening

1. Discovering what a server exposes
2. Secure remote access and SSH configuration
3. Firewalls and default-deny rules
4. Automatic banning of repeated failed logins
5. Security updates, automatic patching and reboots
6. Users, privileges and service isolation
7. HTTPS and security headers

6. Application Hardening

1. Common web application weaknesses
2. Admin login protection and two-factor sign-in
3. Password rules, sessions and cookies
4. Least privilege for accounts, API keys and AI agents
5. Secrets management
6. Dependencies, plugins and third-party code

7. Detection, Response and Recovery

1. Logs and login monitoring
2. Backups, restore tests and rollback
3. Incident response steps
4. Review and testing before every change

8. AI for the Defender

1. AI-assisted security audits
2. Independent AI review of changes
3. Safe limits for AI agents on live systems
4. Personal and organisational security checklists

Disclaimer

This course outline serves as a guideline for training planning and delivery. The trainer reserves the right to amend, reorder, expand, reduce, or otherwise modify the content and delivery approach as he considers appropriate, taking into account participant experience, organisational requirements, technical constraints,

and available time, without prior notice.

© 2026 Faysal Aziz · Kuala Lumpur