TAKE-AWAY DOCUMENT

The AI-Driven Engineering curriculum.

This is the AI-Driven Engineering part of the AI-Driven Cybersecurity Master's. It keeps growing: new sessions are added every week as the field changes. Every session comes with a live workshop, a written workbook, audio read by HAL, and GitHub repos to build with.

1 · The Harness Driving Claude Code

1	**Foundations**	What AI-driven engineering is and why code was never the hard part; the proof, from three shipped products to over 500 job adverts paying for it; the four tiers of building around an AI; the six-step method and the one rule...
2	**The Basics**	Install Claude Code, learn permissions, the tool cycle, CLAUDE.md, git, context memory, models, and sessions.
3	**Beyond the Basics**	MCPs, plan mode, the development protocol, models, permissions, output style, thinking mode, auto-compact, rewind, slash commands.
4	**Commands & Skills**	What commands and skills are, how they differ, where they live, when to reach for which — plus an open brainstorm of what to automate first.
5	**MCP**	What MCP is, where servers live, the three transports, built-in connectors, installing from the CLI, secrets, security, lazy loading, and the integration ladder from MCP to script to API.
6	**Five Builds**	The five hands-on builds you need from Sessions 2 to 5, worked through together on a live call. Use this page as your reference at your own pace: install Claude Code, connect an MCP, write your first command and skill.
7	**Tiers & Hooks**	The four tiers of AI-driven engineering — from raw harness to shared infrastructure — and hooks: the deterministic guardrail that turns a stochastic agent into something you can trust.
8	**Agents**	Agent definitions, sub-agents, team agents, the collision problem and worktrees, Agent View, dynamic workflows and ultracode, plugins and the marketplace, plus the security discussion — can the model itself be broken.

2 · Build It Your AI assistant & infrastructure

9	**Assistant Prep**	The Fable models' security gating and what it means for security people, AI-driven engineering vs the AI security career roles, the move to Tier 2, the assistant prerequisites checklist, which MCPs are worth it, and a live HAL...
10	**Assistant Anatomy**	What an AI assistant actually is, the seven parts it's made of (structure, identity, memory, skills, work, interface, governance), and the order to build them in.
11	**Installing Your Assistant**	Getting the HAL skeleton installed and running, every real install gotcha from the live session, why the struggle is the point, and Structure: the first of the seven parts.
12	**Identity, Memory, Skills**	The four parts that make an assistant yours, plus how to choose your models and keep your data private.

13	**Pentest & Bug Bounty**	What becoming AI-driven actually does to a real job — encoding your expertise into agents that hunt bugs and drive pen-testing tools, shown live.
14	**Work & Interface**	How the job actually gets done — agents, loops and headless runs — and how you reach the thing you've built. Plus a bug bounty update and what changed in the job market.
15	**Governance**	The three questions to ask about any agent you run, the four jobs of a control, and five fixes you can put in place tonight — starting with a deny list that takes five minutes.

3 · One tool, end to end Architect, build, prove, ship

16	**Architect It First**	One small problem — is this domain, IP, URL or file malicious? — architected out loud: research the providers, choose by criteria, screen the risk, pick the stack, and meet the five documents that make building repeatable.
17	**API Client**	The VirusTotal client gets built live — one plain-English prompt, the standards do the steering. What the agent did unprompted, what the watchlist caught, the review that ran, the review that didn't, and the layers of checking...
18	**Proving Your Code**	Six layers of checking on the VirusTotal client, cheapest first: linters and types, tests and mutation, the three scans, specialist reviewers, end-to-end against the real API, exploratory personas, then the one document that runs...
19	**Delivery Pipeline**	The pipeline for code an agent wrote: the six empty seats, CI, CD and DevSecOps in plain words, the delivery protocol from the vault, gates and where they stand, staging and rollback, then Git and GitHub properly: commit, push...
20	**Branches, Hooks, CI**	Main, branch, origin and worktree built up one picture at a time, the pull request, the live repo map, seven rules for every day, two agents on the same code, lefthook on your machine, CI on a clean one, and branch protection...
21	**Safe Deploys**	Dependabot and Trivy watching what you depend on, why an update is only as safe as your tests, dev, staging and production, the deploy and the way back, the commit-count check, what to test at each station, and the words you say.
22	**Assistants & Mods**	The new personal assistants from the big labs, the HAL dashboard in demo mode, then what a mod is, where it shows up, what a settings hook can really do, choosing the lightest tool, building a force-push guard without writing...

4 · The AI-Driven Engineering Lifecycle (AIDEL) A step-by-step way to build with AI you can trust

1	**Screen**	What am I responsible for, and how risky is this?
2	**Frame**	Is this the right problem, and is it worth solving?
3	**Spike**	Is there one real unknown worth testing cheaply first?
4	**Build**	Build it, with the right amount of rigour for the risk.
5	**Verify**	Does it give the right answers? Where's the evidence?
6	**Gate**	Is it allowed to be released?

5 · Technical & security architecture domains

Every area you direct AI across to design and secure a complete system

Architecture Practice

Thinking Like an Architect · Recording & Communicating Decisions · Frameworks & Reference Models · Technology Selection · Directing the AI Executor

Application Architecture

Client & Server · Internal Structure · System Topology · Frontend Delivery · Platform Targets

APIs & Integration

What an API Is · API Specs — the Contract in Writing · API Styles · Events, Webhooks & Queues · Third-Party Services · Living With APIs

Networking & the Web

Internet Fundamentals · Names & Addresses · HTTP & HTTPS · Traffic Management · Private Networks

Reliability, Performance & Operations

Scaling · Performance & Caching · Resilience · Observability · Operations Forever

AI Architecture

Models & Providers · Talking to Models · Retrieval & Knowledge · Agents & Orchestration · AI in Production

Identity & Access

Authentication · Sessions & Tokens · Federation & SSO · Authorization · Machine Identity & Secrets

Data Security & Privacy

Cryptography Essentials · Protecting Data in Each State · Secrets Hygiene · Privacy Duties · Breaches & Custodianship

AI Security & Governance

Attacks on AI Features · Agent Guardrails · AI Data Risks · AI-Built Code & Supply Chain · AI Governance & Rules

Languages & Code

What Code Actually Is · Languages for Web Pages · Languages for Servers & Scripts · Languages for Operating Systems & Infrastructure · Languages for Mobile Apps · Languages for Enterprise Backends · Languages for Data & Databases · Choosing a Language & Stack · Frameworks & Libraries · Packages & Borrowed Code

Data & Databases

Database Kinds · Modelling & Querying · Analytics & Movement · Custodianship

Infrastructure & Cloud

Where Software Runs · The Cloud & Its Providers · Managed & AI-Builder Platforms · Containers & Serverless · The OS Underneath · Infrastructure as Code

Delivery & Deployment

Version Control · Build & Test · Environments & Config · CI/CD & Releases · Why Deploys Break

Cost & Economics

What You're Actually Paying For · Free Tiers & Their Cliffs · Costs That Scale With Success · Cost Controls

Security Principles & Threat Modelling

Core Principles · Know Your Attacker · Threat Modelling · Risk & Compliance for Builders

Application & Supply Chain Security

Named Attacks · The Untrusted Client · Defensive Coding · Security Testing · Supply Chain

Security Operations & Resilience

Defending the Perimeter · Hardening & Cloud Posture · Vulnerability Management · Detection & Monitoring · Incident Response & Recovery

Around the curriculum: a personal roadmap built around your goal, 1,000+ StationX courses for life plus 30,000+ in the Nexus library during the intensive year, your Success Manager, a weekly mastermind group, the Inner Circle, 24/7 expert support, supervised real work experience, and CV, LinkedIn and resume support.