


Fellowship Program in AI Engineering

Ship real AI apps. **Not notebooks.**



6
months

16+ apps
deployed

Industry
mentors

1000+ hiring
partners

Live
cohort

Learn by doing.

Crio.Do is India's work-experience-based learning platform. You don't watch a course — you join a build team, ship real software with mentors who review it like a lead would, and graduate with a portfolio that hiring managers can open.

Founded by ex-Flipkart and ex-Google engineers. Backed by Y Combinator.

95%

placement rate

1000+

hiring partners

81%

average salary hike

60+

live sessions with
industry mentors



index

THE ROLE

- 04** The role
- 05** In demand
- 06** The gap
- 07** What an AI Engineer does
- 08** Where the pay lands

THE PROGRAM

- 09** The program
- 10** Why Crio
- 11** Four pillars
- 12** Your journey
- 13** How a sprint works
- 14** Full curriculum

THE BUILD

- 15** What you'll build
- 16** Six production apps
- 17** Ten mini projects
- 18** Your capstone
- 19** Tools you'll master
- 20** Everything included

THE OUTCOME

- 21** Your outcome
- 22** Career support
- 23** Where graduates work
- 24** Learner stories
- 25** Your certificate

JOIN

- 26** Who it's for
- 27** What you'll need
- 28** How to enrol
- 29** Let's talk



01 The role

Every product team is adding LLM features. Very few engineers can ship them reliably.



The fastest-growing hiring segment of 2026.

33%

year-on-year growth in AI Engineering job postings in India

Naukri JobSpeak, July 2026 (Info Edge)

₹15–50 LPA

AI Engineer pay in India – ₹15–50 LPA at 3–9 years' experience; senior and lead roles at the top of the range • AmbitionBox, Sept 2026

5

roles this program prepares you for

BUILD

DEPLOY

OPERATE



AI Engineer

Builds and operates AI systems end to end – retrieval, agents, evals, serving



Agent Engineer

Agents that act safely over real systems – tools, MCP, guardrails



Applied AI / Solutions Engineer

Fits AI to a business problem – PoC to production



GenAI / LLM Application Engineer

LLM apps and RAG over real company data



Forward Deployed Engineer (AI)


Takes AI into a customer's organisation and owns the outcome



95% of GenAI pilots never reach the P&L.

MIT Project NANDA, The GenAI Divide 2025. The demo works.
Production is where it breaks.

The gap isn't prompting. It's retrieval that stays accurate, evals that catch regressions before users do, agents that act safely over real systems, cost that survives scale, and operations that keep it all running. That gap is where the roles, and the pay, are.




The five stepping stones this program is built around.




Frame it. Build it. Prove it. Run it.

One loop, practised sixteen times — every mini project, every guided app, and your capstone.



Where the pay lands.

AI Engineer pay in India, by experience – reported averages from AmbitionBox (4,700+ salaries, Sept 2026); senior and lead roles reach the top of the range.



Production AI skills are paid like senior engineering skills – from the third year on.

Averages by band from AmbitionBox (4,700+ reported AI Engineer salaries, Sept 2026); senior and lead roles reach ₹50 LPA at the top of the range. Individual offers vary.



02 The program

Built like a job, not a course. Six 4-week sprints, a capstone you present on demo day, and a 2-week interview sprint to finish.



Learn by doing. Literally.

Four things every sprint is built around.




Build-first curriculum

You ship a deployed app every sprint – concepts are taught in the order you need them.



Industry mentors

Practising AI engineers teach the live sessions and run the final demo day.



16+ real apps

Ten mini projects, six guided production apps and a capstone you own.



Lab stack & cloud sandbox included

AWS sandbox, LLM API access and the lab stack, provisioned before Day 1 – free tiers wherever sufficient.



Built on four pillars.

The structure every sprint follows – with a problem-to-production layer that most courses skip.

Fellowship Program in AI Engineering

PILLAR 1 · CORE

Building & deploying AI apps

LLM foundations, grounding with data (RAG), agentic systems, evaluation-driven development, operating in production, ML foundations.

PILLAR 2 · RECAP

Software-engineering fundamentals

System design, data, security & scaling – recapped for a software audience and deepened via LLD/HLD electives.

PILLAR 3 · A FULL MONTH

Using coding agents

Plan→execute→deploy workflow, autonomy & safe permissions, md-files/hooks/MCP, agentic code review, large-codebase work.

PILLAR 4 · FDE LAYER

Shaping the build

Discovery & scoping, spec & success metrics, PoC→MVP→production, demo storytelling, owning outcomes.



Six sprints. One capstone. One demo day to finish.

Every 4-week sprint ends with an app you deployed and evaluated – and the program ends with demo day.

Foundations **SPRINT 1**

How AI apps differ, LLM building blocks, evals 101 – with an ML refresher built in.

01

RAG **SPRINT 2**

Ground models in real data: hybrid search, re-ranking, retrieval evals.

02

Agents & MCP **SPRINT 3**

Agents that act safely over real systems, with your own multi-tenant MCP server.

03

Evals & fine-tuning **SPRINT 4**

Evals as CI gates, LoRA/QLoRA fine-tuning, and your first coding-agent workflow.

04

Production & voice **SPRINT 5**

Observability, cost control, compliance, AWS go-live and a real-time voice agent.

05

Capstone **SPRINT 6**

Ship a feature into a large codebase with coding agents – then your own AI product, end to end.

06

Interview sprint **WEEKS 25-26**

Agentic-system-design and behavioural rounds – all with scored feedback.

★

Sprint contents may be re-sequenced as we finalise the curriculum with subject-matter experts.



Four weeks. One shipped app.

The rhythm every sprint follows.

WEEK 1

Learn & warm up

Live sessions on the sprint's concepts, plus the sprint's mini projects.

WEEK 2

Build

Start the guided production app on the lab stack.

WEEK 3

Build & review

Automated assessment against the rubric; mentor feedback in the live sessions.

WEEK 4

Harden & assess

Eval gate in CI, cost and latency checks — then automated assessment.



A typical week

≈ 5 h
live

— 10 two-hour sessions per sprint with practising AI engineers

≈ 5 h
building

— about 20 hours of take-home builds per sprint, assessed automatically

Rhythm shown is indicative; session days are confirmed at cohort start.



Want the sprint-by-sprint curriculum?

Every sprint, every session theme, every build – on the detailed curriculum page.

View Curriculum →



scan

crio.do/ai-engineering-detailed-curriculum



What you'll build

16+ apps. Every one deployed, evaluated and monitored
– and live on a URL a hiring manager can open.



Built with mentors. Deployed by you.

Each one ships with an eval report, a cost/latency dashboard and a demo.



AskDocs

RAG

A workspace assistant that answers with citations over real company docs.



TriageOps

AGENTS

Multi-agent support operations over a real helpdesk, with a human in the loop.



DomainCopilot

FINE-TUNING

A fine-tuned vertical copilot, measured against a frontier API.



VoiceMate

VOICE


A real-time voice agent with tool use and interrupt handling.



Shipt

CLOUD

An earlier build taken live on AWS with monitoring and cost alarms.



RepoAgent

CODING AGENTS

A real feature shipped into a large open-source codebase with coding agents.

+1


Your capstone — your own AI product, end to end.

Pick a problem, write the spec, build, deploy on AWS, evaluate — and tell the story on demo day.



Small builds that teach the muscle.

Each one ends in a working, deployed artifact and a measured result — before the sprint's guided app.



Ship

a working, deployed artifact



Measure

against the golden set in CI



Assess

automated assessment against the rubric

Every build ships with a lab kit: starter + golden repo, eval harness + golden set, CI auto-grading, rubric, spec & demo templates.



One product you own.

The final sprint. Choose from an industry menu or bring your own problem — then take it all the way to demo day.



Pick from the menu — or bring your own

- ✓ Deep-research assistant with citations
- ✓ CRM record enrichment over a real CRM
- ✓ Payments reconciliation & anomaly triage
- ✓ API documentation & test-suite generator
- ✓ Engineering-delivery risk agent over an issue tracker
- ✓ Data-to-marketing-video generator

Ships as a product you own — spec, code, eval report, live URL and a demo video.




Demo day

Your app live, your eval report on screen, and questions from mentors and peers — the way you'd present to stakeholders.



Everything in orbit around **you.**



Build & serve Models & LLMs Agents, retrieval & evals Ship & integrate
 Concept-first — tools are current reference implementations, swapped as the ecosystem evolves.



You don't buy a thing.

The tools you build with and the cloud you deploy to are provided for the whole program.



AWS cloud sandbox



LLM API access



Vector DB, evals & observability tooling



Voice & deployment



A lab kit for all 17 builds

Plus the standard lab stack and a step-up reading library. **We use free tiers and resources wherever sufficient.**



04 Your outcome

A portfolio that reads like work experience – and the career support to turn it into the next role.



Four directions. One destination.

A 2-week interview sprint, 5 referral interviews, a reviewed profile and industry mentors.

Mock interviews

agentic-system-design & behavioural rounds, with scored feedback

Industry mentors

practising AI engineers who teach every sprint and run the final demo day



5 referral interviews

to SDE, AI Engineer and similar roles

Profile that reads credibly

mentor-reviewed résumé, LinkedIn and GitHub



Where Crio graduates **work.**

Dream jobs at product companies and super-dream jobs at unicorns and tech giants – across Crio's fellowship programs to date.



Hear from our learners.



The AI modules completely changed how I think about development. I went from writing code to architecting systems. The capstone project gave me a real AI product to showcase – it's what got me noticed at Flipkart.

Priya S • SDE at 

↑120%

salary hike



I was intimidated by LLMs, but the curriculum made it approachable. Hands-on with Claude, LangGraph, and building agents – not just watching videos. Three months after graduation, I landed a backend role that explicitly values AI skills.

Rahul K • Full-Stack Dev at 

↑85%

salary hike



Through the Micro-Experiences I learned a bunch of industry-relevant skills. There were several modules where I was stuck, but the team at Crio always kept me motivated.

Ananya M • Backend Engineer at  **Meesho**


↑95%

salary hike



Proof you can put on LinkedIn.

A Crio certificate on completion – backed by a portfolio of deployed apps that proves it.



Sample – the issued certificate carries the fellow's name, the issue date and the signatories' signatures.



Built for engineers who want to **own the system.**

No ML or GenAI experience needed. You do need to be a working programmer.

- 01 Software engineers (2–5 yrs)**
Own AI systems end to end and lead the AI work on your team.

- 02 Senior engineers, leads & architects (5+)**
Decide what to build and prove it works — specs, evals, stakeholder demos.

- 03 ML engineers & data scientists**
Move from notebooks to shipped products: agents, serving, LLMOps.

- 04 BE / BTech graduates**
Add the AI-app layer to substantial projects you've already shipped.

Not sure you're a fit?

Talk to an advisor — we'll tell you honestly whether to start now or prep first.

[Talk to an Advisor](#)

A working programmer. Nothing more.

No ML, deep-learning or LLM experience required – the refreshers are built in.



Python

functions, classes, typing, packaging



Git & GitHub

branching, pull requests, basic CI



HTTP, REST & JSON

calling and building a simple API



SQL basics

joins, aggregates, relational schemas



Command line & Docker

the basics – the lab stack runs in containers



About 10 hours a week

for 6 months, plus the 2-week interview sprint

Before Day 1 · pre-work ≈ 8–10 hours, self-paced

Lab-stack setup, a Python-for-production-AI refresher and an ML concepts primer – all provided with your lab kit.



Three steps to Day 1.

1

Talk to an advisor

Fill the form on the program page — an advisor calls you back for an honest read on fit.

2

Confirm schedule & fees

Start date, weekly commitment and the fee — decided together.

3

Enrol & get your lab kit

Cloud sandbox, API keys and the lab stack — provisioned before Day 1.

PROGRAM FEE

₹3,99,000

MODE

Live online, cohort-based • ~10 hours a week

DURATION

6 months + 2-week interview sprint

[Talk to an Advisor →](#)

scan





#LearnByDoing

Let's build something real.

Speak with a program advisor for an honest read on fit, schedule and fees.

[Talk to an Advisor →](#)



scan

programs@criodo.com | www.crio.do | 080424 53955