

humaine labs

AI Engineering Foundation Enterprise AI FDE Programs

The background features a collage of digital elements. On the left, there are snippets of Python code, including a function `def frostain()` and a class `class Doctor`. In the center, a man and a woman are smiling and looking at a tablet. To their right, there are logic gate diagrams (AND, OR, NOT) and a neural network diagram. Below the neural network, the words 'TRAINING' and 'DEPLOYMENT' are visible. At the bottom, there is a list of names: 'D1', 'C2', 'MM', 'MODEL', 'remavMeasuret', 'Beat.Acordindeessing', 'Algorithm tkia', 'Singoaneotnenone', 'Ashacuard', 'remavMeasuret', 'Beat.Acordindeessing', 'Algorithm tkia', 'Singoaneotnenone', 'Ashacuard'. The overall theme is AI engineering and enterprise solutions.

Humaine's Programs are trusted by **50+ Enterprises**
and **3000+ Executives.**

Lenovo

Capgemini

marico

LARSEN & TOUBRO

LANDMARK
GROUP

aether



AI hasn't just made engineers faster. It's fundamentally changed what engineers are expected to do.

The most in-demand engineers today walk into a client's environment, understand their operational problem, and build an AI solution that works inside their real constraints.

They're not just writing code. They're shaping what gets built and for whom. This is already in the job descriptions. Engineers who see this shift and move early are the ones pulling ahead.



"90% of Fortune 100 companies have adopted AI coding assistants — and developers using them complete tasks up to 55% faster. The productivity gap between AI-capable teams and those without is no longer marginal."

McKinsey & Company

"88% of organisations now use AI in at least one business function — but only 6% qualify as high performers seeing significant enterprise-wide value. The gap is widening."



While you're waiting, your competitors are compounding.

Most engineering organisations are aware of AI. Far fewer are building it into how their teams actually work. That gap has a price, and it's growing every quarter.

amazon

Deployed AI coding tools internally and **saved an estimated \$260 million and 4,500 developer-years** of effort on a single migration project.

J.P.Morgan

Generates \$2 billion in annual savings, with engineering teams seeing 10–20% efficiency gains from internal AI coding assistants.

tcs Infosys wipro

TCS, Infosys, and Wipro have collectively **deployed AI tools to over 300,000 engineers**, with Wipro reporting adoption only peaked after running mandatory hands-on workshops.

JuganU

Used AI-assisted engineering to launch production systems in 6 months instead of 2 years. A delivery timeline their traditional engineering process simply couldn't achieve.

The Workforce Dilemma

Cut engineers to save cost, and you bleed brand trust and knowledge. Hire external AI talent to move fast, and you're bidding for a role the market barely supplies.



Scenario 1

Laying Off Non-AI Engineers

245k+

Tech jobs cut globally in 2025

- Brand and morale damage — customers and staff lose confidence
- Institutional knowledge walks out the door with every exit
- Rebuilding the team costs more than retaining now



Scenario 2

Hiring External Talent

800%

Growth in Forward Deployed Engineer job postings in a single year

- Average AI Engineer salary hits \$206k+, roles unfilled for months
- New hires take 6-12 months to learn systems and stakeholders
- Not enough qualified AI Engineers and FDEs exist in the market to hire

The Alternative



Scenario 3

Upskill the engineers you already trust with Humaine's Programs

Your engineers already know your business, systems, and stakeholders. Humaine's AI Engineering Foundation and Enterprise AI FDE Certification give them the missing skillset — so they ship AI solutions at speed, with zero workforce disruption.

No Layoffs

No Slow ramp-up

Existing Team Knowledge

Ship in weeks

Meet the **AI Engineer** and the **Forward Deployed Engineer**.


The two most in-demand technical roles in the AI era.

An **AI Engineer** builds and ships AI-powered products. They work with generative and agentic AI, design workflows, and take ideas from prompting to production. This is the foundation role, and it's already in high demand.

A **Forward Deployed Engineer** takes that a step further. They embed directly with a client, understand their operational problems firsthand, and deploy AI solutions inside their real environment — with real constraints, real stakeholders, and real accountability for outcomes. It's the advanced seat. The one that sits at the intersection of engineering and delivery.

Humaine's two programmes mirror this path exactly. The AI Engineering Foundation builds you as an AI Engineer. The Enterprise AI FDE Certification takes you into the field.

Companies Currently hiring AI Engineers & FDEs



AI Engineer vs Enterprise AI Forward Deployed Engineer (FDE)

AI Engineers build AI products. Enterprise AI FDEs drive AI transformation.

Metric	AI Engineer	Enterprise AI FDE
Primary Mission	Build and ship AI-powered products	Identify opportunities, deploy AI, and drive transformation inside enterprise environments
Core Focus	Technical implementation	Business problem-solving + technical delivery
Stakeholders	Engineering team, product managers	C-suite, business unit leaders, IT, end users
Success Metrics	Working product, code quality, performance	Business outcomes, ROI, adoption, stakeholder buy-in
Work Environment	Internal product or engineering team	Embedded at client sites, cross-functional teams
Key Skills	LLMs, RAG, agents, APIs, deployment	All AI Engineer skills + discovery, architecture, enterprise integration, stakeholder management
Experience	Basic Python and Git Familiarity	AI Engineering Foundation complete, or equivalent production AI experience
Career Path	AI Engineer → Senior AI Engineer → AI Architect	FDE → Enterprise AI Consultant → AI Solutions Architect → AI Transformation Leader

Wondering how to get
Started?

Reach out to learn@humainelabs.com to get in touch with an expert.



AI Engineering Foundation

6 Modules. Two 2-hour sessions per week
delivered for 6 weeks. 2 Week Capstone.



Bring AI into your Engineering Workflows

Reach out to learn@humainelabs.com to get in touch with an expert.

AI Engineering Foundation

Program Snapshot

What Gets Delivered

This isn't a program your engineers complete and move on from. Every week, they ship something real. By the end, your organisation has engineers who can build across generative AI, agentic workflows, and production-ready systems – and immediately apply that capability to live problems inside your business.

Key Takeaways

- ✓ Practical fluency across the AI engineering stack: generative AI, agentic systems, workflow design
- ✓ Engineers who ship working AI systems every week – not just complete coursework
- ✓ A team that can immediately contribute to live AI initiatives inside your organisation
- ✓ The Humaine Certified AI Engineer credential graded against a real competency standard

How It's Taught

Live working sessions run by practitioners who have built and deployed AI at scale. Not lectures. Not pre-recorded content. Sessions where engineers build alongside the instructor, ask questions in real time, and leave with something shipped. Assessed by AI leaders – not algorithms.



* Learners should get their own OpenAI and Anthropic API keys.

AI Engineering Foundation

Week by Week Breakdown

Week 1

LLM Foundations & Prompting

Understand what an LLM is, make cross-provider API calls, reason about tokens, context windows, latency and cost. Master core prompting and context-engineering patterns.

Deliverable: A CLI script that answers a question via an LLM and prints token usage and estimated cost.

Week 2

Structured Outputs & Embeddings

Make model output reliable and machine-usable with typed schemas (Pydantic). Learn the representation that powers search and RAG — embeddings, cosine similarity, and semantic search.

Deliverable: A semantic-search CLI over a small corpus that prints ranked results and scores.

Week 3

Retrieval-Augmented Generation (RAG)

Build a working RAG pipeline that answers questions over your own documents with source attribution — and measure whether it is actually good. Covers chunking, vector databases, hybrid search, and RAGAS evaluation.

Deliverable: A RAG Q&A script answering 5 questions with source attribution, plus a RAGAS score table.

Week 4

Tools & the Agent Loop

Cross from "text in, text out" to agents that take actions. Covers tool/function calling, the ReAct loop (reason + act), and the Model Context Protocol (MCP) — connecting agents to real systems.

Deliverable: An agent that correctly answers a question requiring a live tool call, with a visible reasoning/tool-call log.

AI Engineering Foundation

Week by Week Breakdown

Week
5

Building Agents with a Framework

Build real, stateful agents and multi-agent workflows using LangGraph and CrewAI instead of hand-rolling everything. Covers state, memory, control flow, multi-agent roles, handoffs, and reflection loops.

Deliverable: A multi-agent pipeline that outputs a report with a visible handoff and at least one revision cycle.

Week
6

Reliability, Safety & Deployment

Make agents trustworthy and shippable. Covers evaluation, observability (Langfuse), guardrails, prompt injection basics, cost control, model routing, and deployment via FastAPI/Streamlit and Docker.

Deliverable: A running app (API or UI) with at least one guardrail, plus a cost-per-request analysis.

Weeks
7
&
8

Capstone: Build & Ship

Design, build, evaluate and demo a complete agentic GenAI application. Required components: RAG backend over a real corpus, an agent using at least two tools (one via MCP), structured outputs, evaluation, tracing, a guardrail, and a UI or API runnable with one command.

Deliverable: Live demo + a tidy Git repository



Enterprise AI Forward Deployed Engineer (FDE) Certification

6 Modules. Two 2-hour sessions per week delivered
for 6 weeks. 3 Week Capstone.



Enter the World of AI Solutioning

Reach out to learn@humainelabs.com to get in touch with an expert.

Enterprise AI FDE Certification

Program Snapshot

What Gets Delivered


The Enterprise AI Forward Deployed Engineer (FDE) Certification starts where the AI Engineering Foundation ends. You take a working AI product into the field. That means client discovery, understanding operational problems at ground level. Enterprise AI integration — working within the security, governance, and infrastructure constraints of a real organisation. And a 4-week capstone: from discovery to a fully deployed, governed AI system with an executive transformation proposal.

Key Takeaways

- ✓ An Enterprise AI Transformation Proposal — executive summary, business case, ROI, deployment roadmap
- ✓ A completed capstone: a real deployed, governed AI system built for a real enterprise environment
- ✓ The Humaine Certified FDE credential — Certified, Professional, or Distinction

How It's Taught

Live working sessions run by practitioners who have built and deployed AI at scale. The capstone is assessed by AI leaders through a mock FDE interview panel (20-minute presentation + 15-minute Q&A).



* Learners should get their own OpenAI and Anthropic API keys.

Enterprise AI FDE Certification

Week by Week Breakdown

Week 1

Enterprise AI Landscape & the FDE Mindset

Understand the FDE role and how AI transformation differs from traditional software delivery. Start a real engagement: select and charter your capstone. Covers discovery frameworks, stakeholder mapping, success metrics, and ROI framing.

Deliverable: Project Charter — problem definition, stakeholders, business objectives, success metrics.

Week 2

Enterprise Discovery & Solution Architecture

Map current-state workflows, identify bottlenecks and AI opportunities, then design an enterprise-grade AI solution architecture. Covers agentic design patterns, Architecture Decision Records (ADRs), and effort estimation.

Deliverable: Enterprise Discovery Report + Solution Blueprint (future-state workflow, agent architecture, user journeys).

Week 3

Vibe Coding & Rapid Prototyping

Build credible, working software fast using AI-assisted coding workflows (Cursor, Claude Code). Learn the discipline that keeps speed from shipping fragile code — spec-driven prompting, rigorous review, and automated testing. Turn a prototype into a customer-ready POC.

Deliverable: Working prototype slice of the core use case + a scripted demo.

Week 4

Enterprise Knowledge & Integration

Build the enterprise knowledge foundation: advanced RAG, knowledge graphs, ontologies. Connect the solution to real enterprise systems — CRM, ERP, APIs — via custom connectors and MCP servers, behind authentication.

Deliverable: Knowledge Architecture Package + Integration Architecture (system landscape, data flows, integration blueprint).

Enterprise AI FDE Certification

Week by Week Breakdown

Week 5

Building & Evaluating Production AI Systems

Develop a robust, enterprise-ready solution with memory architectures, reliability engineering (timeouts, retries, fallbacks), evaluation frameworks (RAGAS, LLM-as-judge), observability (Langfuse), and cost engineering.

Deliverable: Working MVP — functional prototype, agent workflows, enterprise integrations, evaluation results.

Week 6

Deployment, Governance & Enterprise Value

Make the solution deployable, governed, and compliant. Covers containers, Kubernetes, Terraform, CI/CD, security architecture, PII redaction, RBAC, audit logging, EU AI Act compliance, and executive communication of ROI and business value.

Deliverable: Deployment Readiness Package + first draft of the Enterprise AI Transformation Proposal.

Weeks 7 to 10

Capstone: AI Transformation Challenge

Four weeks of working sessions building, integrating, securing, deploying, and presenting a complete enterprise AI solution.

- ✓ Week 7: Build the core agentic workflow and knowledge/data layer
- ✓ Week 8: Integrate with enterprise systems + customer demo and iteration
- ✓ Week 9: Security, governance, compliance controls + production deployment (containers, IaC, CI/CD)
- ✓ Week 10: ROI quantification, executive transformation proposal + final panel presentation (20 min + 15 min Q&A)

Deliverable: Live demo + a tidy Git repository

Tools Covered in the Program

 OpenAI

 Claude

 docker.

 crewAI

 LangChain

 CURSOR

 LangGraph

 GitHub

And more tools...

 **kompass**
by dotkonnekt

Learn how to build AI Solutions
in a cutting edge all in one
Agentic Enterprise AI Tool.

Wondering how to get Started?

Reach out to learn@humainelabs.com to get in
touch with an expert.



The Roles These Programs Power

AI Engineering Foundation

AI Engineer

GenAI
Engineer

Agent
Engineer

LLM
Engineer

AI
Application
Developer



Enterprise AI FDE Certification

Forward
Deployed
Engineer

Enterprise
AI
Consultant

AI
Transformation
Consultant

AI
Delivery
Lead

AI
Solutions
Architect

AI
Engagement
Manager



What Organisations Walk Away With

A team that ships, not just learns.

Every participant delivers a real working product each week. By programme end, your organisation has a portfolio of production-grade AI systems. RAG pipelines, agentic workflows, deployed applications — built by your engineers, on your problems.

AI capability built from within

No lengthy hiring cycles. No expensive external contractors. The AI Engineers you need are already inside your organisation. The Foundation programme turns them into practitioners in two months.

Your engineers ship finished AI products at speed.

AI-capable engineering teams move faster, iterate faster, and deliver more without adding headcount. Training your existing engineers is the highest-leverage investment your organisation can make right now.

Engineers who can talk to the business.

FDE graduates leave able to run discovery sessions, present to C-suite, and communicate technical decisions to non-technical stakeholders, closing the gap between your engineering and business teams.

High Impact AI Learning from Industry Practitioners

At Humaine, we believe AI cannot be taught the old way. It must be experienced through projects, practitioners, and purposeful collaboration.

We are a practitioner-led learning ecosystem built to help professionals and organizations thrive in the AI era.

Our approach blends physical classrooms, live online cohorts, and real-world problem solving to make AI adoption simple, structured, and scalable.



Practitioners – builders who've shipped AI systems in the wild.



Cohort-based and blended –immersive, interactive, and outcome-driven.



Real-world workflows, not tool demos – every concept tied to tangible business outcomes.



Wondering how to get Started?

Reach out to learn@humainelabs.com to get in touch with an expert.



AI Engineering Foundation

8 weeks - 32 contact hours. Build practical fluency in generative and agentic AI — from prompting to a shipped product.

Prerequisites

Basic Python & Git familiarity · Your own OpenAI and Anthropic API keys

- ✓ Generative AI: prompting, structured outputs, RAG
- ✓ Agentic AI: tool calling, MCP, LangGraph, CrewAI
- ✓ Evaluation, guardrails, and deployment basics
- ✓ Capstone: ship a complete agentic GenAI app
- ✓ Humaine Certified AI Engineer credential

Enroll Before Seats Fill



Enterprise AI FDE Certification

10 weeks - 40 contact hours. Take a working AI product into a real enterprise environment.

Prerequisites

Completion of AI Engineering Foundation, or pass the entry assessment · Your own OpenAI and Anthropic API keys

- ✓ Enterprise discovery, workflow analysis, and solution architecture
- ✓ Vibe coding (Cursor, Claude Code) for fast client-side prototyping
- ✓ Enterprise integration via custom MCP servers
- ✓ Cloud-scale deployment, governance, and compliance
- ✓ 4-week capstone: discovery to deployed, governed system
- ✓ Humaine Certified FDE: Certified, Professional, or Distinction
- ✓ 10 guaranteed job offers per FDE cohort

Enroll Before Seats Fill

humaine labs

Powered by



www.humainelabs.com