


FDE PRO

AI Forward Deployed Engineer Program

Enterprise AI, Cloud & Customer Engineering

Learn to take an AI solution from a customer's problem to production. Discovery, architecture, RAG, agents, MCP, enterprise integration, cloud deployment, security and measurable business value, in twelve weeks.



₹30,000

Course fee

3 months

12 weeks, 120+ hours of
live training

5 + 1

Enterprise projects and a
capstone engagement

Hybrid

Classroom in Ameerpet or
live online

Build AI. Deploy AI. Solve enterprise problems.

Program at a glance

FDE PRO is a 3-month fast-track program for one of the most important roles in enterprise AI: the engineer who takes an AI solution from a customer’s problem to production. You learn the engineering (Python, APIs, cloud, RAG, agents, MCP, Kubernetes, Terraform) and the customer work (discovery, architecture, security reviews, ROI and executive demos) together, every week.

Program	FDE PRO: AI Forward Deployed Engineer Program
Duration	3 months (12 weeks)
Course fee	₹30,000
Format	Classroom at Ameerpet, Hyderabad, or live online
Sessions	5 sessions a week, 2 hours each: 4 engineering sessions and 1 Customer Engagement Lab
Training	120+ hours live, 60+ hands-on labs, about 70% hands-on
Projects	5 enterprise projects and the GlobalBank capstone engagement
Primary cloud	AWS, with working knowledge of Microsoft Azure and Google Cloud
Prerequisite	Basic programming in any language. No prior AI experience needed
Outcome	Job-ready for FDE, applied AI, AI solutions and agentic AI engineering roles
Placement	Resume, portfolio, mock interviews and placement assistance (5,500+ alumni placed)

Contents

What is a Forward Deployed Engineer?	3
FDE vs other engineering roles	4
How FDE PRO works	5
The 12-week route	6
Week-by-week curriculum	7
Enterprise projects	13
Capstone: the GlobalBank engagement	14
Tools and skills	15
Careers and placement	16
Fee, batches and enrollment	17
Frequently asked questions	18

The FDE PRO promise

Learn to take an AI solution from a customer’s problem to production.

Every project in this program has to answer six questions, starting with “What customer problem does this solve?” and ending with “What measurable business value did the customer get?”

What is a Forward Deployed Engineer?

Definition

A Forward Deployed Engineer (FDE) is a software engineer who works directly with customers to turn their business problems into working, production-grade solutions. In AI, an FDE runs discovery, designs the architecture, builds with LLMs, RAG and agents, integrates enterprise systems, deploys to the cloud and proves measurable business value.

The role was popularized by Palantir, whose engineers were “forward deployed” to work inside customer organizations. Today AI companies, enterprise software platforms and consulting firms use FDEs to close the gap between a promising AI demo and a system the business actually runs on.

An FDE is not simply a developer and not simply a consultant. An FDE takes an ambiguous business problem, converts it into a technical solution, builds it quickly, integrates it with enterprise systems, deploys it into the customer’s environment, measures how well it performs and stays with it until it runs in production.

Why companies are hiring FDEs now

Building an AI demo takes an afternoon. Getting it into production inside a bank or a hospital takes discovery, integration, identity, security reviews, evaluation and someone who owns the outcome. That gap between a promising pilot and a system the business runs on is where Forward Deployed Engineers work.

FDEs exist to clear the three blockers that stop most AI pilots:

Messy reality	The data lives in ServiceNow, SAP, SharePoint and spreadsheets, behind old APIs, proxies and firewalls.
Security says no	SSO, least privilege, private networking, audit trails and prompt-injection defenses must be in place before go-live.
No proof of value	Without evaluation scores and ROI numbers, pilots never get funded for production.

Current FDE job descriptions ask for exactly this mix: Python, LLMs, RAG and agents, enterprise integration, cloud deployment, security, evaluation and customer discovery, with ownership from the first conversation to post-go-live support.

The FDE way of working

Understand	Design	Build	Integrate	Deploy	Observe	Improve	Deliver value
------------	--------	-------	-----------	--------	---------	---------	---------------

Understand, design, build, integrate, deploy, observe, improve and deliver business value. FDE PRO is organized around this lifecycle.

FDE vs other engineering roles

The FDE role overlaps with many familiar jobs. What makes it different is ownership: an FDE is customer-facing, writes production code and is accountable for the business result in production.

Role	Main job	Customer-facing	Writes production code	Owns
Software Engineer	Build product features	Rarely	Yes	The feature
AI / ML Engineer	Build and tune models and pipelines	Sometimes	Yes	Model quality
Solutions Engineer (pre-sales)	Demo the product and win the deal	Always	Prototypes	The sale
Solutions Architect	Design the target architecture	Often	Sometimes	The design
DevOps / SRE	Run platforms reliably	Rarely	Yes, infrastructure	Uptime
Consultant	Recommend what to do	Always	Rarely	The recommendation
Forward Deployed Engineer	Make it work in the customer's world	Always	Yes	The business result in production

What makes FDE PRO different from a typical AI course

Most AI courses teach a tool list and stop at a demo. FDE PRO makes the customer engagement lifecycle the backbone, and teaches every technology through the customer problem it solves.

- Built around the customer, not the tool list**
 The customer engagement lifecycle is the backbone. Every technology is taught through the customer problem it solves.
- Production, not demos**
 Identity, security reviews, evaluation, observability, cost and incident response are part of every project.
- A weekly Customer Engagement Lab**
 Twelve labs of discovery calls, design documents, security reviews, RCAs, SOWs and executive demos.
- You graduate with a delivery kit**
 Your FDE Accelerator repository, five projects and a full capstone engagement to show employers.

How FDE PRO works

Eight stations, from customer problem to business value

1 Understand Discovery with stakeholders: the real problem and how success is measured.	2 Design Architecture, data flow, security model and trade-offs.	3 Build Production Python, APIs, RAG and agents.	4 Integrate Enterprise systems, MCP tools and the customer's identity.
5 Deploy Containers, Kubernetes, Terraform and CI/CD.	6 Observe Traces, metrics, evaluation scores and cost per request.	7 Improve Fix, tune, harden and cut cost.	8 Deliver value ROI, executive demo and a clean handover.

Twin-track learning

Each week runs on two tracks. The engineering track builds the technical depth; the Customer Engagement Lab makes you use it the way an FDE does, in front of a “customer”.

Engineering track

Four sessions a week

Python, APIs, cloud, RAG, agents, MCP, Kubernetes, Terraform, security and observability, taught hands-on with a lab in every session.

Customer Engagement Lab

One session a week

You apply that week's technology to a simulated customer situation: discovery calls, workflow maps, design documents, security reviews, incident updates, SOWs and executive demos.

Six questions every project must answer

- 1 What customer problem does this solve?
- 2 How would you deploy it inside a real enterprise?
- 3 How do you prove it works?
- 4 What is it allowed to do?
- 5 How will you secure, monitor and operate it?
- 6 What measurable business value did the customer get?

Your FDE Accelerator

Every lab adds a reusable component to your own FDE Accelerator repository. By week 12 you own a delivery kit: discovery templates, service and agent templates, an MCP server, a Terraform module, a CI/CD pipeline, an evaluation harness, dashboards, runbooks and an ROI calculator. It is how real FDE teams work: reuse what is proven, then adapt it to each customer.

The 12-week route

Three months, three phases. Each phase ends with working software in the cloud, and each week ends with a lab, a Customer Engagement Lab and a new component in your FDE Accelerator.

Month 1: Build Engineering & cloud foundation Write production software and ship it to the cloud. W1 The FDE mindset and your developer toolkit W2 Production Python W3 APIs, data and SQL W4 Cloud and containers: AWS first, Azure and Google Cloud ready	Month 2: Integrate Enterprise AI engineering Make AI useful inside real enterprise systems. W5 GenAI and LLM engineering W6 RAG engineering and evaluation W7 Agentic AI: tools, LangGraph and multi-agent systems W8 MCP, enterprise integration and identity	Month 3: Deploy & deliver Production & customer engineering Run it in production and prove the business value. W9 Kubernetes, Terraform and CI/CD W10 AI security, governance and enterprise networking W11 Observability, cost and troubleshooting W12 Customer engineering and capstone delivery
--	---	---

What you add to your Accelerator each week

W1	Repository skeleton, README and discovery templates
W2	Resilient API-client library with retries, logging and config
W3	FastAPI service template with data-quality checks
W4	Dockerfile, AWS deployment template and HLD template
W5	Multi-model AI gateway
W6	RAG module and evaluation harness
W7	LangGraph agent template with an approval step
W8	MCP server template with OAuth and OIDC
W9	Terraform module, Helm chart and CI/CD pipeline
W10	Guardrails, red-team test suite and security-review template
W11	Observability dashboards, runbook and RCA template
W12	SOW template, ROI calculator and executive demo script

How a week runs

Sessions 1–4	Engineering track: concept, live build, hands-on lab
Session 5	Customer Engagement Lab: apply the week in a simulated customer situation
Practice	8–10 hours of guided lab practice outside class
Output	Working code, a customer-facing document and a new Accelerator component

Milestones

Project 1 in week 4, Project 2 in week 6, Project 3 in week 8, Project 4 in week 9, Project 5 in week 10. The GlobalBank capstone runs as four sprints across weeks 9 to 12.

Month 1**Build: Engineering & cloud foundation**

Weeks 1–4

Write production software and ship it to the cloud.

Week 1**The FDE mindset and your developer toolkit**

Month 1: Build

Understand what FDEs do all day, set up a professional workstation and make your first LLM call.

The FDE role	how it differs from software, AI, DevOps and solutions engineers, architects and consultants; pre-sales vs post-sales; POC, pilot, MVP and production
Enterprise vocabulary	stakeholders, champion, decision maker, procurement, security review, architecture review board, UAT, SOW, SLA and SLO, success criteria
Linux for FDEs	filesystem, processes, services and systemd, ports, DNS, SSH, permissions, environment variables and logs; grep, awk, sed, curl and jq
Git and GitHub	branching, pull requests, merge vs rebase, tags and releases, keeping secrets out of repositories
Your first LLM call	call a model API from Python and get structured JSON back on day one

Hands-on lab

Set up the FDE workstation (VS Code, Python virtual environments, Git, CLI tools) and publish a first-LLM-call repository with a README a customer could follow.

Customer Engagement Lab

Decode a vague request: turn “we want AI for support” into a problem statement, a stakeholder map and success criteria.

Adds to your Accelerator: Repository skeleton, README and discovery templates**Week 2****Production Python**

Month 1: Build

Write Python that survives real customer environments: typed, tested, logged and resilient.

Core to advanced	functions, OOP, modules and packages, exceptions, JSON, YAML and CSV processing, regular expressions
Modern Python	type hints, dataclasses, decorators, generators, iterators and context managers
Concurrency	async Python and HTTP clients; threads, processes and queues
Production patterns	structured logging, configuration and .env files, retries with backoff, timeouts, circuit breakers, idempotency
Testing and packaging	pytest, fixtures, mocking external APIs, dependency management

Hands-on lab

Build an enterprise API connector: a paginated, rate-limit-aware REST client with retries, tests and a command-line interface.

Customer Engagement Lab

Run a mock discovery call with the Why, What, Who, Where, When and How question bank, then write up the notes.

Adds to your Accelerator: Resilient API-client library with retries, logging and config

Week 3**APIs, data and SQL**

Month 1: Build

Expose clean, secure APIs over messy enterprise data.

FastAPI	routing, Pydantic models, dependency injection, middleware, background tasks, async endpoints, OpenAPI and Swagger, versioning
API engineering	API keys and JWT, pagination, rate limiting, caching, idempotency keys, webhooks; REST vs GraphQL vs SOAP
SQL FDEs use daily	joins, GROUP BY, CTEs, window functions, indexes, transactions and reading query plans
Databases	PostgreSQL and Redis hands-on; MongoDB overview; systems of record, OLTP vs OLAP, warehouse, lake and lakehouse
Data reality	duplicates, missing values, schema mismatch, PII, freshness and data contracts; pipeline basics with Kafka and Airflow concepts

Hands-on lab

Customer onboarding service: FastAPI, PostgreSQL and Redis with JWT auth, webhooks and an automated data-quality report.

Customer Engagement Lab

Map a business process (support-ticket handling) into a technical workflow diagram the customer can validate.

Adds to your Accelerator: FastAPI service template with data-quality checks**Week 4****Cloud and containers: AWS first, Azure and Google Cloud ready**

Month 1: Build

Package your service and deploy it securely on AWS, with working knowledge of Azure and Google Cloud.

AWS core	IAM roles and policies, VPC, subnets, security groups, EC2, S3, RDS, Lambda, API Gateway, CloudWatch, CloudTrail, Secrets Manager, KMS
Azure essentials	resource groups, VNets and NSGs, Functions, Storage, Key Vault, Azure Monitor
Google Cloud essentials	projects, IAM, VPC, Cloud Run, Cloud Storage, Secret Manager
Docker	images, Dockerfiles, multi-stage builds, volumes, networks, Compose, image scanning and registries (ECR, ACR, Artifact Registry)
Cost guardrails	budgets and alerts so your labs never produce a surprise bill

Hands-on lab

Ship it: Python service to Docker image to Amazon ECR to ECS Fargate behind a load balancer, with secrets in Secrets Manager.

Customer Engagement Lab

Write a one-page high-level design (HLD) with an architecture diagram and a network diagram.

Adds to your Accelerator: Dockerfile, AWS deployment template and HLD template**Project 1 delivered: Customer Integration Service**

Month 2**Integrate: Enterprise AI engineering**

Weeks 5–8

Make AI useful inside real enterprise systems.

Week 5**GenAI and LLM engineering**

Month 2: Integrate

Use LLMs like an engineer: predictable outputs, the right model for the job and a cost you can explain.

How LLMs work for engineers	tokens, context windows, embeddings, attention (intuition), inference, temperature and top-p
Prompt and context engineering	zero- and few-shot prompts, role prompts, structured JSON output, templates, chaining and context management
Models across clouds	Amazon Bedrock (Converse API), Azure OpenAI in Microsoft Foundry, Gemini on Google Cloud and open-weight models run locally
Model selection	quality, cost, latency, context length, tool calling and privacy; when a small model wins; prompting vs RAG vs fine-tuning (LoRA, PEFT) vs distillation
Token economics	input vs output tokens, caching, batching and streaming

Hands-on lab

Build a multi-provider AI gateway: one API, swappable models, with tokens, latency and cost logged for every request.

Customer Engagement Lab

Write a model-selection memo for a customer with data-residency constraints.

Adds to your Accelerator: Multi-model AI gateway

Week 6**RAG engineering and evaluation**

Month 2: Integrate

Build retrieval-augmented generation that answers with citations, and prove its quality with numbers.

RAG pipeline	parsing, chunking (fixed, recursive, semantic, parent-child), embeddings and vector stores: FAISS, Chroma, pgvector, OpenSearch
Retrieval quality	hybrid dense and keyword search, metadata filters, query rewriting, multi-query retrieval, reranking, context compression
Enterprise answers	citations and permission-aware retrieval, so users only see what they are allowed to see; managed RAG with Bedrock Knowledge Bases and Azure AI Search
Evaluation with Ragas	faithfulness, answer relevancy, context precision and recall; golden datasets; regression tests in CI
What comes next	Graph RAG and agentic RAG

Hands-on lab

Enterprise Knowledge Assistant: ingest policy documents, answer with citations and score the system with Ragas.

Customer Engagement Lab

Agree the acceptance criteria: an evaluation plan the customer signs before you build.

Adds to your Accelerator: RAG module and evaluation harness

Project 2 delivered: Enterprise Knowledge Assistant

Week 7**Agentic AI: tools, LangGraph and multi-agent systems**

Month 2: Integrate

Design agents that act, and that stop, ask and stay inside their permissions.

Agent fundamentals	workflows vs agents; the reason, act, observe loop; function and tool calling; structured tools
LangChain and LangGraph	models, prompts, tools and retrievers; state, nodes, edges, conditional routing, loops, checkpoints, threads, persistence and human-in-the-loop
Multi-agent patterns	router, supervisor, hierarchical and parallel specialists
Managed agent platforms	Amazon Bedrock Agents and AgentCore, Microsoft Foundry Agent Service, Google Agent Development Kit (ADK)
Agent guardrails	step limits, budgets, tool allow-lists and approval gates

Hands-on lab

IT-Ops multi-agent system: a supervisor routes to CloudWatch, Kubernetes and Jira agents and waits for human approval before any change.

Customer Engagement Lab

Write the agent permission spec: what it may read, what it may change and who approves.

Adds to your Accelerator: LangGraph agent template with an approval step

Week 8**MCP, enterprise integration and identity**

Month 2: Integrate

Connect AI to the systems enterprises actually run, under the customer's identity and access rules.

Model Context Protocol (MCP)	hosts, clients and servers; tools, resources and prompts; local (stdio) and remote (HTTP) servers; authentication and authorization
Build MCP servers	for a database, GitHub and ServiceNow; tool governance, RBAC and audit trails
Enterprise integrations	ServiceNow, Jira, GitHub, Slack, Microsoft Teams, Salesforce and SharePoint; webhooks and legacy SOAP; API limits, schema mapping, private endpoints, proxies and firewalls
Identity for AI	OAuth 2.0, OIDC, SAML, SSO, SCIM and MFA; Microsoft Entra ID; service principals, managed and workload identities; acting on behalf of a user

Hands-on lab

ServiceNow AI agent: an MCP server that searches knowledge and creates and updates incidents with SSO-aware permissions.

Customer Engagement Lab

Answer the customer security team's questionnaire on data flow, identity and access.

Adds to your Accelerator: MCP server template with OAuth and OIDC

Project 3 delivered: ServiceNow AI Agent via MCP

Month 3

Weeks 9–12

Deploy & deliver: Production & customer engineering

Run it in production and prove the business value.

Week 9**Kubernetes, Terraform and CI/CD**

Month 3: Deploy & deliver

Turn a working prototype into a repeatable, automated production deployment.

Kubernetes	pods, deployments, services, ingress, ConfigMaps, Secrets, namespaces, RBAC, probes, requests and limits, HPA; rolling, blue/green and canary releases
Managed Kubernetes	EKS with AKS and GKE equivalents, private clusters, workload identity and Helm charts
Terraform	providers, resources, data sources, variables, outputs, modules, remote state and locking, import, drift, plan and apply
CI/CD	GitHub Actions with OIDC to AWS (no long-lived keys): test, scan, build, push, deploy, smoke test; environment promotion, approval gates and rollback
GitOps	Argo CD, plus Jenkins and Azure DevOps equivalents

Hands-on lab

Ship the IT-Ops agent to EKS: Terraform for VPC, EKS and IAM, a Helm chart and a GitHub Actions pipeline using OIDC.

Customer Engagement Lab

Write the POC-to-production plan: everything that changes between the demo and go-live.

Adds to your Accelerator: Terraform module, Helm chart and CI/CD pipeline**Project 4 delivered: IT-Ops Multi-Agent Platform****Capstone sprint 0: discovery and architecture****Week 10****AI security, governance and enterprise networking**

Month 3: Deploy & deliver

Pass the customer's security review: attack your own system, then close every gap.

LLM threats	OWASP Top 10 for LLM Applications: direct and indirect prompt injection, sensitive information disclosure, insecure output handling, excessive agency; tool and retrieval poisoning
Defenses	guardrails, input and output filtering, least-privilege tools, human approval, sandboxing, secrets management and audit logs
Red teaming	jailbreaks, data exfiltration and unauthorized tool use against your own agent, then the fixes
Responsible AI and governance	PII masking, data residency, auditability, human-in-the-loop and access governance
Enterprise networking	DNS, TLS, proxies, NAT and firewalls; VPC and VNet design, private subnets, PrivateLink and Private Endpoints, VPN; no-inbound, proxy-only and air-gapped environments

Hands-on lab

Secure Banking AI Assistant: PII masking, RBAC-filtered retrieval, audit trail and approvals, plus a red-team report of attacks and fixes.

Customer Engagement Lab

Present your security controls to a simulated CISO and architecture review board.

Adds to your Accelerator: Guardrails, red-team test suite and security-review template**Project 5 delivered: Secure Banking AI Assistant****Capstone sprint 1: working POC**

Week 11**Observability, cost and troubleshooting**

Month 3: Deploy & deliver

Know what your AI system is doing, what it costs and how to fix it at 2 a.m.

Observability	logs, metrics, traces, alerts and dashboards with CloudWatch, OpenTelemetry, Prometheus and Grafana
AI observability	token usage, model, retrieval and tool latency, agent trajectories, tool failures, evaluation scores and cost per request, with LangSmith- and Langfuse-style tracing
Cost and performance	model routing, prompt and semantic caching, batching, streaming and async processing; the cut-cost-by-60% challenge; the production-readiness gate
Troubleshooting	IAM AccessDenied, DNS failures, timeouts, 403, 429 and 5xx; CrashLoopBackOff, OOMKilled, ImagePullBackOff; hallucinations, bad retrieval, wrong tool, agent loops, context overflow
Incident response	triage, impact, logs, metrics and traces, root cause, mitigation, customer communication and RCA

Hands-on lab

Break-fix gauntlet: ten production failures injected into your deployed stack. Find them, fix them and write the RCA.

Customer Engagement Lab

Write the customer update and the RCA for a P1 outage.

Adds to your Accelerator: Observability dashboards, runbook and RCA template

Capstone sprint 2: deployed, monitored and evaluated

Week 12**Customer engineering and capstone delivery**

Month 3: Deploy & deliver

Deliver like an FDE: architecture, proposal, ROI and an executive demo.

Solution architecture	functional and non-functional requirements, HLD and LLD, sequence and data-flow diagrams, trade-offs
Proposals and SOWs	scope, out of scope, dependencies, risks, timeline, acceptance criteria and success metrics
AI ROI	hours saved, cost avoided, run cost and payback period
Demo engineering	the 5-minute executive demo, the 20-minute technical demo and the 45-minute engineering workshop
Multi-cloud and hybrid	an AI gateway across AWS, Azure and Google Cloud; self-hosted open-weight models for restricted environments
The FDE interview loop	coding, system design, customer-scenario and behavioral rounds; portfolio and resume

Hands-on lab

Capstone sprint 3: UAT with the customer, go-live rehearsal, hypercare drill and executive presentation.

Customer Engagement Lab

Present business value, architecture and ROI to a mock CXO panel.

Adds to your Accelerator: SOW template, ROI calculator and executive demo script

Capstone delivered: the GlobalBank engagement

Enterprise projects

Five projects, each built from a realistic customer brief and each answering the six FDE questions. They build on one another, so by week 10 you have a connected portfolio, not five toy demos.

P1

Weeks 2–4

Customer Integration Service

A customer's CRM and support data sits behind different APIs with different schemas. Pull it, clean it and serve it through one reliable API.

Stack: Python, FastAPI, PostgreSQL, Redis, Docker, Amazon ECS Fargate, RDS, Secrets Manager

Proves: Production Python, API design, data quality and secure cloud deployment.

P2

Weeks 5–6

Enterprise Knowledge Assistant

Employees lose hours hunting through HR and IT policies. Build an assistant that answers with citations, only from documents each user is allowed to see.

Stack: Embeddings, FAISS or pgvector, a reranker, Amazon Bedrock or Azure OpenAI, FastAPI, Ragas

Proves: RAG engineering, permission-aware retrieval and measurable quality.

P3

Week 8

ServiceNow AI Agent via MCP

The service desk wants an agent that searches knowledge and opens and updates incidents, without ever holding more access than the person using it.

Stack: MCP server in Python, ServiceNow developer instance, OAuth and OIDC, Microsoft Teams or Slack

Proves: MCP, enterprise integration, identity and authorization.

P4

Weeks 7 and 9

IT-Ops Multi-Agent Platform

On-call engineers drown in alerts. A supervisor agent investigates with specialist agents, proposes a fix and waits for human approval.

Stack: LangGraph, tool calling, CloudWatch, Kubernetes API, Jira, Amazon EKS, Terraform, Helm, GitHub Actions

Proves: Agent design, safe autonomy and automated production deployment.

P5

Week 10

Secure Banking AI Assistant

A bank wants AI help for transaction investigations under strict PII, audit and approval rules.

Stack: RAG and tools, PII masking, RBAC, guardrails, audit logging, red-team test suite

Proves: AI security, governance and red teaming.

Every project ships like a customer deliverable

A README a customer can follow, an architecture diagram, tests or an evaluation report, a deployment you can demo and a one-page summary that answers the six FDE questions. Hiring managers see how you work, not just what you built.

Capstone: the GlobalBank engagement

The customer says

“We want an AI system that reduces IT support workload by 40%.”

You don't start by coding. You start by asking questions.

Squads of three or four work like a real FDE pod, with trainers playing GlobalBank's stakeholders: the IT head, the security lead, the service-desk manager and the CFO.

Sprint 0 Discover and design

Week 9 Stakeholder interviews, requirements, success criteria and the high-level design.

Sprint 1 Build the POC

Week 10 Knowledge assistant, ServiceNow tools via MCP, the agent workflow, identity and authorization.

Sprint 2 Harden and deploy

Week 11 Terraform and CI/CD to the cloud, observability, evaluation harness, cost model and security review.

Sprint 3 Prove and hand over

Week 12 UAT, go-live rehearsal, a hypercare drill and the executive presentation with ROI.

What you hand over

- Discovery notes
- Requirements document
- HLD and LLD
- Working, deployed system
- Evaluation report
- Security review
- Runbook
- SOW
- ROI model
- Executive deck

Sample production trace

Request ID	FDE-10291
LLM latency	820 ms
Retrieval	180 ms
Tool call	240 ms
Tokens	2,840
Cost	\$0.012
Faithfulness	0.94
Context recall	0.91

How an FDE proves value: an illustrative ROI

Today	100 employees spend 30 minutes a day on a task
With the AI system	the same task takes 10 minutes
Time saved	20 min × 100 people ≈ 33 hours a day, about 8,300 hours a year (250 working days)
Value	at ₹500 an hour, roughly ₹41 lakh a year, before subtracting run cost
Payback	build cost ÷ monthly net savings

The FDE's job is to answer one question with numbers: why should the customer deploy this?

Tools and skills

The FDE PRO technology stack

Languages	Python, SQL, Bash
Backend and data	FastAPI, Pydantic, PostgreSQL, Redis, pgvector; Kafka and Airflow concepts
GenAI and RAG	Amazon Bedrock, Azure OpenAI, Gemini, open-weight models, embeddings, FAISS, Chroma, OpenSearch, rerankers
Agents	LangChain, LangGraph, Model Context Protocol (MCP), Bedrock AgentCore, Microsoft Foundry Agent Service, Google ADK
Evaluation and observability	Ragas, LangSmith, Langfuse, OpenTelemetry, CloudWatch, Prometheus, Grafana
Cloud	AWS (primary), Microsoft Azure, Google Cloud
DevOps	Git, GitHub Actions, Docker, Kubernetes (EKS), Helm, Terraform, Argo CD
Enterprise systems	ServiceNow, Jira, GitHub, Slack, Microsoft Teams, Microsoft Entra ID, Salesforce, SharePoint

By the end of the program, you can

- ✓ Run a discovery call and turn a vague request into requirements and success criteria
- ✓ Map business workflows to architecture and write HLDs and LLDs
- ✓ Write production Python and FastAPI services
- ✓ Build RAG systems with citations, and measure their quality
- ✓ Build tool-calling agents and multi-agent systems with LangGraph
- ✓ Build MCP servers and integrate ServiceNow, Jira, GitHub and Teams
- ✓ Implement SSO, OAuth/OIDC and least-privilege access for AI agents
- ✓ Deploy on AWS with Docker, Kubernetes, Terraform and CI/CD, and find your way around Azure and Google Cloud
- ✓ Secure, red-team, monitor and cost-optimize AI applications
- ✓ Troubleshoot production incidents and write RCAs
- ✓ Write SOWs, calculate AI ROI and demo to executives

Careers and placement

FDE PRO prepares you for roles where companies need engineers who can build AI and make it work inside real organizations.

Roles you can target

- Forward Deployed Engineer (AI)
- Applied AI Engineer
- AI Solutions Engineer / Customer Engineer
- Agentic AI Engineer
- AI Integration Engineer
- GenAI Engineer
- Technical Implementation / Deployment Engineer
- Associate Solutions Architect (AI)

Where FDEs work

- AI product companies and model providers deploying their platforms into large enterprises
- Enterprise software platforms rolling out AI agents inside their customers' systems
- Global consulting and IT services firms building AI delivery practices
- Global Capability Centres (GCCs) in Hyderabad and Bengaluru running internal AI programs
- AI startups building solutions for banking, healthcare, retail and manufacturing

The FDE interview loop, and how we prepare you

Coding	Practical Python: parse, transform, call APIs, handle failure.
System design	Design an AI system with RAG, agents, identity and deployment.
Customer scenario	Run discovery with a 'customer' who is vague, busy or sceptical.
Behavioral	Ownership, ambiguity and delivering under pressure.

Placement support

- ✓ ATS-friendly resume and LinkedIn profile positioned for FDE and applied AI roles
- ✓ Portfolio review: your FDE Accelerator repository, five projects and the capstone
- ✓ Mock interviews covering coding, system design, customer scenarios and behavioral rounds
- ✓ Placement assistance from the team behind 5,500+ placed alumni

Your portfolio on day one of job hunting

A public FDE Accelerator repository, five deployed projects, a capstone with discovery notes, design documents, an evaluation report, a security review, an SOW and an executive deck. That is the evidence FDE hiring managers ask for.

Fee, batches and enrollment

Course fee

₹30,000

Complete 3-month program

What's included

- ✓ 120+ hours of live, instructor-led training
- ✓ 60+ hands-on labs, one in every session
- ✓ 5 enterprise projects and the GlobalBank capstone
- ✓ Your own FDE Accelerator repository
- ✓ Cloud Soft lab guides and handbooks
- ✓ Resume, portfolio and interview preparation
- ✓ Placement assistance

Duration	3 months (12 weeks), 120+ hours of live training
Sessions	5 sessions a week, 2 hours each
Format	Classroom at Ameerpet, beside Ameerpet Metro Station, or live online

Who should join

- Freshers (B.Tech, MCA, M.Sc., BCA) with programming basics who want a premium AI career track
- Software developers in Python, Java, .NET or Node.js moving into applied AI
- DevOps, cloud and SRE engineers adding GenAI and agent skills
- Support, implementation and integration engineers ready to step up to customer engineering
- Data engineers and analysts who want to ship AI solutions
- Pre-sales and solutions engineers who want hands-on build depth
- NRIs and overseas learners through live online batches

Readiness checklist

- ✓ You can write basic programs in any language: variables, loops and functions
- ✓ You are comfortable on the command line, or willing to practise it in week 1
- ✓ A laptop that runs Docker smoothly (16 GB RAM recommended, 8 GB works for most labs)
- ✓ Free accounts: GitHub, AWS free tier, an Azure or Google Cloud trial and a ServiceNow developer instance
- ✓ 8–10 hours a week of practice outside class, because this is a fast-track program

How to join

1

Book a free demo

Call or WhatsApp +91 96660 19191 for batch dates, timings and a demo slot.

2

Choose your format

Classroom at Ameerpet or live online, and the batch timing that suits you.

3

Start week 1

Set up your FDE workstation and make your first LLM call in the first week.

Frequently asked questions

Is Forward Deployed Engineering a good career in India?

Yes, for engineers who enjoy both building and solving customer problems. As companies move AI from pilots into production, they need engineers who can integrate it with real systems, secure it and prove its value. AI companies, enterprise software platforms and global consulting firms now post FDE roles in India, including in Hyderabad.

What is the fee for the FDE course at Cloud Soft Solutions?

The FDE PRO course fee is ₹30,000 for the complete 3-month program. It covers live training, hands-on labs, five enterprise projects, the GlobalBank capstone, lab guides and placement support.

How long is the course and how are classes scheduled?

Three months (12 weeks), with five 2-hour sessions a week, for 120+ hours of live training. Four sessions each week are engineering; the fifth is a Customer Engagement Lab. Call +91 96660 19191 for the next batch dates and timings.

Can freshers join the FDE course?

Yes, if you can already write basic programs in any language. Python is taught from fundamentals to production standard in the first weeks. Because this is a fast track, plan for 8–10 hours of practice a week outside class.

Do I need prior AI or machine learning experience?

No. LLM concepts, prompt and context engineering, RAG, agents and evaluation are all taught in the program. Comfort with programming and a willingness to work with real systems matter more.

How is an FDE different from a Solutions Engineer or Solutions Architect?

A solutions engineer mostly demos the product and helps win the deal; a solutions architect designs the target architecture. An FDE does both, then builds, integrates, deploys and supports the solution in the customer's environment until it delivers results in production.

How is FDE PRO different from a regular GenAI or Agentic AI course?

Most AI courses stop at a working demo. FDE PRO is built around the customer engagement lifecycle: discovery, architecture, integration, identity, deployment, security, evaluation, ROI and executive communication. You learn to take AI all the way to production.

Which cloud platforms are covered?

AWS is the primary hands-on platform, including IAM, VPC, ECS, EKS, Lambda and Amazon Bedrock. You also work with Microsoft Azure (Azure OpenAI, Microsoft Foundry, Key Vault, Entra ID) and Google Cloud (Gemini, Cloud Run, Agent Development Kit).

What projects will I build?

Five enterprise projects: a Customer Integration Service, an Enterprise Knowledge Assistant, a ServiceNow AI Agent via MCP, an IT-Ops Multi-Agent Platform and a Secure Banking AI Assistant. Then the GlobalBank capstone, a simulated end-to-end customer engagement.

Is the course available online?

Yes. Attend in the classroom at Ameerpet, beside Ameerpet Metro Station, or join live online from anywhere, including from outside India.

Do you provide placement support?

Yes. You get an ATS-friendly resume and LinkedIn profile, a portfolio review, mock interviews covering coding, system design and customer scenarios, and placement assistance from the team that has placed 5,500+ alumni.

Will I have to pay for cloud or AI API usage?

Labs are designed around free tiers, local tools and low-cost models, and you set up budgets and alerts in week 4. Some exercises on your own cloud or API accounts may incur small usage charges; the labs show you how to keep them low.

What laptop do I need?

Any laptop that runs Docker smoothly. 16 GB RAM is recommended, especially for running local models; 8 GB works for most labs.

How is FDE PRO different from APEX and NEXUS?

APEX is our broad AI, ML, cloud and cyber security engineering program, and NEXUS focuses on cloud, DevOps, SRE and AIOps. FDE PRO is a focused 3-month fast track for one job: taking AI solutions from a customer's problem to production.

Why learn FDE with Cloud Soft Solutions

Led by a practitioner

The program is led by Sreekanth Bathalapalli (Cloud), Managing Partner & Lead Trainer, with 20 years in IT across cloud and DevOps architecture and AI/GenAI engineering.

5,500+ alumni placed

Our training and placement record across programs, now applied to the FDE role.

A production content library

The same lab guides and handbooks behind our APEX and NEXUS programs: Amazon Bedrock, LangChain, FAISS, RAG pipelines, EKS, Helm, Docker, Jenkins and more.

70% hands-on

Every session ends with something running, not just slides.

Classroom or online

Learn beside Ameerpet Metro Station in Hyderabad, or join live online from anywhere.

Book your free demo

Call or WhatsApp

+91 96660 19191

Call

+91 99496 16388

Email and web

info@cloudsoftsol.com

www.cloudsoftsol.com

Visit

513, 5th Floor, Aditya Enclave, Nilagiri Block
Beside Ameerpet Metro Station, Ameerpet
Hyderabad, Telangana 500016



Scan for the full program page

FDE PRO. Build AI. Deploy AI. Solve enterprise problems.