DATAMARCOS
IT CONSULTING FIRM

DATAMARCOS
GLOBAL TECHNOLOGY LEARNING
AND WORKFORCE UPSKILLING PARTNER

LEARN
BUILD
DEPLOY
GROW

5-MONTH EXECUTIVE PROGRAM | JOB-ORIENTED

Forward Deployed AI Engineer.

Build it. Deploy it. Own it inside the customer's world.

The most in-demand engineering role of the AI era. Learn to scope ambiguous problems, ship secure, production-grade AI systems into real enterprise environments, and own the outcome, guided by a practising architect and four client-style engagements.



5 Months | **22** Weeks | **4** Client engagements | **170+** Hours live and studio

Request a seat →

- Live Classes
- Hands-on Projects
- 1:1 Mentorship

TECH STACK YOU WILL MASTER — FROM IDEAS TO PRODUCTION

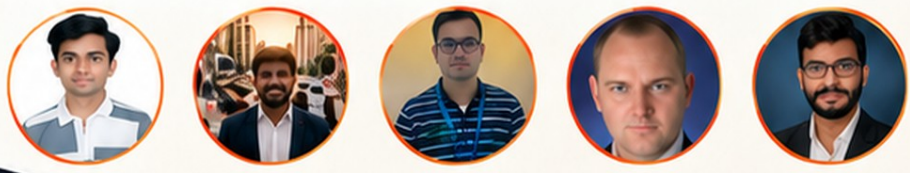
OpenAI · LangChain · LangGraph · CrewAI
n8n · Databricks · Anthropic · Google Cloud
aws · Azure · Snowflake · Pinecone



DEEPESH MISHRA
Founder, DataMarcos | Lead Mentor

"We don't just teach tools. We create engineers who can deliver AI solutions in the real world.

LEARNERS FROM DATAMARCOS PROGRAMS



REAL ENTERPRISE USE CASES | PRODUCTION-GRADE DEPLOYMENTS | MENTORSHIP BY INDUSTRY EXPERTS | JOB-READY SKILL STACK

THE ROLE

The engineer who owns the outcome

A Forward Deployed Engineer (FDE) embeds with customers, turns ambiguous business problems into working software, deploys it inside the customer's real environment, and owns whether it works. The model began at Palantir; AI labs, data platforms and enterprises now hire for it at scale.

Engineer
Writes production code every day: integrations, RAG, agents, pipelines, infrastructure. Not slideware.

Consultant
Runs discovery, frames the problem, writes the SOW, presents ROI and manages stakeholders.

Owner
Deploys into messy, secured, legacy environments and stays accountable until it creates value.

	SOFTWARE ENGINEER	SOLUTIONS ARCHITECT	FORWARD DEPLOYED ENGINEER
Customer proximity	Indirect, via product team	High, mostly pre-sales	Embedded with the customer
Writes production code	Yes	Rarely	Yes, in the customer's environment
Owns deployment and results	Shared	No	Yes, end to end
Handles ambiguity	Defined tickets	Medium	Core of the job
Security and compliance	Team standards	Advisory	Hands-on, per customer

A DAY IN THE LIFE

From stakeholder call to production fix

Morning
Stakeholder sync, clarify the real problem, agree what success looks like this week.

Midday
Build: connect to the client's database, tune retrieval, wire an agent to Jira or ServiceNow.

Afternoon
Debug inside a locked-down environment: networking, permissions, odd data.

Evening
Write the update, log decisions, harden security, plan tomorrow's demo.

WHERE FDES WORK

Hiring across AI labs, platforms and enterprises

Also hiring FDE-style roles: McKinsey QuantumBlack | BCG X | EY | Sarvam AI | Cadence | ServiceNow. Plus Global Capability Centres and enterprise AI teams across India. Logos are for identification; no affiliation implied.

WHY NOW

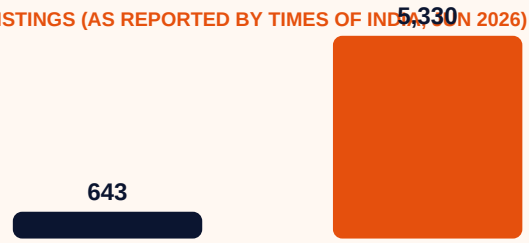
The fastest-growing role in AI engineering

Public data shows FDE hiring accelerating sharply while most AI pilots still fail to reach production, which is exactly the gap this role closes.

643 to 5,330	+729%	Referral-led	Rs 18-90L+
Indeed FDE listings, Apr 2025 to Apr 2026	Year-over-year growth in listings	Most roles filled via targeted outreach	Indicative India pay bands, by tier

Sources: Indeed data as reported by Times of India (Jun 2026); OwnYourCareer, Scaler and fde.directory industry compilations (2026). Geography and methodology differ by source; treat as directional.

INDEED FDE LISTINGS (AS REPORTED BY TIMES OF INDIA, JUN 2026)



+729%

year-over-year growth in listings

Demand is outrunning the supply of engineers who can deploy AI inside real enterprises.

INDICATIVE COMPENSATION

What the market reports (2026)

Estimates from public compilations. Pay varies widely by employer tier and is never guaranteed by any program.

EXPERIENCE	TYPICAL INDIA RANGE (REPORTED)	NOTES
0-2 years	Rs 18-28 LPA	Rare at true entry level; most FDE hires have 2+ years of engineering
3-6 years	Rs 28-55 LPA	The sweet spot for most FDE hiring
Senior / global-remote	Rs 55-90 LPA and above	Global labs and platform companies; larger equity components
IT-services AI deployment roles	Rs 5-30 LPA	Structured roles with lower pay bands than product companies

WHAT HIRING MANAGERS SCREEN FOR

The five signals that get FDE offers

1. Shipped, deployed work
A live system with security, tests and monitoring beats a certificate.

2. Integration depth
Databases, legacy APIs, identity, networking: the messy real world.

3. Evaluation discipline
You measure quality, cost and latency, and can prove it works.

4. Customer communication
Discovery, scoping, pushing back, executive-level updates.

5. Security and judgement
You know where agents fail and how to contain the blast radius.

6. Ownership mindset
You take the problem from vague ask to adopted solution.

WHO THIS IS FOR

Six career journeys, one destination

Application-based admission keeps every cohort focused, practical and peer-strong.

Backend and full-stack engineer
3-10 years of experience

→ **AI FDE at a product company**

Add LLM, RAG, agent and deployment depth to strong software foundations.

IT-services engineer or consultant
Delivery or integration background

→ **Product-company deployment role**

Move from project delivery to owning AI outcomes at the customer.

Data or ML engineer
Pipelines, warehouses, models

→ **Platform FDE (Databricks, Snowflake, cloud AI)**

Turn data skills into agentic, governed enterprise solutions.

Solutions architect or pre-sales
Strong with customers

→ **Hands-on FDE with code ownership**

Back your customer skills with production engineering credibility.

QA, SDET or DevOps engineer
Automation and reliability

→ **AI deployment and reliability engineer**

Bring testing and ops rigor to LLM systems: evals, gates, observability.

Tech lead or architect
Leads teams and designs

→ **AI delivery lead or founding FDE**

Lead enterprise AI engagements with architecture, security and ROI.

You will thrive if

- You code comfortably in Python and have 2+ years of work experience (or strong shipped projects)
- You like customer problems as much as technical ones
- You can commit about 10 to 14 hours a week
- You want to build, deploy and defend real systems

This is probably not for you if

- You want a passive, watch-only video course
- You have never written code
- You expect a guaranteed job or salary outcome
- You prefer theory over shipping

Admission process

1. Apply online • 2. Technical and fit diagnostic • 3. Short interview with the admissions mentor • 4. Offer and cohort allocation. Bridge pre-work (Python, Linux, Git, SQL) is provided before Day 1.

PROGRAM AT A GLANCE

A 5-month executive program for working engineers

Live, mentor-led, and built around the real FDE job: discovery, build, deploy, secure, hand over.

5 Months	170+ Hours	4 Engagements	Small Cohorts	Certificate
22 weeks, weekend-friendly	Live sessions, labs, studio	3 studio + 1 capstone	Application-based entry	DataMarcos Executive FDE

HOW THE PROGRAM RUNS

Format and rhythm

COMPONENT	WHAT HAPPENS	RHYTHM
Live architecture sessions	Mentor-led teaching: the why, trade-offs and live building	Two sessions each week (weekend)
Delivery labs	Hands-on labs and code reviews on your own work	Weekly, weekday evening
Engagement Studio	Client-style engagements with stakeholder role-play and panel reviews	Four engagements over the program
FDE clinics	Discovery calls, executive updates, incident drills, interview loops	Weekly
1:1 mentoring	Architecture reviews, career strategy, feedback	Scheduled through the program
Capstone and Demo Day	Defend your capstone to an industry panel	Final three weeks

OUTCOMES

By the end, you can...

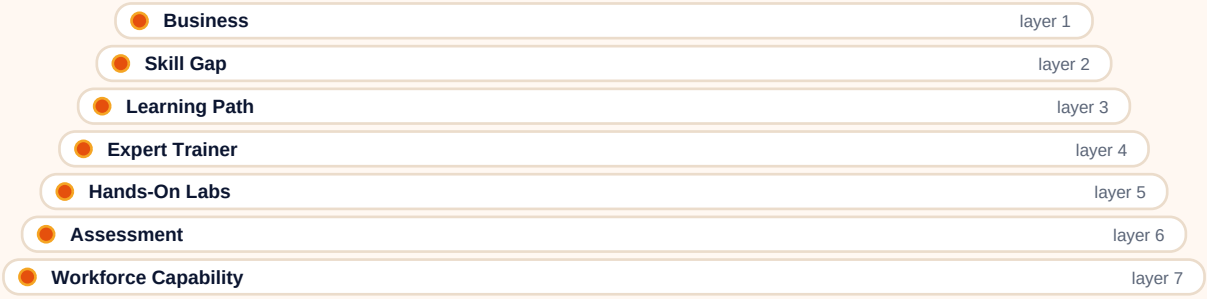
- ✓ Run discovery and turn a vague ask into a scoped, costed engagement
- ✓ Build secure RAG, graph and multimodal systems on messy enterprise data
- ✓ Design single and multi-agent systems with LangGraph, MCP and human approval
- ✓ Integrate with legacy databases, SOAP services, Jira, Slack and identity systems
- ✓ Deploy on AWS, Azure or GCP with Docker, Terraform and CI/CD
- ✓ Defend agents against prompt injection, data leakage and excessive agency
- ✓ Operate AI in production: gateways, tracing, evals, cost and latency SLOs
- ✓ Pass FDE interview loops: system design, live debugging and customer scenarios

OUR APPROACH

How DataMarcos builds capability

The same seven-layer capability system we use for enterprise workforce upskilling, applied to your FDE journey.

OUR CAPABILITY SYSTEM: SEVEN LAYERS THAT TURN SKILLS INTO WORKFORCE CAPABILITY



The FDE Toolkit you keep

Discovery kit and questionnaires • Architecture decision records • SOW and estimate templates • ROI model • Threat model and red-team checklist • Governance and risk register • UAT runbook • Incident runbook and postmortem template • QBR and executive update decks

TEN QUESTIONS TO ASK ANY FDE PROGRAM

And how this one answers

QUESTION	OUR ANSWER
Do you ship into a realistic client environment, not just notebooks?	Yes. Four client-style engagements with stakeholders, constraints and panel reviews.
Do you teach the consulting half: discovery, SOW, ROI, UAT?	Yes. Built in from Week 1 with a reusable FDE toolkit.
Is security and compliance hands-on?	Yes. Threat modelling, red-teaming, guardrails, DPDP, GDPR, HIPAA and EU AI Act essentials.
Do you cover data platforms where many FDE roles sit?	Yes. Databricks, Snowflake and cloud agent platforms in Phase 3.
Do you teach evals and observability, not just building?	Yes. Eval-first from Week 5; gateways, tracing and CI gates in Phase 4.
Who teaches?	A practising architect who has delivered FDE programs for enterprise teams.
Is career support specific to FDE hiring?	Yes. FDE interview loops, referrals strategy and mock rounds. No placement guarantees.

THE ROADMAP

22 weeks from engineer to Forward Deployed Engineer

Five phases, four client-style engagements, and a career track running from Week 1.



Engagement Studio: three client-style engagements, then your capstone engagement, each reviewed by a panel

Career track runs every week: portfolio | resume | LinkedIn | interview loops | referrals

Phase 1: Foundation

Weeks 1-4. FDE craft, APIs, cloud and networking, containers and CI/CD.

Phase 2: GenAI Systems

Weeks 5-9. LLM engineering, RAG, knowledge graphs, multimodal, cost.

Phase 3: Agents and Integration

Weeks 10-14. LangGraph, multi-agent, MCP, legacy systems, data platforms.

Phase 4: Production and Security

Weeks 15-19. Security, compliance, LLMOps, reliability, UAT and ROI.

Phase 5: Capstone and Launch

Weeks 20-22. Full client engagement, Demo Day, interview loops.

Career track

Every week. Portfolio, resume, LinkedIn, interview prep and referrals.

YOUR FDE SKILL STACK, BUILT LAYER BY LAYER



DETAILED CURRICULUM

Week-by-week syllabus

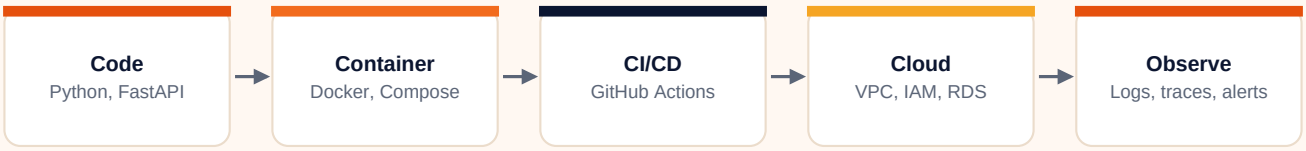
Every week combines live architecture sessions, a hands-on lab, a shipped deliverable and the tools used.

PHASE 1 | Weeks 1-4

Foundation: think like an FDE, ship like an engineer

Fluency in the FDE craft and in deployable software, cloud and DevOps.

EVERYTHING YOU BUILD SHIPS THROUGH THIS PATH



WEEK 1 The FDE Role and Engagement Lifecycle

- What FDEs do: embed, scope, build, deploy, own the outcome
- Discovery craft: stakeholder maps, problem framing, MECE and BLUF
- Artifact toolkit: discovery notes, ADRs, SOW outline, ROI model
- Fluency check: Python, Linux, Git and SQL bridge track

HANDS-ON LAB

Run a mock discovery call and draft a one-page problem brief.

YOU SHIP

Discovery pack for Engagement 1

TOOLS Notion or Docs, Python, Git, SQL

WEEK 2 Backend and API Engineering for Integration

- Async Python, FastAPI, Pydantic, REST and GraphQL, versioning
- Authentication basics, error contracts, pagination, rate limits
- Testing with pytest, contract tests, API design reviews
- Reading and extending systems you did not build

HANDS-ON LAB

Build and test a secured integration API.

YOU SHIP

Tested FastAPI service with OpenAPI docs

TOOLS FastAPI, Pydantic, GraphQL, pytest, Postman

WEEK 3 Cloud, Networking and Security Foundations

- AWS core: VPC, subnets, IAM, KMS, RDS, S3; Azure and GCP mapping
- Private networking, firewalled on-prem access, VPN and bastions
- Secrets management and least privilege
- Linux operations: processes, permissions, shell, logs

HANDS-ON LAB

Stand up a private VPC with RDS and IAM roles.

YOU SHIP

Reference cloud environment as Terraform code

TOOLS AWS, Terraform, Linux, PostgreSQL

WEEK 4 Containers, CI/CD and Deployment

- Docker and Compose; hardened images
- GitHub Actions: tests, scans, build, deploy, rollback
- ECS Fargate, Azure Container Apps and Cloud Run compared
- Observability basics: logs, metrics, traces

HANDS-ON LAB

Pipeline to production on a managed container service.

YOU SHIP

Live service with CI/CD and monitoring

TOOLS Docker, GitHub Actions, ECS Fargate, Grafana

CHECKPOINT **Phase 1 review**

Demo your deployed service and your discovery pack to a mentor panel.

AFTER PHASE 1, YOU CAN

Run a discovery conversation, build and test a secured API, and deploy it with CI/CD and monitoring on a private cloud network.

THE FDE ENGAGEMENT LIFECYCLE YOU WILL PRACTISE



Own the outcome: adoption, ROI, renewal and expansion

PHASE 2 | Weeks 5-9

GenAI Systems: build AI that survives real data

LLM engineering, retrieval, graphs, multimodal documents and cost control, all measured.

THE PRODUCTION RAG PIPELINE YOU WILL BUILD AND MEASURE



Evaluation loop: recall | faithfulness | relevance | cost | latency, gated in CI

WEEK 5 LLM Engineering

- Model landscape (OpenAI, Anthropic, Gemini, open models); selection and cost
- Prompting, structured outputs, tool and function calling
- Context windows, caching, latency and cost budgets
- Eval-first mindset: golden sets and regression tests

HANDS-ON LAB

Build a multi-provider gateway with structured outputs.

YOU SHIP

LLM service with an evaluation harness

TOOLS OpenAI, Claude, Gemini, Ollama, Pydantic

WEEK 6 Retrieval and Core RAG

- Embeddings, chunking strategies and metadata design
- Vector stores: Qdrant, Pinecone, pgvector, Elasticsearch
- Hybrid search, reciprocal rank fusion and reranking
- Ingesting messy enterprise documents

HANDS-ON LAB

Hybrid RAG over a messy document corpus.

YOU SHIP

Retrieval pipeline with measured recall

TOOLS Qdrant, Pinecone, pgvector, LangChain

WEEK 7 Advanced RAG, Knowledge Graphs and Evaluation

- Query rewriting, multi-hop, parent-document and agentic retrieval
- Knowledge graphs with Neo4j and Cypher; GraphRAG patterns
- RAGAS and DeepEval: faithfulness, relevance, regression gates
- Failure analysis and error taxonomies

HANDS-ON LAB

Add graph retrieval and run an evaluation study.

YOU SHIP

RAG system with an evaluation report

TOOLS Neo4j, RAGAS, DeepEval, LangChain

WEEK 8 Multimodal and Document Intelligence

- Vision language models; ColPali and OCR-free parsing
- PDFs, tables, charts, scanned forms and handwriting
- Secure text-to-SQL with read-only access and guardrails
- Speech and voice pipelines

HANDS-ON LAB

Parse invoices and reports into structured data.

YOU SHIP

Document-intelligence service

TOOLS VLMs, ColPali, Whisper, PostgreSQL

WEEK 9 Fine-Tuning, Small Models and Cost

- Fine-tune vs prompt vs RAG; LoRA and PEFT overview
- Small and open models; VPC and on-prem deployment
- Cost engineering: caching, batching, model routing

HANDS-ON LAB

Route requests across models by cost and quality.

YOU SHIP

Cost-quality report and routing policy

TOOLS Hugging Face, Ollama, vLLM concepts

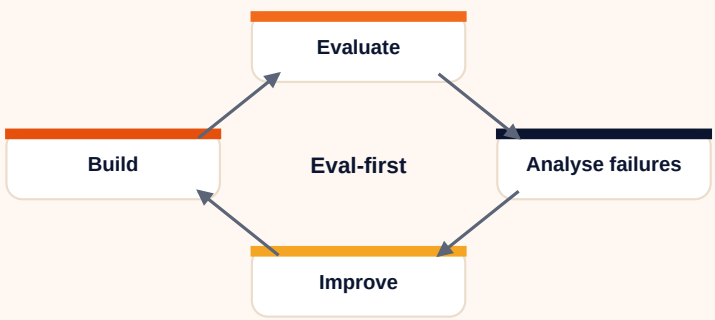
CHECKPOINT **Engagement 1: Project Sentinel**

Deliver a secure knowledge copilot for a BFSI client: hybrid RAG, secure text-to-SQL, RBAC and guardrails, with a discovery pack and a stakeholder demo.

AFTER PHASE 2, YOU CAN

Build, evaluate and cost-optimise RAG, graph and multimodal systems that hold up on messy enterprise data.

EVAL-FIRST ENGINEERING: THE LOOP BEHIND EVERY SYSTEM YOU SHIP



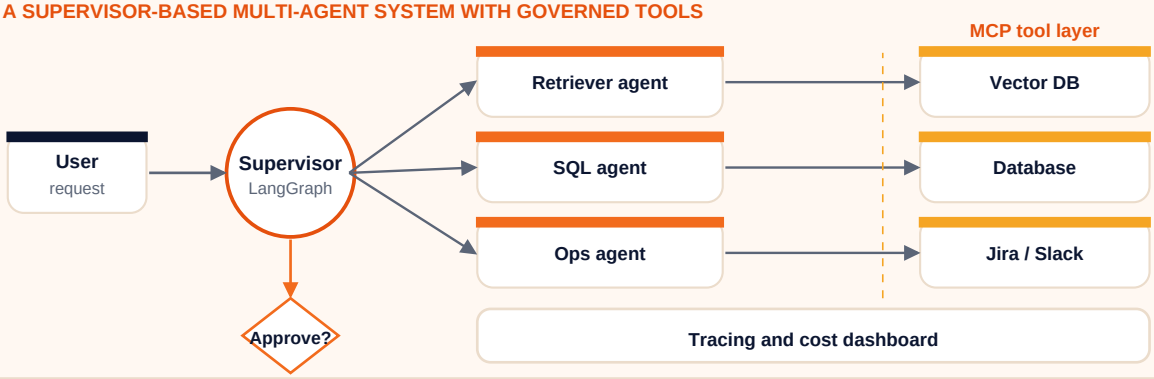
- **Quality** — faithfulness, relevance, task success
- **Cost** — tokens, model routing, caching
- **Latency** — p50 and p95 budgets
- **Safety** — injection, leakage, refusals

PHASE 3 | Weeks 10-14

Agents and Enterprise Integration

Agentic systems that connect safely to the systems enterprises actually run.

A SUPERVISOR-BASED MULTI-AGENT SYSTEM WITH GOVERNED TOOLS



WEEK 10 Agent Design and LangGraph

- Patterns: ReAct, plan-and-execute, reflection, routers
- LangGraph: state, nodes, edges, memory, checkpointing
- Tool design: idempotency, permissions, failure handling

HANDS-ON LAB

Build a stateful, tool-using agent.

YOU SHIP

Agent with evaluation traces

TOOLS LangGraph, LangChain, Langfuse

WEEK 11 Multi-Agent Orchestration and Human-in-the-Loop

- Supervisor, router, parallel and hierarchical teams; sub-graphs
- Human approval gates and escalation design
- Framework trade-offs: LangGraph, CrewAI, OpenAI Agents SDK, Google ADK

HANDS-ON LAB

Supervisor-based multi-agent workflow with approvals.

YOU SHIP

Multi-agent system with cost and trace dashboard

TOOLS LangGraph, CrewAI, ADK

WEEK 12 MCP and Agent Interoperability

- Build MCP servers; tool registries, scopes and governance
- A2A: agent cards and cross-agent handoffs
- Securing tool access: allow-lists, sandboxes, audit logs

HANDS-ON LAB
MCP server in front of a live Jira instance and database.

YOU SHIP
Governed tool layer

TOOLS MCP SDK, Jira, FastAPI

WEEK 13 Legacy Systems and Enterprise Integration

- SQL Server and Oracle, SOAP and XML, flat files, SFTP
- Slack, Jira, ServiceNow, Salesforce and SAP-style integrations
- Event-driven patterns: Kafka, queues, webhooks, idempotency
- Data contracts and schema drift

HANDS-ON LAB
Integrate an agent with a legacy order system.

YOU SHIP
Integration layer with retries and dead-letter handling

TOOLS Kafka, Slack, Jira, SQL Server, Oracle

WEEK 14 Platform FDE: Databricks, Snowflake and Cloud Agent Platforms

- Databricks: Unity Catalog, Mosaic AI and Agent Bricks
- Snowflake: Cortex Search, Analyst and Agents
- Cloud platforms: AWS Bedrock, Azure AI Foundry, Google Gemini Enterprise Agent Platform

HANDS-ON LAB
Build the same assistant on two platforms.

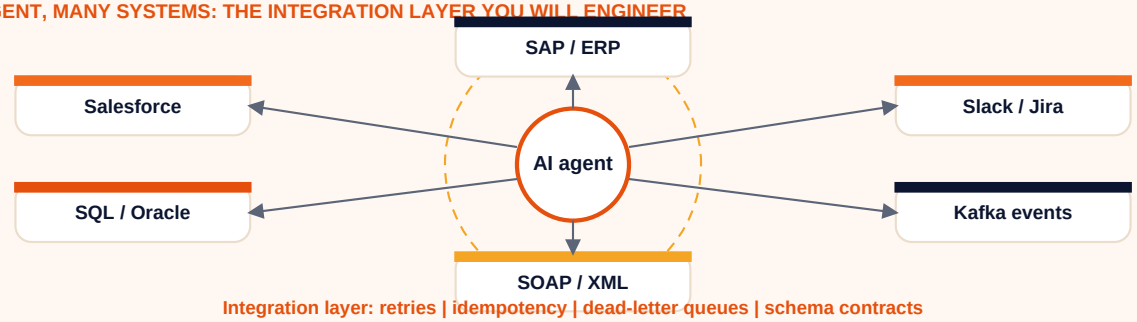
YOU SHIP
Platform comparison and recommendation memo

TOOLS Databricks, Snowflake, Bedrock, Azure, Google Cloud

CHECKPOINT **Engagement 2: Project Atlas**
Deliver a multi-agent operations copilot with MCP tools, Jira and ServiceNow integration, human approval and a cost and trace dashboard.

AFTER PHASE 3, YOU CAN
Design governed multi-agent systems and integrate them with legacy systems, Jira, Slack and data platforms.

ONE AGENT, MANY SYSTEMS: THE INTEGRATION LAYER YOU WILL ENGINEER



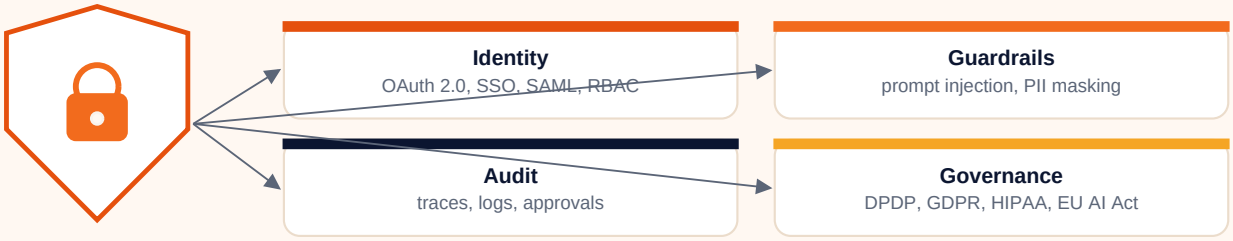
Integration layer: retries | idempotency | dead-letter queues | schema contracts

PHASE 4 | Weeks 15-19

Production, Security and Governance

Operate, secure and hand over AI systems that regulated enterprises will actually approve.

DEFENCE IN DEPTH FOR AI AGENTS



WEEK 15 AI Security Engineering

- OAuth 2.0, SSO, SAML, RBAC and data-level access control
- Prompt injection, jailbreaks, exfiltration; OWASP LLM Top 10
- PII detection and masking (Presidio), NeMo Guardrails, cloud guardrails

HANDS-ON LAB

Red-team your own agent, then patch it.

YOU SHIP

Threat model and hardening report

TOOLS Okta, Auth0, Presidio, NeMo Guardrails

WEEK 16 Governance, Compliance and Regulated Industries

- DPDP Act, GDPR, HIPAA and EU AI Act essentials; SOC 2 and ISO 27001 basics
- Data classification, residency, retention and audit trails
- Model risk management for BFSI and healthcare

HANDS-ON LAB

Compliance review of an AI system.

YOU SHIP

Governance pack: controls, audit log design, risk register

TOOLS Policy templates, audit logging

WEEK 17 LLMOps and Observability

- LLM gateways, rate limits, retries, fallbacks and quotas
- Tracing with LangSmith, Langfuse and OpenTelemetry; token and latency SLOs
- Evals in CI, canary releases and rollback

HANDS-ON LAB
Put your agent behind a gateway with tracing and eval gates.

YOU SHIP
Production-grade operations stack

TOOLS OpenTelemetry, Grafana, Langfuse, GitHub Actions

WEEK 18 Scale, Reliability and Incident Response

- Caching, batching, queues and concurrency; cost at scale
- SLOs, error budgets, runbooks and on-call basics
- Live incident drill and blameless postmortem

HANDS-ON LAB
Load test and fix a failing agent system.

YOU SHIP
Runbook and postmortem

TOOLS Redis, Kafka, Kubernetes, Grafana

WEEK 19 Delivery Craft: UAT, Handoff, Adoption and ROI

- UAT plans, acceptance criteria, cutover, change management
- Handoff to client teams: documentation, training, support model
- Value realisation: ROI, QBRs, expansion and renewal
- Executive communication: BLUF, demos, pushing back

HANDS-ON LAB
Run UAT and a handoff for Engagement 3.

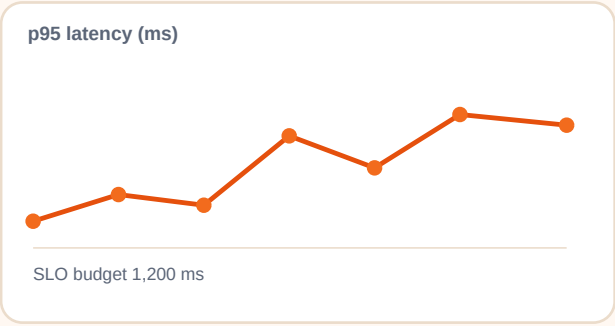
YOU SHIP
UAT runbook and ROI deck

TOOLS Docs, Notion, Grafana dashboards

CHECKPOINT **Engagement 3: Project Helix**
Take Atlas to production for a regulated client: security review, compliance pack, red-team report, UAT and handoff to the client team.

AFTER PHASE 4, YOU CAN
Secure, govern, operate and hand over AI systems that regulated enterprises will approve.

OPERATING AI IN PRODUCTION: SLOs, TRACES, COST AND ALERTS



99.4%	Rs 0.42	94%
Success rate	Cost / request	Eval pass

ALERT
Retrieval recall dropped below threshold: canary rolled back

PHASE 5 | Weeks 20-22

Capstone Engagement and Career Launch

Run a full client engagement end to end, defend it, and convert it into offers.

WEEK 20 Capstone Engagement I: Discover, Design, Build

- Choose a domain track: BFSI, healthcare, or enterprise operations
- Discovery, scoping, SOW, architecture decision records
- Build the core system with RAG, agents and integrations

HANDS-ON LAB

Weekly stakeholder reviews with your mentor panel.

YOU SHIP

Approved architecture and working core system

TOOLS Your stack of choice from the program

WEEK 21 Capstone Engagement II: Secure, Test, Deploy

- Security review, red-team, guardrails and compliance pack
- Evaluation suite, load tests, observability, CI/CD
- UAT, runbooks and handoff documentation

HANDS-ON LAB

Red-team and UAT with an external-style panel.

YOU SHIP

Deployed, documented, secured capstone

TOOLS Docker, Terraform, GitHub Actions, Langfuse

WEEK 22 Demo Day, Interview Loops and Job Launch

- Demo Day: defend your capstone to an industry-style panel
- FDE interview loops: system design with ambiguity, live debugging, customer scenarios
- Resume, LinkedIn and portfolio teardown; referral and outreach playbook
- Offer evaluation and negotiation

HANDS-ON LAB

Two mock interview loops with feedback.

YOU SHIP

Portfolio, resume, LinkedIn and job-search plan

TOOLS Portfolio site, GitHub, LinkedIn

FOUR ENGAGEMENTS, ONE DEMO DAY, ONE OFFER-READY PORTFOLIO



ENGAGEMENT STUDIO

Four client-style engagements, not toy projects

Each engagement comes with a client brief, stakeholders, constraints, security requirements and a panel review. Everything is deployed, documented and portfolio-ready.

E1 Project Sentinel

BFSI client: secure knowledge copilot

Hybrid RAG, secure text-to-SQL, RBAC, PII masking, guardrails. Deliverables: discovery pack, architecture, demo.

E2 Project Atlas

Enterprise operations: multi-agent copilot

LangGraph supervisor, MCP tools, Jira and ServiceNow integration, human approval, cost and trace dashboard.

E3 Project Helix

Regulated deployment: production hardening

Security review, compliance pack, red-team, UAT, runbooks and handoff for a regulated client.

Capstone Engagement

Your domain, your client

BFSI, healthcare or enterprise operations: full lifecycle from discovery to handoff, defended at Demo Day.

THE FDE TOOLKIT

Professional artifacts you will produce

Discovery kit

Stakeholder map, questionnaire, problem brief, success metrics.

Architecture pack

Diagrams, ADRs, integration design, data classification.

Commercial pack

SOW outline, estimate, ROI model, QBR template.

Security and compliance

Threat model, red-team report, governance pack, risk register.

Delivery pack

UAT runbook, handoff guide, training deck, support model.

Operations pack

Runbooks, SLOs, dashboards, incident postmortem.

THE TECHNOLOGY

The stack production FDE teams use

MODELS AND AI PLATFORMS

OpenAI | Claude | Gemini | Hugging Face | Ollama | Databricks | Snowflake | Google Cloud | NVIDIA

AGENTS AND ORCHESTRATION

LangChain | LangGraph | CrewAI | Pydantic | Python

RETRIEVAL, DATA AND EVENTS

Neo4j | PostgreSQL | Elasticsearch | Redis | MongoDB | Kafka

BUILD AND DEPLOY

FastAPI | GraphQL | Docker | Kubernetes | Terraform | GitHub Actions | pytest | Linux | GitHub

SECURITY AND IDENTITY

Okta | Auth0 | OWASP LLM Top 10

OBSERVABILITY

OpenTelemetry | Grafana | MLflow

ENTERPRISE INTEGRATIONS

Slack | Jira | Salesforce | SAP | Stripe

Also covered: AWS (VPC, IAM, KMS, RDS, ECS Fargate, Bedrock) • Azure (Container Apps, AI Foundry) • Qdrant • Pinecone • pgvector • ColPali • RAGAS • DeepEval • LangSmith • Langfuse • Presidio • NeMo Guardrails • MCP SDK • A2A • Oracle • SQL Server • SOAP and XML • ServiceNow • Unity Catalog • Mosaic AI • Agent Bricks • Cortex AI

CAREER LAUNCH

Built around how FDEs actually get hired

FDE roles are often filled through targeted applications and referrals, so we prepare you for the real process.

PHASE	CAREER MILESTONE	OUTPUT
Weeks 1-4	Role targeting: AI-first startups, platform companies, GCCs and consulting	Target list and positioning statement
Weeks 5-9	GitHub portfolio and project write-ups	Architecture-first READMEs for Sentinel and your projects
Weeks 10-14	Resume rewrite for FDE roles; LinkedIn and Naukri profile	ATS-friendly resume with quantified impact
Weeks 15-19	FDE interview loops: system design, debugging, customer scenarios	Practised loops with feedback
Weeks 20-22	Demo Day, mock loops, referral and outreach playbook	Job-search plan and offer-negotiation guide

THE FDE INTERVIEW LOOP

What you will be tested on

1. Technical screen
Python, APIs, data and cloud fundamentals.

2. System design with ambiguity
Scope a vague customer ask and design a secure AI system.

3. Live debugging
Fix an integration or a failing agent under time pressure.

4. Customer scenario
Discovery, pushback, executive update, ROI conversation.

5. Behavioural and ownership
Stories on ambiguity, conflict and delivering under constraints.

6. Portfolio walkthrough
Defend your capstone architecture and trade-offs.

HIRING PATHS

Where graduates aim

AI labs and AI-first startups
Customer-facing deployment and field engineering roles; strong product and customer ownership.

Data and cloud platform companies
Platform FDE roles at data and AI platform vendors and cloud providers.

GCCs, consulting and enterprises
AI deployment teams inside global capability centres, consulting firms and large enterprises.

What you get
- Resume and LinkedIn teardown
- Portfolio and GitHub reviews
- 1:1 mock interview loops with feedback
- Referral and outreach playbook
- Role-targeting guidance

What we do not do
We do not promise placements, salaries or hikes. Hiring outcomes depend on your effort and the market. We promise that you will leave with deployed, secured, documented systems and the ability to explain every design decision like an FDE.

YOUR LEAD MENTOR

Deepesh Mishra

Founder, DataMarcos | Solution Architect | Corporate Trainer | M.Tech, BITS Pilani



Deepesh brings 18+ years of industry experience and has trained more than 25,000 professionals across 200+ organisations in 10+ countries, including teams from Deloitte, EY, Adobe, HSBC, Boeing, Reliance, PwC, Bank of America and Yahoo Singapore.

Over the last year he has delivered GenAI Forward Deployed Engineer programs for Deloitte, Saama Capital and EY, alongside Databricks, data engineering and GenAI cohorts for ZS Associates, Sigmoid and Novartis. He is Databricks-certified and AWS Machine Learning Engineer Associate certified. He teaches the way FDEs work: scope first, then build, secure, deploy and hand over.

25,000+	200+	18+	1,000+
Trained	Orgs	Years	Sessions

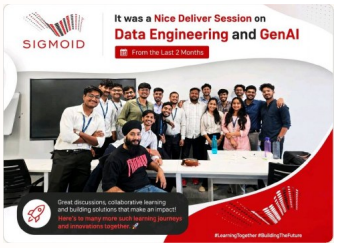
ORGANISATIONS WHOSE TEAMS HAVE BEEN TRAINED

Deloitte. EY Adobe
yahoo! HSBC BOEING
Reliance MICRO FOCUS And 200+ More...

PwC	Bank of America	Novartis	Sigmoid
ZS Associates	Hexaware	Saama Capital	Yahoo Singapore

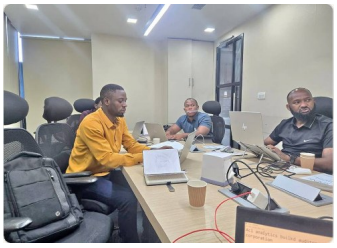
INSIDE THE SESSIONS


ZS Associates cohort



Sigmoid delivery


Live virtual cohort


Hands-on lab

Guest masterclasses

Practitioner sessions with working AI delivery leaders and forward deployed engineers on discovery, security reviews and customer escalations (speakers announced per cohort).

YOUR FDE NETWORK STARTS ON DAY ONE



PROVEN IN THE FIELD

Enterprise delivery, real learner outcomes



Training at Novartis: Analytics, Python, Forecasting, ML



Knowledge sharing and hands-on practice



Cohort wrap-up with the full training group



Learner success story: now a full-time engineer at Sigmoid

WHAT LEARNERS SAY (feedback across Deepesh's data, GenAI and AI programs)

> *"His deep understanding of the subject, combined with a clear and engaging teaching style, made complex concepts easy to grasp, even for those new to the field."*
>
> **Kumar Gaurav, Senior Azure Data Engineer @ PwC**

> *"Many thanks for the great Databricks training at ZS! Let's stay connected."*
>
> **Gourav Thadhani, Data Science and AI @ ZS Associates**

> *"Learning under Deepesh Sir has been a great opportunity. He is truly a great mentor."*
>
> **Ansh Ranjan, Data Engineer Trainee @ Hexaware**

> *"Beyond the technical stack, the biggest lesson I learned was that clear communication is just as critical as writing clean code."*
>
> **Learner, now Software Engineer @ Sigmoid**

ALUMNI AND LIVE SESSIONS

Learners from our programs, in the field

Real learners from DataMarcos programs and snapshots from live cohorts.



Gaurav Taneja
PwC
Senior Azure Data Engineer



Lakshan J
Hexaware Technologies
Software Engineer



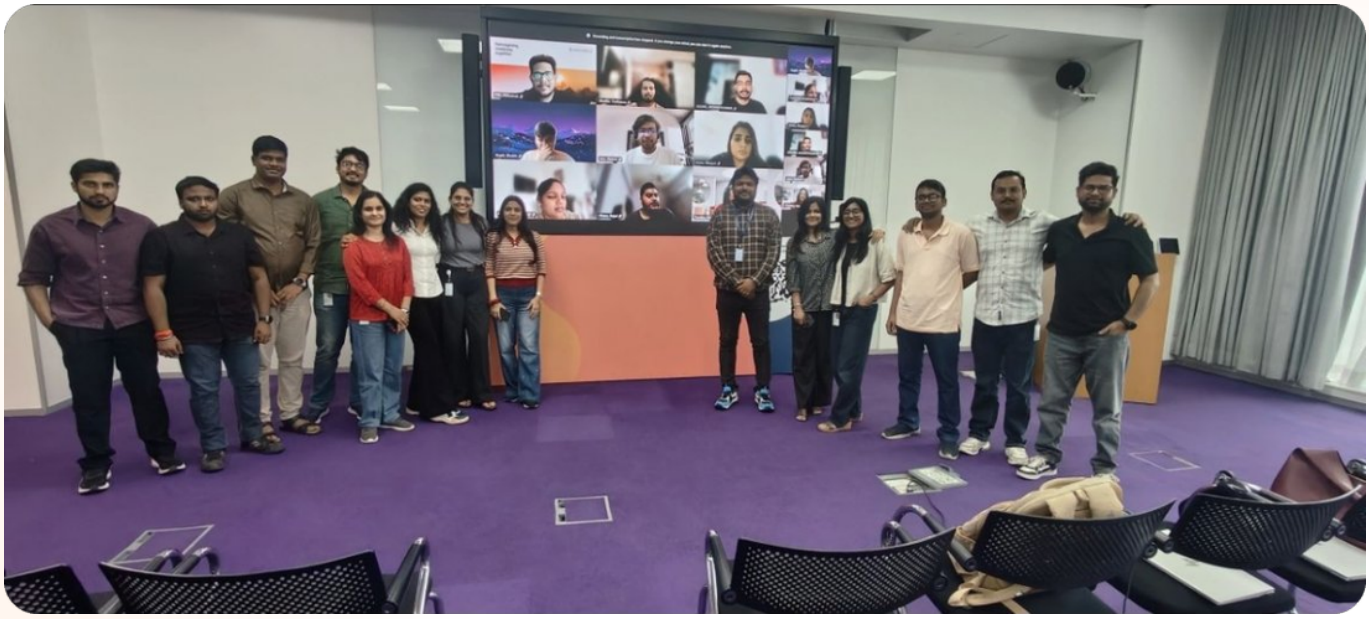
Sharoz Farhan Afridi
Career growth
Technology professional



Richard Leclezio
LinkedIn
Professional learner



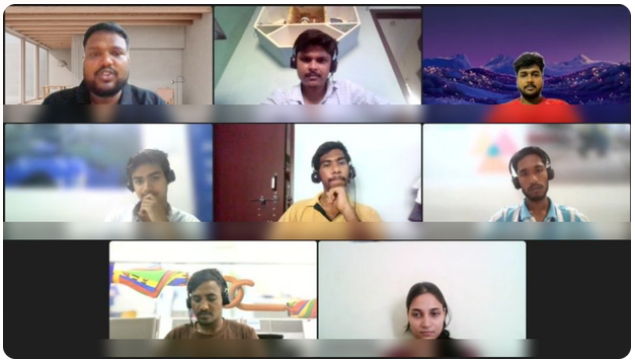
Uphar Tandon
Microsoft
Technology professional



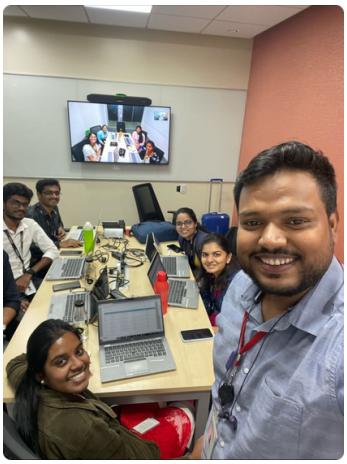
Live cohort with remote participants joining on screen



Instructor-led workshop



Live online cohort



Hands-on lab session

ADMISSIONS

Four steps to your seat

Application-based admission keeps cohorts focused, practical and peer-strong.

1	2	3	4
Apply	**Diagnostic**	**Interview**	**Offer**
Short online application with your background and goals.	Technical and fit assessment: Python, APIs, problem solving.	Conversation with the admissions mentor about your target role.	Cohort allocation, bridge pre-work and onboarding.

DataMarcos Executive Certificate in Forward Deployed AI Engineering

Awarded on completing the program and defending your capstone. It validates hands-on proficiency in AI engineering, enterprise integration, security and delivery, backed by a portfolio of deployed engagements you can show employers.

Investment and cohort dates

Fees and EMI options shared on a counsellor call. Scholarships and upcoming cohort dates: speak to the DataMarcos admissions team at www.datamarcos.com.

FAQ

Your questions, answered

Do I need AI experience?

No. You need solid Python and working engineering fundamentals. A bridge pre-work track covers Python, Linux, Git and SQL before Day 1.

How much time per week?

Plan for about 10 to 14 hours: two live sessions, a weekday lab, and project work.

Will I get a job after this?

We do not guarantee placements. We give you deployed projects, FDE-specific interview preparation, mock loops, and a referral and outreach playbook.

Is security covered in depth?

Yes. A full week on AI security and a week on governance and compliance, plus red-teaming in the studio.

Do you cover Databricks and Snowflake?

Yes. Week 14 is dedicated to platform FDE work on Databricks, Snowflake and cloud agent platforms.

Can my company sponsor me?

Yes. Talk to admissions about corporate sponsorship and team cohorts.

Become the engineer every AI team wants on the ground.

Fees and EMI options shared on a counsellor call. Applications and cohort dates: talk to the DataMarcos admissions team.

www.datamarcos.com | C-908, Sector 62, Noida, Uttar Pradesh, India

DATAMARCOS
IT CONSULTING FIRM