DATAMARCOS

GLOBAL TECHNOLOGY LEARNING
AND WORKFORCE UPSKILLING PARTNER

LEARN
BUILD
DEPLOY
GROW

18-WEEK EXECUTIVE PROGRAM | JOB-ORIENTED

Generative AI & Agentic AI Engineer.

Build it. Ship it. Get hired.

The most in-demand skill set of the AI era. Learn to build LLM apps, RAG systems and multi-agent workflows, then secure, evaluate and deploy them on real enterprise stacks. Guided by a practising architect and three client-style projects plus a capstone.

18	130+	4	1:1
Weeks	Hours	Real projects	Mentorship

Request a seat →

- Live classes
- Hands-on projects
- 1:1 mentorship



DEEPESH MISHRA

Founder, DataMarcos | Lead Mentor

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

TECH STACK YOU WILL MASTER

OpenAI · Claude · Gemini · LangChain · LangGraph · CrewAI · MCP · Databricks · Snowflake · AWS · Azure · Google Cloud · Pinecone · Neo4j · FastAPI · Docker · Terraform

LEARNERS FROM DATAMARCOS PROGRAMS



and 25,000+ professionals trained
across 200+ organisations in 10+ countries

REAL ENTERPRISE USE CASES

PRODUCTION-GRADE DEPLOYMENTS

MENTORSHIP BY INDUSTRY EXPERTS

JOB-READY SKILL STACK

THE OPPORTUNITY

The engineers every AI team is hiring

Companies have moved from AI demos to AI in production. They now need people who can build LLM apps and agents, and the secure, measured systems around them.

Builder
Builds LLM apps, RAG pipelines and agents that run on real, messy data. Not slideware.

Integrator
Connects models to databases, APIs, Jira, Slack and cloud so AI works inside real workflows.

Owner
Evaluates, secures and operates AI in production and stays accountable until it creates value.

ROLES YOU CAN TARGET

GenAI / LLM Application Engineer
Builds copilots, chatbots and generative features on top of foundation models.

RAG and Knowledge Engineer
Designs retrieval, vector search and graph systems that ground answers in company data.

AI Agent Developer
Builds tool-using and multi-agent workflows with LangGraph, MCP and human approval.

AI Solutions / Forward Deployed Engineer
Embeds with customers, scopes problems and ships AI inside their environment.

LLMOps / AI Platform Engineer
Runs gateways, tracing, evals, cost control and CI/CD for AI systems.

AI Security and Governance Engineer
Red-teams agents, applies guardrails and maps systems to compliance needs.

WHAT HIRING MANAGERS SCREEN FOR

The six signals that get AI offers

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

2. Integration depth
Databases, APIs, identity and cloud: the messy real world.

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

4. Agent design judgement
You know when to use RAG, tools, or a multi-agent workflow.

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

6. Ownership and communication
You take a vague ask to an adopted solution and explain it clearly.

WHO THIS IS FOR

Six career journeys, one destination

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

Fresh graduates and career starters

Python basics, strong projects

› **Associate GenAI Engineer**

Turn coursework into deployed, portfolio-ready AI systems.

Backend and full-stack developers

2-10 years of experience

› **LLM and agent engineer**

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

Data, analytics and ML engineers

Pipelines, warehouses, models

› **Applied AI / platform AI engineer**

Turn data skills into agentic, governed enterprise solutions.

IT-services, QA and DevOps engineers

Delivery, automation, reliability

› **AI deployment and reliability engineer**

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

Career switchers with coding basics

Analysts, consultants, support

› **AI solutions engineer**

Back your domain knowledge with hands-on AI engineering credibility.

Tech leads, architects and managers

Lead teams and designs

› **AI delivery lead**

Lead enterprise AI engagements with architecture, security and ROI.

You will thrive if

- ✓ You can code in Python (or have strong shipped projects)
- ✓ You like customer problems as much as technical ones
- ✓ You can commit about 8 to 12 hours a week
- ✓ You want to build, deploy and defend real systems

This is probably not for you if

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

Admission process

1. Apply online
2. Technical and fit diagnostic
3. Short interview with admissions mentor
4. Offer and cohort allocation

Bridge pre-work (Python, Linux, Git, SQL) is provided before Day 1.







PROGRAM AT A GLANCE

An 18-week program for working engineers and job seekers

Live, mentor-led, and built around real work: build, evaluate, secure, deploy and defend.

18 Weeks	130+ Hours	4 Projects	Small Cohorts	Certificate
Two live sessions a week	Live sessions, labs, studio	3 studio + 1 capstone	Application-based entry	DataMarcos Executive

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
Delivery labs	Hands-on labs and code reviews on your own work	Weekly
Project studio	Client-style projects with stakeholder role-play and panel reviews	Three projects + capstone
Career track	Portfolio, resume, LinkedIn, interview loops and referrals	Every week
1:1 mentoring	Architecture reviews, career strategy and feedback	Scheduled through the program
Capstone and Demo Day	Defend your capstone to an industry-style panel	Final two weeks

OUTCOMES

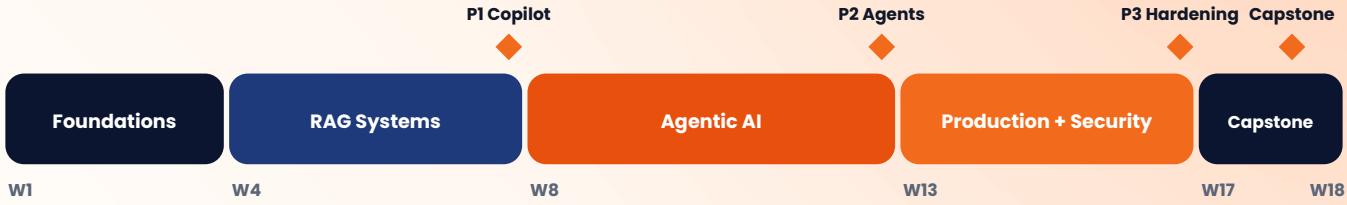
By the end, you can...

- ✓ Build reliable LLM apps with prompts, structured outputs and tools
- ✓ Build secure RAG, graph and multimodal systems on messy data
- ✓ Design single and multi-agent systems with LangGraph, MCP and human approval
- ✓ Integrate with databases, Jira, Slack and enterprise APIs
- ✓ Evaluate quality, cost and latency with measurable evals
- ✓ Defend agents against prompt injection and data leakage
- ✓ Deploy on AWS, Azure or GCP with Docker, Terraform and CI/CD
- ✓ Pass GenAI and agentic AI interview loops

THE ROADMAP

18 weeks from engineer to GenAI and agentic AI builder

Five phases, three projects, one capstone, and a career track running from Week 1.



Project studio: three client-style projects, then your capstone, each reviewed by a panel

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

Phase 1: Foundations

Weeks 1-3. LLM fundamentals, prompting, structured outputs, GenAI APIs.

Phase 2: RAG Systems

Weeks 4-7. Embeddings, RAG, knowledge graphs, multimodal, evals.

Phase 3: Agentic AI

Weeks 8-12. LangGraph, MCP, multi-agent, human approval, integration.

Phase 4: Production

Weeks 13-16. Security, LLMOps, cost, cloud deployment, governance.

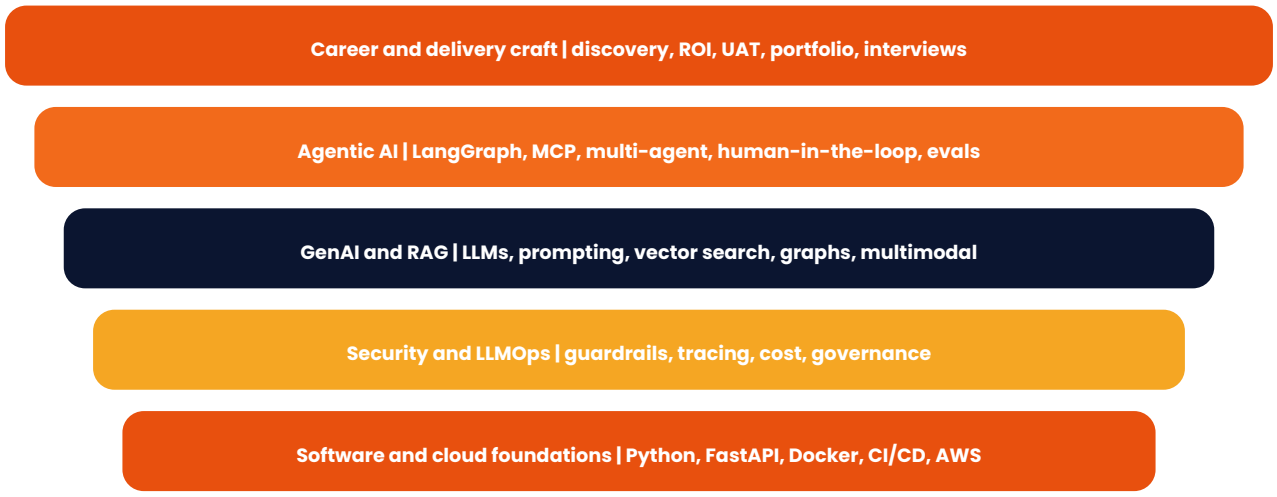
Phase 5: Capstone

Weeks 17-18. Full project, Demo Day, interview loops.

Career track

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

YOUR SKILL STACK, BUILT LAYER BY LAYER

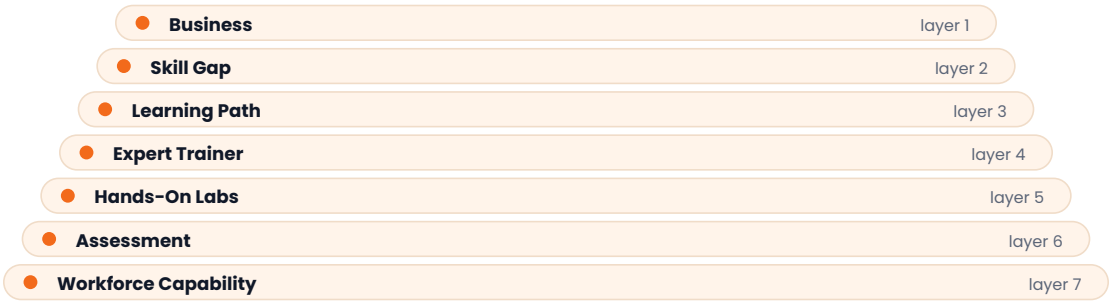


OUR APPROACH

How DataMarcos builds capability

The same capability system we use for enterprise workforce upskilling, applied to your AI engineering journey.

SEVEN LAYERS THAT TURN SKILLS INTO WORKFORCE CAPABILITY



The AI toolkit you keep

Prompt and eval library • RAG and agent templates • Architecture decision records • Threat model and red-team checklist • Governance and risk register • UAT and incident runbooks • ROI model and executive update decks

EIGHT QUESTIONS TO ASK ANY GENAI PROGRAM

And how this one answers

QUESTION	OUR ANSWER
Do you build beyond notebooks into deployed systems?	Yes. Three projects and a capstone, deployed with CI/CD and monitoring.
Do you teach agents properly, not just chatbots?	Yes. LangGraph, MCP, multi-agent orchestration and human approval in Weeks 8-12.
Is security hands-on?	Yes. Red-team your own agent, OWASP LLM Top 10, guardrails, DPDP, GDPR, HIPAA and EU AI Act essentials.
Do you teach evals and observability?	Yes. Eval-first from Week 2; gateways, tracing and CI gates in Week 14.
Do you cover enterprise data platforms?	Yes. Databricks, Snowflake and cloud agent platforms in Week 12.
Who teaches?	A practising architect who has delivered GenAI and FDE programs for enterprise teams.
Is career support specific to AI hiring?	Yes. Resume, LinkedIn and Naukri, mock loops and a referral playbook. No placement guarantees.
Do I keep reusable material?	Yes. The full toolkit of templates, runbooks and your deployed project portfolio.

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 1-3

Foundations: understand LLMs, build like an engineer

Fluency in LLMs, prompting, structured outputs and deployable GenAI APIs.

EVERYTHING YOU BUILD SHIPS THROUGH THIS PATH

Prompt patterns, few-shot → **Structure** JSON, Pydantic → **API** FastAPI, auth → **Container** Docker → **Evaluate** golden sets

WEEK 1 Generative AI and LLM Fundamentals

- How transformers, tokens, embeddings and context work
- Model landscape: OpenAI, Anthropic, Gemini, open models
- Use cases, limits, hallucination and risk
- Python, Git and environment bridge track

HANDS-ON LAB
Call three model providers and compare output, cost and latency

YOU SHIP
Multi-provider comparison notebook

TOOLS Python, OpenAI, Claude, Gemini, Jupyter

WEEK 2 Prompt Engineering and Structured Outputs

- Prompt patterns: role, few-shot, chain-of-thought
- Structured outputs with JSON schema and Pydantic
- Tool and function calling basics
- Eval-first mindset: golden sets, regression tests

HANDS-ON LAB
Build an extraction service with validated outputs

YOU SHIP
Prompt library with an evaluation harness

TOOLS Pydantic, OpenAI and Claude APIs, pytest

WEEK 3 API Engineering for GenAI Apps

- Async Python, FastAPI, streaming responses
- Auth, error handling, rate limits, retries
- Caching, batching and token budgets
- Docker packaging basics

HANDS-ON LAB
Build and test a secured LLM API

YOU SHIP
Tested FastAPI GenAI service with docs

TOOLS FastAPI, Docker, Redis, Postman

CHECKPOINT

Phase 1 review. Demo your LLM service and prompt library to a mentor panel.

AFTER THIS PHASE, YOU CAN Call and compare LLMs, design reliable prompts and ship a tested GenAI API.

PHASE 2 | WEEKS 4-7

RAG Systems: build AI that survives real data

Retrieval, graphs, multimodal documents and evaluation, all measured.

Documents PDF, DB, wiki → **Chunk + embed** metadata → **Vector DB** hybrid index → **Retrieve** rerank → **LLM** grounded → **Answer** citations

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

WEEK 4 Embeddings, Vector Databases and Core RAG

- Embeddings, chunking strategies, metadata design
- Vector stores: Qdrant, Pinecone, pgvector
- Ingesting PDFs, wikis and databases
- Grounded answers with citations

HANDS-ON LAB
RAG pipeline over a messy document set

YOU SHIP
Working RAG app with citations

TOOLS LangChain, Qdrant, Pinecone, pgvector

WEEK 5 Advanced RAG and Evaluation

- Hybrid search, reranking, query rewriting
- Multi-hop, parent-document, agentic retrieval
- RAGAS and DeepEval: faithfulness, relevance, recall
- Failure analysis and error taxonomies

HANDS-ON LAB
Improve retrieval and run an evaluation study

YOU SHIP
RAG system with an evaluation report

TOOLS RAGAS, DeepEval, LangChain, Elasticsearch

WEEK 6 Knowledge Graphs and GraphRAG

- Graph modelling with Neo4j and Cypher
- GraphRAG patterns: when graphs beat vectors
- Combining graph, vector and keyword retrieval

HANDS-ON LAB
Add graph retrieval to your RAG system

YOU SHIP
Hybrid graph-vector retrieval pipeline

TOOLS Neo4j, Cypher, LangChain

WEEK 7 Multimodal AI and Document Intelligence

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

HANDS-ON LAB
Parse invoices and reports into structured data

YOU SHIP
Document-intelligence service

TOOLS VLMs, ColPali, Whisper, PostgreSQL

CHECKPOINT

Project 1: Knowledge Copilot. Deliver a secure copilot with hybrid RAG, text-to-SQL, role-based access and a stakeholder demo.

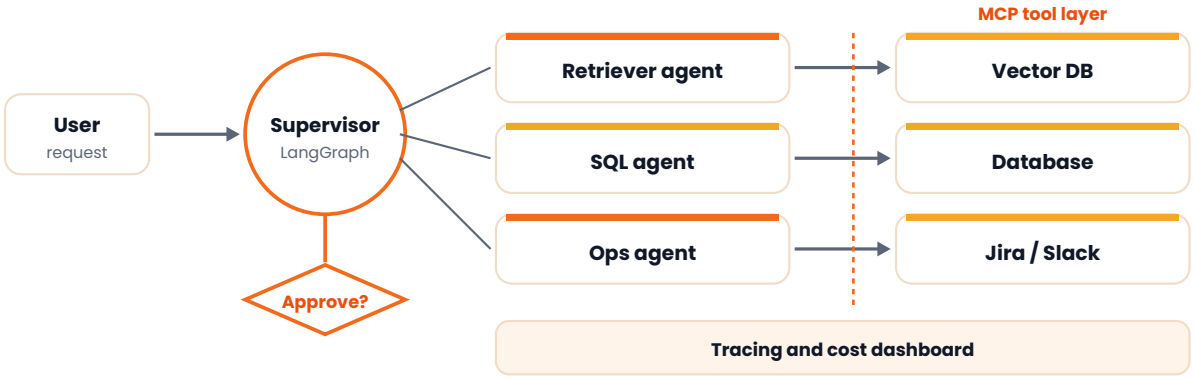
AFTER THIS PHASE, YOU CAN Build, evaluate and improve RAG, graph and multimodal systems on messy data.

PHASE 3 | WEEKS 8-12

Agentic AI: systems that plan, act and connect safely

Agents that use tools, work in teams and connect to the systems enterprises actually run.

A SUPERVISOR-BASED MULTI-AGENT SYSTEM WITH GOVERNED TOOLS



WEEK 8 Agent Design and LangGraph

- Patterns: ReAct, plan-and-execute, reflection, routers
- LangGraph: state, nodes, edges, memory, checkpoints
- Tool design: idempotency, permissions, failures

HANDS-ON LAB
Build a stateful, tool-using agent

YOU SHIP
Agent with evaluation traces

TOOLS LangGraph, LangChain, Langfuse

WEEK 9 Tool Use and Model Context Protocol (MCP)

- Function calling at scale; tool registries
- Build MCP servers; scopes and governance
- Secure tool access: allow-lists, sandboxes, audit logs

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

YOU SHIP
Governed tool layer

TOOLS MCP SDK, FastAPI, Jira, PostgreSQL

WEEK 10 Multi-Agent Orchestration

- Supervisor, router, parallel and hierarchical teams
- LangGraph, CrewAI, OpenAI Agents SDK, Google ADK
- Agent-to-agent handoffs (A2A)

HANDS-ON LAB
Supervisor-based multi-agent workflow

YOU SHIP
Multi-agent system with trace dashboard

TOOLS LangGraph, CrewAI, ADK

WEEK 11 Memory, Human-in-the-Loop and Agent Evaluation

- Short- and long-term memory design
- Human approval gates and escalation design
- Agent evals: task success, tool accuracy, cost, loops

HANDS-ON LAB
Add approvals and an eval suite to your agents

YOU SHIP
Agent evaluation report

TOOLS LangGraph, Langfuse, DeepEval

WEEK 12 **Enterprise Integration and Agent Platforms**

- Slack, Jira, ServiceNow, Salesforce, SQL integrations
- Queues, webhooks, retries, dead-letter handling
- Databricks Agent Bricks, Snowflake Cortex, Bedrock, Azure AI Foundry

HANDS-ON LAB
Build the same assistant on two platforms

YOU SHIP
Integration layer and platform comparison memo

TOOLS Kafka, Databricks, Snowflake, Bedrock, Azure

CHECKPOINT

Project 2: Multi-Agent Ops Copilot. Deliver a supervisor-based copilot with MCP tools, Jira and ServiceNow integration, human approval and a cost and trace dashboard.

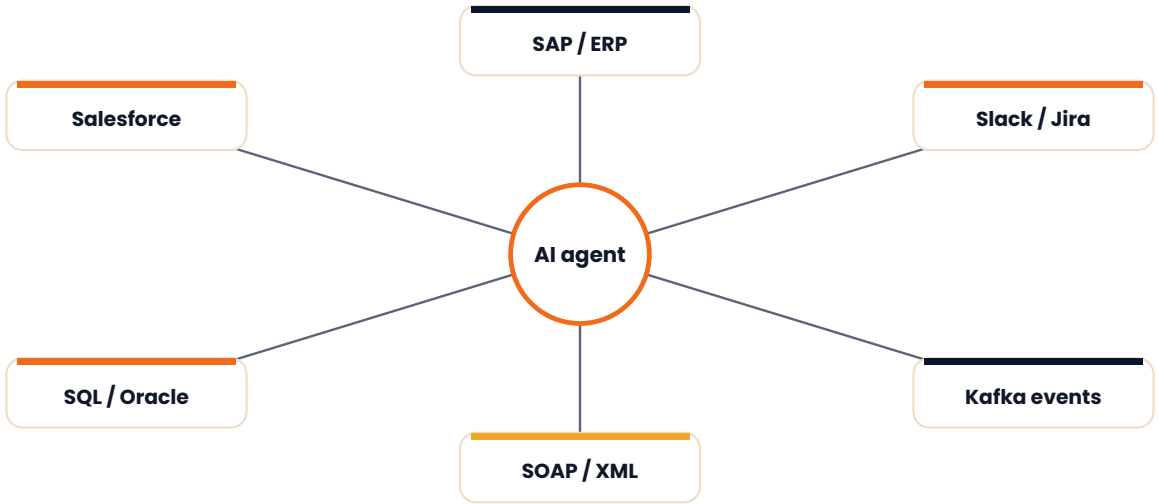
AFTER THIS PHASE, YOU CAN Design governed multi-agent systems and integrate them with enterprise tools.

CHOOSING AN AGENT FRAMEWORK

The trade-offs you will learn to judge

FRAMEWORK	BEST FOR	TRADE-OFF TO WEIGH
LangGraph	Stateful, controllable workflows with checkpoints and human approval	More code and a steeper learning curve
CrewAI	Quickly assembling role-based agent teams	Less low-level control over execution
OpenAI Agents SDK	Lightweight agents with handoffs and guardrails	Fits best inside the OpenAI ecosystem
Google ADK	Multi-agent apps on Google Cloud and Gemini	Fits best inside the Google ecosystem

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



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

PHASE 4 | WEEKS 13-16

Production and Security: ship AI that enterprises will approve

Secure, measure, govern and deploy AI systems that regulated enterprises will actually sign off.

DEFENCE IN DEPTH FOR AI AGENTS

Identity	Guardrails	Audit	Governance
OAuth 2.0, SSO, RBAC	injection, PII masking	traces, logs, approvals	DPDP, GDPR, HIPAA, EU AI Act

WEEK 13 AI Security and Guardrails

- Prompt injection, jailbreaks, exfiltration; OWASP LLM Top 10
- Excessive agency and blast-radius control
- PII masking, NeMo Guardrails, cloud guardrails
- OAuth 2.0, SSO and role-based access control

HANDS-ON LAB
Red-team your own agent, then patch it

YOU SHIP
Threat model and hardening report

TOOLS Presidio, NeMo Guardrails, Okta, Auth0

WEEK 14 LLMOps and Observability

- LLM gateways, fallbacks, rate limits, quotas
- Tracing with LangSmith, Langfuse, OpenTelemetry
- 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 15 Fine-Tuning, Small Models and Cost Engineering

- Fine-tune vs prompt vs RAG; LoRA and PEFT overview
- Small and open models; private VPC deployment
- Caching, batching and 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

WEEK 16 Cloud Deployment, Scale and Governance

- Docker, CI/CD, Terraform on AWS, Azure or GCP
- Scale, SLOs, runbooks and incident response
- DPDP, GDPR, HIPAA and EU AI Act essentials

HANDS-ON LAB
Deploy to cloud, load test and run an incident drill

YOU SHIP
Live deployment, runbook and governance pack

TOOLS Docker, Terraform, GitHub Actions, Kubernetes, Grafana

CHECKPOINT

Project 3: Production Hardening. Take Project 2 to production: security review, compliance pack, red-team report, UAT and handoff.

AFTER THIS PHASE, YOU CAN Secure, govern, deploy and operate GenAI and agentic systems in production.

PHASE 5 | WEEKS 17-18

Capstone and Career Launch

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

WEEK 17 Capstone: Build and Harden

- Choose a track: BFSI, healthcare, enterprise ops or your own
- Discovery, scoping, architecture decision records
- Build with RAG, agents, integrations; add evals, security, CI/CD

HANDS-ON LAB
Weekly reviews with your mentor panel

YOU SHIP
Deployed, documented, secured capstone

TOOLS Your stack of choice from the program

WEEK 18 Demo Day, Interviews and Job Launch

- Demo Day: defend your capstone to an industry-style panel
- GenAI and agentic AI interview loops
- Resume, LinkedIn and portfolio teardown

HANDS-ON LAB
Two mock interview loops with feedback

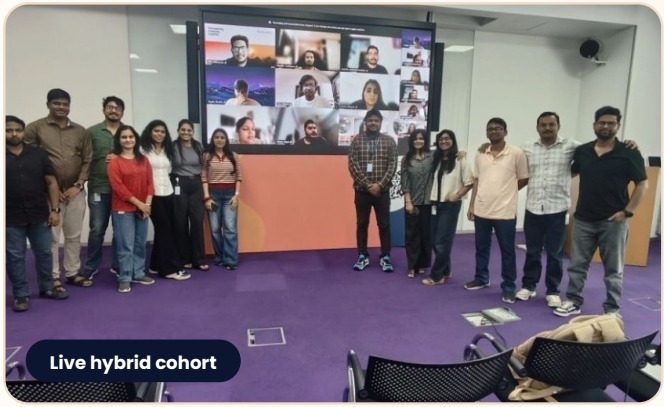
YOU SHIP
Portfolio, resume, LinkedIn and job-search plan

TOOLS GitHub, portfolio site, LinkedIn

THREE PROJECTS, ONE CAPSTONE, ONE DEMO DAY, ONE OFFER-READY PORTFOLIO




Corporate cohort


Live hybrid cohort

PROJECT STUDIO

Four client-style projects, not toy apps

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

P1 Knowledge Copilot

BFSI client: secure knowledge copilot

Hybrid RAG, secure text-to-SQL, role-based access, PII masking and guardrails. Deliverables: discovery notes, architecture and stakeholder demo.

Python | FastAPI | PostgreSQL | Docker

P2 Multi-Agent Ops Copilot

Enterprise operations: multi-agent copilot

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

LangGraph | MCP | Jira | Slack | Kafka

P3 Production Hardening

Regulated deployment

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

Okta | OpenTelemetry | Grafana | Terraform

Capstone

Your domain, your client

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

Databricks | Snowflake | CI/CD | Kubernetes

THE AI TOOLKIT

Professional artifacts you will produce

Prompt and eval kit

Prompt library, golden sets, regression tests, evaluation reports.

Architecture pack

Diagrams, ADRs, integration design and data classification.

RAG and agent templates

Reusable retrieval pipelines, LangGraph agents and MCP servers.

Security and compliance

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

Delivery pack

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

Operations pack

Runbooks, SLOs, dashboards and incident postmortem.

WHAT MAKES THESE PROJECTS DIFFERENT

Real client brief

Stakeholders, constraints and security requirements, like a real engagement.

Panel review

Each project is defended in front of a mentor panel, not just auto-graded.

Portfolio-ready

Deployed, documented and explainable in an interview, line by line.

THE TECHNOLOGY

The stack production AI teams use

MODELS AND AI PLATFORMS

- OpenAI
- Claude
- Gemini
- Hugging Face
- Ollama
- Databricks
- Snowflake
- AWS Bedrock
- Azure AI Foundry
- Google Cloud

AGENTS AND ORCHESTRATION

- LangChain
- LangGraph
- CrewAI
- OpenAI Agents SDK
- Google ADK
- MCP
- A2A
- Pydantic
- Python

RETRIEVAL, DATA AND EVENTS

- Qdrant
- Pinecone
- pgvector
- Neo4j
- PostgreSQL
- Elasticsearch
- Redis
- Kafka
- SQL Server
- Oracle

BUILD AND DEPLOY

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

SECURITY AND IDENTITY

- OWASP LLM Top 10
- Okta
- Auth0
- Presidio
- NeMo Guardrails

EVALS AND OBSERVABILITY

- RAGAS
- DeepEval
- Langfuse
- LangSmith
- OpenTelemetry
- Grafana
- MLflow

ENTERPRISE INTEGRATIONS

- Slack
- Jira
- ServiceNow
- Salesforce
- SAP
- SOAP and XML
- Webhooks

Product and company names are trademarks of their respective owners and are shown for identification only; no affiliation is implied.

CAREER LAUNCH

Job-focused support, every single week

Runs in parallel with the technical modules so your profile is ready when your skills are.

PHASE	CAREER MILESTONE	OUTPUT
Weeks 1-3	Role targeting and skills-gap audit	Target role list and 90-day plan
Weeks 4-6	GitHub portfolio setup	Repo template: README and architecture diagram standard
Weeks 7-9	Resume rewrite for GenAI and agentic AI roles	ATS-friendly resume with quantified project bullets
Weeks 10-12	LinkedIn and Naukri profile building	Optimised profiles and project posts
Weeks 13-14	Interview bank: LLMs, RAG, agents	Practice questions with model answers
Weeks 15-16	AI system-design interviews	Three worked cases: support agent, RAG at scale, LLM gateway
Week 17	Security, ethics and behavioural rounds	STAR stories and security case discussion
Week 18	1:1 mock interviews and job-search sprint	Feedback report, referral and outreach playbook

INTERVIEW-READY

Sample questions we prepare you for

Q. How would you reduce hallucinations in a RAG chatbot?

Q. ReAct vs plan-and-execute: how do you choose?

Q. Design a multi-tenant LLM gateway with cost controls.

Q. How do you protect an agent with database access from prompt injection?

Q. What does your CI/CD pipeline check for an LLM application?

Q. MCP vs function calling: when do you migrate?

Q. How do you evaluate an agent end to end?

Q. LangGraph vs CrewAI: what drives your decision?

What you get

- ✓ Resume review and rewrite
- ✓ LinkedIn and Naukri optimisation
- ✓ GitHub portfolio reviews
- ✓ 1:1 mock interviews with feedback
- ✓ Interview question banks
- ✓ 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 AI systems and the ability to explain every design decision like an engineer. That is what gets offers.

HIRING PATHS

Where graduates aim

AI labs and AI-first startups
Product-building and customer-facing AI engineering roles.

Data and cloud platforms
Applied AI roles at data, AI and cloud platform vendors.

GCCs, consulting and enterprises
AI teams in global capability centres, consulting firms and enterprises.

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 engineers 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



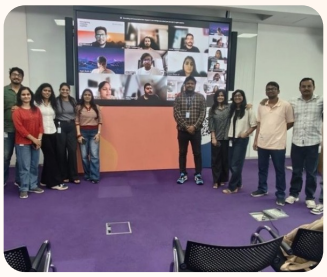
Corporate cohort



Sigmoid delivery



Hands-on lab



Live hybrid cohort

Guest masterclasses

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

PROVEN IN THE FIELD

Career-ready learners from our programs



Gaurav Taneja
PwC
Senior Azure Data Engineer



Lakshan J
Hexaware Technologies
Software Engineer



Sharoz F. Afridi
Career growth
Technology professional



Richard Leclezio
LinkedIn
Professional learner



Uphar Tandon
Microsoft
Technology professional



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**

> *"Learning under Deepesh Sir has been a great opportunity. All my doubts were dispelled and I was able to hone my skills with great hands on activities and projects."*
>
> **Ansh Ranjan, Data Engineer Trainee @ Hexaware**

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



Databricks certification cohort

"The program was well-structured, engaging, and highly practical. The hands-on exercises helped deepen my understanding and build real confidence for the exam."

Anonymous participant feedback

ADMISSIONS

Four steps to your seat

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

1 — Apply
Short online application with your background and goals.

2 — Diagnostic
Technical and fit assessment: Python, APIs, problem solving.

3 — Interview
Conversation with the admissions mentor about your target role.

4 — Offer
Cohort allocation, bridge pre-work and onboarding.

DataMarcos Executive Certificate in Generative AI and Agentic AI Engineering

Awarded on completing the program and defending your capstone. It validates hands-on proficiency in GenAI, agents, enterprise integration, security and delivery, backed by a portfolio of deployed projects 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 Python basics and working engineering fundamentals, or strong projects. A bridge pre-work track covers Python, Linux, Git and SQL before Day 1.

How much time per week?
Plan for about 8 to 12 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, AI interview preparation, mock loops and a referral and outreach playbook.

Is security covered in depth?
Yes. A full week on AI security, a governance and compliance module, plus red-teaming in the projects.

Do you cover Databricks and Snowflake?
Yes. Week 12 covers 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.

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