DATAMARCOS
IT CONSULTING FIRM

FUTUREGYAAN x DATAMARCOS
Flagship Certification 2026

JOB-ORIENTED CERTIFICATION PROGRAM

Applied Agentic AI Engineer

Go from using ChatGPT to architecting, securing and shipping production AI agents. 8 weeks of live, hands-on training in LLM architecture, RAG, multi-agent systems, MCP, Docker, CI/CD and AI SecOps, with a career track built in.

8 Weeks | Live Online · **6 Portfolio Projects** · **Docker + CI/CD + AI SecOps** · **Career Track**

25,000+
Professionals trained

200+
Organisations

18+
Years of experience

10+
Countries

TEAMS OUR MENTOR HAS TRAINED

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



DEEPESH MISHRA

Founder, DataMarcos | Your mentor

Fees and EMI options shared on a counsellor call.

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

PROGRAM AT A GLANCE

Become a job-ready Applied Agentic AI Engineer

An architecture-first, project-driven program for engineers and technology professionals who want to build, secure and ship AI agents in production.

8 Weeks	6 Projects	Intermediate	Live Online	Certificate
Live, mentor-led + weekly labs	5 portfolio builds + capstone	Basic Python + pre-work	Small batches, real-time Q&A	DataMarcos Agentic AI

WHY THIS PROGRAM

Built the way hiring managers hire

Architecture-first
Learn why AI systems are designed the way they are: reference architectures, trade-offs and decision records, not just which buttons to click.

The full production stack
Docker, CI/CD, observability and AI SecOps are core modules, not optional add-ons.

A portfolio that gets noticed
Six deployable projects, each with an architecture diagram and README, ready for GitHub and LinkedIn.

Taught by a practising architect
18+ years of industry experience; has trained 25,000+ professionals across 200+ organisations.

Career track built in
Resume, LinkedIn, GitHub, interview question banks and 1:1 mock interviews run alongside the technical modules.

Small, live batches
Real-time Q&A, code reviews and doubt resolution. Every week ends with something you ship.

OUTCOMES

By the end, you can...

- ✓ Architect an AI application end to end: gateway, RAG, agents, API and UI
- ✓ Build tool-using single and multi-agent systems with LangChain, LangGraph and CrewAI
- ✓ Design production RAG with hybrid search, reranking and RAGAS evaluation
- ✓ Build MCP servers and integrate agents with databases, APIs and n8n workflows
- ✓ Containerise with Docker and ship through CI/CD pipelines with evaluation gates
- ✓ Defend agents against prompt injection, data leakage and excessive agency (OWASP LLM Top 10)
- ✓ Deploy on Streamlit and Vercel and monitor cost, latency and quality with LangSmith / Langfuse
- ✓ Explain your design decisions with confidence in AI system-design interviews

Optional add-on: Google Cloud specialisation
Gemini Enterprise Agent Platform (formerly Vertex AI), ADK, Agent Runtime, Model Armor, MCP and A2A on Google Cloud. See the Google Cloud module pages after the syllabus.

WHO SHOULD ENROLL

Designed for working professionals and career switchers

Software & backend engineers
Add agentic AI to your stack and move into AI engineering roles.

QA & test automation engineers
Build and evaluate AI agents, not just test scripts.

Tech leads & architects
Lead AI initiatives with sound architecture and security.

Data & AI professionals
Move from notebooks to production agentic systems.

Technical PMs & project managers
Scope, estimate and govern AI-native products credibly.

Career changers & founders
With Python basics, build a portfolio or launch your own AI product.

Prerequisites
Basic Python (a self-paced pre-work track is provided before Day 1) • comfort with APIs and Git is helpful • a laptop with Chrome • free accounts on GitHub, Groq, Vercel and Streamlit (we guide you). No prior AI experience required.

ROLES YOU BECOME READY FOR

Map your learning to real job roles

Role readiness, not a placement guarantee. Projects and interview practice are mapped to what employers ask for.

ROLE	WHAT EMPLOYERS LOOK FOR	WHERE YOU BUILD IT
AI Agent Engineer	Tool-using agents, orchestration, reliability, evaluation	Weeks 2, 5 | Support agent, research coordinator
GenAI / LLM Application Developer	RAG, prompt engineering, model APIs, structured outputs	Weeks 1, 3 | LLM gateway, knowledge copilot
AI Solutions Engineer / Architect	Architecture, integration, trade-offs, security reviews	Weeks 1, 4, 7 | Architecture blueprint, threat model
AI Automation Engineer	n8n, MCP, enterprise integration, workflow design	Weeks 5, 8 | MCP server, integrated capstone
LLMOps / AI DevOps Engineer	Docker, CI/CD, observability, evaluation gates	Week 6 | Pipeline with eval and security gates
AI Product Engineer	End-to-end builds, UX, deployment on Vercel/Streamlit	Weeks 4, 8 | Live web app, capstone demo

LEARNING PATH

Eight weeks, one parallel career track

0	1	2	3	4	5	6	7	8
Pre-work	**LLM engineering**	**Agent design**	**RAG at scale**	**Production app**	**Multi-agent + MCP**	**DevOps**	**AI SecOps**	**Integration + capstone**
Python, Git, APIs	Architecture, gateway	Tools, memory, patterns	Retrieval, evals	FastAPI, Vercel	LangGraph, CrewAI	Docker, CI/CD, observability	Threats, guardrails	n8n, launch, interviews

Running in parallel every week: career track (portfolio, resume, LinkedIn, interview preparation). Optional add-on: Google Cloud module (G1-G2) after Week 8.

DETAILED SYLLABUS

Week-by-week curriculum

Every week combines live architecture sessions, hands-on labs, interview focus and a shipped deliverable.

Week 0 · Foundation pre-work (self-paced, before Day 1)

Python for AI: environments (venv), Git and GitHub workflow, requests and REST APIs, JSON, async basics, Pydantic and type hints • LLM vocabulary refresher: tokens, context window, embeddings • VS Code and Jupyter setup • Pre-work quiz so everyone starts on the same page.

WEEK 1 LLM Engineering and AI SaaS Architecture

ARCHITECTURE

1.1 How LLM applications are built

- Tokens, context windows, sampling (temperature, top-p/top-k), failure modes
- Prompting vs RAG vs fine-tuning; SLM vs LLM; LoRA / QLoRA overview

1.2 Model APIs and gateways

- OpenAI, Gemini, Claude, Groq, Hugging Face and Ollama (local models)
- Auth, rate limits, retries, streaming, cost and latency budgets

1.3 Prompt engineering for production

- System prompts, few-shot, chain-of-thought, structured JSON with Pydantic
- Prompt versioning, templates and first prompt-injection defenses

1.4 Reference architecture

- Model gateway, routing, caching, async queues, multi-tenancy
- Architecture decision records (ADRs) and attack-surface mapping

HANDS-ON LABS

- Multi-model playground in Streamlit
- JSON-schema extraction service
- Token and cost calculator

INTERVIEW FOCUS

- How do tokens and context limits shape design?
- Design an LLM gateway for multiple providers

YOU SHIP

Architecture blueprint + multi-model LLM gateway (live on Streamlit)

TOOLS THIS WEEK OpenAI, Gemini, Claude and Groq APIs • Hugging Face • Ollama • Pydantic • Streamlit

WEEK 2 Agent Design Patterns and Tool Architecture

AGENTS

2.1 Agent anatomy

- Goals, planning, tools, memory, reflection; chatbot vs agent
- ReAct, plan-and-execute, Reflexion; when not to use agents

2.2 Tool calling

- Function schemas, structured outputs, validation and error handling
- Idempotent, side-effect-safe tool design

2.3 Memory architecture

- Short-term, summary, long-term and vector memory
- Per-user and per-tenant memory, state persistence

2.4 Frameworks

- LangChain and LCEL, SDK-native agents, AutoGen and Google ADK overview
- Choosing LangGraph vs CrewAI vs plain code

HANDS-ON LABS

- Researcher and summariser agent
- Customer-care agent with live tool calling
- Share a live Streamlit link

INTERVIEW FOCUS

- ReAct vs plan-and-execute: trade-offs?
- How do you make tool calls idempotent?

YOU SHIP

Customer-support agent v1 with tools and memory, deployed and shareable

TOOLS THIS WEEK LangChain + LCEL • SDK function calling • AutoGen • Google ADK • Streamlit Cloud

WEEK 3 RAG Architecture at SaaS Scale

RETRIEVAL

3.1 Embeddings and search
- Embedding models (OpenAI, BGE/E5, Cohere), similarity metrics
- ANN indexes: flat, IVF, HNSW and their trade-offs

3.2 Ingestion pipelines
- Document loaders (PDF, DOCX, HTML), cleaning, metadata
- Chunking: recursive, semantic, structure-aware; versioning knowledge bases

3.3 Advanced retrieval
- Hybrid search, metadata filters, multi-query, parent-document retrieval
- Reciprocal Rank Fusion, cross-encoder and Cohere reranking

3.4 Vector stores and tenancy
- FAISS, ChromaDB, pgvector, Pinecone, Weaviate: when to use which
- Tenant isolation and access-controlled retrieval

3.5 Evaluation
- Precision, recall, MRR, NDCG; RAGAS faithfulness and relevance
- Hallucination mitigation, citations, regression tests

HANDS-ON LABS
- Knowledge-base chatbot with hybrid search + reranking
- RAGAS evaluation report
- Tenant-isolated retrieval demo

INTERVIEW FOCUS
- How do you improve RAG recall and faithfulness?
- Chunking strategy for tables and long PDFs?

YOU SHIP

Enterprise knowledge copilot with citations and an evaluation report

TOOLS THIS WEEK ChromaDB • FAISS • pgvector • Pinecone • Weaviate • Cohere Rerank • RAGAS

WEEK 4 Production Backend and Web App

VERCEL + STREAMLIT

4.1 FastAPI service layer
- Async endpoints, Pydantic models, streaming (SSE), background jobs
- JWT / API-key auth, rate limiting, token-budget enforcement

4.2 Frontends and hosting
- Streamlit for MVPs and internal tools; Gradio and Hugging Face Spaces
- Next.js on Vercel: preview deploys, env vars, custom domains

4.3 Auth and data
- Supabase auth, Postgres with row-level security, Redis caching
- Per-tenant data isolation patterns

4.4 API design
- Versioning, error contracts, OpenAPI docs, CORS
- Load testing and latency profiling

HANDS-ON LABS
- Deploy your RAG agent behind an API
- Vercel app with signup and login
- Load test and optimise

INTERVIEW FOCUS
- Design a multi-tenant AI API
- Streamlit vs Vercel: when to use which?

YOU SHIP

Live SaaS on a Vercel URL with login, your agent behind a secured API

TOOLS THIS WEEK FastAPI • Next.js • Vercel • Supabase • Redis • Gradio • Hugging Face Spaces

WEEK 5 Multi-Agent Orchestration and MCP

ORCHESTRATION

5.1 Orchestration patterns
- Sequential, parallel, hierarchical teams; coordinator-dispatcher
- Role-based specialists vs generalists; consensus and conflict resolution

5.2 LangGraph in depth
- State schema, nodes, edges, conditional routing, cycles, reducers
- Checkpointing, sub-graphs, streaming, human-in-the-loop interrupts

5.3 CrewAI and AutoGen
- Agents, tasks and crews; group chats and handoffs
- Selecting a framework for your use case

5.4 MCP and A2A
- Client-server-transport architecture; building MCP servers
- Tool registries, permission scoping, migration from legacy tool-calling; agent-to-agent protocols overview

5.5 Failure design
- Retries, timeouts, fallbacks, circuit breakers, idempotency

HANDS-ON LABS
- Financial news research coordinator
- MCP server over a live database + internal API
- Human approval workflow

INTERVIEW FOCUS
- MCP vs plain function calling?
- LangGraph vs CrewAI: how do you choose?

YOU SHIP

Multi-agent research coordinator + an MCP server exposing company data safely

TOOLS THIS WEEK LangGraph • CrewAI • AutoGen • MCP SDK • A2A concepts

WEEK 6 DevOps: Docker, CI/CD and Observability

DEVOPS

6.1 Containerisation with Docker
- Images, multi-stage builds, slim images, cold-start tuning
- docker-compose for API, vector DB and workers; env and secrets

6.2 CI/CD pipelines
- GitHub Actions: lint, unit tests, prompt and eval regression gates
- Build and push images, automated deploys, rollbacks, branching strategy

6.3 Environments and cloud
- Dev, staging, prod; config management
- Container deployment on AWS, GCP and Azure (overview)

6.4 Observability and evals
- LangSmith and Langfuse tracing, token, latency and cost dashboards
- Online evaluation, LLM-as-judge, feedback loops

HANDS-ON LABS
- Dockerise the Week 4 app
- Pipeline: test, eval, build, deploy on every push
- Tracing and cost dashboard

INTERVIEW FOCUS
- How do you CI/CD an LLM app?
- What do evaluation gates check?

YOU SHIP

Dockerised app with an automated CI/CD pipeline and live observability

TOOLS THIS WEEK Docker • docker-compose • GitHub Actions • LangSmith • Langfuse • AWS / GCP / Azure (overview)

WEEK 7 AI SecOps and Responsible AI

SECOPS

7.1 Threat modelling
- OWASP Top 10 for LLM apps: prompt injection, insecure output handling, data leakage, excessive agency, supply chain
- Attack trees for agents with tool access

7.2 Guardrails and runtime defense
- Guardrails AI, input validation, output moderation, PII redaction
- Rate limiting, abuse prevention, content policies

7.3 DevSecOps
- Secrets management, secret scanning in CI, dependency and container image scanning
- Least-privilege tools, RBAC, tenant isolation, audit logs

7.4 Responsible AI and compliance
- Bias, hallucination, explainability; human-in-the-loop policy design
- GDPR and India DPDP basics, EU AI Act overview, incident-response runbooks

HANDS-ON LABS
- Threat model for your product
- Red-team test suite against your agent
- Add security gates to the CI pipeline

INTERVIEW FOCUS
- How do you stop prompt injection when tools are connected?
- What is excessive agency and how do you limit it?

YOU SHIP

Threat model + red-team suite + security gates integrated into CI

TOOLS THIS WEEK OWASP LLM Top 10 • Guardrails AI • secret and dependency scanners • red-team tooling

WEEK 8 Integration, Capstone and Job Launch

INTEGRATION

8.1 Enterprise integration
- n8n workflows and webhooks; Gmail, Sheets, Slack, WhatsApp, Telegram, CRM
- n8n vs Zapier vs Make; Stripe billing basics

8.2 Voice and multimodal
- Speech-to-text, LLM, action, text-to-speech pipelines
- Vision and document understanding

8.3 Capstone build
- RAG + agents + multi-agent + MCP in one product, Dockerised, CI/CD, secured
- Architecture review with the mentor

8.4 Demo Day and interview readiness
- Portfolio walkthrough, AI system-design round, final 1:1 mock interviews
- Job-search sprint: referrals, outreach, role targeting

HANDS-ON LABS
- End-to-end integration tests
- Voice-enabled assistant with workflow trigger
- Capstone demo and pitch

INTERVIEW FOCUS
- Walk me through your capstone architecture
- Where would it fail at 100x scale?

YOU SHIP

Fully integrated, deployed capstone, Demo Day presentation and certificate

TOOLS THIS WEEK n8n • Zapier / Make (overview) • Whisper • ElevenLabs • Stripe • Slack, Gmail, Sheets, WhatsApp

ADD-ON MODULE | GOOGLE CLOUD

Agentic AI on Google Cloud

Gemini Enterprise Agent Platform (formerly Vertex AI): build, deploy, secure and govern production agents the Google Cloud way. Two weeks: G1 Build, G2 Govern. Designed for GCP-focused engineers, architects and teams.

Know the names: Google rebranded the platform in April 2026

Vertex AI is now the **Gemini Enterprise Agent Platform**. Older tutorials and job posts still say Vertex AI, so we teach both names. Checked against Google Cloud docs and release notes (Oct 2026); the module tracks the live platform.

EARLIER NAME	CURRENT NAME	WHAT IT DOES
Vertex AI	Gemini Enterprise Agent Platform	Unified platform: models, agents, governance
Vertex AI Agent Builder	Agent Studio (low-code) + ADK (code-first)	Design and build agents visually or in Python, Go, Java, TypeScript
Agent Engine	Agent Runtime	Managed runtime: deployment, scaling, sessions, Memory Bank
Vertex AI Search, RAG Engine, Vector Search	Agent Search, RAG Engine, Vector Search	Enterprise search, managed RAG, vector retrieval
Agentspace	Gemini Enterprise	Where employees use and discover published agents

WEEK G1 Build and Ship Agents on Google Cloud

GOOGLE CLOUD | BUILD

G1.1 Platform landscape and setup

- Agent Studio, ADK, Agent Runtime, Model Garden, governance layer
- Model Garden: Gemini, Claude, Llama, Gemma; quotas and regions
- Projects, IAM roles, least-privilege service accounts

G1.2 Agent Development Kit (ADK)

- Agents, tools, sessions, state, callbacks; Go/Java/TS overview
- Workflow agents (sequential, parallel, loop), multi-agent teams
- adk web / run / api_server; ADK with LangGraph and CrewAI

G1.3 Deploying agents

- Agent Runtime: managed deployment, sessions, Memory Bank
- Cloud Run and GKE paths; choosing between them
- Agents CLI: create, evaluate, deploy, publish; Secret Manager

G1.4 Grounding and RAG on GCP

- RAG Engine, Agent Search, Vector Search; Google Search grounding
- BigQuery, AlloyDB, Cloud SQL (pgvector) as knowledge sources
- Agentic RAG: agents choose retrieval strategy and filters

HANDS-ON LABS

- Multi-agent ADK app (research + analyst)
- Deploy to Agent Runtime and Cloud Run
- RAG Engine knowledge agent with grounding

INTERVIEW FOCUS

- ADK vs LangGraph: when does each win?
- Agent Runtime vs Cloud Run vs GKE for agents?

YOU SHIP

Multi-agent ADK application live on Google Cloud with RAG grounding

TOOLS THIS WEEK Gemini Enterprise Agent Platform • ADK • Agent Runtime • Model Garden • RAG Engine • Agent Search • Cloud Run • GKE • Secret Manager • BigQuery

WEEK G2 **Govern, Secure and Scale on Google Cloud** GOOGLE CLOUD | GOVERN

G2.1 Interoperability: MCP and A2A
- MCP on GCP: McpToolset, remote MCP servers, servers on Cloud Run, MCP Toolbox for Databases
- A2A: agent cards, exposing and consuming remote agents, cross-language pipelines
- Agent Registry: catalog and govern agents and MCP servers

G2.2 AI SecOps on Google Cloud
- Model Armor with Agent Gateway: prompt injection, PII leakage, harmful content
- IAM access policies, VPC Service Controls, CMEK, Secret Manager, audit logs
- Sensitive Data Protection; least-privilege service accounts for tools

G2.3 Observability and evaluation
- Cloud Trace (OpenTelemetry), Cloud Logging, Cloud Monitoring, Agent Observability
- Gen AI Evaluation service and ADK evaluation: multi-turn datasets, regression gates
- Cost, latency and quota monitoring; alerting

G2.4 CI/CD and infrastructure
- Cloud Build or GitHub Actions to Artifact Registry to Cloud Run / Agent Runtime
- Terraform basics; dev, staging, prod projects; canary releases and rollbacks
- Publishing agents to Gemini Enterprise for organisation-wide use

G2.5 Data, voice and multimodal agents
- BigQuery analyst agent: natural language to SQL with human approval
- Gemini Live API and ADK streaming for voice and video agents

HANDS-ON LABS
- MCP server on Cloud Run + A2A agent handoff
- Harden your agent with Model Armor and IAM
- CI/CD pipeline with an evaluation gate

INTERVIEW FOCUS
- How do you secure an agent that can query BigQuery?
- MCP vs A2A: when do you use each?

YOU SHIP

Secured, observable, CI/CD-deployed agent published to Gemini Enterprise

TOOLS THIS WEEK Model Armor • Agent Gateway • Agent Registry • Cloud Trace • Cloud Build • Artifact Registry • Terraform • BigQuery • Gemini Live API

GCP capstone: Enterprise Data Analyst Agent

Input: business questions + company documents
Agents: BigQuery analyst, grounded researcher, compliance reviewer
Features: NL-to-SQL with human approval, RAG Engine grounding, MCP tools on Cloud Run, A2A handoff, Model Armor screening, CI/CD with eval gate
Deploy: Agent Runtime, published to Gemini Enterprise

Who should take this module
- GCP architects and cloud engineers moving into AI agents
- Teams standardising on Google Cloud and Gemini
- Graduates of the core program who want GCP depth

Prerequisites: core program Weeks 1-6 or equivalent (Python, LLM, RAG and agent basics) and a Google Cloud project with billing enabled. We guide setup and cost controls.

PORTFOLIO PROJECTS

Six projects you can show in an interview

Each project ships with an architecture diagram, README, tests and a live deployment.

1. Multi-Model LLM Gateway

Input: user prompts + routing rules

Agents: router, evaluator

Features: provider fallback, caching, cost tracking, JSON outputs

Deploy: Streamlit

Stack: Streamlit • OpenAI / Gemini / Claude / Groq • Pydantic • Redis cache

2. Enterprise Knowledge-Base Copilot

Input: PDFs, docs, wikis

Agents: retriever, reranker, answer composer

Features: hybrid search, citations, RAGAS evals, tenant isolation

Deploy: FastAPI + Vercel

Stack: LangChain • ChromaDB / pgvector • Cohere Rerank • RAGAS • FastAPI • Next.js

3. Customer Support Agent System

Input: customer queries (text, voice)

Agents: intent, retrieval, response, escalation

Features: tool calling, memory, sentiment, human handoff

Deploy: Chat UI + admin dashboard

Stack: LangChain • tools and APIs • Whisper • Supabase • Streamlit chat UI

4. Multi-Agent Financial Research Coordinator

Input: company or topic

Agents: search, summariser, analyst, fact-checker

Features: LangGraph/CrewAI orchestration, structured report, citations

Deploy: Interactive dashboard

Stack: LangGraph • CrewAI • web-search tool • FastAPI • Streamlit dashboard

5. MCP Integration Server

Input: agent tool requests

Agents: MCP client + server

Features: live DB + internal API access, permission scoping, audit log

Deploy: Docker container

Stack: MCP SDK • FastAPI • Postgres • Docker

6. Capstone: AI Agent Platform

Input: business workflow of your choice

Agents: multi-agent team + guardrails agent

Features: RAG + MCP + n8n integrations, CI/CD, security gates, observability

Deploy: Docker + Vercel

Stack: LangGraph • MCP • n8n • Docker • GitHub Actions • Vercel • Langfuse

EVERY PROJECT SHIPS WITH

Architecture	README	Tests	Live link
diagram + ADRs	setup + design notes	unit + eval checks	deployed demo

CAREER TRACK

Job-focused support, every single week

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

WEEK	CAREER MILESTONE	OUTPUT
1	Role targeting and skills-gap audit	Target role list and 90-day plan
2	GitHub portfolio setup	Repo template: README + architecture diagram standard
3	Resume rewrite for agentic AI roles	ATS-friendly resume with quantified project bullets
4	LinkedIn and Naukri profile building	Optimised profiles and project posts
5	Interview bank I: LLMs, RAG, agents	Practice questions with model answers
6	AI system-design interviews	Three worked cases: support agent, RAG at scale, LLM gateway
7	Security, ethics and behavioural rounds	STAR stories and security case discussion
8	1:1 mock interviews + job-search sprint	Feedback report, referral and outreach playbook

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 or salary hikes. We promise that you will be able to show working, deployed, secured AI systems and explain every design decision like an engineer. That is what gets offers.

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?

TOOLS COVERED

The production stack you will master

MODELS AND SDKS


OpenAI

Gemini

Claude

Llama

Hugging Face

Ollama

ElevenLabs

AGENT FRAMEWORKS


LangChain

LangGraph

CrewAI

Pydantic

Python

Jupyter

BUILD AND SHIP


FastAPI

Streamlit

Gradio

Next.js

Vercel

Supabase

Postgres

Redis

DEVOPS AND OBSERVABILITY


Docker

GitHub Actions

GitHub

Google Cloud

SECURITY


OWASP LLM Top 10

Snyk

INTEGRATIONS AND AUTOMATION


n8n


Zapier

Make

Slack

Gmail

Sheets

Telegram

WhatsApp

Stripe

GOOGLE CLOUD (ADD-ON MODULE)


Google Cloud

Gemini

Also covered: Groq API • AutoGen • Google ADK • MCP SDK • Agno overview • ChromaDB • FAISS • Pinecone • Weaviate • RAGAS • LangSmith • Langfuse • Guardrails AI • Whisper • tiktoken • Google ADK • Agent Runtime • Model Armor • BigQuery • Cloud Run

MEET YOUR MENTOR

Deepesh Mishra

Founder, DataMarcos | Solution Architect | Corporate Trainer | BITS Pilani alumnus



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 enterprise teams and hands-on cohorts for Sigmoid, ZS Associates and Novartis. He teaches the way real systems are built: architecture first, then code, then deployment, then interviews.

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

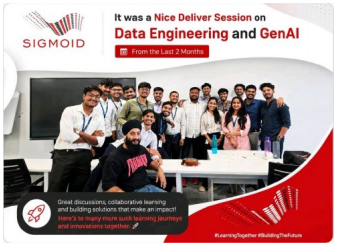
ORGANISATIONS WHOSE TEAMS HAVE BEEN TRAINED



PwC	Bank of America	Novartis	Sigmoid
ZS Associates	Hexaware	Yahoo Singapore	Micro Focus

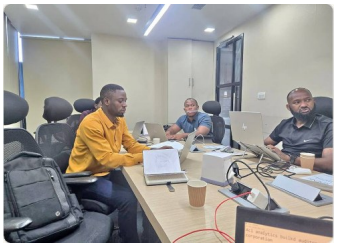
INSIDE THE SESSIONS


ZS Associates cohort



Sigmoid delivery


Live virtual cohort


Hands-on lab

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

"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

"Without you it wouldn't have been possible sir. Thank you for making me a Data Engineer."

Raghul K, Data Engineering learner

CERTIFICATE AND FAQ

Your questions, answered

DataMarcos Applied Agentic AI Certification

Awarded on completing the program and the capstone. It validates hands-on proficiency across LLM engineering, RAG, agents, multi-agent orchestration, MCP, DevOps and AI SecOps, backed by a portfolio of deployed projects you can share with employers.

Do I need to know Python?

Basic Python is enough. A self-paced pre-work track gets you ready before Day 1.

Is this too advanced for me?

The program is architecture-first and hands-on. If you can learn Python basics and follow along with labs, you can do it.

Will I get a job after this?

We do not guarantee placements. We give you projects, a portfolio, resume and LinkedIn support, interview preparation and 1:1 mock interviews to maximise your chances.

What do I need to start?

A laptop with Chrome and free accounts on GitHub, Groq, Vercel and Streamlit. Most labs use free tiers; a few may need small API credits.

How are classes delivered?

Live, mentor-led online sessions in small batches with weekly hands-on labs and real-time Q&A.

Is Google Cloud covered?

Yes. An optional Google Cloud specialisation covers the Gemini Enterprise Agent Platform (formerly Vertex AI), ADK, Agent Runtime, Model Armor, MCP and A2A.

Do I need a cloud account?

Core labs use free tiers. The Google Cloud module needs a GCP project with billing enabled; we guide setup and cost controls.

Batch dates and fees?

Fees and EMI options shared on a counsellor call. Contact the DataMarcos team at www.datamarcos.com.

Ready to build what is next?

Become an Applied Agentic AI Engineer in 8 weeks. Talk to the DataMarcos team about the next batch.

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

DATAMARCOS
IT CONSULTING FIRM