

| COURSE BROCHURE · 2026

Generative AI (Gen AI)

2 months · live + recorded · From prompts to production: ship three Gen AI applications — RAG, multi-tool agent, and fine-tuned model — with proper evals, guardrails, and monitoring.

DURATION

2 mo

MODULES

2

ROLES UNLOCKED

7+

SALARY BAND

6-45 LPA

| ABOUT THIS PROGRAMME

Every product team in 2026 is hiring engineers who can put LLMs into production — not just demo them in a notebook. This two-month deep-dive takes you from prompt engineering through full RAG, agents, fine-tuning, and evaluation. You'll work in Python with OpenAI, Anthropic, LangChain, LlamaIndex, vector databases (Pinecone, Weaviate, pgvector), and the AWS Bedrock + Azure OpenAI stacks. By the end you'll have shipped three production-grade LLM applications — a domain RAG, a multi-tool agent, and a fine-tuned model — with proper evals, guardrails, and observability. The same stack startup CTOs are paying senior salaries for.

| BUILT FOR

- Software engineers transitioning into Gen AI / LLM roles
- Backend developers building AI features into existing products
- ML practitioners moving from classical ML into LLM-powered systems

OUR COMMITMENT

Job-focused training. Real instructors. Lifetime access. We don't sell certificates — we get you hired.

WHY NOW · 2026 MARKET SIGNAL

Every Series A+ startup in India is hiring at least one Gen AI Engineer in 2026. The window for entry-level Gen AI roles narrows when the talent pool catches up — late 2027.

01

LinkedIn India: Gen AI Engineer postings up 320% YoY

02

Median Gen AI engineer salary at product cos: Rs 22 LPA (Naukri 2026)

03

Fewer than 8,000 production-ready Gen AI engineers in India today

TECHNOLOGY STACK

Technologies you'll master

The exact stack senior engineers + consultants are billing for in 2026 — nothing taught for show, everything taught because it lands on a job posting.

- OpenAI GPT-4 / GPT-5 API + function calling
- Anthropic Claude API + tool use
- Open-source LLMs: Llama 3, Mistral, Mixtral, Qwen
- LangChain + LlamaIndex orchestration frameworks
- Vector databases: Pinecone, Weaviate, ChromaDB, pgvector
- RAG pipelines + hybrid search + reranking
- Fine-tuning with LoRA / QLoRA on Hugging Face
- FastAPI + Docker + AWS / Azure deployment
- LangGraph for stateful AI agents
- OpenAI Whisper + ElevenLabs for voice AI

WHY THIS PROGRAMME

What makes this different

- ✓ Three production-grade projects: domain RAG, multi-tool agent, fine-tuned model
- ✓ Full evaluation engineering with Ragas, LangSmith, and custom eval suites
- ✓ Hands-on with OpenAI, Anthropic, Mistral, Llama 3, and open-weight models
- ✓ Vector database engineering — Pinecone, Weaviate, pgvector at production scale
- ✓ Deployment on AWS Bedrock + Azure OpenAI with cost and latency optimization
- ✓ Agent frameworks — LangGraph, CrewAI, AutoGen — built and pressure-tested in class
- ✓ Interview prep with 50+ real Gen AI questions from FAANG and Indian product cos
- ✓ Senior practitioner instructors — engineers who ship LLM systems, not academics

CAREER OUTCOMES

By the end, you will be able to

- ✓ Architect production RAG pipelines that serve 100K+ queries per day
- ✓ Fine-tune open-weight models (Llama 3, Mistral) using LoRA and QLoRA
- ✓ Build evaluation suites that catch hallucinations, leakage, and regressions
- ✓ Deploy multi-tool agents with proper memory, planning, and failure recovery
- ✓ Estimate and optimize LLM API costs to under Rs 5 per 1,000 queries at scale
- ✓ Implement guardrails for PII, prompt injection, and content safety
- ✓ Pass Gen AI Engineer interviews at product companies and AI-first startups
- ✓ Ship features end-to-end — from product spec to instrumented production

SALARY OPPORTUNITIES

Job roles + earning potential

Indian-market salary bands for 2026, sourced from active listings on Naukri, LinkedIn, and Instahyre. Bands assume the certifications and projects this programme delivers.

	JOB ROLE	EXPERIENCE	SALARY
01	Generative AI Developer	Fresher to 2+ Years	Rs 7–12 LPA
02	Prompt Engineer	Fresher to 2+ Years	Rs 6–10 LPA
03	NLP Engineer (LLM-based)	2 to 4 Years	Rs 10–18 LPA
04	Machine Learning Engineer (GenAI focus)	2 to 4 Years	Rs 12–20 LPA
05	AI Application Developer (RAG + LangChain)	2 to 4 Years	Rs 10–18 LPA
06	LLMOps Engineer	3 to 5 Years	Rs 15–25 LPA
07	Vector Database Engineer	3 to 5 Years	Rs 14–22 LPA
08	GenAI Cloud Engineer	4 to 6 Years	Rs 18–28 LPA
09	GenAI Architect	6+ Years	Rs 25–45 LPA
10	Lead Data Scientist (GenAI specialization)	6 to 8+ Years	Rs 25–40 LPA

CAREER PATHS

Roles you can target

Roles ordered from entry → senior. The same role titles you'll see on Naukri, LinkedIn, and Instahyre filtered to the tech stack you'll master.

- Generative AI Engineer
- LLM Application Developer
- AI Solutions Architect
- Prompt Engineer
- AI Product Engineer
- Conversational AI Engineer
- AI Research Engineer

Emerging — high-demand roles for 2026

Roles that didn't exist five years ago, paying premium because the talent pool is small.

- AI Agent Engineer
- RAG Systems Specialist
- LLM Ops Engineer
- Multi-Modal AI Developer

CAREER TRAJECTORY

Where this programme leads, in 5 years

Median 2026 salary bands for the role progression this curriculum unlocks. Numbers are conservative — top performers cross the upper band earlier.

YEAR 1

Junior Gen AI Engineer

Rs 8-15 LPA

Ship first RAG + first agent. Own one production LLM feature.

YEAR 3

Senior Gen AI Engineer

Rs 18-30 LPA

Lead a small LLM platform team. Own evals + cost optimization.

YEAR 5

Principal AI Engineer / Architect

Rs 35-60 LPA

Shape company AI strategy. Design platform for 100+ engineers.

COURSE CURRICULUM

What you'll build, module by module

Every module is taught hands-on with real datasets, instances, or codebases — not slides. By the end, you'll have built, not just watched.

Python & Data Structures

- Installation and Setup
- Installing Python and setting up a development envi...
- Syntax and Basic Constructs
- Variables and data types (integers, floats, strings, b...
- Basic input and output
- Comments and documentation
- Control Structures
- Conditional Statements: `if`, `elif`, `else`
- Loops: `for`, `while`
- Loop control statements (`break`, `continue`, `pass`)
- Functions
- Defining Functions: Parameters and return values
- Scope and Lifetime: Local and global variables
- Lambda Functions: Anonymous functions
- Core Data Structures
- Lists: Creating, accessing, modifying, and iterating ...
- Tuples: Creating and using tuples; Unpacking tuples
- Sets: Creating and using sets; Set operations (union,...
- Dictionaries: Creating and using dictionaries; Dictio...
- File Operations
- Reading and Writing Files: Opening, reading, writin...
- Working with different file modes (`r`, `w`, `a`, `rb`, `...
- Working with CSV and JSON: Reading from and wri...
- OOPs Basics
- Classes and Objects: Defining classes and creating ...
- Class Variables and Methods: Using class variables ...
- Inheritance: Single and multiple inheritance
- Polymorphism and Encapsulation: Method overridin...
- Exception Types: Common exceptions (`ValueError`...
- Try, Except Blocks: Using `try`, `except`, `else`, and `...
- Scientific Computing
- NumPy: Arrays and matrix operations
- Pandas: DataFrames for data manipulation; Readin...
- Data Visualization
- Matplotlib: Plotting graphs and charts
- Seaborn: Statistical data visualization

Generative AI

- GenAI and Its Industry Applications
- Introduction to Generative AI
- AI vs ML vs DL vs NLP vs Generative AI
- Generative AI principles
- What is the role of ML in Gen-AI
- Different ML techniques (Supervised, Unsupervised,...)
- Applications in various domains
- Ethical considerations
- NLP essentials
- Basic NLP tasks
- Different text classification approaches
- Frequency based – Bag of words, TF-IDF, N-gram
- Distribution Models – CBOW, Skipgram (Traditional...)
- Deep learning techniques – CNNs, RNNs, LSTMs, G...
- Auto encoders
- VAE's and applications
- GAN's and its applications
- Different types of GAN's and applications
- Different types of Language models
- Applications of Language models
- Transformers and its architecture
- BERT, RoBERTa, GPT variations
- Applications of transformer models
- What is Prompt Engineering
- What are the different principles of Prompt Engine...
- Types of Different Prompt Engineering Techniques
- How to Craft effective prompts to the LLMs
- Priming Prompt
- Prompt Decomposition
- Generative AI lifecycle
- What is RLHF
- LLM pre-training and scaling
- Different Fine-Tuning techniques
- What is Chunking
- What is the use of chunking the document
- What are the traditional effective chunking techniq...
- What are the problems and limitations with traditio...
- How to overcome the limitations of Traditional chun...
- Advanced Chunking Techniques:
- Character Splitting
- Recursive Character Splitting
- Document based Chunking
- Semantic Chunking
- Agentic Chunking
- What is RAG
- What are the main components of RAG
- High level architecture of RAG
- How to Build RAG using external data sources
- Advanced RAG
- What is LangChain
- What are the core concepts of LangChain
- Components of LangChain
- How to use LangChain agents
- LlamaIndex
- What are Vector Databases
- Why do we prefer Vector Databases over Tradition...
- Different Types of Vector Databases: OpenSource ...
- OpenSource: Chroma DB, Weaviate, Faiss, Qdrant
- Close-Source Vector Databases: Pinecone, Arango...
- Supervised Finetuning
- Repurposing-Feature Extraction
- Advanced techniques in Supervised Finetuning – PE...
- Text based LLMs
- Automatic Evaluation: BLEU Score, ROUGE Score, ...
- Human Evaluation: Coherence, Factuality, Originali...
- Automatic Evaluation: Pixel-level metrics, FID (Frec...
- Human Evaluation: Photorealism, Style, Creativity, ...
- Automatic Evaluation: FAD (Frechet Audio Distance...
- Human Evaluation: Perceptual Quality – PQ, PQ-Na...
- Automatic Evaluation: FVD (Frechet Video Distance...
- Human Evaluation: Visual Quality, Temporal Coher...
- Model Deployment and Management
- Scalability and Performance Optimization
- Security and Privacy
- Monitoring and Logging
- Cost Optimization
- Model Interpretability and Explainability
- Amazon Bedrock, Azure OpenAI
- What is LangSmith?
- Applications and use-cases
- What is agentic AI
- Building single agent
- Multi agents by using MCP
- OpenAI
- Hugging Face
- PyTorch
- TensorFlow
- LangChain & LangGraph
- ChatGPT
- Gemini
- Copilot

I INSTRUCTED BY

Trained by working practitioners

Not academics. Not pre-recorded YouTubers. Engineers and consultants who still ship production code at top companies — the exact people you want to learn from.



Ms. Shaik Razia

Industry Trainer

10+ yrs of experience educator with expertise in training, fine-tuning, and deploying Generative AI models across Text, Image, Audio & Video domains. Skilled in LLMs, Prompt Engineering, RAG systems, and Vector Databases. Expert in tools like LangChain, Streamlit, AWS Bedrock, OpenAI APIs, Hugging Face, and Pinecone. Strong foundation in ML, DL, and NLP, with a focus on real-world AI applications. Proven ability to deliver impactful GenAI training to professionals and students. Approved trainer by Raj Cloud Technologies.

APPROVED TRAINER · RAJ CLOUD TECHNOLOGIES

I WHAT'S INCLUDED

Everything you get beyond the curriculum

The sessions are the tip. Here's the full value stack a single fee unlocks — mentorship, community, career support, and lifetime updates.

- + Lifetime access to recordings — re-watch any session, forever
- + Private community access — get unblocked in minutes, not days
- + 1:1 mentor reviews on every portfolio project you build
- + Weekly doubt-clearing sessions led by the instructor
- + Resume + LinkedIn rewrite by a working hiring manager
- + Two full mock interviews with real-time feedback
- + Interview question bank — 50+ scenario-based questions
- + Real-world datasets + playground accounts you keep
- + Verifiable course completion certificate
- + Lifetime curriculum updates — when modules refresh, you get them

ENROLMENT

Pick your tier and enrol

Two tiers — pick what fits your budget. Both include lifetime access to recordings and a verifiable certificate.

MOST POPULAR

LIVE + RECORDED

Rs 38,499

- ✓ All live sessions + Q&A
- ✓ Lifetime recordings
- ✓ Doubt-clearing sessions
- ✓ Resume + interview prep
- ✓ Certificate of completion

RECORDED ONLY

Rs 26,949

Save Rs 11,550

- ✓ Lifetime access to recordings
- ✓ Self-paced learning
- ✓ Certificate of completion
- ✗ No live sessions
- ✗ No live Q&A

HOW TO ENROL

Three steps to your seat

01 Book a free demo

30-min walk-through with the instructor — ask anything before you commit.

02 Pick your tier + enrol

Live + Recorded or Recorded-only. UPI, card, EMI — all accepted.

03 Start on day one

Onboarding kit, course portal access, and your first live session calendar invite.

| HOW WE STACK UP

Compared to your alternatives

Free tutorials and generic bootcamps both have their place. Here's exactly where they fall short — and where the fee here pays for itself.

WHAT YOU GET	FREE YOUTUBE TUTORIALS	GENERIC BOOTCAMPS	RAJ CLOUD
Live practitioner instructors	×	×	✓
Real datasets + playground accounts	×	partial	✓
1:1 mentor reviews on your portfolio	×	×	✓
Mock interviews with feedback	×	partial	✓
Resume + LinkedIn rewrite	×	×	✓
Job referral network access	×	×	✓
Lifetime updates + recordings	partial	×	✓
Verifiable certificate	×	✓	✓

| QUESTIONS ANSWERED

Common questions, answered upfront

Q. Do I need prior experience to start?

Most students begin with little to no experience. We assume only basic computer literacy. The first two weeks ramp you up to the level the rest of the curriculum needs — no one is left behind.

Q. What if I miss a live session?

Recordings are uploaded within 24 hours of every session. Weekly doubt-clearing sessions catch you up on missed material. Your community is active 24/7. Nothing gets lost, only delayed.

Q. Is there a job placement guarantee?

We don't promise placement — we engineer outcomes. Resume + LinkedIn rewrite, two mock interviews, a referral network, and portfolio reviews are all included. The work happens; the offer is yours to earn.

Q. Can I pay in instalments?

Yes. UPI EMI, credit-card EMI, and a 3-month no-cost split-pay are all supported. Email admissions and we'll set up a plan that fits your cash flow.

Q. Does the certificate hold weight in interviews?

Yes. It's dated, verifiable via our website, and signed by the trainer. Alumni reference it on LinkedIn and inside offer-stage interview rounds — but the portfolio you build matters more, and that's the point.

Talk to admissions

EMAIL

info@rajcloudtech.com

PHONE

[+91 93804 87284](tel:+919380487284)

WEB

rajcloudtech.com

OUR COMMITMENT

Job-focused training. Real instructors. Lifetime access. We don't sell certificates — we get you hired.