Stint
.....transforming careers

COURSE BROCHURE · 2026

TRAINING + PLACEMENT · ONLINE AND BENGALURU

Generative AI

Become a GenAI / RAG Engineer

Build RAG systems that hold up outside a demo: chunking, embeddings, vector search, re-ranking and evaluation, and get placed as a GenAI engineer.

3-month live programme for working professionals (IT and non-IT) and people returning after a career gap

DURATION
3 months, live
Weekday evenings and weekends, alongside your job

FORMAT
Online or classroom
Classroom in HSR Layout, Bengaluru

TARGET SALARY
₹14L–₹35L
For working professionals

PLACEMENT
Support until you're hired
Mock interviews, resume rebuild, referrals

STACK YOU'LL WORK WITH

Python · SQL · PostgreSQL + pgvector · LangChain · LlamaIndex · Embeddings · Vector DBs · Re-ranking · RAG evaluation · LLM APIs · FastAPI

YOUR PATH WITH STINT

Profile assessment → Live training → Projects → Mock interviews → **Get placed**

15+
Placed every month

100+
Hiring companies

80%
Average salary hike

ELIGIBILITY 2+ years of work experience, IT or non-IT. Career breaks welcome.

WHAT YOU'LL LEARN

7 modules. Real, job-ready work.

Seven live modules: Python and SQL first, then retrieval-augmented generation from ingesting documents to retrieval, orchestration, measured quality, guardrails and deployment.

01	**Python for AI engineering**	Core Python: functions, classes, error handling and virtual environments • Working with files, JSON and REST APIs; async calls with asyncio • pandas for preparing documents, tables and evaluation data • Pydantic models for structured LLM output
02	**SQL and data for AI apps**	SQL essentials: joins, GROUP BY, CTEs and window functions • PostgreSQL with pgvector for storing embeddings next to your data • Text-to-SQL: letting an LLM answer questions from a database safely • Logging questions, answers and feedback for evaluation
03	**LLM fundamentals**	Tokens, context windows and how LLMs use retrieved context • Prompting for grounded, cited answers • Calling LLM APIs: streaming, tool calling and retries • Cost, latency and model selection
04	**Ingestion, chunking and embeddings**	Loading PDFs, web pages and tables cleanly • Chunking strategies and chunk-size trade-offs • Choosing and comparing embedding models • Metadata, filtering and keeping the index up to date
05	**Vector databases and retrieval**	Vector indexes and similarity search • Hybrid search: keyword plus vector • Re-ranking retrieved chunks with cross-encoders • Query rewriting and multi-query retrieval
06	**Orchestration with LangChain and LlamaIndex**	RAG chains, retrievers and query engines • Conversation memory and source citations • Tool calling and agentic RAG, and when a plain pipeline is better • Streaming and caching for user-facing latency
07	**Evaluation, guardrails and deployment**	Building an evaluation set from real questions • Measuring retrieval recall, faithfulness and answer relevance • Guardrails: prompt injection, sensitive data and refusals • Deploying with FastAPI, monitoring and cost tracking

Projects you build

Company-style work you can explain in interviews

PROJECT 1

RAG over real documents

Ingest a messy document set, then tune chunking, embeddings and re-ranking, measuring how each change moves answer quality.

PROJECT 2

RAG evaluation harness

Build a test set and scoring pipeline for retrieval recall and faithfulness that shows whether a change made answers better or worse.

PROJECT 3

Production RAG service

Ship a cited-answer assistant behind a FastAPI endpoint with guardrails, streaming, caching and cost tracking.

WHO IT'S FOR

Built for people who already work.

3-month live programme for working professionals (IT and non-IT) and people returning after a career gap.

Working IT professionals

Developers and data engineers who want to build RAG and LLM systems rather than read about them.

Non-IT professionals

Analysts and domain experts with 2+ years of experience who want to move into applied AI work.

Returning after a career gap

Professionals coming back after a break who want the most in-demand skill set and help explaining the gap.

ELIGIBILITY 2+ years of work experience, IT or non-IT. Career breaks welcome.

How it works

From your first call to your offer letter

1	**Profile assessment**	We review your experience and goals and confirm this track fits you.
2	**Live training**	Instructor-led batches on weekday evenings and weekends, online or in our Bengaluru classroom.
3	**Projects and 1:1 mentoring**	Build the projects above with a mentor who reviews your work and unblocks you.
4	**Resume and LinkedIn rebuild**	Rewritten around your projects and the GenAI / RAG Engineer role you are targeting.
5	**Unlimited mock interviews**	Technical and HR rounds with feedback after each one, until you are ready.
6	**Referrals and placement support**	Referrals to hiring partners, interview scheduling and offer negotiation help, until you are hired.

Included in your programme

Live batches
Weekday evenings and weekends

Online or classroom
Join live online or in HSR Layout, Bengaluru

1:1 mentoring
A mentor who knows your goal

Unlimited mock interviews
Technical and HR rounds

Resume + LinkedIn rebuild
Written for your target role

Referrals to hiring partners
Placement support until you're hired

OUTCOMES AND FEES

Most of your fee is paid only after you're placed.

15+	100+	80%	₹14L–₹35L
Learners placed every month	Hiring companies: MNCs, startups, product firms	Average salary hike	Target for GenAI / RAG Engineers

Based on our placements to date. Target range for working professionals. Offers depend on your experience, skills and interviews; ranges are targets, not guarantees.

Choose your plan

All fees include GST

Training + placement	₹3,50,000
Learn the stack, then get placed	

Placement only	₹3,20,000
Already skilled? Skip training	

1

AT ENROLMENT

₹30,000

Training fee

Live batches, projects and 1:1 mentoring.

Placement only: ₹0

2

AT ENROLMENT

₹50,000

Placement token

Starts your resume rebuild, mock interviews and referrals.

Refundable

3

AFTER YOU'RE PLACED

₹2,70,000

Placement fee balance

Due only once you're placed.

Only when placed

Not placed within 3 months? We return your ₹50,000 token.

If we don't place you within **3 months of starting your resume**, your full **₹50,000 token is refunded**. The ₹2,70,000 balance is due only if you're placed. Full terms: stintacademy.com/fees/

You're placed

	With training	Placement only
At enrolment	₹80,000	₹50,000
After placement	₹2,70,000	₹2,70,000
Total you pay	**₹3,50,000**	**₹3,20,000**

Not placed within 3 months

	With training	Placement only
At enrolment	₹80,000	₹50,000
Token refunded	- ₹50,000	- ₹50,000
Total you pay	**₹30,000**	**₹0**

Book your free career assessment

Tell us your experience and goal. We'll confirm the fit, share the next batch dates and answer fee questions.

+91 81974 25757 · Phone and WhatsApp

enquiry@stintacademy.com

stintacademy.com/courses/generative-ai/

1st Floor, #159, SV Arcade, 9th Main Rd, Sector 6, HSR Layout, Bengaluru 560102

Scan for the course page