RAG Engineering on Enterprise Data

Chunking, embeddings, vector databases, hybrid search and grounded answers.

DURATION	FORMAT	LEVEL	LABS
32 hours	Onsite, online or hybrid	Intermediate	Cloud sandbox included

About this programme

Build retrieval-augmented generation (RAG) assistants that answer from your own documents and data — accurately, with citations, and with evaluation you can trust before rollout.

Who it is for: Software engineers, data engineers and ML engineers

What your team will be able to do

- Design an ingestion pipeline for PDFs, wikis and tickets with sound chunking
- Choose embedding models and a vector database for your scale and budget
- Implement hybrid search with re-ranking and metadata filters
- Produce grounded, cited answers and measure faithfulness and recall
- Secure retrieval with document-level access control

Before you start

- Working Python: functions, packages, virtual environments
- Calling a REST API and reading JSON
- Prompt engineering basics, or our Prompt Engineering programme

Course outline

We send the module-by-module outline with your proposal, adapted to your team's tools and goals.

Assessment and certification

How learning is checked

- A short knowledge check on day one and again at the end, so you can see each person's progress
- Graded hands-on labs in every module, reviewed by the trainer
- A final exercise that applies the programme to a problem like your team's own

Certificate and exams

- Participants who complete the programme receive a Cognitive Networks AI certificate of completion.

Formats and cohort size

- Onsite: At your office or a venue you choose, in full days.
- Online, live: Instructor-led over video, usually in half-day sessions so people can keep up with work.
- Hybrid: Some days in a room together, the rest online. Useful when a team works across cities.

Cohort size: 15 to 25 people per trainer. Larger teams run as parallel cohorts on the same outline, so everyone learns the same thing.

Pricing: We quote per cohort, based on duration, format and number of cohorts. There is no per-seat price list: tell us your team and we send a proposal within one business day.

Bring this programme to your team

Request a proposal at cognitivenetworks.ai/contact?programme=rag-engineering

reachus@cognitivenetworks.ai · Bangalore +91 72043 05303 · Dallas +1 (303) 827-7415