AMANAI LAB • DIWALI SPECIAL COHORT • 2026

Diwali AI Accelerator

Learn it. Build it. Get interview-ready.

Generative AI • RAG • Agentic AI • MCP • A2A • Fine-tuning • LLMOps • System Design

3	8	1:1	2	19+
SEATS ONLY	WEEKS	MENTORING	REAL PROJECTS	FREE TOOLS

FREE BONUS ACCESS INCLUDED

- 10 free PDF notes & cheat-sheets (RAG, Agentic AI, LLMOps, System Design...)
- Free System Design practice: 18 problems with AI feedback
- 19+ free AI tools: code lab, resume analyzer, question bank and more

DIWALI COHORT PRICE — **Only 3 seats**

₹49,999 per learner

Regular 1:1 mentorship: ₹1,05,000

Apply online and reserve your seat with a small deposit.

Mentor: Aman Chauhan, Founder, AmanAI Lab

amanailab.com/diwali-accelerator

WHY THIS PROGRAM

Built for one outcome: you clear the interview

Most people learning AI collect certificates and still can't build, evaluate or explain a working system when a hiring panel asks. Companies now hire for **applied** skills: RAG that is measured, agents that use tools reliably, systems that are safe, observable and affordable — and engineers who can reason about trade-offs out loud.

This cohort has **3 learners**. Aman works with each of you directly: reviewing your code, challenging your design decisions, and discussing the interview and system design angles of every topic each week.

Three tracks, running together

Track	What it covers
1. Learn	Generative AI, RAG, Agentic AI, MCP, A2A, fine-tuning, serving, LLMOps, safety and evaluation — depth, not just demos
2. Build	Two end-to-end projects built week by week, with an architecture review and a code review each
3. Interview	Interview discussion every week, GenAI system design, resume and LinkedIn polish

By Week 8 you will be able to

- Design and build a **RAG system** with hybrid retrieval, reranking, citations and measurable evaluation
- Build **tool-using agents** in LangGraph with state, memory, retries and human-in-the-loop
- Orchestrate **multi-agent** systems and connect them with **MCP** and **A2A**
- Decide between **prompting, RAG and fine-tuning**, and explain LoRA / QLoRA, quantization and serving
- Add **guardrails, evaluation, tracing and cost control**; defend against prompt injection
- Lead a **GenAI system design** interview from requirements to trade-offs
- Answer the most-asked LLM, RAG, agent and MLOps questions **clearly and confidently**
- Show **two documented, deployed projects** on GitHub and LinkedIn

Who this is for

Profile	Why it fits
Software engineers / developers	Move into GenAI and agent roles using skills you already have
Data / ML professionals	Go from notebooks and models to LLM applications and agents
Final-year students & freshers	Build a portfolio that stands out from certificate-only candidates
Career switchers with Python basics	Structured, mentor-led path instead of scattered tutorials

Prerequisites: comfortable with Python basics (functions, classes, pip, APIs) and about 8–10 hours a week. No prior AI/ML experience is required — Week 0 fills the gaps.

PROGRAM OVERVIEW

How the 8 weeks work

Item	Details
Duration	Week 0 onboarding + 8 weeks, about 8–10 hours of effort per week
Cohort size	Maximum 3 learners
Mentoring	1 live 1:1 session each week (60–90 min) + chat support for doubts
Projects	2 end-to-end, documented, deployable projects, built incrementally
Interview prep	Weekly interview and system design discussion, career sprint in Week 8
Mentor	Aman Chauhan, AmanAI Lab
Diwali price	₹49,999 per learner (regular 1:1: ₹1,05,000)

Your weekly rhythm

When	What happens
Start of week	Concept briefing, learning material and the week's lab + project milestone
Mid-week	Build time. Ask doubts on chat and get unblocked quickly
Live 1:1 session	Code review, architecture discussion, then an interview and system design discussion on that week's topics
End of week	Deliverable pushed to GitHub; written feedback from Aman

Week 0 — Onboarding (before the cohort starts)

- Diagnostic: a short quiz + call to find your gaps and personalise your plan
- Python refresher for AI work: async, typing, Pydantic, virtualenvs, packaging
- Git / GitHub workflow, project template, API keys and a low-cost compute setup
- ML essentials crash-course: embeddings intuition, loss, overfitting, train/val/test, metrics

Tech stack you will work with

Area	Tools & concepts
LLMs & prompting	OpenAI / Anthropic / open-source models, structured outputs, function calling, prompt versioning
RAG	Embeddings, FAISS / Chroma / pgvector, BM25 + hybrid search, cross-encoder rerankers, RAGAS-style evals
Agents	LangChain, **LangGraph** (state, checkpoints, streaming), tool calling, memory
Protocols	**MCP** (servers, tools, resources) and **A2A** (agent cards, tasks, handoffs)
Fine-tuning & serving	Hugging Face, PEFT / LoRA / QLoRA, quantization, vLLM, semantic caching, model routing
Backend & deploy	Python, FastAPI, Docker, CI/CD, secrets handling, cloud deployment
Quality & safety	LLM-as-judge, golden sets, guardrails, tracing (LangSmith / Langfuse), cost tracking

CURRICULUM MAP

The 8 weeks at a glance

Wk	Learn	Project milestone	Interview + system design
0	Onboarding, Python for AI, ML essentials	Repo + environment ready	Diagnostic assessment
1	LLM internals, prompting, structured outputs	Project 1 scoped & designed	LLM fundamentals
2	Embeddings, vector DBs, chunking, RAG	P1: ingestion + retrieval working	RAG basics, vector search
3	Advanced RAG: hybrid, rerank, query rewriting, evals	P1: evaluation harness	RAG improvement & evaluation
4	Agentic AI, LangGraph, tool calling, memory	P1: agentic RAG layer. P2 scoped	Agents, tool design
5	Multi-agent, guardrails, security, tracing	P2: supervisor + specialists	Safety, prompt injection
6	MCP and A2A protocols	P2: MCP tools + A2A handoffs	MCP / A2A, protocol design
7	Fine-tuning, serving, LLMOps, cost	P1 deployed. P2 hardened	Fine-tuning, MLOps, cost
8	GenAI system design + career sprint	P2 demo day	System design, behavioural, offers

Every week ends the same way: a working deliverable, a code review, and a discussion of the interview and system design angles of that week's topics.

DETAILED CURRICULUM

Week-by-week syllabus

WEEK 1 LLM Foundations & Prompt Engineering

CORE TOPICS
- Tokenization (BPE), embeddings, context window
- Transformer and self-attention, KV cache
- Sampling: temperature, top-k, top-p
- Prompt patterns: few-shot, CoT, role and format control
- Structured outputs and function calling

ADVANCED DEPTH
- Pretraining vs SFT vs RLHF vs DPO
- Why LLMs hallucinate and how to reduce it
- Prompt versioning and regression tests
- Async Python, Pydantic, FastAPI endpoint for an LLM

HANDS-ON LAB
Build a text-to-JSON extractor with validation, retries and a small test set.

PROJECT MILESTONE
Project 1 scoped: use case, data, success metrics.

YOU WALK AWAY WITH
A working LLM utility, project template and your Project 1 design doc.

INTERVIEW DISCUSSION
How LLMs work: attention, sampling, hallucination, prompting vs fine-tuning.

SYSTEM DESIGN DISCUSSION
A prompt-to-JSON service with validation, retries and versioned prompts.

WEEK 2 Embeddings, Vector Search & RAG Fundamentals

CORE TOPICS
- Embedding models, cosine vs dot product
- Vector DBs: FAISS, Chroma, pgvector
- Chunking: fixed, recursive, semantic
- Ingestion pipeline, metadata, citations

ADVANCED DEPTH
- ANN indexes: HNSW and IVF trade-offs
- Parsing PDFs, tables and scanned docs
- RAG vs fine-tuning vs long context
- Embedding model choice and drift

HANDS-ON LAB
Build a RAG app over a real document set, with source citations.

PROJECT MILESTONE
Project 1: ingestion and retrieval working end to end.

YOU WALK AWAY WITH
A running RAG application.

INTERVIEW DISCUSSION
Embeddings, chunking, vector search trade-offs, RAG vs fine-tuning.

SYSTEM DESIGN DISCUSSION
A document Q&A; system: ingestion, index, retrieval and answer generation.

WEEK 3 Advanced RAG & Evaluation

CORE TOPICS
- Hybrid search: BM25 + dense, reciprocal rank fusion
- Cross-encoder rerankers
- Query rewriting, multi-query, HyDE
- Metrics: recall@k, MRR, faithfulness, answer relevance

ADVANCED DEPTH
- Parent-document and contextual retrieval, GraphRAG overview
- LLM-as-judge: biases and calibration
- Golden sets and synthetic test data
- Lost-in-the-middle, context compression, caching

HANDS-ON LAB
Create an evaluation set; improve the pipeline with measured before/after results.

PROJECT MILESTONE
Project 1: evaluation harness and baseline report.

YOU WALK AWAY WITH
An evaluated RAG pipeline with a results table.

INTERVIEW DISCUSSION
Improving retrieval, hybrid search, reranking, evaluating RAG, latency and cost.

SYSTEM DESIGN DISCUSSION
Scaling RAG to millions of documents: freshness, access control, caching.

WEEK 4 Agentic AI with LangGraph

CORE TOPICS
- Agent vs workflow, ReAct, plan-and-execute, reflection
- LangGraph state, nodes, edges, conditional routing
- Tool design: schemas, errors, retries, timeouts
- Short- and long-term memory

ADVANCED DEPTH
- Checkpointing, persistence, streaming
- Human-in-the-loop approvals
- Preventing loops and runaway cost
- Agentic RAG: retrieval as a tool

HANDS-ON LAB
Build a single agent with 3+ tools (search, calculator, your RAG) and persistent state.

PROJECT MILESTONE
Project 1: agentic RAG layer. Project 2 scoped.

YOU WALK AWAY WITH
A reliable single-agent workflow you can demo.

INTERVIEW DISCUSSION
Agent patterns (ReAct, plan-and-execute), tool design, loops, memory.

SYSTEM DESIGN DISCUSSION
An agent service with state, retries, human approval and cost limits.

WEEK 5 Multi-Agent Systems, Guardrails & Security

CORE TOPICS
- Supervisor, hierarchical and swarm patterns
- Specialist agents, shared state, handoffs
- Input / output guardrails and PII handling
- Tracing with LangSmith / Langfuse

ADVANCED DEPTH
- Prompt injection, jailbreaks, data exfiltration, OWASP LLM Top 10
- Evaluating agent trajectories and tool-call accuracy
- Framework trade-offs: LangGraph vs CrewAI vs AutoGen
- Sandboxing and least-privilege tool access

HANDS-ON LAB
Turn your agent into a supervisor + 2–3 specialists with guardrails and tracing.

PROJECT MILESTONE
Project 2: supervisor + specialist agents running.

YOU WALK AWAY WITH
A multi-agent application with safety checks.

INTERVIEW DISCUSSION
Single vs multi-agent, prompt injection, guardrails, evaluating agents.

SYSTEM DESIGN DISCUSSION
A secure multi-agent platform with tracing and least-privilege tools.

WEEK 6 MCP & A2A: Connecting Tools and Agents

CORE TOPICS
- MCP architecture: host, client, server
- Primitives: tools, resources, prompts
- Transports: stdio and streamable HTTP
- A2A: agent cards, tasks, messages, artifacts

ADVANCED DEPTH
- MCP security: auth, tool poisoning, consent
- A2A discovery, streaming and long-running tasks
- MCP vs A2A vs function calling vs REST
- Designing a tool catalogue for an organisation

HANDS-ON LAB
Expose your RAG / data as an MCP server; make two agents collaborate over A2A.

PROJECT MILESTONE
Project 2: MCP tools and A2A handoffs wired in.

YOU WALK AWAY WITH
An MCP-integrated workflow plus an A2A handoff demo.

INTERVIEW DISCUSSION
MCP vs A2A vs function calling, protocol design and security.

SYSTEM DESIGN DISCUSSION
An organisation-wide tool catalogue and agent interoperability layer.

WEEK 7 Fine-tuning, Serving & LLMOps

CORE TOPICS
- When to fine-tune; SFT data preparation
- LoRA, QLoRA and PEFT
- Quantization (INT8 / 4-bit) and vLLM serving
- Prompt/CI-CD, monitoring, cost dashboards

ADVANCED DEPTH
- DPO and preference tuning overview
- Batching, speculative decoding, KV-cache management
- Semantic caching and model routing
- A/B testing, drift detection, rollbacks

HANDS-ON LAB
Fine-tune a small model with LoRA on your data; benchmark against prompting.

PROJECT MILESTONE
Project 1 deployed (Docker + cloud). Project 2 hardened.

YOU WALK AWAY WITH
A fine-tuning notebook, serving benchmark and cost sheet.

INTERVIEW DISCUSSION
LoRA / QLoRA, quantization, serving, monitoring, cutting inference cost.

SYSTEM DESIGN DISCUSSION
An LLM gateway: routing, caching, rate limits, fallbacks, observability.

WEEK 8 GenAI System Design & Career Sprint

CORE TOPICS
- System design framework for GenAI
- Back-of-envelope: tokens, QPS, cost, latency
- Reliability: retries, fallbacks, rate limits
- Resume, LinkedIn, portfolio and demo day

ADVANCED DEPTH
- 8 case studies (see System Design section)
- Behavioural stories and project deep-dives
- Salary negotiation and offer evaluation
- Targeting roles and referrals

HANDS-ON LAB
Design and present 3 full systems under interview conditions.

PROJECT MILESTONE
Project 2 demo day with Aman.

YOU WALK AWAY WITH
Project 2 demo and a polished portfolio, resume and LinkedIn.

INTERVIEW DISCUSSION
Project deep-dives, behavioural stories, resume, offers and negotiation.

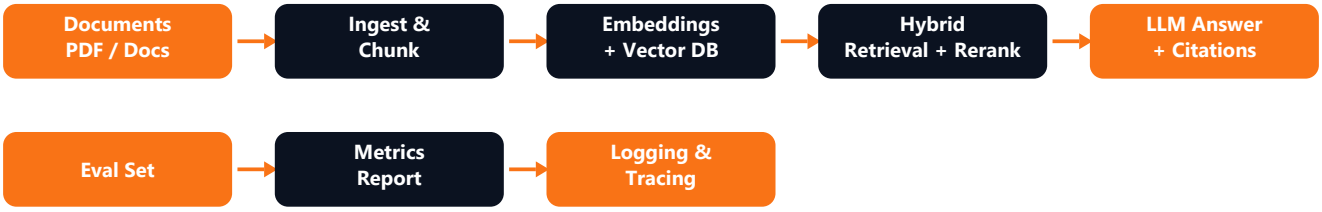
SYSTEM DESIGN DISCUSSION
Full walk-throughs of the 8 case studies, start to finish.

PORTFOLIO PROJECTS

Two systems you can show in any interview

Project 1 — Enterprise Knowledge Assistant

A question-answering system over company documents that gives cited, verifiable answers — one of the most requested GenAI use cases in industry.



You will demonstrate: ingestion design, retrieval quality, citation grounding, evaluation with numbers, deployment and cost awareness.

Project 2 — Multi-Agent AI Assistant

A supervisor agent that plans work and delegates to specialists that use tools through MCP and hand off tasks through A2A, with guardrails and a human approval step.



You will demonstrate: orchestration, tool design, protocol knowledge (MCP + A2A), safety thinking and observability.

Review	What Aman checks
Architecture review (before build)	Requirements, components, failure modes, scale, cost
Code review (before ship)	Structure, tests, error handling, security, README quality
Demo day (Week 8)	A 10-minute walkthrough you can reuse in interviews and on LinkedIn

INTERVIEW READINESS

What we discuss beyond the syllabus

Round	What we prepare	How
Technical concepts	LLMs, RAG, agents, MCP/A2A, fine-tuning, MLOps	Weekly discussion in your 1:1 session
ML / AI coding	Implement attention, cosine similarity, top-k, BM25, a RAG pipeline, a tool-calling loop	Practice sessions in Python
GenAI system design	8 case studies, framework, estimation, trade-offs	Whiteboard-style, Week 8 + earlier practice
Behavioural	STAR stories, project deep-dives, conflict, failure, ownership	Story bank built with Aman
Project deep-dive	Defending every decision in your two projects	Panel-style questioning
Offer stage	Salary negotiation, comparing offers, first 90 days	Playbook + role-play

Practice between sessions

- Use the free **AI Mock Interview**, **Question Bank** (500+ questions) and **Code Lab** on amanailab.com

Resume, LinkedIn & portfolio

- Line-by-line resume rewrite focused on measurable AI-engineering impact
- LinkedIn headline, About and project posts optimised for recruiter search
- GitHub portfolio: README, architecture diagram, results table and demo video
- Role targeting and referral approach for companies hiring GenAI engineers

> **Honest note:** no one can guarantee a job. This program guarantees a mentor's full attention, two real projects and serious interview preparation. The outcome depends on your effort and the market.

GENAI SYSTEM DESIGN

The round that decides senior offers

A 7-step framework

#	Step	What to cover
1	Clarify requirements	Users, use case, accuracy needs, latency, compliance, scale
2	Estimate	Queries/day, tokens per request, storage, cost per query
3	High-level design	Ingestion, retrieval, orchestration, model layer, API, UI
4	Deep dive	Chunking, indexes, prompts, agents, tool design, caching
5	Quality & safety	Evaluation, guardrails, human review, PII, injection defences
6	Scale & reliability	Rate limits, fallbacks, async queues, observability, cost control
7	Trade-offs	Build vs buy, model size, latency vs quality, RAG vs fine-tuning

8 case studies we design together

Case study	Key challenges
Enterprise RAG chatbot for 1M documents	Ingestion at scale, access control, freshness, evaluation
Customer-support copilot	Escalation to humans, tone, tool use, safety, deflection metrics
AI coding assistant	Repo indexing, context selection, latency, code-execution safety
Document processing pipeline (invoices, contracts)	OCR, extraction schemas, validation, human-in-the-loop
Multi-agent research assistant	Planning, parallel agents, citations, cost and loop control
Semantic search for e-commerce	Hybrid ranking, personalisation, latency under load, A/B tests
LLM gateway / platform	Routing, caching, rate limits, quotas, observability, fallbacks
Real-time voice agent	Streaming, latency budget, interruption handling, tool calls

Practice beyond the cohort: amanailab.com/system-design has 18 free system design problems with AI feedback, so you can rehearse between sessions.

FREE BONUS ACCESS

Your complete interview toolkit — free on amanailab.com

These resources are free to use on AmanAI Lab with a free account (optional paid plans only raise daily limits on the heavier AI tools). Inside the cohort, Aman shows you exactly **which one to use each week**.

1. Free PDF notes & cheat-sheets (10 guides + this brochure)

• LLM & GenAI Cheat Sheet (Wk 1–2)	• Prompt Engineering Cheat Sheet (Wk 1)
• Transformer Architecture Deep Dive (Wk 1)	• Python for AI / ML Cheat Sheet (Wk 0)
• RAG Complete Guide (Wk 2–3)	• Agentic AI Guide (Wk 4–6)
• Fine-Tuning Guide 2026 (Wk 7)	• MLOps & LLMOps Guide (Wk 7)
• LLM System Design Guide (Wk 8)	• GenAI & Agentic AI Interview Masterclass Syllabus

2. Free System Design practice

18 GenAI & ML system design problems at amanailab.com/system-design with company tags, difficulty levels and AI feedback on your design.

3. 19+ free AI tools

AI Mock Interview Scored answers across 14 topics	**AI/ML Code Lab** 45 coding problems with a built-in editor
Question Bank 500+ real interview questions	**Skill Gap Analyzer** Find what to learn next
AI Resume Analyzer Score and fix your resume	**Cover Letter Review** Score, keywords and rewrite
LinkedIn Optimizer AI-rewritten profile sections	**JD → Questions** Paste a JD, get tailored questions
Company Prep Google, Meta, OpenAI and more	**Career Roadmap** Week-by-week AI/ML path
Flashcards 5-minute daily revision	**Topic Guides** Deep-dive prep by topic
Skill Quiz AI-generated MCQs on 14 topics	**Paper Explainer** Any arXiv paper in plain English
Code Playground Monaco editor + AI for ML code	**System Design Practice** 18 problems with AI feedback
Job Tracker Track applications and stages	**Daily Challenge** Build a daily practice habit
AI News Stay current for interviews	

INVESTMENT

What’s included and how to join

Included in the Diwali cohort
• 8 weeks (+ Week 0 onboarding) of live 1:1 mentoring with Aman
• Personalized learning roadmap and diagnostic
• Full syllabus: GenAI, RAG, Agentic AI, MCP, A2A, fine-tuning, LLMOps, safety
• 2 end-to-end projects with architecture + code reviews
• Weekly interview discussion tied to each topic
• GenAI system design framework + 8 case studies
• Resume, LinkedIn and portfolio optimisation
• Chat support between sessions
• Free bonus toolkit: 10 PDF notes, 18 system design problems, 19+ tools
• Lifetime access to project resources
• AmanAI Lab mentorship certificate

Diwali cohort fee	₹49,999 per learner
Seat confirmation	₹15,000 online to reserve your seat right after applying (UPI / card)
Balance	Remaining amount before Week 1 (or pay in full)
Seats	3 learners only. When full, applications close

Frequently asked questions

I’m a beginner. Is this for me?
If you know basic Python, yes. Week 0 and Week 1 build the foundations.

Is it recorded or live?
Live 1:1 sessions with Aman each week, plus supporting material.

What if I miss a session?
Sessions are rescheduled within the week based on your availability.

Do I need to pay for APIs or GPUs?
Small API usage is needed. We guide you to keep costs low, with free / open-source options and free GPU tiers for the fine-tuning lab.

Is a job guaranteed?
No. We guarantee mentoring, projects and preparation. Results depend on effort and market.

APPLY NOW

Four simple steps

1	**Apply** Fill the short application form at amanailab.com/diwali-accelerator
2	**Reserve** Pay ₹15,000 online (UPI / card) to confirm your seat, or pay the full fee
3	**Fit call** A short WhatsApp / video call with Aman to plan your roadmap
4	**Start** Receive your diagnostic, personalised roadmap and first-session date

Timeline

Applications open	Now
Applications close	12 November 2026 (or earlier if all 3 seats are filled)
Cohort start	Around 15 November 2026, confirmed on your fit call

Only 3 seats. Apply today.

amanailab.com/diwali-accelerator

About your mentor

Aman Chauhan is the founder of AmanAI Lab, a platform helping engineers learn Generative AI, Agentic AI and interview preparation through practical, project-first learning.