



COHORT 1 • STARTS 16 AUGUST 2026 • 40 SEATS

Ship 33 Real AI Builds. Graduate a **Production** AI Engineer.

*For developers with 2–6 years of experience who already ship code.
Seventeen weeks, two real projects a week, every graded build
reviewed by an engineer — not a certificate mill.*

Week 1 is **free** — and it's part of the course, not a preview of it.

India: **+91 97168 01894**

www.coresmart.ai

Most Courses Teach the Demo. We Teach the **Job**.

Most AI courses stop at "it runs on my laptop." The hard part — the part teams actually struggle with — is production engineering: evaluation, agents, guardrails, deployment, reliability and cost. This program is built around exactly those skills, and you prove each one by shipping a real, working build.

We're not an ed-tech company that hired some engineers. We're an applied-AI engineering company that decided to teach. Our curriculum is designed by people who have built and run real AI systems in production, reviewed against what the industry actually builds — not academic benchmarks alone.

Built for developers with 2–6 years of experience — you already ship code and work with APIs; this course takes you from "I can call an LLM" to "I can design, evaluate, deploy and operate the AI system behind it."

Our mission: turn learners into practitioners, and practitioners into builders.

Build every week

You watch the taught project built, practise it in a guided lab, then ship your own graded project. Same skills, new input. Not lectures.

Reviewed & scored

Your graded project is reviewed by our team every week, with a scorecard and written feedback — never self-graded.

The full production stack

RAG, agents, fine-tuning, guardrails, deployment, reliability and cost: what production teams actually run day to day.

Graduate with a real portfolio

Finish with up to 33 real builds, including a live capstone you present and defend, all sitting on your GitHub.

Learn. Practise. Build. Get **Reviewed.**

You learn on your own schedule, but you are never grading yourself. Each week you watch the taught project built, practise it in a guided lab, then ship your own graded project — reviewed by our team and returned with a scorecard and written feedback.

EVERY WEEK

1**Learn it**

Watch the taught project built live, line by line.

2**Practise it**

Rebuild it yourself in a guided lab, self-checked.

3**Build your project**

Ship your own graded project on a new input.

4**Reviewed + scored**

Returned with a scorecard and written feedback.

SCORED ON

CORRECTNESS

COMPLETENESS

DESIGN

CLARITY

THE CAPSTONE · GRADUATION

A production system you present and defend

1**Build your capstone**

Combine 17 weeks into one production-grade system.

2**Live demo & presentation**

Present and demo your build, graduation-style.

3**Reviewed & scored**

A panel reviews your project and presentation.


4**Feedback + certificate**

Detailed feedback, your score, your certificate.

Presenting a working system live, and defending your choices, is the closest thing to the real job.

Try the Real Course, Not a **Preview** of It.

Week 1 costs nothing — and it is part of the 17 weeks, not a taster built for a brochure. Same live build session. Same graded project. Same scorecard. No card required, and if it's not for you, you've lost nothing.



WHAT YOU SHIP IN WEEK 1 — MESSY NOTES AND MEETING TRANSCRIPTS, TURNED INTO CLEAN, STRUCTURED OUTPUT, LIVE

What you'll learn

- ✓ Onboarding + fast-track foundations — no prior AI assumed
- ✓ How modern AI products are built, layer by layer
- ✓ LLM internals and how RLHF shapes behaviour
- ✓ The model / retrieval / tool / memory stack
- ✓ Streaming responses and structured output
- ✓ Latency, cost and quality tradeoffs

What you build

BUILT LIVE

ReleaseBot

A service that streams messy release notes into a clean, structured changelog — live, over SSE.

GRADED

MinuteMaker

A service that turns a meeting transcript into structured minutes and action items.

Reserve your seat in Week 1 — free

No payment info needed. Attend live, ship both builds, get real written feedback.

REGISTER NOW

India: +91 97168 01894

info@coresmart.ai

Practitioners, Not Professors.

Every credential below is public and linked. Check them before you read the curriculum.



Vinay Bamil

COURSE DESIGNER · PHD IN AI · EX GEN AI COACH, GOOGLE

Designed the course and teaches most of it — bringing a research-grade understanding of LLM internals to a curriculum built entirely around shipped work.

linkedin.com/in/vinay-bamil



Randeep S. Bhatia

CTO AT SPLASH · EX TWITCH, AUDIBLE & EA · AAAI PUBLISHED

Advises on architecture and scale, drawing on a career shipping systems used by millions at Twitch, Audible and EA.

linkedin.com/in/randeepsbhatia



Sanjay Lalwani

INSTRUCTOR · DATA SCIENTIST AT SIEMENS · EX INFOSYS

Runs production machine learning at Siemens, and brings that day-to-day operational discipline directly into the classroom.

linkedin.com/in/sanjaylalwani3

Curriculum informed by Agama Solutions, a US staffing firm with 20 years of hiring intelligence — so the skills taught are the ones clients interview for now.

What Industry Leaders Say About the Programme.

Two senior technology leaders looked at the curriculum and put their names to it.



Most AI courses produce people who can talk about AI. CoreSmart produces people who can actually ship. The curriculum covers the same production-grade stack used in real enterprise systems, and every week delivers a usable artifact rather than just a certificate.



Ashish Mago

CO-FOUNDER & CTO, COMPLIANCE KART

Ex CTO, LMCE · Ex Director Technology, Publicis Sapient



Teams already have AI tools, but delivery velocity remains unchanged. CoreSmart's dual-track model addresses the real gap by helping Business Analysts specify agent behaviour and Developers govern AI systems at scale.



Jignesh Modi

CXO, GLOBAL HOSPITALS & HEALTHCARE TECHNOLOGY

Digital Transformation & AI Leader

Sixteen Weeks.

Four Modules.

Two real projects every week — a taught build and your own graded build, same skills, new input. Every one lands on your GitHub.

A Foundations & Core — Weeks 1–4

- | | |
|-------|---|
| WK 01 | AI Product Anatomy, LLM Internals & RLHF
ReleaseBot + MinuteMaker — streaming notes and transcripts into structured output. <i>Free.</i> |
| WK 02 | LLM Mechanics, Model Choice & Eval Primer
IntentIQ + FeedbackSorter — an intent classifier, and a feedback triage tool. |
| WK 03 | Prompting, Structured Outputs & Tool Calling
TicketStream + ReviewRouter — classification and routing bots, end to end. |
| WK 04 | Multimodal Ingestion, Embeddings & Vector DBs
KnowledgeVault + PaperFinder — semantic search over documents and papers. |

B RAG, Evaluation & Tuning — Weeks 5–8

- | | |
|-------|--|
| WK 05 | RAG Foundations & Grounded Answers with Citations
CitationRAG + DocuRAG — assistants that cite every claim they make. |
| WK 06 | Advanced RAG & Context Engineering
RAGOptimizer + RerankLab — reranking and context compression, measured. |
| WK 07 | Evaluation Science & Adversarial Testing
BreakRAG + RAGBench — stress-testing and benchmarking on a golden set. |
| WK 08 | Fine-Tuning, Open-Source LLMs, LoRA & SLM
SpecialistTuner + DomainTuner — small models, fine-tuned and benchmarked. |

Agents, Safety, and Production.

This is where a demo becomes a product: multi-agent systems, guardrails, real deployment, reliability and cost.

C Agents & Safety — Weeks 9–12

- WK 09** **Agentic AI Foundations & State Management**
OpsAssist + ResearchAgent — agents that complete multi-step tasks with tools.
- WK 10** **Multi-Agent Orchestration & Long-Term Memory**
TriageFlow + DeskOrchestrator — router-plus-specialist agent pipelines.
- WK 11** **Agent Interoperability: MCP, A2A & AGENTS.md**
AgentMesh + ToolBridge — an interoperable layer and an MCP server.
- WK 12** **Guardrails, PII Detection, Responsible AI & HITL**
GuardianAI + SafeAssist — guarded assistants with a human in the loop.

D Production & Ops — Weeks 13–16

- WK 13** **Developer Tooling Agents & Streaming UI**
WorkbenchAI + PRCopilot — a live agent that reviews pull requests.
- WK 14** **Deployment, Observability, Versioning & Async**
DeployCore + ObserveOps — deployed live, with tracing and metrics.
- WK 15** **AI Application Reliability Engineering**
ReliabilityKit + ChaosProbe — fault injection and a resilience report.
- WK 16** **FinOps, Prompt Caching, SLM Routing & Self-Healing**
CostGuard + CachePilot — caching and routing, with a savings report.

Reliability engineering and cost control are the difference between an AI demo and a system a company can actually run at scale. Very few courses go here — it's exactly what production teams hire for.

All Seventeen Weeks, at a Glance.

For quick reference — every week, its focus, and the two builds that come with it.

01	AI Product Anatomy & RLHF	ReleaseBot · MinuteMaker (<i>free</i>)
02	LLM Mechanics & Eval Primer	IntentIQ · FeedbackSorter
03	Prompting & Tool Calling	TicketStream · ReviewRouter
04	Embeddings & Vector DBs	KnowledgeVault · PaperFinder
05	RAG Foundations & Citations	CitationRAG · DocuRAG
06	Advanced RAG & Context Eng.	RAGOptimizer · RerankLab
07	Evaluation & Adversarial Testing	BreakRAG · RAGBench
08	Fine-Tuning, LoRA & SLM	SpecialistTuner · DomainTuner
09	Agentic AI Foundations	OpsAssist · ResearchAgent
10	Multi-Agent Orchestration	TriageFlow · DeskOrchestrator
11	Interoperability: MCP & A2A	AgentMesh · ToolBridge
12	Guardrails & Responsible AI	GuardianAI · SafeAssist
13	Dev Tooling Agents & UI	WorkbenchAI · PRCopilot
14	Deployment & Observability	DeployCore · ObserveOps
15	Reliability Engineering	ReliabilityKit · ChaosProbe
16	FinOps & Prompt Caching	CostGuard · CachePilot
17	Capstone & Live Demo Day	Your production system, defended live

A Real Production System. Not a Demo.



Week 17 brings the whole course into one production-grade system. Choose a flagship track below — every one ships with its dataset — or bring your own idea and we help you scope it.

C

CaseCompass

Ask 40,000 real court opinions anything: grounded, cited answers over public US case law.

LEGAL

DATA INCLUDED

M

MediGuard

A patient-info assistant with medical-grade guardrails, grounded in official health sources.

HEALTHCARE

DATA INCLUDED

F

FilingScout

An analyst agent that reads SEC filings and writes cited earnings briefs on any company.

FINANCE

DATA INCLUDED

S

StackSage

A developer support copilot trained on real Stack Exchange Q&A and open-source docs.

DEV TOOLS

DATA INCLUDED

T

TravelGenie

An agent that plans complete day-by-day trip itineraries, cited from real travel guides.

CONSUMER

DATA INCLUDED

P

PulseBrief

A scheduled agent that turns the week's world events into one clean, cited briefing.

MEDIA & OPS

DATA INCLUDED

Every capstone ships with its dataset

Each project comes with a cleaned, licensed dataset, so you build from day one. Prefer your own idea? Bring it — we help you source and scope it. Present it live on Demo Day to a panel, graduation-style.

Build It. Present It. Defend It.


Presenting a working system live, and defending your choices to a panel, is the closest thing to the real job.

What you deliver

- ✓ A real production system, deployed and running
- ✓ A live demo and architecture walkthrough
- ✓ 33 real builds, every one inspectable on GitHub
- ✓ A portfolio pack — case study and diagram
- ✓ 16 weeks of scorecards showing your arc

After you finish

- ✓ **Portfolio review.** We go through your GitHub and sharpen how each build reads.
- ✓ **Production-ready skills.** Design, evaluate, deploy and operate real AI systems.
- ✓ **Interview-ready.** Positioning and prep to explain and defend everything you built.



Roles this prepares you for

AI Engineer

GenAI Engineer

LLM Application Engineer

Applied AI Engineer

AI Platform Engineer

AI Solutions Engineer

This is skills and positioning support, not a job guarantee — we make sure you can answer for the build.

The Things People Ask Us **First.**

Q. Do I need coding experience to join?

No prior AI background is required — Week 1 starts with onboarding and fast-track foundations, including a Python and API refresher. Comfort with basic coding helps, but you don't need to know AI or ML going in.

Q. How much time does it take a week?

170+ hours across 17 weeks — roughly 10 a week. Weekday content is recorded, so you fit it around work. Build sessions are live on weekends and recorded, so missing one doesn't put you behind.

Q. Is this self-paced?

Partly. Weekday content is recorded, so you learn on your own schedule. Weekend sessions are live — that's where Project 1 gets built alongside the instructor. You're never grading yourself: every graded project is reviewed and returned with a scorecard.

Q. What do I actually build?

Two real projects every week for 16 weeks, plus a capstone — 33 builds in total, all on your GitHub. They move from streaming services and classifiers, through RAG, evaluation and fine-tuning, then agents, guardrails, deployment, reliability and cost.

Q. Who's teaching this?

Vinay Bamil designed the course and teaches most of it. Randeep S. Bhatia advises on architecture and scale. Sanjay Lalwani runs production machine learning at Siemens. Every credential is public and linked in this brochure.

Q. What happens at the end of the programme?

Week 17 is capstone week — one production-grade system, presented and defended live to a panel, graduation-style. Afterwards we review your portfolio and help you prepare to explain everything you built.

JOIN COHORT 1

Start With Week 1. Free.

Same live build session, same graded project, same scorecard
as the paid weeks — just nothing to lose.

COHORT DETAILS

Cohort 1 · Starts 16 August 2026 · 40 Seats

Recorded weekday content, live weekend build sessions. Reach
out using the details below to reserve your seat in Week 1.

Get in touch

CoreSmartai Inc.

Email: info@coresmart.ai

India: +91 97168 01894

US: +1 (650) 499-6634

Web: www.coresmart.ai

Our offices

India Office

B-4/5, Suite 03, Sector 63
Noida, Uttar Pradesh – 201301

US Office

39159 Paseo Padre Pkwy, Ste 311
Fremont, CA 94538