

InfoQ Certified AI Engineering Program syllabus

5-week online certification cohort

Week-by-week syllabus

Week 1: Becoming an AI-Native Engineering Team

- **Focus:** We define what "AI-native" engineering means for your team today. You will examine product thinking, daily habits, architectural tradeoffs, and where AI adds real value beyond the hype. You will walk away knowing how to spot "AI ambiguity" in Hilary Mason's terms and why Naomi Saphra warns that language models operate via pattern continuation rather than true reasoning.
- **Sample discussion task:** You will debate how to maintain your existing risk management and resilience practices when adopting AI. Your group will assess your team's current position on the maturity scale and define one specific, viable AI engineering opportunity for your organization.
- **Weekly homework:** Map out your current Software Development Life Cycle and propose a strategy to integrate AI tools. You will detail the organizational friction you expect to encounter, identify one boundary where you will deliberately keep AI out, and define a measurable signal to track success.

Week 2: Designing & Building RAG & Context Pipelines

- **Focus:** RAG is the workhorse pattern of enterprise AI, and its quality lives or dies on data, retrieval design, and memory. You will move beyond prompt engineering to design and build production-grade RAG systems and context engineering pipelines, deploy knowledge graphs,

and manage both ephemeral and long-term memory so that enterprise AI stays grounded and fresh rather than confidently wrong.

- **Sample discussion task:** You will identify exactly where traditional RAG falls over in production. Your group will compare vector RAG against Graph RAG across contrasting use cases, defending when each approach wins, and design pipelines that separate ephemeral context from durable knowledge.
- **Weekly homework:** Build a retrieval architecture for a complex query scenario at your organization. You will justify your design choices at each layer, explain how the system handles underlying data changes, and name the single biggest failure mode you anticipate.

Week 3: Designing & Building AI Agents

- **Focus:** You will design agentic systems without turning them into fragile demos. You will leave with a clear breakdown of the spectrum from single-purpose tools to multi-agent systems. We cover orchestration, control, memory, and the tradeoffs between autonomy and certainty.
- **Sample discussion task:** You will critique a real agent architecture for common failure modes like tool errors and runaway loops. Your group will then pressure-test the decision criteria for when a single, well-tooled agent beats a multi-agent orchestration.
- **Weekly homework:** Design an AI agent for a real task at your organization. You will outline the tool inventory, identify expected failure modes, design escape hatches to prevent runaway behavior, and build an evaluation plan to verify the agent actually works.

Week 4: AI Platforms & Infrastructure

- **Focus:** We focus on the platform and infrastructure patterns that make AI systems production-ready. You will review inference gateways and GPU cost controls. You will decide what should be centralized and what should remain federated.
- **Sample discussion task:** You will review a provided system architecture to identify inference waste. Your group will defend the hardware and stack choices needed to reduce token costs, and whiteboard the gateway policies (rate limits, cost caps, PII handling) your organization actually needs.
- **Weekly homework:** Outline an internal AI platform strategy for a mid-sized organization. You will detail workload routing for latency-critical tasks, specify the core observability signals

required for system health, and identify which capabilities should remain federated versus centralized.

Week 5: AI Operational Excellence: Evals, Trust & Reliability

- **Focus:** The final week covers the operational disciplines that turn AI systems into dependable production assets. We examine evaluation metrics, trust, and rollout readiness. Most of this session is devoted to the capstone project presentations.
- **Capstone Project:** The culminating project for the certification. Your group will choose one of two tracks: design a new AI-powered feature for your organization, or outline exactly what it takes to get an existing fragile prototype running in production.
- **Deliverables:** Your presentation must cover the following architectural requirements:
 - Walk through the full stack design, including SDLC boundaries, context retrieval pipelines, and platform policies.
 - Define the failure-containment strategies and escape hatches you built to prevent runaway-agent behavior.
 - Integrate security policies and a rigorous evaluation loop spanning the model, the system, and the user experience.
 - The final deliverable is a 20-minute group presentation followed by a 30-minute peer discussion. You must end your presentation with an unresolved architectural question that your team is still working out.