Agentic AI PM Certification & Anthropic Claude Certification Prep

Course syllabus · Instructor: Mahesh Yadav · Agentic AI Institute

A hands-on certification for PMs and builders who want to ship, not just talk about, agentic AI.

You build a working AI agent in Week 1, either in code with Claude Code or visually with n8n, then in further weeks you will ground it in your own data with RAG and Agentic RAG, put it under real evaluation and red-team discipline, secure it, price it, and deploy it to production across Azure, finishing with a public Pitch Day in front of industry judges and real customers in Week6.

Five weeks, with live sessions every week (all recorded), a hands-on lab every week with mentors to unblock you, and two project tracks: code (Claude Code, APIs) or low-code (n8n, v0).

Throughout the course, every module is mapped to the official Anthropic Claude certifications. Certification prep is our offering: no other AI PM course prepares you for them.

Week 1: Foundations: The AI PM, the Model, and Your First Agent

- **The AI PM & AI problem framing.** What companies now expect from a "builder PM": writing skills, claude.md and agents.md files, and framing problems actually suited to agentic AI — separating hype from shippable value.
- **Agentic AI & ML concepts.** How LLMs actually work — tokens, context windows, function calling, structured outputs — plus supervised/unsupervised/reinforcement learning and deterministic vs. probabilistic product logic.
- **Evaluating and selecting the model.** A model-selection framework for cost, latency and quality: proprietary vs. open source, token economics, time-to-first-token, and when a small language model beats a frontier model.

Hands-on lab

Build your first AI agent in n8n, understand Claude's models, context and capabilities, a Google Colab lab on Claude models, prompting and APIs, and create your Anthropic API key.

1. Build Your First AI Agent in n8n
2. Understanding Claude: Models, Context and Capabilities
3. Google Colab Lab: Claude Models, Context, Prompting and APIs
4. Create your Anthropic API key

Anthropic certification alignment: Claude Certified Associate — Foundations (CCAO-F): core concepts and practice sets begin here.

Week 2: Roadmaps, RAG and Context Engineering

- **Product roadmap & prioritization.** AI-specific prioritization tenets, roadmaps, and the AI PRD: technical requirements, data requirements, risks, guardrails and success metrics defined before a line of code.
- **Understanding the AI stack.** Model providers, tools, platforms — and where your product sits in the modern stack.
- **RAG, Agentic RAG, context & knowledge graphs.** Vector databases, chunking, embeddings, retrieval strategies, context engineering beyond prompts — where vanilla RAG fails and Agentic RAG wins.
- **Memory & connectors.** Build custom knowledge connectors for short- and long-term memory.

Hands-on lab:

Ground your agent in your own data. Why we need RAG (context window failure), build your first RAG pipeline, Agentic RAG, and Claude Skills for PMs.

1. Why We Need RAG (Context Window Failure)
2. Build Your First RAG Pipeline
3. Agentic RAG
4. Claude Skills for PMs

Anthropic certification alignment: CCAO-F completes · Claude Certified Developer — Foundations (CCDV-F) prep begins with Claude Code and connector labs.

Week 3: Evals, Error Analysis and Benchmarks

- **KPIs for AI agents.** Success metrics and golden datasets for qualitative quality — evals, not vibe checks.
- **Evaluating GenAI applications & error analysis.** A disciplined failure-analysis process: read traces, build your own failure taxonomy, then measure each failure mode with code checks and LLM-as-judge evaluators — and validate the judges themselves.
- **Automating evaluations for scale.** Auto-evaluation pipelines, regression suites that gate deployment (CI/CD for agents), and continuous eval loops that catch regressions and hallucinations before your users do.
- **Benchmarks & leaderboards.** How to read (and distrust) public benchmarks and leaderboards for AI agents.

Hands-on lab

Building subagents with connectors, evaluating your n8n AI agent with Azure AI Foundry, and the judge model in Azure AI Foundry

1. Building Subagents with Connectors
2. Evaluating Your n8n AI Agent with Azure AI Foundry
3. Judge Model in Azure AI Foundry

Anthropic certification alignment: CCDV-F: evaluation and testing scenarios drawn directly from these labs.

Week 4: Safety, Security and Multi-Agent Systems

- **AI safety & security.** The agentic threat surface: prompt injection, tool misuse, PII leaks and memory poisoning — guardrails, policy controls, human-in-the-loop approval for irreversible actions, and governance frameworks (NIST AI RMF, EU AI Act) your enterprise buyers will ask about.
- **Multi-agent AI systems (MCP, tools).** MCP integrations, Agent-to-Agent patterns, orchestration (manager/worker), state tracking and termination rules that kill task-looping.
- **Skills & connectors in Claude Code.** Build skills and get real work done via connectors in Claude Code — the exact workflow Anthropic's ecosystem is standardizing on.
- **Observability in production.** Trace instrumentation, monitoring dashboards and alerts — observability-first development, so you see failures before customers do.

Hands-on lab:

Ship a multi-agent workflow. Azure AI Foundry, multi-agent workflow in Azure AI Foundry, and a multi-agent lab in n8n with Agentic RAG

1. Azure AI Foundry
2. Multi-Agent Workflow in Azure AI Foundry
3. Multi-Agent Lab in n8n with Agentic RAG

Anthropic certification alignment: CCDV-F completes · Claude Certified Architect — Foundations (CCAR-F) prep begins: production system design with Claude Code and MCP.

Week 5: Business of AI: Pricing, ROI, GTM and Fine-Tuning

- **Pricing AI products for growth.** Pricing models for agentic products, cost profiling, prompt caching, model routing and token-cost control.
- **Build vs. buy & ROI.** Frameworks for choosing your AI strategy — build your own harness vs. leverage an existing one — and finding ROI in AI investments you can defend to a CFO.
- **GTM challenges, secrets & MOAT.** Why AI products stall in distribution, and how to build a durable moat.
- **Training loops & PEFT (LoRA).** Building training loops via agents; when fine-tuning pays off and the trap when it doesn't; running local models.

Hands-on lab

1. Specs-Based Development with Claude

2. Connecting Specs-Based Development with Foundry
3. Fine-Tuning Your Agent in Azure AI Foundry
4. Model Fine-Tuning with LoRA

Anthropic certification alignment: CCAR-F: architecture reviews and multi-cloud scenario drills.

Week 6 - Pitch Day

Pitch Day. Present your deployed agent to judges and real customers, win prizes, and get your first users and feedback. Several alumni have turned Pitch Day into interviews.

Open labs. More labs are released as needed to complete your projects, with the help of mentor support sessions.

Post-Cohort

Alumni community for life. Public project portfolio pages, the alumni Success Page, Friday live sessions (3+ years running), and referrals from 10,000+ alumni.

Back to course page: https://maven.com/mahesh-yadav/genaipm