✦ ARTIFICIAL INTELLIGENCE · ADVANCED

Building AI-Powered APIs & Microservices

Put AI behind a clean API.



56	8	SAM-AIAPI
HOURS	MODULES	CERTIFICATION

✦ INSIDE THIS BROCHURE

Contents

Program overview 03

Who it's for 04

Learning outcomes 05

Program runway 06

Detailed curriculum 07

Tools & platforms 10

Projects & capstone 11

Certification & method 12

Learning path 13

FAQs 14



AT A GLANCE

Level	Advanced
Duration	56 Hours
Modules	8
Certificate	SAM-AIAPI

✦ PROGRAM OVERVIEW

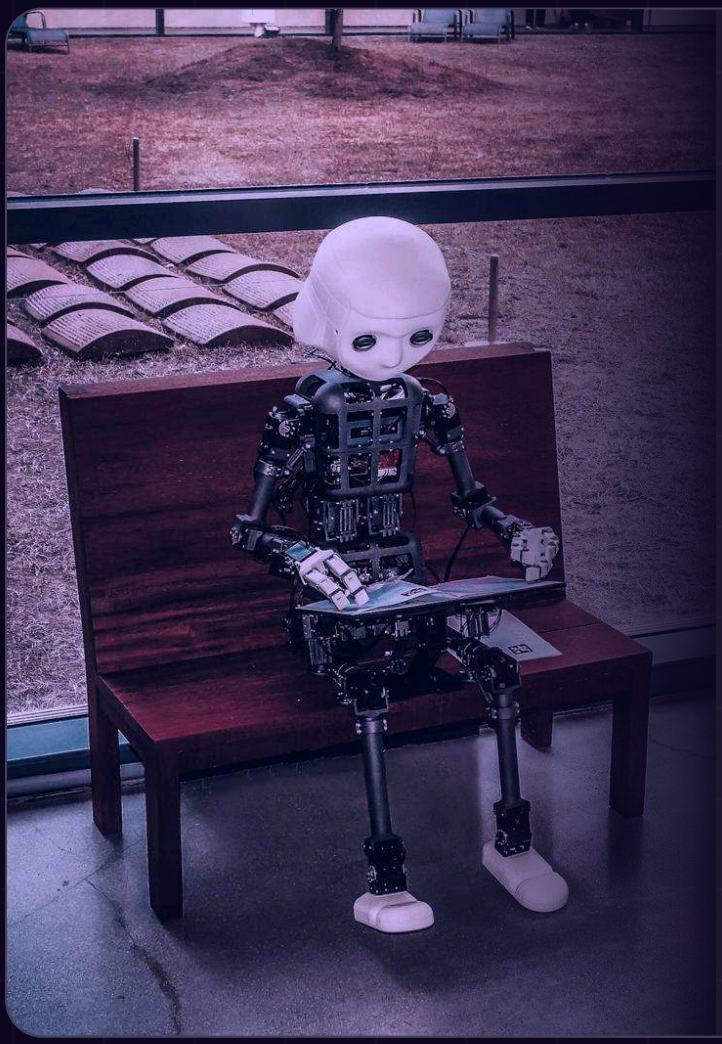
Put AI behind a clean API.

Design, build and deploy APIs and microservices that wrap AI models for real products.



AI features reach users through APIs and services, so they must be fast, secure, scalable and cost-controlled like any other production system. Wrapping models well means handling streaming, timeouts, queues, caching, observability and deployment. This advanced program teaches backend engineers to design, build and operate AI-powered APIs and microservices using FastAPI, Docker, Kubernetes and modern model-serving tools.

LEVEL	DURATION	CERTIFICATION	DELIVERY
Advanced	56 Hours	SAM-AIAPI	Instructor-led · Live virtual · Blended



✦ WHAT YOU'LL PRACTISE

Design AI service APIs	01
Build services with FastAPI	02
Containerise and deploy AI services	03
Handle scale, cost and latency	04

✦ WHO IT'S FOR

Built for people ready to grow.

01

Every professional

AI Fundamentals, ChatGPT, Copilot, Gemini, Claude

02

Developers & data teams

ML, deep learning, NLP, computer vision, MLOps

03

AI engineers

LLM fine-tuning, RAG, vector DBs, AI agents

04

Leaders & public sector

AI for C-suite, government and governance

PREREQUISITES

Strong Python and backend development experience	01
Understanding of REST APIs and HTTP	02
Basic Docker and Linux knowledge	03

ROLES YOU CAN TARGET

- AI Backend Engineer ↗
- Machine Learning Engineer (serving) ↗
- Platform / DevOps Engineer for AI ↗
- API Developer ↗
- AI Solutions Architect ↗

✦ LEARNING OUTCOMES

What you'll be able to do.

01 Design clean, versioned API contracts for AI services

02 Build async FastAPI services with validation and streaming

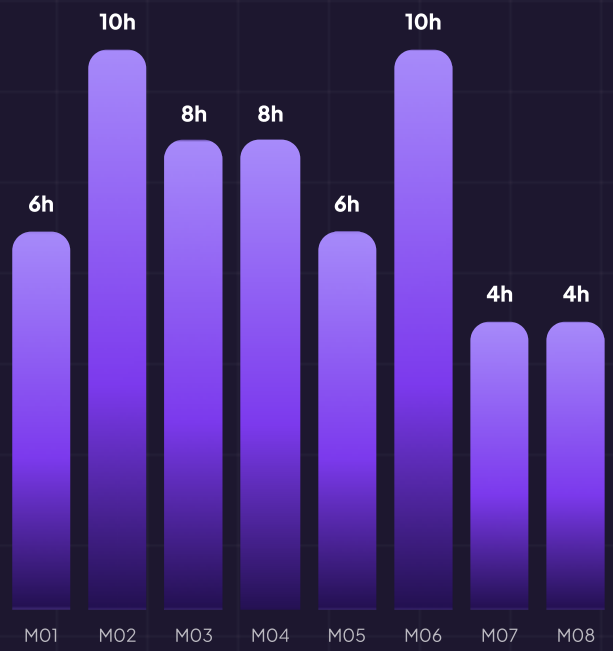
03 Integrate hosted LLMs and self-hosted models behind APIs

04 Implement queues, caching and rate limiting for scale

05 Containerise and deploy services on Kubernetes

06 Monitor latency, cost and errors with observability tools

HOURS PER MODULE



✦ PROGRAM RUNWAY

Your journey, module by module.

8
MODULES · 56 HRS

01	**AI Service Architecture** Monolith versus microservices for AI, Sync, async and batch patterns, API contracts with OpenAPI	6 hrs
02	**FastAPI for AI Services** Async endpoints, Pydantic models and validation, Dependency injection	10 hrs
03	**Integrating Models** Hosted LLM API clients, Self-hosted models with vLLM, Ollama for local inference	8 hrs
04	**Background Jobs and Data** Task queues with Celery or RQ, Redis for caching and queues, PostgreSQL and pgvector	8 hrs
05	**Security and Reliability** Authentication with OAuth2 and JWT, API keys and rate limiting, Input validation and prompt safety	6 hrs
06	**Containers and Kubernetes** Dockerfiles for AI services, Docker Compose stacks, Kubernetes deployments and services	10 hrs
07	**Observability and Cost Control** Structured logging, OpenTelemetry tracing, Prometheus and Grafana	4 hrs
08	**CI/CD and API Lifecycle** GitHub Actions pipelines, Contract and load testing, Versioning and deprecation	4 hrs
★	**Capstone · Scalable AI Microservices Platform** Assessment, capstone review and certification	

✦ DETAILED CURRICULUM

The syllabus.

Modules 01–03

MODULE 01

AI Service Architecture

6 hours

Hands-on lab
Design an architecture diagram and API contract for an AI service.

TOPICS COVERED

- Monolith versus microservices for AI
- Sync, async and batch patterns
- API contracts with OpenAPI
- Model gateway patterns
- Latency and cost budgets

MODULE 02

FastAPI for AI Services

10 hours

Hands-on lab
Build a FastAPI service that streams LLM responses with validation.

TOPICS COVERED

- Async endpoints
- Pydantic models and validation
- Dependency injection
- Streaming responses with SSE
- WebSockets basics
- Testing with pytest

FastAPI

MODULE 03

Integrating Models

8 hours

Hands-on lab
Serve one hosted and one local model behind a unified API.

TOPICS COVERED

- Hosted LLM API clients
- Self-hosted models with vLLM
- Ollama for local inference
- Classic ML models with ONNX
- Provider fallbacks and LiteLLM

Ollama

LEARN BY DOING

8 guided labs, one in every module.



✦ DETAILED CURRICULUM

Syllabus, continued.

Modules 04–06

MODULE 04

Background Jobs and Data

8 hours

Hands-on lab
Add a Redis-backed queue for long-running document summarisation jobs.

TOPICS COVERED

- Task queues with Celery or RQ
- Redis for caching and queues
- PostgreSQL and pgvector
- Long-running job status APIs
- Idempotency and retries

PostgreSQL · Redis

MODULE 05

Security and Reliability

6 hours

Hands-on lab
Secure an AI API with JWT auth and per-user rate limits.

TOPICS COVERED

- Authentication with OAuth2 and JWT
- API keys and rate limiting
- Input validation and prompt safety
- Timeouts and circuit breakers
- Secrets management

MODULE 06

Containers and Kubernetes

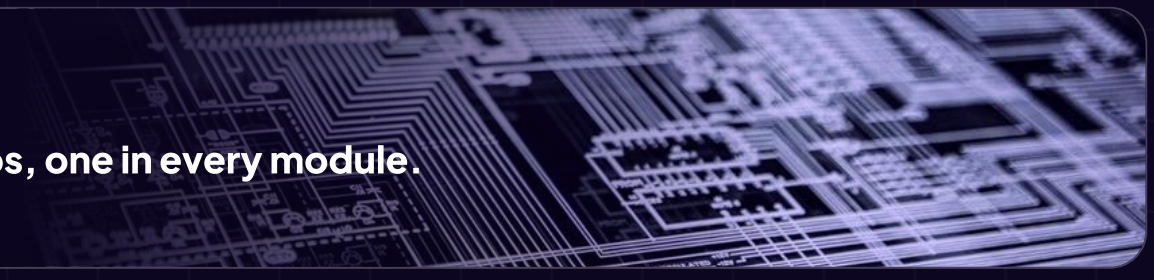
10 hours

Hands-on lab
Deploy a multi-service AI stack to Kubernetes with autoscaling.

TOPICS COVERED

- Dockerfiles for AI services
- Docker Compose stacks
- Kubernetes deployments and services
- Autoscaling with HPA
- GPU workloads overview
- Helm basics

Docker · Kubernetes · Helm



LEARN BY DOING

8 guided labs, one in every module.

✦ DETAILED CURRICULUM

Syllabus, continued.

Modules 07–08

MODULE 07

Observability and Cost Control

4 hours

Hands-on lab
Instrument services with OpenTelemetry and build a latency dashboard.

TOPICS COVERED

- Structured logging
- OpenTelemetry tracing
- Prometheus and Grafana
- Token usage and cost tracking
- Semantic and response caching

Prometheus · Grafana

MODULE 08

CI/CD and API Lifecycle

4 hours

Hands-on lab
Create a CI pipeline that tests, builds and deploys the service.

TOPICS COVERED

- GitHub Actions pipelines
- Contract and load testing
- Versioning and deprecation
- Blue-green and canary releases
- API documentation

GitHub



CAPSTONE

Scalable AI Microservices Platform

✦ YOUR STACK

Tools you'll build with.

Real platforms and frameworks — the same stack modern teams ship with, practised in cloud-hosted labs.

Python
Automation & scripting

FastAPI
Modern Python APIs

Docker
Containers & images

Kubernetes
Container orchestration

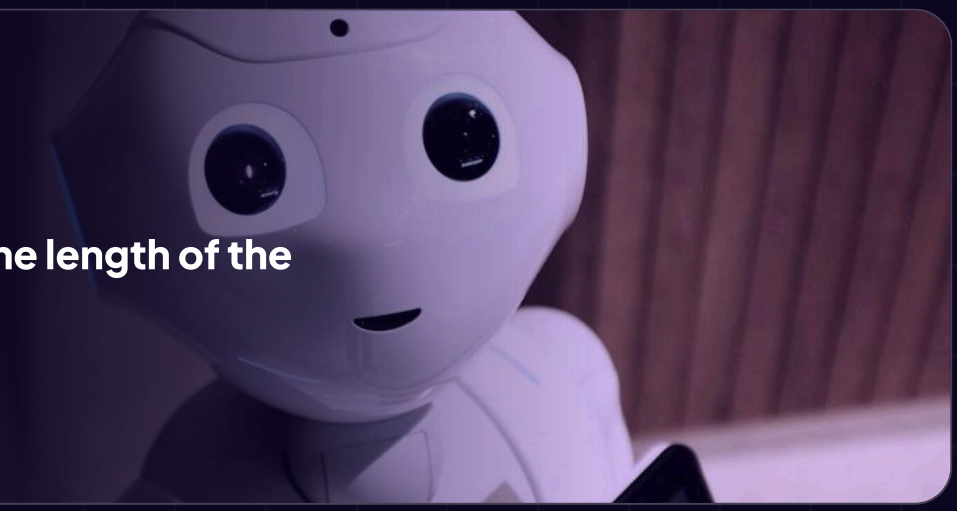
GitHub
Code hosting & Actions CI

Helm
Kubernetes packaging

Prometheus
Metrics & alerting

Grafana
Observability dashboards

PostgreSQL
Relational database



LAB ACCESS

24×7 cloud labs for the length of the program

✦ HANDS-ON WORK

Projects that prove you can do it.

01

Streaming Chat API

Build a secured FastAPI chat service with streaming, conversation storage, rate limiting and provider fallback.

02

Async Document Processor

Create a microservice that queues uploaded documents for AI extraction and exposes job status endpoints.

03

Model Gateway Service

Develop a gateway that routes requests across multiple models with caching, logging and cost tracking.

CAPSTONE PROJECT

Scalable AI Microservices Platform

Design, build and deploy a set of AI microservices with FastAPI, queues, caching, authentication and model integrations on Kubernetes, including CI/CD, observability dashboards, load test results and architecture documentation.

Real-world scenario | Mentor review | Portfolio ready



CERTIFICATION

SAM-AIAPI

Awarded on successful completion of the program assessments.

- MCQ assessments
- Practical lab evaluation
- Project & capstone review
- Final presentation



✦ HOW YOU LEARN

The SAM Learning Framework™

PHASE 01

S

Scenario

Real-world attack and use-case scenarios open every module.

PHASE 02

A

Apply

Hands-on lab work, CTF challenges and live projects.

PHASE 03

M

Master

Assessment, peer review and certification prep.

PHASE 04

M

Mentor

Post-program career guidance and support.

DELIVERY MODELS

- Instructor-Led (ILT)
- Virtual Instructor-Led
- Self-Paced Online
- E-Learning Licensing
- Blended Learning
- Corporate Onsite
- Bootcamp Intensive
- Hackathon / CTF
- Mentorship

✦ LEARNING PATH · AI AUTOMATION & ENTERPRISE INTEGRATION

Where this program fits.

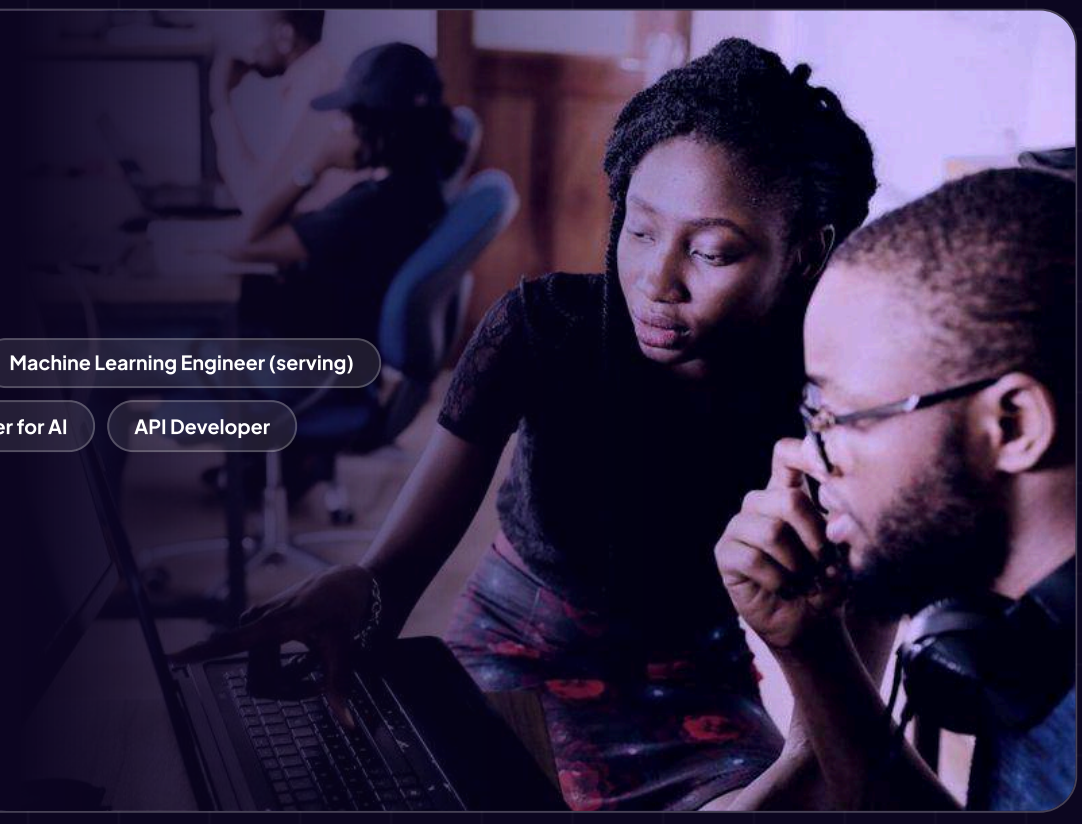
Intermediate	Advanced
2 programs	You are here

BEFORE THIS

AI Automation with n8n / Make / Zapier	32 Hours
Microsoft Power Platform + Copilot Studio	40 Hours

UP NEXT

The top of this pathway — explore related tracks.



CAREER ROLES

- AI Backend Engineer
- Machine Learning Engineer (serving)
- Platform / DevOps Engineer for AI
- API Developer
- AI Solutions Architect

✦ FAQS

Questions, answered.

Q01

Do I need prior experience to join?

Yes. This advanced program expects solid Python backend experience, understanding of REST APIs and HTTP, and basic Docker and Linux skills. Prior exposure to LLM APIs is helpful but key integration patterns are covered.

Q02

How is the program delivered and scheduled?

The program is instructor-led and delivered as live virtual or blended sessions. Batch timings, weekday or weekend schedules and session lengths are shared at enrolment, and recordings or notes support revision between sessions.

Q03

What certification will I receive?

Learners who complete the sessions, labs and final assessments are awarded the Samcommunity Academy SAM-AIAPI certificate. It is issued after successful assessment, not on attendance alone, and reflects the practical skills covered in the program.

Q04

Will I get access to labs and tools?

Yes. Labs use FastAPI, Docker, Kubernetes (local clusters such as kind or minikube), Redis, PostgreSQL and observability tools, all runnable locally. LLM API credits may be needed for some exercises.

Q05

Is there career support after the program?

Learners receive guidance on portfolio projects, resume positioning and interview preparation related to the skills covered. Career support is advisory; the program does not guarantee job placement, specific roles or salary outcomes.



✦ LET'S TALK

Start Building AI-Powered APIs & Microservices with Samcommunity.

INDIVIDUALS

academy@samcommunity.in

Batches, enrolment & career guidance

CORPORATE & INSTITUTIONS

corporate@samcommunity.in

Team batches & custom programs

Xylyn Samcommunity Private Limited · DPIIT-recognized startup · Surat, Gujarat, India