



Certified ISO 9001:2015

Gen and Agentic AI BootCamp



Table of Contents

About the Program	3

Key Highlights of the Program	4

Who is this Program Ideal for?	5

About ActiveSkyGate	5

ActiveSkyGate Advantages	7

Certificates	9

Program Outcomes	10

Learning Path	11

Program Curriculum	12

Career Assistance	25

Tools Covered	26

Certification Process	28

About The Program

This 8-week Agentic AI Bootcamp is designed for working professionals or experienced IT professionals who want to move beyond learning AI concepts and start building real, production-ready AI agents.

You will learn how to design, build, and deploy autonomous AI agents using Azure AI and Azure AI Foundry, working on real-world use cases across development, testing, and business operations.

Unlike traditional AI courses that focus on theory or isolated tools, this Agentic AI program is built around real implementations - helping you automate workflows, integrate with enterprise systems, and deliver measurable business outcomes.

By the end of the program, you won't just understand AI - you will be able to build and deploy AI agents that solve real problems and create real value.



Key Highlights Of The Program



Build and deploy enterprise-grade AI agents in a structured 8-week program



Automate workflows across development, testing, and operations



64+ hours of live, hands-on training focused on real-world implementations



Learn complete lifecycle: architecture → development → deployment



Develop 10+ AI agent solutions using Azure AI & Azure AI Foundry



Get career support: mock interviews, resume, LinkedIn & GitHub optimization



Execute a capstone project aligned with real business use cases



Unlock monetization opportunities with AI agent implementations



Work within the Microsoft AI ecosystem used in enterprise environments

Who Is This Program Ideal For?

This course is for developers, AI enthusiasts, and product teams ready to build real-world AI agents.

✓ AI/ML Engineers

✓ Backend Engineers

✓ Data Scientists

✓ Developers

✓ Data Analysts

✓ Full Stack Developers

✓ Cloud Engineers

✓ Recent College Graduates

About ActiveSkyGate

ActiveSkyGate is a globally recognized, ISO-certified EdTech leader. We specialize in delivering live, career-defining training in emerging and enterprise-ready technologies — bridging the gap between what the industry demands and what available Talent.

Our brand name captures our essence:

- Agile: A reflection of our relentless focus on adaptability, speed, and evolution — the pillars of modern professional excellence.
- Fever: The intense energy and passion that drives our learners, trainers, and partners to stay ahead in a fast-moving tech landscape.
- Together, ActiveSkyGate represents a burning commitment to real-world, agile learning that powers future-ready careers.

We focus on three core pillars:

Certification Programs: Co-delivered with global certifying bodies like Microsoft, Scaled Agile, and Kanban University, ensuring credibility and global recognition.

Industry-Led Bootcamps: Designed in collaboration with top-tier practitioners and academic leaders, our bootcamps offer project-based learning aligned with real job roles.

Career Acceleration & Interview Prep: Including personalized coaching, mock interviews, and portfolio development to help learners land high-impact roles across industries.

Our Expertise - From Artificial Intelligence, Agentic AI, and Generative AI to Cloud, Cybersecurity, Project Management, and Agile methodologies, AgileFever offers a catalog of 200+ cutting-edge courses — all taught 100% live, supported by mentors, capstones, and real-world use cases.

Whether you're a corporate team looking to upskill at scale, or a new graduate preparing for your first role, our programs are built to deliver industry-readiness from Day 1.

Our Impact - With over 25,000 professionals trained and 1,000+ enterprise teams empowered, ActiveSkyGate is not just another EdTech company — we are a strategic capability partner to some of the world's most ambitious individuals and organizations.

We Deliver Trainings of



ActiveSkyGate Advantages

✔ **World-Class Faculty&Mentors**

Learn from top university professors (Caltech, Michigan State, IITs, UT Dallas) and seasoned FAANG professionals.

✔ **Industry-Backed, Job-Ready Curriculum**

Curriculum designed by experts from leading companies, focused on real-world applications, not just theory.

✔ **Fully Live, Instructor-Led Training**

Engage in 100% live sessions with instant doubt resolution and real-time interaction-no pre-recorded content.

✔ **Hands-On Learning With Projects & Labs**

Work on 10+ projects and multiple labs to build practical, job-ready skills.

✔ **Exclusive Lifetime Access**

Get lifetime access to updated course content, so your skills stay relevant long after training ends.

✔ **Private Peer Learning Community**

Join an exclusive network for collaboration, discussions, and professional growth.

✔ **Affordable Premium Education**

Gain top-tier training and mentorship at highly competitive pricing.

✔ **Mentorship From Industry Leaders**

Get personalized guidance from experts at Fortune 500 companies and fast-growing startups.

☑ **Career Coaching & Soft Skills Development**

- Personalized career coaching from certified professionals to align your goals with industry expectations
- Behavioral coaching to help you excel in HR and cultural fit interviews
- Multiple weekly office hours for ongoing support, doubt resolution, and expert guidance

☑ **Job Market Readiness & Visibility**

- Resume refinement with direct input from top recruiters to craft a standout profile
- Mock interviews conducted by real hiring managers to build interview confidence and readiness
- LinkedIn profile optimization to boost visibility and attract top opportunities



Certificates

Upon successful completion of the Gen and Agentic AI BootCamp, you will receive a Certificate of Completion from ActiveSkyGate. Additionally, you will be awarded certificates for all applicable modules within the program, recognizing your dedication, initiative, and commitment to advancing your expertise in AI. These certificates serve as a testament to your skills and growth as an AI professional.



Program Outcomes

After completing this course, learners will be able to:

- ✓ Build intelligent AI agents from scratch using Python.
- ✓ Understand the full agent architecture and roles of each part.
- ✓ Choose and use frameworks like AutoGen and CrewAI to accelerate development.
- ✓ Using tools to fetch the data from the Internet and build some use cases.
- ✓ Create agents with persistent vector memory.
- ✓ Develop agents that dynamically adapt plans and tool usage.
- ✓ Design multi-agent systems with defined roles.
- ✓ Deploy scalable and maintainable agent applications.
- ✓ Implement safety constraints and perform ethical checks.
- ✓ Design agent SaaS applications for real-world use cases.
- ✓ Deliver a functional multi-agent system ready for practical application.

Learning Path



Module 1

LLM Internals, GenAI Landscape & Reasoning Models

GenAI Capabilities & Business Applications

- What GenAI can actually do in enterprise — text, code, images, search, summarisation, classification. Real-world use cases across industries

Limitations & Challenges of GenAI

- Hallucinations, context limits, cost, latency, data privacy, bias, stale knowledge — what agents are built to address

Transformer Architecture

- Self-attention, multi-head attention, encoder-decoder structure, positional encoding — intuitive understanding for practitioners

Tokenisation & Context

- Windows BPE tokenisation, token costs, why “1 token $\neq$ 1 word”, context length limits and practical implications

Sampling & Temperature

- Temperature, top-p, top-k, log probabilities — controlling determinism vs creativity in agent outputs

Reasoning Models ★

- GPT-o1, o3, Claude 3.7 extended thinking, Gemini 2.0 Flash Thinking — how they differ, when to use them, prompting strategies

LLM Providers & APIs

- OpenAI GPT-4o, Anthropic Claude 3.5/3.7, Google Gemini 1.5/2.0 — API setup, model selection, cost comparison

Prompt Caching ★

- Anthropic and OpenAI prompt caching — 60–80% cost reduction for agents with repeated context

Labs: GenAI capability demos · Tokenisation experiments · Reasoning model comparison · API cost calculator

Project: LLM Explorer — compare outputs across models, temperature settings, reasoning vs standard mode

Tools: OpenAI API · Anthropic API · Google Gemini API · Jupyter · tiktoken · Pandas · Matplotlib

Module 2

Prompt Engineering + Context Engineering

Advanced Prompting Patterns

- Chain-of-Thought (CoT), Tree-of-Thought (ToT), ReAct, self-consistency, prompt chaining, few-shot design

Prompt Debugging & Benchmarking

- Hallucination detection, A/B testing prompts, log probability analysis, reliability scoring, prompt benchmarking

Context Engineering ★

- What to put in the context window — compression, pruning, routing, summarisation, context budgets

Structured Outputs ★

- JSON mode, Pydantic v2 validation, function calling schemas, output parsers, retry on parse fail

Prompt Versioning

- Managing prompts as code — versioning, templates, LangSmith prompt hub AI-

Native Dev Tools ★

- Cursor + Claude Code as development accelerators — used throughout the bootcamp from this module forward

Labs: Prompt A/B testing · Hallucination rate measurement · Context compression experiment · JSON schema design

Project: Prompt Evaluation Platform — compare prompts, measure hallucination rates, benchmark structured outputs

Tools: LangChain · LangSmith Prompt Hub · Pydantic v2 · OpenAI API · Cursor · Claude Code

Module 3

Advanced RAG Engineering + Agentic RAG

Embeddings & Similarity

- Metrics Embedding model comparison (OpenAI, Cohere, BGE), cosine/dot product/Euclidean metrics, dimensionality tradeoffs

Chunking Strategies

- Semantic chunking, sliding window, recursive splitting, metadata tagging, parent-child chunks

Vector Databases

- FAISS, Pinecone, Weaviate, Chroma, pgvector — selection criteria, index management, production operations

Advanced Retrieval

- Hybrid search (dense + BM25), reranking (Cohere/BGE), query rewriting, HyDE, multi-query retrieval

Agentic RAG ★

- RAG inside an agent loop — agent decides WHEN to retrieve, WHAT to retrieve, iterative/self-correcting retrieval, routing between strategies

RAG Evaluation

- RAGAS framework, context precision/recall, faithfulness, answer relevance, LangSmith eval datasets

Multi-Document Reasoning

- Cross-document retrieval, citation chains, conflicting source resolution

Labs: Embedding benchmarking · Chunking comparison · Hybrid retrieval · Agentic RAG loop · RAGAS eval

Project: Enterprise Knowledge Intelligence System — agentic RAG pipeline with iterative retrieval and evaluation dashboard

Tools: LangChain · LlamaIndex · Pinecone · FAISS · Weaviate · Cohere · LangSmith · RAGAS

Module 4

Fine-Tuning LLMs, Multimodal AI & Document Intelligence

Parameter-Efficient Fine-Tuning

- LoRA, QLoRA, PEFT — when to fine-tune vs RAG vs prompt engineer, cost-benefit analysis

Dataset Preparation

- Instruction datasets, ShareGPT/Alpaca format, data cleaning, quality filtering with Pandas

Model Evaluation

- BLEU, ROUGE, perplexity, domain-specific evals, model versioning with Weights & Biases, Matplotlib visualisation

Vision-Language Models

- GPT-4o vision, Claude 3.5 vision, multimodal embeddings, image + text reasoning chains, multimodal reasoning

Diffusion Models +

- How diffusion works — noise schedules, denoising process, Stable Diffusion, DALL-E, text-to-image APIs, use in agent pipelines

Document AI / ADE ★

- Agentic Document Extraction — PDFs, scanned docs, tables, handwriting, mixed-language content. Unstructured.io, LlamaParse

Labs: LoRA adapter training · Vision-language chain · Diffusion model API demo · Document extraction pipeline

Project: Multimodal Document Intelligence Agent — extract structured data from PDFs + images

Tools: Hugging Face · PEFT · LoRA · OpenAI Vision · Anthropic Vision · Unstructured.io · W&B; · Pandas · Matplotlib

Module 5

Agentic AI Foundations & Architecture Patterns

GenAI to Agentic Evolution

- From chatbots to agents — autonomy levels, decision loops, agent vs workflow vs pipeline distinction

Agent Components

Planner, Executor, Tools, Memory — roles, interactions, failure modes of each component

Architecture Patterns

- ReAct, Plan-and-Execute, Tree of Thought, Reflection loops, self-critique patterns

Goal Decomposition & Task Queues +

- Breaking high-level goals into subtasks, task queue management, dependency resolution, execution ordering

Decision Making Under Constraints +

- Agents operating within budget limits, time constraints, tool availability — constraint-aware planning

Agent Types

- Reactive vs deliberative vs hybrid agents, goal-driven vs utility-based, single vs multi-agent contexts

Labs: Rule-based planning agent · Goal decomposition exercise · Constraint-aware planning demo

Project: Goal-Seeking Assistant — deliberative agent with constraint-aware planning, task queues, and trace logging

Tools: LangGraph · LangSmith · OpenAI API · Python

Module 6

Agent Frameworks: LangGraph, CrewAI, AutoGen & OpenAI Agents SDK

LangGraph Deep Dive

- Stateful graph design, cyclic graphs, node types, conditional edges, checkpointing, state reducers

CrewAI Framework

- Role-based agents, task delegation, crew orchestration, process types (sequential/hierarchical)

AutoGen / AG2

- Conversational agents, group chat managers, code execution agents, agent negotiation patterns

OpenAI Agents SDK ★

- Native OpenAI approach — agents, handoffs, guardrails, tracing. Rapidly becoming enterprise standard in 2026

Framework Tradeoffs

- LangGraph vs CrewAI vs AutoGen vs OpenAI SDK — complexity, flexibility, enterprise readiness, community size

Semantic Kernel

- Microsoft SK for enterprise orchestration, plugin design, Azure AI Foundry integration

Labs: LangGraph stateful pipeline · CrewAI crew setup · OpenAI Agents SDK handoff demo · Framework tradeoff comparison

Project: Same use case built in 2 frameworks — compare developer experience, output quality, cost

Tools: LangGraph · CrewAI · AutoGen · OpenAI Agents SDK · Semantic Kernel · LangSmith

Module 7

Tool Use, Memory Architecture & Structured Outputs

Function Calling & Tool Schemas

- OpenAI/Claude function calling, tool schema design, Pydantic models, output validation, parallel tool calls

Tool Integration

- REST/GraphQL API calling, browser tools, file system read/write, code execution sandboxes

Short-Term Memory

- Conversational context, window management, token-aware summarisation buffers

Long-Term / Vector Memory

- FAISS/Pinecone/ChromaDB for persistent memory, episodic vs semantic vs procedural memory

Memory Architecture Patterns ★

- mem0, Zep, MemGPT — purpose-built agent memory systems, choosing the right architecture per use case

Memory Persistence

- LangGraph memory nodes, cross-session state, memory retrieval and update strategies

Labs: Tool schema design · API-calling agent · Vector memory chatbot · mem0 integration
· Cross-session persistence test

Project: Chatbot with Persistent Memory — vector-backed + mem0 user context across multiple sessions

Tools: LangChain · LangGraph · FAISS · ChromaDB · Pinecone · mem0 · Zep · OpenAI Tools
· Pydantic v2

Module 8

Multi-Agent Systems, HITL & Communication Protocols

Multi-Agent Orchestration

- Role assignment, task decomposition, agent handoffs, parallel vs sequential execution, supervisor patterns

Message Passing & Negotiation

- Inter-agent communication, shared state management, conflict resolution, consensus mechanisms

Human-in-the-Loop (HITL) ★

- Approval gates, escalation paths, oversight checkpoints for healthcare/finance/legal — fail-safe patterns

MCP — Model Context Protocol

- Architecture, tool/resource exposure, server/client design, interoperability with external services

A2A & ACP Protocols

- Agent-to-Agent protocol, Agent Collaboration Protocol — comparison, when to use each, implementation with CrewAI/AutoGen

Labs: Buyer-seller multi-agent simulation · HITL approval gate · MCP server setup · A2A handoff demo

Project: Multi-Agent Collaboration System — role-based agents with HITL checkpoint and protocol-based communication

Tools: CrewAI · AutoGen · LangGraph · OpenAI Agents SDK · MCP SDK

Module 9

Evaluation Framework, LLMOps, Safety & Guardrails

Evals Framework ★

- Defining eval metrics, building eval pipelines, regression testing agent behaviour — RAGAS, Braintrust, UpTrain

LangSmith Deep Dive

- Tracing, evaluation datasets, prompt versioning, A/B testing, production monitoring dashboards

Agent Monitoring & LLMOps

- Latency logging, retry mechanisms, fallback logic, cost/token tracking, alerting on failures

Safety & Guardrails

- Prompt injection defense, loop prevention, constraint design, goal alignment, Guardrails AI / NeMo Guardrails

RBAC & Compliance

- Permission scoping, API key lifecycle, audit logging, trust boundaries, GDPR/HIPAA/SOC2 considerations

Labs: Eval pipeline setup · LangSmith trace analysis · Guardrails AI integration · RBAC policy design

Project: Goal-Chasing Agent with Safety Checks — guardrails, eval pipeline, LangSmith monitoring dashboard

Tools: LangSmith · RAGAS · Braintrust · Guardrails AI · NeMo Guardrails · Pandas · Matplotlib

Module 10

Agent Deployment, Infrastructure & Enterprise AI on Azure

Containerisation & CI/CD

- Dockerising agents, GitHub Actions pipelines, environment variables, secrets management

Cloud Deployment Options

- HuggingFace Spaces, AWS Lambda, GCP Cloud Run — selection criteria, cost comparison, scaling strategies

Azure AI Foundry

- Azure AI Studio setup, Azure OpenAI Service, Azure AI Search for enterprise RAG pipelines, agent development on Azure

Semantic Kernel on Azure

- Orchestrating agents with SK, connecting to enterprise business data, Copilot Stack integration

Azure Monitoring & Performance ★

- Azure Monitor, Application Insights for agents, logging agent decisions, performance optimisation on

Azure Azure Cost Management ★

- Using Azure free credits efficiently, avoiding unnecessary API costs, cost alerts, right-sizing deployments

Microsoft Applied Skills Badge

- Guided MS Learn lab completion, step-by-step instructor guidance, official MS Applied Skills badge from Microsoft

Labs: Docker deployment · HF Spaces live link · Azure AI Studio setup · Azure Monitor setup · MS Learn labs

Project: Cloud-Deployed Agent — Dockerised LangGraph agent on HF Spaces + Azure enterprise knowledge assistant

Tools: Docker · GitHub Actions · HuggingFace Spaces · Azure AI Foundry · Semantic Kernel · Azure Monitor · Application Insights

Module 11

Agent-as-a-Service, Scheduling & Cost Optimisation

SaaS Pricing Models

- Per-call, subscription, usage-based — token budgeting, cost modelling, margin calculation for production agents

Prompt Caching in Production

- GPTCache, Anthropic/OpenAI caching APIs — 60–80% cost reduction for repeated context in agent loops

Scheduling & Persistent Jobs ★

- Running agents on a schedule (cron-style), background job management, queue systems (Celery/Redis), agent status dashboards

Model Routing & Cost

- Optimisation Routing between cheap/expensive models dynamically, semantic caching, budget guardrails per agent

Agent Productisation

- Multi-tenant architecture, API gateway design, rate limiting, packaging agents as commercial products

Streamlit / Gradio

- UIs Building interactive demo UIs for rapid prototyping, client demos, and SaaS MVPs

Labs: GPTCache setup · Celery job scheduler demo · Model routing logic · Streamlit agent UI · Cost calculator

Project: Agent-as-a-Service — Travel Planner or Sales Bot with scheduling, caching, pricing model, and live demo UI

Tools: Streamlit · Gradio · GPTCache · FastAPI · Celery · Redis · OpenAI API · Anthropic API

Module 12

AI Coding Agents & Agentic Development Workflows

AI Coding Agent Landscape

- GitHub Copilot Workspace, Devin-style agents, Claude Code, Cursor Agent mode — what enterprises are deploying in 2026

Agents that Write & Test Code

- Using LLM agents to generate, refactor, debug and test code — the #1 enterprise use case in 2026

Building a Coding Agent

- Hands-on: build a simple code-generation + test-execution agent using OpenAI Agents SDK or LangGraph

Labs: Coding agent quickbuild · Automated test generation demo

Project: Code Review Agent — takes a repo, generates inline documentation and test suggestions

Tools: OpenAI Agents SDK · Claude Code · LangGraph · GitHub API · Python

Module 13

Capstone Project (4-Track Group Model)

- **Capstone Track A:** Enterprise Knowledge Assistant — Agentic RAG + multi-agent + Azure deployment + safety + evals
- **Capstone Track B:** AI GTM Strategy Agent — multi-agent market intelligence, document analysis, strategic planning report
- **Capstone Track C:** DevOps / Code Review Agent — automated code review, test generation, CI/CD pipeline integration
- **Capstone Track D:** Customer Support + CRM Agent — memory, HITL, Zoho/Salesforce integration, Agent SaaS pricing

Labs: End-to-end capstone build · Peer code review · Mock interview · Portfolio demo day (4 live demos)

Project: Full capstone: RAG + multi-agent + memory + deployment + safety + evals + monitoring in one system

Tools: All tools from prior modules

Career Assistance

■ **AI- Powered Resume & Profile Building**

Your resume, LinkedIn, and GitHub — optimized by industry professionals to stand out to recruiters and land interviews faster.

■ **Mock Interviews**

Practice real technical and behavioural interviews with honest feedback from people who actually hire for AI roles.

■ **1:1 Career Mentoring**

Work directly with industry veterans to position yourself for AI roles — covering job search strategy, communication, and career planning.

■ **Hiring Exposure**

Get direct visibility with active hiring managers. Understand what they actually look for — and how to stand out in competitive AI hiring pipelines.

Tools Covered

Gain hands-on experience with leading tools:



Visual Studio Code



Google Antigravity



OpenAI



Gemini



LangChain



LangGraph

crewai



Ollama

FAISS
(FACEBOOK AI SIMILARITY SEARCH)



n8n



Streamlit



GitHub Copilot



LangSmith



docker



Hugging Face



Guardrails AI



Azure AI Foundry



Render



Google
colab

Complete Tool Coverage

- **LLM Providers:** OpenAI GPT-4o / o1 / o3 · Anthropic Claude 3.5 / 3.7 · Google Gemini 1.5 / 2.0 · Hugging Face
- **Agent Frameworks:** LangGraph · CrewAI · AutoGen / AG2 · OpenAI Agents SDK · Semantic Kernel · CAMEL
- **RAG & Retrieval:** LangChain · LlamaIndex · RAGAS · Cohere Rerank · HyDE · BM25
- **Vector Databases:** FAISS · Pinecone · Weaviate · ChromaDB · pgvector Agent Memory mem0 · Zep · MemGPT · LangGraph Memory Nodes · ChromaDB
- **Observability:** LangSmith · Braintrust · W&B; · LLMOps Space · UpTrain · Azure Monitor · Application Insights
- **Safety & Evals:** Guardrails AI · NeMo Guardrails · RAGAS · Pydantic v2 · LangSmith Evals · Braintrust
- **Deployment:** Docker · GitHub Actions · HuggingFace Spaces · AWS Lambda · GCP Cloud Run
- **Azure Ecosystem:** Azure AI Foundry · Azure OpenAI Service · Azure AI Search · Semantic Kernel · Azure Monitor
- **Document AI:** Unstructured.io · LlamaParse · Azure Document Intelligence · PDFPlumber
- **Dev Tools:** Cursor · Claude Code · GitHub Copilot · Jupyter · VS Code · tiktoken
- **Frontend / UI:** Streamlit · Gradio · FastAPI · GPTCache
- **Scheduling:** Celery · Redis · APScheduler · FastAPI Background Tasks
- **Fine-Tuning:** PEFT · LoRA · QLoRA · Hugging Face Trainer · Weights & Biases
- **Data & Utils:** Python · NumPy · Pandas · Matplotlib · JSON · GitHub API

Certification Process

Candidates can register and complete to this program in 4 simple steps:

STEP
1

Register for the Program

Enroll in the Gen and Agentic AI BootCamp

STEP
2

Attend and Complete the Training

Attend all live session sand complete the course curriculum

STEP
3

Complete the Capstone Project

Apply your learning in a hands-on project to demonstrate your practical skills

STEP
4

Receive Course Completion Certificate

Upon successful completion, receive an industry-recognized certificate from ActiveSkyGate

Pre-requisites:

No prior experience in Agentic AI is required. However, to ensure everyone starts with the same foundation, participants must complete the following:

Python for AI Fundamentals (6 Hours)

Covers essential Python concepts used in AI applications, automation workflows, and agent development. This should be completed before the training begins.

n8n Automation Fundamentals (4 Hours)

Introduces workflow automation, integrations, triggers, and AI-powered automation using n8n. After completing the program, participants will receive access to recordings, and learning materials will be provided for self-paced learning.

Note: These foundational topics will not be covered during the live training and are essential for participating in hands-on labs and agent-building exercises.

Exam Details

There is no exam for this Agentic AI bootcamp. Completion of the capstone project serves as the final evaluation for the program.



ActiveSkyGate

United States

ActiveSkyGate Americas, Inc.
13107, Risen Star Road, Frisco, Texas,
75035 United States
Email us at: dibo@activeskygate.com
Call us at: +1 (469) 444-0274

Australia

ActiveSkyGate LLC
Unit 12/19-21 Veron, St. Wentworthville
NSW 2145, Australia
Email us at: dibo@activeskygate.com