Microsoft

PTI

Applied Agentic AI

Systems, Design & Impact Program

Microsoft

Why Agentic AI Matters Now

AI has moved beyond experimentation. Companies are now deploying it at scale to boost productivity and reinvent products—creating urgent demand for professionals who can lead, not just observe, this shift.

The next edge is agentic AI and multi-agent systems—AI that can plan, decide, and act across complex workflows. Organizations are rapidly increasing investment in AI agents, reshaping how products are built and how work gets done. This is creating a powerful opportunity for mid to senior product managers and tech leaders, as roles that combine product thinking with agentic AI are becoming some of the most sought-after and best-compensated in the AI ecosystem.

By Numbers

66%

of companies using AI agents report measurable productivity gains

88%

of executives plan to increase AI budgets by May 2026 due to agentic AI

24% → 67%

Organizations where AI agents take independent action today vs. by 2027

28% → 57%

Executives expecting autonomous decision-making from agentic AI in workflows by 2027

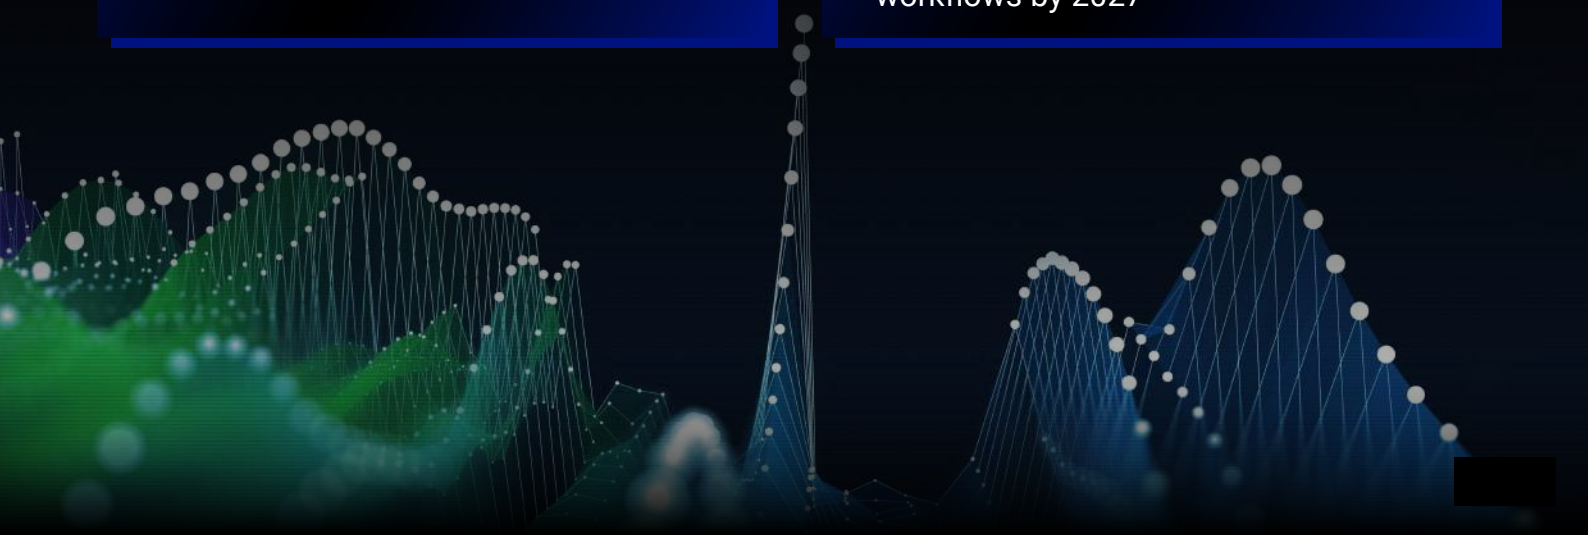

Microsoft

Learning Path

Course Name	Learning Hours
Python Refresher With AI (Optional)	12 Hrs
Foundations of AI & Agentic AI	3 Hrs
Generative AI Tech Stack and Prompt Engineering	6 Hrs
LLM Internals & Planning Systems	6 Hrs
Multi-Agent Ecosystems (I)	6 Hrs
Multi-Agent Ecosystems (II)	9 Hrs
Model Context & Tooling Protocol	6 Hrs
Metrics, GTM & ROI	3 Hrs
Agentic UX Design & Transparency	3 Hrs
Dev Tools & Product Readiness	3 Hrs
Develop AI Agents on Azure (Microsoft)	8 Hrs
Capstone: Product Strategy Simulation	6 Hrs

Elective Courses

Develop Generative AI Apps in Azure (Microsoft)	7.5 Hrs
Masterclass on Agentic AI Solutions Using Copilot Studio and AutoGen	2 Hrs

Microsoft

COURSE 1

Python Refresher With AI (Optional)

This course revisits Python programming essentials tailored for AI/ML applications. It covers core programming constructs, environment setup including IDEs and cloud platforms, data structures, control flows, OOP basics, file handling, and introduces AI-powered code generation tools like GitHub Copilot. Practical exercises focus on real-world data manipulation and culminate in a capstone project comparing traditional and AI-assisted coding approaches, preparing learners to effectively integrate Python skills with evolving AI workflows.

Guided Practices:

Building Employee Financial Insights With Python: Build a Python system to process employee salary data using variables, data types, operators, inputs/outputs, and error handling while following coding best practices to generate basic financial insights

Storing and Analyzing Customer Feedback in Python: Store, organise, and analyse customer feedback using Python lists, tuples, dictionaries, and sets. Use loops and conditionals to examine records and uncover meaningful insights

Creating Object-Oriented Programs in Python: Create reusable functions and classes in Python, apply single-level inheritance, and build a working script that demonstrates core object-oriented concepts with clean, efficient design

Building a Number-Guessing Game With GitHub Copilot: Build a number-guessing game in Python using GitHub Copilot in VS Code. Use inline comments and docstrings to guide code generation and produce a clean, fully functional game

Microsoft

COURSE 2

Foundations of AI & Agentic AI

Introducing the Learn > Build > Deploy framework, this course covers AI hierarchy distinctions (AI, ML, DL, GenAI, Agentic AI), transformer architectures, and autonomous AI agents. Topics include key papers like "Attention is All You Need," CoT prompting, ReAct frameworks, and a 4-layer GenAI stack analogy. It establishes foundational knowledge critical to technical product management of AI systems with an emphasis on theoretical and applied agentic AI concepts.

Skills Developed:

- AI Literacy
- GenAI Fundamentals
- Transformer Understanding
- Prompting Techniques

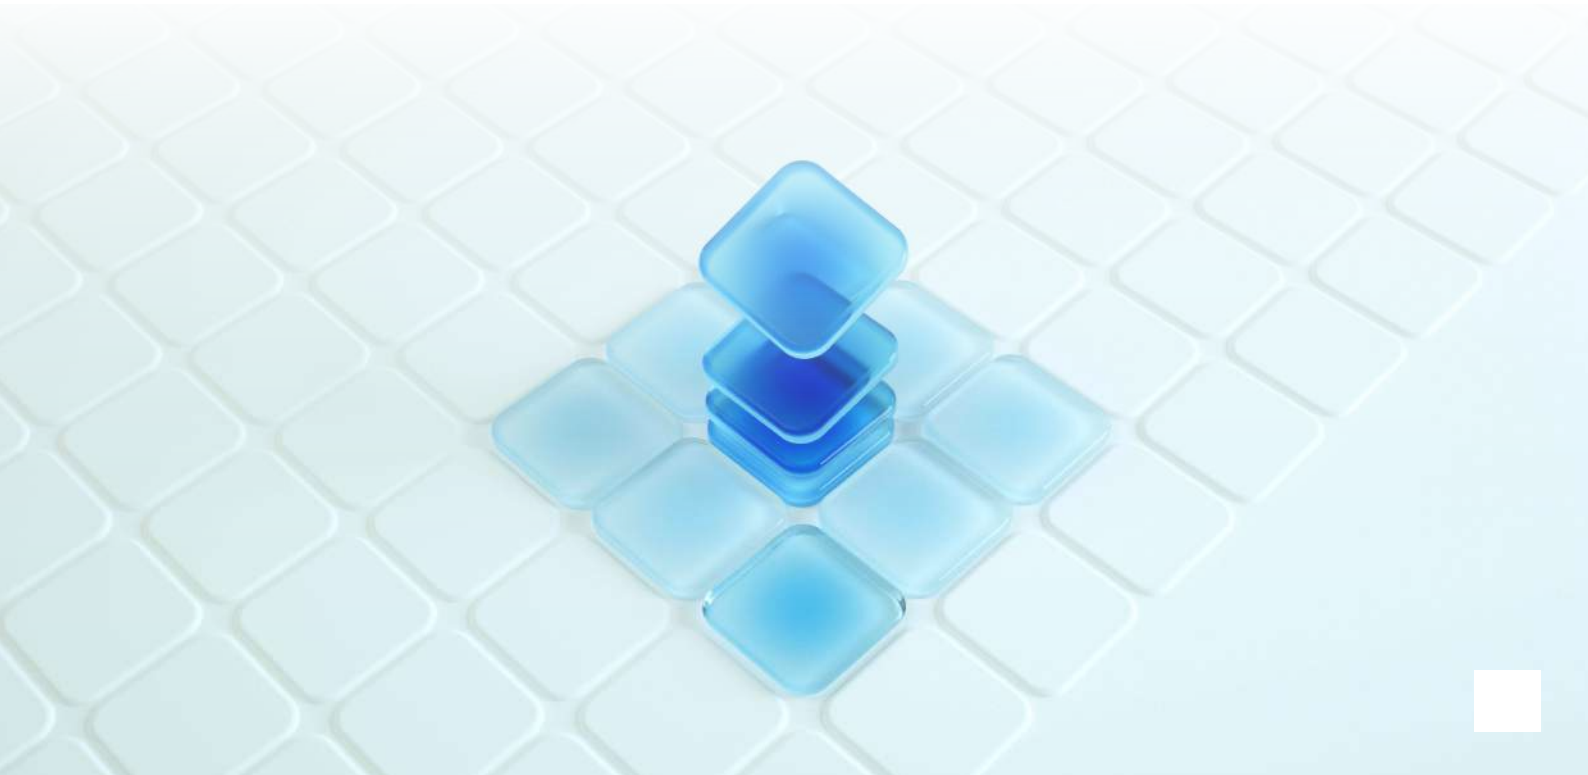

Microsoft

COURSE 3

Generative AI Tech Stack and Prompt Engineering

Introducing the Learn > Build > Deploy framework, this course covers AI hierarchy distinctions (AI, ML, DL, GenAI, Agentic AI), transformer architectures, and autonomous AI agents. Topics include key papers like "Attention is All You Need," CoT prompting, ReAct frameworks, and a 4-layer GenAI stack analogy. It establishes foundational knowledge critical to technical product management of AI systems with an emphasis on theoretical and applied agentic AI concepts.

Tools and Frameworks Covered:

Lovable | emergent | LangChain | crewai | AutoGen | LangGraph

Skills Developed:

- GenAI Tech Stack
- Vector Databases
- AI UX Design
- Agent Orchestration
- Low-Code Prototyping

Microsoft

COURSE 4

LLM Internals & Planning Systems

Focusing on PM productivity with AI, this course teaches prompt engineering principles and planning systems using LangChain and function calling APIs. It includes live sessions building Q&A bots integrating APIs, planning workflows with agents, and advanced prompt strategies to optimize interaction with language models. Labs guide the development of multi-step agents and contextual tool integration, enhancing practical skills in agent-based product development.

Guided Practices:

Prototyping Product Ideas With AI Tools: Rapidly prototype product ideas using AI tools like Lovable, Uizard, and Firebase Studio to explore how AI accelerates value definition, UI creation, and feature simulation

Building a Multi-Agent System for Feature Prioritisation: Build a LangChain-powered multi-agent system using Azure OpenAI to simulate feature prioritisation by orchestrating domain agents to score features and generate a unified priority list

Demos:

- Debugging prompt failures in product feature summaries
- Building a RAG-powered FAQ agent with custom knowledge
- Applying CoT and ReAct prompting for drafting user stories and analyzing problems
- Building a finance Q&A agent using LangChain

Microsoft

COURSE 5

Multi-Agent Ecosystems (I)

This course explores advanced retrieval-augmented generation (RAG) systems and multi-agent architectures through hands-on implementation with CrewAI and LangGraph. It details agent collaboration patterns, role-based architectures using YAML, memory strategies, and real-world orchestration frameworks. Learners build modular multi-agent teams focusing on scalability, state management, and autonomous information synthesis. Deliverables include pitches advocating modular agent architectures.

Guided Practices:

Generating Product Specs With MetaGPT: Use MetaGPT's multi-agent system to generate a complete product specification from a single business prompt and observe how coordinated agent roles enable structured product development

Orchestrating AI Agents With LangGraph: Orchestrate cross-framework AI agents using the A2A protocol with LangGraph to understand how modular agents share context, collaborate on research tasks, and produce unified outputs

Demos:

- Understanding the Anthropic Research Agent
- Executing serial and parallel workflows using CrewAI agents
- Building a CrewAI multi-agent system with RAG (chatbot)

Course-End Project:

Building an Agentic RAG Router System

Design and implement an Agentic RAG system with role-based agents. Use a Router Agent and Retriever Agent to route queries across PDF vector search or web tools and generate accurate, source-grounded answers

Tools and Frameworks Covered:

crewai | LangGraph | A2A | MetaGPT | AUTOGPT | Agentic RAG Design Frameworks

Skills Developed:

- Multi-Agent Design
- Memory Strategies
- Workflow Orchestration
- RAG Architecture
- Agent Collaboration
- Analytical Reasoning

Microsoft

COURSE 6

Multi-Agent Ecosystems (II)

Building on foundational multi-agent knowledge, this course delves into enterprise-grade agent orchestration using Microsoft AutoGen and n8n workflow automation. It covers communication protocols, database integration, and production deployment strategies. Projects include developing marketing agent pipelines with attention to scalability, performance, security, and compliance. Visual workflows and protocol deep dives support mastery of complex distributed agent ecosystems.

Guided Practices:

Building a Multi-Agent Support Triage System: Build a multi-agent AutoGen workflow that triages support tickets by coordinating role-based agents and a planner to deliver transparent, data-driven prioritization decisions

Designing Human-in-the-Loop Publishing Workflows: Create an n8n workflow that combines agentic content generation with human review, routing drafts through approval and SEO enhancement for controlled, safe publishing

Demos:

- Installing and running AutoGen locally
- Response format negotiation for multi-agent communication
- Integrating n8n for automated customer support ticket routing
- Using AutoGen for autonomous internal document QA & policy compliance
- Implementing AutoGen agent chain for automated code testing & debugging

Course-End Project:

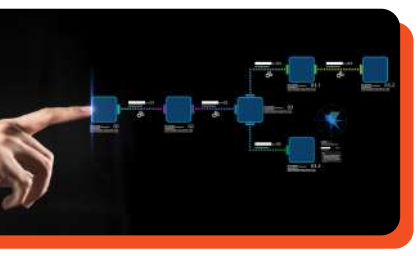

Building AI-Powered LinkedIn Automation With n8n

Design an AI-powered content automation system using n8n and an AutoGen-style microservice to ideate, draft, review, and publish LinkedIn posts with guardrails for scale and consistency

Tools and Frameworks Covered:

AutoGen · n8n · docker · Phoenix Framework · Gmail · LangSmith · Google Docs · mongoDB · asana · slack · PostgreSQL · Pinecone · A2A · Agentic RAG Workflow Patterns

Skills Developed:

- Agent Orchestration
- Multi-Agent Communication
- Compliance
- Workflow Automation
- Database Integration

Microsoft

COURSE 7

Model Context & Tooling Protocol

This course introduces the Model Context Protocol (MCP) for integrating and standardizing AI tools. Topics include structured context binding, interoperability standards, JSON schema design, secure tool hosting, and memory persistence. Labs develop contextual AI agents chaining outputs across tools with authentication and performance optimization. Emphasis is on enterprise readiness, security best practices, and tool discoverability through standardized protocols.

Demos:

- Building a MarketIntel MCP server for market research
- Response format negotiation for multi-agent communication
- Integrating n8n for automated customer support ticket routing
- Building a research agent workflow with MCP server in n8n
- Implementing AutoGen agent chain for automated code testing & debugging

Tools and Frameworks Covered:

n8n | GitHub | Model Context Protocol | FASTMCP

Skills Developed:

- Context Binding
- Schema Definition
- Performance Optimization
- Tool Interoperability
- Secure Tool Hosting

Microsoft

COURSE 8

Metrics, GTM & ROI

Offering a comprehensive framework, this course teaches measurement of AI agent performance using OKRs, key indicators like success rate and latency, and ROI calculations. It covers observability tooling with LangSmith and Phoenix, real-time logging, and conversational analysis. Business strategy topics include pricing, go-to-market planning, and deployment of agent MVPs with analytics dashboards. Practical instrumentation and monitoring equip learners for operational excellence.

Demos:

- Setting up LangSmith tracing on a simple AI agent
- Creating a Lean Canvas on Miro (collaborative exercise)
- A/B testing AI responses with prompt variants

Tools Covered:

LangSmith

miro

Skills Developed:

- Agent Metrics Analysis
- A/B Testing
- GTM Strategy
- Observability Setup
- ROI Evaluation
- Agent Monitoring

COURSE 9

Metrics, GTM & ROI

Centering user experience for AI products, this course covers interaction design patterns for agentic UX, including flexible, probabilistic flows, ambiguity handling, and human-in-the-loop checkpoints. It addresses ethical risks such as hallucinations and bias and teaches guardrail implementations and transparency techniques such as confidence disclosures and explainability interfaces. Learners create complete UX prototypes emphasizing trust, user control, and fail-soft design.

Guided Practice:

Designing Proactive Agent Behaviours in Figma: Design a proactive agent confirmation flow in Figma that presents AI actions, reasoning, and confidence, enabling users to approve or override decisions through a transparent interface

Demos:

- Creating a chat-style prototype in Figma
- Creating a reasoning tooltip and a planned action in Figma
- Designing an interface with AI-suggested actions and user control buttons
- Designing human-in-the-loop approval and feedback interfaces in Figma

Course-End Project:

Designing an Agentic UX Trust Prototype

Design a trust-focused Agentic UX prototype for a financial advisor AI that previews actions, explains reasoning, and enables override or cancellation, with LangChain simulating agent planning and logic

Tool Covered:

Figma

Skills Developed:

- UX Prototyping
- Human-in-the-Loop Design
- UX for Human & AI Interaction
- AI-First Design
- UX Design for AI Risk Mitigation

Microsoft

COURSE 10

Dev Tools & Product Readiness

Focused on deployment and live operations, this course examines cloud vs edge hosting, serverless and containerized environments, and model hosting strategies. It includes hands-on Firebase and n8n automation workflows, feedback and testing system integrations for user insights, alert configurations for monitoring, and infrastructure-as-code introductions with Terraform and Pulumi. The course prepares learners for scalable, maintainable AI product readiness.

Guided Practice:

Building a Full-Stack Agent MVP in Colab: Build a full-stack agent MVP in Google Colab using Python, ipywidgets, and the OpenAI API to create a chat interface, log user interactions, and capture feedback for continuous improvement

Demos:

- Log and visualize user–agent interaction events in Firestore
- Run a self-hosted open-source LLM locally using Ollama
- Embed a Lovable feedback widget and track Firebase Analytics events

Course-End Project:

Designing a Multi-Agent Workflow Planner

Design a multi-agent workflow planner for a startup accelerator where role-based agents collaborate to complete tasks like generating a pitch deck outline using coordinated planning and feedback loops

Tool Covered:

Hugging Face

Skills Developed:

- Workflow Automation
- Infrastructure-as-Code
- Product Readiness
- Event Logging
- Model Hosting

Microsoft

COURSE 11

Develop AI Agents on Azure (Microsoft)

This course focuses on building AI agents leveraging Microsoft Azure's cloud infrastructure and toolset. It covers Azure-specific frameworks, deployment workflows, security integration, and scalable orchestration techniques. Learners gain hands-on experience developing and hosting AI agents in the Azure ecosystem with attention to enterprise-grade reliability and compliance.

Skills Developed:

- Deployment on Azure
- Security Integration
- Enterprise Compliance
- Cloud Orchestration
- Scalable Agent Design
- Cloud Architecture

Microsoft

Capstone

Exclusive Mentor-Led Project Support

The capstone integrates multi-agent system design and go-to-market planning through a production-grade project building a 4-agent market research and GTM framework using n8n and CrewAI with MCP integration. It emphasizes business strategy development, including Lean Canvas, pricing models, and acquisition strategies, alongside instrumented agent performance data and real-world chatbot deployment. This synthesis project prepares learners for practical AI product leadership.

Tools and Frameworks Covered:

n8n | crewai | MCP

Skills Developed:

- Multi-Agent System Design
- MCP Integration
- GTM Planning
- Pricing Strategy
- Performance Analytics
- Real-World Deployment

Microsoft

Elective Courses

Develop Generative AI Apps in Azure (Microsoft)

This course prepares you to build AI solutions on Azure using Microsoft Foundry. You'll plan and set up AI environments, select and deploy models from the model catalog, build apps with the Foundry SDK, use prompt flow, develop RAG solutions with your data, fine-tune models, apply responsible AI practices, and evaluate generative AI performance using Azure AI Studio tools.

Masterclass on Agentic AI Solutions Using Copilot Studio and AutoGen

This masterclass offers exposure to designing and deploying agentic AI solutions using modern low-code and open-source frameworks. Delivered through live sessions led by industry experts, it showcases how Copilot Studio and AutoGen can accelerate development and enable rapid deployment of agentic systems in real-world business environments.

**Please note: The program curriculum is subject to change to ensure it always reflects the latest, in-demand topics.*

Microsoft

Program Advisors

Ricardo Calix

Associate Professor,
Computer IT and Graphics

Ricardo Calix is an Associate Professor of Computer Information Technology and Graphics at Purdue Northwest. He has a Ph.D. in Engineering Science from Louisiana State University in 2011. His research areas include ML, biometrics, intrusion detection systems, and NLP.

Armando Galeana

Founder and CEO at
Ubhuru Technologies

A seasoned data science leader, with extensive experience in digital transformation. Throughout his career, Armando has leveraged his vast expertise in AI & ML to build infrastructure, create new lines of business and drive global implementations.

Amitendra Srivastava

Chief Data Scientist at
Intelytica

Amitendra's expertise lies in utilizing data analysis and machine learning techniques to solve complex business problems and drive strategic decisions. As Chief Data Scientist, he leverages the power of data to create value and drive innovation.

Raghav Goel

Generative AI & Data
Science Consultant

A passionate and successful corporate trainer who has delivered 150+ training sessions for corporates in India, Middle East, USA, and South East Asia for corporate clients like Publicis Sapient, KPMG, Capgemini, Coforge, ITC, DXC, Huawei, and IBM.

Program Trainers

Timothy Henize

18+ years of experience

AI & Machine Learning Engineer

Isaac Whittaker

9+ years of experience

Sr Applied Research Scientist

Bassel Dakhlallah

13+ years of experience

Data Analytics & Engineering Leader

Jaime Saldarriaga

9+ years of experience

AI & Machine Learning Engineer

Sohail Hosseini

15+ years of experience

DevOps Engineer

Edinson Medina

7+ years of experience

Senior Data Scientist

Simplilearn's Authorized Partner for Hong Kong Special Administrative Region (HKSAR)

PTI Professional Development Ltd.
20/F. Wellable Commercial Building,
513 Hennessey Road, Causeway Bay, Hong Kong
Website: www.the-pti.com