

AI Capability Pathway: Understand, Use, Build and Govern

FA-0813

Agentic & Generative AI · AI for Leaders & Business · Automation, RPA & Power Platform

Introduction

A four-level pathway separates AI awareness, business-user skills, introductory enterprise implementation and advanced engineering. Select the level that matches your responsibilities and prior experience rather than treating every participant as an AI developer.

Level 1 is approximately two hours; Level 2 and Level 3 are one day each. These introductory levels use concise explanations and selected prepared activities. Level 4 is a separately scoped advanced technical course without a fixed duration specified here.

The pathway covers prompting, context, reliability, information handling, governance, RAG, agents, APIs, automation, evaluation and lifecycle operations. Microsoft platform demonstrations depend on current feature availability, configured permissions and appropriate licences; preview capabilities are identified and are not assumed production-ready.

Learning outcomes

- Level 1: explain AI/generative-AI capabilities, limits and basic risks in everyday language.
- Level 2: structure prompts/context, review output and apply responsible business-use practices.
- Level 3: create a basic prepared agent/flow example and explain APIs, grounding, identity and architecture choices.
- Level 4: evaluate production-oriented AI architecture, RAG, tool-using/multi-agent patterns, lifecycle controls, observability and governance.

Prerequisites

Entry requirements differ by level. Level 1 has no technical prerequisites. Level 2 assumes basic AI awareness and workplace-application experience. Level 3 assumes generative-AI and general IT familiarity, with Microsoft 365/Power Platform and API experience beneficial. Level 4 requires equivalent implementation knowledge, cloud/API/authentication familiarity and preferably Power Platform/Azure and software-lifecycle experience.

Training outline

1. Level 1 – Awareness and demystification (approximately 2 hours)

AI literacy for nontechnical participants.

Level prerequisites

- None.
- No technical or programming knowledge is required.

Level learning goals

- Explain AI and generative AI in everyday language.
- Understand broadly how modern AI systems produce responses.
- Distinguish AI capabilities from common misconceptions.
- Recognize major benefits and limitations of AI.
- Identify basic privacy, security and reliability risks.
- Understand why human oversight remains important.
- Participate more confidently in organizational discussions about AI.

Curriculum

Understanding Artificial Intelligence

- What artificial intelligence means
- Traditional software versus AI systems
- Machine learning and neural networks
- Generative AI

- Large language models
- Multimodal AI
- AI assistants, copilots and agents

How Modern Generative AI Works

- Training and model knowledge
- Tokens and language prediction
- Prompts and responses
- Context
- Why AI appears intelligent
- Model knowledge, uncertainty and the limits of generated answers

Demystifying AI

- Intelligence versus prediction
- AI reasoning capabilities
- Creativity and generation
- Common AI myths
- Current limitations

Opportunities Created by AI

- Productivity
- Knowledge work
- Research and communication
- Automation
- Decision support
- Emerging workplace applications

Understanding AI Risk

- Hallucination
- Bias
- Privacy
- Confidential information
- Misinformation

- Cybersecurity misuse
- Employment and societal impact
- Overreliance on AI

The Wider AI Safety Discussion

- Geoffrey Hinton and modern AI development
- Near-term AI risks
- Long-term AI concerns
- Increasing model capabilities
- Human control and oversight
- Responsible adoption

Becoming AI Aware

- Questioning AI output
- Protecting organizational information
- Knowing when AI should not be used
- Human accountability
- Continuous AI literacy

2. Level 2 – Applied AI for business users (1 day)

Practical prompting, context, review and responsible departmental use.

Level prerequisites

- Level 1 or equivalent basic AI awareness.
- General experience using workplace applications.

Level learning goals

- Use generative AI more effectively in daily work.
- Construct clear and structured prompts.
- Understand context windows and context limitations.
- Apply context-engineering techniques.
- Validate and refine AI-generated output.
- Recognize confidential and inappropriate information.

- Identify useful AI opportunities within business departments.
- Understand personal and departmental AI assistants.
- Apply responsible-use and governance principles.
- Recognize prompt injection and other common AI threats.

Curriculum

Working Effectively with Generative AI

- AI conversation fundamentals
- Prompts and instructions
- Tasks, objectives and constraints
- Output requirements
- Iterative interaction
- Verification and refinement

Prompt Engineering

- Prompt structure
- Roles and perspectives
- Context
- Constraints
- Output formatting
- Prompt refinement
- Reusable prompt patterns

Context Engineering

- Understanding context windows
- Relevant versus irrelevant context
- Context quality
- Instruction hierarchy
- Document context
- Conversation context
- Context limitations

Improving AI Reliability

- Recognising and investigating possible hallucinations
- Verification
- Source checking
- Assumptions
- Ambiguous instructions
- Human review

AI for Business Productivity

- Human Resources
- Finance
- Operations
- Administration
- Sales
- Customer service
- Management
- Research and documentation

Working with Organizational Information

- Public information
- Internal information
- Confidential information
- Personal information
- Regulated information
- Intellectual property

Responsible AI Usage

- Bias
- Privacy
- Security
- Accuracy
- Accountability
- Automation bias

- Appropriate human oversight

AI Governance for Users

- Organizational AI policies
- Acceptable-use requirements
- Approved versus unapproved AI tools
- Shadow AI
- Data classification
- Auditability
- Compliance responsibilities

Personal and Departmental AI Assistants

- Personal AI assistants
- Department-specific assistants
- Instructions and boundaries
- Knowledge sources
- Permissions
- Human escalation

AI Security Awareness

- Prompt injection
- Indirect prompt injection
- Data leakage
- Malicious content
- Unsafe instructions
- Excessive trust
- Credential and secret exposure

3. Level 3 — Microsoft enterprise AI implementation (1 day)

Selected prepared implementation examples; not a complete advanced engineering programme.

Level prerequisites

- Level 2 or equivalent generative AI knowledge.

- General IT knowledge.
- Familiarity with Microsoft 365 or Power Platform is beneficial.
- Basic knowledge of APIs is beneficial.

Level learning goals

- Understand the components of Microsoft's enterprise AI stack.
- Create and configure a basic Copilot Studio agent using a prepared example
- Integrate agents with business actions and workflows.
- Understand Power Automate and agent-flow architectures.
- Consume AI capabilities through APIs.
- Understand Microsoft Foundry and model endpoints.
- Explain RAG architecture and enterprise grounding.
- Understand authentication and permission requirements.
- Compare different approaches to enterprise agent development.
- Apply basic security and governance controls.

Curriculum

Enterprise Generative AI Architecture

- Models
- Prompts
- Context
- Knowledge
- Tools
- Actions
- Agents
- Enterprise systems

Microsoft Enterprise AI Ecosystem

- Microsoft 365 Copilot
- Copilot Studio
- Power Automate
- Microsoft Foundry: selected current platform capabilities
- Azure AI/Foundry services: current naming and availability

- Azure AI Search
- Microsoft Entra ID
- Microsoft Purview

Copilot Studio Fundamentals

- Agent creation
- Agent instructions
- Knowledge sources
- Topics
- Generative orchestration
- Tools and actions
- Connectors
- Testing and publishing

Agent Flows and Power Automate

- Agent flows
- Power Automate cloud flows
- Workflow triggers
- Business-system connectors
- Agent actions
- Human approvals
- Returning results to agents

API-Based AI Integration

- REST API concepts
- AI endpoints
- Authentication
- HTTP actions
- Request and response handling
- Enterprise API integration

Microsoft Foundry

- Models and deployments

- Model endpoints
- Foundry projects
- Foundry Agent Service
- Tools and enterprise integrations
- Identity and access

Retrieval-Augmented Generation

- RAG concepts
- Enterprise knowledge
- Document ingestion
- Indexing
- Azure AI Search
- Vector and hybrid retrieval
- Grounded responses
- Citations

Selecting an Agent Architecture

- Copilot Studio agents
- Agent flows
- Power Automate
- API-driven solutions
- Foundry agents
- Architecture selection

Basic Agent Governance

- Authentication
- Authorization
- Least privilege
- Data loss prevention
- Connector controls
- Knowledge boundaries
- Human approval
- Logging and auditing

4. Level 4 — Advanced enterprise engineering and agentic systems

A separately scoped advanced technical track; duration determined separately.

Level prerequisites

- Level 3 or equivalent enterprise AI implementation knowledge.
- Familiarity with cloud architecture.
- Understanding of APIs and authentication.
- Power Platform or Azure experience is strongly recommended.
- Basic software-development lifecycle knowledge is recommended.

Level learning goals

- Design production-oriented enterprise AI architectures.
- Apply AI SDLC and lifecycle-management principles.
- Understand MLOps and GenAIOps practices.
- Engineer and evaluate enterprise RAG systems.
- Design secure tool-using agents.
- Apply identity and least-privilege architectures.
- Monitor AI quality, safety, usage, latency and cost.
- Manage prompt, model and knowledge changes.
- Understand agentic retrieval and advanced RAG.
- Design multi-agent architectures.
- Implement human-in-the-loop and approval controls.
- Establish production governance for enterprise AI.

Curriculum

AI Solution Architecture

- Enterprise AI architecture patterns
- Model layer
- Orchestration layer
- Knowledge layer
- Tool layer

- Identity layer
- Governance layer
- Observability layer

AI Software Development Lifecycle

- Requirements and risk assessment
- Design
- Development
- Evaluation
- Security testing
- Deployment
- Monitoring
- Continuous improvement
- Retirement

MLOps and GenAIOps

- Model lifecycle
- Prompt lifecycle
- Knowledge lifecycle
- Configuration management
- Versioning
- Environment promotion
- CI/CD concepts
- Rollback strategies

AI Change Management

- Model changes
- Prompt changes
- Knowledge changes
- Tool changes
- Connector changes
- Permission changes
- Regression testing
- Production approvals

Advanced RAG Engineering

- Chunking strategies
- Embeddings
- Vector retrieval
- Hybrid retrieval
- Metadata filtering
- Reranking
- Query transformation
- Agentic retrieval and version-specific capabilities, including preview query planning where applicable
- Knowledge-base architecture
- Permission-aware retrieval

RAG Governance and Security

- Source authorization
- Document permissions
- Sensitive information
- Data leakage
- Index security
- Knowledge freshness
- Assessing groundedness; retrieval does not eliminate incorrect answers
- Retrieval auditing

Advanced Agent Engineering

- Tool-using agents
- Planning and reasoning
- Structured tool invocation
- API tools
- MCP-based integration and explicit identity/action boundaries
- Agent memory
- State management
- Action boundaries

Agent Security

- Direct prompt injection
- Indirect prompt injection
- Tool poisoning
- Data exfiltration
- Excessive agency
- Privilege escalation
- Identity propagation
- Credential management

Multi-Agent Systems

- Specialized agents
- Supervisor agents
- Delegation
- Task routing
- Agent-to-agent communication
- Shared knowledge
- Shared state
- Permission boundaries
- Failure management

Deterministic Automation versus Agentic Automation

- Traditional workflows
- AI-assisted workflows
- Autonomous agents
- Agent flows
- Power Automate
- Copilot Studio
- Foundry agents
- Architecture trade-offs

AI Evaluation

- Response quality

- Groundedness
- Relevance
- Safety
- Retrieval quality
- Tool-selection accuracy
- Agent task completion
- Regression evaluation

AI Observability and Monitoring

- Usage telemetry
- Agent activity
- Model consumption
- Token consumption
- Cost monitoring
- Latency
- Failure rates
- Quality degradation
- Security events

Enterprise AI Governance

- Environment strategy
- Identity and RBAC
- Data loss prevention
- Connector governance
- Microsoft Purview
- Microsoft Entra
- Audit requirements
- Risk classification
- Production controls

Production AI Operations

- Development-to-production promotion
- Testing gates
- Release management

- Incident management
- Rollback
- Usage governance
- Service ownership
- Continuous evaluation