www.logicalimagination.com
800.657.1494

# Designing, Building & Governing AI Assistants with Python

**Course #:** AI-204 **Duration:** 3 days

## Prerequisites

Completion of AI Foundations & Risk for IT Professionals, Designing Reliable AI Workflows & Interactions, AI Architecture & Agentic Systems, and Building MCP-Based AI Systems with Python, or equivalent experience designing and implementing AI-enabled systems using Python.

## Details

This capstone course brings together AI architecture, workflow design, MCP-based implementation, and operational governance into a single, end-to-end experience. Participants design and build a production-ready AI assistant that uses structured tools and backend services, then deploy, monitor, and govern that assistant over time. Rather than focusing on demos or chatbots, the course emphasizes assistant behavior, boundaries, reliability, accountability, and lifecycle management, ensuring AI assistants remain trustworthy and supportable after initial deployment.

After attending this course, students should be able to:

Design AI assistants with clear roles, scope, and behavioral boundaries
Implement a tool-using AI assistant backed by MCP-based services
Manage multi-turn interactions, confirmations, and error handling
Deploy AI assistants responsibly into organizational environments
Establish ownership, governance, monitoring, and change-management practices
Evaluate assistant effectiveness and retire or evolve assistants safely

This course is designed for technical professionals responsible for delivering AI assistants that must be functional, reliable, secure, and sustainable in real organizational environments. This course is hands-on and implementation-focused, with governance and operational considerations integrated throughout.

## Software Needed

Participants must have a laptop or desktop computer (Windows, macOS, or Linux) with Python 3.10 or later installed, a modern web browser, and reliable internet access. The ability to create and activate Python virtual environments, install packages, and run local development services is required. Access to an organization-approved AI assistant runtime that supports tool integration (such as an MCP-compatible client) is required. Participants should follow all organizational security, privacy, and confidentiality guidelines when building, testing, or discussing AI assistants.

## Outline

Designing, Building & Governing AI Assistants with Python

- **What an AI Assistant Really Is**
  - Assistants as persistent system actors
  - How assistants differ from workflows and agents

    - Why assistant design affects trust and risk
- **Defining Assistant Purpose, Role, and Scope**
    - Clarifying what the assistant is responsible for
    - Explicitly defining what the assistant must not do
    - Aligning assistant behavior with organizational needs
- **Designing Assistant Behavior**
    - System instructions and role definition
    - Managing tone, clarity, and expectations
    - Avoiding misleading authority and overconfidence
- **Wiring the Assistant to MCP Tools**
    - Connecting the assistant to MCP servers
    - Selecting appropriate tools for assistant use
    - Enforcing boundaries through schemas and validation
- **Multi-Turn Interaction Design**
    - Clarification and confirmation patterns
    - Managing state across interactions
    - Handling incomplete or ambiguous requests
- **Error Handling and Safe Failure**
    - Anticipating assistant failure modes
    - Designing recovery and fallback behavior
    - Preventing silent failures
- **Testing Assistant Behavior**
    - Functional testing of assistant flows
    - Edge cases and misuse scenarios
    - Validating assistant responses and actions
- **Deployment Models for AI Assistants**
    - Local, internal, and enterprise deployments
    - Configuration, secrets, and permissions
    - Managing environments and access
- **Observability and Monitoring**
    - Logging assistant actions and decisions
    - Detecting drift and degradation
    - Identifying misuse and unexpected behavior
- **Ownership and Accountability**
    - Defining who owns an AI assistant
    - Human accountability for assistant behavior
    - Escalation and intervention models
- **Governance and Risk Management**
    - Aligning assistants with policy and compliance
    - Managing data access and privacy
    - Responding to incidents and near misses
- **Change Management and Versioning**
    - Updating prompts, tools, and behavior safely
    - Communicating changes to users
    - Avoiding breaking changes
- **Scaling and Managing Multiple Assistants**
    - Preventing assistant sprawl
    - Standardizing governance without blocking innovation
    - Platform approaches to assistant management
- **Evaluating and Retiring Assistants**
    - Measuring usefulness and trust
    - Deciding when to redesign or retire an assistant
    - Lessons learned from real deployments