

Building and Deploying AI Agents in Microsoft Copilot Studio: Copilot Chat, Standard, and GitHub Copilot Harnesses

This is a proprietary, advanced course covering the entire Microsoft Copilot Studio platform – all three harnesses (Copilot Chat, standard, and GitHub Copilot) – together with the operational layer that determines whether an agent ever makes it to production. This is what sets it apart from the two 2-day courses, each of which focuses on a single approach: process-driven (standard harness) or goal-driven (GitHub Copilot harness).

Participants start by understanding what an agent is and what the Microsoft agent ecosystem looks like, and then build on each harness in turn: they configure an agent in Microsoft 365 Copilot and Agent Builder, move on to the agentic reasoning loop with Instructions, Skills, Knowledge, Tools, and Workflows, and then to the standard harness with topics, generative orchestration, multi-agent systems, event triggers, and Adaptive Cards. The course closes with an operations module: responsible AI and content moderation, application lifecycle management (ALM), and AI FinOps – Copilot Credits, consumption estimation, and cost control.

The key competency this course delivers is choosing the right harness deliberately before you start building, and being able to justify that choice with business requirements: predictability, channel, required message wording, multi-step reasoning capability, and cost. Changing the harness after the fact means a rebuild, not a refactor – which is why the decision must be made consciously.

The course consists of 6 modules and 13 step-by-step labs (approx. 9 hours of hands-on work).



Training recipients

This course is intended for:

- AI solution architects and consultants, who design Copilot Studio implementations and must deliberately choose the harness, orchestration pattern, and cost model for a specific process.
- Experienced Power Platform developers and low-code makers, who have already built an agent using one approach and want to learn the others and how to combine them.
- IT professionals and Microsoft 365 and Power Platform administrators, responsible for environments, permissions, data loss prevention (DLP) policies, publishing, solution lifecycle, and capacity management.
- Technical leads and AI product owners, who are responsible for moving pilots into production – including governance, quality evaluation, and budget.
- Trainers and Center of Excellence (CoE) teams, who build internal agent development standards for their organization.



Benefits

After completing this course, participants will be able to:

- Choose the right harness – decide between the Copilot Chat, standard, and GitHub Copilot harnesses based on requirements rather than habit, and justify that decision to the business.
- Work across the entire platform – build agents in Agent Builder, in the agentic reasoning loop, and in the classic topic-based model, learning the strengths and weaknesses of each approach.
- Master Instructions and Skills – apply the “always vs. sometimes” principle and progressive disclosure, and write descriptions that actually invoke the right component.
- Connect Knowledge, Tools, and Workflows – connect the agent to enterprise data (SharePoint, Dataverse, connectors, MCP) and automate event-driven processes with the new Workflows experience.
- Build multi-agent systems – design a team of specialists (child agents and connected agents) using the hub-and-spoke pattern, with deliberate context passing.
- Build autonomous agents – run an agent from an event trigger and hand the result over to a human as an interactive Adaptive Card in Microsoft Teams.
- Apply responsible AI and governance – layer three content-blocking mechanisms, write a complete AI disclosure, and place the agent in the appropriate risk zone.
- Implement ALM and measurable quality – move a solution between environments using solutions, environment variables, and connection references, and prove improvements with evaluations rather than gut feeling.
- Control costs (AI FinOps) – understand Copilot Credits, price a conversation, and learn why fewer turns beat a cheaper model.

In practice, you will leave the course able to design, build, secure, deploy, and cost an agent in any of the three approaches – rather than knowing only one of them



Training program

Module 1 – Microsoft 365 Copilot and the Copilot Chat Harness

- Why AI deployments stall; the operating model and the Frontier Firm; Microsoft 365 E7 licensing
- What an agent is: orchestrator, model, memory, knowledge, tools, skills, and governance; the agent spectrum and five types of agents
- The Microsoft agent ecosystem; choosing the harness before you build; orchestration patterns and structured reasoning
- Microsoft 365 Copilot: Work IQ, memory, model selection, Agent Store, and the Researcher and Analyst agents
- Copilot Chat harness: orchestration loop, component model, first-party capabilities
- Four ways to build: Agent Builder, Copilot Studio, Microsoft 365 Agents Toolkit, Work IQ Dev Tools; MCP servers and MCP Apps
- Labs: M365 Copilot and Frontier Agents and Build AI Assistants with Agent Builder

Module 2 – GitHub Copilot Harness: Instructions and Skills

- Why the agentic reasoning loop (thought – action – observe – decide) closes the gaps of classic orchestration
- Component model: instructions, knowledge, tools, skills, memory, agent sandbox, connected agents
- Knowledge retrieval: search, ranking, and downloading full files into the sandbox; per-user and per-agent memory
- Migrating from the standard harness: component mapping, rebuilding instead of one-to-one porting
- Instructions: what always applies vs. what descriptions decide; names, inputs, outputs, and dynamic chaining
- Skills: anatomy of a SKILL.md package, progressive disclosure, the determinism dial from goal to Python script
- Labs: The Agentic Reasoning Loop and Deep Dive: Skills

Module 3 – GitHub Copilot Harness: Knowledge, Tools, and Workflows

- Knowledge sources: SharePoint and OneDrive, files, Dataverse, public websites, Microsoft 365 Copilot connectors, Azure AI Search
- Retrieval pipeline: query rewriting, retrieval, summarization – with validation at every step; security trimming
- Seven types of tools: prebuilt and custom connectors, REST APIs, MCP servers, agent flows, prompts, computer use
- Connector vs. REST API vs. MCP decision; RPA vs. computer use; combining multiple tools in a single agent
- Workflows: event-driven automation, trigger and payload, building blocks, determinism vs. AI on a

single canvas

- Agent node: inline vs. referenced; human in the loop; testing nodes, Activity and Monitor
- Labs: Copilot Studio Tools and Workflows

Module 4 – Standard Harness: Orchestration and Multi-Agent Systems

- When the standard harness is the right choice: mandated predictability, channel constraints, required message wording
- Component model: topics, knowledge, tools, variables, agents, channels, analytics
- Topics: triggers and nodes; variables and their scope (topic, global, cross-session)
- Multi-agent systems: child agents vs. connected agents; hub-and-spoke, pipeline, and collaborative patterns
- Passing context between agents and instruction best practices for an agent team
- Classic vs. generative orchestration; inputs, outputs, and dynamic chaining; debugging incorrect routing
- Labs: Create the Agent, Adding the Child Agent, Adding the Connected Agent

Module 5 – Standard Harness: Event Triggers and Adaptive Cards

- Topic trigger vs. event trigger; event sources in Microsoft 365, Dataverse, and on a schedule (recurrence)
- Payload: everything the agent knows at the start; filtering on the trigger instead of in the flow
- Where autonomy pays off and where it is just noise; identity, connection references, DLP, and publishing
- Adaptive Cards as a data structure: body, elements, actions, and schema version
- Live data binding – one card definition for every record; designing and debugging cards
- Labs: Automating Event Triggers and Notify a Teams Channel with an Adaptive Card

Module 6 – Operations: Responsible AI, ALM, and Cost Management

- The six Microsoft Responsible AI principles; distinguishing AI safety, security, and governance
- Three independent content-blocking mechanisms; complete AI disclosure; prompt injection and Agent Runtime Protection
- The zone model (green, yellow, red): capabilities vs. potential blast radius
- Why ALM for AI is a loop, not a line; solutions, environment variables, and connection references
- Pipelines in Power Platform, source control, and component collections; analytics vs. evaluations; test coverage zones
- AI FinOps: Copilot Credits, conversation pricing, purchasing models, three layers of control, and the “rediscovery tax”
- Labs: Responsible AI and content moderation (classic or new experience version to choose from) and Monitor Performance and Evaluate Agent Quality



Expected preparation of the participant

- Hands-on experience with Microsoft 365 services (Teams, SharePoint, Outlook) at an advanced user

level.

- Basic knowledge of Microsoft Power Platform: environments, solutions, and connectors.
- General understanding of artificial intelligence and large language model (LLM) concepts – tokens, context window, retrieval-augmented generation (RAG).
- Prior experience building at least one agent or flow (Power Automate, Copilot Studio, Agent Builder) is an advantage.
- No programming experience is required – the course follows a low-code / no-code approach.
- Completion of AI-901 or PL-900, or equivalent knowledge, is an advantage.



Training Includes

- manual in electronic form available on the platform: <https://www.altkomakademia.pl/>
- access to Altkom Akademia's student portal

Training method:

- Lecture: 30%
- Workshops: 25%
- Labs: 45%



Duration

3 days / 21 hours

Language

- Training: English
- Materials: English