

AI Agent Governance in Microsoft 365: Secure, Govern, and Scale

Duration: 8 x 90 min.

Audience: Dirigeant et cadre|Professionnel TI

Who is this training for?

- Administrators
- Solution and enterprise architects
- Governance, security, and compliance leaders

Working in the public or education sector? Discover this training program: [AI Agents in the Public and Education Sectors: Ensuring Responsible Governance in Microsoft 365 - Training Courses | Afi U.](#)

Training objectives

By the end of the training, participants will be able to:

- Explain the hosting and operating architecture of AI agents within Microsoft 365 and Power Platform
- Design a coherent environment strategy (POC, pilot, production) and govern the use of Managed Environments
- Document a secure AI agent deployment model based on selected channels
- Apply rigorous governance of identities, roles, permissions, and consent following the principle of least privilege
- Structure data, connector, and flow governance using appropriate DLP policies
- Configure and interpret audit, compliance, traceability, and monitoring mechanisms
- Establish ALM, operational, and cost-control practices (licensing and PAYG) to support AI agent industrialization

Summary

With the introduction of AI agents across Microsoft 365, Teams, SharePoint, and Power Platform, organizations are entering a new phase: **orchestration rather than experimentation**.

Business-led experimentation alone is no longer sufficient. A **structured governance framework** has become essential. Administration teams must now establish the conditions for a **controlled, secure, and sustainable** deployment of AI agents: environment readiness, access management, data governance, guardrails, usage traceability, and cost management. This learning path is specifically designed for teams responsible for these activities.

It is not intended to train AI agent developers. Instead, it equips **administrators, architects, and governance, security**, and compliance leaders with the tools needed to structure, secure, and scale AI agent usage across the Microsoft ecosystem.

Course outline

Session 1 – AI Agent Architecture and Operating Framework

- Overview: Microsoft 365, Teams, SharePoint, Entra ID, Power Platform, Copilot Studio, Purview
- AI agent typologies (productivity, data-assisted, public-facing) and deployment models
- Target architecture principles for POC, pilot, and production

Session 2 – Identity, Roles, Permissions, and Consent

- Administrative roles (tenant, environment, Dataverse) and separation of duties
- User vs. admin consent and application permission review
- Conditional Access, MFA, Zero Trust, and least privilege

Session 3 – Environment Strategy and Managed Environments

- Default Environment: role, risks, and governance considerations
- Dev/test/prod/sandbox separation, security groups, and environment templates
- Managed Environments: control, insights, and enterprise-scale governance

Session 4 – Data, Connectors, and DLP Policies

- Backends (SharePoint, Dataverse, Azure SQL) and source validation

- Standard, premium, and custom connectors and data-mixing risks
- Power Platform DLP (Business/Non-Business/Blocked) vs. Purview DLP

Session 5 – Deploying a Copilot Studio agent on the appropriate channels

- Publishing and republishing processes and channel-specific prerequisites
- Deployment on Teams, SharePoint, websites (demo vs production), and Microsoft 365 Copilot
- Authentication models (Microsoft, manual, none) and their implications
- Licensing, consumption limits, and operational considerations by channel

Session 6 – Compliance, audit, traceability, and monitoring

- Microsoft Purview Audit and Unified Audit Log
(activation, key events, search scenarios)
- Data classification, sensitivity labels, retention, and eDiscovery
- Log forwarding to Microsoft Sentinel and SOC integration strategies
- Monitoring and oversight practices for agent usage and behavior

Session 7 – Center of Excellence (CoE), tenant hygiene, lifecycle, and ALM

- Power Platform CoE Starter Kit
(inventory, visibility, adoption, and governance)
- Tenant hygiene practices
(orphaned resources, inactivity policies: 30/60/90 days)
- Application Lifecycle Management (ALM) foundations
(Dataverse solutions and promotion to managed environments)

Session 8 – Licensing, PAYG, budgeting, operations, and readiness review

- Trials vs subscriptions and Copilot Studio PAYG models
- Capacity consumption, usage tracking, and cost visibility
- Budgeting and alerts for AI agent operations
- Go / no-go checklist and criteria for POC industrialization

Approach and methodology

The program is built around an **AI agent governance tracking tool**, used as the pedagogical backbone throughout the course.

This tool serves as a shared reference to structure learning, document key decisions, and track AI agent evolution (POC, development, testing, production). The tracking tool enables participants to:

- Reduce cognitive load by centralizing concepts, decisions, and reference points in a single workspace
- Foster learning by doing, with each session enriching the tool through real-world scenarios and hands-on exercises
- Make governance decisions visible and traceable (environments, data, security, channels, lifecycle)
- Support workplace transfer, as the tool remains usable after the training as an operational and decision-support asset

An individual **quiz** is administered at the end of each session to validate learning outcomes and ensure pedagogical traceability.

Prerequisites

Read-only (minimum) access to the customer tenant or a controlled demonstration environment

Recommendations

- Strong working knowledge of Microsoft 365 and Entra ID
- General understanding of Power Platform
- Familiarity with security principles (roles, permissions, least privilege)