

AI Agents in the Public and Education Sectors: Ensuring Responsible Governance in Microsoft 365

Duration: 9 X 90 min.

Audience: Dirigeant et cadre|Professionnel TI

Who is this training for?

Administrators, architects, and governance, security, and compliance leaders working in the public or education sectors

Do you work in another industry? Take a look at this training program: [AI Agent Governance in Microsoft 365: Secure, Govern, and Scale - 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 arrival of AI agents in Microsoft 365, Teams, SharePoint, and the Power Platform, public sector organizations and educational institutions are entering a new phase: one of orchestration, where responsibility, compliance, and trust are essential.

User-led experimentation is no longer sufficient. A structured governance framework has become indispensable.

Administration teams must now create the conditions for deploying AI agents in a way that is secure, compliant, and aligned with the specific requirements of the public and education sectors: protection of sensitive data, access management, usage governance, traceability, auditability, and risk management. This cohort is specifically designed for the teams carrying these responsibilities within ministries, public agencies, and educational institutions.

This program is not intended to train AI agent developers.

Instead, it empowers administrators, architects, and leaders in governance, security, compliance, and privacy to design and implement a rigorous, transparent, and sustainable framework for the responsible use of AI agents.

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
-

Session 9 – Governed deployment of AI agents in high-constraint environments

- Governance challenges of AI agents in regulated, large-scale environments
- Segmentation of access and usage across multiple user populations
- Channel selection strategies and associated risk management
- Data governance, compliance requirements, and organizational obligations
- Licensing constraints, cost management, and operational sustainability
- Operational synthesis and decision criteria prior to large-scale deployment

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)