zenity slalom

Compliance Readiness for Adopting AI Agents

Table of Contents

zenity slalom

Compliance Readiness for Adopting AI Agents

AI agents are being embedded across all areas of business, touching sensitive workflows, regulated data, and performing actions autonomously. As agents proliferate, enterprises must also ensure they are in compliance with organizational and legislative policies. To protect the institution and satisfy regulators, financial services organizations need a structured, compliance-first approach to AI agents that does not get in the way of innovation.

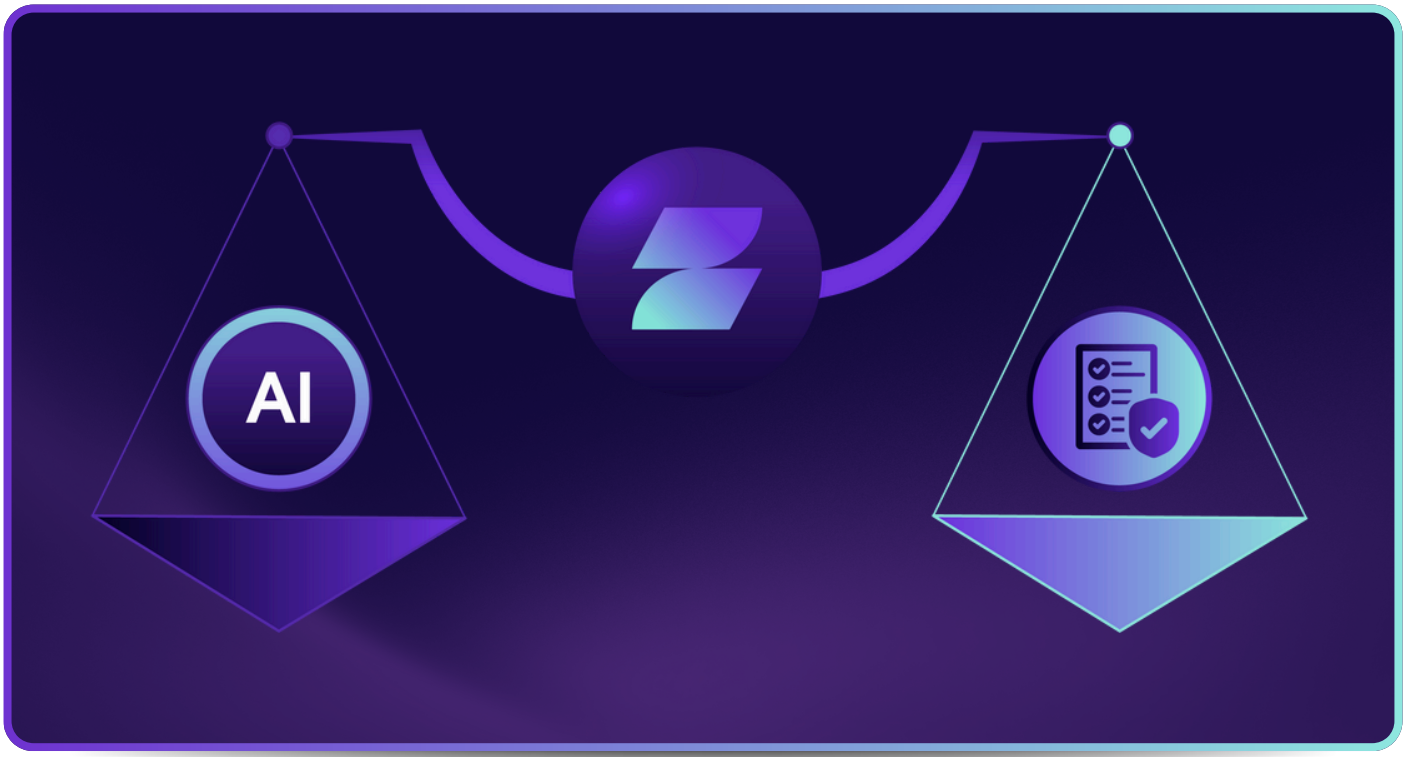

How Microsoft Agentic Platforms Are Used

Understanding the nuances of different agentic platforms; personas, attack surfaces, and risk profiles, helps lay the foundation for compliance readiness. Consider three Microsoft platforms, Microsoft Foundry, Copilot Studio, and Microsoft 365 Copilot.

Platform	Users	Use Cases
 Microsoft 365 Copilot	• Across the workforce for productivity: summarizing emails, drafting documents, analyzing spreadsheets, and surfacing insights from SharePoint, Teams, and OneDrive.	• Content generation, meeting and email summaries, data pull from documents, and decision support.
 Copilot Studio	• "Business developers" in operations, service, and line-of-business teams to create task-focused copilots and chatbots.	• Customer service flows, internal helpdesk, process automation, and workflow orchestration using low-code tools.
 Microsoft Foundry	• Engineering, data science, and platform teams to build highly customized, code-first agents and copilots.	• Complex workflows, integration-heavy automations, domain-specific copilots that touch core banking systems, data warehouses, and risk engines.

Each platform varies dramatically in terms of who uses them, and the types of agents that emerge, and enterprises should consider which policies, controls, and guardrails can help these agents remain compliant and be secure, without hindering velocity.

Leverage Existing Guidance and Guardrails

In understanding the different platforms and the types of agents that are built on them, we can dive into applying lessons from regulation and compliance that fit the profile for each platform. Regulators have made it clear: existing expectations apply whether or not "AI" is explicitly mentioned. There are four assumptions to hold when thinking about how compliance mandates should be met in regards to agentic adoption

Assumption 1:

AI agent usage is in-scope for existing guidance on models, automation, and third-party risk.

Assumption 2:

Existing guidance generally applies to financial institution activities regardless of AI use. Treat agents as extensions of existing models, decision engines, and automated systems.

Assumption 3:

Even when guidance does not explicitly reference AI or agents, its principles, governance, validation, documentation, controls, and testing map directly to AI agent lifecycle risks.

Assumption 4:

There is an open opportunity (and expectation) to clarify how AI fits within existing risk programs, rather than waiting for bespoke "AI-only" rules.

These assumptions mean model risk management (MRM), operational risk, third-party risk, and governance frameworks should explicitly include agentic use cases across Foundry, Copilot Studio, and 365 Copilot.

Understand Key Regulatory Bodies and Expectations

Behind these assumptions lie a variety of regulatory bodies that carry with them certain expectations on how enterprises should be securing and governing agents. In the United States, the FDIC, FRB, and OCC jointly conduct horizontal cybersecurity reviews of the eight U.S. global systemically important banks as part of an Interagency Coordinated Cybersecurity Review program to support effective cybersecurity supervision across these systemically important financial institutions.

In practice, each institution's chartering regulator, in conjunction with FFIEC guidance, is the primary authority assessing your cybersecurity.Here are a few to be cognizant of:

Federal Deposit Insurance Corporation (FDIC):

The FDIC is the U.S. federal agency that insures customer deposits at banks and ensures safety, soundness, consumer protection, and cybersecurity in the banking system

Federal Reserve:

The Federal Reserve is the U.S. central bank responsible for monetary policy, financial stability, and supervising many banks and bank holding companies, including how they manage model risk, operational resilience, and cybersecurity as they adopt and secure AI-driven systems.

Office of the Comptroller of the Currency (OCC):

The OCC is the U.S. regulator that charters, regulates, and supervises national banks and federal savings associations, setting expectations for safety and soundness, model governance, third-party risk management, and robust controls over AI and other technologies used in banking operations.

Consumer Financial Protection Bureau (CFPB):

The CFPB creates rules to ensure transparency, accuracy, and fairness in financial products and can take action and is looking for organizations to prove that agents meet fairness, explainability, and consumer protection obligations.

Securities and Exchange Commission (SEC):

The SEC is concerned with any tool that touches investor decisions, disclosures, trading, or client interactions; much of which is now being handled and augmented by agents.

National Credit Union Administration (NCUA):

The NCUA plays a role similar to the SEC only for credit unions, and its relationship to AI agents is focused on safety, soundness, cybersecurity, consumer protection, and compliance, but not within investment markets. The NCUA regulates federally insured credit unions and oversees how they deploy technology, protect consumer data, and manage operational risk.

Due to the importance of financial services institutions and the role they play in our society, there are many organizations and agencies that provide frameworks and tools that FSIs can use as a way of making sure that agents are protected and safe. When an FSI uses AI agents, whether it be for member service, lending, underwriting, internal operations, etc., these organizations deeply care about how those agents affect risk, compliance, and member protection and enterprises must be able to prove that their agents are in-line with those requirements. Two specific examples of guidance for how FSIs use and adopt agents:

1. The FDIC, Federal Reserve, OCC) issued model risk management (e.g., SR 11-7) and third-party risk management guidance that can apply when AI agents influence decisions, recommendations, or customer outcomes. These principles for model risk management, meaning sound development, validation, monitoring, and governance, align with NIST's AI-related practices and can be extended to AI agents used for credit, fraud, AML, and operational decisions. The main takeaway is that firms must tailor application of these frameworks to specific AI use that are inherently tied to the takeaways in our earlier section, namely that a Copilot summarization of a meeting is not the same as Foundry-based decisioning agent tied to underwriting.

2. NIST's AI and cybersecurity guidance reinforces the need for documentation, testing, monitoring, and control over AI behavior, which maps directly to agent design and runtime oversight. This is meant to apply a consistent framework across all agentic platforms to prevent fragmented practices, audit gaps, and weak links between "shadow" and sanctioned AI.

Recent enforcement makes clear that regulators care about outcomes and governance, and that whether a system is internally referred to as an agent, an automation, a copilot, or an LLM, that they must be safe, controlled, and compliant. Particularly when paired with emerging AI-focused laws and regulatory commentary, these actions signal that FSIs must:

- Have robust governance, documentation, and testing around AI agents that influence decisions or customer treatment.
- Accurately represent AI capabilities and limitations in disclosures and marketing.
- Ensure agents do not bypass needed controls (verification, approvals, human-in-the-loop) just because the experience is "copilot-driven."

Aligning Key Stakeholders Around Agent Compliance

A successful compliance readiness program for agents across Foundry, Copilot Studio, and 365 Copilot is cross-functional by design.

CIO / Head of IT & Platforms

- Owns environment strategy and standardization across Microsoft tenants and AI platforms.
- Responsibilities include:
  - Define where and how agents can be built and run (dev/test/prod environments, sandboxes).
  - Ensure central visibility into all agents, their configurations, and integrations.
  - Implement scalable management patterns (templates, guardrails, and platform services) to prevent fragmentation.

Operations & SRE / Reliability Teams

- Focus on operational excellence of agents.
- Responsibilities include:
  - Observability: logs, metrics, traces, and event streams for agent runs and failures.
  - Data retention and audit trails for regulatory review, customer dispute resolution, and incident analysis.
  - Test automation and regression testing for agent behaviors, prompts, and toolchains.
  - High availability and graceful degradation for business-critical agentic workflows.

AI Organization (CAIO, AI CoE)

- Owns responsible AI strategy and value realization.
- Responsibilities include:
  - Catalog and understand agent and copilot usage across Foundry, Copilot Studio, and 365 Copilot.
  - Define ROI metrics per use case (time saved, error reduction, revenue impact) and evaluate pilots versus production.
  - Ensure “responsible innovation”: create patterns and guardrails for agent creation, tool integrations, memory usage, and data sources.
  - Govern billing and cost management by linking agent consumption to ownership and business value.

CISO and Security Organization

- Own the security and compliance posture. :
- Responsibilities include:
  - Security posture management for agents: least privilege access, identity management, and secure configurations.
  - Threat detection for prompt injection, data exfiltration, misuse of tools, and agent hijacking.
  - Data protection and privacy: control over what data agents can access, store, and share.
  - Overarching compliance: ensure that AI agent use aligns with model risk policies, cybersecurity rules, privacy laws, and sector-specific regulations.

Phased Approach to Compliance Readiness

Agents are clearly complex systems that have a set of complex compliance considerations that need to be carefully managed to reduce the risk of non-compliance. A pragmatic roadmap for how FSIs can adopt agents safely across Microsoft platforms could look something like this:

Phase 1

Controlled Experimentation (Sandboxes with Guardrails)

- Stand up dedicated dev and pilot environments in various agentic platforms
- Allow selected teams to experiment with clear rules and boundaries:
  - Authentication and authorization standards (who can build, publish, and consume agents).
  - Basic sharing rules for data and agent access (no sensitive data in early-stage experiments).
  - Initial policies for which tools and data sources can be invoked.
- Define initial ROI metrics (time saved, volume of tasks automated, reduction in manual steps) and require simple documentation of each experiment.

Phase 2

Structured Pilot Programs

- Promote successful prototypes into more robust, monitored pilots.
- Apply stronger compliance and security controls:
  - Formal review of agent purpose and intent, data sources, and tool access.
  - Alignment with model risk and third-party risk frameworks where applicable.
  - Enhanced observability with logging, audit trails, and incident workflows.
- Expand authentication and sharing rules: role-based access, group scoping, and environment-specific policies.
- Introduce standard patterns and templates for agent builds such as approved toolsets for finance employees vs. operations teams

Phase 3

Scaled, Governed Adoption

- Enable broader access across business units, with tiered environments per risk level.
- Institutionalize governance processes:
  - Regular governance evaluation of agent inventory, risk posture, and policy adherence.
  - Fine-tuning of rules as regulators publish new guidance and internal findings emerge.
- Implement continuous protection:
  - Runtime monitoring for anomalous or unsafe behavior.
  - Ongoing configuration reviews, access audits, and guardrail updates.
- Tie agent portfolios to formal compliance mappings (e.g., how specific agent classes align with model risk guidance, privacy obligations, sector rules).

Phase 4

Continuous Compliance and Optimization

- Treat agentic AI as a living program, not a one-off project.
- Integrate feedback loops from audits, incidents, and metrics into design, development, and operations.
- Maintain living documentation of:
  - Where agents run (Foundry, Copilot Studio, 365 Copilot).
  - What they do, what data they touch, and who is accountable for making sure agents stay compliant, even as they evolve.
  - How controls map to relevant regulations and internal policies.
- Optimize for both protection and value: refine ROI metrics, retire low-value agents, double down on high-value ones—with compliance as a default constraint, not an afterthought.

By treating AI agents across the Microsoft ecosystem as governed, observable, and continuously managed systems, financial institutions can meet regulatory expectations, reduce risk, and still unlock the transformational upside of agentic AI.

To learn more about how to ensure agents are compliant, visit us at https://www.zenity.io