TeamsFox

MICROSOFT ENABLES AI.

WE GOVERN IT.

The IT Leader's Guide to AI Governance in Microsoft 365

Executive Summary

AI agents are already inside your Microsoft 365 tenant. Copilot Studio has been generally available since 2024. Power Automate flows connected to AI models are active across thousands of organisations. Business users are building automated tools without IT involvement, and often without IT awareness.
The governance frameworks to manage this have not kept pace.

Microsoft has introduced Agent 365 (GA May 2026) and expanded Purview and Defender controls. These are useful starting points, but they address the surface of the problem. They provide a control plane for agents.
They do not fix the data those agents access, the permissions they inherit, or the costs they generate.

This paper sets out what the AI governance problem actually looks like inside a real Microsoft 365 environment, what the risks are, and what a practical governance approach looks like. The competitor landscape is covered separately in the appendix for reference.

Find us | www.teamsfox.com

Key numbers
at a glance

01 | **Microsoft Security Blog,** Feb 2026:
80% of Fortune 500 companies are already running active AI agents

02 | **Gartner,** Apr 2026:
by 2028, the average Fortune 500 enterprise will run more than **150,000 AI agents**, up from fewer than 15 in 2025

03 | **Gartner:**
60% of AI projects will be abandoned due to missing or inadequate AI-ready data

04 | **Gartner:**
63% of organisations either do not have AI-ready data or are not sure if they do

05 | **Gartner**
by 2027, **40%** of enterprises using consumption-priced AI will see costs more than double their budget

06 | **KPMG:**
91% of leaders say data security, privacy, and risk will influence their AI strategy over the next six months (Q1 2026 AI Quarterly Pulse Survey)

07 | **Gartner:**
more than **40%** of agentic AI projects will be canceled by the end of 2027, due to escalating costs, unclear business value, or inadequate risk controls

80%
Fortune 500 running active AI agents

60%
AI projects fail due to poor data foundations

40%+
Agentic AI projects to be canceled by 2027

Why AI Governance **Cannot Wait**

There is a temptation to treat AI governance as a future problem. AI agents are new, they are complex, and most organisations are still working out how to use Microsoft Copilot. Agents feel like the next problem after this one.
That framing is wrong. The window to establish governance discipline is not before the rollout. It is now, during it. Every week that agents operate without visibility, lifecycle management, and access controls is a week of compounding risk. Ungoverned agents accumulate permissions. Shadow AI spreads. Sensitive data gets indexed by AI systems that were never meant to see it.

The organisations that move now and establish governance foundations before AI scales will have a real advantage over those trying to retrofit controls onto an already complicated environment.

The scale of the problem

Faster than IT can track

Microsoft's agentic AI platform is expanding faster than most IT teams can track. Copilot Studio lets any user build a custom agent from templates and connect it to SharePoint, Teams, Dataverse, and hundreds of third-party services. Power Automate flows with AI models are already live across thousands of tenants. Agent 365, Microsoft's new centralised control plane, went to general availability on May 1, 2026, at $15 per user per month.
Gartner predicts that 40% of enterprise applications will include task-specific AI agents by 2026, up from less than 5% in 2025. It is not a gradual trend. It is a step change, and it is happening in environments where the governance infrastructure was not designed for non-human identities acting autonomously.

Not coming from IT-led rollouts

Much of this growth is not coming from IT-led rollouts. Microsoft's own Cyber Pulse research found that 29% of employees had already used an unsanctioned AI agent for work tasks in early 2026, and Gartner estimates that 41% of employees now qualify as citizen developers, building automations and agents without a formal development background. Organisations without a Centre of Excellence or equivalent governance function see agent sprawl at three to four times the rate of those that have one. The agents most likely to slip past IT are not the ones Microsoft ships. They are the ones a business user builds on a Friday afternoon.

A fragmented environment

At the same time, the operational burden is already real before a single AI agent is added. Organisations now manage an average of 12 different data security solutions, and 21% of security decision-makers cite the lack of consolidated, comprehensive visibility across those tools as their primary security challenge (Microsoft, Data Security Index Report, 2024). Adding AI agents to an already fragmented environment without governance infrastructure is not an acceleration. It is a risk multiplier.

Why speed creates risk

AI deployment in enterprise environments follows a **predictable pattern**

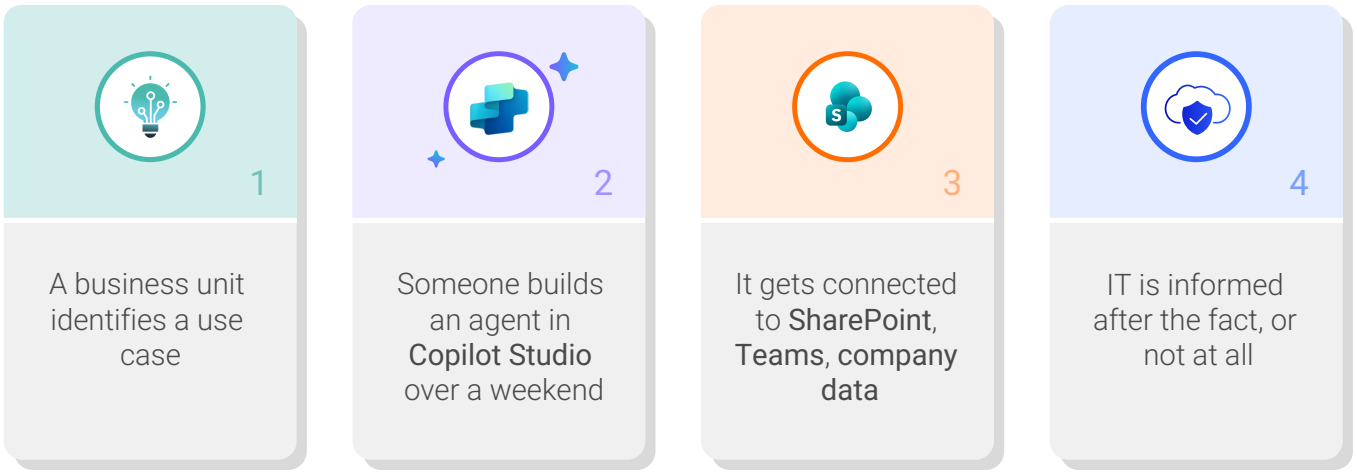

This is not theoretical. The Microsoft Security Blog noted in March 2026 that growing visibility gaps are increasing the risk of agents becoming attack vectors. Agents inherit broad permissions from their provisioning identity. Those permissions accumulate over time. Nobody reviews the original grant.

The organisation ends up with an identity that has significant data access and no one formally accountable for it.

You cannot govern what you cannot see, and Microsoft Copilot can see a lot.

What AI Agents Actually Do in Your M365 Tenant

An AI agent is not a smarter chatbot. A modern Microsoft agent can read emails, query SharePoint, interact with business applications through connectors, send messages, create calendar entries, and trigger other automated workflows. It acts on behalf of a user, or entirely on its own, within whatever permissions its identity has been granted.

The permission problem

The key issue is permissions. An agent inherits the access rights of the identity it runs under. If that identity has broad access, the agent has broad access. It will use all of it, automatically, at scale, without the judgement a human would apply.
In practice, agents are often provisioned with wide permissions so they can handle multiple scenarios. An agent built to analyse financial reports might be given access to all financial data in SharePoint rather than the specific folder it actually needs. As use cases expand, nobody revisits the original access grant. The permissions stay. The agent keeps running with far more access than any individual task requires.

What Copilot sees in your tenant

AI amplifies whatever foundation you have. The good and the bad. If your Microsoft 365 environment has overpermissioned content, stale data, and governance debt built up over years, AI will accelerate all of it. It will find files that were technically accessible but practically invisible. It will surface sensitive information that nobody went looking for. It will treat obsolete data as current and authoritative.
Microsoft Copilot respects existing permissions. It will not surface content a user cannot already access. That sounds reassuring. The problem is that in most Microsoft 365 tenants, permissions are far wider than they should be.

Oversharing is common. Teams sites created with default access settings. SharePoint libraries shared with all authenticated users. Sensitive documents in departmental sites with overly broad membership. Guest access granted for a project and never removed. The consequences are immediate. IT architects doing Copilot rollouts have reported finding sensitive files such as redundancy lists and financial documents accessible to the entire organisation, surfaced instantly the moment Copilot is switched on. In many tenants, a significant portion of sensitive content is accessible to far more people than intended, and often no one has noticed because no one has ever mapped it.
This concern is widely shared. 80% of risk leaders name sensitive data leakage as their top AI-related concern (Microsoft, Data Security Index Report, 2024), and the underlying cause is almost always the same one described above: permissions that were already too broad before any AI agent arrived. Copilot will surface all of it. An agent connected to your SharePoint will draw on everything that the identity running it can access. The AI readiness question is not whether Copilot works. It is what Copilot can see and whether it should see it.

Five common AI data risks in a typical M365 tenant:

1 **Overpermissioned content:**
sensitive data accessible to users and agents who should not see it

2 **ROT data:**
redundant, obsolete, trivial files that AI treats as current and authoritative

3 **Stale permissions:**
guest access, inactive accounts, legacy service accounts with broad rights still in place

4 **Ungoverned agents:**
tools built in Copilot Studio without IT oversight, lifecycle policies, or access review

5 **Agent-scale exposure:**
a single agent can query across SharePoint, mail, and Teams in seconds, surfacing in one search what would once have taken a person months of manual digging to find

The consumption billing trap

Microsoft's AI billing is not one product at one price. Agent 365 is per-seat at $15 per user per month, which is predictable. Copilot Studio is billed in credits ($200 per pack of 25,000 credits per month), where the cost depends entirely on what agents actually do. Azure AI Foundry runs on tokens, with rates that scale significantly at higher tiers. Finance governance was built around per-seat predictability. When consumption meters run against usage patterns no one modelled in advance, the bill always wins. Gartner forecasts that by 2027, 40% of enterprises using consumption-priced AI will see unplanned costs exceeding twice their budget.

This is not a hypothetical risk. Gartner projects that AI agent software spending will reach $207 billion in 2026, up 139% from $86.4 billion in 2025, and a 2025 Mavvrik survey found that 85% of companies missed their AI cost forecasts by more than 10%.

The pattern shows up at enterprise scale too. In May 2026, Fortune reported that a major technology company had exhausted its entire annual AI budget in four months after rolling out a consumption-priced AI coding assistant to thousands of engineers with no usage cap, with per-engineer costs reaching $500 to $2,000 a month. Around the same time, a large enterprise software vendor scaled back its own internal deployment of a comparable tool once the cost profile became untenable at scale. Cheaper tokens will not fix this on their own: Gartner expects inference costs to fall 90% by 2030, but also notes that agentic workflows consume 5 to 30 times more tokens per task than simple prompts, so the unit price drops while total usage, and total spend, keeps climbing.

Copilot Cowork, which reached general availability on June 16, 2026, adds a fourth meter on top of these three. It requires the per-seat Microsoft 365 Copilot licence, then bills usage separately in Copilot Credits, with the price of every task calculated from four variables: model used, context retrieved, tool calls made, and runtime. Microsoft's own rollout materials are a useful admission of the forecasting problem this creates. The company published three task-size tiers (light, medium, heavy), four user personas, and a downloadable spreadsheet so customers can estimate spend before committing budget, because seat count alone tells finance nothing about what a usage-metered agent will actually cost (Microsoft 365 Blog, June 2026). That is the consumption billing trap in its purest form: a fourth bill, on a fourth cycle, in a fourth unit, layered on top of three that already do not reconcile.

The Governance Gaps Microsoft's **Native Tooling Leaves Open**

Microsoft has invested in AI governance. Agent 365 provides a control plane for agent visibility and basic policies. Purview DSPM surfaces data security posture. Defender's Shadow AI pane identifies local agent activity on Windows devices. These are meaningful capabilities. They are not enough on their own.

The risk does not go away by adding more tools on top. Microsoft's own research found that organisations running 11 or more data security tools experienced 202 security incidents, compared to 139 for organisations running 10 or fewer (Microsoft, Data Security Index Report, 2024). Stacking point solutions without consolidating visibility does not reduce exposure. It redistributes it across more dashboards that nobody has time to check.

Gap 1 | File-level data visibility

Agent 365 tells you which agents are running and what policies apply. It does not tell you, at the file level, what data those agents are accessing, whether that data is current and accurate, or whether the permissions around it are fit for AI use. Purview DLP provides oversharing alerts. TeamsFox provides the file-level detail: exactly which files are inactive, redundant, overpermissioned, or stale, ranked by risk and savings potential.

Gap 2 | Pre-AI data hygiene

AI agents and Copilot are only as reliable as the data they draw from. Most Microsoft 365 tenants carry significant ROT, which stands for redundant, obsolete, and trivial content built up over years without systematic cleanup. On average, 40% of M365 storage falls into this category. AI systems that ingest ROT data produce unreliable outputs, surface outdated information, and create compliance risk when stale data is treated as current. Microsoft's native tooling does not address this. Fixing it before an AI rollout is a prerequisite, not an optional step.

Gap 3 | Agent identity & lifecycle management

Agent identities need the same lifecycle discipline as user identities: provisioning with least-privilege access, regular access reviews, and decommissioning when no longer needed. Most organisations have no process for this. Agents are created, connected to data, and then forgotten. The result is agent sprawl: redundant agents with outdated permissions sitting live in the tenant with no active business purpose.

Agent 365 helps here, but it has real limits. Its registry is bounded to the Microsoft ecosystem, so third-party agents that reach M365 data through OAuth grants or direct Graph API calls are not automatically discovered.

Gap 4 | Cross-tenant governance for MSPs

MSPs and CSPs managing Microsoft 365 across multiple client environments face the governance problem at scale. Each tenant has its own agent sprawl, its own permissions debt, its own ROT. Native tooling is per-tenant. TeamsFox's multi-tenant visibility gives MSPs a single consolidated view across all managed environments, making AI governance manageable without multiplying admin overhead.

Gap 5 | AI cost governance

Who is tracking actual AI consumption costs across the organisation? Agent 365 is per-seat. Copilot Studio runs on credits. Azure Foundry bills per token. Copilot Cowork adds a fourth: task-based Copilot Credits, billed on top of the per-seat Copilot licence it requires. These costs arrive through different billing channels, on different cycles, with different units. Without a view that correlates AI usage with licence entitlement and consumption spend, finance teams cannot manage what they cannot see. TeamsFox applies the same licence governance discipline it uses for human identities to AI consumption.

The picture is wider than Microsoft alone. Most enterprises are now running Copilot alongside Google Gemini for Workspace, ChatGPT Enterprise, or both, each procured through a different channel, priced under different terms, and governed by a different set of data residency and compliance commitments. Native Microsoft tooling has no visibility into what those other tools cost or access. A governance approach that only covers the Microsoft estate leaves the rest of the AI footprint, and often the fastest-growing part of it, completely unmanaged.

The Regulatory Picture

AI governance in Microsoft 365 is not purely a best practice. Regulatory frameworks are creating real compliance requirements for enterprise AI deployments in Europe. We will note where these are relevant without overstating what is currently clear for any specific organisation.

EU AI Act

The EU AI Act is already partly in force, and the next deadline is close. Prohibited AI practices, such as social scoring and real-time biometric surveillance in public spaces, have been banned since February 2025. Obligations for general-purpose AI models, the foundation models that power systems like Microsoft Copilot, have applied since August 2025.

The remaining provisions, including the full high-risk AI framework, take effect on 2 August 2026, with no grace period afterwards. Fines for high-risk violations can reach 35 million euros or 7% of global annual turnover, whichever is higher.
Classification follows the workflow, not the tool. Microsoft Copilot used for everyday drafting and summarising sits in the limited-risk category. The same Copilot deployment used to inform hiring decisions, credit assessments, or other high-risk use cases brings the full set of obligations: documented risk classification, human oversight, audit logging, and in some cases registration in the EU AI database. Citizen-built Copilot Studio agents carry the same exposure, often with less visibility, since they can inherit permissions and execute multi-step actions without systematic IT review.

Most organisations are not ready. 85% of security leaders at organisations actively developing AI say they are not prepared to comply with AI regulations (SAP LeanIX, AI Survey Results, 2024), and the inventory gap underneath that number is just as stark: over 50% lack a systematic inventory of the AI systems running in their environment, which means many cannot yet say with confidence which of their AI workflows are in scope, let alone demonstrate compliance if a regulator asks. Building that inventory, and the documented oversight that sits behind it, is the practical starting point regardless of how enforcement unfolds.

NIS2 and DORA

NIS2 sets cybersecurity standards for critical infrastructure across the EU, and DORA applies to EU financial sector organisations. Both frameworks are relevant to the extent that AI agents create new attack surfaces and new operational dependencies that IT governance needs to address.

The specific reporting and documentation requirements for individual organisations depend heavily on their sector, size, and the nature of their AI deployments.

TeamsFox recommends working with legal and compliance teams to assess specific obligations rather than applying generic checklists. What we can say with confidence is that any governance work done on agent access, permissions hygiene, and audit trails will support compliance readiness regardless of which specific regulation applies.

Practical guidance:

With the EU AI Act's full high-risk framework applying from 2 August 2026, the safest starting point is building a complete picture of what AI agents can access in your tenant and who is accountable for them. That single step supports every regulatory framework currently in scope.

The TeamsFox Framework: 5 Pillars of M365 AI Governance

TeamsFox

TeamsFox's approach to AI governance is built on five pillars that cover the full lifecycle of AI deployment: from data readiness through to ongoing cost control and measurable adoption. These pillars reflect capabilities TeamsFox has already proven in M365 environments, applied now to the AI governance context where they are most urgently needed.

Pillar 1: Data Readiness

1

Clean before you connect

AI systems are only as reliable as the data they draw from. Before Copilot or any agent is connected to your tenant, the data foundation needs to be sound. That means removing ROT content that AI will treat as current, fixing overpermissioned files that Copilot will surface to people who should not see them, and identifying stale data that will corrupt AI outputs.

TeamsFox provides file-level ROT analysis across Teams, SharePoint, and OneDrive. It identifies inactive files, excessive versions, redundant content, and obsolete data. Organisations using TeamsFox before their Copilot rollout achieve on average a 40% storage cost reduction and a materially cleaner data surface for AI ingestion.

TeamsFox result: 40% storage cost reduction. Applied to AI: 40% less ROT means 40% less bad data for agents to work with.

Pillar 2: Permissions Hygiene

2

Least privilege before AI is given access

The principle of least privilege is the most fundamental security control for AI systems. An agent should have access only to the data and functions it strictly needs for its specific task. Broad permissions create data exposure risk, compliance risk, and AI output quality problems when agents draw on data they were not designed to use.

TeamsFox's governance and lifecycle management capabilities identify overpermissioned accounts, stale guest access, and service accounts with excessive rights. Applied to AI governance, this means identifying overpermissioned agents specifically, not just the human identities behind them, and mapping the full data access path for every agent: which sites, mailboxes, and data stores it can reach, and through which permission grant. That mapping is what turns least-privilege from a policy statement into something IT can actually enforce, audit, and defend to a regulator. It is also the permissions hygiene work that makes Agent 365's least-privilege policies actually meaningful in practice.

Pillar 3: Lifecycle Governance

3

Manage agents like applications

AI agents are not one-time deployments. They are live systems that need maintenance, access review, and eventual retirement. Without lifecycle governance, organisations accumulate agent sprawl: redundant agents with outdated permissions sitting live in the tenant, forgotten agents still connected to sensitive data sources, orphaned agents that represent pure risk with no active business purpose.

TeamsFox's lifecycle management approach, already proven for user identities and M365 workspaces, applies directly to agent governance: build an inventory of all agents, assign owners, define review cadences, and automate decommissioning for inactive agents. A named owner and review date for every agent, a trigger to review after 90 days of inactivity, and a rule to disable access within 24 hours of an owner leaving the organisation turn lifecycle governance

from a policy statement into an enforced process. The same discipline extends agent, a trigger to review after 90 days of inactivity, and a rule to disable access within 24 hours of an owner leaving the organisation turn lifecycle governance from a policy statement into an enforced process. The same discipline extends to behaviour: monitoring agents for unusual access volume, out-of-scope API calls, or unexpected data writes catches risk that a one-time access review would miss. The key is that controls run continuously, not as a one-time cleanup.

This shift in ownership is already underway. 87% of data security, governance, compliance, and privacy leaders now hold responsibilities that span more than one of these areas (Microsoft, Customer Requirements Research, 2024), and some organisations are going further, building dedicated AI governance roles inside the security team itself. As one Financial Services CTO put it: "We're specifically adding roles around AI data governance within the security team." The named-owner model TeamsFox enforces is not a new idea being imposed on IT. It is where accountability is already moving; the gap is tooling that keeps up with it.

Pillar 4: Visibility and Cost Control

4

Control AI spend, not just see it

AI spending in Microsoft 365 is fragmented across four billing models: per-seat for Agent 365, credit-based for Copilot Studio, consumption-based for Azure Foundry, and now task-based Copilot Credits for Copilot Cowork, billed on top of the per-seat licence it requires. Without a unified view, organisations cannot understand what AI is actually costing them, which agents are generating value, and where consumption budgets are drifting.
Microsoft has acknowledged the problem directly. At Copilot Cowork's general availability in June 2026, Microsoft named cost management as one of the most important new capability areas in the release. Cowork ships off by default, requiring an admin to deliberately turn it on and decide who gets access, alongside spending limits, usage alerts, and usage reporting at the tenant, group, and user level. That is a meaningful step, and it validates the thesis of this paper: usage-based, agentic AI cannot be governed after the fact. But the scope is narrow by design. Those controls cover Copilot Cowork's own credit meter and its own access toggle. They say nothing about Agent 365 seats, Copilot Studio credits, Azure Foundry tokens, or the Google Gemini and ChatGPT Enterprise spend most enterprises are already carrying alongside Microsoft's stack. Microsoft provides controls for the product it sells. TeamsFox provides control across all of them.

Pillar 5: Adoption and Value Optimization

5

Govern for value, not just risk

Governance answers whether an agent is safe. It does not answer whether an agent is useful. Most organisations have no way to tell the difference between an agent that is actively solving a business problem and one that quietly consumes licence, storage, and compute while almost nobody opens it. Treating every agent as equally important, regardless of how it is actually used, wastes budget and obscures the agents that genuinely need attention.

TeamsFox extends the same usage analytics it applies to Microsoft 365 licensing to the agent layer: which agents are used daily, which are touched once a month, and which have not been opened since they were built. Low-adoption agents are flagged before they become a cost and compliance liability, not after a renewal cycle has already locked in another year of unused spend. The same usage data feeds back into design decisions: patterns of actual use show which agents need retraining, narrower permissions, consolidation, or retirement, so governance is based on how agents really behave rather than how they were assumed to behave at launch.

30%	60%	40%	25%
Avg. licence cost reduction	Avg. admin time reduction	Avg. storage cost reduction	Avg. AI Agents cost reduction

A Practical AI Governance **Checklist for IT Leaders**

The questions below give IT leaders a structured starting point for assessing AI governance maturity in their Microsoft 365 environment. Each section surfaces immediate priorities and identifies where the highest-risk gaps typically sit.

Data Foundation

- ☐ Have you completed a file-level ROT analysis across Teams, SharePoint, and OneDrive?
- ☐ Do you know which content is accessible to more users than intended?
- ☐ Have stale files, duplicate versions, and obsolete data been removed or archived?
- ☐ Is your SharePoint metadata structure current and accurate enough for AI to work from?

Permissions and identity

- ☐ Do you have a current map of which identities (both human and non-human) have access to which data?
- ☐ Have guest accounts been reviewed and scoped to their minimum required access?
- ☐ Are service accounts and API connections following least-privilege principles?
- ☐ Has Entra ID been audited for dormant accounts and over-privileged roles?

Agent governance

- ☐ Do you have an inventory of all AI agents running in your tenant, including Copilot Studio, Power Automate, and citizen-built tools?
- ☐ Does each agent have a named owner and a documented business purpose?
- ☐ Are there approval processes for new agents before they connect to production data?
- ☐ Do you have sunset policies for inactive or orphaned agents?

Cost and licence governance

- ☐ Do you have visibility into all four AI billing streams: Agent 365 per-seat, Copilot Studio credits, Azure Foundry tokens, and Copilot Cowork's task-based credits?
- ☐ Have you set consumption budgets and alert thresholds for Copilot Studio and Foundry usage?
- ☐ Are Agent 365 licences assigned only to users who actively need them?
- ☐ Is AI licence rightsizing part of your regular M365 licence governance cycle?

Adoption and value

- ☐ Do you track which agents are actively used versus dormant?
- ☐ Is there a process to flag and retire low-adoption agents before renewal?
- ☐ Are usage patterns reviewed to optimise agent design, permissions, or consolidation?
- ☐ Can you demonstrate which agents are delivering measurable business value?

Conclusion

The AI agent era in enterprise Microsoft 365 is not approaching. It is here. Agent 365 is generally available. M365 E7 bundles Copilot and agent controls into a single suite. The EU AI Act's full high-risk framework applies from 2 August 2026, with no grace period after that date. The market is moving fast.

The organisations that establish governance foundations now will be in control when agents become standard infrastructure. Those that wait will find themselves with a backlog of ungoverned agents, overpermissioned data, and consumption costs that no one authorised.

This is also the direction the market is already moving. 95% of security, governance, compliance, and privacy leaders say unifying their tools and teams is now a priority, and 90% expect their organisation to adopt a unified platform approach (Microsoft, Audience Research / Customer Requirements Research, 2024). TeamsFox is built for that shift, not against it.

TeamsFox's position is straightforward. We do not just provide governance tooling. We provide the foundational M365 hygiene that makes AI governance possible: file-level data quality, permissions architecture, lifecycle automation, and tenant-wide visibility that Microsoft's own platform has not fully addressed.

Microsoft enables AI. We govern it.

TeamsFox

TeamsFox provides the governance layer that turns Microsoft 365's AI capabilities from an organisational risk into a manageable asset.

About TeamsFox

TeamsFox GmbH is a Franco-German Microsoft 365 management and optimisation platform. Headquartered in Dusseldorf, Germany, TeamsFox is trusted in 20+ countries across Europe, MENA, and Asia.

Website | www.teamsfox.com

Email | info@teamsfox.com