

Getting the most out of generative AI at Microsoft with good governance

Oct 19, 2023 | [Lukas Velush](#)



Our smart approach to governance at Microsoft is helping us safely unlock the benefits of next-generation AI in tools like Microsoft 365 Copilot.

Since generative AI exploded onto the scene, it's been unleashing our employees' creativity, unlocking their productivity, and up-leveling their skills.

But we can fly into risky territory if we're not careful. The key to protecting the company and our employees from the risks associated with AI is adopting proper governance measures based on rigorous data hygiene.

Technical professionals working within Microsoft Digital (MSD), our internal IT organization, have taken up this challenge. [They include the AI Center of Excellence \(AI CoE\) team](#) and the Microsoft Tenant Trust team that governs our Microsoft 365 tenant.

The endgame here is acceleration. AI accelerates employees' ability to get questions answered, create things based on dispersed information, summarize key learnings, and make connections that otherwise wouldn't be there.

—David Johnson, tenant and compliance architect, MSD

Since the widespread emergence of generative AI technologies over the last year, our governance experts have been busy ensuring our employees are set up for success. Their collaboration helps us ensure we're governing AI through both guidance from our AI CoE and a governance model for our Microsoft 365 tenant itself.

[\[Learn how Microsoft is responding to the AI revolution with a Center of Excellence. Discover transforming data governance at Microsoft with Purview and Fabric. Explore how we use Microsoft 365 to bolster our teamwork.\]](#)

Generative AI presents limitless opportunities—and some tough challenges

Next-generation AI's benefits are becoming more evident by the day. Employees are finding ways to simplify and offload mundane tasks and focus on productive, creative, collaborative efforts. They're also using AI to produce deeper and more insightful analytical work.

"The endgame here is acceleration," says David Johnson, a tenant and compliance architect with MSD. "AI accelerates employees' ability to get questions answered, create things based on dispersed information, summarize key learnings, and make connections that otherwise wouldn't be there."

There's a real urgency for organizations to empower their employees with advanced AI tools—but they need to do so safely. Johnson and others in our organization are balancing the desire to move quickly against the need for caution with technology that hasn't yet revealed all the potential risks it creates.

"With all innovations—even the most important ones—it's our journey and our responsibility to make sure we're doing things in the most ethical way," says Faisal Nasir, an engineering leader on the AI CoE team. "If we get it right, AI gives us the power to provide the most high-quality data to the right people."

We're going to be one of the first organizations to really get our hands on the whole breadth of AI capabilities. It will be our job to ensure we have good, sensible policies for eliminating unnecessary risks and compliance issues.

—Matt Hempey, program manager lead, MDE

But in a world where AI copilots can comb through enormous masses of enterprise data in the blink of an eye, security through obscurity doesn't cut it. We need to ensure we maintain control over where data flows throughout our tenant. It's about providing information to the people and apps that have proper access and insulating it against ones that don't.

To this end, our AI CoE team is introducing guardrails that ensure our data stays safe.

Tackling good AI governance

The AI CoE brings together experts from all over Microsoft who work across several disciplines, from data science and machine learning to product development and experience design. They use an AI 4 ALL (Accelerate, Learn, Land) model to guide our adoption of generative AI through enablement initiatives, employee education, and a healthy dose of rationality.


"We're going to be one of the first organizations to really get our hands on the whole breadth of AI capabilities," says Matt Hempey, a program manager lead on the AI CoE team. "It will be our job to ensure we have good, sensible policies for eliminating unnecessary risks and compliance issues."

As Customer Zero for these technologies, we have a responsibility for caution—but not at the expense of enablement.

"We're not the most risk-averse customer," Johnson says. "We're simply the most risk-aware customer."

The AI CoE has four pillars of AI adoption: strategy, architecture, roadmap, and culture. As an issue of AI governance, establishing compliance guardrails falls under architecture. This pillar focuses on the readiness and design of infrastructure and services supporting AI at Microsoft, as well as interoperability and reusability for enterprise assets in the context of generative AI.

Microsoft AI CoE's operational pillars



We've created four pillars to guide our internal implementation of generative AI across Microsoft: Strategy, architecture, roadmap, and culture. Our work on AI governance falls under architecture.

Building a secure and compliant data foundation

Fortunately, Microsoft's existing data hygiene practices provide an excellent baseline for AI governance.

There are three key pieces of internal data hygiene at Microsoft:

1. Employees can create new workspaces like Sites, Teams, Groups, Communities, and more. Each workspace features accountability mechanisms for its owner, policies, and lifecycle management.
2. Workspaces and data get delineated based on labeling.
3. That labeling enforces policies and provides user awareness of how to handle the object in question.

With AI, the primary concern is ensuring that we properly label the enterprise data contained in places like SharePoint sites and OneDrive files. AI will then leverage the label, respect policies, and ensure any downstream content-surfacing will drive user awareness of the item's sensitivity.

AI will always respect user permissions to content, but that assumes source content isn't overshared. Several different mechanisms help us limit oversharing within the Microsoft tenant:

1. Using site labeling where the default is private and controlled.
2. Ensuring every site with a "confidential" or "highly confidential" label sets the default library label to derive from its container. For example, a highly confidential site will mean all new and changed files will also be highly confidential.
3. Enabling company sharable links (CSLs) like "Share with People in <name of organization>" on every label other than those marked highly confidential. That means default links will only show up to the direct recipient in search and in results employees get from using Microsoft's various Copilots.
4. All Teams and sites have lifecycle management in place where the owner attests that the contents are properly labeled and protected. This also removes stale data from AI.
5. Watching and addressing oversharing based on site and file reports from Microsoft Graph Data Connect.

Microsoft 365 Copilot and other Copilots respect labels and display them to keep users informed of the sensitivity of the response. They also respect any rights management service (RMS) protections that block content extraction on file labels.

To make the Copilot platform as successful and securely extensible as possible, we need to ensure we can control data egress from the tenant.

—Keith Bunge, software engineering architect for employee productivity solutions, MSD

If the steps above are in place, search disablement becomes unnecessary, and overall security improves. "It isn't just about AI," Johnson says. "It's about understanding where your information sits and where it's flowing."

From there, our various Copilots and other AI tools in question can then safely build a composite label and attach it to its results based on the foundational labels it used to create them. That provides the context it needs to decide whether to share its results with a user or extend them to a third-party app.

“To make the Copilot platform as successful and securely extensible as possible, we need to ensure we can control data egress from the tenant,” says Keith Bunge, a software engineering architect for employee productivity solutions within MSD.

We can also use composite labels to trigger confidential information warnings to users. That transparency provides our people with both agency and accountability, further cementing responsible AI use within our culture of trust.

Process, people, and technology are all part of this effort. The framework our team is developing helps us look at data standards from a technical perspective, as well as overall architecture for AI applications as extensions on top of cloud and hybrid application architecture.

—Faisal Nasir, principal architect, MSD

Ultimately, AI governance is similar to guardrails for other tools and features that have come online within our tenant. As an organization, we know the areas we need to review because we already have a robust set of criteria for managing data.

But since this is a new technology with new functionality, the AI CoE is spending time conducting research and partnering with stakeholders across Microsoft to identify potential concerns. As time goes on, we’ll inevitably adjust our AI governance practices to ensure we’re meeting [our commitment to responsible AI](#).



From left to right, David Johnson, Faisal Nasir, Matt Hempey, and Keith Bunge are among those working together here at Microsoft to ensure our data estate stays protected as we adopt next-generation AI tools.

“Process, people, and technology are all part of this effort,” Nasir says. “The framework our team is developing helps us look at data standards from a technical perspective, as well as

overall architecture for AI applications as extensions on top of cloud and hybrid application architecture.”

As part of getting generative AI governance right, we’re conducting extensive user experience and accessibility research. That helps us understand how these tools land throughout our enterprise and keep abreast of new scenarios as they emerge—along with the extensibilities they need and any data implications. We’re also investing time and resources to catch and rectify any mislabeled data, ensuring we seal off any existing vulnerabilities within our AI ecosystem.

Not only does this customer zero engagement model support our AI governance work, but it also helps build trust among employees through transparency. That trust is a key component of the employee empowerment that drives adoption.

Eight steps for getting tenant data governance right with Microsoft 365 Copilot

1 New workspaces

Let employees create new workspaces: We empower our employees to take full advantage of Microsoft 365 Copilot by making sure all our data is on our Microsoft 365 tenant.



2 Label containers

We label our containers for data segmentation to ensure our data is not over exposed by default. To do this, we set our container label defaults with the "private/no guests" setting.



3 Layered consistency

We derive file labels from their parent container labels. This layered consistency boosts security at multiple levels.



4 Verify labels

We trust employees to apply sensitivity labels, but we also verify them. We verify by checking against our Data Loss Prevention (DLP) standard and by using auto-labeling and quarantining when needed.



5 Employee training

We train our employees how to handle and label sensitive data to increase our labeling accuracy.



6 Lifecycle management

We use strong lifecycle management protocols that require our containers to be attested.



7 Limit oversharing

We limit oversharing at the source by enabling company shareable links. (Better than forcing people to add large groups for access.) Our standard is even higher for highly confidential items, which we only "share with specific people" using our "need to know" standard.



8 Data extraction

We use our Microsoft Graph Data Connect extraction to catch and report oversharing after the fact. When we find it, we drive employees to fix.



Realizing generative AI's potential

As our teams navigate AI governance and drive adoption among employees, it's important to keep in mind that these guardrails aren't there to hinder progress. They're in place to protect and ultimately inspire confidence in new tools.

"In its best form, governance is a way to educate and inform our organization to move forward as quickly as possible," Hempey says. "We see safeguards as accelerators."

We know our customers also want to empower their employees with generative AI. As a result, we're discovering ways to leverage or extend these services in exciting new ways for the organizations using our products.

"As we're on this journey, we're learning alongside our industry peers," Nasir says. "By working through these important questions and challenges,

we're positioned to empower progress for our customers in this space."



Key takeaways

Consider these tips as you think about governing the deployment of generative AI at your company:

- Understand that IT organizations have inherently cautious habits.
- Leverage what industry leaders like the Responsible AI Initiative are sharing.
- Recognize that employees will adopt these tools on their own, so it's best to prepare the way beforehand.
- Consider your existing data hygiene and how it needs to extend to accommodate AI.
- Make sure you have an enterprise plan for ensuring labeling and security, because AI tools will provide the most complete access by default.



Try it out

Get started on your own next-generation AI revolution—[try Microsoft 365 Copilot today](#).



Related links

- [Learn how Microsoft is responding to the AI revolution with a Center of Excellence.](#)
- [Discover transforming data governance at Microsoft with Purview and Fabric.](#)
- [Explore how we use Microsoft 365 to bolster our teamwork.](#)



We'd like to hear from you!

[Please share your feedback with us—take our survey and let us know what kind of content is most useful to you.](#)

Tags: [AI](#), [AI and Machine Learning](#), [automation](#), [security](#)