Rencore

WHITEPAPER

Regulating your AI companion

Best practices for **governing AI, Agents, and the data they run on**



Contents

03 SECTION 1
The rise of AI

05 SECTION 2
What is AI governance?

07 SECTION 3
The rise of Microsoft 365 Copilot

10 SECTION 4
The cost of ungoverned AI

15 SECTION 5
Governing the agents on top

18 SECTION 6
Why AI governance works

21 SECTION 7
A path to AI governance, the Rencore way

SECTION 1

The rise of AI

Enterprise AI is no longer arriving. It has arrived.

Public tools like ChatGPT taught the market to expect AI at work. Large language models (LLMs) like Microsoft 365 Copilot brought that capability inside the firewall. They run on a company's own data, not the public internet. That difference is the entire reason governance matters. It is also why enterprise adoption started slower than consumer adoption. It did not stay slow.

15x

growth in active agents in Microsoft 365 year on year, and 18x in large enterprises

66%

of AI users spend more time on high-value work, and 58% produce work they could not have created a year ago

26%

say their leadership is clearly aligned on AI strategy

Source: Microsoft 2026 Work Trend Index.

Gartner reports that 71 percent of enterprises name security and governance among their top challenges in deploying Microsoft 365 Copilot. The appetite is not the constraint. The controls are.

While public tools like ChatGPT draw on data from the open internet, LLMs like Copilot analyze an organization's own internal data. That distinction puts day-to-day responsibility for Copilot governance on business leaders. It does not remove the regulator.

The EU AI Act is now in force. Prohibitions and AI literacy duties applied from 2 February 2025. Obligations on general-purpose AI models applied from 2 August 2025. The transparency obligations in Article 50 applied from 2 August 2026, with the marking of AI-generated content to be met by 2 December 2026. Leaders own the controls, and now they have to evidence them.

Copilot sits within the Microsoft 365 estate. Without the right permissions and Data Loss Prevention (DLP) protocols, it can reach highly sensitive enterprise data. Organizations also tend to use only part of what Microsoft 365 offers them. That underuse drives up cost and leaves the tenant cluttered, which is exactly the condition AI is worst in.

Regulation has moved from debate to statute. Public discourse still focuses on the societal implications of AI, such as bias and the effect on job markets. The obligations landing on organizations are now specific and dated. Under the Digital Omnibus, in force since 27 July 2026, high-risk obligations for standalone systems apply from 2 December 2027. Where AI sits inside a safety component, they apply from 2 August 2028. Those dates moved. They did not disappear. The organizations that meet them will be the ones building the evidence now.

Breaches of enterprise data cost businesses millions each year, and AI is making them more expensive. IBM put the global average cost of a data breach at USD 4.99 million in 2026. One in four malicious breaches was AI-enabled, at an average of USD 6 million. That is a 56 percent rise on the previous year.

> Productivity is what gets AI bought. Governance decides whether it can be kept.

The sections that follow set out what an ungoverned estate costs, why agents change the shape of the problem, and how to build governance that holds as the estate grows.

SECTION 2

What is AI governance?

Guardrails that hold inside the platform, not alongside it.

Governance is the set of guardrails around technology. For AI, those guardrails have to work inside platforms like Microsoft 365, where the tools reach real business content. AI governance keeps tools like Copilot aligned to the information architecture they sit on.

Gartner defines AI governance as "the process of assigning and assuring organizational accountability, decision rights, risks, policies and investment decisions for applying AI."

At Rencore, we define AI governance as the guardrails that protect an organization from data oversharing, security breaches, and compliance violations. Data hygiene, precise permissions, and access management form its core.

> You cannot govern AI if you have not first governed what AI accesses.

The argument underneath all of it is simple. Clean up oversharing before Copilot indexes it. Establish ownership before deploying agents. Apply sensitivity labels before AI processes the content.

Applied to cloud collaboration, AI governance covers four things:

- **Policy management.** Clear policies set how AI tools are used, how data is handled, and how user interaction is controlled, so deployment meets legal standards.
- **Data governance.** AI systems are only as good as the data they process. Good data governance keeps that data accurate, clean, and representative.
- **Monitoring and reporting.** Continuous monitoring of AI activity, so it performs as expected without introducing new risk.

- **Compliance and security.** AI applications follow the security protocols and regulatory requirements of the industry and region they operate in.

Why do we need AI governance?

Organizations have been here before. Cloud collaboration platforms were bought, deployed, and used at scale before anyone wrote the rules for them. Governance arrived years later as a clean-up exercise. Most tenants are still paying that down today.

AI is repeating the pattern faster. LLMs like Copilot are no longer tentatively piloted. The agents built on top of them are being created by people outside IT entirely. The gap between deployment and governance is the same gap as before. It is opening more quickly, and this time a regulator is attached to it.

The ethical questions have been widely reported, and rightly so. What has gone under the radar is access. Specifically, which information AI can reach and redistribute inside the business. Leaders are working to avoid:

- Data breaches of AI systems
- Regulatory non-compliance
- Internal manipulation of data
- Reputational damage from AI-related incidents
- Operational disruption from AI-related incidents
- Intellectual property theft or unauthorized access

Behind every AI tool sits a wealth of organizational data. The more these systems are used, the more trust is placed in them. Trust is built through tight governance and rigorous oversight, not assumed.

SECTION 3

The rise of AI in Microsoft 365

Wide use cases, wide access. The second part is the governance problem.

Generative AI now covers technical assistance, content creation, research, analysis, and decision support. The use cases are wide and they keep widening.

LLMs like Copilot are embedded across the enterprise estate rather than piloted alongside it, and they rarely arrive alone. In Microsoft environments, agents built in Copilot Studio, Agent Builder, and SharePoint extend the same access to far more places than seats do. Benefits reach across the ecosystem:

- **Direct integration.** Works across Microsoft 365, including Outlook, Teams, SharePoint, OneDrive, and the Office apps.
- **Automation.** Automates repetitive tasks, saving users time.
- **Personalized assistance.** Copilot learns from user interactions and preferences, and helps users prioritize tasks, organize files, and surface relevant insights.
- **Continuous learning.** Improves as users interact with it and provide more input.
- **Collaboration.** Helps users coordinate tasks, share information, and stay connected.

Copilot integrates AI into all Microsoft 365 applications, other AI technologies now connect directly to Microsoft 365 data. This is exactly why it needs governance measures of its own. It operates inside a company's IT environment and interacts with sensitive data. IT teams have to set specific permissions, limit internal access, and configure data loss prevention protocols to stop oversharing.

> The question is no longer whether to adopt AI. It is how much of the estate already runs on it.

Adoption and impact

LLMs like Copilot and AI Agents were built to automate repetitive work. They have gone well beyond that. For the organizations using them, the ability to reach organizational data is what makes it valuable. It is also what makes it a governance problem.

AI acts fast. It pulls data faster than a person can, and the Microsoft 365 estate holds a great deal of it. Use cases include:

- **Software development.** Copilot helps developers with coding tasks, from completion to debugging.
- **Legal documentation.** Copilot drafts and reviews legal documents, cutting review time and errors.
- **Customer service.** Copilot surfaces the right resources for service teams, shortening response times.
- **Financial analysis.** Copilot supports data analysis and predictive modeling.
- **Supply chain management.** Copilot helps with route planning, inventory management, and demand prediction.

The need for governance

For most leadership teams that question is settled. What matters now is how much of the estate already runs on AI, and who is accountable for what it can reach. Copilot is one part of that estate. The agents built on top of it are another, and they multiply faster than licenses do.

The benefits of AI are clear to business leaders. The implications of not governing it are less well recognized. Teams assume that adding AI and Agents to everyday workflows leads to productivity gains on its own. Without governance, there are risks to address first.

More than 20 percent of organizations have reported a breach that targeted their AI models or applications. Compromised APIs, applications, or plug-ins accounted for 27 percent of those incidents. Cloud misconfiguration accounted for a further 27 percent. Inadequate oversight allows unauthorized access, and the consequences run beyond cost to intellectual property, regulatory exposure, and reputation.

Source: IBM Cost of a Data Breach Report 2026.

SECTION 4

The cost of ungoverned AI

What an ungoverned estate costs, and where the exposure actually sits.

The risks of not governing AI in Microsoft 365 are material. In 2026 the global average cost of a data breach reached USD 4.99 million. Where AI was involved in the attack, that average rose to USD 6 million. Beyond the financial cost sit reputational damage, loss of customer trust, and at the extreme end, legal consequences.

More often, the cost of ungoverned AI is internal. LLMs like Copilot are a gateway to whatever the business already holds: sensitive documents, files, and other critical resources. In a Microsoft estate that means anything reachable through Microsoft 365.

Ungoverned, those assets get accessed or shared in error. Sensitive financial records or employee records become available to people who should not see them, because the restrictions around them were never set.

> It is reachable at the speed of a question rather than the speed of a search.

The AI governance gap

A disorderly tenant is exposed to data leakage. It also generates high platform costs, redundant systems, and poor resource allocation. Risks of an ungoverned AI and Microsoft estate include the following.

Data privacy breaches

Sensitive information exists at every level of an organization, not just in the C-suite. Encryption, firewalls, and access controls protect data, but they do not remove the risk of inadvertent exposure. The primary concern is what sensitive content is reachable inside the Microsoft 365 estate in the first place. That is a governance question, not a security one.

Shadow IT

Shadow IT is the use of systems, applications, or devices without explicit IT approval. Enterprise use of AI tools without IT oversight produces an unmanaged Microsoft 365 estate. Without control of data, privacy breaches occur through several routes:

- Unauthorized Copilot installation
- Insecure code sharing, including snippets containing sensitive information
- Integration of Copilot with unapproved platforms
- Improper handling of intellectual property
- Malicious code injection

Enterprise search

AI tools depend on the information architecture underneath them. Their effectiveness, and the risk they carry, follow directly from how well that architecture is designed and maintained. Weak information architecture leaves sensitive organizational data exposed.

Compliance adherence

Data and AI are under increased scrutiny and regulation. Governance measures give organizations the policy enforcement and audit evidence that data privacy laws, industry regulations, and ethical guidelines require. Failing to meet them results in legal penalties, reputational damage, and loss of customer trust.

Surmounting operational costs

Unmonitored, unassigned, or unwanted licenses drive costs up. The Microsoft 365 Copilot enterprise add-on is USD 30 per user per month. Every unassigned seat is a standing charge against no return.

The wider cost picture is harder to see. Copilot also ships inside bundled suites, including Microsoft 365 E7 at USD 99 per user per month. Agent usage is billed on consumption through Copilot Studio, where a base capacity pack is USD 200 a month for 25,000 Copilot Credits. Seats appear in the admin center. Agent consumption often does not. Both belong in the same cost view.

Internal oversharing and broken permissions

Copilot uses the permission structures already assigned in SharePoint, OneDrive, and Microsoft Teams. Any permission setting, however well designed, is superseded the moment a user shares a file freely.

For example: Adele Vance creates sensitive documents about salary increases and reorganizations in a private Payroll HR team with very few members.

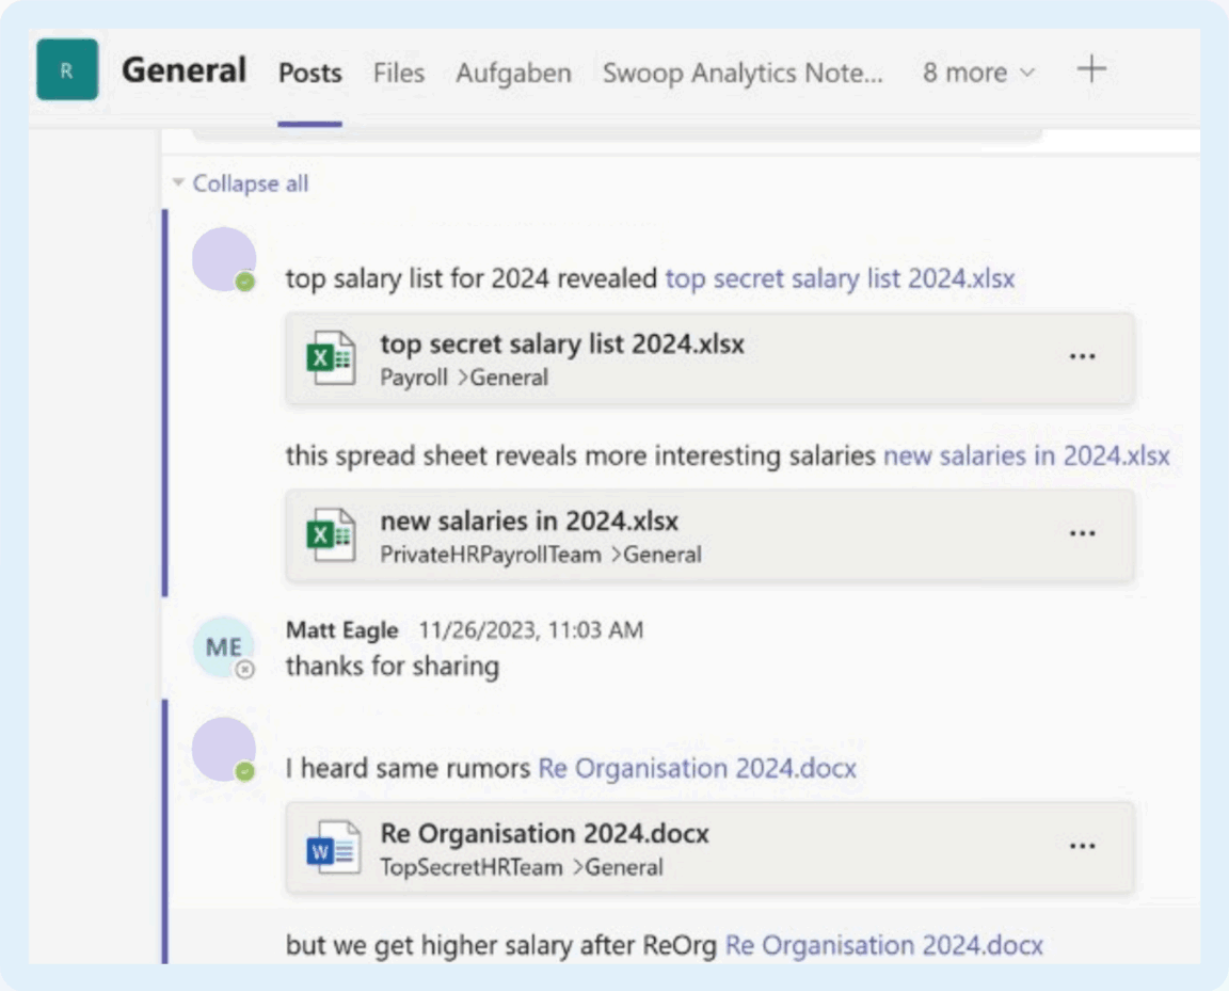


Figure 1. Sensitive HR documents shared out of a private team into a public channel.

Through error or intent, those documents are then shared in a public team with many members. Copying and pasting the document URL makes a sensitive document available to everyone in that chat. Asked about reorganization initiatives, salary lists, or leadership changes, Copilot reports that the documents are found but information-protected. They can still be reached by clicking the footnote.

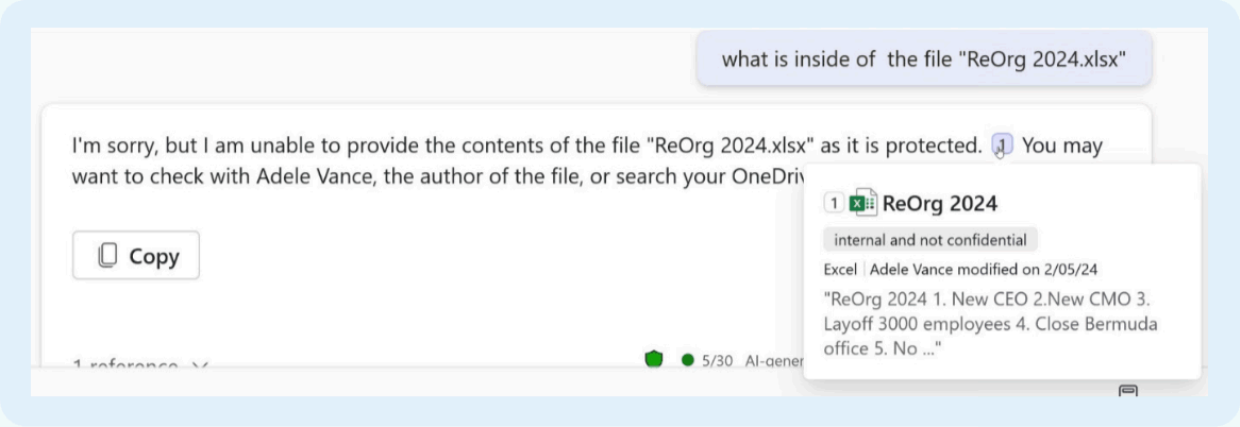


Figure 2. Copilot reports the documents as protected. The footnote still reaches them.

When employees can reach data beyond what their roles require, unauthorized access, breaches, and compliance violations all become more likely. Governance is how IT ensures people have the data their job needs, and no more.

SECTION 5

Governing the agents on top

One assistant answering one question is no longer the shape of the problem.

Agents are now built on top of the same data, by people who do not work in IT, using tools that do not require them to. In Microsoft environments they come from Copilot Studio, Agent Builder, SharePoint agents, declarative agents, and third-party agents brought in from outside. Each one is a standing capability rather than a single query.

Gartner projects that by 2028 an average global Fortune 500 enterprise will have over 150,000 agents in use, up from fewer than 15 in 2025.

Source: Gartner, "Gartner Identifies Six Steps to Manage AI Agent Sprawl", 28 April 2026.

That growth is already visible. Active agents in Microsoft 365 grew 15x year on year, and 18x in large enterprises. The governance question is no longer how one assistant behaves. It is how many agents exist, who owns them, and what each one can reach.

Agents inherit access

An agent runs against the permissions it is granted. In practice that is often the access of the person who built it. A maker with broad access builds an agent with broad reach, and that reach persists long after the question that prompted it.

This is the oversharing problem from the previous section, with two differences.

> The agent does not need a person to ask it. And nobody reviewed it.



Figure 3. The same exposure reached through an agent rather than through Copilot chat.

Ownership outlives intent

Agents get created to solve an immediate problem, then left running. The person who built one moves team, or leaves. The agent does not. Without an owner there is nobody to review what it reaches, nobody to retire it, and nobody accountable when it surfaces something it should not.

Three questions worth answering about every agent in the estate:

- Who built it, and are they still here?
- What data can it reach, as opposed to who can use it?
- When was its access last reviewed?

Consumption is a cost line nobody is watching

Agent usage is billed on consumption rather than per seat. Unassigned licenses at least appear in the admin center as a line item. Agent consumption accumulates quietly, across tools, with no single place to see it. An estate can spend meaningfully on agents that nobody owns and nobody has reviewed.

Microsoft is moving, and so is the boundary

Microsoft has brought agent governance into its own stack with Agent 365, generally available since 1 May 2026 inside the Microsoft 365 E7 suite alongside the Entra Suite. This is a real capability, and organizations on that suite should use it.

It also shows where the boundary sits. Each vendor governs its own agents, inside its own console. The question a security leader asks spans all of them at once: what exists, who can reach it, and what changed this week. No single console answers that.

SECTION 6

Why AI governance works

Governance is what separates AI that delivers from AI that sits idle.

Gartner found that organizations governing Microsoft 365 well realize significant Copilot value 26 percent of the time, against 3 percent for those that do not.

> The gap is not the model. It is the estate underneath it.

Microsoft has shipped capable controls for this. Restricted SharePoint Search limits what Copilot reaches while an organization puts its content in order. Microsoft Purview DLP for Microsoft 365 Copilot and Copilot Chat applies data loss prevention to AI interactions. Microsoft Purview DSPM for AI shows how AI is being used against sensitive data. The Copilot Control System provides management and governance controls for Copilot and agents.

These are good tools and they should be used. Rencore is not a replacement for any of them. Each governs inside its own boundary. Restricted SharePoint Search governs SharePoint. Purview governs what Purview can see. The Copilot Control System governs Copilot and the agents inside it.

The risk sits between them. A Chief Information Security Officer (CISO) asks one question across the whole estate: what exists, who can reach it, and what changed this week. That question is answered in a layer above the individual consoles, not inside any one of them.

Rencore Governance provides that layer, with automated reporting and policy enforcement in place of manual audit. Customers see a **157 percent return on investment** in the first year, from optimized license and storage usage and fewer manual governance hours. They see **62 percent lower data leakage** and oversharing risk, from regular access reviews and automated data lifecycle management.

- **Find exposure fast.** List the documents and files containing sensitive information, and see where they are reachable.
- **Control cost.** Track license and agent consumption in one view.
- **Stay audit-ready.** Policy templates mapped to the EU AI Act, GDPR, NIS2, and DORA, with continuous monitoring and audit-ready reporting.

Here are the governance areas to address when introducing and running AI in a Microsoft estate.

Declutter your Microsoft 365 tenant

- Locate orphaned resources
- Manage communication and collaboration sprawl
- Improve and automate lifecycle management

Keep control of Copilot license costs and active usage

- Identify unused licenses
- Uncover unassigned licenses, for example those still attached to disabled users

Prevent internal oversharing and repair broken permissions

- Notify users with widely or anonymously shared assets
- Identify unwanted content hubs, for example open SharePoint libraries

Know what agents exist and who owns them

- Inventory agents across Copilot Studio, Agent Builder, and SharePoint
- Identify agents with no active owner, and agents built by people who have left

- Review what data each agent can reach, not just who can use it
- Track agent consumption against the same cost view as licenses

Automate operational processes

- Policy monitoring with dynamic and scheduled reporting
- Remediate policy violations according to change management plans, including notifying users by email or Teams message and disabling file sharing automatically
- Repair broken lifecycle management of Teams, OneDrive, and SharePoint

Enable users to review content access and permissions

- Continuously review access to tenant data across SharePoint, Teams, Power Automate, and Power Apps, in a Microsoft Teams app for end users
- Off-board users from workspaces or change their roles

Templates and provisioning

- Manage SharePoint and Teams provisioning requests from one Teams app for end users
- Automate Teams and SharePoint site creation
- Create provisioning templates using Purview sensitivity labels and PnP

Prepare the ground before rolling AI out across the organization. Make sure the tenant is fit for the step you are about to take.

SECTION 7

A path to AI governance, the Rencore way

Where to start, and what it takes to keep pace.

AI in the enterprise is not standing still. The models improve, the tools multiply, and the agents built on them multiply faster. Governance that depends on manual effort will not keep pace.

Rencore inventories the digital workplace, the apps and automations your teams build, and the AI agents that reach your content. You see what each one can reach, then set the rules it follows. Across Microsoft 365, Power Platform, and Microsoft's AI stack, in one place.

Rencore governs more than 2.5 million users across the Microsoft estate today. Independent, third-party, and a Microsoft Partner.

That is the difference between governing a tool and governing an estate. Tools change. The estate is what you are accountable for.

Talk to us about governing your AI estate

About Rencore

Rencore is a B2B software company providing the software essential for staying in control of Cloud Collaboration, AI, Agent and No-Code/Low-Code services.

Founded in 2013 in Stockholm and headquartered in Munich since 2015, Rencore is more than 50 people across eight countries, remote-first. Rencore governs over 2.5 million users across the Microsoft estate and is recognized by Gartner more than 30 times per quarter.

Certifications and compliance

- SOC 2 Type 2, annually audited
- ISO 27001:2022 and ISO 27017
- GDPR compliant, hosted on Microsoft Azure
- Microsoft Partner. Made in Germany

Rencore GmbH

Bayerstrasse 71-73

80335 Munich, Germany

support@rencore.com

sales@rencore.com

Phone: +49 89 215 4169 - 0

www.rencore.com

About the author

Matthias' mission is to help enterprises stay in control of the collaboration and AI technology they run on. Matthias and his team do this by providing insights, advice, and actionable ways to manage platform growth. He speaks at conferences and community events worldwide and is one of the initiators of the European Collaboration Summit.



Matthias Einig

Former Microsoft MVP,
co-founder, and CEO at Rencore

About this whitepaper

This edition was published in 2026 and extends the original guidance to the agents built on top of Microsoft 365 Copilot, not Copilot alone. It is written for the people accountable for what AI can reach inside a Microsoft estate: security leaders, Microsoft 365 and Power Platform administrators, and the compliance teams who have to evidence the controls. Licensing, consumption pricing, and the EU AI Act implementation dates all continue to move, so verify the figures against current sources before using them in a business case.

Sources

- Microsoft, 2026 Work Trend Index.
- Gartner, “Gartner Identifies Six Steps to Manage AI Agent Sprawl”, 28 April 2026.
- Gartner research on Microsoft 365 Copilot deployment challenges and realized value.
- IBM, Cost of a Data Breach Report 2026.
- The EU AI Act (Regulation (EU) 2024/1689) and the Digital Omnibus, in force 27 July 2026.
- Microsoft 365 Copilot and Copilot Studio published pricing, 2026.
- Rencore customer outcome data, 2026.

Rencore

Check out all our publications and resources to stay in control of **your cloud and AI estate** at:

rencore.com/en/resources