ELEVATE
SOLUTIONS

GUIDE · BUSINESS & INDUSTRY

AI Adoption Guide for Regulated Industries

Legal, healthcare, financial

How to use AI safely when you answer to regulators — data governance, vendor risk, compliance mapping, and an acceptable-use framework your firm can adopt this quarter.

01 — AI RISKS IN REGULATED WORK

What can actually go wrong

AI tools are already inside your firm — in browsers, in phones, in the products you already license. The question is not whether your staff will use them. It is whether they will use them under rules you wrote, with data boundaries you chose. Five risk areas deserve attention before any rollout.

1. Confidentiality. A prompt is a disclosure. When an attorney pastes a draft brief into a consumer chatbot, when a clinician summarizes a patient encounter, when an advisor uploads a client's financials — that content leaves your environment and lands under whatever terms govern that tool. For law firms, disclosing privileged material to a third party can put privilege at risk; for healthcare organizations, patient information in a tool without a business associate agreement is a HIPAA problem regardless of how useful the output was. Consumer-tier AI tools are generally not contracted for regulated data. Treat every prompt as if it were an email to an outside vendor, because functionally it is.

2. Accuracy and fabrication. Language models produce fluent text, not verified facts. They invent case citations, misstate dosages, and fabricate figures with complete confidence. Courts in multiple jurisdictions have sanctioned attorneys who filed briefs containing AI-fabricated citations. The failure in those cases was not using AI — it was filing output nobody checked. The control is simple and non-negotiable: a qualified human reviews and verifies every AI-assisted work product before it reaches a client, a court, or a record.

3. Data residency and training terms. Where do prompts go, where are they stored, who can read them, and are they used to improve the model? The answers differ sharply between consumer and enterprise tiers of the same product, and they change as vendors update terms. Enterprise agreements from major AI vendors generally provide that customer prompts and outputs are not used to train the vendor's foundation models — but that is a contract term, not a law of nature. Confirm what your tenant's agreement actually says, in writing, before approving a tool for client data.

4. Shadow AI. Staff are using AI tools today, with or without permission. A blanket ban does not stop this; it just moves usage to personal devices and personal accounts, where you have no visibility, no logs, and no contract. The realistic posture is to govern: provide approved tools that are good enough that people want to use them, set clear data boundaries, and make the prohibited paths explicit. An approved tool with logging beats a banned tool used in secret.

5. Regulator and client expectations. The rules are evolving faster than most policies. Bar associations have issued guidance on lawyers' duties of competence and confidentiality when using generative AI. HHS has issued guidance touching AI in healthcare contexts. Clients — especially institutional ones — increasingly ask in security questionnaires whether you have an AI use policy. "We have not thought about it" is becoming an answer that costs work. You do not need to predict where regulation lands; you need a documented, defensible governance posture that shows you took the question seriously.

THE BALANCED VIEW

This is a governance problem, not a reason to abstain. Organizations that deployed security AI and automation extensively saw breach costs average $1.9M lower and lifecycles roughly 80 days shorter than those that did not. IBM Cost of a Data Breach Report 2025 The firms that get value from AI are the ones that put boundaries around it early — not the ones that banned it, and not the ones that ignored it.

02 — GOVERNANCE FRAMEWORK

Seven elements of an AI governance program

AI governance does not require a new department. It requires seven concrete artifacts, most of which extend processes you already run for any vendor or any data-handling change.

1. An AI acceptable-use policy. Short, specific, signed. Section 04 of this guide contains one you can adopt as written. A policy nobody can remember is a policy nobody follows — keep it to one page of clauses plus the approved-tools table.

2. An approved-tool list with data-class boundaries. The core of the program. Do not approve "AI" — approve specific tools for specific data classes. A simple four-class model covers most regulated firms:

DATA CLASS	EXAMPLES	WHERE IT MAY BE USED
Public	Published content, marketing copy, public filings	Any approved tool, consumer or enterprise tier
Internal	Templates, procedures, non-client drafts	Approved enterprise-tier tools only
Client confidential	Matter files, deal documents, client financials	Enterprise tools inside your tenant, with confirmed no-training terms and access controls
Regulated	PHI, material nonpublic information, privileged work product	Only tools explicitly contracted for it (e.g., BAA in place for PHI) — and for some matters, none

3. A vendor review gate. An AI tool is a vendor. Run it through the same third-party risk review you would apply to any system that touches client data: data-handling and retention terms, training-on-your-data language, subprocessors, breach notification obligations, security attestations, and — for healthcare — whether the vendor will sign a BAA. Worth the effort: third-party involvement appeared in 30% of breaches, double the prior year. Verizon DBIR 2025

4. A human-review requirement. No AI-assisted output goes to a client, a court, a regulator, or a permanent record without review by a person qualified to verify it. This is the single control that prevents the most embarrassing and most expensive failures. Write it down, train on it, enforce it.

5. Logging and audit trail. Prefer tools that live inside your identity perimeter — sign-in through your tenant, activity captured in your audit logs. If a tool offers an enterprise audit log, turn it on at deployment. When a client or regulator asks "who used AI on this matter," you want an answer that takes minutes, not a forensic project.

6. Training. One hour, role-specific, repeated annually: what the approved tools are, what data classes they may touch, what review is required, and what fabrication looks like in your discipline. The human element is involved in 60% of breaches Verizon DBIR 2025 — policy without training is paper.

7. Periodic review. Vendors change terms; staff find new tools; regulators publish new guidance. Review the approved list and the policy on a fixed cycle — quarterly is realistic for the first year, semiannual after that —

and assign the review to a named role (managing partner, compliance officer, practice administrator), not to "the firm."

> **SEQUENCE MATTERS**
>
> Adopt the policy first, then the tool list, then training — in that order, within one quarter. A tool rollout that precedes the policy teaches staff that the policy is optional.

03 — SECURE AI DEPLOYMENT (COPILOT)

Deploying Copilot without surfacing what you forgot you shared

For firms already on Microsoft 365, Copilot is usually the first enterprise AI deployment. It is also the one where preparation matters most — because Copilot is only as well-behaved as your permissions.

Why enterprise tenancy beats consumer tools. Microsoft 365 Copilot operates inside your tenant boundary: prompts and responses stay within your Microsoft 365 service boundary, inherit your compliance configuration, and appear in your audit logs. Enterprise agreements generally provide that your prompts, responses, and tenant content are not used to train the underlying foundation models — confirm the data-protection terms attached to your own tenant's agreement, since terms vary by product and tier and consumer versions are governed differently. That combination — contractual data boundaries plus your own audit trail — is what consumer chatbots cannot give you.

The permission-inheritance point, stated precisely. Copilot does not bypass permissions and does not grant anyone new access. It surfaces content the signed-in user can *already* reach. That is exactly the problem: most tenants carry years of quiet oversharing — site-wide links, "everyone" groups, old matter folders shared broadly and never cleaned up. Before Copilot, that oversharing was hidden by obscurity; nobody searched for it. Copilot removes the obscurity. A user who asks "summarize what we know about the Henderson matter" will get whatever their account can touch, including the folder someone overshared in 2021. The fix is not a Copilot setting. The fix is the permission cleanup you have been deferring.

> **RULE OF THUMB**
>
> If you would be uncomfortable with every employee running a perfect search across everything their account can access, you are not ready to enable Copilot tenant-wide. Fix the access first; the AI rollout becomes safe as a byproduct.

Pre-deployment checklist.

- [ ] **Tighten org-wide sharing defaults** 1
  SharePoint admin center → Policies → Sharing: set the default sharing link to "Specific people" and restrict or disable "Anyone" links unless a documented case requires them.
- [ ] **Run the data access governance reports** 2
  SharePoint admin center → Reports → Data access governance: review sites with broad sharing links and "Everyone except external users" exposure. Triage the worst sites first.
- [ ] **Audit high-sensitivity sites by hand** 3
  Matter files, HR, finance, and executive sites: review site membership and broken-inheritance folders. Remove "everyone" style groups from anything client-confidential.

- [ ] **Apply Purview sensitivity labels as guardrails** 4
  purview.microsoft.com → Information protection → Sensitivity labels: label regulated and client-confidential content. Labeled, encrypted content carries its protections into Copilot interactions and gives you a policy hook beyond folder permissions.
- [ ] **Clean up stale content** 5
  Archive or delete abandoned sites and former-employee OneDrives. Copilot reasons over what exists; outdated drafts surface alongside current versions and pollute answers.
- [ ] **Pilot with a small, watched group** 6
  10–25 licenses across roles and practice groups for 30–60 days. Ask pilots to actively hunt for content they should not see — and log what they find as permission defects to fix.
- [ ] **Enable and review audit logging** 7
  Confirm Copilot interaction events appear in Purview Audit. Decide retention before rollout, not after the first incident question.
- [ ] **Expand in waves, re-checking as you go** 8
  Broaden licensing by department after pilot findings are remediated. Re-run the data access governance reports before each wave.

The same sequence — fix access, label data, pilot, log, expand — applies to any tenant-integrated AI tool, not just Copilot. Vendors will differ; the discipline does not.

04 — ACCEPTABLE USE POLICY

AI acceptable use policy — ready to adopt

Replace the bracketed fields, complete the approved-tools table, and adopt. Keep it to this length — a policy staff can hold in their heads is the one they follow.

[Company] — Artificial Intelligence Acceptable Use Policy

Version 1.0 · Effective [date]

1. PURPOSE AND SCOPE

1.1 This policy governs the use of artificial intelligence tools — including generative AI, chat assistants, transcription, and AI features embedded in other software — in connection with [Company] work.

1.2 It applies to all employees, contractors, and temporary staff, on any device, whenever [Company] or client information is involved.

1.3 Where a client agreement, court rule, or regulator imposes stricter requirements, the stricter requirement controls.

2. PERMITTED USES

2.1 Approved tools (Section 4) may be used for drafting, summarizing, research assistance, brainstorming, and routine productivity tasks, within the data boundaries in Section 4.

2.2 Public and internal information may be used in any approved tool. Client-confidential and regulated information may be entered only into tools the table explicitly approves for that data class.

2.3 AI output may be used as a starting draft or research aid. It is never a finished work product.

3. PROHIBITED USES

3.1 Entering client-confidential, privileged, or regulated information (including PHI and material nonpublic information) into any tool not approved for that data class — including personal accounts on otherwise-approved products.

3.2 Delivering AI-generated content to a client, court, regulator, or permanent record without review and verification by a person qualified to assess it.

3.3 Submitting any citation, quotation, statistic, or factual assertion produced by an AI tool without independently verifying it against the original source.

3.4 Using AI tools to record, transcribe, or summarize meetings or calls without the consent required by law and by [Company] policy.

3.5 Disabling, bypassing, or circumventing security controls, logging, or data-loss-prevention measures applied to AI tools.

3.6 Representing AI-generated work as independently authored where a client, court, or regulator requires disclosure.

4. APPROVED TOOLS AND DATA BOUNDARIES

4.1 Only the tools listed below are approved. Requests to add a tool go to [role/email] and require completion of the vendor review process before any use with company or client data.

TOOL / TIER	HIGHEST DATA CLASS PERMITTED	ACCOUNT TYPE	NOTES

4.2 Approved tools must be accessed through [Company]-provisioned accounts. Personal accounts are out of scope of [Company]'s contracts and are not approved for any non-public data.

5. REVIEW AND DISCLOSURE

5.1 The author who submits or sends AI-assisted work product is responsible for its accuracy, as if they had written it entirely themselves.

5.2 Where a court, regulator, client agreement, or professional rule requires disclosure of AI use, staff must disclose. When in doubt, ask [role] before submitting.

5.3 Material AI involvement in client deliverables is disclosed to the client where [Company]'s engagement terms or the client's policies require it.

6. REPORTING AND VIOLATIONS

6.1 Suspected data exposure through an AI tool — yours or a colleague's — must be reported to [role/email] immediately. Prompt self-reporting is treated as a mitigating factor.

6.2 Violations may result in suspension of tool access and disciplinary action up to termination, consistent with [Company]'s employment policies.

6.3 This policy is reviewed at least [quarterly/semiannually] by [role], and whenever a tool, vendor term, or regulatory change warrants it.

Employee name _____

Signature _____

Date _____

Policy owner / approved by _____

ELEVATE SOLUTIONS

Elevate Solutions provides managed IT and cybersecurity for businesses that answer to regulators, clients, and courts.

Elevate Solutions

6230 Wilshire Blvd Suite 2120, Los Angeles CA 90048

(424) 400-6625 · support@elevatesolutions.io · elevatesolutions.io