M E M O

Example memo written as a portfolio sample. The audience is illustrative; all figures are sourced.

TO: Executive Team
FROM: Rashon Alsberry, Supply Chain and Third-Party Risk Analyst
DATE: September 27, 2026
RE: AI risk lives in the permissions, not the model

Bottom line

Your largest AI exposure is already inside the building. It's the tools your people adopted without approval, and the agents you're about to hand standing access nobody reviews. Governance is running behind adoption, and the losses land in that gap. Fix the inventory and the permissions first. Model risk comes after.

The signal

Adoption came in through employees, not procurement. 75% of knowledge workers use generative AI at work, and 78% of those users bring their own tools.[1]

Breaches follow the gap. Among organizations that had a breach involving their AI models or applications, 97% lacked proper AI access controls. 63% of breached organizations had no AI governance policy or were still developing one, and high levels of shadow AI added $670,000 to the average cost of a breach.[2]

Courts put the liability on you. Air Canada had to pay out on a refund policy its chatbot made up.[3] Attorneys were sanctioned for filing case citations ChatGPT invented.[4] "The AI said it" isn't a defense anyone's accepted.

Data leaves by paste. Samsung engineers put proprietary source code into ChatGPT. The ban came after the leak.[5]

None of these losses came from a model going rogue. Every one came through an ordinary permission (a login, a paste, a customer-facing answer) that nobody owned.

Why the old playbook misses it

Risk management was built for systems that do what they're told. AI does what it infers. And the paradox you'll have to hold is that the same tools that can watch risk for you are a new surface for it. Don't pick a side. Run both.

I've seen this shape before. In 2021 I wrote up the SolarWinds breach for a security course at George Mason. The attackers came in through a trusted vendor update, posed as legitimate users, and gave themselves admin rights.[6] The lessons the industry took from it were zero trust and least privilege, and assuming you're already breached. AI agents are that same shape with a friendlier interface: trusted software, acting as a user, holding more access than the task needs.

Recommendations

By day 30: Inventory. List every AI tool in use, approved or not, and every place company data enters one. Run it as an amnesty, not a crackdown. A ban only drives usage somewhere you can't see it. You can't govern a list you don't have. Treat every unapproved tool on that list as what it is, an unvetted vendor holding company data, and put it through the same third-party review as any other supplier.

By day 60: Permissions. Every AI system and agent gets its own identity, the least access its job requires, and a named human owner. No agent holds standing admin rights. Log what each one does.

By day 90: Verification gates. Any AI output that reaches a customer, a court, a contract, or a payment gets a human check before it leaves the building. Everything else can move fast.

Ongoing: Borrow the framework, don't invent one. NIST's AI Risk Management Framework (Govern, Map, Measure, Manage) is free and maps cleanly onto the steps above.[7] ISO/IEC 42001 gives you a certifiable standard if customers start asking.[8] If you sell into Europe, EU AI Act obligations are phasing in now: most rules have applied since August 2026, and the high-risk rules follow in December 2027 and August 2028.[9]

What would change this read

- The inventory turns up only a handful of unapproved tools. Then permissions become the whole job.
- Your insurer starts pricing AI exclusions or offering governance discounts. Then this moves from operations to finance, quickly.
- A regulator in our sector publishes AI-specific rules. Re-scope against those.

The question for this team

If an AI agent took an action tomorrow that cost us a client, could we name the person who approved its access?

Sources

1. Microsoft & LinkedIn, *2024 Work Trend Index Annual Report: "AI at Work Is Here. Now Comes the Hard Part,"* May 8, 2024. https://www.microsoft.com/en-us/worklab/work-trend-index/ai-at-work-is-here-now-comes-the-hard-part

2. IBM & Ponemon Institute, *Cost of a Data Breach Report 2025*, July 30, 2025. https://www.ibm.com/reports/data-breach. Key findings: https://newsroom.ibm.com/2025-07-30-ibm-report-13-of-organizations-reported-breaches-of-ai-models-or-applications,-97-of-which-reported-lacking-proper-ai-access-controls

3. *Moffatt v. Air Canada*, 2024 BCCRT 149 (British Columbia Civil Resolution Tribunal), February 14, 2024. https://www.canlii.org/en/bc/bccrt/doc/2024/2024bccrt149/2024bccrt149.html

4. *Mata v. Avianca, Inc.*, No. 22-cv-1461 (PKC), Opinion and Order on Sanctions (S.D.N.Y. June 22, 2023). https://law.justia.com/cases/federal/district-courts/new-york/nysdce/1:2022cv01461/575368/54/

5. Mark Gurman, "Samsung Bans Staff's AI Use After Spotting ChatGPT Data Leak," *Bloomberg*, May 2, 2023. https://www.bloomberg.com/news/articles/2023-05-02/samsung-bans-chatgpt-and-other-generative-ai-use-by-staff-after-leak

6. CISA, Advisory AA20-352A, *Advanced Persistent Threat Compromise of Government Agencies, Critical Infrastructure, and Private Sector Organizations*, December 17, 2020 (updated 2021). https://www.cisa.gov/news-events/cybersecurity-advisories/aa20-352a

7. NIST, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, January 2023. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf. *Generative Artificial Intelligence Profile*, NIST AI 600-1, July 2024. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

8. ISO/IEC 42001:2023, *Information technology — Artificial intelligence — Management system*, December 2023. https://www.iso.org/standard/42001

9. Regulation (EU) 2024/1689 (AI Act), in force August 1, 2024. https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng. Amended by Regulation (EU) 2026/1744 ("AI Omnibus"), in force July 27, 2026, which moved high-risk obligations to December 2, 2027 (Annex III) and August 2, 2028 (Annex I). https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng. Timeline: https://ai-act-service-desk.ec.europa.eu/en/ai-act/timeline/timeline-implementation-eu-ai-act