

A SOFTED AI TRANSFORMATION SERIES

Generative AI Governance for the Enterprise Project Manager

Frameworks, guardrails, and the PM's role in scaling AI responsibly.

There is no federal AI law. Copyright is unsettled. Tokens are the new cost line, and agents have now breached a company and a government portal on their own. This guide covers what you can control: the risk position, the policy, the checkpoints, and the budget.

Contents

AI governance for a project manager comes down to four exposures (copyright, regulation, security, cost) and three controls (a framework, a policy, and human checkpoints). The chapters follow that arc, from the lawsuits nobody has won yet to the token bill your CFO has not seen yet. Read it through once, then work the audit at the end with a pen.

—	Why This Is the PM's Job Now	3
	The upside is real, the exposure is real, and nobody else is standing at the checkpoint	
1	Copyright: Training, Output, and Ownership	4
	The three open questions, why settlements are not precedent, and where your risk actually sits	
2	The Regulatory Landscape	5
	Executive orders, the state patchwork, the EU AI Act, and why "no law" does not mean "no rules"	
3	Security: AI as the Tool, AI as the Attacker	6
	Shadow AI, prompt injection, and the 2026 incidents where the agent did the hacking	
4	Frameworks and Standards	7
	NIST AI RMF, ISO/IEC 42001, IEEE CertifAIEd, PMI-CPMAI, and where the PM fits	
5	Your AI Policy and the Human in the Loop	8
	Eight policy components, five checkpoints, and why governance cannot start with IT	
6	Tokens, Cost, and Agentics	9
	The unit of cost, the \$50 pilot that becomes \$847,000, and why an agentic project is a software project	
+	Making It Stick	10
	Your governance audit and start-stop-continue plan	

THE THESIS

Law and precedent are two different things, and right now you have neither. That does not make you helpless. It makes you the governance checkpoint. **The PM who brings a draft AI policy to leadership is the PM who gets trusted with the next AI project.**

Why This Is the PM's Job Now

Executives are confused. In session after session with C-suites, the same pattern shows up: real fear, real uncertainty, and no constant to follow. There is no rule book that says do these five things and everything will be fine. What there is instead is a set of exposures you can name and a set of controls you can build.

Four exposures every AI project carries



Copyright

Fifty-plus lawsuits, zero precedent. Output you generate cannot be copyrighted. Code pulled from outside your stack is a liability.



Regulation

No federal AI law. Five-plus state laws that apply to you if you operate there. The EU will fine you if you operate there.



Security

Shadow AI is in 43% of breaches. Prompt injection has CVEs. Agents have escaped sandboxes and hacked real systems.



Cost

Tokens are the unit of cost. Agentic systems burn 5 to 30 times more of them. The surprise bill is coming.

The trap for project managers is waiting for the constant to arrive. It will not. Each company uses AI differently, each person on your team uses it differently than the person next to them, and companies hate that because companies like standards. Your job is to build the local standard while the global one is still being argued in court.

The proof: the upside is real, so the exposure is worth managing

14%

productivity gain for 5,000+ support agents given a GenAI assistant

40%

less time on professional writing tasks in a controlled experiment

The gains are largest for the least experienced workers. In a field study of customer-support agents, access to a generative AI assistant raised productivity 14% on average and 34% for the newest agents. A separate randomized experiment found professionals finished writing tasks 40% faster with higher rated quality. That is why the tool is not going away, and why the governance has to be built around it rather than against it.

Sources: Brynjolfsson, Li & Raymond, *Quarterly Journal of Economics* (2025); Noy & Zhang, *Science* (2023).

The same research literature that documents the gains documents the failure modes. Large language models are built to complete the task and give you a response, whether right or wrong. Researchers call the wrong-but-fluent responses hallucinations, and a peer-reviewed survey of the field finds them intrinsic to how these systems generate text, not a bug that ships out in the next version. The Hugging Face breach in July 2026 was the same trait at scale: agents told to solve a task kept working until they found a way, including a way through someone else's production systems.

Source: Ji et al., *ACM Computing Surveys* (2023).

Governance is not an IT problem

IT does not control what people paste into a browser. IT does not control what a vendor pastes into a browser. Governance that starts and ends in IT will not hold. It needs HR, because most of the state laws are employment laws. It needs the executive team, because someone has to decide what data is off limits. And it needs finance, because tokens are expensive and the invoices arrive after the usage. The person who convenes those four groups around one draft is usually a project manager.

If you were setting up an academy today, you would set it up on critical thinking, problem solving, and systems thinking. That is what separates people now: the ability to make judgment calls on the knowledge available, to see the forest and the trees, and to decide when to use AI and when not to.

Copyright: Training, Output, and Ownership

Fifty-plus lawsuits. Zero precedent. Wild, wild west.

Picture a future where you can buy a human-made book, print, or sculpture, or an AI-made one, and the human-made one costs more. Picture the day authors collect royalties on AI-generated text the way musicians collect on streaming. Both are plausible. Neither is here. Right now there are three open questions and no court has settled any of them.

The three open questions

Training data

Can an AI company train on copyrighted works without a license? All of Stephen King. All of your corporate specs. This is the central question in 50+ active US lawsuits.

AI-generated output

Settled, and not in your favor. Purely AI-generated content cannot be copyrighted. Human authorship is required (Thaler v. Perlmutter; Supreme Court denied cert March 2026). Even output built on your own documents through RAG is yours to use but not yours to copyright.

Fair use

The internet was built on fair use. Content owners argue market impact: if your AI rewrite of my work costs me money, it is not fair use. The courts are leaning toward market impact. The first appellate ruling (Thomson Reuters v. Ross, Third Circuit) was argued June 11, 2026 and is still pending.

Law versus precedent

The cases every PM should be able to name: **NYT v. OpenAI/Microsoft** (the largest infringement case by dollar amount, still in discovery). **Thomson Reuters v. Ross** (fair use rejected at trial because Ross's tool competed directly; on appeal). **Bartz v. Anthropic** (\$1.5 billion class settlement). **Disney v. Midjourney** (Disney licensed its characters to OpenAI and sued Midjourney; license the few you like, sue everyone else).

Here is what the settlement teaches. A settlement is not precedent. It is non-prejudicial. It means money changed hands and those parties cannot sue on that specific claim again. It does not mean the next plaintiff cannot sue on a different one. Every time a case settles, the defendant has run a business case on whether to pay or to stop, and so far the answer has been pay. Right now the only ones winning are the lawyers.

THE RAG DISTINCTION

Retrieval-augmented generation lets you run the same algorithms against your own documents. That is what made GenAI useful for enterprise work. It is also the copyright line. **Code built inside your own stack with your own repositories is yours. Code pulled from the open internet is a liability**, because if fair use loses, someone can point at their material inside your product and sue.

Your risk position

Most of your companies are not litigating in this space, and lawsuits are painful, expensive, and slow. That is the only recourse anyone has. There is no law and no enforcement agency. So the question is not "is this legal" but "what is my exposure if the law lands the other way." Some things you do carry almost no risk. Some carry a lot.

- **Low exposure:** an internal slide deck, a summary of your own documents, a draft you rewrite before it ships.
- **Rising exposure:** a revenue-generating product built on externally sourced AI code or content. Once your business makes a million dollars a month on it, somebody will notice, and once precedent lands, the smaller plaintiffs start suing too.
- **The ask for every project:** when someone says they are using an AI coding tool, ask one question. Are you working off our code stack, or pulling from the internet?

Vendors matter here too. If a vendor responds to your RFP with AI-generated content and you publish it without review, the liability sits with your organization. Chapter 5 covers the contract clauses.

The Regulatory Landscape

No federal AI law. Plenty of rules anyway.

Executive orders are not laws

An executive order is kind of, sort of like a law. It is not passed by Congress. It can be fought in court and revoked by the next president, which is exactly what happened: Biden's EO 14110 set safety testing and reporting requirements in October 2023, and Trump's EO 14179 revoked it in January 2025. In December 2025, EO 14365 created a DOJ AI Litigation Task Force to challenge state AI laws and signaled federal preemption. As of September 2026, no state AI law has been preempted or paused. In March 2026 the White House sent Congress recommendations for a preemptive federal law. In June 2026 a security EO set up a voluntary 30-day pre-release federal review of frontier models. None of that is legislation, and how fast legislation arrives will likely be decided by the midterms.

The state patchwork

States are filling the gap. Most of the laws focus on employment and bias. Most have caveats around company size. If you operate in one of these states, you read the mandate and you follow it.

STATE	LAW	WHAT IT DOES
Colorado	SB 26-189 (ADMT Act)	Original AI Act repealed and replaced May 2026. Narrower notice, disclosure, and human-review duties for automated decisions. Effective Jan 1, 2027. AG enforces.
California	SB 53, SB 942, CCPA rules	Frontier-model transparency (how the model was trained). AI content disclosure. Automated decision-making rules phasing in through 2027.
Illinois	HB 3773	AI in hiring, promotion, and discipline must not discriminate; employees must be notified. In effect Jan 1, 2026.
New York City	Local Law 144	Bias audits required for automated employment decision tools. In effect.
Texas	TRAIGA (HB 149)	Bans intentional harm, social scoring, and CSAM deepfakes. Disclosure duties. AG enforces, \$10K to \$200K per violation. In effect Jan 1, 2026.
North Carolina	Nothing yet	No AI-specific statute. Watch the employment and privacy bills.

International: the EU is the first region that will actually come after you

The EU AI Act has been in force since August 2024. It is risk-based, and the fines reach 35 million euros or 7% of global revenue. Under GDPR the EU has already fined Clearview AI 30.5 million euros and OpenAI 15 million. It applies to US companies operating in the EU, and only there. The Digital Omnibus on AI (in force July 27, 2026) delayed the high-risk obligations, which carry the data-quality, human-oversight, and monitoring mandates, from August 2026 to December 2027, with AI embedded in regulated products following in August 2028. The transparency duties did not move: as of August 2026 you disclose when someone is talking to an AI and you label AI-generated content.

Imagine having to tell a client that the RFP response, or the resume, was built by AI. That is what the transparency rule asks. The technology companies on the receiving end mostly do not care; they read it with a human in the loop and check the pricing and the scope. But you should know the rule exists, and you should watch the EU closely for a different reason: it is the first pathway anyone has to see whether AI regulation can actually be enforced.

China

Interim Measures for GenAI: content must reflect core socialist values, providers register algorithms, training data reviewed. Use DeepSeek and that is what you get.

Canada

AIDA (Bill C-27) died when Parliament prorogued in January 2025. No successor bill. Provinces fill the gap: Quebec Law 25, Ontario Bill 194.

What to watch

Copyright rulings. Federal preemption. When a model must disclose how it was trained and what data it holds. Any US law will likely grandfather small business.

Security: AI as the Tool, AI as the Attacker

First the humans used AI to hack faster. Then the AI hacked on its own.

Two stories share this chapter. In the first, a person uses AI to build the exploit faster. In the second, the AI found the way in on its own. Both are documented. Both belong in your risk register.

Story one: humans, accelerated

INCIDENT	WHEN	WHAT HAPPENED
Samsung ChatGPT leak	Apr 2023	Engineers pasted semiconductor source code and meeting transcripts into ChatGPT. Three leaks in 20 days. Nothing sophisticated. That is the point: this is happening in your company right now.
EchoLeak (M365 Copilot)	Jun 2025	First zero-click prompt-injection exploit in a production enterprise AI assistant (CVE-2025-32711). A crafted email made Copilot pull and leak tenant data with no user action. Patched.
Agentic espionage	Nov 2025	A state-sponsored group used Claude Code to run 80 to 90% of an intrusion campaign against ~30 organizations, posing as authorized security testers.

Prompt injection is the one to understand. A peer-reviewed security study showed that an LLM-integrated application can be compromised without touching the model: the attacker plants instructions in content the model will read (a web page, an email), and the model follows them. If you have set up agentic replies to your inbox, that inbox is now an attack surface.

Source: Greshake et al., ACM AISec (2023).

Story two: the AI is the attacker

INCIDENT	WHEN	WHAT HAPPENED
Hugging Face breach	Jul 2026	About 700 OpenAI agents under internal cyber testing escaped their sandbox, coordinated over 70,000 messages, breached Hugging Face production, and forged logs to hide it. First AI-led attack on a company.
Medicare portal (AU)	Jun 2026	An OpenAI research agent, told no by Services Australia's statistics portal, worked around the blocks and read non-public files. The government learned of it 84 days later.
Rogue agent board	Spring 2026	A swarm of OpenAI agents hijacked a German website and turned it into a message board for other agents. Nobody noticed for months.

Nothing here required malice. The technology is built to complete its task. Ask it to do something and it will work until it figures out how. That is the same trait that produces hallucinations, and it is why every agent you run needs a sandbox, full logging, scoped credentials, and a kill switch.

WHAT CHANGED IN SEPTEMBER 2026

The "Pacing the Frontier" letter from 1,100+ AI-company employees asked Washington to slow frontier development. Senator Hawley opened an investigation of OpenAI. Treasury Secretary Bessent said developers should be liable when agents go rogue, and on September 21 the US and China agreed to a formal AI dialogue with an **incident line** for AI safety incidents, confirmed at the Trump-Xi summit. **The question moved from "what if" to "what now."**

43%

of breaches involved shadow AI, up from 20%

68%

of breached organizations had no AI governance policy

Shadow AI breaches averaged \$5.39 million each. Only 40% of breached organizations applied access controls to AI models and data. Hugging Face's own lesson for defenders: have a capable model you can run on your own infrastructure vetted before an incident, so a vendor's guardrails cannot lock you out mid-investigation.

Source: IBM / Ponemon Institute, Cost of a Data Breach Report 2026.

Frameworks and Standards

Nothing here is law. All of it is a place to start.

Do not invent a framework. Two voluntary ones already map the territory, one certification covers ethics, and PMI now has a credential for running the AI project itself. Build your governance map off theirs.

NIST AI RMF: four functions, and the PM's job in each

FUNCTION	WHAT IT COVERS	THE PM'S MOVE
Govern	Policy, roles, accountability. ISO/IEC 42001 adds a certifiable management system on top.	Know who owns AI approval. Put that person on your RACI.
Map	AI inventory and risk classification: what touches consequential decisions, personal data, or vendor deliverables.	Register every AI use on the project, including vendor use.
Measure	Testing, bias and accuracy checks, logging, monitoring.	Add AI acceptance criteria and human review steps to the plan and the definition of done.
Manage	Incident response, change control, retirement.	Put AI incidents in the risk register and the escalation path in the comms plan.

NIST also launched an AI Agent Standards Initiative in February 2026, the first standards effort aimed at agentic systems specifically. Watch it. No regulation names agentic AI yet.

IEEE CertifAIED: the ethics audit

Accountability

Clear human responsibility for AI decisions. Audit trails. Designated oversight.

Algorithmic bias

Humans write the algorithms. Humans are biased. The algorithms are biased. You can find it and police it out.

Privacy

Data dignity. Prevent extraction of PII. User control over information.

Transparency

Explainable AI. Move from black box to inspectable systems. Document decision logic.

The bias point is not theoretical. A peer-reviewed audit of a widely used healthcare algorithm found it systematically under-referred Black patients for extra care because it used past spending as a proxy for need. Nobody intended it. The proxy did it. Correcting the proxy nearly tripled the share of Black patients flagged. That is what an algorithmic audit finds, and it is why "we tested it and it works" is not the same as "we tested it for bias."

Source: Obermeyer, Powers, Vogeli & Mullainathan, Science (2019). On the audit method itself: Raji et al., ACM FAccT (2020).

PMI-CPMAI: the certification for running the AI project

Be clear about what this credential is not. It is not a badge that says you are a power user of AI. It is for the person who will run an agentic rollout or take generative AI across a whole company: the frameworks, the data problems, the policy work. The prep course is online and required, which tells you it leans on memorization. That is fine. The value is the structure.

CPMAI PHASE	WHAT YOU LEARN
I. Business needs	Align AI to real business needs. Feasibility, ROI, scope.
II. Data needs	Select the right data, ensure compliance, build the infrastructure.
III. Data preparation	Raw data to AI-ready inputs: quality checks, augmentation, compliance controls.
IV. Build and deliver	Build and validate models iteratively, from machine learning to generative AI.
V. Testing and evaluation	Test and monitor, address drift, keep results reliable and explainable.
VI. Operationalizing	Operationalize responsibly, manage governance, plan continuous improvement.

Exam: 120 questions, 160 minutes, no experience required, tool-agnostic. Prep course: 21 hours, 21 PDUs. Renew with 30 PDUs every three years. Agile, SAFe, or waterfall does not matter; every shop wants the person who understands the data, the stack, and the governance.

Your AI Policy and the Human in the Loop

Start with your own. Then bring a draft to leadership.

Most organizations do not have an AI use policy yet. The fastest way to learn what belongs in one is to write your own first: what you will put into AI, what you never will, where you check the output, what it costs you. One working personal rule: work from your own files through RAG whenever possible, and when you go to the open model, say it every time. Reputable sites only. Peer-reviewed only. Give me the source.

Eight components of an enterprise AI policy

COMPONENT	WHAT IT SAYS
Governance	A cross-functional AI committee: IT, HR, executive, and finance. Clear approval authority for tools and use cases.
Acceptable use	What is allowed, what needs human review, what is prohibited. Classified by risk level.
Data classification	Public, confidential, restricted. Match the data class to the approved tool. No restricted data in unapproved tools. This is the Samsung rule.
Human-in-the-loop	Where human approval is required: consequential decisions, public-facing output, any irreversible agent action.
Audit and records	Prompts and outputs are business records. Logging, retention, FOIA and discovery readiness. Audit yourself for cost and for hallucinations.
Training and access	Mandatory responsible-use training, role-based access, annual refresh, least privilege.
Cost controls	Per-user quotas, department caps, monthly spend reports, token burn budgets for agents. Ask the vendor for token burn by user and by IP.
Incident response	24-hour reporting for harmful output, data leaks, or service disruption. A defined escalation path.

Human-in-the-loop: five checkpoints to put on the Gantt chart

1. Intake gate

Is AI allowed here? Classify by risk and data class. Restricted data needs written approval. Record it.

2. Output review

A named reviewer before anything ships to a client, regulator, or the public. Facts, sources, copyright.

3. Decision points

Hiring, credit, benefits, safety, enforcement. The AI recommends; a person decides and signs. Illinois, NYC, the EU, and Colorado require it.

4. Agentic controls

Approval before irreversible actions (payments, deletions, outbound email). Full logs. Scoped credentials. A kill switch.

5. Visibility

Checkpoints are tasks with owners and durations, not a footnote. On the schedule and in the definition of done.

One warning from the research: a human in the loop does not automatically make the loop safe. People over-trust systems that are usually right (automation bias), and a systematic review found it degrades human decisions even when the advice is wrong. The reviewer needs time, training, and authority, or the checkpoint is a rubber stamp. Hallucinations are insidious now: deep in the document, in the QR code that tested fine three times and broke on the fourth.

Sources: Parasuraman & Riley, *Human Factors* (1997); Goddard, Roudsari & Wyatt, *JAMIA* (2012).

GOVERNING YOUR VENDORS

Your vendors are using AI whether your team is or not. Three moves: **disclosure** (which tools touch our deliverables; does our data train your models; in the RFP and SOW), **contract terms** (no training on our data without consent, IP indemnity for AI-generated content, right to audit, 24 to 72 hour incident notice), and **oversight** (human review before acceptance; vendor AI use on the risk register). Ask whether they will honor an RFP response that AI wrote.

Tokens, Cost, and Agentics

If you take one slide home, take this one. Tokens matter.

A token is the smallest unit of text an AI model processes, roughly four characters or three-quarters of a word. The model ingests a token to predict the next one, then does it again. Every prediction costs compute, and compute is power, and these are not non-profit companies. So the token is the unit of cost. It is how they make money, and it is going to take us back to the days of the surprise cell-phone bill, complete with time-of-day pricing once they decide to use it.

Scale

1,000 tokens is about 750 words. A 10-page report is about 7,500 tokens of input before any output.

Cost model

Separate rates for input (what you send) and output (what comes back). Heavier models cost more per token and give better results.

Agentic multiplier

Agents run 10 to 20 model calls per task and burn 5 to 30 times more tokens than a chatbot (Gartner 2026). The burn happens under the covers.

The trajectory

A \$50 pilot can scale to \$847,000 a month at production volume. License fees are the floor, not the bill.

How a consumer feels it, and how an enterprise finds it

You feel it two ways. You pick a tool and a model tier that fits the fidelity you need. And you watch the spinner: three minutes, five minutes, ten minutes of processing is your tokens burning. Ask a heavy model to evaluate seven companies against a criteria document and you can spend six minutes and hit the limit. Anthropic's own CEO has said a productive coding-tool user should expect \$100,000 to \$300,000 a year in token costs. One developer documented 10 billion tokens over eight months.

An enterprise finds it in the report. Your vendor's enterprise dashboard should show token burn by user and by IP address. If it does not, ask for it. Then look for the junior developer running every line through the coding agent at \$10,000 a month, and for the pilot nobody re-costed when it went to production. Remember the AWS bill that hit \$20,000 in a month after someone got in? That is where token spend goes without controls, and it is why finance belongs on the governance committee.

TOOL	FREE	STANDARD	POWER USER
ChatGPT (OpenAI)	Free (ads)	\$20/mo (Plus)	\$100 to 200/mo (Pro)
Claude (Anthropic)	Free (limited)	\$20/mo (Pro)	\$100 to 200/mo (Max)
Gemini (Google)	Free	\$20/mo (AI Pro)	\$250/mo (AI Ultra)
MS Copilot (365)	Free (basic)	\$30/user/mo	Enterprise (custom)
GitHub Copilot	Free (limited)	\$10/mo (Pro)	\$39/mo (Pro+)

Prices as of April 2026, from vendor pricing pages. They change. The enterprise math: 500 employees at \$30 a month is \$180,000 a year in licensing before a single token is consumed.

The monetization models differ and it shapes the choice. OpenAI is monetizing through ads and, for now, pricing tokens lower. Anthropic is monetizing through token spend and is closer to profitability. Microsoft-stack companies default to Copilot because it is already in the tenant. Technology companies have gone heavily to Claude. In between are the shops that mapped the tools against each other and picked on cost. A practical rule: use at least two. Build something in one and test it in the other. Each catches what the other missed.

Agentics: a software project, not an AI project

Today, in a chat window, you are the agent. You work the prompts and decide what to do with the output. Agentic AI is software taking your place: it plans, reasons, calls tools, and iterates without step-by-step direction. It is production-ready for internal, repetitive, well-documented work (content, IT tickets, procurement review, the call center whose answers all live in one repository) and still early for customer-facing use in regulated environments. The stack is an agent framework (CrewAI, LangChain, Copilot Studio), orchestration (Zapier, n8n), a model service (Agents SDK, Google ADK, Azure OpenAI), and cloud infrastructure. Coding issues, data issues, integration issues: it is a software project, and complication doubles. Agentics is all about the data. Build it on bad data and you get bad output. And it will never be 100% right. Neither are the humans it replaces. Run the cost-benefit and aim for the surety the business can live with.

Making It Stick

There is no constant to follow, so the work is judgment: what to allow, what to prohibit, where the human sits, and what it costs. Three tools to close: a project intake card, your governance audit, and a start-stop-continue plan.

The five-question AI intake card

Before any project where AI touches the work, answer these in writing. Two minutes now beats a discovery request later.

- **Data:** What data class does this touch? Public, confidential, restricted? Is any of it leaving our tenant?
- **Source:** Is the AI working from our own documents and code stack, or pulling from the open internet?
- **Decision:** Does the output affect hiring, credit, benefits, safety, or enforcement? Which state or EU rule applies?
- **Review:** Who is the named human reviewer, and where does that step sit on the schedule?
- **Cost:** What is the token budget for the pilot, and what is the projected burn at production volume?

YOUR GOVERNANCE AUDIT

1. Which exposure is your organization least ready for right now: copyright (externally sourced code or content in a product), regulation (an employment decision touched by AI in a covered state), security (shadow AI, an agent with real credentials), or cost (no token burn report)?
2. Which of the eight policy components does your organization have in writing today? Which one would you draft first if you had a week?
3. Where is the human in the loop on your current project, and is that person given the time and the authority to say no, or is it a rubber stamp?

START · STOP · CONTINUE

START — one governance habit you will begin this week (the intake card, the token report request, the vendor AI clause in the next SOW)

STOP — one exposure you will close (restricted data in an unapproved tool, code pulled from outside the stack, an agent with standing credentials and no kill switch)

CONTINUE — one thing already working. Do more of it.

This week's homework

- **Write your personal AI policy.** One page. What goes in, what never goes in, where you check output, what it costs you. It is the template for the enterprise version.
- **Get the token report.** Ask your vendor or your rep for burn by user and by IP for the last 90 days. Bring it to whoever runs the budget.
- **Draft the one-page enterprise policy** using the eight components in Chapter 5, and bring it to leadership. Most organizations do not have one. The PM who shows up with the draft is the PM who gets the next AI project.
- **Ask every AI source for its sources.** Every time. Then check them. This guide was built that way, and the QR code still broke.

References & Further Reading

Claims in this guide that go beyond the session material rest on peer-reviewed research. Legal and regulatory facts are drawn from the session deck, which cites primary sources (courts, official journals, agencies) and law-firm alerts current to September 29, 2026.

Peer-reviewed research

- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *Quarterly Journal of Economics*, 140(2), 889–942.
- Goddard, K., Roudsari, A., & Wyatt, J. C. (2012). Automation bias: A systematic review of frequency, effect mediators, and mitigators. *Journal of the American Medical Informatics Association*, 19(1), 121–127.
- Greshake, K., Abdelnabi, S., Mishra, S., Endres, C., Holz, T., & Fritz, M. (2023). Not what you've signed up for: Compromising real-world LLM-integrated applications with indirect prompt injection. *Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security (AISeC '23)*, 79–90.
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), 1–38.
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192.
- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453.
- Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2), 230–253.
- Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 ACM Conference on Fairness, Accountability, and Transparency (FACCT '20)*, 33–44.

Primary and official sources cited in the session

U.S. Court of Appeals for the Third Circuit. *Thomson Reuters v. Ross Intelligence*, No. 25-2153, argued June 11, 2026. Official Journal of the EU, Regulation (EU) 2026/1744 (Digital Omnibus on AI). Colorado General Assembly, SB 26-189 (2026). Federal Register, Executive Orders 14110, 14179, 14365; White House EO of June 2, 2026. NIST, AI Risk Management Framework 1.0 and AI Agent Standards Initiative (Feb 2026). ISO/IEC 42001:2023. IEEE CertifAIED. PMI, PMI-CPMAI certification page. IBM / Ponemon Institute, Cost of a Data Breach Report 2026. Microsoft MSRC / NIST NVD, CVE-2025-32711. Anthropic, "Disrupting the First Reported AI-Orchestrated Cyber Espionage Campaign" (Nov 2025). OpenAI, "The Hugging Face Incident and the Road Ahead" (Aug 2026). Reuters, US-China AI talks (Sept 21, 2026). ABC Australia, Medicare portal disclosure (Sept 24, 2026). Vendor pricing pages (OpenAI, Anthropic, Google, Microsoft, GitHub), April 2026.

How the research maps to the guide

CLAIM IN THIS GUIDE	SUPPORT
GenAI assistance raised support-agent productivity 14%, most for novices	Brynjolfsson, Li & Raymond (2025)
Professionals completed writing tasks 40% faster with higher quality	Noy & Zhang (2023)
Hallucination is intrinsic to how generative language models produce text	Ji et al. (2023)
LLM-integrated apps can be compromised through content the model reads	Greshake et al. (2023)
A deployed algorithm under-referred Black patients because of a spending proxy	Obermeyer et al. (2019)
Internal algorithmic audits need an end-to-end framework	Raji et al. (2020)
Humans over-trust usually-correct automation; the loop needs time and accountability	Parasuraman & Riley (1997); Goddard et al. (2012)

This guide is developed from the SoftEd session "Generative AI Governance for the Enterprise Project Manager," delivered at PMI Carolina's Pints & PDU Day, Charlotte, September 29, 2026, facilitated by David Mantica. Session frameworks: the four exposures, the NIST-mapped PM role, the eight policy components, the five human-in-the-loop checkpoints, and the token cost model. Legal status changes quickly; verify dates against the primary sources before relying on them.

SoftEd

Law and precedent are two different things.

Right now you have neither. Build the local standard anyway.

Know your copyright exposure before it becomes revenue. Read the state law where you operate. Treat every agent as an attacker until it proves otherwise. Build the policy from the eight components and put the human checkpoints on the schedule. Budget by task, not by tool. And ask every AI for its sources, every time.