THE AI READINESS GAP REPORT

2026

What the World's Top Consulting Firms Say Organizations Need —
And What Mid-Market Leaders Are Actually Doing About It

By Kim Nelson-Wright
Co-Founder & CEO | PreFlight AI™
Clarity Before Automation™

June 2026

EXECUTIVE SUMMARY

Nine of the world's most respected research organizations — McKinsey, Deloitte, KPMG, PwC, BCG, Accenture, Stanford HAI, EY, and Microsoft — have each spent tens of millions of dollars surveying thousands of executives across dozens of countries. Their conclusions are remarkably consistent. And they point to the same problem: organizations are deploying AI faster than they can govern it.

This report synthesizes those findings through a single lens that the consulting giants cannot provide: the practitioner perspective. After 20 years of sitting inside healthcare, government, and nonprofit organizations as the person responsible for making operational decisions hold up under pressure, I can tell you that the readiness gap these reports describe is real, urgent, and entirely solvable — if organizations act before deployment rather than after the first incident.

The five findings every report agrees on:

- **Only 1 in 5 organizations has mature AI governance,** even as 73% cite data privacy as their top concern.
- **Adoption is racing ahead of readiness.** Worker access to AI has grown by 50% in one year. Governance infrastructure has not kept pace.
- **Named accountability is the single highest-impact governance action.** Organizations with clear, named owners of AI decisions achieve measurably better outcomes than those without.
- **The agentic AI governance gap is the most urgent.** 23% of organizations use agentic AI today. 74% plan to within two years. Only 21% have a mature governance model for it.
- **Governance is a revenue variable, not a cost center.** 74% of all AI-generated economic value is captured by the 20% of organizations that invest most heavily in governance.

> *The organizations that will lead in AI are not the ones that move fastest. They are the ones that move with clarity. Clarity Before Automation™ is not a slogan. It is a discipline.*

WHO THIS REPORT IS FOR

The nine reports synthesized here were written for enterprise audiences — Fortune 500 companies with dedicated AI governance teams, legal departments, and risk functions staffed with specialists. The findings, however, apply with equal or greater urgency to mid-market organizations: companies with 50 to 500 employees in healthcare, government, nonprofit, and regulated industries that face the same compliance obligations with a fraction of the infrastructure.

This report is for the COO who is being asked to 'lead the AI strategy' without a roadmap. The CHRO who just learned their applicant screening tool qualifies as high-risk under the EU AI Act. The executive director whose board approved an AI initiative before anyone asked whether the organization was ready to run it.

If you are making AI decisions inside a mid-market organization without a dedicated governance function — this report was written for you.

WHAT 9 GLOBAL REPORTS AGREE ON: THE FIVE FINDINGS

FINDING 1: Only 1 in 5 Organizations Has Mature AI Governance

21%	**73%**	**8%**
of organizations have mature AI agent governance models	cite data privacy and security as their top AI risk concern	maintain a comprehensive AI governance framework
— Deloitte & KPMG, 2026	*— Deloitte, 2026*	*— Economist Impact, 2026*

The gap between concern and action is where the liability lives. Nearly every organization surveyed by Deloitte, KPMG, and the Economist Impact is aware that AI poses real governance risks. Fewer than one in ten has built the comprehensive framework to address them. The rest are governing by intention, not infrastructure.

FINDING 2: Adoption Is Racing Ahead of Readiness

Worker access to AI has grown from under 40% to roughly 60% in a single year — a 50% expansion in reach. Investment is surging, with organizations planning an average of $186 million in AI spending over the next 12 months according to KPMG. Yet Deloitte's 2026 report finds that only 30% of organizations believe their risk and governance infrastructure is highly prepared for what they're deploying.

> *"Enterprise AI adoption is broadening faster than enterprise AI integration. Many organizations now have approved tools, live use cases, and senior-level support. Fewer have made the organizational changes required to turn those ingredients into consistent business value." — Deloitte State of AI in the Enterprise, 2026*

The pattern is consistent across all 10 reports: strategy and tool access improve quickly. Data management, governance structure, and talent readiness lag significantly behind. Organizations are building the plane while flying it — and governance is the part that hasn't been assembled yet.

FINDING 3: Named Accountability Is the Single Highest-Impact Governance Action

Across McKinsey, Deloitte, and KPMG research, one factor consistently separates organizations with high governance maturity from those without: the presence of named, accountable ownership for AI decisions.

> *"Organizations with clear, named accountability for responsible AI achieve measurably higher maturity scores than those with diffuse or unclear ownership. Accountability structure correlates with governance outcomes." — McKinsey AI Trust Maturity Survey, 2026*

This is the finding that most mid-market organizations underestimate. They believe governance requires a team, a platform, or a budget. What McKinsey's data shows is that it begins with a name. One person, with explicit authority, named as the accountable owner of AI governance decisions inside the organization.

KPMG adds further evidence: organizations with strong board-level AI coverage (89% vs. 76%) and deeper board AI expertise (45% vs. 20%) report significantly better governance outcomes. Governance doesn't start with tools. It starts with people and accountability.

FINDING 4: The Agentic AI Governance Gap Is the Most Urgent Problem of 2026

23%	74%	35%
of organizations use agentic AI today *— Deloitte, 2026*	plan to deploy agentic AI within two years *— Deloitte, 2026*	admit they could not shut down a rogue AI agent if one emerged *— Writer, 2026*

> *"The average enterprise is running agentic AI. The average enterprise is not ready to govern it. When an AI agent makes a wrong decision autonomously, that decision has already executed. The damage is done before a human ever sees the log." — McKinsey, AI Trust Maturity Survey, 2026*

This represents a qualitative shift in what governance failure looks like. Traditional AI governance addressed outputs — a model that produced a biased recommendation, an inaccurate result, a problematic response. Agentic AI governance addresses actions. When an autonomous system books a vendor, sends a communication, or triggers a workflow, the error has already executed. The governance window that existed between decision and action no longer exists.

For mid-market organizations in healthcare and regulated industries, this is particularly acute. AI agents are being deployed in scheduling, triage, case management, and HR — all categories where autonomous action without proper accountability creates immediate legal and operational exposure.

FINDING 5: Governance Is a Revenue Variable, Not a Cost Center

74%	30%
of all AI-generated economic value is captured by just 20% of organizations — those investing most in governance *— PwC, 2026*	higher operating profit growth among firms investing 10%+ of AI budget in ethics and governance *— IBM, 2026*

The case for governance is not compliance. It is performance. Organizations that govern well do not just reduce risk — they capture more of the AI value they have already invested in. That is the case for readiness infrastructure. EY's Global Responsible AI Pulse Survey, Wave 2, October 2025 (975 C-suite leaders, 21 countries) finds that organizations with real-time monitoring and oversight committees are far more likely to report improvements in revenue growth, employee satisfaction, and cost savings — the exact areas where most organizations report struggling to see AI return on investment.

Nearly every company in EY's survey has already suffered financial losses from AI-related incidents. Average damages are conservatively estimated at $4.4 million per organization. For a mid-market company, a single AI-related compliance failure, legal action, or data breach in that range is not an abstraction. It is an existential event.

WHAT THE GAP LOOKS LIKE IN PRACTICE

A 180-person regional healthcare network deployed an AI scheduling and triage tool across three departments. The rollout was smooth. The vendor was credible. The technology performed well in the pilot.

Six months later, during a routine compliance review, the auditor asked: Who inside this organization is accountable for the decisions this system makes?

Three different people gave three different answers. None of them were wrong. The accountability had never been named.

The organization was not in regulatory violation — yet. But they had no incident response protocol, no monitoring framework, and no documented decision rights. They had a tool that worked and a governance structure that did not exist.

This scenario, drawn from composite experience across 20 years of operational leadership in regulated industries, is not unusual. It is the norm. And it is entirely preventable.

THE MID-MARKET TRANSLATION: WHAT THIS MEANS WITHOUT A GOVERNANCE TEAM

Every finding described above was surfaced in research conducted primarily among large enterprises — organizations with annual revenues exceeding $1 billion, dedicated AI governance functions, legal departments, and risk teams. For those organizations, the readiness gap is real but addressable with scale.

For the mid-market organization — 50 to 500 employees, a COO or CHRO who owns AI strategy alongside everything else, no dedicated governance infrastructure — the gap is the same, but the margin for error is smaller and the time to recover from a failure is longer.

> *"For mid-size companies, the Deloitte report is a warning signal: access to AI can be expanded quickly, AI maturity cannot. Companies that expand AI access without building governance in parallel will stack technical debt. That debt becomes visible when the first productive workloads move from pilot to scale." — Deloitte Mid-Market Analysis, 2026*

Here is the mid-market translation of each enterprise finding:

Enterprise Finding	Mid-Market Reality
21% have mature governance	If fewer than 1 in 5 large enterprises has mature governance, the figure for mid-market organizations — with fewer resources and less dedicated infrastructure — is almost certainly lower. The question is not whether your organization has a gap. It is how large the gap is.
Named accountability is the differentiator	This is actually good news for mid-market organizations. Named accountability doesn't require a team, a platform, or a budget. It requires a decision. Who, inside your organization, is explicitly responsible when an AI system makes a wrong call? If you can't answer that question in 10 seconds, you don't have governance infrastructure — you have good intentions.
Adoption outpacing readiness by 50%	Mid-market organizations are not immune to the pressure to move fast. Boards are asking about AI ROI. Leadership is experimenting with tools. HR is evaluating AI-powered applicant screening. Finance is trialing AI in reporting. All of this is happening — in most cases — without a corresponding investment in the governance infrastructure that determines whether any of it is defensible.
74% of AI value captured by governance-investing organizations	This is the business case. Organizations that govern well capture more value from the AI they deploy. The return on governance investment is measurable, documented, and consistent across McKinsey, PwC,

	EY, and IBM research. Governance is not the cost of doing AI responsibly. It is the precondition for doing AI profitably.

THE FOUR QUESTIONS: THE PREFLIGHT AI™ PRACTITIONER FRAMEWORK

AI governance doesn't begin with an audit. It begins when leaders can answer four questions about every AI system they operate or plan to deploy. These questions were developed from 20 years of operational experience inside regulated organizations — not from a framework document, but from the situations where governance actually broke down.

1. WHO OWNS IT?

Not which team, which vendor, or which department. Which person, by name, is accountable when this system makes a wrong call? In most mid-market organizations, the honest answer is: nobody. AI governance begins at the moment a named person accepts that accountability — with the authority to act on it.

McKinsey's 2026 AI Trust Maturity Survey confirms that named accountability is the highest-correlation factor with governance maturity. It costs nothing. It requires only a decision.

2. WHAT DATA DOES IT USE?

Where does this system's data come from? Is it current? Is it biased? Who has access to it? Who is responsible for its quality? For mid-market organizations using AI in hiring, performance management, or patient care, data provenance is not a technical question — it is a legal and operational one.

The EU AI Act, HIPAA, GDPR, and emerging US state AI laws all require organizations to demonstrate that the data feeding AI systems in high-risk categories is fit for purpose and governed appropriately.

3. HOW IS IT MONITORED?

What happens between deployments? Who is watching whether the system is performing as expected — not in testing, but in production? Most organizations have a deployment moment and then silence. Real governance is what happens in the silence.

EY's Responsible AI research finds that organizations with real-time monitoring and oversight committees report measurably better outcomes across revenue, satisfaction, and cost savings compared to those without.

4. WHAT HAPPENS WHEN SOMETHING GOES WRONG?

This is the question that reveals whether governance is real or performative. Does your organization have a documented incident response plan for AI failures? Is there a named person who activates it? Has it been tested? For most organizations, the honest answer reveals that governance exists only in the moments before anything goes wrong.

Deloitte finds that 35% of organizations admit they could not shut down a rogue AI agent if one emerged. For conventional AI, the equivalent question is: do you know how to turn it off, roll it back, and communicate what happened to the people affected?

Compliance asks: 'Do we have a policy?' Readiness asks: 'Can we operationalize it?' These four questions are how you find out.

THE READINESS GAP BY SECTOR

HEALTHCARE

- High-risk AI systems under EU AI Act include: AI-assisted clinical decision support, patient triage tools, diagnostic imaging analysis, and care coordination platforms.
- HIPAA requirements for data governance and access controls apply directly to any AI system processing protected health information.
- The most common governance gap: healthcare organizations deploy AI tools at the department level (radiology, scheduling, billing) without a named organizational accountability owner. Each deployment is locally approved but not centrally governed.
- The four-question test most often fails on Question 3 (monitoring) and Question 4 (incident response). Clinical teams know how to respond to a clinical failure; they rarely have a parallel plan for an AI system failure that affects clinical decisions.
- Regulatory pressure: CMS, ONC, and state regulators are increasingly requiring AI-related disclosures in federally funded programs.

GOVERNMENT & PUBLIC SECTOR

- NSPM-11 (National Security Presidential Memorandum) establishes AI governance requirements for federal agencies and contractors. State and local governments are increasingly subject to AI transparency requirements.
- The most common governance gap: AI tools deployed by one division or department without enterprise-level accountability. Workforce analytics, constituent service chatbots, and procurement AI are often running without formal AI risk classification.
- Procurement risk is the emerging pressure point: CAIDP's statement to the G7 Leaders' Summit explicitly called for Hiroshima AI Process principles to be embedded in public procurement requirements. When governance requirements enter procurement, they cascade to every vendor and contractor in the supply chain.
- Public sector organizations also face an accountability paradox: political leadership changes, but AI systems remain. Without documented governance infrastructure, the accountability chain breaks at every transition.

NONPROFIT & SOCIAL SECTOR

- Nonprofits are among the fastest-adopting organizations for AI — driven by resource constraints, the OpenAI People-First Fund (now open, $500K grants available at peoplefirst.fluxx.io/apply/rfp — as of June 2026, confirm current availability before sharing), and board expectations.
- The most common governance gap: mission alignment is assumed rather than tested. A nonprofit deploying AI in case management, client screening, or impact measurement rarely audits whether that AI's decisions are consistent with the organization's stated equity and justice values.
- Grant-funded AI initiatives create a specific governance risk: the AI is governed by the grant period, not by an organizational accountability structure. When the grant ends, governance often ends with it.
- The regulatory exposure that nonprofits underestimate: many receive federal funding that now carries AI-related compliance requirements. Grant agreements increasingly include AI governance provisions that organizations accept without having the infrastructure to fulfill them.

PROFESSIONAL SERVICES & CORPORATE

- Law firms, accounting firms, HR advisory firms, and management consultancies are deploying AI in client deliverables — often without disclosing this to clients or building the governance infrastructure to stand behind those AI-assisted outputs.
- The most common governance gap: vendor dependence treated as governance. 'Our AI vendor has SOC 2 certification' is not an answer to the question of who inside your organization owns accountability for the AI-assisted output delivered to a client.
- The EY Responsible AI Pulse survey finds that nearly all organizations surveyed have already suffered financial losses from AI-related incidents, with average damages conservatively exceeding $4.4 million. For a professional services firm, a single AI-related error in a client deliverable carries reputational and liability exposure that far exceeds the cost of governance infrastructure.
- The competitive advantage is real: KPMG finds that organizations confident in their AI talent pipeline and governance infrastructure are nearly four times as likely to report meaningful business outcomes. Governance is becoming a differentiator in client pitches.

WHAT TO DO IN 30 DAYS: YOUR GOVERNANCE QUICK START

Every report synthesized here recommends building governance infrastructure before AI deployment — not as a reactive compliance exercise, but as the operational foundation that makes AI investment defensible and sustainable. Here is a practical 30-day outline for mid-market leaders who are ready to close the readiness gap.

Week 1 — Name the Owners

1. **Name the accountability owner.** One person, by name, responsible for AI governance decisions across the organization. This is not a committee. It is a person with authority.
2. **Inventory what's already running.** Conduct an informal audit of every AI tool currently in use across the organization — including tools deployed at the department level without enterprise approval. You cannot govern what you haven't mapped.
3. **Complete the PreFlight AI™ Readiness Scorecard.** Six questions. Three minutes. Instant score. Identifies your highest-priority governance gaps before you do anything else. Free at bit.ly/3QoVPRo

Week 2 — Assess the Risk

4. **Apply the four questions to every AI system you inventoried.** For each: Who owns it? What data does it use? How is it monitored? What happens when something goes wrong? Document the answers — or document that you don't have them yet.
5. **Classify your risk exposure by regulatory framework.** Which of your AI systems touch functions covered by EU AI Act, NIST AI RMF, HIPAA, NSPM-11, or state AI laws? Prioritize governance infrastructure for high-risk categories first.

Week 3 — Build the Accountability Structure

6. **Document decision rights.** Who approves new AI deployments? Who can halt a deployment if a governance concern is raised? What is the escalation path?
7. **Establish monitoring checkpoints.** For each AI system currently in production, define what 'performing as expected' looks like and who reviews it, and on what schedule.
8. **Draft an incident response protocol.** One page is enough to start: what happens when an AI system fails, who is notified, who communicates externally, and who makes the decision to take it offline.

Week 4 — Close the Gaps

9. **Book a PreFlight AI™ Governance Readout & Roadmap Session.** A 60-minute working session that walks through your scorecard results, identifies your top 3-5 governance priorities ranked by risk and urgency, and delivers a written 30/60/90-day roadmap within 48 hours. $1,000 for organizations up to 100 employees; $2,000 for 100-200 employees. Fee credited toward a full Founding Pilot Diagnostic if you move forward within 30 days.
10. **Communicate the governance structure to your board.** Named accountability, risk classification, monitoring protocol, and incident response. KPMG finds that strong board-level AI coverage correlates directly with governance maturity outcomes.

ABOUT PREFLIGHT AI™

Kim Nelson-Wright *Co-Founder & CEO, PreFlight AI™* 20+ years of senior operational leadership across healthcare (Planned Parenthood, Concentra Medical), government (WMATA Senior Strategic Advisor), and nonprofit (Save the Children) sectors. Oxford Saïd Business School AI Governance Certificate. UPenn/Wharton AI For Business Specialization.	**PreFlight AI™** The independent AI readiness and governance diagnostic platform — the assessment layer organizations complete before deploying any AI vendor or tool. Platform-agnostic, deployment-independent, vendor conflict-free. *We are not the tool. We are the check you run before the tool goes live.*

Not sure where you stand? Start with the free 3-minute scorecard: https://kim-ftxkiz5k.scoreapp.com
AI Governance Readout & Roadmap Session: calendly.com/kim-getpreflightai/ai-governance-readout
AI Governance Readiness Playbook ($47): knelsonwright.gumroad.com/l/AIGovernanceReadinessPlaybook
getpreflightai.com | kim@getpreflightai.com

SOURCES

This report synthesizes publicly available research published between 2025 and 2026. All statistics are drawn from the primary sources listed below.

McKinsey & Company: State of AI Trust in 2026: Shifting to the Agentic Era (March 2026) — https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/state-of-ai-trust-in-2026-shifting-to-the-agentic-era

Deloitte: State of AI in the Enterprise 2026 (January 2026) — https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html

KPMG: Global AI Pulse Survey Q1 2026 (March 2026) — https://kpmg.com/xx/en/our-insights/ai-and-technology/ai-pulse.html

EY (Ernst & Young): Global Responsible AI Pulse Survey Wave 2 (October 2025) — https://www.ey.com/en_gl/insights/ai/how-can-responsible-ai-bridge-the-gap-between-investment-and-impact

PwC: AI Predictions 2026 — https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-predictions.html

BCG: As AI Investments Surge, CEOs Take the Lead (2026) — https://www.bcg.com/publications

Accenture: The Age of Co-Intelligence (2026) — https://www.accenture.com/us-en/insights/technology

Stanford HAI: AI Index Report 2026 — https://aiindex.stanford.edu/report/

Microsoft WorkLab: Agents and Human Agency (2026) — https://www.microsoft.com/en-us/worklab

CAIDP: Statement to the 2026 G7 Leaders Summit (June 2026) — https://www.caidp.org/statements/