CRIT CYBER

critcyber.com

The State of AI Security & Governance in Organizations

Findings from a primary research study of 23 cybersecurity, IT, GRC (Governance, Risk, Compliance), and AI professionals

23	4	8
PROFESSIONALS INTERVIEWED	MONTHS OF RESEARCH	KEY FINDINGS

Joshua Khokhar · Founder & CEO, CRIT Cyber

Study period: November 2025 – February 2026

Published July 2026 · Version 1.0

Contents

Executive Summary 3

About This Study 4

Research Questions 4

Methodology 4

Analysis 4

Participant Overview 5

Ethics, Consent, and Conflict of Interest 6

FINDING 1: Data Leakage and Shadow AI Were Near-Universal Concerns, Yet Rarely Subject to Technical Control 7

FINDING 2: A Gap Between Stated Policy and Technical Enforcement Was the Central Governance Failure Observed 9

FINDING 3: AI Adoption Was Predominantly Bottom-Up, While Effective Governance Required Executive Sponsorship 11

FINDING 4: Restriction Without Education Was Reported to Fail and to Introduce New Risks 13

FINDING 5: Governance Maturity Varied Widely Across Participating Organizations 14

FINDING 6: Autonomous AI Agents Emerged as an Unaddressed Governance Frontier 16

FINDING 7: Established Audit Cycles Were Perceived as Unable to Keep Pace with AI's Rate of Change 18

FINDING 8: Mid-Market Organizations Reported Being Underserved by Existing Tools 20

What Organizations Say They Need 22

Participant Voices: Extended Perspectives 24

On the Nature of AI Risk 24

On Shadow AI as a Structural Inevitability 24

On Model Security Beyond Data Privacy 24

On What Good Governance Looks Like 2

Limitations 2

Conclusion 2

Appendix A: Interview Guide 2

How to Cite This Report 2

Executive Summary

This report presents findings from a primary research study conducted between November 2025 and February 2026, during which 23 cybersecurity, IT, GRC (Governance, Risk, Compliance), and AI professionals were interviewed across industries including financial services, healthcare, government and defense, higher education, technology consulting, engineering, and AI software. Interviews were conducted by Joshua Khokhar, Founder and CEO of CRIT Cyber, as part of the market discovery and validation research underpinning CRITGuard, the AI security gateway from CRIT Cyber. It is a platform designed to help organizations monitor, govern, and control AI usage.

The study was designed to answer a core question: how are organizations actually navigating the security, compliance, and governance challenges that arise from adopting generative AI tools internally? What the study found was more urgent than anticipated.

22 of 23

participants cited data leakage and shadow AI as top concerns

0 of 23

organizations had technical enforcement covering the full scope of their stated AI policy

Across every organization in the study, from a Fortune 100 compliance team to a founder-led product design studio, the same structural gaps appeared repeatedly. The organizations studied have adopted AI tools faster than they have built governance for them. Policies exist on paper but are unenforceable without technical controls. Among the organizations studied, shadow AI is not an edge case; it is the norm: one mid-market financial institution identified 94 distinct AI products in active use, with roughly 600 employees using ChatGPT against only 200 officially approved users, out of 1,500 staff.

The study surfaced eight cross-cutting findings, each supported by data from multiple participants across different organizational contexts. These findings point toward a consistent conclusion: there is a substantial and growing unmet need for accessible AI governance capabilities, particularly among mid-market organizations that lack the resources to build enterprise-scale security infrastructure but bear the same data protection obligations as larger peers.

About This Study

Research Questions

This study was designed around four guiding research questions:

> **RQ1. How are organizations currently adopting and using generative AI and LLM tools internally, and how did that adoption begin?**

> **RQ2. What security, compliance, and governance concerns arise from internal AI usage, and how are organizations currently addressing them?**

> **RQ3. What gap exists between stated AI policy and technical enforcement, and what visibility do organizations actually have into usage?**

> **RQ4. What capabilities do security, IT, and compliance leaders say they need in order to govern AI usage effectively?**

Methodology

This study used a qualitative, interview-based methodology. Each participant took part in a 30–60 minute video or voice interview conducted remotely between November 2025 and February 2026. Interviews followed a semi-structured question guide covering role and AI usage, security and compliance concerns, governance approaches, tools and policies, what is and is not working, and desired capabilities. Three participants (Mic Merritt, the Fortune 100 GRC Director, and a security and risk analyst at a pediatric healthcare organization) responded via written channels or without audio recordings; all remaining interviews were transcribed and analyzed in full.

Participants were recruited through LinkedIn outreach and professional networks. The study was not commercially motivated during the interview phase (no product was being sold), and the purpose (academic-style market research for a forthcoming product) was disclosed upfront. All participants provided informed consent to participate; consent to be named by name and employer in this report was obtained per participant and is reflected in the participant table below.

Analysis

Analysis proceeded in three passes. First, every interview response was tabulated question-by-question in a structured matrix aligned to the interview guide (Appendix A), producing a directly comparable record across all 23 participants. Second, tabulated answers were cross-checked against the source transcripts to correct misattributed responses and fill gaps; eight corrections and four additions were made during this verification pass. Third, per-participant analytic memos were synthesized across cases to identify recurring themes.

A theme was retained as a finding only when it appeared independently across multiple participants in different organizational contexts. Counts (for example, 22 of 23, or 19 of 23) are reported wherever responses were directly comparable across the sample; where they were not, qualitative language is used

instead. Analysis was performed by the author as a single coder, supported by AI-assisted transcription (Tactiq) and AI-assisted transcript cross-referencing. Throughout the study, an independent cybersecurity professional served as an advisor: they observed a subset of the interviews as they were conducted, provided methodological input during the study, and reviewed the completed findings and this report. That review functioned as peer debriefing, a check on interpretation, rather than independent coding; no second coder separately coded the transcripts, and this is noted in the Limitations section.

Participant Overview

The 23 participants represent a cross-section of organizational types, seniority levels, and security maturity levels, a deliberate design choice to capture variance, not consensus.

Participant	Role	Organization	Sector
GRC Director (anon.)	GRC Director	Fortune 100 Company	Finance / Corporate
Security & Risk II Analyst (anon.)	Security & Risk II Analyst	Pediatric Healthcare Organization	Healthcare
Alec Harrison	Co-Founder & CEO	VisualBoston	Product Design & Development
Mark Stenquist	Co-Founder	VisualBoston	Product Design & Development
Jason McDonald	Assoc. Director of DevOps	Harvard Medical School	Higher Ed / Medical Research
Nick Pope	Chief AI Officer	AI Pope LLC / Train in Your Lane	AI Education / Consulting
Caleb Cole	AI Governance and Security Program Manager	Credit Union (unnamed)	Financial Services
Mic Merritt	AI Security Consultant & Researcher	Merritt Collective / Colorado Tech Univ.	Cybersecurity Consulting / Research
Josh Cole	Senior Data Analyst	Univ. System of NH	Higher Ed
Jonathan Hornbeck	Founder & Consultant	Efficient Futures LLC	Consulting
GRC Professional (anon.)	Senior Manager, GRC	Enterprise AI Software Co.	SaaS
Jonathan Bicknell	VP of Software	Concepts NREC	Turbomachinery Engineering & Software
IT Director (anon.)	IT Director	Engineering Software Firm	Engineering Software
Philip Magnuszewski	Chief Innovation Officer	Infused Innovations	Tech Consulting
Managing Partner (anon.)	Managing Partner, AI GRC	AI Governance & Audit Consultancy	Audit / Compliance

Participant	Role	Organization	Sector
Consultant (anon.)	Responsible AI & Sustainability Consultant	ESG Consultancy	ESG / Sustainability
Aleksandr Babich	Sr. Manager, Cloud Engineering	Fintech Company (withheld)	Financial Technology
Andy Stevens	IS Security Manager	Chenega / AFRL (DoD contractor)	Defense / R&D
Gavin Lawton	Director of Development	CrowdNow Solutions	IT Services / AI
Neeraja Karuppiah	Full-Stack Developer	Daylily AI	AI / E-Commerce SaaS
Fractional CAIO (anon.)	Fractional CAIO / Founder	Digital Legacy Startup	AI
Nicole Meng	Graduate Researcher	Tufts University	Higher Ed / AI Research
Maman Ibrahim	Principal Partner, Cyber & Digital Risk	EugeneZonda	Cybersecurity Consulting

Note: Seven participants are referenced in this report by role and industry only, either at their request or because explicit naming consent was not obtained. Aleksandr Babich and Caleb Cole are named personally, but their employers are not disclosed at their request. Each participant named in this report consented to attribution.

Ethics, Consent, and Conflict of Interest

Disclosure: the author is the founder and CEO of CRIT Cyber LLC, which is developing CRITGuard, a commercial product in the AI governance category this study examines. This is a material conflict of interest, and readers should weigh the findings with it in view. Two safeguards apply. First, the product context was disclosed to every participant at the start of each interview, and no sales activity took place during the interview phase; product demonstrations, where they occurred, followed the close of the research questions at the participant's option. Second, findings are grounded in participants' own words, quoted with consent, so that readers can evaluate the evidence directly rather than relying on the author's characterization of it.

All participants provided informed consent to take part and were told the interviews would inform both a published report and product development. Consent to be named (by name, or by name and employer) was sought separately per participant after the interviews concluded. Participants who did not respond to the attribution request are anonymized by default and described only by role and industry. Participants named in this report will receive the draft for review prior to publication.

FINDING 1

Data Leakage and Shadow AI Were Near-Universal Concerns, Yet Rarely Subject to Technical Control

Of the 23 professionals interviewed, all but one identified data leakage or shadow AI as a primary concern. This was the most consistent single finding in the study, appearing across every sector, every organizational size, and every seniority level. Yet despite the near-universal recognition of the risk, effective technical controls are the exception rather than the rule.

The Problem: Everyone Knows, Few Act

Across this sample, the gap between awareness and action is striking. Organizations understand that employees are uploading sensitive data to public LLMs. They know that shadow AI tools are proliferating beyond what IT can see. Yet the response in most cases is policy-based rather than technical: organizations publish an acceptable use policy, conduct a training session, and hope for compliance.

> *Users don't understand what to enter. People are using it and most of it's shadow.*
>
> **Caleb Cole, AI Governance and Security Program Manager, Credit Union (unnamed)**

PLAIN LANGUAGE Employees are using AI tools without understanding what data they're putting at risk. DLP tooling is not yet in place for AI use cases, and most of the usage identified is unauthorized.

> *What I found when I go into companies is that the company infrastructure has been uploading to ChatGPT. That's a really critical issue.*
>
> **Managing Partner, AI Governance & Audit Consultancy**

PLAIN LANGUAGE In audit after audit, I find the same thing: employees are feeding sensitive internal data to public LLMs without any awareness that this is a problem.

The scale of shadow AI can be enormous. Caleb Cole's organization identified 94 distinct AI products in active use across 1,500 employees, with roughly 600 employees actively using ChatGPT against only 200 users officially approved to do so. These figures reflect a fundamental failure of the assumption that organizations can control AI adoption through policy alone.

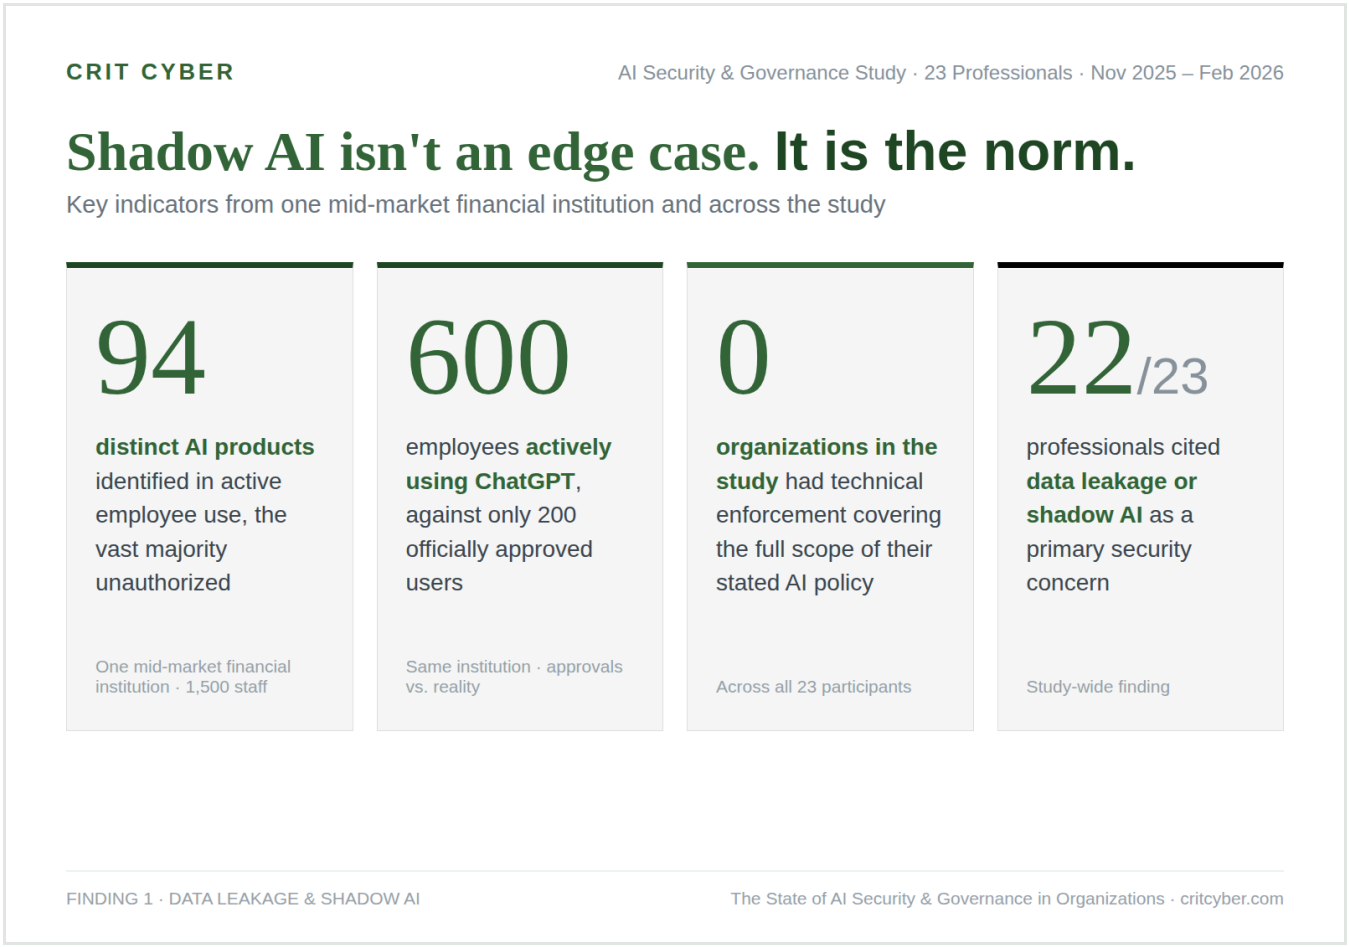


Figure 1. Shadow AI at one mid-market financial institution, and the study-wide control gap.

94

AI products identified in use at one mid-market financial institution

65%

of AI consulting firm clients now have an AI policy, up from near-zero six months prior (Gavin Lawton, CrowdNow Solutions)

The rapid policy adoption rate, from near-zero to 65% among one practitioner’s client base in six months, signals a market in rapid maturation. Organizations are beginning to acknowledge the problem. The next phase is technical enforcement.

FINDING 2

A Gap Between Stated Policy and Technical Enforcement Was the Central Governance Failure Observed

Across the study, the most common governance posture was: policy exists, enforcement does not. Organizations have written AI acceptable use policies, circulated them to employees, and in some cases required signatures, but have implemented no technical controls to verify or enforce compliance. The result is a governance structure that is visible from the outside and ineffective from the inside.

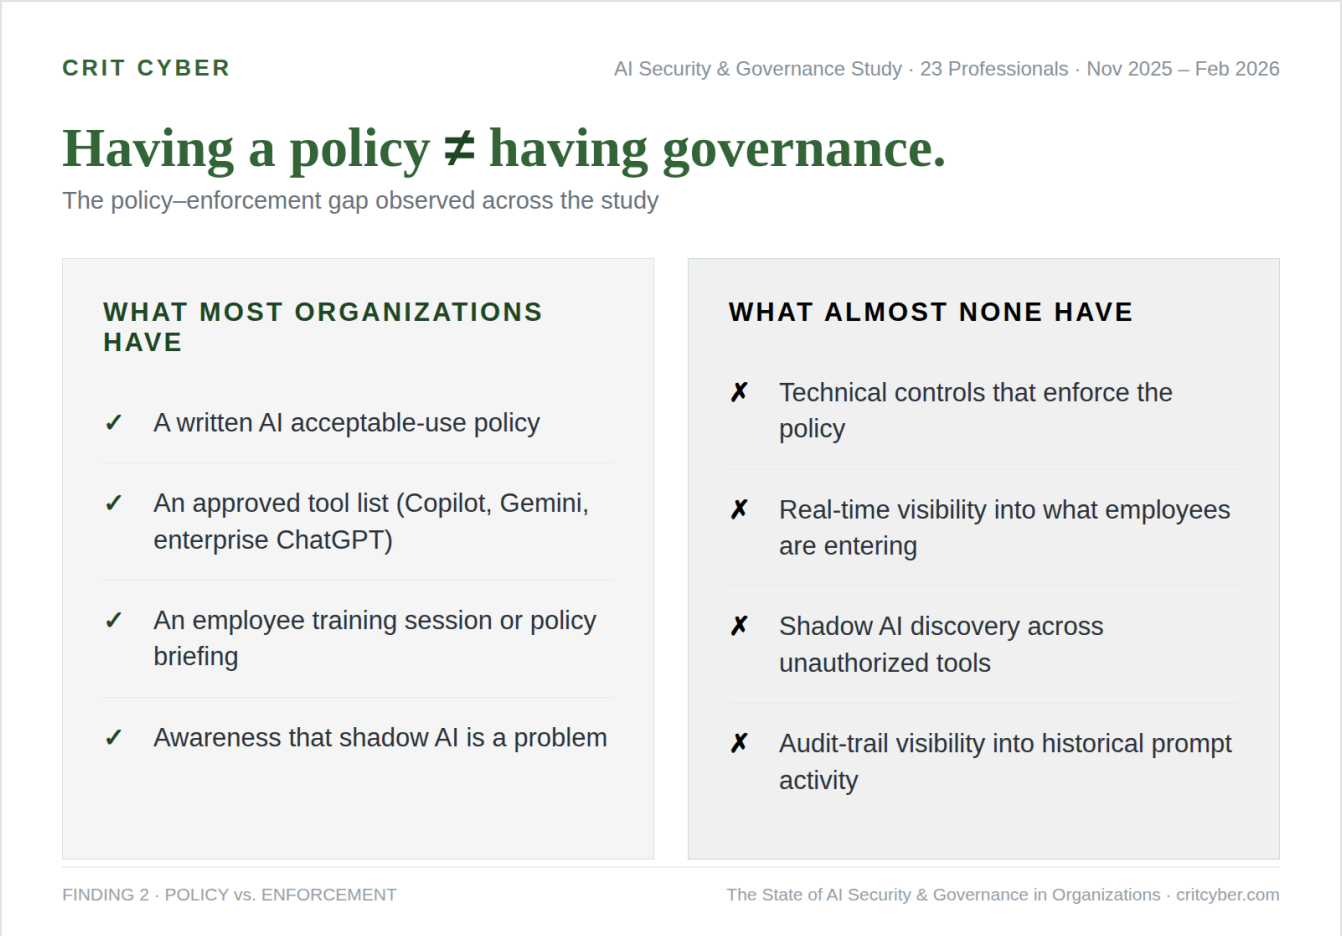


Figure 2. What organizations have versus what almost none have: the policy–enforcement gap.

> *AI policies are expected to be followed by the employees. Most of the time there are no technical restraints preventing them from accessing public AI tools and entering sensitive data.*
>
> **Jonathan Hornbeck, Founder & Consultant, Efficient Futures LLC**

PLAIN LANGUAGE We have a policy. Employees have seen it. But there is nothing stopping anyone from opening a browser, going to ChatGPT, and uploading whatever they want. The policy is a document, not a control.

> *There's this policy locked-down, not physically locked-down.*
>
> **Jason McDonald, Associate Director of DevOps, Harvard Medical School**

PLAIN LANGUAGE This is the policy-without-enforcement posture in miniature: the rules exist on paper, but no technical controls enforce them. Employees may know the rules; whether they follow them in practice is a separate question.

The Contradiction: Brand Trust Substitutes for Governance

A recurring pattern across the interviews was the phenomenon of strict general security postures coexisting with ungoverned AI tools, specifically Microsoft Copilot. Because Copilot arrives pre-integrated with Microsoft 365 environments, organizations treat it as implicitly safe. They do not apply the same vetting process they would apply to any other new software tool. The result is that the one AI tool employees use most is also the one most exempt from governance review.

> *Major name-brand AI tools get more of a pass because they’re a bigger company. There’s a higher level of trust versus, hey, this new cool model. Big companies aren’t going to do that.*
>
> **Nick Pope, Chief AI Officer, AI Pope LLC**

PLAIN LANGUAGE Brand familiarity is standing in for security diligence. Organizations assume that if it’s Microsoft or Google, it’s safe. That assumption is not a security posture.

FINDING 3

AI Adoption Was Predominantly Bottom-Up, While Effective Governance Required Executive Sponsorship

In 19 of 23 interviews, participants described AI adoption as having originated from employees before leadership. In many cases, AI arrived not through a deliberate organizational decision but through a product update: a Copilot notification in Microsoft 365, a new feature in an existing tool, an employee's personal ChatGPT subscription used on a work device. Leadership sanctioned what was already happening, rather than directing what would begin.

> *One day, you just receive a notification from Microsoft: Copilot is there, it can help you. No one woke up one day and asked, 'Oh, by the way, I heard there's something called AI.' It comes from your computer.*
>
> **Maman Ibrahim, Principal Partner, Cyber & Digital Risk, EugeneZonda**

PLAIN LANGUAGE AI adoption did not begin with a decision. It began with a notification. Employees were using AI before organizations had any governance framework in place, because the tools arrived pre-installed.

This dynamic creates a structural lag: the governance program is always chasing a deployment that already happened. The standard posture, prohibition first and then managed adoption, fails because prohibition arrives after employees are already using the tools. By the time a policy is published, the shadow AI library is already substantial.

The Governance Paradox: Bottom-Up Adoption Requires Top-Down Accountability

> *The number one key indicator if a company can successfully pivot into being an AI-first company is: does the CEO use AI daily?*
>
> **Nick Pope, Chief AI Officer, AI Pope LLC**

PLAIN LANGUAGE Executive behavior is the strongest predictor of organizational AI success. Policy comes from the top down. If leadership doesn't use AI themselves, they cannot set the cultural tone that makes adoption genuine rather than performative.

The organizations that were farthest ahead in governance were uniformly those where a senior executive had made AI governance a personal mandate. At one credit union, the SVP of Enterprise Security hired a dedicated manager to build the AI governance program. At an enterprise AI software company, the CISO was responsible for AI strategy. At AFRL, a former Google C-suite director drove the institutional push. At the same credit union, a governance committee was being chartered. In every case, bottom-up adoption became structured governance only when a senior leader made it a priority.

> *To implement AI the right way, you have to have that top cover. Without that, you're never going to make it.*
>
> **Andy Stevens, IS Security Manager, Chenega / AFRL**

PLAIN LANGUAGE Technical implementation alone is insufficient. AI governance requires executive mandate. Without it, every initiative is fighting against institutional inertia.

FINDING 4

Restriction Without Education Was Reported to Fail and to Introduce New Risks

One of the most unexpected and consistent findings of the study was the strong consensus against purely restrictive AI governance. Across multiple participants with very different organizational contexts, from a government R&D contractor to a founder-led product design studio, the view was the same: over-restriction does not prevent AI use; it drives it underground.

> *Overly restrictive policies do way more harm than good. It actually becomes much more likely that you're going to leak stuff if you are iron-fisted, because they're just gonna use their cell phones. Now you don't have control. Now you don't have insight.*
>
> **Nick Pope, Chief AI Officer, AI Pope LLC**

PLAIN LANGUAGE The irony of hyper-restrictive AI governance is that it produces the exact outcome it is trying to prevent. Employees don't stop using AI when you block it; they use it on personal devices, where you have zero visibility and zero control.

> *Let's not put a parking boot on it. It's secure, you're not gonna get a speeding ticket, but you're not gonna enjoy it either. You gotta be able to enjoy AI and the wonderful things it can do.*
>
> **Josh Cole, Senior Data Analyst, University System of NH**

PLAIN LANGUAGE Security that kills utility is not security; it is obstruction. The goal is AI that is safe to use, not AI that is impossible to use.

The Education Imperative

The preferred alternative to pure blocking is education at the point of action: tell employees not just that a prompt was blocked, but why, and what to do instead. This approach was independently endorsed by multiple participants, unprompted.

> *My KPI for this quarter is: I lowered the amount of blocks in ChatGPT. That shows us the tool is working and it's worth the value. Showing the behavioral change of the team is more significant than just showing the numbers of issues and blocks.*
>
> **Alec Harrison, Co-Founder & CEO, VisualBoston**

PLAIN LANGUAGE The right metric for AI governance tooling is not the number of blocks; it is the reduction in blocks over time. That trend line is the signal that employees are learning. That is the real measure of whether governance is working.

FINDING 5

Governance Maturity Varied Widely Across Participating Organizations

One of the most striking dimensions of the study was the range of AI governance maturity across participants, not just between large and small organizations, but within comparable organizational sizes. Two organizations with 200 employees might have completely different governance postures based on their industry, regulatory exposure, leadership orientation, and whether a dedicated security hire had been made.

The Maturity Spectrum

Level	Characteristics	Examples from Study
Wild West	No policy. No controls. Free use of any tool on any device.	Small businesses, agencies, and early-stage startups
Policy Without Enforcement	Written policy exists. No technical controls. Relies entirely on employee compliance.	Most SMBs
Managed Adoption	Approved tool list. Some web filtering. Enterprise licensing for approved LLMs. Shadow AI still present.	Credit union (unnamed, earlier), most mid-market
Structured Governance	Executive security mandate. Governance committee. Shadow AI detection. Active monitoring. DLP in progress.	Credit union (unnamed), Chenega/AFRL
Security-First Mature	Zero Trust applied to AI. Strict vendor vetting (7 POCs). All non-vetted LLMs blocked. Full policy documentation.	Fintech company (unnamed), Merritt Collective

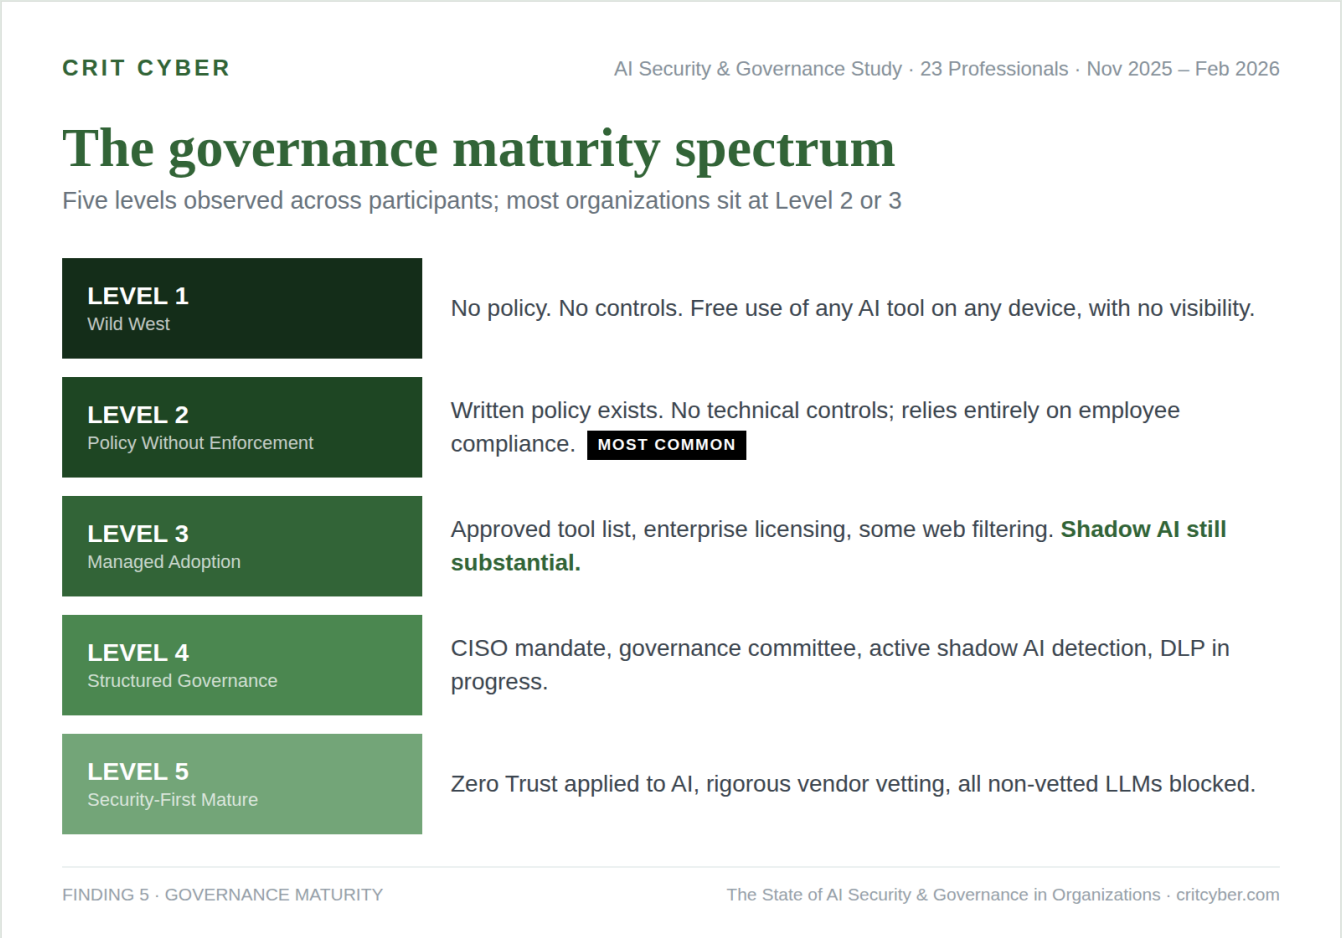


Figure 3. The five-level AI governance maturity spectrum observed across the study.

The implication of this maturity spectrum extends beyond any single vendor. The organizations most underserved are those at Level 2 or 3: they have acknowledged the problem and taken initial steps (a policy, perhaps an enterprise LLM license), but lack the technical enforcement to move forward. Level 1 organizations have not yet begun the journey; Level 4 and 5 organizations command resources most others lack.

A note on how this spectrum squares with the study's headline statistic: even the organizations at the top of it achieved control primarily through prohibition (rigorous vendor vetting and the blocking of unvetted tools) rather than through data-level enforcement inside the tools they sanction. No organization in the study, including the most mature, had technical controls covering the full scope of its stated AI policy. This is why the zero-of-23 finding holds even at Level 5.

> *We're moving from CMMI Level 1 to 2. I'd really like to get us to a solid 3.5. I think that's an 18-month journey.*
>
> **Caleb Cole, AI Governance and Security Program Manager, Credit Union (unnamed)**

PLAIN LANGUAGE Governance maturity is a multi-year journey. Most organizations are at the beginning. The path from 'we have a policy' to 'we have real controls' takes sustained investment and organizational will.

FINDING 6

Autonomous AI Agents Emerged as an Unaddressed Governance Frontier

While the majority of the study focused on LLM usage (employees typing prompts into ChatGPT or Copilot), a significant subset of participants raised autonomous AI agents as a distinct and emerging risk category. Unlike LLM chat interfaces, agents take actions: they access files, call APIs, send messages, make decisions, and operate without real-time human oversight. The governance frameworks that apply to LLM chat do not adequately address agentic behavior.

> *If I have to pull one risk, only one, it's the ability of AI agents to be autonomous. You cannot define everything. When there is something that doesn't fit into what you tell them to do, they hallucinate. In a business context, if they can't predict, they invent. And then they act upon what they invented.*
>
> **Maman Ibrahim, Principal Partner, Cyber & Digital Risk, EugeneZonda**

PLAIN LANGUAGE The defining risk of autonomous agents is that they take real-world action under uncertainty. Unlike a chatbot that gives a wrong answer, an agent that encounters something outside its instructions will improvise, and that improvisation might involve modifying a file, sending an email, or executing a transaction.

This is not a theoretical concern. Participants working with agentic systems had already seen governance failures in practice:

> *I've heard some crazy stuff, like unlocking full agentic capabilities through OpenAI agents that had access to a Google Drive that overrode things or deleted files. Going off the rails, not necessarily from a security standpoint, but from just bad governance and not understanding risk.*
>
> **Nick Pope, Chief AI Officer, AI Pope LLC**

PLAIN LANGUAGE The first wave of agentic AI incidents is already here. Over-permissive agents accessing file systems they shouldn't, modifying content, deleting data. These are governance failures, not security breaches, and the tools to prevent them don't yet exist.

The Orchestration Gap

Gavin Lawton, who builds AI agent platforms for client organizations, identified the absence of an orchestration layer as the primary gap in the current AI governance tooling landscape:

> *There's no tool for that right now to even just manage what's going on with agents. Co-pilot's starting to put some stuff out, but when you're actually building agents, there doesn't seem to be a lot of visibility and ability to backtrack.*
>
> **Gavin Lawton, Director of Development, CrowdNow Solutions**

PLAIN LANGUAGE Tooling for AI agent governance is still emerging; nothing yet is well-known or widely adopted. You can build agents, deploy them, and watch them operate, but it is hard to meaningfully see what they are doing, audit what they have done, or understand why they made the decisions they made.

> *Who is accountable for what it's doing? When you're using it for business, what is it exactly doing? Which process is it supporting? And if the context changes, it will not follow the change in the environment as fast as we think.*
>
> **Maman Ibrahim, Principal Partner, Cyber & Digital Risk, EugeneZonda**

PLAIN LANGUAGE Agent accountability is a governance gap without a solution. Organizations deploying agents cannot clearly answer basic questions: what is it doing, why, on whose behalf, and with what data? These are the foundational accountability requirements that traditional security frameworks demand, and that agent tooling doesn't yet provide.

FINDING 7

Established Audit Cycles Were Perceived as Unable to Keep Pace with AI's Rate of Change

A recurring frustration across security-mature participants was the mismatch between traditional audit and compliance cycles (typically six months to a year) and AI's rate of change. New models, capabilities, attack surfaces, and regulatory guidance are emerging on timelines that make annual or biannual assessments effectively obsolete before they are complete.

> *By the time you do audits and assessments, which might take six months, AI is way past where we were six months ago. We have to adapt to a faster method.*
>
> **Andy Stevens, IS Security Manager, Chenega / AFRL**

PLAIN LANGUAGE The six-month audit cycle is fundamentally incompatible with AI's pace of change. An assessment that takes six months to complete arrives describing a technology landscape that no longer exists. Governance frameworks built around annual cycles are not fit for purpose in an AI context.

The same dynamic applies to compliance frameworks. Several participants noted that existing frameworks (NIST, CMMC, SOC 2) were not built with generative AI in mind and require significant interpretation to apply to LLM-specific risk scenarios. Newer frameworks like ISO 42001 and NIST AI RMF are emerging, but are not yet widely adopted or well understood.

> *AI is advancing so quickly. The compliance processes are a nightmare, particularly because by the time you start trying to do audits and assessments, we're way past where we were. We have to adapt to a faster method.*
>
> **Andy Stevens, IS Security Manager, Chenega / AFRL**

PLAIN LANGUAGE The compliance machinery is built for a slower world. AI doesn't wait for the audit cycle to complete before evolving. Every governance program is running behind by definition, and the gap is getting wider.

Continuous Monitoring as the Alternative

Mic Merritt articulated the response that the most technically sophisticated participants converged on: AI governance cannot be a point-in-time exercise. It requires continuous monitoring, continuous evaluation, and a posture that assumes model behavior will change.

> *AI security is not a one-time setup. These systems change quickly, and their behaviors can shift as context and usage evolve. Ongoing testing and monitoring are necessary, even for mature deployments.*
>
> **Mic Merritt, AI Security Consultant & Researcher, Merritt Collective**

PLAIN LANGUAGE You cannot audit AI once and declare it compliant. The model you approved six months ago may behave differently today due to updates, new capabilities, or changed context. Continuous monitoring is not optional; it is the only posture that reflects the actual risk.

The consequence is structural rather than procedural. An annual or semi-annual assessment cadence produces documentation of an environment that no longer exists by the time it is signed. Participants operating under formal compliance regimes described governance work as permanently trailing deployment: approvals, risk assessments, and audits completing one or two model generations behind the tools employees actually use. The organizations in the study coping best had begun shifting from point-in-time certification toward continuous, monitoring-based assurance.

FINDING 8

Mid-Market Organizations Reported Being Underserved by Existing Tools

The enterprise AI security market has an upper-market bias. The most capable tools (Cybereason, CrowdStrike, Microsoft Purview, Palo Alto Networks AI-SPM) are priced and configured for large enterprises with dedicated security teams and significant IT budgets. Mid-market organizations, typically defined as 50 to 500 employees, face the same data protection obligations as their enterprise counterparts but cannot access or justify enterprise-scale tooling.

> *We're basically Cybereason for mid-market. The price tag on Cybereason is very enterprise-focused. A huge price tag. But the concept, an agent on the laptop finding data leakage into free AI tools, if you can get that at half the price, I think you've got something.*
>
> **Senior Manager, GRC, Enterprise AI Software Company**

PLAIN LANGUAGE The enterprise AI security concept is right. The pricing and complexity are wrong for mid-market buyers. A mid-market-accessible version of what the enterprise tools do is a genuine market gap.

Multiple participants made the same observation from different angles: the mid-market has AI usage, has data protection obligations, and lacks purpose-built tooling. The one-IT-person organization is a specific and common profile: a company with enough complexity to need governance tooling but not enough resources to deploy enterprise software.

Ease of Implementation as a Buying Criterion

Ease of implementation was cited as a top or near-top adoption criterion by nine of the 23 participants. This is not merely a preference; it is a necessity for organizations without dedicated implementation resources. A product that requires a six-week enterprise deployment is simply not accessible to the mid-market buyer.

> *Data handling and effectiveness, ease of implementation.*
>
> **Jonathan Hornbeck, Founder & Consultant, Efficient Futures LLC**

PLAIN LANGUAGE The two things that matter most for adopting an AI governance tool: does it actually work, and can we get it running without a major project?

The mid-market squeeze has three parts. Regulatory obligation is size-blind: a 200-person credit union or clinic carries data-protection duties comparable to an enterprise. Existing AI security tooling is priced and architected for enterprise deployment: dedicated security engineering, long implementations, six-figure contracts. And the staffing reality of the mid-market, often a single IT administrator wearing the security hat,

means any workable tool must be deployable and operable by one person. Taken together, these three constraints define an underserved segment whose exposure is growing at the same pace as everyone else’s.

What Organizations Say They Need

Across 23 interviews, participants were asked directly and indirectly about the capabilities that would make the most meaningful difference in their AI governance. The following capabilities were cited most frequently and most urgently:

1. Real-Time Visibility into AI Usage

The #1 named need across the study. Organizations want to know what tools employees are using, how often, and what data is going in, in real time, not through periodic audits.

> *Real-time visibility and observability.*
>
> **GRC Director, Fortune 100 Company**

PLAIN LANGUAGE The single capability that would make the biggest difference is knowing what is happening right now, not what happened last quarter.

2. Shadow AI Discovery

Nearly every participant wanted a clear view of what AI tools are in use in their environment, including tools they did not authorize and tools they did not know existed, including vendor-embedded AI (e.g., Atlassian's Rova in Jira) that bypasses traditional software procurement.

3. Data Leakage Prevention with Policy Enforcement

Participants want to prevent specific categories of data from entering LLMs (PII, client IP, financial data, PHI, CUI) with configurable policies that reflect their specific regulatory context.

4. Audit Trails and Incident Investigation

For compliance-driven organizations, the ability to answer 'what data was submitted to which model, by whom, and when' is a core audit requirement.

> *Being able to determine the contents of what is being distributed, having the ability to see at least a gist of what was submitted to ChatGPT, so we know what our potential exposure is.*
>
> **IT Director, Mid-Size Engineering Software Firm**

PLAIN LANGUAGE AI governance requires after-the-fact investigability: if something went wrong, we need to be able to determine exactly what happened, when, and who was involved.

5. Educational Feedback at the Point of Intervention

Multiple participants independently endorsed the concept of a blocking mechanism paired with an explanation, telling the user not just that their prompt was rejected, but why and what to do instead. This approach supports behavioral change over time.

6. Behavioral Change Tracking

The reduction in blocks over time, demonstrating that employees are learning, was identified as the most meaningful KPI for AI governance tooling. This is distinct from total block count, which can increase even as governance improves.

Participant Voices: Extended Perspectives

The eight findings above synthesize patterns across the sample. This section steps back and lets individual participants speak at greater length: four perspectives that did not reduce cleanly to a single finding but that shaped the study's conclusions. Each is presented with minimal editing and a plain-language gloss.

The following quotes represent some of the most analytically significant statements from the study, selected for their depth, specificity, or novelty. They are included in full because they contain nuance that summary cannot capture.

On the Nature of AI Risk

> *AI can do many things, including the wrong ones, perfectly. Everything that you ask AI to do, whether you do it explicitly or by accident because you don't understand the context, AI will do it.*
>
> **Maman Ibrahim, Principal Partner, Cyber & Digital Risk, EugeneZonda**

PLAIN LANGUAGE The core risk of AI is not failure. It is competent execution of the wrong instruction. AI succeeds at what it is asked to do; the problem is that what it is asked to do is often wrong, incomplete, or dangerous without the person asking it knowing.

> *We forget the ethics piece. We can control, we set all these technical guardrails, but we forget the ethics. Ethics, explainability, fairness, bias: that's one of the biggest risks of AI.*
>
> **Maman Ibrahim, Principal Partner, Cyber & Digital Risk, EugeneZonda**

PLAIN LANGUAGE Technical governance and ethical governance are not the same thing. Organizations that implement DLP and access controls have addressed compliance. They have not addressed the bias, fairness, and explainability dimensions that determine whether AI is safe and just in its outputs, not just its data handling.

On Shadow AI as a Structural Inevitability

> *There's nothing we can do about shadow AI. For years we tried to contain Shadow Tech in enterprise, and today there's no organization that succeeded. Employees are human. Before you grant them access to the mailbox, they already configured their phone to access it.*
>
> **Maman Ibrahim, Principal Partner, Cyber & Digital Risk, EugeneZonda**

PLAIN LANGUAGE Shadow AI is not a governance failure. It is a human constant. No organization has successfully prevented shadow technology adoption. Employees will always find workarounds. The correct response is not to try to prevent it; it is to build visibility into it and manage the risk.

On Model Security Beyond Data Privacy

> *Making generative models safe is a comprehensive task. It's not just one aspect, it's not just protecting the user's privacy. It's also the model's privacy. An attacker could come and break the model just as much as they can extract information from it.*
>
> **Nicole Meng, Graduate Researcher, Tufts University**

PLAIN LANGUAGE AI security has two directions: protecting data from the model, and protecting the model from attack. Adversarial manipulation, backdoor insertion, and model extraction are equally important as data leakage, and receive far less governance attention.

On What Good Governance Looks Like

> *An ideal solution would combine testing, monitoring, access control, and data governance in one place. It would give organizations a clear view of what the model is doing, what data it's touching, and how it behaves under stress. It would also evolve as models change.*
>
> **Mic Merritt, AI Security Consultant & Researcher, Merritt Collective**

PLAIN LANGUAGE The ideal AI governance platform is not a point solution. It integrates monitoring, access control, data governance, and adversarial testing in a single interface, and it updates continuously as models and threat landscapes evolve.

Limitations

This study's findings should be read with the following limitations in view.

Sampling. The 23 participants form a convenience sample recruited through LinkedIn outreach and the author's professional network. The sample skews toward New England, toward professionals already engaged with AI governance questions, and toward people willing to speak with a startup founder. It is not a random or representative sample, and the counts reported (22 of 23, 19 of 23, and so on) describe this sample only; they are not population estimates.

Sector coverage. Industry representation is uneven: consulting and technology are heavily represented; financial services, healthcare, government and defense, higher education, and manufacturing appear through one to three participants each; several major sectors (retail, legal) are absent entirely.

Data quality. Three of the 23 interviews were conducted through written channels or without audio recording, limiting them to summary-level analysis. All data is self-reported by participants and was not independently verified against organizational records; figures such as tool counts and user counts reflect what participants stated in the interview.

Analysis. Thematic coding was performed by a single analyst (the author). An independent cybersecurity professional advised the study, observed a subset of interviews, and reviewed the findings and report, providing a check on interpretation; however, no second coder independently coded the transcripts, so the coding itself was not formally validated. The author's commercial position (see Ethics, Consent, and Conflict of Interest) is a potential source of confirmation bias in theme selection, partially mitigated by the verbatim-quote evidence structure of this report and the independent review described above.

Timeliness. Interviews were conducted between November 2025 and February 2026 in a field that changes monthly. The findings are a snapshot of that window; specific tool names, adoption figures, and policy postures will age quickly, though the structural patterns identified are likely to persist longer.

Conclusion

The state of AI governance in organizations today is characterized by a widening gap between the speed of AI adoption and the maturity of the systems designed to manage it. Employees are using AI tools, many of which were never formally approved, to handle data that carries real legal, regulatory, and reputational risk. Policies exist but cannot be enforced. Visibility into what is happening is limited or nonexistent. And the pace of change in both AI capability and AI risk is outrunning the audit and compliance cycles organizations have built.

This is not primarily a failure of intent. The organizations in this study are not indifferent to risk. They are trying to build governance programs with limited resources, against a technology environment that is moving faster than their processes can adapt. What they are missing is tooling: purpose-built, accessible, and deployable by a single IT administrator without a six-week implementation project.

The eight findings of this study point toward a consistent and urgent need. Organizations at every maturity level in this study, from the founder-led product design studio with no AI policy to the financial institution with 94 AI products in active use (most never formally approved), need help making their governance intentions real. Closing that gap, with accessible tooling, education-first enforcement, and governance that moves at the speed of AI, is the defining challenge this research surfaces.

About CRIT Cyber

CRIT Cyber LLC is building CRITGuard, an AI security gateway that helps organizations monitor, govern, and control AI usage internally. CRITGuard provides shadow AI discovery, real-time data loss prevention, policy enforcement, and compliance reporting for mid-market organizations.

critcyber.com · Joshua Khokhar, Founder & CEO

© 2026 CRIT Cyber LLC. All rights reserved. This report may be shared for informational and educational purposes with attribution.

Appendix A: Interview Guide

Interviews were semi-structured: every participant received the core questions below, with probes used as needed and the order adapted to the natural flow of conversation. Interviews ran 30–60 minutes.

1. Tell me about your role and what you're responsible for day-to-day. How does AI/LLM usage fit into your responsibilities?
2. How is your company currently using GenAI/LLM tools?
3. Which specific LLM tools or platforms are you using?
3a. How long ago did you start using LLMs?
4. How did your organization's AI usage start? Was it top-down or bottom-up? (Probe: officially sanctioned? Shadow AI?)
5. What are your security concerns around using LLMs?
5a. If using LLMs already: what security or compliance concerns came up when your organization started using them?
5b. If planning to: what security or compliance concerns come up when your organization wants to use LLMs?
5c. Probe: data leakage/privacy? Trackability? Access control?
6. How do you approach security in the organization?
6a. How does AI/GenAI security fit into that picture?
7. Have you experienced any security incidents or close calls related to AI usage?
8. What techniques or tools are you using to mitigate these issues and manage AI security?
9. What policies or guidelines have you put in place around AI usage?
9a. How do you enforce these policies?
10. How do you currently track or monitor LLM usage in your organization?
10a. What visibility do you have? What are you blind to?
10b. How granular is your tracking: can you see who's using what, when, and for what purpose?
11. What's working well with your current approach?
11a. What's not working? Where are the gaps?
11b. Can you give me a specific example of when your current approach failed or fell short?
12. What specific capabilities would make your job significantly easier?
12a. Probe: real-time monitoring? Policy enforcement? Audit trails? Cost tracking?
13. What would the ideal solution look like for managing AI security and governance?
14. What would need to be true for you to adopt a new tool in this space?
14a. Have you evaluated or tried any tools for AI governance or security? What happened?

15. Is there anything we didn’t cover that you think is important for us to understand?

How to Cite This Report

Khokhar, J. (2026). The State of AI Security & Governance in Organizations: Findings from a Primary Research Study of 23 Cybersecurity, IT, GRC (Governance, Risk, Compliance), and AI Professionals (Version 1.0). CRIT Cyber LLC. Available at critcyber.com.

Questions about this research: joshua@critcyber.com