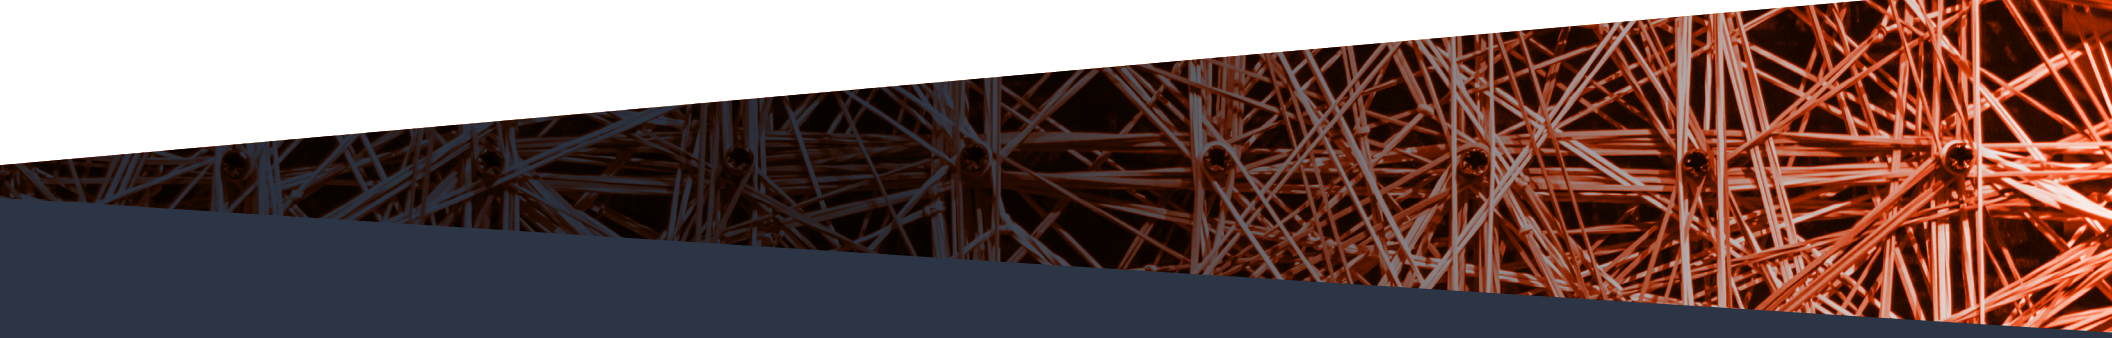

Prepared for

protegrity

The State of AI Friction:

Why Enterprise AI Deployment is Slower, Costlier, and More Limited Than Expected

July 2026 EMA White Paper

By **Christopher M. Steffen, CISSP, CISA, CCZT;** VP of Research
Information Security, Risk, and Compliance Management

EMA™
IT & DATA MANAGEMENT RESEARCH,
INDUSTRY ANALYSIS & CONSULTING

Table of Contents

Executive Summary: The Great Transformation

Enterprise AI adoption is not being held back by lack of investment, interest, or technical ambition. According to EMA's April 2026 primary research survey of 152 IT and security leaders, 82.9% of organizations report that security or compliance reviews have delayed AI projects from moving into production — and among those experiencing delays, most are stalled for a month or longer. More than 80% report deploying AI agents in a diminished state at least occasionally because of security concerns. The problem is not starting AI projects. The problem is the friction they face getting to full production value.

AI friction — the accumulated drag imposed by security reviews, compliance processes, sensitive data access challenges, autonomy constraints, and trust concerns — is measurably slowing deployment, reducing functionality, and increasing the operational cost of AI governance. Organizations have the ambition and the budget. What they lack is a clear path from prototype to production that doesn't stall in the reviews, approvals, and data access bottlenecks that define the gap between AI initiative and AI impact.

Taken together, these frictions add up to what amounts to an AI friction tax — a cost enterprises pay not in a single line item but across three: time lost to delay, capability lost to restricted deployment, and budget lost to compliance overhead.

The State of AI Friction

Enterprise AI has moved well beyond experimentation. EMA's survey data shows that 37.5% of organizations are actively scaling AI across multiple departments and production workflows, and another 20.4% have already moved into autonomous agentic deployments. Fewer than 2% report that AI initiatives are paused due to risk or compliance concerns.

How would you describe the current stage of your organization's AI or generative AI initiatives?

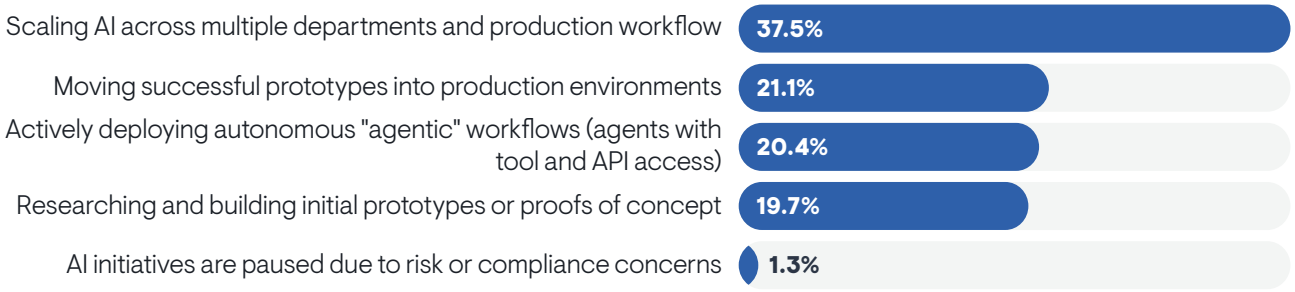

Sample Size = 152

Despite that momentum, production deployment is running into a new and largely underappreciated class of blockers. The issue is no longer whether organizations want to use AI. The issue is whether they can safely and efficiently move AI into real workflows without stalling in security, compliance, and data access reviews.

EMA's data makes the scale of this problem clear. When asked to rate the statement "Data security or compliance reviews have delayed AI projects from moving into production," 82.9% of respondents agreed or strongly agreed – with 32.2% strongly agreeing.

Please rate following statement: Data security or compliance reviews have delayed AI projects from moving into production.

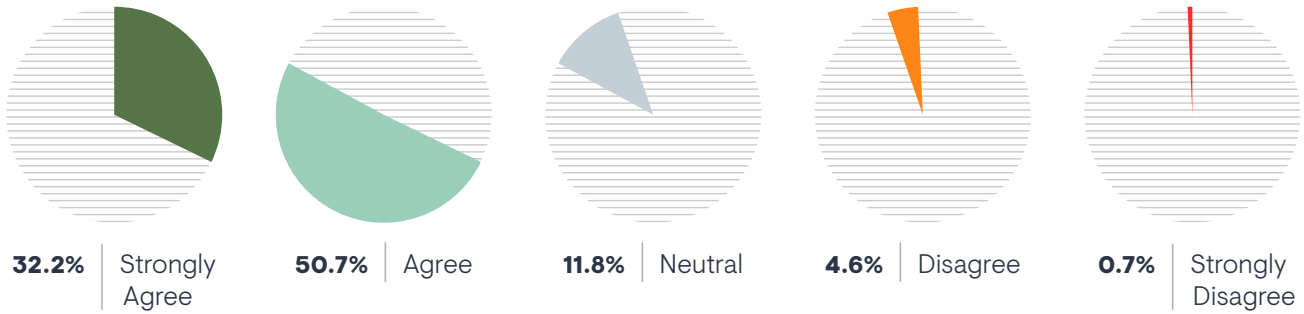

Sample Size = 152

Among organizations experiencing these delays, the duration is significant — and in a category where model capabilities shift each quarter, duration has compounding consequences. When asked how long their AI projects were delayed, 37.3% reported delays of one to three months, 22.2% reported four to six months, and an additional 7.2% reported delays of seven months or more. Combined, 66.7% of delayed organizations experienced production setbacks of one month or longer. Nearly 30% faced delays of four months or more.

In most software categories, a four-month delay is an annoyance. In AI, it can mean launching against a fundamentally different competitive baseline than the project was originally scoped against. A system delayed four months doesn't just arrive late — it may arrive obsolete, built around model assumptions, cost structures, or capability benchmarks that have already moved. The friction tax is not only a cost of lateness. It is a cost of relevance.

This is the deeper story in EMA's data. Enterprises are not failing to adopt AI. They are adopting compromised AI — launching delayed, diminished versions of the systems they originally built, while simultaneously increasing investment in the next round of projects facing the same friction. Organizations increasing agentic AI budgets by 10% or more account for 88.2% of respondents. The ambition is not the problem. The infrastructure to safely realize it is.

How long were your AI projects delayed from moving into production due to data security or compliance reviews?

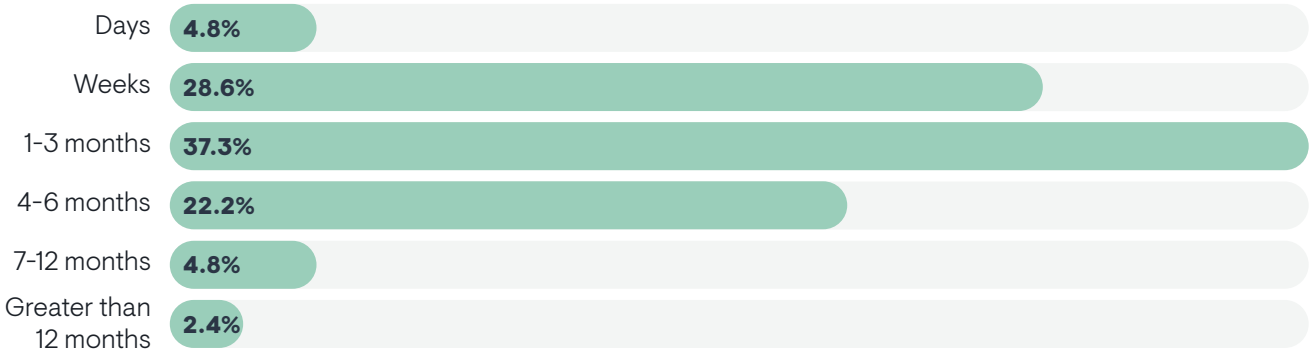

Sample Size = 126

The cost of AI friction: most delayed AI projects are not slowed by days. They are delayed by a month or more.

Security reviews emerged as the single largest bottleneck to AI production deployment, representing 20.0% of total bottleneck mentions – the highest share of any response category. Difficulty securing sensitive data for model training followed at 17.2% of mentions, with high infrastructure costs (16.9%) and inability to meet audit and compliance requirements (14.8%) rounding out the top four.

Which of the following represent the primary bottlenecks preventing your AI projects from reaching production?

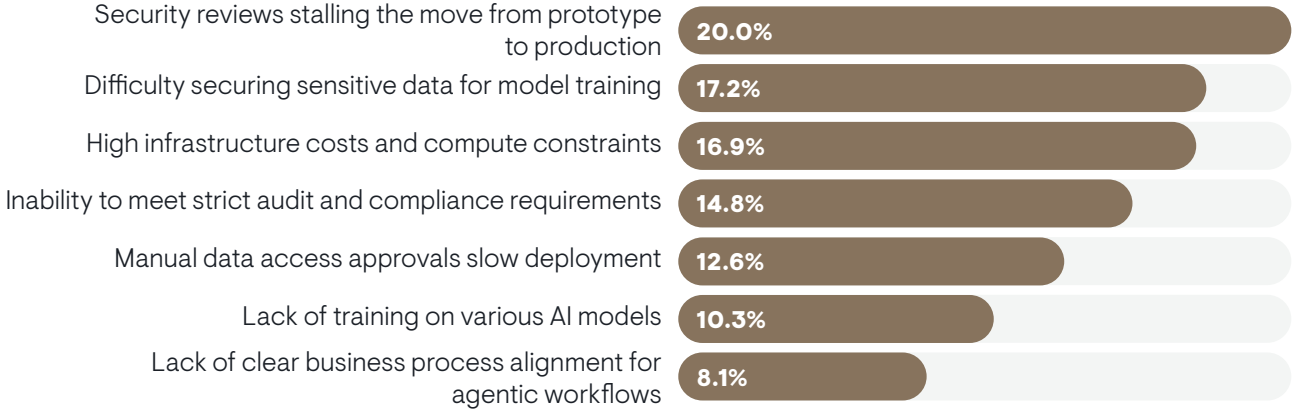

Sample Size = 152, Valid Cases = 152, Total Mentions = 419

Together, these findings frame AI friction as a measurable friction tax — assessed not at a single point but across the full deployment lifecycle. Organizations may have the AI ambition, the budget, and the technical momentum, but their ability to move AI into production is being constrained at precisely the moment it matters most: when a working prototype must clear the security reviews, compliance gates, and data access approvals required to become a production system. That bottleneck mix — security reviews at 20% of total mentions, sensitive data access at 17.2%, infrastructure costs at 16.9%, and compliance requirements at 14.8% — is the friction tax being assessed at the point of deployment.

Ranking the Top Five Sources of AI Friction

AI friction is not a single problem. It manifests across five distinct dimensions, each with its own mechanism and business cost. The following table ranks them by impact, drawing directly on EMA's primary research data.

The Five Structural Sources of AI Friction

Source of Friction	Why It Matters	Cost of Friction
#1: Security Review	Security reviews are the most direct blocker between prototype and production. AI systems cannot move forward until they clear a process that was designed for a different era of software deployment.	**55.3% of organizations cite security reviews as a top AI production bottleneck.** *Delayed launches, slower time-to-value, and missed productivity gains compound with every review cycle.*
#2: Compliance & Audit	Governance and audit processes designed for static systems are being applied to dynamic AI deployments, creating manual overhead that slows approvals and increases operational cost at scale.	**47.4% say reducing manual audit and compliance cost is their single most desired outcome from improved AI data security.** *This ranks above faster revenue, brand protection, and partner data sharing.*
#3: Sensitive Data Access	AI systems require high-value enterprise data to deliver meaningful outputs. When that data is difficult to secure and approve for use, organizations substitute lower-quality inputs — and get lower-quality AI.	**47.4% cite difficulty securing sensitive data for model training as a top production bottleneck.** *The result is AI that works in the demo and underperforms in production.*
#4: Autonomy	Organizations want autonomous agents. They cannot yet fully trust them. The gap between agentic AI ambition and the governance infrastructure required to authorize and monitor autonomous action keeps agents narrower than designed.	**Only 19.7% of organizations report highly autonomous AI agents.** *Nearly 80% operate at moderate, limited, or no autonomy — not by design, but by default.*
#5: Trust & Risk	Data leakage, prompt injection, privilege escalation, and agent impersonation are active deployment constraints that force security teams to restrict what agents can do before they go live.	**62.5% cite data leakage through autonomous agent actions as their top security concern. 42.8% frequently deploy agents in a diminished state as a result.**

Source: EMA Primary Research Survey, April 2026. n=152. Data references correspond to survey question identifiers.

Security review friction ranks first because it is the most direct and consistently cited blocker between prototype and production. Compliance and audit friction ranks second because it is both delaying deployment and creating measurable operational cost that enterprises are actively trying to reduce. Sensitive data access friction ranks third because it affects not just deployment timelines, but also the quality and utility of AI outputs. Autonomy and trust friction round out the list because they reveal how earlier-stage friction forces enterprises to deploy narrower, more constrained AI systems than they originally designed – a capability loss that compounds over time.

The Cost of Friction: Delayed and Diminished AI

AI friction does not only slow deployment. It changes what actually gets deployed.

One of the most differentiated findings in EMA's research is the prevalence of AI deployed in a diminished state — agents launched with restricted permissions, reduced autonomy, or manual workarounds substituting for the capabilities that security reviews prevented. When asked whether their organization had experienced deploying AI agents in a diminished state because of security concerns, 42.8% said this was a frequent occurrence for high-value projects. An additional 38.8% said it happened occasionally. Combined, more than 80% of organizations are regularly delivering less AI than they built.

This is an economic problem that currently has no name and no line item. Every AI business case is built around the agent at full capability — the system that can query production data, trigger external APIs, execute multi-step workflows, and act with meaningful autonomy. The ROI calculation assumes that agent. What many organizations are actually deploying is a constrained version: tools disabled because data access couldn't be cleared, autonomy limited because governance infrastructure wasn't ready, permissions restricted because security review ran out of time. Enterprises are paying full price for AI capability and operationalizing a fraction of it.

The autonomy data makes this concrete. Only 19.7% of organizations report AI agents operating at high autonomy — capable of planning and executing multi-step workflows independently. Nearly half (49.3%) report moderate autonomy, limited to preapproved scopes for specific actions. Another 27.0% report limited autonomy, requiring human-in-the-loop approval for every action. That distribution is not a technology limitation. It is a governance limitation wearing the appearance of one.

The security concerns driving this restraint are real and specific. Data leakage through autonomous agent actions tops the list at 62.5%, followed by regulatory compliance considerations (48.7%), prompt injection leading to unsafe execution (44.1%), tool manipulation and privilege escalation (39.5%), and agent impersonation or spoofing (32.2%). These are not edge-case risks. They are the reasons capable AI systems are being held back from their own potential.

Have you experienced instances in which AI agents were deployed in a "diminished" state (restricted utility) because of security concerns?

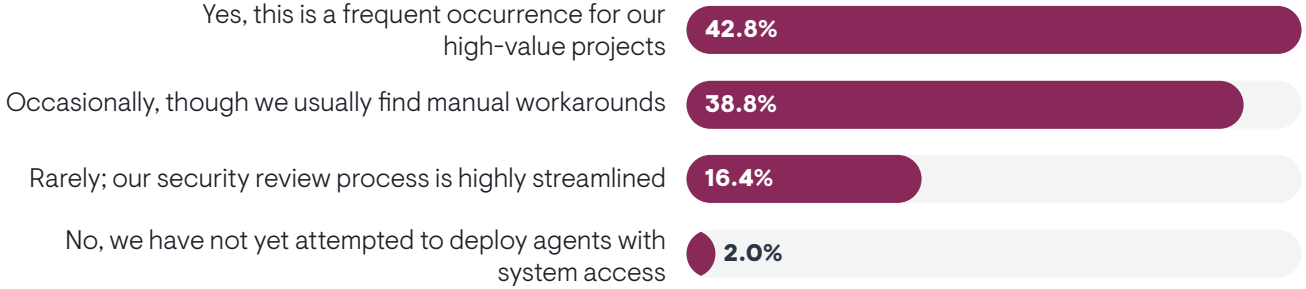

Sample Size = 152

AI friction creates a double cost: delayed deployment and diminished capability.

The autonomy data reinforces this picture — and reveals a protection gap that sits directly beneath it. Asked to characterize the current autonomous action capabilities of their AI agents, nearly half of respondents (49.3%) reported moderate autonomy, with agents operating within preapproved scopes for specific automated actions. Another 19.7% reported high autonomy, with agents capable of planning and executing multi-step workflows independently. Combined, 69% of organizations are running agents that write to databases, trigger external APIs, and take action in production environments.

The protection models securing those environments were not built for that. The two most common data protection methods currently in use are traditional encryption at the storage layer (57.2%) and dynamic data masking (56.6%) — perimeter-style controls designed around human queries and structured access patterns, not autonomous agents operating across systems at machine speed. The autonomy has outrun the protection model designed to contain it. Organizations are not failing to protect their data. They are protecting it with tools calibrated for a threat model that no longer reflects how their AI systems actually behave.

To what extent do your AI agents currently have "autonomous action" capabilities, such as writing to databases or triggering external APIs?

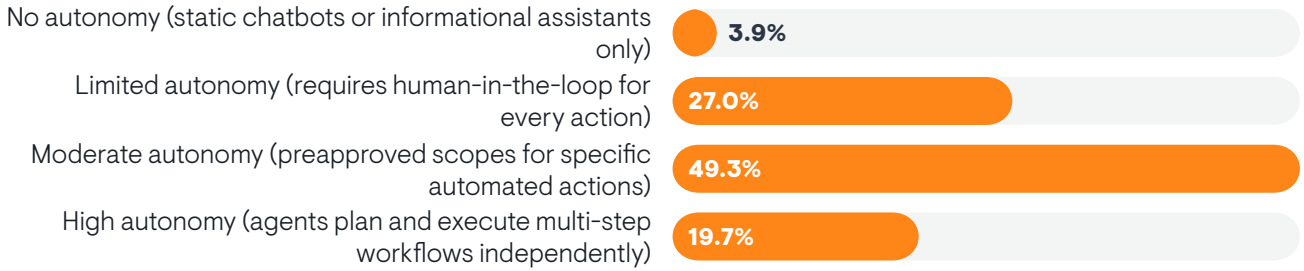

Sample Size = 152

The gap between AI ambition and operational autonomy is not primarily technical. EMA's data suggests it is governance. Organizations want autonomous AI agents. They are actively increasing agentic AI budgets – 68.4% plan budget increases of 10% to 25% in the next 12 months, and 19.7% plan increases exceeding 25%. The security and governance infrastructure required to authorize, monitor, and trust autonomous agent actions is not keeping pace with deployment ambition.

Over the next 12 months, how much is your team planning to increase its AI-related budget, specifically for agentic AI?

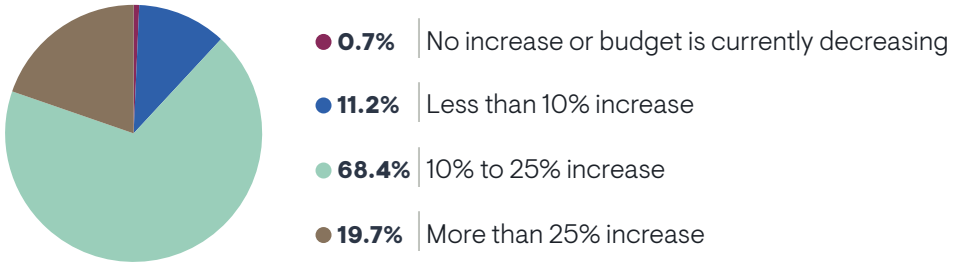

Sample Size = 152

The security concerns driving this restraint are concrete. When asked which challenges concern them most regarding autonomous agent deployment, 20.2% of respondents cited data leakage through autonomous agent actions, making it the top concern by a significant margin. Regulatory compliance considerations followed at 15.7%, with prompt injection leading to unsafe execution (14.2%), tool manipulation and privilege escalation (12.7%), and agent impersonation or spoofing and industry risk frameworks tied at 10.4% each. Supply chain risks from LLM plug-ins and MCP integrations were cited by 10.0% of respondents — a figure that will likely rise as agentic architectures grow more complex.

Which security challenges concern you most regarding the deployment of autonomous AI agents?

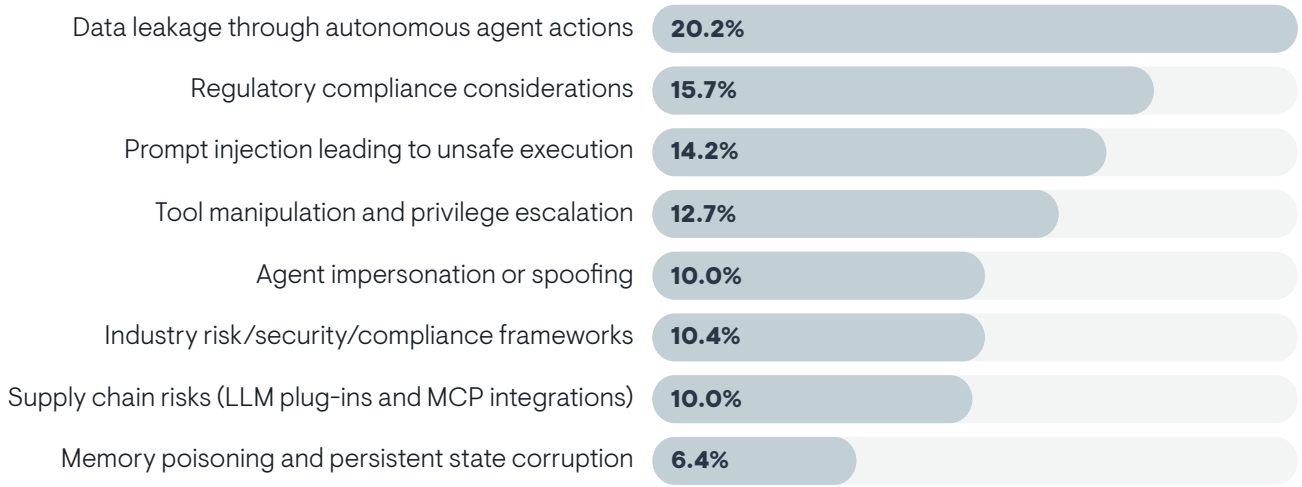

Sample Size = 152, Valid Cases = 152, Total Mentions = 471

Compliance Burden is Becoming a Business Constraint

EMA's survey data signals something important about how enterprises are experiencing AI governance: it has become an operational cost problem, not merely a security risk problem. Organizations are not simply worried about breach prevention. They are increasingly concerned about the manual burden created by audit, compliance, and approval processes that slow AI deployment and increase the cost of operating at scale.

When asked to identify the primary business outcome they seek from improved AI data security and governance, 47.4% selected "significant reduction in manual audit and compliance operational costs" – the top response by a significant margin. This outcome ranked above faster time to revenue for new AI products (24.3%), prevention of brand damage from data leaks or biased AI outputs (17.8%), and empowerment of safe data sharing with external partners and vendors (10.5%).

What is the primary business outcome you seek from improved AI data security and governance?

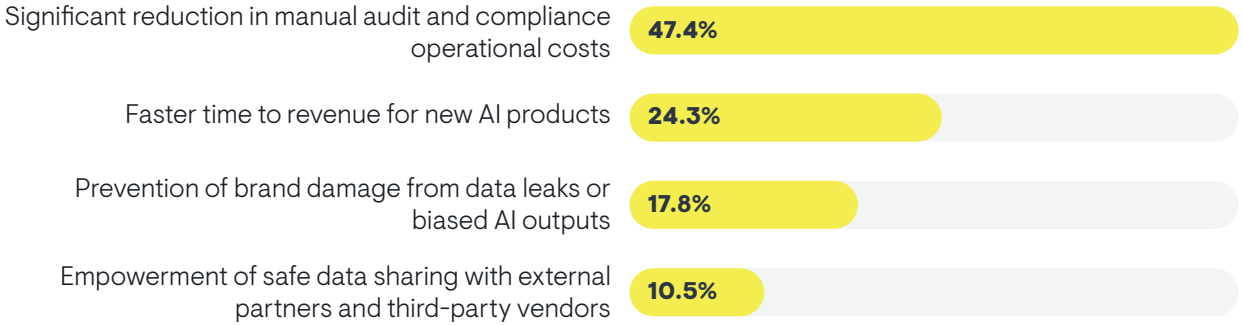

Sample Size = 152

Compliance cost reduction ranked as the primary desired outcome of AI data security improvements – nearly twice the rate of faster revenue as a motivator.

The appetite for a better approach is correspondingly high. When asked how valuable they would find a solution that provides "AI-ready data without roadblocks" by embedding protection inside the workflow itself, 40.8% of respondents selected "extremely valuable" and 50.7% selected "very valuable." Combined, 91.5% rate this capability as extremely or very valuable – a near-consensus signal that embedded data protection has become a priority, not a preference.

How valuable would you find a solution that provides "AI-ready data without roadblocks" by embedding protection inside the workflow itself?

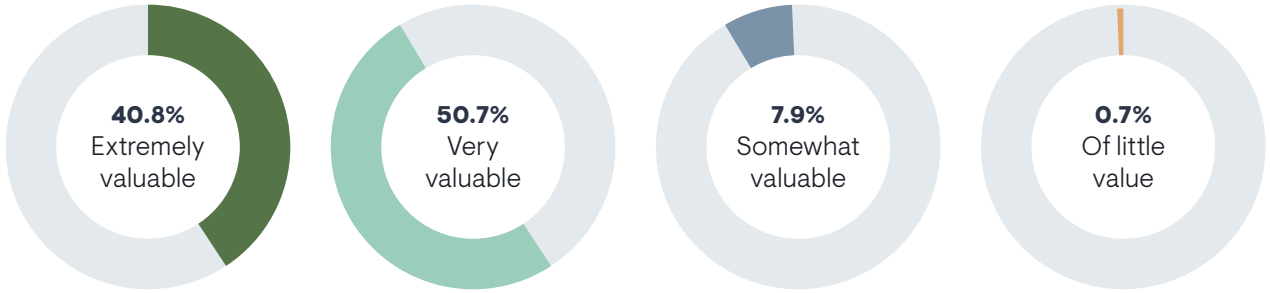

Sample Size = 152

The governance frameworks in scope reflect the breadth of the compliance challenge. When asked which regulatory frameworks currently impact their AI and data governance strategy, respondents cited an average of three frameworks each, signaling that most enterprises are navigating overlapping and sometimes conflicting compliance obligations simultaneously. GDPR is the dominant framework at 25.4% of respondents, followed by California Privacy Laws (CCPA/CPRA) at 14.8%, PCI-DSS at 14.6%, HIPAA at 13.5%, and DORA at 10.8%. Notably, the EU AI Act — despite significant industry attention — was cited by only 8.0% of respondents, suggesting that enterprises are currently anchoring AI governance to established privacy-era frameworks rather than waiting for purpose-built AI regulation to mature. The compliance burden is not arriving with the EU AI Act. It is already here, running on infrastructure designed for a different problem. (respondents could select multiple responses).

Which of the following regulatory frameworks currently impact(s) your AI and data governance strategy?

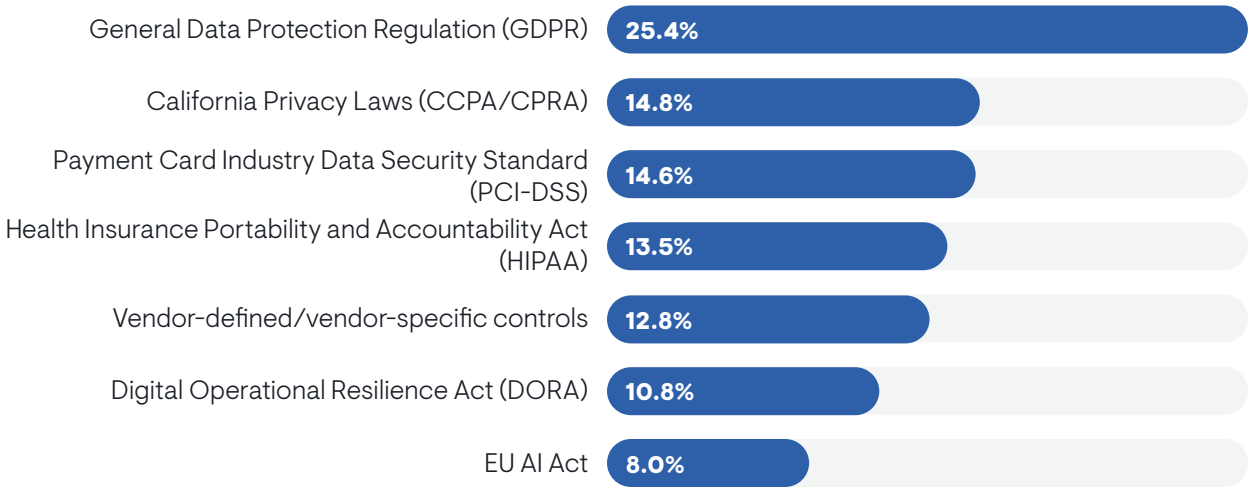

Sample Size = 152, Valid Cases = 152, Total Mentions = 452

The approval structure surrounding AI deployment helps explain why security review friction ranks as the single largest production bottleneck. When asked who in their organization approves AI deployments into production, 47.4% of respondents identified the CISO — the most common answer by a wide margin. Adding CIO and CTO approvers brings the combined share to 82.9%. The same function that owns breach liability is also the primary gate through which AI must pass before it reaches production. That is not a dysfunction. It is a rational structural response to real risk. But it is also the mechanism that generates much of the delay data documented elsewhere in this paper. When caution is the job, caution is what the process produces — and in environments where AI projects already face a 66.7% probability of month-plus delays, the cost of that caution is no longer abstract.

Who in your organization approves AI deployments into your production environment?

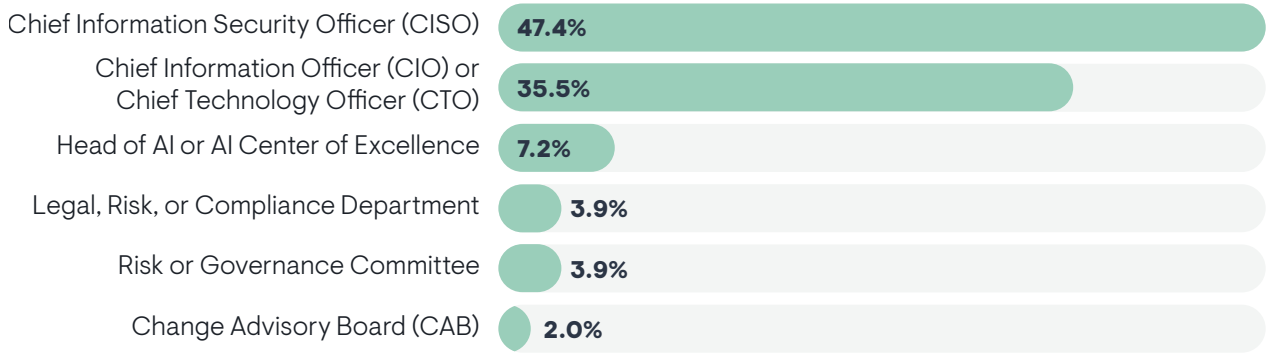

Sample Size = 152

AI Friction Has Become an Operational Pain Point

The data EMA collected in April 2026 does not describe a future governance challenge. It describes a present operational constraint with real business costs already affecting AI deployment at scale.

Security reviews are delaying production. Compliance processes are increasing operational cost. Sensitive data access challenges is weakening AI outputs. Trust concerns are forcing teams to deploy diminished agents rather than the systems they originally designed. The organizations experiencing these frictions are not edge cases; they represent mainstream enterprise adoption, with AI initiatives actively scaling across departments and significant budget increases already committed.

The path forward requires a fundamental shift in how data security and governance are positioned in the AI delivery lifecycle. Current models — in which security and compliance reviews function as gates that AI systems must pass through before production deployment — introduce the friction the survey data documents. The enterprises reporting the most friction are those in which security controls remain external to the workflow, requiring human review, manual data access approvals, and sequential audit steps that slow deployment and increase cost.

EMA's survey data points toward where the market is heading, and the signal gets clearer when the response options are grouped by what respondents actually want rather than read as competing choices. Grouped one way, the results look like a clear preference for central control: exactly half of respondents (50.0%) chose a security-team-owned model, either a centrally managed security platform (28.9%) or centralized appliance clusters run by a dedicated team (21.1%). Grouped the other way, nearly four in ten (38.8%) want protection embedded at the developer level, delivered as an SDK or container (23.0%) or as decentralized modules running directly in developer workflows (15.8%). The remaining 11.2% would outsource the problem entirely through public SaaS or a managed provider.

What is your preferred model for deploying data security controls into your AI pipeline?

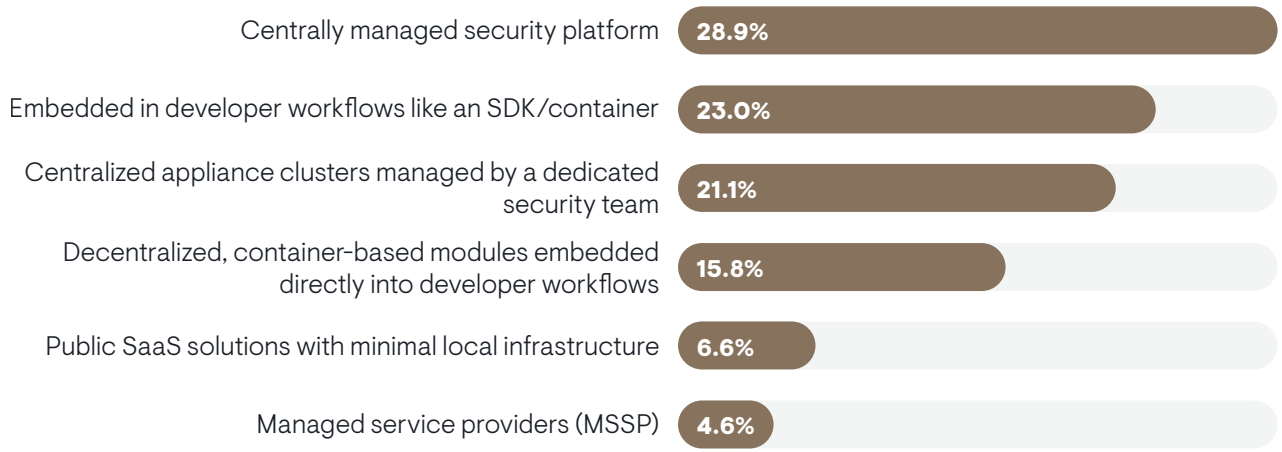

Sample Size = 152

The developer-embedded number is the one to watch. Data protection has traditionally lived with centralized security teams, even as other parts of security moved closer to developers, so a 38.8% preference for pushing data-security controls down into the workflow itself marks a real shift. Its size reflects how AI has changed the problem: when autonomous agents and pipelines touch data continuously, protection that waits at a central checkpoint is protection the workflow routes around.

Read together, these two camps are not in conflict — they are answering two different questions. The centralized half is expressing a preference about policy: who sets and governs the rules. The embedded group is expressing a preference about enforcement: where protection actually runs.

Those are not mutually exclusive, and the model the data points toward satisfies both — central policy management with enforcement distributed to where the data lives, rather than a single perimeter gate. That convergence, more than any one bar in the chart, is where the category is heading. Read together, these preferences resolve to a single architectural requirement: central policy management with distributed enforcement. Organizations want one place to set and audit policy, and they want that policy enforced at the point of data use — inside the pipeline, inside the workflow, inside the agent — rather than applied as a gate before deployment begins.

When protection is embedded rather than external, the gate disappears and the friction tax with it. Current data protection strategies in use reflect a market still anchored to perimeter-era controls. When asked which data protection and management methods are currently implemented in their high-throughput environments and respondents most commonly cited traditional encryption at the storage layer (18.7%) and dynamic data masking (18.5%). Role-based access control without field-level protection was selected by 14.0%, and anonymization models by 13.8%. Vaultless tokenization (8.8%) and tabular synthetic data generation (8.8%) trail the leaders but represent the direction the category is moving — approaches that protect sensitive data while preserving its utility for AI training and inference, rather than simply restricting access to it. (respondents could select multiple responses).

The dominant methods secure data at rest. The emerging methods secure data in use. That distinction is where the protection gap lives.

Which data protection and management methods are currently implemented in your high-throughput environments?

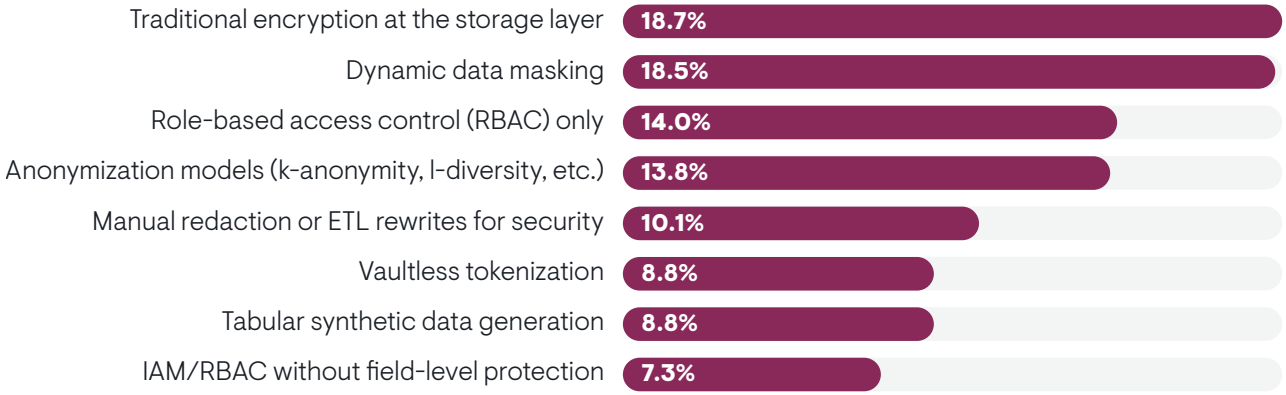

Sample Size = 152, Valid Cases = 152, Total Mentions = 465

For enterprise leaders, AI friction is no longer a nice-to-have governance improvement. It is a business-critical constraint on AI value — one with a measurable cost already visible in delayed timelines, restricted agent capability, and compliance overhead that consumes resources without accelerating deployment. The friction tax is real, it is quantifiable, and it is being paid right now by organizations that have already committed the budget and built the systems.

The protection gap at the center of this problem is structural. Sixty-nine percent of organizations are running agents that write to databases and trigger external APIs. The dominant controls securing those environments — storage-layer encryption, dynamic data masking, role-based access control — were designed for an era when humans, not autonomous systems, queried the data. The autonomy has outrun the protection model. And until the protection model catches up, the diminished AI finding will persist: capable systems deployed at a fraction of their designed capacity, business cases built on full-capability ROI delivering partial-capability results.

The next phase of enterprise AI will be defined not by which organizations invest most, but by which organizations can close that gap — moving AI safely and efficiently from prototype to full production without stalling in the reviews, approvals, and data access bottlenecks that currently define the distance between AI ambition and AI value.

Conclusion

Enterprise AI is no longer constrained by appetite, budget, or technical readiness. It is constrained by friction between AI systems and the governance infrastructure meant to secure them. Nearly 83% of organizations report production delays tied to security or compliance review, and 66.7% of those delays stretch a month or longer. Even when AI does reach production, 81.6% of organizations report deploying it in some diminished form — narrower in scope, more dependent on human oversight, or stripped of the autonomous capability it was originally designed to deliver. This is not a temporary growing pain. It is the operating reality of AI at enterprise scale in 2026, and it will persist until security and governance move from gating mechanisms into embedded, workflow-native controls.

For enterprise leaders, four actions follow directly from this research:

1. Audit where security review sits in the AI delivery pipeline. If protection is applied as a late-stage gate rather than a built-in property of the data and workflow, friction is structural, not incidental.
2. Treat compliance operating cost as a board-level metric. With 47.4% of respondents citing reduced audit and compliance burden as their top desired outcome, governance efficiency is now a measurable line item, not a soft risk-management goal.
3. Match autonomy ambition to governance capacity before scaling agent budgets further. Budget increases are outpacing the trust infrastructure needed to support full agent autonomy; closing that gap should precede, not follow, additional investment.
4. Evaluate data protection methods that travel with the data — tokenization, dynamic masking, synthetic data — rather than relying solely on perimeter encryption and access controls, which the data shows are no longer sufficient for environments where autonomous agents, not humans, are querying production system.

Most enterprises now run multiple AI projects in parallel. A median delay of one to three months, multiplied across an AI portfolio, compounded by the capability lost to diminished deployment, and recurring compliance overhead that consumes resources without accelerating delivery, makes friction the largest hidden cost on AI investment today — larger than compute, larger than talent, larger than model quality. Those costs have no single line item, which is precisely why they have gone unmeasured. Naming them, quantifying them, and tracing them to their structural sources is the contribution this data makes. Organizations that act on it will deploy AI faster and at full capability. Those that do not will keep paying the friction tax — on every project, in every quarter, indefinitely.

About Enterprise Management Associates, Inc.

Founded in 1996, Enterprise Management Associates (EMA) is a leading IT research and consulting firm dedicated to delivering actionable insights across the evolving technology landscape. Through independent research, market analysis, and vendor evaluations, we empower organizations to make well-informed technology decisions. Our team of analysts combines practical experience with a deep understanding of industry best practices and emerging vendor solutions to help clients achieve their strategic objectives. Learn more about EMA research, analysis, and consulting services at **www.enterprisemanagement.com** or follow EMA on **X** or **LinkedIn**.

2770 Arapahoe Road, Ste 132 - 544, Lafayette, CO 80026 +1 303.543.9500 **www.enterprisemanagement.com**

4752.072326