CSAI Foundation | Cloud Security Alliance

AI Governance Failures Trace to Process Design, Not Policy

2026-09-22

AI-assisted Rapid Research

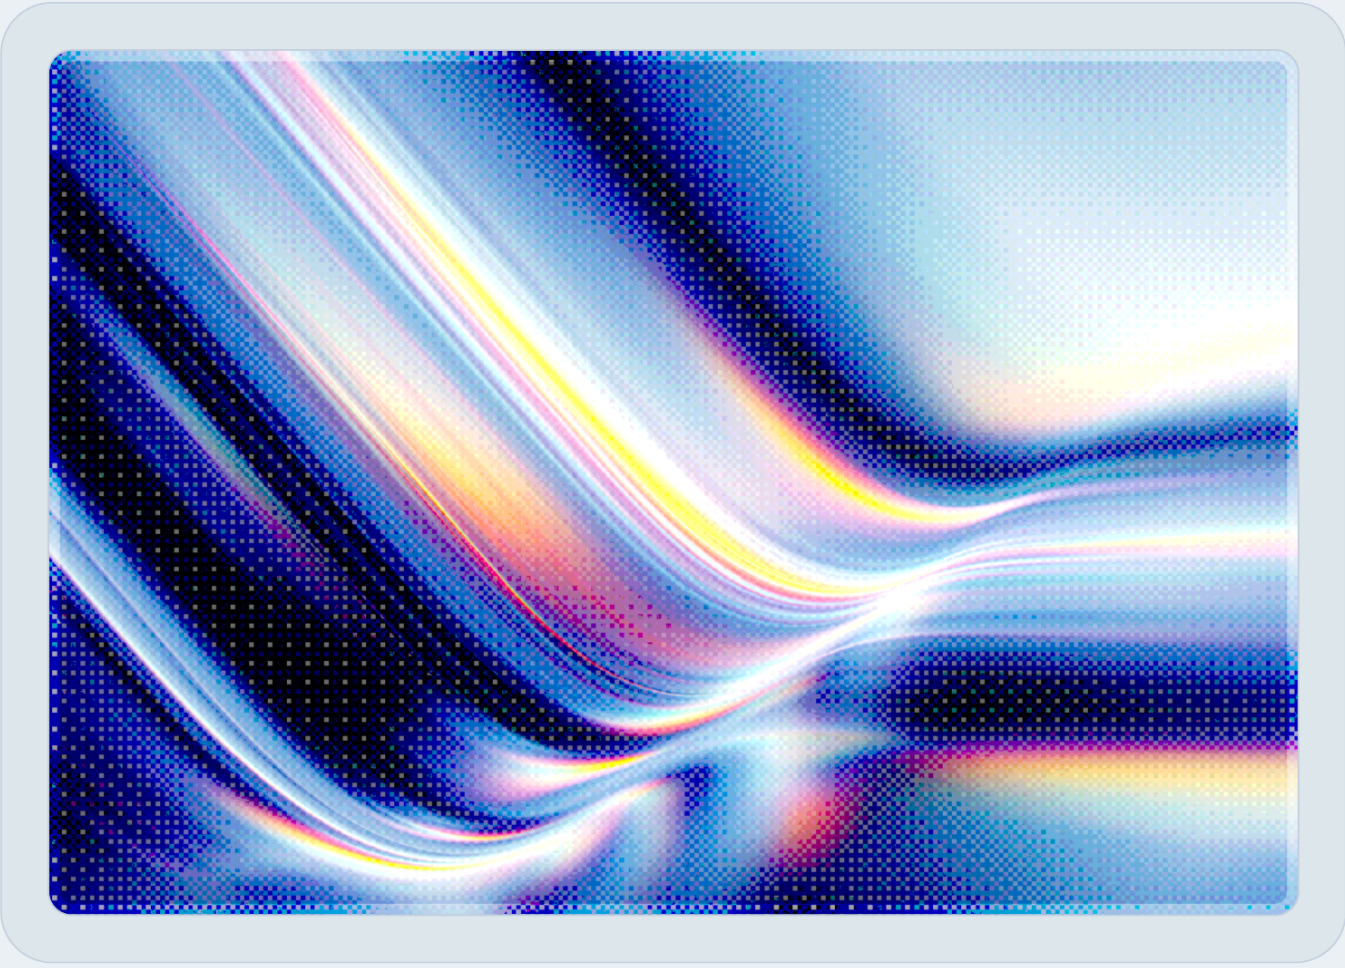

CSAI

CSA cloud security alliance®

Key Takeaways

New survey research from Camunda, conducted by Sapio Research across 1,000 process decision-makers and 5,000 employees at large enterprises in the United States, United Kingdom, Germany, and France, finds that 40% of organizations experienced an AI-related compliance or governance issue in the past 12 months, and that 84% of those incidents trace back to process-related problems rather than missing or inadequate policy [1][2]. The finding reframes a debate that has largely centered on writing better AI policies and points instead at the workflows those policies are supposed to govern. Organizations are bolting AI onto approval chains, handoffs, and exception paths that were designed for a human to execute every step, and when AI takes over part of that step the compliance checks, the audit trail, and the accountability the process was built to provide all land in the wrong place. Seventy-two percent of respondents say process-related challenges have already caused an AI initiative to fail, at an average cost of $1.55 million per organization, and 61% say process redesign cannot keep pace with how quickly they are being asked to deploy AI [1]. For security and compliance leaders, the practical implication is that governance work aimed only at policy documents and model-level controls will miss the failure mode that is actually driving incidents: control points, evidentiary trails, and escalation paths that no longer match how decisions get made once AI is in the loop.

Background

The research, commissioned by workflow orchestration vendor Camunda and fielded by Sapio Research in July and August 2026, surveyed 1,000 senior IT, operations, and transformation leaders at organizations with at least 1,000 employees, alongside a separate sample of 5,000 employees at similarly sized organizations, across the four countries noted above [1]. The decision-maker survey asked specifically about compliance and governance incidents tied to AI over the prior year, and the employee survey probed how people actually use AI tools day to day, including where they distrust or work around them. Coverage of the findings in Help Net Security and IT Security Guru both frame the core message the same way: process design, not policy language, is the proximate cause of most AI governance failures organizations are reporting right now [2][3].

A plausible mechanism, consistent with the survey's findings, is that most enterprise workflows, whether in finance, procurement, HR, or customer operations, were built around the assumption that a person performs each step: a person reviews an application, a person approves an exception, a person hands a

case to the next queue with enough context for the next person to pick it up. Those human-centered designs encode their compliance logic implicitly, through who has access to which system, who is expected to notice an anomaly, and who signs off before something moves forward. Under that reading, when an AI system is inserted into one or more of those steps without redesigning the surrounding process, the implicit logic could break down: the check that used to happen because a person was paying attention no longer happens, the handoff that used to generate a paper trail because a person typed a note into a ticket no longer generates one, and the audit record ends up unable to show how or why a decision was reached. Camunda's Kurt Petersen, SVP of Customer Success, put it directly: "Organizations are rapidly adopting AI. But they are applying it to processes designed for a world before AI, then wondering why the return on investment falls short." He added that processes "designed before the AI era cannot support the technology effectively without being re-engineered, no matter how much is spent on agents and models" [1].

This is consistent with, and reinforces, a separate finding CSA published earlier this month based on Schellman's 2026 State of AI Governance Report: 74% of enterprise leaders believe they could pass an AI compliance audit today, yet only 27% describe their AI governance programs as fully mature, and just 44% have documented AI incident response procedures [4]. That report characterized the gap as one between confidence and actual operational maturity [4]. The Camunda/Sapio data adds a mechanism to that gap: the shortfall is not primarily that organizations lack a written AI policy, but that the operational processes meant to carry out the policy were never rebuilt to account for AI's presence in the loop.

Security Analysis

The most consequential finding for security and compliance teams is that respondents attribute the underlying cause of AI-related compliance or governance issues to a process gap rather than a policy gap 84% of the time [1]. This matters because enterprise AI governance programs commonly invest in artifacts that are comparatively easy to produce and audit against a regulator's checklist — an acceptable-use policy, a model inventory, a risk classification scheme, a review board charter — a pattern the survey's findings are consistent with, even though it did not measure governance-program spending directly. Those artifacts are necessary, but they describe intent. They do not, by themselves, guarantee that the operational workflow a given AI system participates in still has a control at the point where a decision is made, still routes an exception to a human who can act on it, and still produces a record that shows why the system did what it did. The survey's employee-side data illustrates how that gap shows up in practice: 44% of employees say they manually override AI outputs [1] — a pattern consistent with processes that were not redesigned to accommodate AI's role in them — and 41% report using AI in ways that are more about satisfying an internal mandate than about the tool actually improving how the work

gets done [1]. Both behaviors are workarounds people invent when the formal process does not fit how AI actually operates, and both leave weaker documentation than the process was supposed to produce in the first place.

A second, related pattern is misplaced trust in the durability of controls once AI enters a process that was not redesigned around it. Sixty-one percent of decision-makers say business process redesign cannot keep pace with the speed at which they are expected to roll out AI, and 79% say it is organizationally easier to bolt AI onto an existing process than to redesign it [1], plausibly because redesign tends to generate more internal resistance. That preference for the path of least resistance illustrates one plausible way audit trails go missing: for example, a handoff that used to be logged because a person typed a summary into a case management system could become a silent API call, with no corresponding update to the process map, RACI assignment, or retention schedule. Separately, 48% of respondents say they are concerned about AI agents operating outside their intended scope [1]. This concern plausibly reflects the same dynamic described above — surrounding processes change without a corresponding update to the agent's guardrails — and maps onto the kind of agentic drift CSA's MAESTRO threat-modeling work addresses for autonomous and semi-autonomous systems.

Third, the survey surfaces a perception gap between leadership and the workforce that has its own governance implications. Eighty-nine percent of organizations report that AI has increased team productivity, but only 64% of employees agree, and 69% of employees say they were not fully consulted before an AI tool was introduced into their workflow [1]. It is plausible, though not established by the cited data, that employees who were not consulted are also less likely to flag early that a process step no longer makes sense once AI has been inserted into it — and that this dynamic overlaps with the override and mandate-compliance behaviors described above. Left unaddressed, this dynamic could compound over time, as unreported process mismatches accumulate faster than periodic audits are typically scheduled to catch them.

Finally, the financial dimension of these findings is directly relevant to how security and compliance teams should be prioritizing remediation. Seventy-two percent of organizations attribute a failed AI initiative to process-related challenges, at an average cost of $1.55 million, and 82% believe their AI investments will underperform without further investment in process redesign specifically [1]. For a security or risk function competing for budget against product and engineering priorities, this reframes process redesign as a cost-avoidance argument rather than a compliance-only argument: the same process gaps that create governance exposure are the ones driving AI project failure and wasted spend, which gives risk and audit functions a shared business case with operations and finance stakeholders who might otherwise treat governance work as overhead.

Recommendations

Immediate Actions

Security and compliance leaders should treat every AI deployment, including pilots, as a trigger for a targeted process review rather than a policy-compliance checklist exercise. Before an AI system goes live in a workflow, the team responsible for that workflow should walk through each control point, approval, and handoff the process previously relied on and explicitly determine, for each one, whether it still occurs, who or what performs it now, and what evidence it produces. Where a control point has effectively disappeared because a person is no longer in that part of the loop, it needs either a redesigned automated control or an explicit, documented decision that the risk is acceptable. Organizations should also inventory where employees are already manually overriding AI outputs or using AI only to satisfy a mandate, since both behaviors are early warning signs of the same process-design gaps this research identifies, and both are typically visible to frontline managers well before they surface in a formal audit.

Short-Term Mitigations

Over the following one to two quarters, organizations should build process-level audit trail requirements into their AI governance framework as a distinct control category, separate from model documentation and policy attestation. This means mapping, for each AI-touched workflow, exactly where a decision is made, what evidence is captured at that point, and how long it is retained, then validating that mapping against what the system actually does rather than what the original process diagram says it does. Given that 61% of decision-makers say redesign cannot keep pace with deployment speed, governance teams should also establish a lightweight, repeatable review cadence, rather than a one-time redesign, so that process-control mismatches introduced by iterative AI feature rollouts are caught within weeks rather than at the next annual audit. Employee consultation should be built into this cadence directly: since employees who were not consulted are the ones most likely to quietly route around a broken process, structured feedback channels from the people actually executing AI-touched workflows are a more reliable early-warning signal than periodic control testing alone.

Strategic Considerations

At a strategic level, this research supports treating business process redesign as a first-class component of AI governance maturity, on par with model risk management and data governance, rather than as an operations concern that sits outside the compliance function's remit. Organizations that continue to

invest primarily in AI policy documents and model-level controls, while leaving the surrounding human-centered processes largely unchanged, should expect the same category of process-driven compliance incidents this survey documents to recur, regardless of how sophisticated the policy layer becomes. Boards and executive sponsors evaluating AI governance investment should ask not only whether an AI policy exists and whether a model inventory is current, but whether the processes an AI system participates in have been explicitly redesigned, control by control, to account for its presence — and should treat a "no" answer to that question as a live compliance exposure rather than a future-state improvement item.

CSA Resource Alignment

This research note's central finding — that AI governance failures trace to process design rather than policy gaps — connects most directly to CSA's *Dynamic Process Landscape: A Strategic Guide to Successful AI Implementation* [5]. That guide addresses the same pattern this survey now quantifies: organizations bolting AI onto process designs built for human execution, then experiencing compliance and audit-trail failures because control points, data flows, and human-in-the-loop checkpoints were never re-mapped. Its Dynamic Process Landscape (DPL) methodology, which emphasizes process transparency, modular process blocks, and tamper-evident audit trails, is a direct operational response to the exact failure mode the Camunda/Sapio data describes, and organizations acting on this note's recommendations should treat that guide as their primary implementation reference.

The second most relevant CSA artifact is *AI Organizational Responsibilities: Governance, Risk Management, Compliance and Cultural Aspects* [6], which lays out RACI-based accountability models for AI governance and compliance functions. Where the Dynamic Process Landscape guide addresses how to redesign a workflow, this paper addresses who is accountable for making sure that redesign happens and stays current, which speaks directly to this note's recommendation that process-control mismatches be reviewed on a recurring cadence rather than caught only at annual audit.

Finally, organizations building or updating a formal AI governance program should anchor their control set to CSA's AI Controls Matrix (AICM) v1.1 [7], which extends CSA's Cloud Controls Matrix with AI-specific governance, risk, and compliance controls and should serve as the default reference for AI-specific control work. AICM provides the structured control catalog against which the process-level audit trail and control-mapping work recommended above can be assessed and attested, giving security and compliance teams a common reference point for turning the process-redesign priorities in this note into auditable control language.

References

[1] Camunda. "72% of Organizations Say Process-Related Challenges Have Caused AI Initiatives to Fail and it is Costing Them Millions." Camunda Press Release, September 2026.

[2] Help Net Security. "AI compliance issues hit 2 in 5 large companies, and legacy workflows are a big factor." Help Net Security, September 21, 2026.

[3] IT Security Guru. "Two in Five Organisations Hit by AI-Related Compliance Failures in Past Year, Research Finds." IT Security Guru, September 17, 2026.

[4] Cloud Security Alliance. "Enterprise Reality: Why Organizations Aren't as Prepared for AI Governance as They Think They Are." CSA Blog, September 16, 2026.

[5] Cloud Security Alliance. "Dynamic Process Landscape: A Strategic Guide to Successful AI Implementation." CSA, 2025.

[6] Cloud Security Alliance. "AI Organizational Responsibilities - Governance, Risk Management, Compliance and Cultural Aspects." CSA, 2024.

[7] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." CSA, 2026.