IRIS · ESSAY TWO · 2026

One Constraint Lifts. Another Endures.

BY JAY DUNNE SERIES PART 2 OF 3 DOMAIN ENTERPRISE RISK INTELLIGENCE

The field has added new technical capabilities year upon year and decade after decade. Better tools, more data, sharper analytics, more comprehensive dashboards, faster reporting cycles. And yet the problems have remained the same. Security and risk functions still cannot communicate effectively with the C-suite. The different functions and risk domains still don't play nice with each other. The function still lives in constant fear of budget cuts and workforce reductions, and struggles to produce meaningful metrics. It is still seen by the business as a cost center rather than as a source of strategic value.

This is the disconnect that should be holding the attention of every senior executive paying attention to the current AI wave. The capability story is real. Vendors are not lying about what the technology now enables — continuous monitoring, cross-domain signal correlation, translation throughput, pattern detection at scales no human analyst could process. The industry literature has been documenting it for the past two years. New platforms launch monthly. The procurement conversations are happening across nearly every enterprise large enough to have a risk function.

And yet the pathologies persist. The integration mandate that ISO 31000[1] and COSO[2] have called for since 2017 — that Part 1 of this series examined as a structural diagnosis — is no closer to being fulfilled in 2026 than it was in 2017, despite a decade of capability gains.[3] That tells you something important about where the actual binding constraint lives. The capabilities have been improving. The pathologies have remained constant. They cannot both be true unless the constraint shaping the field is somewhere other than where the industry conversation has been looking.

The capability gains are worth describing more precisely. The dominant industry conversation tends to flatten and narrow them into vendor claims, but the underlying technical shift is more substantive than the marketing suggests.

AI deployment in enterprise risk functions has compressed analyst workloads dramatically. Tasks that previously required weeks of manual data structuring — tagging imagery with location and time stamps, identifying entities across unstructured chatter, parsing regulatory documents at scale, correlating signals across financial, operational, and cyber data feeds — now resolve in hours or minutes. The intelligence community has been articulating this shift in their own voice. John Urby, a former Army intelligence analyst now at GDIT, has captured the underlying problem plainly: *"There is too much data out there and not enough analysts available to sift through it. Without artificial intelligence and automation, it would be impossible for analysts to get through every piece of data to verify what is valid or invalid and what's worth continuing to look at."*[4] The same dynamic now operates in enterprise risk. Synthesis across volumes of data that would have required Fortune 50 analyst budgets is becoming accessible to organizations several scales smaller.

The capability gains are real and economically meaningful, but the greatest return from these new capabilities has yet to break into the conversation.

The evidence sits in plain view, as the same conditions that have defined the enterprise risk function for decades persist alongside the new capabilities. The COSO 2026 study cited in Part 1 documented the integration gap with hard numbers: only 7% of ERM programs are fully integrated into strategy decisions, 54% are still perceived as compliance or assurance functions, 98% of risk practitioners believe ERM should play a more strategic role than it does.[3] Those figures describe the field as it stands today, not as it stood before the AI wave. The capabilities have been arriving for several years now. The pathologies have not moved.

That is the puzzle Part 2 is going to explain. It is not enough to observe that capability has improved and pathology has persisted. The argument requires naming why the two are decoupled — why decades of technical advancement have left the structural conditions untouched. And the answer requires looking at what the field has been mistaking as its binding constraint.

The diagnosis is uncomfortable, and it cuts against the dominant narrative. Part 1 named the visible scarcity. Part 2 names the binding one. The constraint that has shaped enterprise risk intelligence for a generation has not been the absence of analytic capability. It has been the absence of organizational integration. The analyst-hour economics that Part 1 described were, on closer examination, downstream of something more fundamental. The reason most enterprises could not afford continuous cross-domain synthesis was not just that analytic labor was expensive. It was that no organizational structure existed to absorb the synthesis productively, no decision pathway existed to route it into capital allocation and strategic choice, no executive sponsor existed to defend the function's authority. The labor was expensive, in part, because the institutional architecture surrounding it was unbuilt. The constraint was structural all along. The analyst-hour problem was just the visible scarcity that made it tractable to name.

This is the insight Geof Kahn has been articulating for the intelligence community, where AI deployment has been running ahead of corporate enterprise adoption by several years. Writing in The Cipher Brief, Kahn — a former senior CIA and HPSCI staff intelligence professional — names the pattern plainly: *"It's tradecraft,*

not technology, that is the primary constraint on intelligence performance in the AI era."[5] His argument is that AI is compressing parts of the intelligence cycle but modernization is occurring unevenly across collection, analysis, validation, dissemination, and policymaker integration. The result is what he calls asynchronous modernization: faster collection overwhelming coordination, faster analysis outpacing dissemination, automated insight generation challenging validation and trust. The technology has moved. The system around it has not.

The parallel to enterprise risk intelligence is direct. The technical capabilities that the industry is now selling — continuous monitoring, cross-domain signal correlation, pattern detection at scale, translation throughput — are real. They are, for the first time, economically accessible to organizations that are not Fortune 50 firms or intelligence agencies. And yet the structural conditions that have defined enterprise risk for decades remain unchanged. The synthesis still does not reach executive decisions in time to shape them. Domain owners still operate from incompatible frameworks. The function still defends its budget every cycle and still struggles to demonstrate its impact in language the business recognizes.

The reason is that the binding constraint has moved. It is no longer analytic throughput. It is the organizational architecture surrounding the analytic function — the decision rights, the reporting lines, the executive sponsorship, the cross-domain governance, the translation capability between threat language and business impact language. None of those are technology problems. None of them will be solved by procurement. And none of them are receiving the strategic attention the AI shift actually requires.

The current moment is not subtle in its risks. The dominant industry conversation around AI in enterprise risk is being shaped by procurement decisions, vendor narratives, and tooling deployments. Most enterprises are treating the shift as a question of what to buy and what to install. That framing will reproduce the same pathology the field has been carrying for a generation — just with new technology.

The failure mode is structurally predictable. A risk function adds AI-enabled monitoring tools to its existing workflow. The tools generate more signals at greater speed. The function now produces more alerts, more dashboards, more reports. The executives who were not reading the old reports are not reading the new ones either. The cross-domain governance problem persists because no procurement decision can touch decision rights. The budget defensiveness persists because the function still cannot demonstrate impact in business language. The cost-center perception persists because the function's deliverables remain operational rather than strategic. The capability has improved, but the position of the function has not.

This is not a hypothetical concern. The intelligence community has been articulating this risk explicitly. Jeff Baker, writing about generative AI as a force multiplier in military and intelligence contexts, has made the point directly: *"You cannot simply drop a commercial AI model into this environment and expect it to work."*[6] His argument is that systems integration into existing workflows is the harder problem than the model itself, and that organizations that miss this distinction will spend significantly on capability without changing how decisions actually get made. The enterprise risk function is heading directly toward the same trap.

There is also a peer-reviewed academic literature documenting that AI-enabled analytical systems do detect signals human analysts would miss. Work in the *NIH/Scientific Reports* literature has demonstrated this in domains from food safety to crisis event detection.[7] The capability is real and the research grounding is real. But none of that research addresses the organizational question. The studies establish that the technology can detect; they do not establish that organizations will act on what gets detected. The research is necessary grounding for the capability claim. It is not sufficient grounding for the transformation claim.

The vendor literature obscures this distinction. Workday's framing of AI helping risk teams *"move faster, see further, and act sooner"*[8] is true at the technical level and misleading at the organizational level. ChatFin's claims of *"30% reduced risk exposure and 50% reduced compliance costs"*[9] describe outcomes that some organizations may achieve and that most will not, because the outcome depends on organizational design, not procurement. The enterprises that read these claims as descriptions of what AI delivers — rather than as descriptions of what AI enables under specific organizational conditions — will deploy capability and find the structural problems unchanged.

The enterprises that read the moment correctly will recognize that sole focus on the procurement question is the wrong strategy. The strategic question is whether the organization is positioned to absorb the capability productively — and most enterprises, on current trajectory, are not.

The executives who navigate this shift successfully will not be the ones who move fastest on AI procurement. They will be the ones who recognize that the AI wave has not created the integration problem in enterprise risk intelligence, but rather exposed it. The structural pathologies that the field has been carrying for a generation — the C-suite disconnect, the domain silos, the budget defensiveness, the metrics problem, the cost-center perception — are now visible in a new and consequential way. The technology has made the integration question impossible to avoid, because capability is no longer the bottleneck and there is nowhere left to hide the organizational design problem.

The strategic question that follows is not what tools to deploy, but whether the enterprise's organizational architecture can absorb cross-domain synthesis productively. Where does the function sit when it stops being defined by what it monitors and starts being defined by what it integrates? Who staffs it? How do its outputs reach the rooms where capital, strategy, and operational decisions actually get made? What authority does it carry when it enters those rooms? Those are governance questions, design questions, and leadership questions. They are the questions the procurement conversation cannot touch.

The function that has spent decades defending its budget and translating threat language into reports that executives do not read will not, by adding AI-enabled tooling, suddenly become the strategic synthesis layer the moment requires. The capability gains will arrive on top of unchanged organizational pathology. The dashboards will become more numerous, the reports will produce faster, and the cost-center perception will remain.

The enterprises that translate this moment into durable advantage will be the ones that treat the AI shift as the trigger to address the organizational design question that has been quietly defining the function for a generation. The capability is finally available. The question is whether the institutional architecture can rise to meet it.

A follow-up piece will look more directly at what the redesigned function should become.

WORKS CITED

[1] International Organization for Standardization. (2018). *ISO 31000:2018 — Risk Management — Guidelines.*

[2] Committee of Sponsoring Organizations of the Treadway Commission. (2017). *Enterprise Risk Management — Integrating with Strategy and Performance.*

[3] Luttenton, R., Samp, S., & Stone, A. (2026). *From Guidance to Action: Exploring Practical Enterprise Risk Management.* Committee of Sponsoring Organizations of the Treadway Commission (COSO).

[4] Scoop News Group. (2025, November 12). Too much data, too few analysts: How AI offers a “force multiplier” for intelligence analysts. *DefenseScoop.* Sponsored by General Dynamics Information Technology (GDIT).

[5] Kahn, G. (2026). AI is Speeding Up Intelligence, But Not the System Around It. *The Cipher Brief.*

[6] Baker, J. *GenAI as a Force Multiplier: Turning Data into Decision Advantage.* SAIC.

[7] Selected studies in *Scientific Reports* (Nature Portfolio) on AI-enabled detection of crisis events and food safety risks. See for example Gao, Q., Levi, R., & Renegar, N. (2022). Leveraging machine learning to assess market-level food safety and zoonotic disease risks in China. *Scientific Reports* 12, 21650.

[8] Workday. *AI in Enterprise Risk Management: What You Need to Know.* Workday Perspectives.

[9] ChatFin. (2026, January). *Top 10 Best AI Tools for Enterprise Risk Management (ERM) — 2026 Edition.*