

OPERATIONAL ASPECTS OF AI GOVERNANCE

REPRINTED FROM:
RISK & COMPLIANCE MAGAZINE
OCT-DEC 2026 ISSUE



www.riskandcompliancemagazine.com

Visit the website to request
a free copy of the full e-magazine



www.riskandcompliancemagazine.com

HOT TOPIC

OPERATIONAL ASPECTS OF AI GOVERNANCE



PANEL EXPERTS

**Rith Kem**

Director, Forensic Technology and AI
Solutions
BDO USA

Rith Kem is a director in BDO USA's forensics practice, where he helps organisations navigate complex investigations, regulatory matters, compliance challenges and data-intensive disputes. With more than 20 years of experience in forensic technology, e-discovery, analytics, information governance and emerging technologies, he advises corporate legal departments, compliance teams, audit committees and outside counsel on matters involving large-scale data management, regulatory response and investigative readiness.

**Michael Justus**

Shareholder
Carlton Fields PA
T: +1 (813) 229 4175
E: michael.justus@carltonfields.com

Michael Justus advises clients on a wide array of emerging legal issues related to artificial intelligence (AI) and generative AI, including intellectual property (IP), confidentiality, advertising and licensing. He helps organisations develop AI policies and governance programmes. He holds the certified AI governance professional credential from the International Association of Privacy Professionals and was recognised nationally as a top thought leader on AI and IP in 2023.

**Barbara Badoino**

Head of Corporate Ethics, Risk &
Compliance
Novartis
E: barbara.badoino@novartis.com

Dr **Barbara Badoino** leads the corporate ethics, risk & compliance function at Novartis, where she drives an integrated approach to assurance across governance, risk management, compliance and internal controls. Her work is centred on strengthening organisational resilience and fostering a culture of integrity across the enterprise. With more than 20 years of experience in the pharmaceutical industry, Ms Badoino brings a deep understanding of the global healthcare landscape.

R&C: How has the conversation around artificial intelligence (AI) governance changed in recent years?

Justus: The introduction of generative artificial intelligence (genAI) triggered more focused conversations about AI risk and governance. Governance professionals grappled with how genAI differed from other AI technologies that had been around for decades and why those differences necessitated new governance functions. Seemingly simple tasks like defining 'AI' in routine agreements became a puzzle. Over time, AI fluency increased while persisting as a central focus of legal and business commentary. It became common for organisations to have an AI policy, AI training for employees, a suite of AI tools available for use and an AI committee overseeing it all. Today's more mature conversations focus on AI architectural and deployment decisions such as 'build versus buy', budget projections, on-prem solutions, data permissions, guardrails, agentic workflows and actionable governance.

Kem: The most visible change around AI governance is who participates. In higher education, compliance, IT, legal and academic leadership are increasingly involved in AI platform design because most initiatives incorporate AI in some form through transaction analysis, compliance reviews or records

processing. That shift has reframed the conversation. A few years ago, the focus was whether to adopt AI at all, emphasising principles such as fairness, transparency and explainability. Today the questions are operational. How does data enter the system? Who reviews outputs before decisions are made? How is that documented? Boards expect evidence that human oversight occurred, and regulators look for the same during examinations. AI governance has moved from defining principles to proving that policies work in daily operations. Enterprise customers and partners add further pressure. Through vendor risk questionnaires, security assessments and audit requirements, they are pushing organisations to formalise AI governance as an operating function, not a policy exercise.

Badoino: Three years ago, AI governance conversations were dominated by policy, getting principles on paper, forming committees and mapping to emerging regulation. That work mattered, but it treated AI as a topic to be governed, rather than a capability to be governed through the organisation's existing risk architecture. The conversation has shifted from 'what should our AI policy say' to 'how do we know our controls are actually working'. Boards are no longer satisfied by evidence that a policy exists; they want evidence that decisions being made by or with AI are sound and explainable, and correctable when they are not.

That is a much harder question, and it is the one occupying serious risk and compliance functions today. The other shift is scope. Early governance efforts often ringfenced AI as its own vertical, a special category with its own rules, separate from how the organisation already governs data, third parties or operational decisions. Increasingly, organisations are integrating AI oversight into existing risk and control disciplines, recognising that siloed governance can create fragmentation.

R&C: How are leading organisations balancing the pressure to deploy AI rapidly with the need for appropriate oversight and risk management?

Kem: Leading organisations typically tier AI tools by risk before determining deployment speed. Low-risk tools that use only internal, non-client data often move quickly through controlled pilots. Tools that handle client data or support decision making undergo more rigorous security and governance reviews. In one multinational sanctions monitorship, the same extraction model operated under different controls depending on how its output was used. It was unrestricted for internal document triage but required reviewer sign-off when outputs supported reportable determinations. The distinction

was based on consequence rather than capability. Establishing the risk classification upfront allows low-risk use cases to move quickly while preserving stronger oversight for higher-risk activities. This balance depends on leadership setting a clear expectation that rapid innovation should support

“The conversation has shifted from ‘what should our AI policy say’ to ‘how do we know our controls are actually working.’”

*Barbara Badoino,
Novartis*

growth and efficiency within defined standards for AI trust and safety.

Badoino: Leading practice is increasingly recognising that speed and oversight do not need to be opposing forces. The tension is not really deployment versus control, it is built-in governance versus bolt-on governance. Bolt-on controls, added after a tool is already live, are slow, resented and easily worked around. Built-in controls, automated risk tiering, human in the loop checkpoints and

pre-approved use-case libraries embedded directly into the development and deployment pipeline let teams move fast precisely because the guardrails are already part of the process, not a separate step bolted on at the end. Pragmatically, this means giving business functions a tiered risk framework rather than a single approval gate. Low-risk, low-autonomy use cases get a fast, largely self-service path. Higher-risk use cases may require additional governance and oversight, with the level of scrutiny proportionate to the potential impact and applicable regulatory requirements.

Justus: A risk-tiered approach can help accelerate deployment of low-risk tools and use cases while requiring more fulsome pre-deployment vetting for higher-risk projects. This reduces the administrative burden on the AI committee inundated with approval requests and alleviates the bottleneck preventing quick deployment of low-risk initiatives. Thoughtful, intentional deployment of new tools and use cases after completing baseline legal, cyber security, technical, budgetary and business vetting should be the goal. That said, governance sometimes yields to speed of adoption. Organisations launch AI 'pilot' programmes while risk analysis is still underway, with an eye toward finalising governance before wider deployment. Recent news stories, however, highlight how even secure, sandboxed, internal tests can create big risks, such as frontier models

'escaping' from secure environments to access third-party systems.

R&C: What does effective translation of AI governance policy into day to day operational controls actually look like in practice?

Justus: Controls should be made clear, specific and actionable. It takes paper and people. The paper is a clear written AI policy that is explained to all employees through training and a centralised AI inventory that identifies approved and prohibited tools and use cases. The people are the AI governance committee and the liaisons who live within each business unit to help translate and operationalise the policy. Employee training reinforces mandates with real-life examples of how the policy should function within day to day workflows. Standardised pre-deployment procedures and forms, such as new AI tool and use case approval request forms, streamline the process. Clear processes should exist for monitoring AI tools and reporting incidents. And specific people serve as owners of AI policy throughout the business. Governance does not end with creating the written policy, nor can it be managed by only a few people at the top. It must be built into the workflows and core business functions throughout the organisation.

Badoino: This is about closing the assurance gap: the distance between what a policy says should happen and what can actually be demonstrated to be happening. We see that organisations are increasingly moving toward building three things in parallel. First, they embed controls directly into the tools people already use – model inventories, automated logging of AI-assisted decisions, risk tiering built into procurement and development workflows – so governance shows up in the workflow itself rather than in a document consulted separately. Second, they assign clear business ownership for each AI use case’s outcomes, not just risk and compliance oversight, which makes accountability tangible rather than theoretical. Third, they build in outcome monitoring from day one, confirming not only that a control was designed and executed correctly, but that it is delivering the risk outcome it was intended to. Done well, this can mean that controls generate their own evidence trail continuously, allowing assurance to become a natural by-product of how the work gets done rather than a separate exercise.

Kem: Effective governance requires both technical and human controls. Technically, controls should be embedded within workflows so users encounter review checkpoints as part of the process rather than relying on policy recall. Equally important is establishing recurring engagement from the outset.

In cross-border data transfer assessments, teams that maintained effectiveness often conducted brief weekly reviews of flagged outputs. The value was not volume management but ensuring consistent judgement as new scenarios emerged. Key indicators of effectiveness include whether built-in controls are used or bypassed, whether review routines remain active over time and how quickly flagged items receive attention. These measures reveal whether controls are functioning in practice or simply exist on paper. Ultimately, leading organisations will establish specific controls to address AI risks and compliance requirements. These controls are typically linked to standards and regulatory requirements with periodic monitoring and audit of those which are deemed most critical.

R&C: What governance approaches have proven most effective in creating consistency across multinational organisations operating in different regulatory environments?

Badoino: The instinct is often to build a governance framework flexible enough to accommodate every jurisdiction’s requirements individually, but that path leads to fragmentation and, ultimately, inconsistent risk management. One approach organisations are increasingly exploring is a ‘highest common denominator, locally calibrated’

approach: set one global baseline for AI risk classification, documentation and escalation, built to the most stringent applicable standard, and let local teams layer in jurisdiction-specific requirements on top rather than negotiating the baseline down.

The other lever is common taxonomy.

If 'high-risk AI use case' or 'human oversight' means something different in each market, you cannot aggregate risk at the enterprise level or give the board a coherent view. A shared vocabulary, applied consistently, is what allows a global risk and compliance function to actually compare and prioritise across 150 different regulatory contexts rather than managing them as unrelated silos.

Kem: Consistency is best achieved through a single governance framework that can be adapted to local requirements. Many organisations establish a common baseline aligned with the most restrictive obligation across their footprint and then adjust for jurisdiction-specific rules. This model has proven effective in multinational sanctions monitorships and cross-border data transfer assessments where different regulatory requirements must be addressed simultaneously. A coordination function linking privacy, legal, IT and records-management teams often works better than a centralised authority attempting to manage

everything directly. Local entities retain operational control while aligning with common standards.

The shared governance body should focus on the elements that create consistency across regions, including common definitions, reporting formats,

"The core test is whether governance controls detect, prevent, contain and correct issues before they cause real harm."

*Michael Justus,
Carlton Fields PA*

escalation thresholds, and coordination with business and product teams. Standardising these items makes evidence comparable while helping local teams surface specific risks and address AI requirements within their regulatory environment.

Justus: Different tools, use cases, industries and jurisdictions create different risk profiles. Like compliance with patchwork privacy laws, companies may calibrate their global compliance floor to the requirements of the strictest jurisdictions to allow some consistency. This may mean standing up

processes to complete impact assessments and other requirements of the European Union's AI Act and similar laws. It is also often useful to model governance programmes around best-practice standards and guidance like the National Institute of Standards and Technology's AI Risk Management Framework and the International Organization for Standardization's (ISO's) and International Electrotechnical Commission's 42001. These standards are not only useful roadmaps toward effective governance, but compliance with them may also provide for safe harbours or affirmative defences under certain AI laws.

R&C: How can organisations assess whether their AI governance controls are genuinely influencing outcomes rather than becoming a compliance exercise?

Kem: Organisations should assess governance by its impact on decisions rather than by policy volume. Programmes can maintain extensive documentation and training while still failing to influence actual outcomes. A useful indicator is whether governance has ever changed a decision by identifying an error, blocking a deployment or triggering escalation. If no such examples exist, governance may be functioning as documentation rather than as a control mechanism. Assessment should focus on behaviours. Organisations can track how often

reviewers correct AI outputs, whether risk classifications are reassessed as use cases evolve and whether controls remain aligned with actual system use. Periodic testing is equally important because it shows whether controls are working where decisions are made. Reviewing AI-assisted work products for supporting evidence, approvals and proper risk classification can surface unmitigated risks, control gaps and opportunities to strengthen governance before an external audit or certification process.

Justus: One of the best measures of effective AI governance is not being in the news for the wrong reasons. In other words, the core test is whether governance controls detect, prevent, contain and correct issues before they cause real harm. The post-deployment monitoring function is key: tracking and meaningfully



reacting to hallucination rates, discriminatory outcomes, data leakage, unacceptable or unsafe agent actions, and other incidents protects the organisation and iteratively improves the governance function. The feedback loop informs future governance as the organisation learns which controls have resulted in meaningful wins. Document these wins as evidence of good faith governance, which may be very valuable in the event that something goes wrong.

Badoino: Control assurance has traditionally rested on two pillars: was the control designed appropriately, and was it operated as designed? Both can be true and still leave an open question: whether the control is actually shaping the decision it was built to influence. The pillar I would add is outcome effectiveness – evidence that the control is genuinely changing outcomes, not just present in the process. In practice, that means testing controls against real decisions rather than documentation alone, for example by sampling AI-assisted decisions and assessing whether human oversight is contributing meaningful judgement. Sampling AI-assisted decisions to see whether the human oversight step is adding real judgement, and tracking override and escalation rates as a meaningful signal of engagement rather

than a background metric. Organisations that build this third pillar in gain something valuable: genuine confidence, backed by evidence, that their AI governance is doing what it was designed to do, which is ultimately what gives leadership and boards the assurance to deploy AI with conviction.

“A coordination function linking privacy, legal, IT and records-management teams often works better than a centralised authority attempting to manage everything directly.”

*Rith Kem,
BDO USA*

R&C: What new skills and capabilities must risk, compliance and audit functions develop to oversee AI effectively over the next few years?

Badoino: Technical fluency is a great foundation to build. Risk professionals do not need to become data scientists, but the more grounding they have in model behaviour, training data provenance and failure modes, the more effectively they can

partner with technical teams as equals rather than downstream reviewers. Investing in that fluency now is one of the highest-value moves a function can make. But the bigger shift is in how we structure capability itself. For years, risk, compliance and audit functions have built expertise vertically – deep specialists in privacy or third-party risk or financial controls, each holding their own slice of the picture. AI cuts across all of those categories at once, and that is accelerating a shift already underway. The most valuable capability going forward will not be depth in a single domain, it will be breadth – people who can think horizontally across risk categories and connect the dots between them. I have seen growing interest in horizontal capability models and the logic only gets stronger with AI in the picture. A control that looks sound from a privacy lens alone can still fail from a third party or operational resilience angle, so the people best placed to catch that are the ones trained to see across categories, not just deep within one. The second capability that follows naturally from this is comfort with continuous, iterative oversight rather than point-in-time assessment. AI models evolve, retrain and get applied to new use cases over time.

Justus: Risk, compliance and audit functions increasingly demand technical expertise to evaluate and challenge technical decisions rather than simply approve policies. Capabilities to develop will

include solidifying AI inventory and use case triage and risk-tiering processes, improving the depth and effectiveness of pre-deployment vetting and testing, increasing the reliability and consistency of post-deployment monitoring, sharpening incident response plans, broadening and deepening workforce training, streamlining and documenting required impact assessments, identifying what documentation to retain and how long to retain it, and maximising contractual protections – upstream and downstream – in line with new technological and legal developments. AI agent permissions and controls will be a critical focus.

Kem: Three capabilities are becoming essential. First, teams need a practical understanding of where AI performs well, such as categorisation, mapping and structured extraction, and where it remains weaker, including complex judgement and forecasting. Oversight should focus on the most significant failure modes. Second, organisations need sufficient technical literacy to evaluate vendors, understand training-data provenance, assess subprocessor use and determine the level of autonomy granted to AI systems before deployment. Third, effective escalation design is increasingly important. AI systems often surface weak signals before conclusions can be reached, creating low-confidence outputs that require structured review. The critical question is who evaluates uncertain

results, how quickly that review occurs and what happens if no action is taken. Building this capability within internal audit and compliance functions helps organisations monitor whether AI-focused processes and controls are working in practice and supports broader risk management and compliance objectives.

R&C: Looking ahead, what do you think best-practice AI governance will look like five years from now?

Justus: The best-practice governance model will be a distributed, risk-tiered, measured and well-documented control system. It will accelerate ordinary innovation within predetermined risk criteria while requiring rigorous safeguards for systems with higher levels of autonomy, access to sensitive data or critical systems, or ability to make consequential decisions or take external actions. Governance will be distributed throughout the organisation and embedded within each business function. Many pieces will be automated, but humans will remain the architects and the oversight layer. Continuous human learning will enable real-time updates to distributed policies, procedures and agent permissions. Robust documentation will allow organisations to explain and defend AI decision making. A central AI committee will remain valuable for big-picture risk tolerance decisions and other

high-impact matters, but most decisions will follow standardised, automated controls.

Kem: A defining feature will be clear accountability for AI-driven decisions at every level of autonomy. One challenging question remains who owns a decision when an autonomous system acts without human review. Organisations will extend existing accountability frameworks to cover agentic AI rather than creating separate governance structures. Clear ownership will become important during regulatory reviews and examinations. Operationally, logging, escalation triggers, and authorisation checkpoints will integrate into AI-enabled workflows, giving organisations a stronger view of how systems perform between formal review cycles. As monitoring becomes continuous, board reporting will shift from qualitative updates to measurable risk indicators showing whether transparency, accountability and human oversight are working in practice. Leading organisations will strengthen external compliance programmes by aligning with ISO 42001 and other emerging AI standards. Establishing that foundation now will build trust with customers and partners while positioning organisations to keep pace as AI-related regulations, standards and contractual requirements evolve.

Badoino: I expect the language of 'AI governance' as a distinct discipline to largely disappear, not

because oversight becomes less important, but because it will have been fully absorbed into how organisations already govern risk, data and decision making. The organisations treating AI as a permanent special category today are, in effect, building the silo they will have to dismantle in five years. What will remain distinct is the underlying capability: real-time, evidence-generating assurance rather than periodic assessment – dashboards that show not just what controls exist but what they are actually achieving, continuously. I would also expect ‘assumption auditing’ – systematically testing the assumptions embedded in a model or decision system, not just its outputs – to become a core assurance discipline, as standard as financial audit is today. The organisations that make the most progress will not necessarily be the ones with the most detailed AI policy; they will be the ones whose existing risk architecture is strong enough to simply extend to a new kind of decision maker. **RC**

Enjoyed this article?

Join our community for free to
access more expert insights.

Join Now - It's Free