MINIMUM VIABLE GOVERNANCE FOR GENERATIVE AI

Nick van der Meulen, *Research Scientist*
MIT Center for Information Systems Research (CISR)
Jennifer Jewer, *Associate Professor of Information Systems*, Memorial University of Newfoundland and *Research Collaborator*, MIT CISR
Nadège Levallet, *Associate Professor of Management and Information Systems*, University of Maine and *Research Collaborator*, MIT CISR

Organizations investing in generative artificial intelligence (GenAI) face a dual challenge: understanding emergent risks from the technology, and putting mechanisms in place to manage these risks without stifling innovation.[1] Many leaders are discovering that their organizations' existing governance approaches are failing the second challenge. When applying traditional governance to GenAI, organizations find themselves oscillating between extremes: uncontrolled risk without governance, and paralysis with governance.

"Governance designed for technologies with 50-year life cycles doesn't work when the technology itself transforms every 18 months."

What organizations need instead is **minimum viable governance**: the least amount of governance required to manage risk effectively while enabling the organization to sense and seize opportunities. Drawing on an in-depth case study of an organization in a highly regulated environment and prior MIT CISR research,[2] this briefing identifies what distinguishes minimum viable governance from prior governance approaches and explores the characteristics that leaders can use to evaluate and redesign their governance mechanisms.

WHY TRADITIONAL GOVERNANCE FALLS SHORT

Organizations govern technology through mechanisms that can be organized into five domains: **principles**, **policies**, **people**, **processes**, and **platforms**. Principles articulate guiding values; policies translate those values into enforceable rules. People are organized in structures with clear decision rights. Processes define repeatable workflows. And platforms provide the infrastructure to enforce governance and monitor the use of technology at scale. When leaders perceive a governance gap, the instinct is to fill every domain as quickly and comprehensively as possible—which carries its own risks.

Consider the experience of FinCo, a global diversified financial services firm.[3] When GenAI emerged in late 2022, employees across FinCo's federated business units quickly began experimenting with it independently—drafting client communications, summarizing reports, generating marketing copy—using unsanctioned GenAI tools. Risk-sensitive executives became alarmed. In an industry where data breaches and regulatory violations carry existential consequences, the executives moved swiftly to establish formal governance.

FinCo's response was broad in scope, with the firm taking action in all five of the governance domains. Corporate data, IT, and ethics units collaborated on *principles* for responsible AI use, articulating guiding values such as no automated decision-making and privacy by design. FinCo's board then requested a comprehensive enterprise AI *policy* to translate those values into enforceable rules, which took close to a year and involved hundreds of stakeholders. For use case oversight, FinCo established AI Review Committees (ARCs): *people* organized in formal cross-functional structures with clear decision rights. The ARCs reviewed proposals through a repeatable *process* in which risk ratings determined routing. Regional ARCs handled low- and medium-risk cases monthly; a corporate ARC reviewed high-risk proposals quarterly. Finally, FinCo also built controls into its *platforms*, such as FinGPT, a secure internal wrapper for approved large language models designed for safe experimentation.

By any conventional measure, FinCo's governance was thorough. A year later, innovation had ground to a halt. The comprehensive policy was already outdated by the time

1 For a detailed mapping of where GenAI risks emerge, see N. van der Meulen, H. Lefebvre, and B. H. Wixom, "Mapping the Generative AI Risk Space," MIT CISR Research Briefing, Vol. XXVI, No. 1, January 2026, https://cisr.mit.edu/publication/2026_0101_GenerativeAIRisk_VanderMeulenLefebvreWixom.

2 See Van der Meulen, Lefebvre, and Wixom, "Mapping the Generative AI Risk Space;" and N. van der Meulen, "Realizing Decentralized Economies of Scale," MIT CISR Research Briefing, Vol. XXIII, No. 1, January 2023, https://cisr.mit.edu/publication/2023_0101_DecentralizedDecisionMaking_VanderMeulen.

3 FinCo is a pseudonym. By protecting the organization's identity, we can share lessons from the firm's experience candidly. The authors conducted interviews with seventeen leaders at this organization in 2025.

© 2026 MIT Center for Information Systems Research, Van der Meulen, Jewer, and Levallet. MIT CISR Research Briefings are published monthly to update the center's member organizations on current research projects.

it was complete. Low-risk proposals stalled for months in ARC review cycles (for example, a low-risk agent prototype involving no sensitive data took six months to be approved). Employees needed sign-off from both the legal team and the relevant ARC just to access FinGPT, a platform that was already designed for safe experimentation. Business sponsors of GenAI initiatives faced committees where risk-oriented voices consistently outnumbered those advocating for pursuing new opportunities. Frustrated teams reverted to unsanctioned GenAI tools. GenAI vendors pitched solutions directly to business leaders who couldn't get traction through official channels. Shadow GenAI—the unauthorized use of GenAI tools and solutions—spread more widely than before.[4] The board, recognizing that FinCo had swung from unchecked adoption to paralysis—only to see unsanctioned use return—issued an urgent mandate to redesign the organization's governance approach.

FinCo had mechanisms in every domain. The question is whether those mechanisms were suited to GenAI's unique demands. Traditional governance assumes stable technologies, predictable risks, and manageable demand for a new technology. GenAI upends these assumptions. Its natural language interface, ubiquitous availability, and broad applicability fuel adoption that outpaces any organization's capacity for centralized review. Simultaneously, probabilistic models, performance drift, and compounding risk across interdependent components create a risk space that shifts faster than leaders can anticipate. As one FinCo executive observed, "Governance designed for technologies with 50-year life cycles doesn't work when the technology itself transforms every 18 months."

FOUR CHARACTERISTICS OF MINIMUM VIABLE GOVERNANCE

The answer is governance that calibrates to what it governs. Minimum viable governance builds on the MIT CISR concept of minimum viable policy, which showed that foundational principles can reduce the need for comprehensive policies, giving teams greater operational decision rights while safeguarding business continuity. In our research, organizations with well-developed minimum viable policy practices cut the average time to make complex decisions in half, and their teams identified new opportunities at three times the rate of peers without such practices.[5] Minimum viable governance applies this philosophy of governing at the minimum level required across all five governance domains. Achieving this requires four characteristics that distinguish minimum viable governance from traditional governance: it is **structurally agile**, **trustworthy by design**, **integrated end-to-end**, and **opportunity-sensitive**.

Structurally Agile

For governance to keep pace with GenAI, leaders must design mechanisms that can be introduced, adjusted, or retired as conditions change. The mechanisms FinCo leveraged in its principles were sound, but those it used in its other governance domains were too rigid. The firm's policy cycle took close to a year, producing a document already outdated upon completion. Meanwhile, the ARCs' fixed cadences could not accommodate the volume of proposals, creating queues that gave low-risk requests no speed advantage over high-stakes ones.

Structurally agile governance matches review intensity to proposal complexity. For low-risk initiatives, leaders should consider replacing committee review with lighter mechanisms that can be adjusted as the organization's understanding of GenAI matures, such as pre-approved categories or delegated authority. One highly regulated organization we studied classified initiatives into three tiers: high-risk initiatives received full committee oversight; mid-tier initiatives could proceed on self-service platforms with pre-configured controls; and teams with low-risk initiatives proceeded autonomously with approved tools. Routing decisions were made based on simple intake forms that could be reviewed asynchronously. The organization also incorporated a mandatory "look back" mechanism into their approach, periodically revisiting initiatives to determine whether their risk level had changed. As one leader explained, "In twelve months' time, if this initiative really grows legs, we're going to revisit it." The organization also consolidated fragmented governance committees as they learned what worked, retiring structures that had made sense earlier but were no longer fit for purpose.

Trustworthy by Design

For governance to scale without creating bottlenecks, leaders must build oversight into governance mechanisms themselves. FinCo's FinGPT blocked outbound internet traffic, logged conversations, and masked personally identifiable information; it was designed for governed experimentation. Nonetheless, FinCo gated access to FinGPT behind sign-offs from both the legal team and the ARC, wrapping a platform designed for experimentation in a process demanding permission.

4 For more on shadow GenAI, see N. van der Meulen and B. H. Wixom, "Managing the Two Faces of Generative AI," MIT CISR Research Briefing, Vol. XXIV, No. 9, September 2024, https://cisr.mit.edu/publication/2024_0901_GenAI_VanderMeulenWixom.

5 See N. van der Meulen, "Realizing Decentralized Economies of Scale."

Providing secure, ready-to-use GenAI platforms with embedded controls is more effective than restricting access.[6] One healthcare organization we studied built proxy services directly in front of its large language models that allow the organization to automatically log every interaction, analyze outputs for hallucinations, and filter for policy violations. When platforms enforce controls and flag anomalies in this way, approval processes (and the committees that carry them out) can shift from gatekeeping before action to monitoring and intervening as needed. Governance becomes trustworthy by design when it rests on an auditable trail of prompts, outputs, and human decisions—making oversight continuous and verifiable.

Integrated End-to-End

For governance to address the full GenAI risk space, mechanisms must function as a unified system. FinCo brought multiple functions together in its ARCs, but outside the committee room, governance fragmented. Risk, compliance, legal, procurement, and architecture teams each developed their own assessment criteria. Teams proposing GenAI initiatives had to navigate these parallel processes, with no coordination between governance mechanisms.

In another financial services firm, the enterprise data platform team embedded a dedicated risk officer directly into cross-functional teams from day one, participating in all operations meetings. As the platform leader described it, the goal was "bringing them along for the ride rather than intersecting with them at some point." This approach gave the team a direct link to legal and audit functions throughout development, enabling the team to pass a full-scale audit with zero findings. GenAI governance must span the entire risk space, which demands that governance functions become part of the initiatives they oversee.

Opportunity-Sensitive

For governance to serve opportunity as well as manage risk, leaders must account for the cost of inaction. FinCo's conservative posture meant that business sponsors encountered committees oriented only toward identifying concerns. When initiatives stalled, the shadow GenAI that FinCo's governance was designed to prevent returned.

Effective GenAI governance reframes what counts as risk. As a governance leader at another large financial services firm told us, "The biggest risk is that we move too slowly, because a slow and cumbersome oversight process creates a vacuum filled by other actors who may not have our clients' best interests at heart." Some organizations operationalize this by adopting a solutions-first posture: teams develop proposals assuming no restrictions, then refine them with input from legal and compliance to identify genuine constraints. As one executive put it, "If we start with restrictions, we're going to end up with very narrow proposals." Tracking time-to-decision alongside risk incidents makes these trade-offs visible—and reveals whether governance is protecting the organization or holding it back.

BUILDING YOUR GENAI GOVERNANCE PLAYBOOK

The four characteristics describe what minimum viable governance looks like. Achieving it requires mechanisms across all five domains—principles, policies, people, processes, and platforms—functioning as a mutually reinforcing system. Principles that guide judgment enable action when policies haven't anticipated the scenario. People with standing to advocate for opportunity prevent governance from becoming overly risk averse. Processes that differentiate initiatives by risk level match scrutiny to stakes. Platforms that enforce controls reduce the need for lengthy approval processes.

For each mechanism, ask: can it adapt as conditions change? Does it build oversight in? Does it integrate with mechanisms in other domains? And does it account for the cost of delay alongside risk?

Leaders should examine each existing governance mechanism against the four minimum viable governance characteristics. For each mechanism, ask: can it adapt as conditions change? Does it build oversight in, or fall back on approvals? Does it integrate with mechanisms in other domains? And does it account for the cost of delay alongside risk?

Ultimately, minimum viable governance exists between two failure modes. Its ceiling is the point at which governance impedes innovation more than it reduces risk, as signaled by growing shadow GenAI and lengthening time-to-decision. Its floor is the point at which governance exposes the organization to unacceptable risk, as signaled by rising risk incidents or gaps in audit trails. Every mechanism that pushes governance outside this range should be questioned. Specific mechanisms will vary by industry, risk tolerance, and organizational structure. What matters is that they reinforce each other, and that leaders treat governance as a capability to develop continuously—with the same urgency they bring to the technology it governs.

6 See N. van der Meulen and B. H. Wixom, "Bring Your Own AI: How to Balance Risks and Innovation," *MIT Sloan Management Review*, October 3, 2024, https://sloanreview.mit.edu/article/bring-your-own-ai-how-to-balance-risks-and-innovation/.

MIT CENTER FOR INFORMATION SYSTEMS RESEARCH (CISR)

MIT CISR helps executives meet the challenge of leading increasingly digital and data-driven organizations. We provide insights on how organizations effectively realize value from approaches such as digital business transformation, data monetization, business ecosystems, and the digital workplace. Founded in 1974 and grounded in MIT's tradition of combining academic knowledge and practical purpose, we work directly with digital leaders, executives, and boards to develop our insights. Our research is funded by member organizations that support our work and participate in our consortium.

CISR RESEARCH PATRONS

AlixPartners
Avanade
Cognizant
Collibra
IFS
PwC

CISR SPONSORS

ABN Group (Australia)
Alcon Vision
ANZ Banking Group (Australia)
AustralianSuper
Banco Bradesco S.A. (Brazil)
Barclays (UK)
BNP Paribas (France)
Bupa (Australia)
CalSTRS
Caterpillar, Inc.
Cemex (Mexico)
Cencora
CIBC (Canada)
Commonwealth Superannuation Corp. (Australia)
Cuscal Limited (Australia)
Dawn Foods
DBS Bank Ltd. (Singapore)
Doosan Corporation (Korea)
Ericsson (Sweden)
Fidelity Investments
Fomento Economico Mexicano, S.A.B., de C.V.
Genentech
HCF (Australia)
Hunter Water (Australia)
International Motors
JERA Co., Inc.
Jewelers Mutual
JPMorgan Chase
Kaiser Permanente
Keurig Dr Pepper
King & Wood Mallesons (Australia)
Mater Private Hospital (Ireland)
Nasdaq, Inc.
National Australia Bank, Ltd.
Nomura Holdings, Inc. (Japan)
Nomura Research Institute, Ltd. Systems Consulting Division (Japan)
Novo Nordisk A/S (Denmark)
OCP Group
Pentagon Federal Credit Union
Posten Bring AS (Norway)
Principal Life Insurance Company
Ralliant
Reserve Bank of Austrailia
RTX
Saint-Gobain
Scentre Group Limited (Australia)
Schneider Electric Industries SAS (France)
TabCorp Holdings (Australia)
Telstra Limited (Australia)
Terumo Corporation (Japan)
UniSuper Management Pty Ltd (Australia)
Uniting (Australia)
Vanguard
WestRock Company
Xenco Medical
Zoetis Services LLC

Information as of March 2026

MIT CISR is funded by our members, and we gratefully acknowledge their financial support and their many contributions to our work.

Membership and benefits: cisr.mit.edu/content/member-benefits

MIT CISR research publications: cisr.mit.edu/research-library

MIT Sloan School of Management
Center for Information Systems Research

245 First Street, E94-15th Floor
Cambridge, MA 02142
t 617-253-2348 | e cisr@mit.edu

cisr.mit.edu | in

Team | Isobela Byerly-Chapman, Margherita Di Pinto, Christine G. Foglia *Associate Director*, Dorothea Gray-Papastathis, Cheryl A. Miller, Ina M. Sebastian, Alan Thorogood, Nick van der Meulen, Peter Weill *Chairman*, Barbara H. Wixom, Stephanie L. Woerner *Director*