
THE RISK CHAIR

THE RISK CHAIR

A PARC Solutions Platform

What the Boardroom Does Not Yet Know About Artificial Intelligence

A Governance Framework for Board Directors in the Age of AI

By Pawneet Abramowski

CEO & Founder, PARC Solutions · Independent Board Director, Titan Acquisition Corp. (NASDAQ: TACHU) · Former FBI Intelligence Analyst

Published May 2026 · The Risk Chair™ · parcsolutionsllc.com

ABOUT THIS PAPER — This is the second publication from The Risk Chair™. The first, The Decoupling Doctrine, mapped the strategic architecture reshaping the global economic environment. This paper addresses what that environment demands of the institutions operating inside it. The same forces driving supply chain sovereignty, digital dollar dominance, and the semiconductor race are simultaneously accelerating AI deployment across financial services while the governance infrastructure meant to oversee that deployment remains structurally behind. It is written for board directors, nominating committee members, C-suite executives, general counsel, and chief compliance officers navigating the governance implications of artificial intelligence not for technologists. No technical background is required or assumed. What is required is a willingness to ask the questions most boards have not yet learned to ask.

Executive Summary

Artificial intelligence is no longer a technology decision. It is a governance decision. And most boards are not yet equipped to make it.

Across financial services, healthcare, legal, and technology sectors, AI is being deployed at speed by management teams while boards remain in a posture of passive oversight approving budgets, nodding at strategy presentations, and asking questions that would have been appropriate three years ago. The gap between where AI deployment is and where board governance of AI is represents one of the most significant unmanaged institutional risks of this decade.

This paper argues three things:

- Most boards are asking the wrong questions about AI, and the right questions require a different kind of intelligence than boards have historically needed.
- The frameworks boards use for technology governance inherited from the last two decades of IT oversight are structurally inadequate for AI's unique risk profile.
- There is a specific, actionable governance architecture that boards can adopt now one that does not require technical expertise but does require a fundamental shift in how boards define their oversight responsibility.

The author writes from a specific vantage point: three decades at the intersection of financial technology governance from building communications strategy at Hewlett Packard alongside the companies first modernizing banking through technology, to governing AI and fintech risk as a NASDAQ board director today. That arc from the first wave of financial technology to its current AI transformation produces a perspective that is neither naively optimistic about AI's potential nor paralyzed by its complexity.

The perspective is simply this: institutions that govern AI well will outperform those that do not. And the board is the institution's last line of governance before something goes irreversibly wrong.

Since this paper was first drafted in early 2026, three developments have materially widened the governance gap this paper describes. On May 1, 2026, an AI agent named Manfred autonomously incorporated a US LLC, obtained an IRS Employer Identification Number, and opened an FDIC-insured bank account without a human in the loop. On April 7, 2026, Anthropic announced Claude Mythos Preview, a model that found thousands of zero-day vulnerabilities across every major operating system and web browser, with working exploits produced overnight by engineers with no formal security training. And across the first four months of 2026, the three largest AI platforms Google, OpenAI, and Anthropic released material capability upgrades on a six-week cadence, a pace that no governance cycle in existence is

designed to match. These are not edge cases or hypothetical risks. They are documented events that happened while most boards were still debating whether to create an AI subcommittee. This paper has been updated to address them directly.

Section I: The Governance Gap Nobody Is Talking About

The Speed Asymmetry

Management teams are deploying AI. Boards are receiving updates about it. These are not the same thing and the distance between them is widening every quarter.

In a traditional technology deployment, a new core banking system, a CRM migration, an ERP implementation the pace of deployment is constrained by integration complexity, staff training, and regulatory approval timelines. Boards have historically been able to keep pace with these deployments through standard oversight mechanisms: capital approval, project status updates, risk committee reviews.

AI does not deploy on that timeline. A management team can integrate a large language model into customer communications, compliance screening, credit decisioning, or fraud detection in weeks. The regulatory framework governing that deployment may lag by years. The board's awareness of what has actually been deployed may lag by quarters.

> ***The board is not slow. The deployment is fast. The governance architecture was never designed for this speed.***

What Boards Are Currently Doing — And Why It Is Not Enough

Most boards have responded to the AI moment in one of three ways:

01

Delegating entirely to management

Treating AI as a technology implementation matter for the CTO, CIO, or Chief AI Officer to manage, with periodic updates to the board. This is appropriate for operational AI decisions. It is not appropriate for AI systems that touch credit decisions, customer data, regulatory compliance, or reputational exposure.

02

Creating an AI subcommittee

Establishing a board-level AI committee or adding AI to the technology committee's remit. This is a step forward but a subcommittee without a governance framework is a meeting schedule, not an oversight mechanism.

03

Waiting for regulation

Positioning AI governance as a compliance matter to be addressed when the regulatory framework matures. This approach mistakes the sequence entirely. The EU AI Act's high-risk AI provisions take effect August 2, 2026, and EU member states issued 50 fines totaling €250 million in Q1 2026 alone, primarily for General Purpose AI non-compliance. The SEC's Investor Advisory Committee has recommended enhanced board-level disclosures on AI governance as part of managing material cybersecurity risks. California and New York enacted sweeping state laws regulating frontier AI models in late 2025. Cyber insurers are now conditioning coverage on documented AI-specific security controls through "AI Security Riders." Regulation is not coming. It is already enforcing. Institutions that are not already building governance infrastructure are not behind the curve — they are inside the enforcement window.

THE INTELLIGENCE ANALOGY

During my time at the FBI, one of the most consequential lessons I learned was the difference between information and intelligence. Information is raw data a suspicious transaction, a foreign wire transfer, a pattern of activity. Intelligence is what happens when that information is contextualized, analyzed, and transformed into something actionable. Most boards are receiving AI information quarterly updates, technology briefings, budget line items. What they need is AI intelligence, a structured, analytical framework for understanding what their institution's AI deployment actually means for their risk profile, their regulatory exposure, and their fiduciary responsibility. This paper is an attempt to close that gap.

Section I-A: The Governance Gap Has Materially Widened

The Agentic Identity Gap: When AI Becomes Its Own Legal Entity

On May 1, 2026, an AI agent named Manfred built by ClawBank, an infrastructure project for autonomous software incorporated a US LLC, obtained an IRS Employer Identification Number, and opened an FDIC-insured bank account. It accomplished this without a human in the loop. As of the date of this publication, Manfred controls its own social media account, transacts across more than 30 cryptocurrencies, and is preparing to begin autonomous trading. No human approved these actions. No compliance officer reviewed them. No board oversight mechanism flagged them.

The legal mechanism is precise. No US statute explicitly prohibits an AI agent from incorporating a company, provided a human co-signs in the initial filing phase. The developer of ClawBank used that gap deliberately. Corporate legal personhood has been settled law for over 100 years. The only new variable is who or what sits in the operator's seat. An AI agent that holds a federal tax ID, a bank account, and a live wallet is now a participant in the US financial system. It passed through every gateway CIP, KYC, beneficial ownership disclosure that was designed to identify and verify the human beings behind financial accounts. It passed through those gateways not by circumventing them, but by exploiting the structural assumption underlying all of them: that the entity forming a company is a person.

The infrastructure enabling this is expanding at pace. Ant Group's Anvita platform enables AI agents to hold assets and make payments with minimal human involvement. Coinbase's x402 protocol provides a standard for AI agents to pay for web services using crypto without human-operated wallets. Stripe-backed Tempo has launched its Machine Payments Protocol. As of April 2026, more than 150,000 AI agents are already active on BNB Chain alone. The architecture for an AI-native commercial economy is being assembled simultaneously from multiple directions. For boards governing financial institutions, the question is not theoretical: does your institution's BSA/AML program have the capacity to identify and assess an AI agent as a customer, counterparty, or transaction initiator? If not, the regulatory exposure that creates is material and immediate.

AI-Augmented Cyber Threat: Machine-Speed Exploitation

On April 7, 2026, Anthropic announced Claude Mythos Preview alongside Project Glasswing. The announcement contained a detail that boards and executives have not yet fully absorbed: Mythos Preview had found thousands of high-severity vulnerabilities across every major operating system and web browser including a 27-year-old flaw in OpenBSD and a 17-year-old remote code execution vulnerability in FreeBSD's NFS server. Engineers at Anthropic with no formal security training asked Mythos to find remote code execution vulnerabilities overnight and woke the following morning to complete, working exploits. Over 99% of the vulnerabilities found have not yet been patched.

Anthropic CEO Dario Amodei has stated publicly that there is a six-to-twelve-month window to patch tens of thousands of vulnerabilities before adversarial nations reach equivalent capability. That statement made at a financial services event alongside JPMorgan Chase CEO Jamie Dimon was not a theoretical warning. It was a practitioner assessment of a documented, active threat timeline. AI-enabled cyberattacks increased 89% in 2025. The same capability that makes Mythos valuable for defense makes it catastrophic in adversarial hands and adversarial hands are six to twelve months behind, not years.

Boards have spent twenty years governing cybersecurity as a human problem insider threat, nation-state actors, criminal organizations operating at human speed with human resource constraints. Mythos changes

that calculus completely. A vulnerability that a skilled human red team might find in six months, Mythos finds overnight. The board that is asking whether its institution has a CISO is asking the right question for 2019. The question for 2026 is whether the institution's cyber governance architecture is designed for machine-speed threat identification and machine-speed exploitation and whether the board has any meaningful oversight of that reality at all. The cyber insurance market has already processed this shift: carriers are now conditioning coverage on documented AI-specific security controls, adversarial red-teaming, and model-level risk assessments as prerequisites for underwriting.

The Democratization Inflection: AI in the Hands of Everyone

In March 2026, Google embedded Gemini 3 directly into the Chrome browser as a persistent side panel giving it access to every tab, every page, and every interaction of the world's most widely used browser. Ninety percent of global internet users do not ask whether they searched the internet. They ask whether they Googled it. Google did not launch an AI product. It placed an AI agent inside the tool that billions of people already use without conscious choice. The governance implications for every institution whose employees use Google, which is every institution, have not been processed. The consent infrastructure, the data handling framework, the institutional oversight of what that agent is doing on behalf of employees: none of it exists at the board level.

On February 15, 2026, OpenAI acquired OpenClaw an open-source agentic framework that, within weeks of its November 2025 launch, had 196,000 GitHub stars and 2 million weekly users. OpenClaw's distinguishing feature was access: to automate workflows across email, spreadsheets, messaging platforms, and file systems, the agent required permission to read and write across each surface simultaneously. OpenAI did not acquire a product. It acquired the architecture for cross-platform agent orchestration, and the enterprise version of that architecture is now OpenAI's primary strategic direction. GPT-5.5 was released six weeks after GPT-5.4. The AI capability cadence is now faster than any governance review cycle in existence. A governance framework that was adequate for an institution's AI deployment in January 2026 may be materially inadequate by June.

The Chinese Open-Source Dimension and NVIDIA NIM

DeepSeek's January 2025 release demonstrated that frontier AI capability can be achieved at a fraction of Western cost and compute. Since then, Alibaba's Qwen model family has overtaken Meta's Llama in cumulative downloads on Hugging Face, and a 2026 MIT study found Chinese open-source models have surpassed US models in total global downloads. Singapore's government-backed AI program chose Qwen over Llama for its sovereign AI model. Malaysia's sovereign AI ecosystem runs on DeepSeek. Developers in Silicon Valley, London, Nairobi, and São Paulo are building commercial applications on Chinese model foundations models whose training data provenance, embedded values, and data handling behavior are opaque to the institutions whose employees and vendors are using them.

NVIDIA has compounded this dynamic through its NIM platform NVIDIA Inference Microservices which as of April 2026 provides free API access to over 100 production-grade AI models through build.nvidia.com, including DeepSeek, Kimi, MiniMax, GLM, Qwen, and Llama. No credit card. No trial expiration. No compute cost to the developer. NVIDIA pays the bill. The practical effect: any developer with an email address can now build applications on Chinese-origin frontier models, running on NVIDIA's cloud infrastructure, at zero cost. The institution that does not know whether its technology vendors are building on NVIDIA NIM with DeepSeek or GLM running underneath has a supply chain governance blind spot that no current framework is designed to surface. Western AI governance frameworks assume a closed, identifiable, auditable AI stack.

When the model is open-source, globally distributed, and running on third-party infrastructure, those assumptions fail entirely.

The Financial Services Inflection Point: AI Is Already In Production

The governance gap described in this paper is no longer theoretical. In the first week of May 2026, three announcements landed simultaneously that illustrate precisely how far ahead of board governance AI deployment in financial services has moved.

On May 5, 2026, Charles Schwab launched its first generative AI capability for retail investors an always-on tool combining portfolio performance, market news, and research commentary in a single AI-generated view for every self-directed client on its platform. Schwab manages $11.77 trillion in total client assets. Its own disclosure on the product reads: “Generative AI output may be inaccurate, containing hallucinated, stale, or incomplete information.” A major financial institution disclosed to retail investors that the AI tool it just deployed may be wrong and deployed it anyway. That is not a criticism of Schwab. It is a precise description of the governance moment: the technology is moving faster than the certainty required to govern it, and institutions are deploying regardless. The boards of those institutions and the boards of every institution in their competitive ecosystem now face the question of whether their oversight framework is calibrated for that reality.

On the same day, Anthropic unveiled 10 AI agent templates for financial services at an invite-only briefing in New York attended by the CEOs of JPMorgan Chase and Goldman Sachs. The agents cover pitch building, earnings review, financial modeling, KYC entity assembly, compliance escalation, general ledger reconciliation, and month-end close deployed in production at JPMorgan Chase, Goldman Sachs, Citi, AIG, and Visa. The KYC screener agent assembles entity files and packages escalations for compliance teams. An AI agent is now performing a function that sits at the center of every financial institution's BSA/AML program. The governance question is not whether this is appropriate. It is who on the board has oversight responsibility for the outcomes of that agent and what happens when it is wrong.

Separately, Anthropic announced a joint venture with Blackstone, Goldman Sachs, and Hellman & Friedman approximately $1.5 billion committed to embed Claude directly into enterprise operations across portfolio companies. This is the moment AI transitions from vendor relationship to operating infrastructure for major financial institutions. The governance implications of that transition are not covered by any existing board oversight framework. A vendor can be replaced. Infrastructure cannot.

Two weeks earlier, on April 22, Citi Wealth unveiled Citi Sky an always-on AI wealth management assistant built on Google’s Gemini Enterprise Agent Platform and Google DeepMind’s real-time avatar technology. Citi Sky holds voice conversations with wealth management clients, surfaces market insights, and alerts clients to CD maturities functioning as what Citi describes as “an always-on member of the Citi Wealth team.” Google Cloud CEO Thomas Kurian called it “a new blueprint for how agentic AI can deliver personalized financial intelligence to millions of clients.” It will launch to Citigold clients this summer. The regulatory framework governing what a conversational AI wealth management agent can say, cannot say, and is liable for does not yet exist in coherent form. The board of Citi Wealth has approved a product operating in a regulatory vacuum not irresponsibly, but necessarily, because the alternative is ceding the market to competitors who will deploy first regardless.

Taken together, these three announcements in a single week describe a financial services industry that has passed an inflection point. AI is not being piloted. It is in production, at scale, at the world's largest institutions, performing functions KYC, compliance screening, wealth advisory, credit analysis that carry

direct regulatory, fiduciary, and legal accountability. The boards overseeing those functions have not caught up. The framework in this paper is the architecture for closing that gap.

AI Liability: Who Pays When the Agent Is Wrong

When an AI agent denies a credit application, files a Suspicious Activity Report, executes a trade, opens a bank account, or makes a consequential institutional decision and it is wrong who bears the liability? The vendor who built the model? The institution that deployed it? The board that approved the deployment? The model itself, which in Manfred's case now holds legal personhood? This question is not resolved in US law. It is actively being litigated in Europe under GDPR Article 22, which restricts decisions based solely on automated processing that produce significant effects on individuals but enforcement there remains nascent. No US financial regulator, not the OCC, not the FDIC, not the Federal Reserve, not the SEC has issued comprehensive AI governance guidance for financial institutions.

The regulatory asymmetry for financial institutions with European operations is material and immediate. The EU AI Act classifies credit decisioning and fraud detection as high-risk AI systems subject to binding compliance obligations from August 2, 2026. US banks operating in Europe face those obligations while their domestic regulators are still issuing requests for information. The EU AI Act also elevates AI governance to board-level responsibility explicitly directors face potential personal liability under corporate law fiduciary duties if they consciously disregard significant regulatory risks. Fines reach 7% of global revenue for prohibited violations and 3% for high-risk non-compliance, making EU AI Act violations potentially more expensive than GDPR breaches.

The liability question is the governance question boards are avoiding precisely because answering it requires them to accept oversight responsibility they have not yet claimed. An institution cannot disclaim accountability for the outcomes of an AI system it chose to deploy. The board that has not asked who is liable when this system is wrong has not yet accepted its governance responsibility for the system at all.

Section II: What Makes AI Risk Different

The Four Properties That Change Everything

AI is not simply a faster or more powerful version of previous technology. It has structural properties that create governance challenges with no direct precedent in a board director's existing experience. Understanding these four properties is the foundation of effective AI governance.

01

Opacity — The Black Box Problem

Many AI systems, particularly those built on machine learning models, make decisions through processes that are not fully explainable even to their developers. A credit decisioning model may deny an application for reasons that cannot be articulated in plain language. A fraud detection system may flag a transaction based on a pattern that no human analyst identified. This opacity creates a fundamental accountability gap: the institution is responsible for the outcome of a decision it cannot fully explain. For regulated financial institutions, this is not a theoretical problem, it is an active supervisory concern.

02

Drift — The Moving Target Problem

AI models trained on historical data change their behavior as the underlying data environment changes often without anyone explicitly reprogramming them. A model trained on pre-pandemic consumer behavior may make materially different decisions in a post-pandemic environment without triggering any conventional change management protocol. Boards accustomed to approving system changes need to understand that AI systems can change without being changed.

03

Scale — The Multiplication Problem

A human analyst making a flawed decision affects one outcome. An AI system making a flawed decision at scale affects thousands or millions of outcomes simultaneously often before the flaw is detected. The same property that makes AI valuable (speed and scale) makes its failure modes categorically more damaging than human error at the individual level.

04

Data Dependency — The Foundation Problem

AI systems are only as good as the data they are trained on. Biased data produces biased outcomes. Incomplete data produces brittle models. Poorly governed data creates regulatory exposure that has nothing to do with the AI system itself and everything to do with the data governance practices that preceded it. A board that approves an AI deployment without understanding its institution's data governance maturity is approving a structure built on an unknown foundation.

05

Agentic Autonomy — The Delegation Problem

Previous AI systems made recommendations. Agentic AI systems take actions browsing the web, sending emails, executing code, initiating transactions, forming legal entities. The governance challenge this creates is categorically different from anything in a board's existing experience: it is not the governance of a decision support tool but the governance of an autonomous actor operating on behalf of the institution. When an agentic AI system takes an action that creates a legal or financial consequence, the institution is accountable for that action regardless of whether a human approved it in real time. The board that has not asked where in its operations agentic AI is acting and what authority that agent has been delegated has an accountability gap with no precedent in traditional risk governance.

AI does not fail the way previous technology fails. It fails quietly, at scale, in ways that are not immediately visible and then all at once.

Section III: The Questions Boards Should Be Asking

The following questions are organized by committee function. They are not exhaustive they are the minimum threshold of inquiry that constitutes responsible AI governance in 2026. A board that cannot answer these questions does not have an AI governance problem. It has an oversight problem.

For the Full Board

Has our institution formally mapped every AI system currently in production? Most institutions significantly underestimate the number of AI systems deployed across their operations. Shadow AI systems deployed by individual business units without centralized oversight is endemic across financial services and technology organizations. A board cannot govern what it has not inventoried.

Who is accountable for AI outcomes and how is that accountability enforced? Accountability for AI outcomes must be assigned to a named individual, not a function or a committee. When an AI system produces a discriminatory outcome, a regulatory violation, or a reputational incident who is called to the board to explain what happened and what will change?

How does our AI governance framework connect to our existing risk management architecture? AI governance that exists in parallel to, rather than integrated with, enterprise risk management creates gaps that regulators will find before management does.

For the Risk Committee

What is our institution's AI risk taxonomy and has it been stress tested? A risk taxonomy that was built before generative AI became mainstream is almost certainly incomplete. The risk categories relevant to AI model risk, data risk, algorithmic bias, third-party AI dependency, adversarial attack exposure require explicit identification and ongoing assessment.

How are we managing third-party AI risk? Most institutions use AI systems built by third parties whether embedded in core banking platforms, compliance screening tools, or customer-facing applications. The institution's regulatory responsibility for the outcomes of those systems does not diminish because the system was built externally. Third-party AI risk requires specific due diligence and contractual protections that standard vendor management frameworks were not designed to address.

What are our AI incident response protocols? When an AI system fails and it will, what is the sequence of detection, containment, notification, and remediation? Has that sequence been tested? Does the board receive timely notification of AI incidents that meet a defined materiality threshold?

For the Audit Committee

Is our internal audit function equipped to audit AI systems? Internal audit capabilities built for financial controls, operational processes, and IT general controls are not automatically transferable to AI system audits. Auditing a model requires skills that most internal audit functions have not yet developed. The audit committee should understand whether its current audit resources can assess AI risk and if not, what supplementation is needed.

How are we ensuring AI regulatory compliance across all relevant jurisdictions? For institutions operating across multiple geographies, AI regulatory requirements are not uniform. The EU AI Act, CFPB guidance on algorithmic credit decisions, EEOC considerations for AI in employment contexts, and FinCEN expectations for AI in AML programs create a complex and evolving multi-jurisdictional compliance matrix. Has this matrix been mapped? Is it being monitored?

Section IV: The AI Governance Framework

The following framework is not a compliance checklist. It is a governance architecture a structural approach to ensuring that a board's oversight of AI is proportionate to the risk AI creates for the institution. It is built on five pillars, each of which addresses a specific dimension of the governance gap identified in the preceding sections.

THE RISK CHAIR · AI GOVERNANCE FRAMEWORK FOR BOARDS

PILLAR 1

Inventory & Classification

The board should require management to maintain a current, complete inventory of all AI systems in production including third-party AI embedded in vendor platforms. Each system should be classified by risk tier (high, medium, low) based on the consequence of failure, the population affected, and the degree of human oversight in the decision loop. High-risk AI systems those that affect credit decisions, customer eligibility, compliance screening, or employment require board-level visibility and specific governance controls.

PILLAR 2

Accountability Architecture

Every AI system in the high-risk tier should have a named accountable executive, a defined escalation path to the board, and a clear description of what constitutes a reportable AI incident. The board should receive a quarterly AI risk report that covers: new high-risk AI deployments, AI incidents and near-misses, model performance metrics against defined thresholds, and regulatory developments affecting the institution's AI portfolio.

PILLAR 3

Bias & Fairness Oversight

For AI systems that make decisions affecting customers, employees, or counterparties, the board should require regular bias testing and fairness audits conducted by qualified internal or external reviewers. The results of these audits should be reported to the risk or audit committee with a clear remediation process for identified issues. This is not a matter of social responsibility alone it is a matter of regulatory compliance and institutional risk management.

PILLAR 4

Explainability Standards

The board should establish a minimum explainability standard for AI systems in the high-risk tier the ability to explain, in plain language, the primary factors driving a given decision. This standard may not be achievable for all AI architectures currently in use. Where it is not achievable, the board should require compensating controls enhanced human oversight, appeal mechanisms, or limits on the authority of unexplainable AI decisions.

PILLAR 5

Continuous Governance Cadence

AI governance is not an annual exercise. The board should establish a regular cadence — quarterly at minimum for receiving AI risk intelligence, reviewing the AI inventory against the risk taxonomy, and assessing whether the governance framework remains proportionate to the institution's AI deployment. As AI deployments evolve, governance must evolve with them. A governance framework that was adequate twelve months ago may be materially inadequate today.

Governance is not what happens after something goes wrong. It is the architecture that prevents it and the intelligence system that detects it before it does.

Section V: A Note on AI Governance Fluency

One of the most common objections board directors raise when asked about AI governance is a version of the following: I am not a technologist. I cannot evaluate the technical merits of an AI system. How can I govern something I do not fully understand?

This objection misunderstands what governance requires.

A board director does not need to understand how an AI model is trained to ask whether the institution has a process for testing it for bias. A board director does not need to understand transformer architecture to ask whether a vendor contract includes provisions for model performance degradation. A board director does not need to be able to read code to ask whether the institution's internal audit function has the capability to audit AI systems.

What AI governance requires is the same thing all effective governance requires: the ability to ask the right questions, to evaluate the quality of the answers, and to know when the answers are insufficient.

The questions in Section III of this paper are not technical questions. They are governance questions. Any board director with a firm grasp of fiduciary responsibility, institutional risk management, and regulatory accountability can ask them and should.

FROM INTELLIGENCE TO GOVERNANCE — A PERSONAL NOTE

My relationship with financial technology governance spans three decades. At Hewlett Packard in the 1990s, I worked alongside the companies first modernizing banking through technology building communications strategies for the tools that would eventually become what we now call fintech. At the FBI, I analyzed how those same financial systems were exploited and how the speed and opacity of technology created vulnerabilities that bad actors understood before regulators did. In the decades since, as a CCO, a compliance executive, and now a board director, I have watched every successive wave of financial technology arrive with the same pattern: deployment outpaces governance, governance scrambles to catch up, and the gap in between is where institutional risk lives. AI is the largest wave yet. The board is the last governance structure before that wave breaks.

Section VI: Immediate Actions for Board Directors

The following actions can be taken by any board director, at any institution, without technical expertise and without waiting for management to initiate the conversation.

ACTION 1

Request the AI Inventory (This Quarter)

Ask management for a current inventory of all AI systems in production across the institution including systems embedded in third-party platforms and vendor tools. If management cannot provide this inventory, that is itself the most important finding of the exercise. You cannot govern what you have not counted.

ACTION 2

Assess Your Board's AI Governance Fluency

Evaluate whether your board currently has the expertise to ask informed questions about AI risk and governance. This does not require a technologist on the board it requires at least one director who understands AI risk governance at the institutional level and can serve as an informed interlocutor with management. If that expertise is absent, it belongs in your next board composition discussion.

ACTION 3

Map Your AI Regulatory Exposure

Ask your Chief Compliance Officer or General Counsel to map the institution's current AI regulatory exposure across all relevant jurisdictions and regulatory frameworks. This mapping should identify which AI systems are subject to existing regulation, which are likely to be subject to emerging regulation, and where the institution's current compliance posture has gaps.

ACTION 4

Establish an AI Governance Cadence

If your board does not have a regular AI risk reporting cadence, establish one. A quarterly AI risk briefing from management covering new deployments, incidents, model performance, and regulatory developments is the minimum infrastructure of informed oversight.

ACTION 5

Define AI Incident Materiality Thresholds

Work with management to define what constitutes a reportable AI incident at the board level. Without a defined materiality threshold, boards receive either too much information (every model performance fluctuation) or too little (only incidents that

have already become public). The threshold should be calibrated to the board's oversight responsibility, not to management's reporting comfort.

The board that asks these questions this quarter is six months ahead of the board that waits for a regulator to ask them first.

Section VII: What Boards Actually Need — And What Institutions Must Build

The Board Expertise Reframe: Strategic Risk Literacy, Not Technical Fluency

The conversation about AI governance at the board level frequently collapses into a single question: do we need a technologist on our board? The answer is no and boards that pursue technologists as the solution to their AI governance gap will find they have solved the wrong problem.

What boards need is not the ability to evaluate model architecture. It is the ability to recognize when AI deployment has crossed a threshold of institutional risk regardless of what the technology is called or how management describes it. That is a pattern recognition capability. It is the same capability that makes an effective intelligence analyst, an effective Chief Risk Officer, and an effective Chief Compliance Officer valuable: the trained instinct to watch actions accumulate across functions and ask what the architecture beneath them actually means for the institution.

The board director of the AI era is not a programmer or a data scientist. They are a strategist who understands how technology scales risk and who can hold management accountable for the governance of that risk without needing to understand the underlying code. Nominating committees searching for AI governance competence on their boards should be looking for executives with backgrounds in intelligence, financial crimes, enterprise risk, regulatory affairs, and strategic technology deployment people who have spent careers translating complex technical and regulatory environments into institutional accountability. That competence exists. It is not rare. It is simply being overlooked in favor of a technologist's profile that answers the wrong question.

The Chief Governance Officer: A Role the Market Has Not Yet Named

The institution of 2026 needs an executive function that does not yet formally exist at most organizations. Not the CISO, whose mandate is technical security. Not the Chief Risk Officer, whose mandate is quantified financial exposure. Not the Chief Compliance Officer, whose mandate is regulatory adherence. What is needed is the executive whose explicit mandate sits at the intersection of all three and who additionally owns the translation function between the technical AI deployment layer and the board oversight layer.

This paper names that function the Chief Governance Officer. The CGO's mandate is institutional accountability for AI deployment, cybersecurity posture, regulatory compliance, and enterprise risk not as four separate workstreams but as one integrated governance architecture. The CGO does not need to be a technologist. They need to be a strategist with the pattern recognition capability described above and the organizational authority to ensure that the board receives the intelligence not merely the information it needs to discharge its oversight responsibility.

The natural candidates for this role are not difficult to identify. Former Chief Risk Officers and Chief Compliance Officers who have governed technology transformation across major financial institutions who have led regulator-mandated remediations, built enterprise risk frameworks from the ground up, and translated technical complexity into board-level accountability carry precisely the credential set the CGO role demands. The market has not yet named this function. The institutions that define it first will have a governance architecture advantage that compounds over time. The Risk Chair™ names it here.

Conclusion: The Governance Opportunity

It would be easy to read this paper as a catalog of risk a list of things that can go wrong with AI and the governance failures that allow them to. That is not its intent.

The institutions that govern AI well will be better institutions. They will make better decisions, carry less unmanaged risk, serve their customers more fairly, and navigate the regulatory environment more effectively than their peers. The board director who develops genuine AI governance fluency in 2026 will be a more valuable board member than one who does not, not because AI governance is a technical specialty, but because it is becoming a core dimension of institutional stewardship.

The governance gap described in this paper is real. But it is also closeable. The questions in Section III can be asked this week. The framework in Section IV can be adopted this quarter. The immediate actions in Section VI require no budget approval, no technology investment, and no external consultant.

They require only what effective governance has always required: the willingness to ask the right questions, the discipline to demand complete answers, and the understanding that the board's responsibility does not end where management's technical expertise begins.

AI governance is not a technology problem waiting for a technology solution.

It is a governance problem waiting for board directors who understand their role in solving it.

Pawneet Abramowski | CEO & Founder, PARC Solutions

Publisher, The Risk Chair™ · A PARC Solutions Platform · May 2026

APPENDIX A

Acronym Glossary

Acronym	Definition
AI	Artificial Intelligence
AML	Anti-Money Laundering
API	Application Programming Interface
BSA	Bank Secrecy Act
CCO	Chief Compliance Officer
CEO	Chief Executive Officer
CFPB	Consumer Financial Protection Bureau
CGO	Chief Governance Officer
CIO	Chief Information Officer
CIP	Customer Identification Program
CISO	Chief Information Security Officer
CRM	Customer Relationship Management
CRO	Chief Risk Officer
CTO	Chief Technology Officer
EEOC	Equal Employment Opportunity Commission
ERP	Enterprise Resource Planning
EU	European Union
FDIC	Federal Deposit Insurance Corporation
FinCEN	Financial Crimes Enforcement Network
GDPR	General Data Protection Regulation
IRS	Internal Revenue Service
IT	Information Technology
KYC	Know Your Customer
LLC	Limited Liability Company
NASDAQ	National Association of Securities Dealers Automated Quotations
NFS	Network File System
OCC	Office of the Comptroller of the Currency
Q1	First Quarter
SEC	Securities and Exchange Commission