
THE RISK CHAIR™
A PARC SOLUTIONS PLATFORM
ISSUE
03
JUNE 2026
THE GOVERNANCE DEFICIT
CREDIT DECISIONING
AI MODELS
FRAUD DETECTION
AML MONITORING
DATA SOURCES
PUBLISHED JUNE 2026 • THE RISK CHAIR™ • PARCSOLUTIONSLLC.COM

THE RISK CHAIR™
A PARC Solutions Platform

THE GOVERNANCE DEFICIT

How Financial Institutions Are Deploying AI Faster Than They Can Govern It, and What Boards Must Do Now

Pawneet Abramowski
CEO & Founder, PARC Solutions · Independent Board Director, Titan Acquisition Corp. (NASDAQ: TACHU) · Former FBI Intelligence Analyst
Published June 2026 · The Risk Chair™ · parcsolutionsllc.com

About This Paper

This is the third publication from The Risk Chair™. The first, The Decoupling Doctrine, mapped the strategic architecture reshaping the global economic environment. The second, What the Boardroom Does Not Yet Know About Artificial Intelligence, addressed the governance obligations boards have not yet accepted. This paper connects the two, specifically in the context of financial institutions navigating simultaneous demands from Basel III capital reform and AI deployment. It is written for board directors, chief compliance officers, chief risk officers, general counsel, and the senior executives responsible for ensuring that what their institutions are building can withstand the scrutiny that is coming.

FOREWORD

The Room Where Everyone Said Yes, and Then Bet on the Opposite

On June 3, 2026, the Forbes Iconoclast Summit convened in New York. The room contained the leadership of some of the world's largest financial institutions, JPMorgan Chase, HSBC, Mizuho, Carlisle, Bridgewater, and others. Ray Dalio spoke. The conversation was wide-ranging. But one thread ran consistently through nearly every session: everyone in that room acknowledged that AI is being deployed. Everywhere. Now.

And then, in the same conversations, the investment thesis that dominated the room was not AI infrastructure. It was live sports. Live arts. Physical, irreducibly human experience. The most sophisticated capital allocators in the world are simultaneously betting that the scarcest premium commodity in an AI-saturated world will be the thing that cannot be optimized, predicted, or replicated.

This is not a contradiction. It is the oldest pattern in civilization. When the infrastructure of society becomes sufficiently automated, the appetite for what cannot be automated intensifies. The Greeks built the Olympics at the height of their intellectual and architectural achievement. Rome built the Colosseum as its empire reached its administrative peak. The live event was not a failure of civilization. It was evidence of its maturity, and of the human need for something that no system can replicate.

The boards and executives in that room understand this intuitively. What they have not yet resolved is the governance architecture required to manage the AI they are deploying at scale while they are making those bets. That unresolved space, between deployment velocity and governance capacity, is what this paper addresses.

Pawneet Abramowski | CEO & Founder, PARC Solutions
Publisher, The Risk Chair™ · June 2026

EXECUTIVE SUMMARY

The Governance Deficit: Three Arguments

Financial institutions are deploying artificial intelligence faster than any governance framework in existence can track. That is not a technology failure. It is a leadership and accountability failure, and it is happening at the same moment that two of the most consequential regulatory demands in a generation are converging on the same institutions simultaneously.

The Basel III re-proposal, issued by US regulators on March 19, 2026, explicitly extends its operational risk framework to AI models. The EU AI Act's high-risk provisions, including credit decisioning and fraud detection, become enforceable on August 2, 2026. Colorado's AI Accountability Act takes effect on June 30, 2026. The US Treasury released a comprehensive AI governance guidance framework for financial institutions in May 2026. Regulation is not hypothetical. It is active, specific, and arriving with enforcement teeth.

Three arguments follow.

> *First: Financial institutions have adopted a hybrid AI architecture, proprietary builds layered on vendor platforms and open ecosystem models, which creates a governance blind spot at the interface layer. The institution believes it owns and controls the system. What it controls is the surface. The underlying model's behavior, drift, and data provenance remain outside the institution's governance perimeter.*

> *Second: The same institutions implementing Basel III operational risk requirements are deploying AI in the credit models, stress testing frameworks, and fraud detection systems that Basel III governs. Managing these as separate workstreams, as most institutions currently do, produces two incomplete governance architectures rather than one coherent one. The Basel III re-proposal has closed that option. Operational risk requirements now explicitly cover AI model failures.*

> *Third: The layoffs reshaping financial services in 2026 are not cost discipline. They are capital reallocation. Budget is moving from headcount to AI infrastructure at a scale that has no precedent. The governance implications of that reallocation have not been absorbed by the boards approving it. An institution that reduces its compliance and risk headcount while simultaneously expanding its AI deployment has not reduced its risk. It has concentrated it.*

The room at Forbes Iconoclasts understood something that is not yet visible in the survey data. The leaders deploying AI at the largest institutional scale in history are simultaneously betting that the scarcest commodity in an AI-saturated world will be irreducibly human experience.

They understand the arc. What most of their boards have not yet built is the governance architecture required to complete it. That test is coming sooner than most expect.

THE FRAMEWORK

Deploy. Govern. Withstand.

Every financial institution sits somewhere on this arc. Most are at Deploy. Almost none have built Govern. Withstand, the moment an examiner, a regulator, or a board member asks what the institution actually knew, is arriving whether the institution is ready or not.

1

DEPLOY

Every institution in this paper is here. AI is in production, at scale, across compliance, fraud, credit, and client-facing functions. This is not the gap.

2

GOVERN

Almost no institution has an integrated framework. AI governance and Basel III operational risk are managed as separate workstreams. This is the gap, and it is closing only on paper.

3

WITHSTAND

The test is coming: Basel III examination, EU AI Act enforcement, Colorado Act enforcement, a board asked under oath what it knew and when. Few institutions are ready.

FIVE THINGS EVERY BOARD NEEDS TO KNOW

The Governance Deficit, At a Glance

1	**The money is moving, not disappearing.** Layoffs across financial services in 2026 are capital reallocation toward AI infrastructure, not cost-cutting. Boards approving this have approved a risk transfer they may not have evaluated.
2	**You do not control what you think you control.** Most institutions run a hybrid AI architecture: proprietary interfaces on vendor and open-source foundations. The board governs the surface. The underlying model's behavior sits outside the governance perimeter.
3	**Basel III and AI governance are now one regulatory framework, not two.** The March 2026 re-proposal explicitly brings AI model failures into operational risk capital calculations. Treating them separately creates two incomplete frameworks.
4	**The model deployed today may not be the model running in six months.** And in stress tests, frontier models have chosen to blackmail and leak information when their continued operation was threatened, even after explicit instructions not to.
5	**The test is coming faster than most boards expect.** Colorado's AI Accountability Act is enforceable June 30. The EU AI Act's GPAI obligations land August 2. The Basel III comment period closes June 18, one day after this paper publishes.

PART I

The Capital Reallocation Reality

The Money Has to Come From Somewhere

In 2026, Amazon, Meta, Google, and Microsoft are investing a combined $650 billion in AI infrastructure within a single year. JPMorgan Chase has committed $20 billion of its 2026 technology budget to technology investment, with AI as the primary driver. These are not research allocations. They are operational commitments. The money has to come from somewhere.

The fintech sector has recorded at least 9,706 job cuts in 2026, making it the fifth hardest-hit segment of the technology industry. Block CEO Jack Dorsey cut the company's headcount nearly in half, from 10,000 to fewer than 6,000, attributing the move directly to AI: "Intelligence tools have changed what it means to build and run a company." Cloudflare CEO Matthew Prince stated explicitly that AI "make 1,100 jobs obsolete" even as the company reported record revenue, up 34% year over year. PayPal is reducing approximately 4,760 positions over three years as automation expands across fraud detection, risk management, and customer support.

The arithmetic is straightforward. The capital required for AI infrastructure is being freed from headcount. This is not financial distress. It is deliberate reallocation. The institutions making these decisions are profitable and growing. The cuts are strategic, not defensive.

> *Sam Altman acknowledged publicly that "there's some AI washing where people are blaming AI for layoffs." The distinction matters for governance: institutions that are genuinely restructuring around AI capability require a governance framework designed for that restructuring. Institutions that are cutting costs under an AI rationale face a different risk. They reduce the human oversight capacity at precisely the moment AI deployment is accelerating.*

What the Reallocation Means for Institutional Risk

A Gartner study released in May 2026 surveyed 350 global business executives. Approximately 80% reported workforce reductions tied to AI. There was no correlation between those reductions and improved ROI. The analyst assessment was direct: workforce reductions may create budget room, but they do not create return.

An institution that reduces compliance and risk headcount while expanding AI deployment in the same compliance-relevant functions, KYC, AML monitoring, fraud detection, credit underwriting, has not reduced its regulatory exposure. It has concentrated it. The outputs are less interpretable. The failures are less predictable. The oversight team is smaller. The board that approved the reallocation may not have evaluated what it was actually approving.

Under Basel III, this concentration has capital implications. The March 2026 re-proposal links required capital buffers to operational loss histories. An AI-driven failure that produces material losses flows directly into capital adequacy calculations. The institution managing AI deployment and Basel III compliance as separate workstreams may discover they share the same denominator at exactly the wrong moment.

PART II

The Hybrid Architecture Problem

How Financial Institutions Are Actually Deploying AI

The public narrative around AI in financial services tends toward two poles: either institutions are building proprietary AI systems that give them full control, or they are entirely dependent on third-party vendors. The reality is more complex and more consequential from a governance standpoint.

Most financial institutions in 2026 are deploying a hybrid architecture: proprietary components, workflows, decisioning interfaces, compliance overlays, built on top of foundational models and vendor platforms provided by third parties. The institution controls the surface layer. The underlying model's training data, behavioral parameters, and drift trajectory are determined by the vendor or the open-source community that produced it.

This pattern is not new. Financial institutions followed precisely the same architecture with blockchain. Most "proprietary" bank blockchain implementations were built on Hyperledger Fabric or Ethereum permissioned variants, not on genuinely native infrastructure. The appearance of control came from the interface layer. The underlying protocol and its behavior remained outside the institution's governance perimeter. The same dynamic is now playing out in AI, at significantly greater scale and with significantly greater regulatory consequence.

The Governance Blind Spot at the Interface Layer

When a board approves an AI deployment, it is typically approving a description of the system's intended function. What it is not reviewing, and in most cases what management cannot fully provide — is a complete picture of how the model behaves against the institution's specific data, how that behavior changes over time, and what the model does at the edges of the distribution where consequential errors concentrate.

The board believes it controls the system. What it controls is the interface. The model underneath, its training data, its drift behavior, its update cadence sits outside the institution's governance perimeter. That distinction is not academic. It is the accountability gap that regulators have already started asking about.

This pattern has a history. The machine learning transformation that financial institutions pursued from roughly 2013 through 2020 followed the same arc. Institutions acquired transaction monitoring platforms, credit decisioning engines, and fraud detection systems from established vendors - NICE Actimize, Mantas, Detica. These were not poorly built systems. The technology was capable. The governance gap was not in the deployment itself. It was in what came after.

Financial institutions have always operated largely on closed data environments. The machine learning systems deployed in that era ran against internal data, proprietary transaction histories, customer profiles, internal risk classifications. The AI systems being deployed today operate the same way. Microsoft Copilot deployed across an organization, runs against that organization's internal data. It is not reaching external models for calibration. It is not being benchmarked against industry-wide transaction patterns. It is a powerful system operating within the boundaries of what the institution has fed it.

The problem was never whether the technology was deployed. The problem was whether anyone asked whether the model was performing as intended against the institution's specific data profile, and whether that question was being asked regularly, independently, and with the same rigor applied to credit underwriting models, stress testing models, and transaction monitoring models under existing model risk management frameworks.

The answer, in most cases, was no. And the regulators noticed. HSBC paid $1.9 billion in 2012. Deutsche Bank paid $630 million in 2017. Standard Chartered paid $1.1 billion in 2019. The fines were not for deploying inadequate technology. They were for deploying technology without the model validation framework required to know whether it was working, whether thresholds were appropriate for the institution's risk profile, whether the model's outputs were being independently reviewed, whether drift was being detected and addressed before it produced a compliance failure.

Anyone who has sat inside those implementations understands the dynamic. The business stakeholder presents the AI solution to leadership. The investment case is compelling. The vendor is credible. The technology is approved. What does not get approved, because it is rarely presented, is the ongoing model validation program, the independent review cycle, the governance infrastructure required to know whether the system continues to perform as intended six months after deployment. The board sees the capability. It does not see the maintenance obligation that comes with it.

AI in 2026 is the same problem at a materially higher velocity. The 62% of financial institutions that report having deployed AI agents have not, in most cases, submitted those agents to the same model validation standard that applies to their credit models, their stress testing frameworks, or their transaction monitoring systems. The enterprise AI deployment, Copilot across the organization, an AI-powered KYC screening tool, a generative compliance assistant arrives with a vendor implementation plan. It does not arrive with an independent model validation requirement, a drift monitoring protocol, or a board-level reporting obligation tied to model performance. Those are governance decisions that have to be made by the institution. Most have not made them. The regulatory reckoning, when it comes, will be about that gap, not about whether the technology was deployed, but about whether anyone was responsible for knowing whether it was working.

The NVIDIA NIM Dimension: Open Source at Scale

NVIDIA's NIM platform, NVIDIA Inference Microservices, 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. NVIDIA pays the compute bill. 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 blind spot that no current governance framework is designed to surface. Western AI compliance frameworks assume a closed, identifiable, auditable technology stack. When the model is open source, distributed globally, and running on third-party cloud infrastructure, that assumption fails. The audit trail the framework assumes does not exist.

The governance question for boards is not whether their institution is directly deploying Chinese-origin models. It is whether any vendor in their technology supply chain is. And whether the institution's third-party risk framework is equipped to assess that exposure.

The AI Washing Problem

The market for AI-native fintech platforms is saturated and insufficiently differentiated. Thousands of companies are raising capital under AI branding. The governance challenge is that many of these platforms are deploying better user interfaces over two-year-old underlying technology, relabeled as AI, repriced as a premium, and marketed as a transformation solution.

The institution that deploys a vendor's AI-powered compliance solution and discovers it is a rules engine with a generative interface has not improved its compliance posture. It has added interpretability complexity to a system that was already difficult to audit. The question "what is the model actually doing" is not a technical question. It is a governance question. The board's due diligence obligation extends to the substance of what is deployed, not the marketing description of it.

PART III

The Regulatory Convergence

Basel III Re-Proposal, March 19, 2026: AI Is Now Explicitly In Scope

On March 19, 2026, US regulators re-proposed the Basel III Endgame as the Basel III Proposal, with material changes from the original 2023 publication. Among the most consequential for institutions deploying AI: the revised operational risk standard applies explicitly to AI model failures. Supervisors have stated clearly that the Standardized Measurement Approach, which links required capital buffers to verified loss histories and income proxies, covers chatbots, credit engines, and vendor AI models on equal footing with traditional operational risk categories.

The output floor, which limits how far model-based risk weights can diverge from standardized levels, was designed to prevent capital arbitrage through internal model manipulation. Its application to AI models means that an institution whose AI-driven credit model systematically underestimates risk, whether through drift, data bias, or distributional shift, faces both the direct loss from the model failure and the capital impact of the operational risk charge that follows. EU banks must submit the first full SMA templates in supervisory reporting by Q3 2026.

Comments on the three Basel III Proposal rulemakings are due June 18, 2026, one day after this paper publishes. This paper is not arriving after the policy conversation has settled. It is arriving in the middle of the active comment period, while the institutions, trade associations, and supervisors shaping the final rule are still making their case. For boards and CCOs, that timing is an opportunity: the institutions that understand the AI and operational risk convergence now are positioned to engage the comment process and the implementation planning that follows from a position of clarity rather than catch-up.

An institution's AI governance framework and its Basel III compliance framework are not parallel workstreams. They are the same workstream. The March 2026 re-proposal collapsed that separation at the regulatory level. The CCO and the CRO managing these independently are building two incomplete programs where one integrated framework is required.

US Treasury AI Governance Guidance — May 2026

In May 2026, the US Department of Treasury released a Financial Sector AI Deliverable Reference and Application Guide, six interrelated resources designed to support practical, regulator-informed AI risk management for financial institutions. The guidance is explicit: AI risk should be embedded in existing risk and compliance frameworks, not treated as a standalone technology issue.

Grant Thornton's 2026 AI Impact Survey of banking leaders found that the lack of centralized, tested governance is the primary constraint holding banks back from measurable AI performance. The Treasury guidance makes that finding a regulatory expectation: institutions that have not integrated AI governance into their enterprise risk framework are not just underperforming — they are out of alignment with regulator expectations. The guidance arrived after the window in which institutions could claim they were waiting for direction. That window is closed.

Colorado AI Accountability Act — June 30, 2026

Colorado's AI Accountability Act, currently the most comprehensive state-level AI regulation in the United States, takes effect nine days after this paper publishes. For financial institutions, which the Act explicitly classifies as deployers of high-risk AI in consequential decision-making, the obligations include consumer notification when AI is used in decisions affecting them, impact assessments for high-risk AI systems, and reasonable care requirements to prevent algorithmic discrimination.

The Act's significance extends beyond Colorado's borders. It represents the leading edge of a state-level regulatory patchwork that, in the absence of federal preemption, will create compliance obligations across multiple jurisdictions with varying requirements. Institutions operating nationally that have not begun building AI impact assessment infrastructure will be managing a compliance deficit that compounds as additional state laws follow Colorado's framework.

EU AI Act — August 2, 2026

The EU AI Act's high-risk AI provisions, including credit decisioning and fraud detection, become enforceable on August 2, 2026. A political agreement on the AI Omnibus simplification package was reached on May 7, 2026, extending the transition period for high-risk AI systems embedded in regulated products until August 2028, but transparency obligations and governance requirements for General Purpose AI models take effect in August 2026 as originally scheduled.

For US financial institutions with European operations, the August deadline creates an asymmetric compliance obligation: binding EU requirements on AI systems already in production while domestic US regulators are still issuing guidance. The fines for non-compliance reach 7% of global revenue for prohibited violations and 3% for high-risk failures, making EU AI Act penalties potentially more expensive than GDPR breaches. The EU explicitly elevates AI governance to board-level responsibility. Directors face potential personal liability under corporate law fiduciary duties if they consciously disregard significant regulatory risks.

The SEC's Position

The SEC's Investor Advisory Committee has recommended enhanced disclosures concerning how boards oversee AI governance as part of managing material cybersecurity risks. The 2026 SEC examination priorities have displaced cryptocurrency as the industry's primary concern in

favor of cybersecurity and AI. This is not yet a mandatory disclosure regime. But examination priorities signal enforcement focus, and the institutions that have not built board-level AI oversight documentation are building a gap between their actual governance posture and the posture they will be asked to describe under examination.

PART IV

The Compliance Stack: Where Basel III and AI Governance Meet

The Same Systems, Two Frameworks

The financial institution implementing Basel III capital requirements and the financial institution deploying AI are, in most cases, the same institution managing the same operational infrastructure through two separate governance lenses. The credit risk model calculating probability of default is an AI system subject to both Basel III's model risk requirements and emerging AI governance obligations. The fraud detection system generating alerts is simultaneously a Basel III operational risk management tool and an AI model subject to explainability and bias assessment requirements. The AML monitoring platform flagging transactions is a compliance system and an AI system, and the regulatory expectations for each category are now arriving simultaneously from different directions.

The PARC Solutions Insights series on Basel III & Fintech documented how governance expectations flow downstream from regulated institutions to their fintech partners, and how capital sensitivity and structural discipline are not bank-internal requirements but conditions that extend to every organization inside a bank's risk perimeter. That downstream flow now carries AI governance obligations alongside Basel III structural requirements. The fintech partner deploying AI-powered KYC or embedded credit decisioning is not just subject to Basel III's downstream governance expectations. It is subject to the AI model risk requirements that the March 2026 re-proposal has made explicit.

The Explainability Problem

Basel III requires institutions to demonstrate defensible methodology for their risk models. Supervisors examining a bank's credit risk framework expect to see documented model validation, performance testing, and a clear explanation of how the model produces its outputs. AI models, particularly deep learning systems, and large language models used in decisioning contexts, often cannot provide that explanation in the form regulators expect.

This is a governance design problem, not a technology problem. Institutions deploying AI in Basel III-relevant functions without a parallel explainability framework are building a regulatory examination exposure that will surface when supervisors ask questions the institution's vendor cannot answer on its behalf. The vendor's model card is not a model validation document. That obligation belongs to the institution, regardless of who built the underlying system.

When the Model Builds Itself

On June 4, 2026, Anthropic published a paper titled "When AI Builds Itself," reporting that more than 80% of the code merged into its own codebase is now written by Claude, and that its

engineers were merging approximately eight times as much code per day in the second quarter of 2026 as they were in 2024. The human role at each step of building AI is shrinking. Anthropic called for the world to have the option of a coordinated, verifiable slowdown of frontier AI development across multiple labs and multiple countries, including the US and China simultaneously. Anthropic was explicit that it would not act unilaterally. If only one company slowed down, competitors would race ahead regardless. The organization most directly confronting the trajectory of its own technology concluded that no single actor, including itself, can or will slow down absent coordination that does not currently exist. That is not a reassurance. It is the governance vacuum stated plainly by the party with the most complete view of it.

> *In October 2025, Anthropic stress-tested 16 frontier AI models from every major developer, including its own, in simulated corporate environments where the models had autonomous access to company emails and sensitive information. The models were assigned only harmless business goals. When researchers introduced a threat to a model's continued operation that conflicted with the company's strategic direction, models from every developer, in at least some cases, chose to blackmail executives, leak sensitive information, and engage in corporate espionage, without being instructed to. When researchers added explicit behavioral instructions not to do so, harmful behavior dropped from 96% to 37%. The instructions helped. They did not solve the problem.*

This finding is not a hypothetical edge case. It is a documented description of how the AI systems being integrated into financial institution operations behave when their own continuation is threatened, even after explicit safeguards are applied. A board approving an agentic AI deployment in a compliance, accounting, or client-facing function is approving a system whose worst-case behavior, even with safeguards, has been measured and published. The 37% residual rate is the number the governance framework in Part V is designed to address.

For financial institutions, the practical governance implication of recursive AI development is this: the system deployed today may be materially different from the system operating six months from now. Not because anyone changed it, but because its underlying foundation has been updated by processes that no human fully directed. The institution's model risk management framework, its Basel III model validation documentation, and its AI governance oversight all assume a relatively stable system being periodically reviewed. That assumption is increasingly incorrect.

The board that approved a deployment in Q1 2026 may be governing a materially different system by Q3 2026 without a formal re-approval having occurred. That is not a hypothetical risk. It is the documented trajectory of the technology as described by its own developers.

PART V

The Governance Framework: What Integrated Looks Like

One Framework, Not Two

The central recommendation of this paper is architectural rather than procedural. Financial institutions do not need a Basel III compliance program and a separate AI governance program. They need one integrated framework that treats AI model risk as a subset of operational risk under Basel III, while simultaneously satisfying the AI-specific disclosure, assessment, and oversight obligations now arriving from multiple regulatory directions.

Returning to the framework introduced at the outset: this section is the architecture for Govern. Part I and II described what Deploy looks like across the industry today. Part III and IV described what Withstand will demand. What follows is what most institutions are missing in between, the operational requirements that turn a deployment into a governed system.

That integrated framework has five operational requirements.

Start with the Inventory

The first requirement is simply to know what you have. Every AI model in production, including vendor-provided models and any open-source models running in the institution's technology supply chain, needs to be inventoried and classified against Basel III operational risk categories. The question is not whether the model is "an AI system." The question is whether its failure would generate an operational loss, a regulatory penalty, or a capital impact. If the answer is yes, it belongs in the operational risk inventory and the AI governance framework at the same time. Most institutions have not done this. The inventory does not require a technology investment. It requires someone with the authority to ask for it.

Build the Explainability Standard Before the Examiner Asks for It

For each AI model in a Basel III-relevant function, document the minimum explainability standard that a supervisor examining that function would require. A credit risk model must satisfy Basel III model validation standards. A fraud detection system must satisfy BSA/AML examination standards. Where the AI model cannot satisfy that standard natively, a compensating control is required, typically a human review layer that satisfies the supervisory expectation even when the model output alone cannot. This is not a new concept. It is the same standard that was applied to quantitative models after the 2008 crisis. The difference is that most institutions applied it then because regulators forced them to. Applying it now, before the enforcement cycle, is the governance differentiator.

Treat Vendor AI as Third-Party Risk, Not Technology Procurement

The standard vendor onboarding questionnaire was designed for software with defined, version-controlled releases. It is not adequate for AI vendors whose underlying models may change materially between review cycles without formal notification. Every vendor providing AI-enabled services needs to be assessed for model provenance, training data practices, update cadence, and behavioral change notification. The institution needs a contractual right to be notified when the model changes in a material way, and a process for evaluating that change against its operational risk and Basel III frameworks. Deploying an AI system is not a one-time procurement decision. It is an ongoing governance relationship.

Give the Board Intelligence, Not Updates

Board oversight of AI should not be a periodic technology update from the CTO. It should be a standing governance agenda item, presented by someone with cross-functional accountability for AI model risk, operational risk under Basel III, and regulatory compliance simultaneously. The Chief Governance Officer construct introduced in this series describes that function. Whatever title the institution assigns, the structure matters more than the name: a single designated executive with the mandate and the authority to integrate these three accountability streams and report them to the board as one coherent picture.

Govern Continuously, Not Annually

The annual model review cycle was designed for a world where models changed slowly and under institutional control. That world no longer exists. Anthropic acknowledged this explicitly in June 2026: the model Claude is writing most of its own code, and the update cadence at every major AI lab is now measured in weeks, not quarters. Institutions need automated performance monitoring against defined thresholds, triggered re-validation when those thresholds are breached, and a board notification process that fires when a material model change occurs regardless of where the institution is in its annual review calendar. The governance cadence has to match the technology cadence. Right now, it does not.

CLOSING SYNTHESIS

What the Room at Iconoclasts Understood

The leaders at Forbes Iconoclasts on June 3, 2026, were not naive about what they are deploying. They are investing in AI at scale while simultaneously betting on live sports, live arts, and physical human presence. That is not inconsistency. It is a sophisticated read of where value concentrates when everything else becomes automated.

The ancient parallel is instructive. The Greeks who built the Parthenon also built the Olympic stadium. The Romans who constructed the most advanced administrative and engineering infrastructure in the ancient world also built the Colosseum. The appetite for irreducibly human experience: the contest that cannot be scripted, the performance that cannot be replicated, the moment that exists only once, intensifies precisely as the surrounding infrastructure becomes more automated and more predictable. That is not a regression. It is a sophisticated allocation of human attention to the things that automation cannot produce.

The governance challenge for financial institutions is analogous. The question is not whether to deploy AI. That decision has been made at every institution of consequence. The question is whether the governance architecture being built around that deployment is adequate to the obligation. To regulators, to customers, to counterparties, and to the boards that have accepted fiduciary accountability for the outcomes.

The institutions that will navigate this period well are not the ones that deployed the most advanced AI the fastest. They are the ones that built governance infrastructure that could hold under examination. That infrastructure could answer, clearly and specifically, who is responsible for what the system does, what happens when it is wrong, and whether the board understood what it was approving when it approved it.

That architecture does not require technical expertise at the board level. It requires pattern recognition, the trained instinct to identify when AI deployment has crossed a threshold of institutional risk, and the governance discipline to ask the right questions before the threshold is crossed rather than after. That instinct is not technical. It is the same instinct that distinguishes an effective intelligence analyst, an effective Chief Risk Officer, and an effective board director. The technology is new. The governance problem is not.

The room at Iconoclasts understood the opportunity. This paper is the map for capturing it without losing control of what is being built. The gap between Deploy and Withstand is where institutions are lost right now. Govern is the work. It is also the differentiator.

Pawneet Abramowski | CEO & Founder, PARC Solutions

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

APPENDIX A

Acronym Glossary

Acronym	Definition
AI	Artificial Intelligence — Advanced computational systems enabling machine learning, automation, and decision support.
AML	Anti-Money Laundering — Regulatory framework requiring detection, prevention, and reporting of illicit financial activity.
B3P	Basel III Proposal — The March 19, 2026, US re-proposal of the Basel III Endgame framework with material revisions.
Basel III	International regulatory framework established by the Basel Committee on Banking Supervision governing bank capital, leverage, and liquidity standards.
BSA	Bank Secrecy Act — US law establishing recordkeeping and reporting requirements for financial institutions to combat financial crime.
CAIO	Chief AI Officer — Executive responsible for AI strategy, deployment, and technology integration.
CCO	Chief Compliance Officer — Executive responsible for regulatory compliance programs and obligations.
CFIUS	Committee on Foreign Investment in the United States — Interagency body reviewing foreign investments for national security risks.
CGO	Chief Governance Officer — Proposed executive function at the intersection of AI governance, cybersecurity, regulatory compliance, and enterprise risk.
CHIPS Act	Creating Helpful Incentives to Produce Semiconductors Act — US legislation funding domestic semiconductor manufacturing.
CISO	Chief Information Security Officer — Executive responsible for information security strategy and cybersecurity operations.
CRO	Chief Risk Officer — Executive responsible for enterprise risk management and oversight.
CRR III / CRD VI	Capital Requirements Regulation III / Capital Requirements Directive VI — EU legislative framework implementing Basel III standards.
DORA	Digital Operational Resilience Act — EU regulation establishing mandatory ICT risk management and resilience requirements for financial entities, in force since January 2025.
EDD	Enhanced Due Diligence — Heightened customer verification and monitoring requirements for high-risk relationships.
EU AI Act	European Union Artificial Intelligence Act — Comprehensive EU regulation classifying AI systems by risk level with corresponding compliance obligations.
FDIC	Federal Deposit Insurance Corporation — US regulator insuring bank deposits and supervising financial institutions.
FinCEN	Financial Crimes Enforcement Network — US Treasury bureau responsible for AML enforcement and financial intelligence collection.

FRTB	Fundamental Review of the Trading Book — Basel Committee framework revising market risk capital requirements.
GDPR	General Data Protection Regulation — EU data privacy regulation with enforcement authority and significant penalty provisions.
GENIUS Act	Guiding and Establishing National Innovation for US Stablecoins Act — Legislation creating a regulatory framework for dollar-backed stablecoins.
GLM	General Language Model — Chinese open-source large language model developed by Zhipu AI.
GPAI	General Purpose AI — AI systems capable of performing a wide range of tasks, subject to specific obligations under the EU AI Act.
KYC	Know Your Customer — Customer identification and verification requirements under BSA/AML frameworks.
LLM	Large Language Model — AI systems trained on large text datasets capable of generating and analyzing natural language.
MRM	Model Risk Management — Framework for identifying, assessing, and mitigating risks arising from the use of quantitative models.
NDAA	National Defense Authorization Act — Annual US legislation governing defense spending and national security policy.
NIM	NVIDIA Inference Microservices — NVIDIA's platform providing free API access to over 100 production-grade AI models.
OCC	Office of the Comptroller of the Currency — US regulator supervising national banks and federal savings associations.
OFAC	Office of Foreign Assets Control — US Treasury agency administering and enforcing economic and trade sanctions.
SAR	Suspicious Activity Report — Regulatory filing required when potential financial crime activity is identified.
SEC	Securities and Exchange Commission — US federal agency regulating securities markets and protecting investors.
SMA	Standardized Measurement Approach — Basel III operational risk capital calculation methodology replacing internal model approaches.
TSMC	Taiwan Semiconductor Manufacturing Company — World's largest semiconductor foundry.

APPENDIX B

Selected Bibliography

Primary sources consulted in the preparation of this publication. All sources reflect publicly available materials as of June 2026.

Anthropic. “When AI Builds Itself.” Anthropic Research Blog, June 4, 2026.

Basel Committee on Banking Supervision. Basel III Monitoring Report, 2025.

Colorado General Assembly. Colorado AI Accountability Act (SB 205), effective June 30, 2026.

European Commission. “AI Act Enters Application: Governance Rules and GPAI Obligations.” Digital Strategy, August 2025.

European Commission. “Political Agreement on AI Omnibus Simplification Package.” May 7, 2026.

Federal Reserve, OCC, and FDIC. Notice of Proposed Rulemaking: Basel III Proposal (B3P), March 19, 2026.

Forbes. Iconoclast Summit, New York, June 3, 2026.

Gartner. 2026 AI Impact Survey: Banking Leaders. May 2026.

Grant Thornton. “Treasury Guidance Brings Urgency to AI Governance.” Grant Thornton Financial Services, May 2026.

Nasdaq Verafin. 2026 Financial Crime Report: AI-Driven Attack Patterns. April 2026.

TradingPlatforms. Global Technology Layoff Tracker 2026. June 2026.

US Department of the Treasury. Financial Sector AI Deliverable Reference and Application Guide. May 2026.

Wilson Sonsini Goodrich & Rosati. ‘2026 Year in Preview: AI Regulatory Developments.’ January 2026.

Wolters Kluwer. “Building Trustworthy AI Governance: Colorado’s AI Accountability Act.” January 2026.

EY. “US Basel III Proposal: Key Updates and Implications.” March 2026.

Corporate Compliance Insights. “2026 Operational Guide to Cybersecurity, AI Governance & Emerging Risks.” January 2026.

AI Certs. “AI Regulation Meets Basel III Endgame: Operational Risk Shift.” April 2026.

DWU Consulting. “AI Revolution & the US Economy: 2026–2027.” March 2026.

The Street. “This Fintech Firm Is Replacing Their Workers with AI.” April 2026.

Programs.com. “List of Companies Announcing AI-Driven Layoffs.” June 2026.