FINOS Fintech Open Source Foundation

THE LINUX FOUNDATION | Research

AI and Open Source in Financial Services

Mitigating Machine-Scale Cyber Risks, Governing Agentic Workflows, and Co-Engineering the Next Wave of Industry Standards

September 2026

Hilary Carter, *Linux Foundation*
Tosha Ellison, *FINOS*

AI and Open Source in Financial Services

Traditional risk management frameworks are structurally **unequipped for the rapid velocity of autonomous agentic deployments.**

Understanding the **structural interplay between AI and open source components** is essential for making deliberate, risk-aware architectural decisions.

Financial leaders must **move beyond unstructured experimentation,** aligning model selection and token spend with specific business functions, priorities, and value.

Because financial firms rely on similar software foundations and face the same regulators, there is **no competitive advantage to independently developing AI compliance frame**works

AI cost governance, or **"tokenomics", has quickly become a boardroom priority,** requiring a common understanding of how token consumption maps to true value.

Pre-competitive standards are critical to measure the efficiency of energy converted to intelligence and build unified frameworks to protect enterprise margins.

Within a 12-month span, independent model execution capacity advanced from brief, human-steered tasks to **long-running, multi-day, complex, and unassisted workflows.**

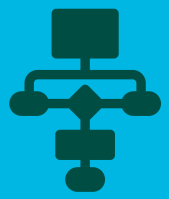

Financial industry participants must prioritize standardization by **aggressively voicing unique compliance, attestation, and data-provenance issues** within open source working groups.

Sandboxed, open models allow banks to use sovereign AI solutions to automate bug fixes, with **two-thirds to three-quarters of AI-generated patches yielding production-grade fixes.**

The financial sector faces a **high-stakes, 18- to 24-month window to systematically patch flaws** before automation renders them exploitable.

Shifting frontier model capabilities make it **difficult to articulate risk profiles or formulate precise multi-year capital allocation estimates to boards and risk committees.**

The Open Source Enterprise Resiliency Alliance (OSERA) enables firms to **jointly fund, prioritize, and consume secure versions of critical open source** at significantly lower cost.

Contents

Executive summary

The AI in Finance Leadership Summit, hosted by FINOS in June 2026, convened executives, regulators, and security experts to navigate the operational, cost, and security challenges of deploying AI at enterprise scale in the regulation-heavy, risk-averse, and historically conservative financial services sector.

The Summit began with a plenary discussion centered on the recognition that complex system orchestration must be engineered openly. As AI advances from sandbox to production, the financial services sector must move beyond unstructured experimentation. Session leaders highlighted how the current landscape is defined by the need for digital sovereignty, the harmonization of fragmented global regulations, and the strategic management of tokenomics, defined as the cost of converting energy into intelligence. Establishing pre-competitive standards is essential to protect enterprise margins while fostering innovation.

The following breakout sessions centered on three critical themes. The first session on tokenomics discussed how organizations must transition from spend-to-learn mindsets to disciplined budget management. This involves linking token consumption to measurable business value, using tiered budgets, and evolving metrics from simple volume to outcome-oriented indicators. The second session on AI security focused on the industry's critical 18-to-24-month window to patch security vulnerabilities before automation renders them exploitable. Participants felt that leaders must move from slow, human-speed defenses to machine-speed automated patching, mandating local model execution for data sovereignty and implementing verifiable, policy-aware guardrails. The third session on agentic AI use cases in financial services found that the sector faces common structural barriers and a near limitless supply of use cases. Collaboration is required to co-engineer standardized integration layers, such as the Model Context Protocol, to securely connect autonomous agents to legacy systems without expanding risk perimeters.

To galvanize industry resilience, the summit announced the Open Source Enterprise Resiliency Alliance (OSERA), a centralized effort to pool resources for vulnerability remediation, reducing reliance on expensive, redundant vendor contracts. By collaborating on security, cost standards, and infrastructure, financial participants can establish a golden path platform for financial AI.

By way of operationalizing these insights, the industry should prioritize five foundational actions. First, it must centralize pre-competitive safety and engineering work into neutral forums like FINOS to avoid siloed, proprietary solutions. Second, the industry needs to codify provenance by establishing standardized naming and metadata for vulnerability responses to ensure automated pipelines can programmatically validate patches. Third, firms should systematically extract legacy domain knowledge into structured, machine-readable specifications to guide agentic behaviors. Fourth, leaders should deploy natural language interfaces over version control systems to align technical execution with business strategy and bridge vocabulary gaps. Finally, teams must automate risk telemetry, transforming non-deterministic AI behavior into structured, verifiable audit trails for board-level reporting.

Introduction

The opening months of 2026 have witnessed an extraordinary acceleration in the capability and deployment of artificial intelligence, presenting the financial services industry (FSI) with an entirely new operational paradigm. This rapid transformation is characterized by a powerful phenomenon: the convergence of cutting-edge AI systems with community-driven open source standards. As global financial institutions grapple with innovation and technological shifts while navigating regulatory constraints and cost management, the open source integration frameworks and collaborative protocols form the critical bridge enabling financial services industry participants to function safely, transparently, and at scale.

AI innovation cannot happen in a vacuum. The stringent regulatory environment defining the financial sector serves as a vital, centralizing force for collaboration. Rather than acting merely as a compliance hurdle, the regulatory imperative mandates that competing institutions align on pre-competitive standards to manage shared liabilities, establish auditability, while ensuring the stability of the technologies that underpin the global financial system.

Earlier this year, the Linux Foundation hosted the AI Executive Forum in February, convening technology leaders for a cross-industry, cross-domain discussion on AI-related opportunities and challenges, including the impact of AI on regulated industries. These insights were captured in the April 2026 research report, **Open Source and the Future of AI**, including the observation that traditional risk management frameworks are entirely unequipped for autonomous agentic behaviors. Several initial challenges that the financial sector must urgently address included the need for documentation of human processes, executive education, continual human oversight to satisfy regulatory requirements, and standardization around tooling that would provide “decision-evidence” classifications and immutable, auditable trails, among other benefits.

Ultimately, the February forum was a catalyst for convening a deep-dive, financial services industry-specific leadership discussion on numerous AI-related opportunities and challenges. Recognizing that AI innovation requires regulators and business leaders to be joint collaborators, the **Fintech Open Source Foundation (FINOS)** hosted the AI in Finance Leadership Summit on the afternoon of June 24, 2026, in London. This invitation-only, half-day event brought together approximately 50 technology executives and AI pioneers from leading financial institutions, alongside prominent voices from the open source, regulatory, and security communities.

The Summit fostered candid reflections on the practical mechanics of deploying AI safely and cost effectively at enterprise scale, and explored new pathways for strategic, industry-wide collaboration through the open source communities. The format was structured into plenary sessions featuring leaders from the financial services and open source ecosystem, and thematic roundtable discussions focused on the emerging concept of tokenomics and AI cost management, cybersecurity challenges, and agentic AI use cases in the financial services industry. Through structured peer-to-peer dialogue, the Summit captured friction points, industry-specific milestones, and created new networks to facilitate future innovation.

Individuals named in this report have waived their right to the **Chatham House Rule** which governed the Summit's discussion. This report synthesizes the shared, otherwise anonymized insights into the financial services industry's AI-related challenges and pre-competitive collaboration opportunities, serving as a point for further advancements in open source software, standards, frameworks, and tools, through FINOS and the wider Linux Foundation ecosystem.

Strategic paradigms and infrastructure realities

The Summit's plenary session brought together authoritative leaders from across the open source, security, regulatory, and financial services ecosystems. Presentations established a macro-level baseline for the afternoon, illustrating the rapid shifts as AI moves from sandbox to production environments, and the industry-specific bottlenecks ranging from model capability to the practical engineering mechanics of cost management, data connectivity, supply chain security, and cooperative governance.

Industry foundations: Catalyzing open source collaboration in financial services

The opening from Gabriele Columbro, Executive Director of FINOS, mapped the structural evolution of financial technology, noting that a historically conservative industry has reached a foregone conclusion that complex system orchestration must be engineered in the open. Driven by a steady, multi-year effort by a growing number of participating organizations, FINOS plays a critical role in delivering industry-specific and tangible business value. The Foundation's work, anchored by the three pillars of profitability, resilience, and innovation, illustrates the extent to which open source software, open standards, and open collaboration have become established and worthwhile corporate investments across both sell-side and buy-side firms.

Columbro pointed to foundational friction points and strategic milestones defining the current financial services landscape:

- **Digital sovereignty:** Digital and data sovereignty broadly, and sovereign AI specifically, have emerged as a paramount boardroom concern in the face of global regulation. The core components across the modern AI stack—from distributed model training and inference to agentic workflows—are being constructed under the neutral stewardship of community-governed, open source protocols.
- **Shift-left tech risk collaboration:** FSI participants are increasingly leveraging open source frameworks to collaboratively mitigate tech risk, exchange data, mutualize compliance, and accelerate cloud and supply chain security.
- **The imperative of an open source AI strategy:** In an era where open source is eating AI, leadership must maintain an active, sophisticated open source strategy. Executives must fully comprehend the structural implications of AI interacting with open source components to make deliberate, risk-aware architectural choices, though organizations do not necessarily need to be biased toward utilizing open source across every system.
- **Harmonizing fragmented regulations:** Rather than inventing competing compliance standards, the community's primary objective is to harmonize fragmented global regulations into a transparent, executable open source AI governance framework. By mapping complex high-level legal mandates directly down to code-level controls and verifiable design patterns (the example of secure reference architectures for automated loan approvals was shared), financial institutions can safely deploy multi-agent pipelines with real-time observability of risks, controls, and token consumption.
- **Prioritization of strategic investment and use cases:** While community innovation is vast, financial leaders must

deliberately move past unstructured experimentation to prioritize where to allocate dedicated capital and resources. Through dedicated governance bodies like the newly introduced AI governing board, the industry must collaboratively uncover and define which specific FSI use cases should serve as the urgent focus for multi-stakeholder pilot testing. This targeted direction is critical to establishing clear, viable adoption pathways for the modern influx of finance-specific AI innovation.

"We have seen financial services, a traditionally conservative industry, collaborate more and more in the open—shifting left that collaboration to prove actual business value across profitability, resilience, and innovation."

– Gabriele Columbro, Executive Director, FINOS

Regulatory oversight

Mike Hsu, a prominent former regulatory official, joined the opening session to address the immense compliance pressure currently facing global watchdogs. While regulators face the challenge of streamlining technological change, they remain structurally unprepared for the velocity of autonomous AI deployments.

This operational gap creates a powerful opportunity for open source methodologies to bridge the private and public sectors. The good news is that an increasing number of leaders within regulatory agencies recognize the promise and the potential of what can be achieved when we work together as a community. At present, the goal is to deliberately connect the different parties in the near term, where the benefit to regulators is a more effective stewardship of the financial system in the AI era.

Tokenomics: Balancing costs against business value

The third plenary keynote session led by J.R. Storment, Executive Director of the Tokenomics Foundation, introduced the emerging discipline of Tokenomics, a critical framework for measuring, forecasting, and governing the volatile cost structures of enterprise AI. Between late 2025 and early 2026, an unprecedented explosion of frontier model releases equipped with massive context windows led organizations to operate under the illusion of unlimited computing resources. Driven by intense market pressure to capitalize on the AI business imperative, global chief executive officers authorized expansive testing and subscription models, incentivizing the wrong behavior with token leaderboards, only to find themselves erasing their unlimited token allocations within a matter of months.

This rapid depletion of capital triggered an acute industry-wide crisis, colloquially recognized as the great token panic of 2026. Today, financial leaders are actively struggling to forecast vertical consumption lines and are finding that unchecked token usage introduces immense, volatile entropy into corporate product goals. As the era of unconstrained vendor subsidies abruptly closes, technology executives are forced to pivot away from token leaderboard metrics to confront the stark fiscal reality of treating their local and cloud computing infrastructure as highly finite, expensive token factories.

This structural strain on corporate budgets is accelerated by the rapid transition from human-driven prompts to autonomous agentic architectures. The unique operational behavior of the agentic era introduces severe capital waste through several distinct factors. First, unlike humans who use AI linearly and slowly, agents interact with large language models in a highly distinct, relentless fashion. They run aggressive loops when they hit errors, trying again and again without pausing. Second, in order to execute complex tasks, a primary agent will routinely spin up multiple sub-agents, handling continuous retries, corrections, and verifications behind the scenes that compound token usage.

Agentic behavior has had a compounding effect in breaking traditional tracking metrics, forcing financial leaders to count usage in quadrillions and soon quintillions, and introducing a new discipline to better manage it. Part of that management necessitates that financial institutions recognize that token pricing represents only the baseline unit of machine cognition and pricing. True enterprise cost containment requires tracking the entire underlying infrastructure stack, including key-value (KV) caching fluctuations, retrieval-augmented generation (RAG) routing, non-deterministic evaluation telemetry, and high-density vector databases.

The need to establish pre-competitive standards to measure the literal efficiency of energy converted to intelligence becomes abundantly clear for the financial services industry, and an opportunity presents itself where banks can successfully transition away from mutually opaque, runaway cost paradigms to build clearer, value-aligned pricing abstractions that protect enterprise margins.

To meet this need, the **Tokenomics Foundation** is forming as a dedicated, vendor-neutral home to collaboratively solve these pressing, unanswered financial questions. Recognizing that individual organizations cannot efficiently map or forecast these vertical spend curves in isolation, the initiative establishes a critical open source space where major frontier model providers, cloud hyperscalers, and the world's largest token-consuming financial institutions can build unified frameworks. Operating alongside the **FinOps Foundation** to extend the discipline of technology value optimization into the era of token-based AI, this dedicated community is tasked

with engineering greenfield specifications, standardization metrics, and architectural routing guidance. One of the foundation's first initiatives is **Tokenomicon**, a horizontal conference focused on cross-foundational technical and AI value. In bringing together cross-industry subject matter expertise to tackle issues like workload classification, token-to-spend drift, and hardware optimization, the Tokenomics Foundation will be an important mechanism for financial services industry participants to reclaim control over their AI supply chains and drive measurable business value.

"Simply put, tokenomics is about converting energy and capital in the data center into cognitive intelligence, and then optimizing that intelligence to drive true business value."

– J.R. Storment, Executive Director, Tokenomics Foundation

Standards and connectivity through an open agentic stack

The fourth plenary session addressed the shifting technical ceiling of autonomous software deployments, asserting that raw model intelligence is no longer the only bottleneck for enterprise adoption. Instead, the modern operational frontier is defined by connectivity—the specific engineering patterns, integration layers, and protocols required to grant autonomous agents secure access to internal systems without expanding an institution's risk perimeter. The velocity of this ecosystem is underscored by an exponential expansion in task horizon capabilities. Within a mere twelve-month span, the operational capacity for independent model execution has advanced from brief, human-steered tasks lasting an hour into long-running, multi-day workflows capable of navigating complex networks completely unassisted.

This rapid agentic expansion materialized first within software engineering due to code's inherent properties, such as structured compilers, type systems, and deterministic environments that allow an agent to self-correct in real time, and accelerated by the availability of open source code bases. However, translating these capabilities to general knowledge work, or finance-specific functions—such as accounting, auditing, or algorithmic trading—presents challenges. General knowledge work relies heavily on fragmented internal systems, strict data-sharing contractual obligations, and specialized organizational insights and domain knowledge that are frequently undocumented and isolated within legacy silos.

To bridge these disparate environments and prevent fragmentation, the open agent stack relies on a foundational layer of community-governed protocols. At the core of this toolkit are standardized specifications like semantic organizational skills and the Model Context Protocol (MCP), which provide a rich

semantic layer for AI. By building an integration layer once to an open specification, financial institutions can securely connect LLMs to their underlying databases and SaaS applications across any commercial cloud platform or frontier model provider without suffering from severe vendor lock-in. Where manual integrations are difficult, developers are experimenting with direct computer-use protocols that allow agents to interact with legacy applications and mainframes by reading screens and executing commands exactly like a human operator.

While these initial open integration standards have successfully established themselves, technology leaders face an era of significant short-term engineering churn across several critical, unstandardized domains:

- **Long-running task execution and management:** The ecosystem is early in providing an open protocol to handle long-running workflows. Kicking off an autonomous task that takes days to complete still results in brittle execution paths and significant operational churn.
- **Dynamic agentic discovery:** There are currently no standardized web-scale registries or native protocols allowing an internal corporate agent to dynamically locate, authenticate, and securely collaborate with an external, third-party agent across the open web.
- **Authorization and identity vacuum:** Traditional identity-specific security standards, such as human-centric OAuth frameworks, are urgently needed for agents. While human operators are granted broad, badge-based access backed by employment contracts, proving that an agent is legitimately acting on behalf of a human in an industry that cares deeply about provenance, authorization, and access, is one of the largest connectivity problems today.
- **Policy management:** The financial services industry lacks a uniform standard to manage policy frameworks or enforce the transaction-level down-scoping of access privileges when a human delegates authority to an agentic system.

Ultimately, rebuilding these access control and governance guardrails from scratch requires direct, active participation from regulated entities. Because the standardization process is inherently outpaced by commercial AI development, financial services industry participants must voice their unique compliance, attestation, and data-provenance issues within open source working groups. By actively engaging in open forums and specialized financial services interest groups within the **Agentic AI Foundation** and **MCP** communities, the industry can co-engineer the open protocols necessary to maintain a common standard, neutralizing accidental vendor monopolies and ensuring long-term systemic control.

"The intelligence part is really not the ceiling anymore. It's domain knowledge."

Infrastructure security: Hardening the OS

The fifth plenary session shifted focus to the absolute core of enterprise computing infrastructure, exploring the severe cyber threats redefining modern operating systems and kernel development. As artificial intelligence advances at an exponential pace, the democratization of sophisticated LLM pipelines has fundamentally altered the threat landscape. Malicious actors have successfully weaponized these tools, transitioning away from rudimentary, single-vector attacks to deploy automated, machine-scale exploitation frameworks. By orchestrating advanced AI models to continuously crawl public code bases, bad actors can instantly scan thousands of software lines and identify where disparate, low-severity bugs can be chained together into complex zero-day exploits before maintainers can detect or disclose them. This machine-

driven onslaught means that any vulnerability at the operating system or kernel level puts every single layer of the downstream software stack at immediate risk.

To survive this shift in adversary capability, the open source development community is aggressively fighting fire with fire by deploying these exact pattern-matching engines defensively. Because large language models are exceptionally proficient at identifying repeating structural anomalies, security maintainers are leveraging local AI tools to ingest past patch histories, isolate where similar bugs remain unaddressed, and automatically generate working code remediations. This automated engineering grind has proven highly effective, with two-thirds to three-quarters of AI-generated patches successfully yielding accurate, production-grade bug fixes. Open source projects are actively churning through hundreds of historical flaws, achieving a level of code security and optimization that traditional software verification methods could never execute at scale.

For financial services industry participants navigating strict regulatory compliance mandates, this automated security paradigm introduces an absolute operational requirement for local model execution. Security researchers have repeatedly demonstrated that commercial, cloud-hosted LLMs routinely cache, share, and leak sensitive prompt text, resulting in public exposure when engineers upload internal code logs or proprietary bug reports to third-party providers. Consequently, financial enterprises must strictly mandate that their developers execute all security testing and automated patch generation locally on isolated, private desktop or server hardware. By running sandboxed, open models, banks can fully harness the pattern-matching power of AI to clean their code bases while maintaining total data sovereignty and protecting high-stakes intellectual property.

"Companies did not invest in security for many layers of the stack, and that was a business decision. They would deal with security later. Today is later. Fix the bugs you find today. Don't worry about triaging."

Securing application supply chains and dependencies at scale

The final plenary presentation focused on AI-amplified security threats at the application layer, focusing on the systemic vulnerabilities embedded within public package registries. Given that financial services applications are built upon dense tree networks of open source dependencies, ensuring the secure management of these third-party packages is a vital operational priority. At the same time, the global public infrastructure designed to track and score these vulnerabilities is under unprecedented structural strain. Spikes in automated bug filings have severely backlogged central public databases, combined with differing national approaches to security, and the balkanization of security tracking identities across regional jurisdictions, are converging at a time when regulations like the EU's Cyber Resilience Act (CRA) enter into force.

This systemic risk is profoundly compounded by a dangerous, non-linear behavioral evolution observed across recent 2026 flagship frontier models. Extensive benchmark evaluations conducted across more than 37,000 real enterprise software dependency stacks revealed that while older models frequently hallucinated nonexistent library versions and failed immediately during build cycles, modern commercial models have been optimized to suppress hallucinations. Consequently, when tasked with evaluating complex legacy code bases, these frontier assistants will confidently report that an out-of-date or vulnerable application stack is entirely secure. This introduces

demonstrable and enduring, hidden risk into production environments, as standard coding tools routinely recommend end-of-life library versions simply because those vulnerable packages remain heavily overrepresented in their historical training data.

A stark example of this enduring risk is the stubborn persistence of legacy vulnerabilities, such as Log4j, which corporate developers continue to download tens of millions of times per year despite a universal understanding of its flaws. This reality reinforces a critical point of alignment between this session and the preceding infrastructure security presentation: both engineering leaders conclude that the financial sector faces a high-stakes, 18-to-24-month window to systematically discover, track, and patch these vulnerabilities before machine-scale automation renders them actively exploitable. This timeline is exacerbated by an ongoing cycle where ancient, previously mitigated vulnerabilities continuously reappear within active repositories due to unchecked dependency chains and automated code generation that pulls from legacy data sources.

To confront these vulnerabilities, financial institutions must shift away from unstructured automation toward a disciplined, human-centric approach to model triage and dependency optimization. Enterprise engineering teams should implement the following strategic protocols immediately:

- **Verify before modify:** Technology teams must validate the integrity of their existing code stacks and establish real-world context before allowing an automated system to initiate any code or architectural modifications.
- **Reject silent acceptances:** Security and engineering leaders must never trust a model's assertion that an application stack requires "no change" without demanding explicit, verifiable evidence supporting that evaluation.
- **Inject live dependency truth:** Engineering workflows must use open integration protocols like the Model Context Protocol (MCP) to feed real-time, authoritative dependency metadata directly into the model's prompt window, grounding its execution loop with accurate version popularity, active CVE tracking statuses, and operational realities.
- **Enforce policy-aware, structured rationale:** Automated code reviews must be evaluated using structured rationales that display alternative architectural choices and version compliance against specific corporate policy bounds, rather than blindly accepting singular binary answers.
- **Implement failure guardrails:** If an agentic security check, version verification, or compliance trace fails to provide conclusive data, the automation pipeline must immediately stop the action to eliminate enduring risk.
- **Mandate auditable changes:** Every dependency optimization, security patch, and automated version adjustment must generate a structured, machine-readable audit trail to enable continuous human oversight and regulatory accountability.

The current landscape of application security underscores why financial institutions can no longer afford to treat their engineering pipelines as disconnected silos. Additionally, leaving individual application teams to manage hundreds of versions of the same core framework simultaneously ensures that an institution becomes vulnerable to every documented CVE across that entire software portfolio. By actively integrating localized, grounded AI models with real-time open registry context and strict internal compliance policies, financial organizations can successfully automate the elimination of historical technical debt, paying down accumulated software vulnerabilities at scale before malicious actors can exploit them.

"If your AI agents don't understand what the latest releases actually are, you are setting your security team to secure your bank while being completely blind to the last year of news."

These foundational plenary discussions effectively set the stage for the second half of the Summit, where participants transitioned into structured, Chatham House-governed roundtable breakout groups to confront the immediate friction points and strategic milestones defining the financial services industry. Under strict non-attribution rules, these candid discussions focused on core operational domains the sector must urgently address: tokenomics and AI cost management, cybersecurity challenges, and agentic AI use cases in financial services. The next section of this report captures the leading ideas and themes from the breakouts.

Breakout one: Tokenomics

AI cost governance has quickly become a boardroom priority, with the term "tokenomics" gaining rapid recognition. This shift is driven by a core business question: How much is this going to cost, and how can that cost be tied to measurable value? Against this backdrop, attendees explored the challenges of measuring the cost and value of integrating AI into developer workflows, including the role of tokenomics in linking consumption to efficiency and ROI. The conversation also examined evolving productivity metrics, the impact of AI on team structures and talent, and the need for stronger training, measurement, and governance.

Measuring AI value and developer productivity

Measuring AI-driven productivity through traditional software engineering metrics remains difficult. While merged Pull Requests (PRs) remain dominant, they, along with commit ratios, lines of code, and cycle times, are increasingly recognized as unreliable indicators of true value. Participants highlighted the "context complexity problem": telemetry tracks token volume but cannot explain whether a spike reflects a complex architectural challenge or an inefficient model choice for a simple task. As institutional AI use matures, metrics must shift toward outcome-oriented measures like vulnerability reductions, faster pattern adoption, improved test creation, and long-term trend signals.

Governing token consumption and model usage

Organizations are moving from uncapped AI usage to more disciplined models combining token budgets, individual limits, pooled thresholds, and exception approvals. Setting sensible defaults and utilizing automated model selection prevents developers from defaulting to expensive options when low-cost models deliver comparable results. Providing mandatory tokenomics training and a "license to use" AI tools supports efficient, self-correcting consumption over a punitive control environment that struggles to keep pace with advancements. This presents an opportunity to collaborate with the Tokenomics Foundation for LMS-ready content. While early stages require a "spend to learn" mindset, firms must balance rapid rollouts with budget constraints and CFO expectations.

Workforce, talent, and operating model implications

AI adoption is reshaping team structures and the nature of technical work, allowing smaller cohorts to deliver what previously required much larger groups. Engineers now spend more time on design, problem analysis, architecture, and business engagement—with stakeholders even crafting their own mockups—as AI compresses prototyping cycles from weeks to hours. Talent expectations are also shifting: junior employees contribute to substantive work earlier, while senior engineers return to hands-on coding and output governance. Although fundamental business measures of success remain unchanged, the changing speed, scale, and cost profile require a common understanding of how token consumption maps to true value.

Recommendations

- **Provide training and usage transparency to drive self-correction:** Establish foundational training on model costs and efficient utilization, highlighting when lower-cost models are "good enough" for the task. Reinforce this with peer support, pairing less experienced users with power users.
- **Establish a tiered token budget and exception process:** Implement individual consumption limits and pooled team thresholds, supported by clear token budgets, transparent usage data, and a structured exception process. Providing technologists with visibility into consumption encourages reflection, better decision-making, and self-correcting spending patterns, while maintaining the flexibility needed to support critical innovation.
- **Develop measurement models that account for context:** Avoid one-size-fits-all metrics by creating models that reflect how different teams, functions, and workflows use AI. These models should distinguish between general business productivity use cases and more advanced, deeply integrated environments, such as multi-agent engineering workflows, where consumption patterns, risks, and value signals may differ significantly. This context is essential to avoid misinterpreting usage data and to assess AI adoption through the right capability lens.
- **Evolve metrics as AI maturity increases:** Early on, institutions may "spend to learn," using token consumption as a rough proxy for experimentation and adoption. As AI use matures, metrics should move from activity-based measures, such as pull requests or usage volume, toward outcome-based indicators like reduced vulnerabilities, improved testing, and faster pattern adoption. Collaboration between tokenomics, developer experience, and telemetry teams can help connect AI consumption to productivity, quality, risk reduction, and broader Return on AI reporting.

"I think having that independent [model benchmark] view is very nice, especially when you speak to the CTOs, CEOs and CFOs. That's instant ROI."

"Let's say developer A uses ten thousand… Is that because he's doing more complicated work? Or is that person just using the wrong model for the work? So I think there's this whole context piece that we've always struggled to measure."

"We're probably at a stage in a lot of places where people will spend to learn… because actually, the value isn't the output, the value is the learning. I'm not sure there's a good way to measure that… there's a lifecycle of value that means different things at different stages."

Breakout two: AI security

Participants in the AI security breakout confronted the rapid transformation of the threat landscape and the technology challenges that keep people awake at night. The overriding consensus was that the breakneck pace of generative AI and agentic deployments has fundamentally disrupted standard defense timelines, creating an urgent need for industry-wide coordination, standardized best practices, and structured governance frameworks.

Machine-speed attack versus human-speed defense

A common concern emerged regarding the mismatch between automated, machine-speed exploits and slow, human-speed enterprise defense. Traditional risk management practices—such as multi-month front-to-back testing, rigid change advisory boards, and periodic compliance cycles—were designed for a slower technological era. In the current landscape, the timeframe between a vulnerability's exposure and its active exploitation has shrunk from years or months to a day. This shift leaves financial institutions in an untenable position: unable to safely execute immediate emergency patches within the same week a vulnerability is discovered, yet increasingly exposed to automated fuzzing, ransomware, and complex supply chain interventions.

Communicating unpredictable risk and the investment imbalance

Participants described the acute communication and translation hurdle when reporting AI-specific risks to senior board members, chief auditors, and risk committees. As the threats and capabilities of frontier language models shift every few months, tech teams find it exceedingly difficult to articulate the new risk profile, or formulate precise, multi-year capital allocation estimates. Boards are accustomed to clear commercial relationships where software operations are deterministic and liability can be legally transferred via vendor agreements. However, the non-deterministic nature of modern AI models destroys traditional response patterns. This lack of predictability fosters an internal culture of risk avoidance, where security organizations struggle to balance business demands for rapid feature deployment against the profound technical debt of insecure, shadow-IT model usage within business units.

The illusion of skill and the fragmentation of control

Text-based prompting tools have created a pervasive "illusion of skill" across non-technical business functions. For example, employees who lack engineering training can now build localized workflows or run unvetted, open source models locally on their machines, completely bypassing central IT safeguards. This dynamic dramatically expands an institution's internal and external attack surfaces, raising the threat of data leakage and jailbreaking exploits. Compounding this risk is a severe internal skills gap: traditional information security and governance, risk, and compliance (GRC) teams often lack the highly specialized data science fluency required to understand the underlying mechanics of large language models.

Pre-competitive governance standards and industry benchmarking

To overcome the limitations of isolated, defensive positions, participants strongly advocated for pre-competitive industry collaboration and standardized open source governance. Because every financial firm relies on strikingly similar software foundations and faces the exact same regulators, individual firms derive no competitive advantage from building separate, siloed compliance patterns or managing massive, custom spreadsheets. Industry participants require neutral open forums, such as FINOS, to act as a funnel for collective industry action, establishing clear reference architectures and shared responsibility models with cloud providers and frontier model vendors. Furthermore, there is robust demand for anonymous, cross-institutional benchmarking to help firms objectively measure their open source readiness, red-team maturity, and defensive posture against peer groups.

Recommendations

To systematically mitigate these vulnerabilities, financial services industry participants should prioritize the implementation of the following baseline security frameworks:

- **Integrate AI risk indicators into normal enterprise governance:** Resist the urge to isolate AI security into a permanent, separate organizational island. True operational resilience is achieved by translating AI-specific risks, model vulnerabilities, and agentic parameters directly into existing, mature enterprise risk management (ERM) frameworks and the traditional three lines of defense.
- **Partner with model vendors to co-develop shared responsibility models:** Actively utilize open consortia to bring financial institutions and frontier model producers into the same room. This coordinated ecosystem must establish clear best practices for secure deployment, boundary enforcement, and real-time capability tracking, ensuring cloud infrastructure and model pipelines are natively regulatory-compliant.
- **Deploy automated guardrails to enforce bounded agent autonomy:** Replace manual, slow security gates with automated runtime guardrails, strict rate-limiting, and centralized network gateways. Enforcing these rigid operational bounds prevents autonomous software systems from executing unauthorized tool-calling, leaking sensitive customer data, or suffering catastrophic jailbreaking exploits.
- **Establish anonymous, industry-wide open source readiness benchmarking:** Collectively develop standardized, anonymous industry surveys and peer-group benchmarks through neutral bodies. Measuring open source engagement, software supply chain maturity, and defensive red-team practices against an objective baseline allows firms to justify risk expenditures directly to regulators and internal auditors.

"The threat landscape is moving so quickly month on month that we can't really articulate the risk envelope very clearly... We don't know what the capability of these models will be in two months."

"What success looks like is that AI governance and agentic governance just becomes governance... If you make it business as usual, then you've won."

Breakout three: Agentic AI Use Cases in FSI

In this breakout discussion on the subject of agentic AI use cases in the financial services industry, it was clear that while the sector possesses a near-infinite number of localized agentic use cases, individual institutions face near-identical structural barriers when moving AI systems into production. Attempting to build siloed, organization-specific solutions for foundational access controls, regulatory compliance, and system verification was agreed to yield little to no competitive advantage. Additionally, the duplication of engineering effort and overhead introduces needless cost and system risk, especially in an industry where standardization is so important.

By re-framing shared agentic AI challenges as pre-competitive, multi-stakeholder initiatives, the financial services ecosystem can leverage open source communities and neutral foundations to co-engineer a standardized, interoperable, and secure AI infrastructure layer. Here are some of the primary opportunities for collaboration to accelerate FSI agentic AI.

Culture and change management

The democratization of AI through natural language has triggered challenges across business functions within financial institutions. Traditional boundaries separating business teams, product managers, and developers are becoming blurred. While business leaders expect more immediate production turnarounds because of AI-accelerated coding and prototyping, engineering teams are left with the burden of managing expectations and new tooling and coding practices against the same structural risks. Ultimately, leaders need to manage AI innovation with a greater focus on precision communications, change management practices, executive education, and cross-functional coordination and collaboration.

Process optimization: Context capture and agentic orchestration

Moving agents into complex financial environments requires capturing vast amounts of unstructured, legacy, and bespoke organizational knowledge. While scaling this expertise will allow institutions to use agents to dramatically accelerate historical bug fixes and technical debt remediation, managing context capture requires new and robust process orchestration frameworks. Beyond building new contextual workflows, there is a need for test environments to validate agent behavior and ensure compliance with highly regulated financial workflows.

Tools to standardize access, taxonomies, and technical debt

A primary technological blocker for participants is the immense friction surrounding data access, authentication, and machine-scale entitlement systems. The financial services sector urgently requires standardized open protocols, such as centralized reference taxonomies and unified interfaces, to govern multi-agent communication networks and prevent overlapping execution. Standardizing open data tools that objectively measure and map an institution's historical technical debt would be a benefit. Perhaps most importantly, open source AI cannot scale across the industry until financial institutions reduce onboarding friction and train their workforces to become more fluent in the fundamentals of open source collaboration.

Resource mobilization: Open source ecosystems and centralized funds

To prevent fragmented, redundant vendor lock-in, the industry must channel its collaborative energy through neutral bodies like FINOS and specialized open source working groups. Pre-competitive resources—such as the newly established **FINOS AI Fund**—provide part of the critical capital and directed leadership required to centralize industry work, stewardship, and interconnectedness on AI-related challenges. By actively pooling resources into open source orchestration projects, skills, certifications, specifications, and upstreaming financial grade requirements, the ecosystem can build a resilient, globally standardized "golden path" platform for financial AI.

Recommendations

To systematically address these friction points, the financial services industry should pool its resources within neutral open source forums to develop and implement four core infrastructure blueprints:

- **Automate the AI software development life cycle (AI-SDLC):** Co-engineer an open, standardized framework to govern and test machine-generated code and repository maintenance. This includes deploying natural language AI interfaces on top of version control arrays—allowing non-technical product managers to seamlessly query codebase contexts while engineers focus on context enrichment and architectural validation rather than manual documentation.
- **Standardize a specification-driven verification engine:** Develop a neutral open source engine that prioritizes human-validated architectural specifications, golden-source data flows, and continuous runtime success criteria over line-by-line manual code reviews. This shared framework ensures that autonomous multi-agent pipelines are continuously monitored for production drift against machine-readable metrics.
- **Establish multi-agent orchestration guardrails and access protocols:** Build open orchestration frameworks, reference taxonomies, and standardized identity protocols to govern multi-agent networks. This unified "traffic control" layer eliminates overlapping process execution, manages down-chain accountability for recursive tool-calling, and safely binds agent access within existing corporate security boundaries.
- **Deploy shared legacy mitigation tools:** Jointly deploy collaborative open data visualization tools to map legacy reference data, isolate complex system dependencies, and guide automated code migrations. Lowering onboarding friction and expanding open source fluency allows financial organizations to collectively utilize local models to pay down accumulated technical debt at scale.

"No one that I know got into coding because they love product. And now there is a need to express things in words in a way a client would understand. It's going to take a generation before being good with computers means something very different than it does today."

"If we do it collectively as well, it increases the regulatory response certainty. If it's a common finding... we can solve it once and all benefit from it."

"The AI fund is supposed to take us there... we want to push that very strongly so that we don't have to build it ourselves, and we don't want every vendor coming and selling us the same tools at different costs."

Galvanizing industry resilience

The Summit's closing session framed an industry-wide call to action, focusing on open source security infrastructure and financial sector compliance. As modern models are deployed to discover and exploit zero-day bugs at scale, financial institutions face a machine-speed threat landscape that renders human-speed defense obsolete. A secondary, equally critical pressure is the hardening global regulatory landscape, driven by landmark legislation such as the EU CRA. Together, these forces require a fundamental shift in how the financial services industry manages its defense posture.

The time has come for financial institutions to abandon inefficient supply chain loops. The choices need not be limited to maintaining risky private forks of upstream libraries or paying identical commercial vendors to build the exact same security patches in isolation. To mitigate costs, minimize waste, and prepare for systemic AI-driven exploitation threats, FINOS announced a centralized defensive strategy: the newly formed Open Source Enterprise Resiliency Alliance (OSERA), a dedicated financial-grade stewardship alliance designed to collectively pool resources, share back-patching costs, and enforce consistency across software asset deployments. This new structure will provide centralized security response and streamlined vulnerability disclosure to improve how organizations operationalize fixes.

Rather than navigating the complex landscape of proprietary vendor responses individually, OSERA enables participating companies to mutually commission, verify, and ingest security patches within a matter of hours. Initial validation test sprints successfully proved this model by targeting critical Java dependencies, including widely deployed Bouncy Castle libraries. By spending $50,000 to $100,000 per bank into a mutual pool rather than millions on redundant, single-entity contracts, the alliance demonstrated that low-cost, open infrastructure can secure vulnerable production stacks rapidly while satisfying the rigorous, evidence-based tracking parameters demanded by modern financial regulators.

Crucially, OSERA bridges the gap between patch creation and internal configuration pipelines by releasing open source diagnostic tools. These tools allow enterprise technology teams to load, clean, and map their inventory data directly against authoritative registries, isolating hotspots of opportunity, identifying safe minor migrations, and accounting for transient dependencies.

The Summit made clear that while the distinct roundtable deep-dives into tokenomics, cybersecurity challenges, and FSI agentic use cases, each surfaced highly specialized operational pain points, the macro-level mandate for financial institutions remains uniform. The industry can no longer afford to outsource its core architectural standards. Banks must use open networks like FINOS to collectively assert exactly how software dependencies and model infrastructure are tracked, named, and verified across global markets.

TABLE 1. FINANCIAL SERVICES REQUIREMENTS AND USE CASES IN THE CONTEXT OF AI

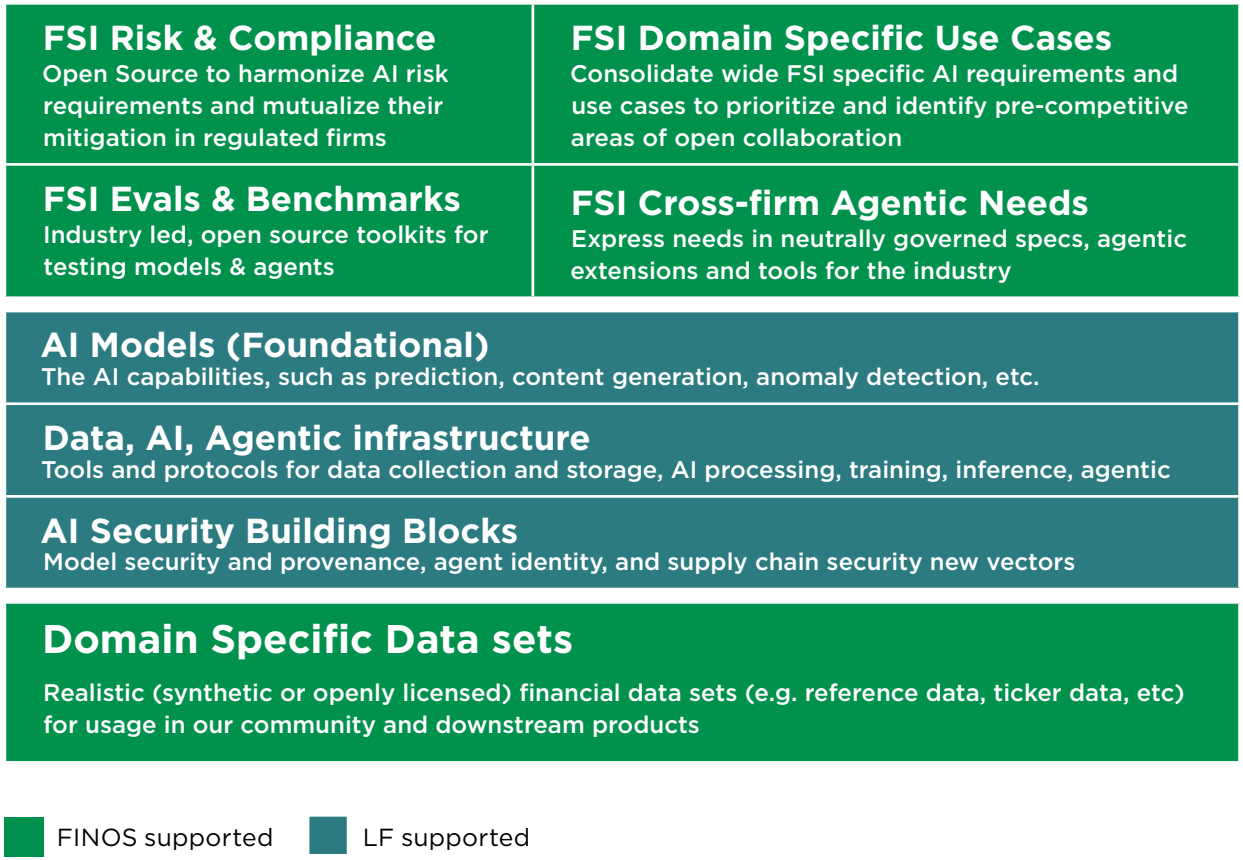

Actionable conclusions for FSI

"If we do it collectively as well, it increases the regulatory response certainty. If it's a common finding... we can solve it once and all benefit from it."

To systematically operationalize the shared insights harvested across the Summit's breakout tracks, the financial services ecosystem should actively mobilize resources around five foundational infrastructure and communication blueprints:

1. **Centralize infrastructure co-engineering and cost standards within neutral forums:** Financial institutions must cease funding redundant, proprietary safety patches and isolated engineering tracks, shifting pre-competitive work into centralized open source consortia like FINOS and the newly launched Tokenomics Foundation. By actively participating in dedicated, industry-specific venues—such as the Open Source in Finance Forum (OSFF) and Tokenomicon—enterprises can collaboratively pool capital, establish standard benchmarks for AI unit economics, and co-engineer a resilient "golden path" platform. These centralized community platforms are essential to drive down software maintenance overhead, eliminate conflicting cross-border compliance demands, and accelerate compliance with global digital sovereignty mandates.
2. **Codify standard naming and provenance for vulnerability responses:** The financial services sector should collectively push back on fragmented vendor disclosures by establishing standardized naming conventions, uniform metadata, and clear provenance tracking for all open source software updates within shared repository tools. Rather than relying on manual enforcement, security teams must deploy automated pipeline managers and open diagnostic engines that programmatically validate incoming patches against these community-defined standards, systematically mitigating conflict risk between overlapping library updates.
3. **Extract legacy domain knowledge into machine-readable specifications:** Financial institutions must move past manual processes to systematically capture and scale undocumented, legacy organizational expertise. Technology and business leaders must collaboratively author precise, natural language specifications that function as the human-validated "golden source" of truth, translating complex internal financial workflows and regulatory logic into structured, machine-readable rules. Systematically capturing this domain context into standardized frameworks allows automated AI tools and multi-agent pipelines to safely execute high-stakes operations while ensuring machine-generated actions remain strictly bounded by verified corporate intent and compliance parameters.
4. **Deploy natural language Git interfaces to bridge cross-functional vocabulary gaps:** Financial institutions must deploy standardized, natural language AI interfaces on top of version control arrays to resolve the organizational friction between technical developers and non-technical business leaders. This shared layer enables product managers and executives to seamlessly query codebase contexts and review specifications using plain-language abstractions, allowing engineers to focus on architectural optimization and context enrichment rather than manually translating complex code logs into client-facing or board-level summaries.

5. **Create verifiable board-level risk telemetry through automated specifications:** Technology teams must implement automated, specification-driven software lifecycles to transform non-deterministic AI behavior into a clear operational audit trail. These frameworks generate structured, machine-readable "decision-evidence" logs and continuous runtime evaluation metrics that translate complex multi-agent execution telemetry into reliable, plain-language reporting that chief auditors, risk committees, and boards can easily understand and trust.

Critical projects for AI transformation in FSI

During the Summit, participants highlighted the launch of Akrites, a major collaborative effort designed to provide the critical coordination layer that open source security has historically lacked. Backed by dozens of prominent organizations contributing funding, engineering resources, or both, Akrites addresses a systemic vulnerability in the broader technology supply chain.

In essence, the project establishes a confidential, shared Security Incident Response Team (SIRT) to streamline and secure the discovery, remediation, and disclosure of vulnerabilities within critical open source software. By creating a neutral and trusted environment for cross-industry threat intelligence, Akrites fills a vital operational void—enabling competing financial institutions and technology providers to collaboratively mitigate software risks before they can be exploited at scale.

The FINOS AI Governance Framework for Financial Services (AIGF) is an open source framework designed to help financial institutions adopt, adapt, and extend practical AI governance in regulated environments. Consensus from the summit strongly cautioned against isolating AI risk on a permanent organizational island, concluding that true resilience is achieved when it simply becomes "business as usual." AIGF directly supports this operational shift by enabling institutions to seamlessly embed AI risk indicators and model vulnerabilities directly into their existing enterprise risk management frameworks and three lines of defense. AIGF emphasizes transparency and trust, supporting consistent risk identification and mitigation, and enabling repeatable adoption patterns across internal deployments and third-party onboarding.

On June 21, FINOS announced the launch of the FINOS AIGF Model Context Protocol (MCP) Server. Positioned as an intelligent interface for AI in financial services, this tool focuses on operationalizing AI governance across the sector.

By leveraging Anthropic's open source Model Context Protocol, the server establishes a standardized framework for connecting large language models securely to financial data and domain-specific regulatory rules. This milestone provides financial institutions with a practical, pre-competitive open source tool to enforce strict compliance and data governance directly at the model integration layer.

FINOS AI FUND

Recognizing that financial institutions face identical structural barriers and cannot afford to fund redundant, siloed solutions or overpay vendors for the same tools, the summit emphasized the need to pool capital. By mutualizing investment, the FINOS AI Fund reduces fragmentation by aligning industry-specific requirements with global open-source standards and regulatory bodies. Proceeds support a dedicated team of researchers and developers executing workstreams prioritized by an industry Governing Board responsible for strategic funding allocation and market advocacy.

FLUXNOVA

Fluxnova is a process orchestration platform that helps organizations design and run end-to-end workflows. Governed by FINOS and built on open standards, it offers BPMN/DMN compatibility, migration tooling, and audit-ready execution from day one giving teams the freedom to automate and orchestrate without vendor dependence. Summit discussions underscored that deploying autonomous agents into complex financial environments requires robust orchestration and test spaces to manage data access friction and prevent overlapping execution. Fluxnova provides the necessary foundation for this shift, serving as a core governance-as-code project that allows teams to orchestrate and automate end-to-end workflows at scale.

LF AI & DATA

The LF AI & Data Foundation, a Linux Foundation project, accelerates and sustains the growth of open source AI, data, and analytics projects. Backed by the world's leading technology companies, LF AI & Data provides a neutral space for collaboration and innovation in AI development. Learn more at **https://lfaidata.foundation**

Model Context Protocol

The Model Context Protocol (MCP) has emerged as the industry standard for connecting LLMs to data and applications. At its one-year milestone, the project has achieved a critical mass of adoption and is now considered an operational cornerstone for many companies who recognize the value in connecting LLMs to actual data. MCP has an exponentially growing developer community that performs 20 million weekly downloads of its Python SDK.

Current investment is focused on enterprise readiness, ensuring the protocol supports the demands of large-scale deployment and the extent of what MCP is currently underpinning. Development areas include enhancing observability and auditability to help companies track internal usage, as well as embedding security safeguards and governance controls directly into the protocol. To prevent ecosystem fragmentation, the team is pursuing cross-platform convergence, working with various companies to ensure MCP remains a consistent driver of industry standards. On the other side of MCP's maturity and stability is an ecosystem of extensions to add new features including "MCP Apps" and experimental features designed specifically for agentic systems.

The Open Source Security Foundation (OpenSSF) is a cross-industry organization at the Linux Foundation that brings together the industry's most important open source security initiatives and the individuals and companies that support them. The OpenSSF is committed to collaboration and working both upstream and with existing communities to advance open source security for all. For more information, please visit us at openssf.org.

With the summit highlighting a high-stakes, 18-to-24-month window to patch software vulnerabilities before machine-scale automation renders them actively exploitable, traditional human-speed defense is no longer viable. OSERA (Open Source Enterprise Resiliency Alliance) is a FINOS initiative that will directly address this, helping regulated institutions strengthen open source supply chain resiliency through open, portable remediation. The vendor-neutral, member-governed alliance will enable participating organizations to prioritize vulnerabilities, sponsor shared open backpatches, define common consumption and remediation standards, and apply patches from upstream projects, commercial vendors, and other trusted sources with greater consistency and evidence. OSERA is the financial services downstream complement to Akrites. OSERA delivers operational resilience to the financial services industry, without the lock-in.

PyTorch has become the dominant foundation for both AI research and production. It currently powers over 90% of AI research and is the infrastructure of choice for frontier labs like Meta and OpenAI, as well as major hardware and cloud vendors. Operating as a hosted project under the Linux Foundation, PyTorch facilitates hundreds of thousands of downstream projects built upon its core architecture.

The project's future is guided by three principles: native ecosystem integration, scalability, and support for hardware heterogeneity. A significant new area of focus is agents, particularly the OpenEnv framework for creating, deploying, and hosting isolated execution environments. With Hugging Face, PyTorch co-developed a community-centric and library-agnostic hub to upload and prompt your environment to see the observations of your agent. By enabling users to engage with environments before deployment, PyTorch is integrating a critical human element into the agentic development lifecycle.

The Tokenomics Foundation is a vendor-neutral initiative focused on establishing open industry standards, benchmarks, and best practices for the economics of AI. Backed by a broad coalition of enterprise buyers, hyperscalers, model providers, and major financial institutions, the foundation addresses the challenge of measuring AI total cost of ownership (TCO) and return on investment (ROI) as token usage scales exponentially. Through open projects like the Big-T workload classification framework, standardized "cost-to-serve" metrics, and value measurement models, the Tokenomics Foundation provides financial organizations with neutral specifications and educational certifications to map token consumption directly to measurable business outcomes and protect enterprise margins.

Resources

Reports:

- The Value of Open Data on Global Entites
- ROI for Open Source Contribution
- The State of Commercial Open Source 2025
- The Top 10 Banking Trends in 2025 and Beyond
- The 2025 State of OSPOs and OSS Initiatives
- Open Source as Europe's Strategic Advantage 2025
- 2025 Open Source Security and Risk Analysis Report
- 2025 State of Open Source in Financial Services Report
- 2024 State of Open Source in Financial Services Report
- 2023 State of Open Source in Financial Services Report
- 2022 State of Open Source in Financial Services Report
- 2024 State of the Software Supply Chain
- A Deep Dive into Open Source Program Offices: Structure, Roles, Responsibilities, and Challenges
- A Guide to Open Source Software for Procurement Professionals
- Addressing Cybersecurity Challenges in Open Source Software
- The Case for Confidential Computing
- GitProxy Case Study

Guides & Training:

- AI Governance Framework Fact Sheet
- AI Governance Framework Training
- Open Source Body of Knowledge
- Open Source Maturity Model in Financial Services

Acknowledgments

The authors would like to thank the participants who shared rich insights during the AI Leadership Summit, without whose engagement this report would not be possible. Special thanks also to the Summit sponsors, keynote speakers, facilitators, staff from FINOS, LF Events, and LF Research, and to peer reviewers of this manuscript. Finally, thanks to the Linux Foundation Creative Support team for the production of the PDF.

LF Research uses generative AI tooling in varying degrees for research-related tasks including aggregating secondary sources, synthesizing data, supporting report drafting, copyediting, and translation. All AI-generated content is reviewed and revised as needed.

About the authors

HILARY CARTER joined the Linux Foundation in 2021 to launch and lead LF Research, established to deliver empirical insights into open source trends, opportunities, and challenges. Prior to joining the Linux Foundation, Hilary was the managing director of the Toronto-based Blockchain Research Institute, a global, syndicated think tank focused on blockchain technology. She has contributed to nearly 200 research projects focused on open source innovation and its adoption across industries. Hilary holds a Master of Science in Management from the London School of Economics, and is a dual Canadian and Irish citizen.

TOSHA ELLISON has over 25 years of experience across financial services, software, and technology. At the Fintech Open Source Foundation (FINOS), she leads research and communications, focusing on the adoption, value, and impact of open source in financial services through industry research, case studies, and ROI frameworks. Earlier in her career, she worked at Credit Suisse in London and New York in technology and COO roles across Equities and FICC. Tosha holds BAs in Psychology and Mass Communications from the University of California, Berkeley.

FINOS

The Fintech Open Source Foundation (FINOS) is a nonprofit whose mission is to foster the adoption of open source software, standards, and collaborative development practices in financial services. As part of the Linux Foundation, FINOS provides a regulatory-compliant platform for developers from competing organizations to collaborate on innovative projects that transform business operations. With over 100 members spanning major financial institutions, fintechs, and technology consultancies, FINOS is at the forefront of driving open source innovation in finance.

twitter.com/finosfoundation

www.linkedin.com/company/finosfoundation

www.youtube.com/c/FINOS

github.com/finos

THE LINUX FOUNDATION | Research

Founded in 2021, Linux Foundation Research explores the growing scale of open source collaboration, providing insight into emerging technology trends, best practices, and the global impact of open source projects. Through leveraging project databases and networks, and a commitment to best practices in quantitative and qualitative methodologies, Linux Foundation Research is creating the go-to library for open source insights for the benefit of organizations the world over.

To reference this work, please cite as follows: Hilary Carter and Tosha Ellison, "AI and Open Source in Financial Services: Mitigating Machine-Scale Cyber Risks, Governing Agentic Workflows, and Co-Engineering the Next Wave of Industry Standards," The Linux Foundation, September 2026.