This Week in AI for Financial Services

Week of May 1, 2026 | S. Andersen, CTI | Creative Technology and Innovation

Executive Summary

This week was the operational test of the productivity-versus-risk duality that the prior week's earnings calls set up. Visa and Mastercard both reported and used their calls to shift the agentic commerce conversation from press release to financial guidance, with Visa posting its strongest revenue growth since 2013 at 17 percent and Mastercard reporting 16 percent net revenue growth. Both networks now quantify stablecoin and agent-initiated transaction volume on the call rather than in side decks, and both are betting that authenticated tokens plus value-added services will preserve unit economics as agent volume grows. Visa added five blockchains to its stablecoin settlement pilot on April 29, taking the total to nine and the annualized run rate to roughly 7 billion dollars, up 50 percent in a quarter. Mastercard formalized BVNK as its planned stablecoin acquisition and rolled out Verifiable Intent for agent payments in deeper partnership with OpenAI.

On the bank side, the headline event of the week was Customers Bank CEO Sam Sidhu using an AI clone of his own voice to deliver part of the firm's Q1 earnings call on April 24, then announcing a multi-year OpenAI partnership embedding OpenAI engineers on-site to automate lending, deposits, and payments. The bank targets compressing commercial loan closings from 30 to 45 days down to roughly seven, and complex account openings from over a day to under 20 minutes. The stunt is less interesting than the architecture: a 26 billion dollar regional bank explicitly arguing that smaller institutions can move faster on AI than the universal banks can, and locking in a co-development deal with OpenAI to do it. Banking Circle launched a stablecoin settlement service for USDC, USDG, and EURI on April 27. Meta began paying creators in USDC via Stripe and Circle on April 29. India's Finance Minister convened bank CEOs to step up cybersecurity preparedness specifically against Anthropic's Mythos. The center of gravity this week was the move from announcement to throughput, and from Q1 narrative to Q2 execution.

1. Banking and Financial Institutions

Customers Bank deploys CEO AI clone on Q1 earnings call and signs OpenAI partnership. On April 24, Customers Bank CEO Sam Sidhu told analysts roughly 30 minutes into the firm's Q1 2026 earnings call that the prepared remarks they had been listening to were delivered by his AI clone, not him. He used the moment to announce a multi-year strategic partnership with OpenAI under which OpenAI engineers will be embedded at the 26 billion dollar regional bank to co-develop agentic workflows for lending, deposits, and payments. Sidhu told CNBC the partnership is expected to drive the bank's efficiency ratio from roughly 49 percent to the low 40s, with commercial loan closings compressing from 30 to 45 days down to roughly seven, and complex commercial account openings from over a day to under 20 minutes.

Approximately 75 percent of Customers Bank employees already use OpenAI-powered tools, and the bank says AI now writes about half of its software code, saving roughly 28,000 hours of work or the equivalent of about 15 full-time hires. The strategic premise is that smaller institutions face lower regulatory complexity than the universal banks and can therefore move faster, and the partnership puts that thesis on the record.

Banks renegotiate cloud contracts as AI workloads break legacy assumptions. PYMNTS reported on April 28 that financial institutions are systematically reopening cloud agreements written for transactional, predictable workloads, on the grounds that those agreements no longer fit AI's compute-heavy and bursty consumption patterns. Banks are pushing for interoperability, multi-cloud portability, regulatory compliance terms, and reduced single-provider dependence rather than just cost optimization. Customers Bank's April 27 earnings-call demonstration of its AI avatar reinforced the same point from the operational side: AI is no longer a discretionary application sitting on top of the core, it is moving into the core itself. For CIOs, the practical consequence is that procurement, model risk, and resilience programs that used to live in separate review tracks now have to be reconciled at the contract layer.

Wall Street Q1 cuts come into sharper focus as bank executives begin connecting headcount to AI. Bloomberg's April 15 analysis of Q1 reporting found that the largest U.S. banks shed roughly 5,000 jobs in the quarter, a more than seven-fold increase from the 707 employees the same group cut in Q1 2025, even as the firms posted record first-quarter profits. Most CEOs declined to attribute the cuts to AI directly. Bank of America's Brian Moynihan was the most explicit, telling analysts that AI gives the firm "places to go" and that the bank is still in the early stages of what the technology will do. The Q1 disclosures continued to land through the week of April 27 in analyst notes and in workforce-strategy commentary, with the practical question shifting from whether AI is replacing roles to how quickly the savings show up in efficiency ratios. Customers Bank framed it most directly: more revenue per employee, slower hiring, and AI agents handling work that previously required dozens of staff.

Regional banks confront the AI readiness gap that the Customers Bank deal makes visible. A Grant Thornton 2026 AI Impact Survey cited in coverage of the Customers Bank announcement found that half of banking executives say governance and compliance are already limiting AI performance, while only 18 percent are confident they could pass an independent audit of their AI controls. Tier 1 banks above 100 billion dollars in assets are reported to be at roughly 75 to 80 percent full AI integration. Regional and community banks sit at 30 to 40 percent, mostly stuck in pilot or limited-deployment stages. Customers Bank's deal is being read by regional-bank CIOs and boards as a template for how to leapfrog the gap by pairing directly with a frontier lab and co-developing licensable solutions, rather than buying off-the-shelf agent products from the existing core-banking incumbents.

Further reading:

- CNBC (April 27, 2026): This Bank CEO Let His AI Clone Handle an Earnings Call — Now He's Signing an OpenAI Deal
- PYMNTS (April 28, 2026): Banks Confront the Cloud Contracts Built for Yesterday
- American Bazaar (April 27, 2026): Customers Bank Partners with OpenAI; CEO Uses AI Clone on Earnings Call
- Bloomberg (April 15, 2026): Wall Street Banks Cut 5,000 Jobs Even as They Notched Record Profits

2. Payments, Stablecoins, and Agentic Commerce

Mastercard Q1 2026 earnings put agentic commerce and BVNK acquisition at the center of the strategy. Mastercard reported Q1 2026 net revenue of roughly 8.4 billion dollars, up 16 percent on a reported basis and 12 percent on a currency-neutral basis, with adjusted diluted earnings per share of 4.60 against a 4.41 consensus. CEO Michael Miebach used the April 30 call to position Agent Pay and virtual cards as the architecture for AI-driven commerce, naming OpenAI, Google, Microsoft, the FIDO Alliance, and Craftsman as integration partners. Mastercard launched Verifiable Intent, a tamper-resistant record of what a consumer authorized when an AI agent acts, designed to make agent-initiated transactions resemble traditional card transactions in dispute and protection terms. Miebach also formalized the planned BVNK acquisition as the company's bet on stablecoin liquidity, interoperability, and the licensing perimeter for digital-asset settlement. CFO Sachin Mehra noted that value-added services and solutions grew 22 percent year over year, or 18 percent currency-neutral, which is the line management wants investors anchoring to as cards-only growth slows.

Visa Q2 2026 earnings deliver the strongest revenue growth since 2013 with stablecoins and agentic commerce as named drivers. Visa reported fiscal Q2 net revenue of 11.2 billion dollars, up 17 percent year over year and the strongest growth rate since 2013 excluding the post-pandemic rebound and the Visa Europe acquisition effect. GAAP EPS rose 36 percent and non-GAAP EPS climbed 20 percent to 3.31 dollars, beating consensus of 3.18. Payments volume grew 9 percent in constant dollars to 3.7 trillion. Cross-border volume excluding intra-Europe rose 11 percent. CEO Ryan McInerney told analysts on the April 28 call that AI and agentic commerce will expand Visa's addressable market in four ways: digitization, agent-initiated transaction multiplication, micro-transactions where agents pay for their own data and resources, and the spread of authenticated tokens that improve fraud and dispute economics. Management raised full-year guidance for adjusted net revenue growth to low double-digit to low teens and EPS growth to low teens.

Visa expands its stablecoin settlement pilot to nine blockchains as run rate hits 7 billion dollars. On April 29, Visa added Arc, Base, Canton, Polygon, and Tempo to its stablecoin settlement pilot, joining existing support for Avalanche, Ethereum, Solana, and Stellar. The pilot now runs at a 7 billion dollar annualized run rate, up roughly 50 percent quarter over quarter, and Visa now operates more than 130 stablecoin-linked card programs across more than 50 countries. Each new chain serves a distinct settlement profile: Arc is Circle's Layer-1, Base is Coinbase's Ethereum scaling network, Canton brings configurable privacy for regulated capital markets, Polygon and Tempo target low-cost and high-throughput stablecoin liquidity. Read alongside Mastercard's BVNK move, the practical signal is that the two largest card networks are now making explicit, named, multi-chain settlement commitments rather than running parallel pilots quietly.

Meta returns to crypto payments through USDC, Stripe, and Circle. On April 29, Meta began rolling out stablecoin payouts to selected creators in Colombia and the Philippines, with payments processed through Stripe and USDC issued on Solana and Polygon by Circle. Stripe's Link wallet handles the payout. The pilot is the most credible institutional return to crypto-native creator economics that Meta has attempted since the Libra and Diem projects were shut down, and notably it is structured as an infrastructure integration rather than a proprietary token. The choice of USDC reinforces Circle's positioning as the compliance-credentialed default for cross-border platform payouts, and Stripe's role as processor confirms its push to become an "AWS for money" payments platform on the back of stablecoin rails.

Banking Circle launches a regulated stablecoin settlement service in Luxembourg. Banking Circle announced on April 27 a stablecoin settlement service supporting USDC, USDG, and EURI, with two-way fiat-to-stablecoin settlement aimed at cross-border use cases. The launch follows Banking Circle's April 15 receipt of a Crypto-Asset Service Provider license from Luxembourg's CSSF under MiCA. For European corporate treasurers, the launch matters because it places stablecoin settlement inside a fully licensed bank rather than in a fintech wrapper, materially reducing the counterparty and rehypothecation questions that have kept stablecoin payouts off institutional roadmaps in EU jurisdictions.

x402 transaction milestones and Agentic.Market expansion continue. Cobo's April 28 protocol overview reported the x402 protocol now processes more than 75 million transactions per month with over 24 million dollars in monthly volume across Base, Solana, Ethereum, and Polygon. Daily x402 transaction counts hit roughly 151,000 on April 24, more than double the prior baseline, with Coinbase's April 20 launch of Agentic.Market named as the proximate catalyst. As of late April, Agentic.Market and x402 cumulatively cover more than 165 to 167 million transactions and roughly 50 million dollars in cumulative volume, with about 85 percent settled on Base. Read against Visa's and Mastercard's earnings calls, the practical observation is that the agent-to-agent economy now has both a permissioned card-network rail with embedded fraud protection and a permissionless on-chain rail with sub-cent fees, and it is now an empirical question which one wins which segment.

Further reading:

- PYMNTS (April 30, 2026): Mastercard Sees Agentic Commerce Growth Despite Travel Headwinds

- PYMNTS (April 28, 2026): Visa Earnings Spotlight Agentic Commerce as Next Growth Engine
- Visa (April 29, 2026): Visa Accelerates Stablecoin Momentum — Adding Five Blockchains for Settlement
- CoinDesk (April 29, 2026): Tech Giant Meta Starts Paying Some Creators in Stablecoin with Stripe's Support
- Phemex News (April 27, 2026): Banking Circle Launches Stablecoin Settlement Service for USDC, USDG, EURI

3. AI Infrastructure for Finance

Cloud-contract renegotiation becomes the binding constraint on AI deployment. PYMNTS' April 28 piece on cloud-contract reopenings is the clearest articulation yet of the pattern that has been building across the Q1 earnings cycle. Banks negotiated their original cloud agreements around predictable transactional workloads, storage tiers, and uptime guarantees. AI workloads are bursty, compute-intensive, and require data movement between governance domains in ways that those agreements did not anticipate or price. Banks are now renegotiating around interoperability, regulatory compliance terms, and reduced single-provider dependence rather than around cost optimization, and the open question is whether hyperscaler pricing power survives the shift. The signal for vendor selection is that banks that have built portability into their contracts now have leverage they did not have a year ago.

Customers Bank reframes the build-versus-buy debate at the regional level. Sidhu's April 24 framing of the OpenAI deal as a co-creation rather than a license is the operational frontier of the AI-infrastructure conversation. Both sides are contributing resources to build new tools, with OpenAI gaining real-world use cases inside a regulated financial institution and Customers Bank getting access to frontier engineering. The structure mirrors the early days of cloud-native deployment, when reference customers became the design partners for what the platforms eventually shipped to everyone else. Sidhu's stated goal of being one of the first truly AI-native regional banks by the end of 2026, with bankers spending more of their time on work that creates value for clients, sets a benchmark that competitors will have to either match or explain why not.

SR 26-2 leaves agentic AI explicitly out of scope, and banks now have to self-govern with no rulebook. The April 17 joint Federal Reserve, OCC, and FDIC issuance of SR 26-2, which replaced SR 11-7 after 15 years, is now being absorbed across model risk programs. The guidance modernizes the framework for traditional models with an Exposure plus Purpose proportionality framework, but the agencies explicitly excluded generative and agentic AI on the grounds that those technologies are "novel and rapidly evolving." The agencies have also signaled they will issue an interagency request for information on AI specifically. The practical effect for banks deploying agentic AI for fraud detection, customer service orchestration, or trading support is that they have no specific supervisory guidance and must extrapolate from a framework that openly does not cover their use case. This is the regulatory vacuum that BPI's October 2025 OSTP RFI response anticipated, and it now sits as the single most consequential governance question heading into Q2 board meetings.

FSSCC AI Risk Management Framework completes its rollout for member firms. The Financial Services Sector Coordinating Council has now released the final four of six resources in its Financial Services AI Risk Management Framework, a financial-services-tailored retooling of the NIST AI RMF. The framework includes an AI adoption stage questionnaire, a risk-and-control matrix, a user guidebook, and a control objective reference guide. For banks looking to assemble a defensible internal control narrative under SR 26-2's principles-based approach, the FSSCC framework, the FS-ISAC April 21 advisory, and the BPI joint trade comments together provide the working substitute for the agentic AI guidance that the regulators have explicitly declined to issue.

Further reading:

- PYMNTS (April 28, 2026): Banks Confront the Cloud Contracts Built for Yesterday
- OCC (April 17, 2026): Bulletin 2026-13 — Model Risk Management: Revised Guidance
- AI2Work (April 23, 2026): Fed and OCC Overhaul Bank Model Risk Rules but Leave AI Uncharted
- ABA Banking Journal (March 26, 2026): FSSCC Releases Additional AI Resources for Financial Institutions

4. Fraud, Security, and Identity

India's Finance Minister convenes bank CEOs over Mythos cybersecurity preparedness. On April 27, the Indian Centre's Finance Minister Nirmala Sitharaman chaired a meeting with public-sector and private-sector bank CEOs urging them to step up cybersecurity preparedness amid global concerns about Anthropic's unreleased Claude Mythos model. Banks were asked to take pre-emptive steps to secure IT systems, protect customer data, and guard financial resources against AI-linked threats. The Centre's framing was that the gap between vulnerability disclosure and exploitation has effectively collapsed once advanced models are in adversaries' hands. India's intervention is the first time a G20 finance ministry has formally convened its banking CEOs specifically on a single named frontier-model risk, and it sharpens the question of how fast the rest of the G7 banking supervisors will move from convenings to written expectations.

Help Net Security's April 22 sector report quantifies the AI-era threat baseline. Help Net Security published a financial-sector cyber-threats analysis on April 22 finding that approximately 90 percent of breaches affecting financial institutions in 2025 carried a financial motive, with data breaches accounting for roughly 64 percent and ransomware for 36 percent. The average cost of a data breach in financial services reached 5.56 million dollars, putting the sector second among all industries by breach cost. Personal data was compromised in 54 percent of cases, internal organizational data in 35 percent, and credentials in 22 percent. Read alongside the FBI's 2025 Internet Crime Report, FS-ISAC's April 21 sector advisory, and the BPI testimony to House Homeland Security on April 23 citing more than 20 billion dollars in cyber-enabled fraud losses for 2025, the empirical baseline for the AI-versus-AI cyber arms race is now sharper than at any point in this cycle.

Visa positions agentic-commerce fraud architecture around authenticated tokens. On the Q2 earnings call, McInerney told analysts that Visa cardholders will continue to be protected against fraud in agent-initiated transactions and that the use of authenticated tokens in agentic commerce will help reduce fraud rather than increase it. Issuers and merchants will get richer transaction data, including agent-side intent, which Visa argues will improve dispute resolution. Mastercard's Verifiable Intent product is the parallel construct on the other network. The structural argument the networks are now making, in unison, is that agent-initiated transactions are more, not less, suited to network-grade fraud control because the data trail is denser. The counter-argument, which auditors and regulators will eventually have to settle, is whether token-level data is actually the right artifact for ECOA, UDAAP, and dispute-rights purposes when the consumer's authorization happened in a chat window rather than at a checkout.

Further reading:

- Business Standard (April 27, 2026): Why Indian Govt Is Warning Banks Against Anthropic's Claude Mythos AI
- Help Net Security (April 22, 2026): Shadow AI, Deepfakes, and Supply-Chain Compromise Are Rewriting the Financial Sector Threat Playbook
- BPI (April 25, 2026): BPInsights — Industrialization of Fraud Testimony Roundup
- Motley Fool (April 30, 2026): Mastercard Q1 2026 Earnings Transcript — Verifiable Intent for Agent-Initiated Payments

5. Funding, Mergers, and Market Signals

Rogo raises 75 million dollars Series C and opens London office. Agentic financial-workflow platform Rogo announced a 75 million dollar Series C on April 29 led by Sequoia Capital, with participation from Henry Kravis, Wells Fargo, Thrive Capital, Khosla Ventures, Tiger Global, and J.P. Morgan, taking total funding to over 165 million dollars. Capital is earmarked for international scaling, particularly in Europe through the new London office, and continued development of Rogo's end-to-end agent for deal origination and execution. The fact that two universal banks, Wells Fargo and J.P. Morgan, are now in the cap table of an applied-

AI banking workflow platform alongside frontier-tech VCs is the practical signal: bank corporate development is no longer just a procurement vehicle for AI, it is a strategic-investment vehicle.

Manifest OS lands 60 million dollars Series A at a 750 million dollar valuation for legal AI. On April 28, Manifest OS announced a 60 million dollar Series A at a 750 million dollar valuation, with the explicit pitch of "ending the billable hour" by automating legal research and document drafting. The round is relevant for fintech because the same applied-AI-into-regulated-workflow pattern is what's driving valuation expansion in tax, compliance, insurance, and procurement. Performativ, a wealth-management infrastructure platform, also closed a notable round the same day with Deutsche Börse Group leading and Rabo Investments and FinTech Collective participating, focused on unifying fragmented adviser systems for the 63 trillion dollar wealth-management market. The two deals reinforce that vertical AI in regulated finance workflows continues to capture a disproportionate share of late-April fintech VC flow.

Capital concentration into AI-native fintech remains the dominant 2026 pattern. Crunchbase data through the end of Q1 confirmed total global venture investment of 300 billion dollars in 6,000 startups in Q1 2026, an all-time high and roughly 70 percent of all 2025 venture spending in a single quarter. Fintech globally drew about 12 billion dollars across 751 deals year to date through April 6, with vertical AI applied to regulated workflows capturing a disproportionate share. The two largest M&A deals of Q1 were Savvy Games Group's planned 6 billion dollar acquisition of ByteDance's Moonton and Capital One's now-closed 5.15 billion dollar Brex acquisition. The April 24 to 29 funding cycle continued the pattern of capital concentrating into a smaller number of larger rounds, with seed deal counts down 30 percent year over year even as seed dollars rose. Treasurers and corporate-development teams reading these numbers should expect both higher entry valuations and faster consolidation pressure in 2026 strategic planning.

Visa and Mastercard's stock reactions price in the agentic-commerce thesis. Mastercard shares rose modestly post-earnings as management raised full-year 2026 net revenue guidance to the high end of low double-digit growth. Visa raised its full-year guidance to low double-digit-to-low-teens revenue growth and low-teens EPS growth. The combined market-cap reaction across the two networks confirms that public investors are now willing to underwrite agentic-commerce and stablecoin-settlement growth rather than treat it as optionality. For bank treasurers running competitor-set comparisons in 2026 strategic planning, the implication is that the cost of capital advantage that the networks have over the banks on AI investment may be widening rather than narrowing.

Further reading:

- Tech Startups (April 29, 2026): Top Startup and Tech Funding News — Rogo, Aidoc, Manifest OS, Performativ

- FinTech Global (April 29, 2026): Rogo Raises 75M Series C to Scale AI Finance Platform
- Crunchbase (April 1, 2026): Q1 2026 Shatters Venture Funding Records as AI Boom Pushes Startup Investment to 300B
- Investing.com (April 30, 2026): Mastercard Q1 2026 Slides — Revenue Beats but April Trends Soften

6. Regulation, Governance, and Policy

SR 26-2 absorption shifts model-risk programs into self-governance for agentic AI. The April 17 joint Fed, OCC, and FDIC issuance of SR 26-2 continues to drive rewrites of internal model-risk and AI-governance documentation. The shift from prescription to judgment, with the Exposure plus Purpose proportionality framework replacing the SR 11-7 annual-validation cadence, lands well with traditional model-risk teams. The explicit exclusion of generative and agentic AI is the harder question. Banks that read the exclusion as a green light to deploy without supervisory friction risk the opposite: examiners will fall back on the 2023 interagency third-party guidance and on safety-and-soundness expectations, and the absence of bright-line rules likely tightens, not loosens, the practical compliance bar. The forthcoming joint RFI on AI specifically, signaled by the OCC in the April 17 release, is now the most-watched piece of supervisory paper of 2026.

India treats Mythos as a financial stability matter and asks banks to harden defenses. The April 27 Sitharaman convening with bank CEOs is the strongest signal yet that frontier AI cyber risk is now being treated as a financial-stability concern by sovereign regulators outside the AI-developer jurisdictions. The Bank of England's October 2025 AI Consortium minutes had already raised the agentic-workflow systemic-risk concern, and Governor Andrew Bailey's New York remarks earlier this month framed AI cyber risk in stability terms. India's intervention adds the world's third-largest banking system by population to the explicit-supervisory-attention list. U.S. banking regulators remain conspicuously silent on Mythos exposure on the record, with the Treasury convenings continuing in private. The trans-Atlantic and now Asia-Pacific gap in supervisory posture is itself worth tracking.

BPI testifies to House Homeland Security on AI as a fraud force multiplier. The House Homeland Security Committee's April 23 hearing, "Online Scams, Crypto Fraud, and Digital Extortion," featured TRM Labs, USTelecom, and other witnesses citing more than 20 billion dollars in cyber-enabled fraud losses in 2025, up 26 percent year over year, and describing AI as a force multiplier for criminals through deepfakes, voice cloning, and phishing at scale. The House Financial Services Committee held a parallel hearing on sanctions effectiveness. Both hearings strengthen the policy floor for AI-specific supervisory expectations and align the legislative branch's framing of AI risk with the FS-ISAC, FBI, and Treasury framings. They also reinforce that the BPI joint comment program on the NIST AI RMF and Security Considerations for AI Agent Systems is now landing on a Hill that has its own data on the AI-fraud cost.

FCC and FTC enforcement tempo on AI-adjacent fraud continues. On the conduct side, the FCC's April 2 proposed 4.5 million dollar fine against Voxbeam Telecommunications for facilitating fraud scams targeting U.S. banks, and the April 23 New York merchants' suit against Visa and Mastercard on interchange liability under the 2019 settlement, both signal that the agencies are not waiting for an AI-specific framework before tightening the conduct envelope around payments. Neither action is framed around AI directly, but together they shape the operating perimeter inside which AI-driven commerce has to land. The FTC also issued letters to four payment service providers warning against alleged "debanking" activities, which is procedurally adjacent but doctrinally important for anyone building agentic-commerce onboarding flows.

Further reading:

- OCC (April 17, 2026): Updated Model Risk Management Guidance — Bulletin 2026-13
- Business Standard (April 27, 2026): Indian Banks, Govt Raise Anthropic Threat — Claude Mythos AI Cybersecurity Risks
- BPI (April 25, 2026): BPInsights — House Homeland Security and FSC Hearings Roundup
- ABA Banking Journal (April 22, 2026): FS-ISAC Releases Advisory on Hardening Cybersecurity from AI

Outlook: Watch Items for the Coming Week

Three threads carry forward. First, the Q1 and Q2 earnings cycle for the payment networks and for Customers Bank establishes a new disclosure floor: agentic-commerce volume, stablecoin settlement run rate, and AI-driven efficiency-ratio trajectory are now investor-conversation items. Other regional banks will face pressure on the Q2 call to either match Customers Bank's specificity or explain why not. Second, the joint OCC-Fed-FDIC RFI on AI in model risk management, telegraphed in the SR 26-2 release, is the single most consequential piece of forthcoming supervisory paper, and its drafting is being read against the SR 26-2 exclusion of generative and agentic AI. Banks should expect the RFI to land in Q2 or early Q3 and should be staging their AI-governance documentation for it now. Third, the Anthropic 90-day public Glasswing report is still due in roughly two weeks, and India's Mythos convening raises the probability that other G20 banking ministries follow with on-the-record cybersecurity advisories before that report lands.

The deeper story of the week is that the productivity-versus-risk duality the prior week's earnings calls set up is now being adjudicated by the market in real time. Visa and Mastercard both have raised guidance and put numbers on agentic and stablecoin volume. Customers Bank has put a multi-year operating-model bet on the table with frontier-lab co-development as the architecture. India has put a sovereign-supervisor warning on the table about the same model class that the U.S. regulators are still privately convening about. The absence of a coherent, explicit supervisory posture on agentic AI in the U.S. is no longer a pacing issue, it is an active asymmetry that other banking systems are now trading against. That, more than any single announcement this week, is what is worth tracking.

Prepared by S. Andersen, CTI. Creative Technology and Innovation. Views are the author's own and do not constitute legal, financial, or regulatory advice.