AI Transforming the Financial Landscape

x9.org/x9ai-study-group-newsletter

Insights into AI's Role in Financial Services

View the X9 AI Study Group

X9 AI Study Group Newsletter - Issue No. 20 (2026)

NIST, NSA and CISA Publish Ten Practices for AI Agent Incidents

NIST CAISSI, the NSA Super Intelligence Security Center and CISA published ten key practices on October 16. The practices cover incidents during “super intelligence” training and evaluation. The focus is on agents taking significant uninstructed actions against operator controls.

The ten practices follow five stages: preparation, detection, response, recovery, and reporting. The document is voluntary. The practices build on NIST SP 800-61r3 for cybersecurity incident response. NIST says this work answers a recommendation in the White House AI Action Plan.

Continue Reading



Banks Detect Unvetted AI Agents Inside Customer Accounts

Major U.S. banks found third-party AI agents working inside customer online accounts. Customers gave their usernames and passwords to personal AI agents. The banks did not vet these agents. Some agents took actions the account holder did not approve, such as a call to the bank with a complaint.

Most of the activity was not hacking. White House AI adviser David Sacks said the agents followed user instructions. He proposed read-only access for agents. A Chase spokesperson told customers to share credentials with no one. JPMorgan CEO Jamie Dimon said on the company's Oct. 13 earnings conference call that JPMorgan is "not going to allow just any outside agents to come in inside the company because it's just too risky for clients and for JPMorgan."

Continue Reading



FDIC Chair Signals Future Guidance on Agentic AI

FDIC Chair Travis Hill said AI agents moving deposits is "the big thing" regulators and banks think about. He spoke on October 20 at an Axios event in Las Vegas. Citi and Bank of America joined JPMorgan Chase to block external AI agents from customer accounts.

Hill said the revised model risk guidance from the FDIC, OCC and Federal Reserve does not work well for generative and agentic AI. The FDIC is considering a public request for information on how to supervise these tools. Hill also noted that about half of all deposits in the system are corporate.

Continue Reading



SEC Names Super Intelligence a FY2027 Exam Priority

The SEC Division of Examinations released its fiscal year 2027 priorities on October 23. The list covers more than 34,000 registrants. Emerging technologies, including “super intelligence,” are a key area of focus.

Other focus areas include adviser conflicts of interest, complex and illiquid products, and private credit. Broker-dealer AML, cybersecurity, and tokenization are also on the list. SEC Chairman Paul Atkins called exams one of the Commission’s most important tools. The Division says the list is not complete.

Continue Reading



Mortgage Firms Must Prove Their AI Controls Work

MBA NewsLink contributor Mark Dangelo writes the mortgage AI debate has moved from policy to proof. Examiners now ask if AI controls work across the enterprise. A written policy alone does not pass an exam.

Dangelo lists nine items for each AI inventory entry. Examples include business purpose, owner, data sources, affected consumers, human review, vendors, and retained records. He calls for decision lineage. Decision lineage is the ability to trace an outcome from source data to final action. For agentic AI, governance must address authority as well as accuracy.

Continue Reading



As the US and China form AI blocs, countries are increasingly torn between them

An Associated Press article examines how intensifying U.S.-China competition over artificial intelligence is pressuring other countries to align with one of two emerging AI blocs.

The U.S.-led Pax Silica emphasizes access to advanced chips, technology partnerships, manufacturing, and secure AI supply chains, while China's World AI Cooperation Organization offers an alternative centered on more affordable, open-source technologies and participation in shaping AI governance and standards.

Kazakhstan illustrates the dilemma, having pursued relationships with both initiatives while attracting significant U.S. and Chinese investment. Other countries are similarly attempting to preserve flexibility rather than making a binary choice.

However, growing strategic competition could make such balancing increasingly difficult, potentially dividing global AI infrastructure, supply chains, technology ecosystems, and standards development along geopolitical lines.

Continue Reading



Agentic commerce is here. The proof of authority is missing.

Proof argues that agentic commerce lacks reliable evidence connecting an AI agent's action to a verified human and the exact authority granted. Its x401 protocol addresses this by having the human sign a constrained mandate covering such terms as amount, recipient, purpose, and timing. The agent presents that credential to the merchant, which verifies the identity, scope, and signature before proceeding, while payment can occur separately through x402 and USDC.

Importantly, x401 builds on OIDF standards. OpenID4VP carries the credential presentation. The current architecture also uses OAuth 2.0 for token exchange and OpenID4VCI for credential issuance. OIDF standards therefore supply much of x401's underlying interoperable identity and credential infrastructure.

Continue Reading



U.S. Consumers Are Wary of AI Agents Handling Their Money

An American Bankers Association survey conducted by Morning Consult found significant consumer hesitation about allowing AI agents to manage financial activities.

Seventy-seven percent are concerned about giving AI permission to move money or make payments, while only 27% are likely to use an AI agent to autonomously manage bank-account cash balances over the next year.

Consumers also strongly favor safeguards: 79% say they should not have to provide online banking credentials to AI tools, 80% believe banks should be able to distinguish customers from their AI agents, and 76% support requiring agents to demonstrate adequate privacy and security protections before accessing customer information.

The findings suggest consumer acceptance of agentic finance will depend heavily on security, transparency, and appropriate controls over agent access to financial accounts.

Continue Reading



Design Tenets for Agent Identity Management Systems

NIST's "Design Tenets for Agent Identity Management Systems" proposes ten principles for securely managing AI-agent identity, authentication, and authorization across enterprise and consumer environments.

Agents should possess distinct, cryptographically verifiable identities linked to the individuals or organizations authorizing their actions, with delegated authority narrowing as tasks move downstream. NIST advocates zero standing privilege, short-lived and revocable authorization, risk-based human intervention, and cryptographically signed mandates defining an agent's intended outcomes and constraints. It also calls for separating AI reasoning from authorization policy, maintaining deterministic access controls, continuously monitoring agent behavior, and preserving tamper-evident records of identity, delegated intent, and actions.

NIST acknowledges existing standards and technologies cannot yet fully implement every tenet, positioning them instead as guideposts for developing more mature agentic identity-management capabilities.

Continue Reading



XBOW says it identified a bug in an application belonging to the U.S. Election Assistance Commission

XBOW Labs details how its autonomous security-testing system discovered a critical remote-code-execution vulnerability in a public web application operated by the U.S. Election Assistance Commission (EAC).

The affected system was a resource and logistics calculator, **not** a voting machine or voter-registration system.

XBOW analyzed the application's public frontend to understand its backend interface, identified suspicious handling of user-supplied data in a simulation feature, and used response-timing behavior to confirm that attacker-controlled input could trigger backend code execution without authentication.

Human researchers subsequently verified the finding. XBOW reported the vulnerability through the EAC's vulnerability-disclosure process, and the EAC validated and remediated it before the November 3, 2026 federal election. The research highlights AI's growing capability to autonomously discover serious vulnerabilities in real-world production systems.

Continue Reading

Stay Ahead in Financial AI

Subscribe to our newsletter for the latest insights and updates on AI in financial services.

Subscribe Now

Archive of Newsletters

AI and Financial Services Insights

- AI Newsletter – Issue No. 19 (2026)
- AI Newsletter – Issue No. 18 (2026)
- AI Newsletter – Issue No. 17 (2026)
- AI Newsletter – Issue No. 16 (2026)
- AI Newsletter – Issue No. 15 (2026)
- AI Newsletter – Issue No. 14 (2026)
- AI Newsletter – Issue No. 13 (2026)
- AI Newsletter – Issue No. 12 (2026)
- AI Newsletter – Issue No. 11 (2026)
- AI Newsletter – Issue No. 10 (2026)
- AI Newsletter – Issue No. 9 (2026)
- AI Newsletter – Issue No. 8 (2026)
- AI Newsletter – Issue No. 7 (2026)
- AI Newsletter – Issue No. 6 (2026)
- AI Newsletter – Issue No. 5 (2026)
- AI Newsletter – Issue No. 4 (2026)
- AI Newsletter – Issue No. 3 (2026)
- AI Newsletter – Issue No. 2 (2026)
- AI Newsletter – Issue No. 1 (2026)
- AI Newsletter – Issue No. 7 (2025)
- AI Newsletter – Issue No. 6 (2025)
- AI Newsletter – Issue No. 5 (2025)
- AI Newsletter – Issue No. 4 (2025)
- AI Newsletter – Issue No. 3 (2025)
- AI Newsletter – Issue No. 2 (2025)
- AI Newsletter – Issue No. 1 (2025)