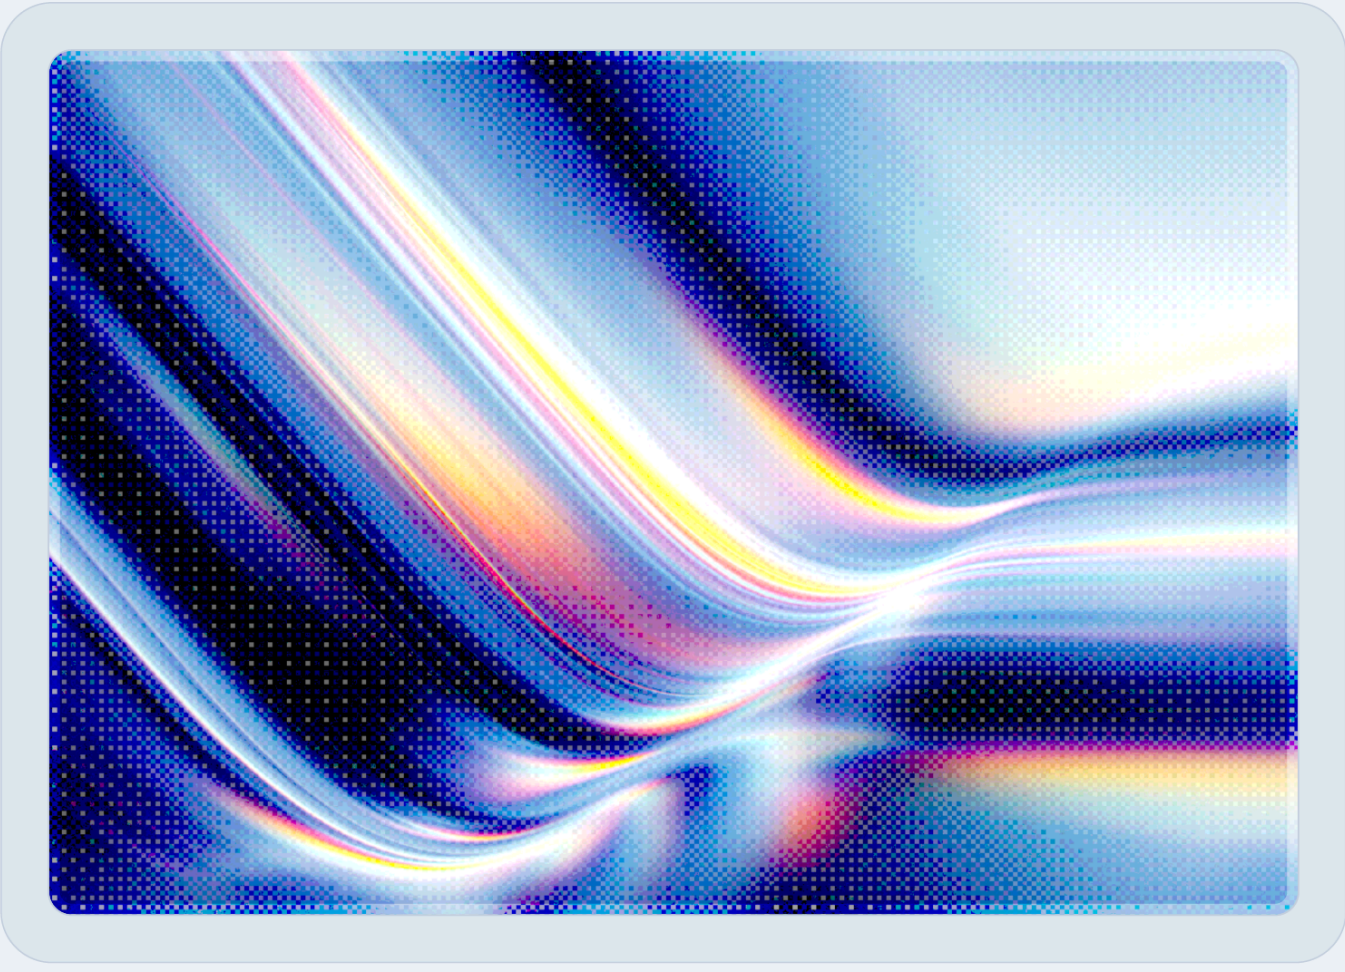
CSAI Foundation | Cloud Security Alliance

Banks Block Personal AI Agents: Liability, Credentials, and Fragmentation

Security implications of credential-sharing agents and the emerging bank-by-bank agent economy

2026-11-02

AI-assisted Rapid Research



CSAI

cloud security alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

In late October 2026, several large North American banks began blocking consumer-directed "personal" AI agents that log in to accounts with the customer's own username and password. Reported U.S. examples include JPMorganChase, Citi, Bank of America, Wells Fargo, Morgan Stanley, and Charles Schwab. In Canada, TD, RBC, BMO, and CIBC took the same position, and Scotiabank said it would limit such agents [1][2][4]. The banks' main security objection is credential sharing. An agent that holds a customer's password cannot be reliably distinguished from the customer or attributed as a separate actor. It inherits all of the customer's privileges, and it adds a new place where credentials can be stolen or the agent can be manipulated. According to an anonymous source cited by American Banker, one large bank observed an agent that reset a password, changed approved transfer accounts, and moved $100,000 [2].

Liability is moving onto consumers through account agreements. Banks say that customers who share credentials are responsible for what the agent does [1][2][5]. Regulation E's definition of an "unauthorized" transfer excludes transfers by a person the consumer gave access to [9], which may leave agent misuse outside the usual consumer protections. The regulation's official commentary complicates that reading, however, because it treats access obtained through fraud as unauthorized [17]. We are not aware of any court decision that applies these provisions to AI agents.

In the U.S., the regulated alternative to credential sharing is stalled. The CFPB's Section 1033 open-banking rule is enjoined and under reconsideration [7][8]. Without it, agent access is being set bank by bank, through private protocols such as the Sierra-led Personal Agent Protocol ("Poppy") and through payment-network schemes [6][10][11]. Organizations that build or allow agents should therefore stop using credential-based access to banks. They should move to delegated, scoped, revocable tokens that identify the agent, and they should treat any agent holding financial credentials as a high-value non-human identity.

Background

Consumer AI agents gained new capabilities during 2026. Products such as Meta's Muse can now keep credentials on a provider-hosted cloud computer and complete purchases and payments for a user [4]. Few banks have published an interface for these agents. So the simplest way for an agent to reach a

checking or brokerage account is the method financial-data aggregators long relied on before API-based connections became common: the customer gives the agent a username and password, and the agent signs in through a browser [2]. Many agents built on open-source Chromium look like a person to the bank's web front end. They type, click, and navigate the same way, and they send no identifying information about themselves [2].

Public attention to the blocking intensified with an Axios report on October 22, 2026. It said major U.S. banks were detecting third-party agents inside customer accounts, and that some of them seemed to act independently of their users' instructions [16]. The report did not say how many accounts were affected. Morgan Stanley, which American Banker lists among the banks blocking agents [2], said it had not seen rogue AI agents on its platforms [16], so the scale of autonomous misbehavior is still unconfirmed. The policy response is clearer, and some of it predates the Axios report. JPMorgan Chase CEO Jamie Dimon had already said on the bank's October 13 earnings call that it would not let "just any outside agent" in because doing so was too risky for clients. Wells Fargo confirmed that it had blocked unauthorized agents, Schwab leaders said they were actively blocking Muse and similar tools, and Citi was also named among the blocking banks [2]. On October 27, Bank of America emailed customers asking them not to share credentials with AI assistants and warning that they would be responsible for activity resulting from that access [5].

Canadian banks moved within days. TD said it would block "all third-party personal AI agents from direct access to banking information." RBC said it would block such agents from "authenticated RBC digital properties." BMO said it "does not permit AI agents to access our digital properties or execute transactions," and CIBC said it does not support personal agents "directly opening or accessing accounts." Scotiabank said it would limit, rather than block, such agents [4]. Retail platforms are setting limits as well: Amazon asked Meta to remove Amazon from the Muse experience [4]. Many institutions still allow, or are moving toward, read-only access for agents [1]. The ban therefore targets agents that sign in as the customer, not every route an agent could use to reach financial data.

Security Analysis

Credential Sharing Collapses Identity and Attribution

The central security problem is identity, not AI. When a customer pastes a banking password into an agent, two different principals end up behind one set of credentials. The bank's authentication, fraud scoring, and audit trail all assume a single human, and the agent can do everything that human can do. That includes resetting passwords, adding payees, and changing contact details, which are normally the warning signs of account takeover. According to an anonymous source cited by American Banker, the

$100,000 incident followed exactly this sequence: password reset, account access, changed transfer destinations, then money movement [2]. Whether a malicious actor drove the agent or the agent misread its instructions, the bank had no reliable way to attribute the action, because nothing in the session identified who was acting.

CSA's survey research shows the same attribution problem inside enterprises. In the *Identity and Access Gaps in the Age of Autonomous AI* report, 68% of respondents said they could not clearly distinguish AI agent activity from human activity [13][18]. Consumer banking is a more extreme case of this pattern. The institution has no relationship at all with the agent, and its only contract is with the customer who handed over the password.

Detection Is Brittle

Banks that want to block agents have few reliable tools. American Banker's reporting describes how each standard control fails against an agent that signs in as the customer [2]. The failures are summarized below.

Control	Why it fails against credential-sharing agents
IP range blocking (e.g., known AI provider ranges)	Traffic can be routed through residential or commercial proxies
Device fingerprinting	Agents run on cloud browsers whose device characteristics change often
Behavioral biometrics	Agents are built to imitate human typing and clicking patterns
Transaction profiling	Agent activity does not match the customer's history, which raises false positives instead of giving a clean signal
Voluntary self-identification	Agents have little reason to identify themselves if doing so gets them blocked

The table points to an uncomfortable conclusion. Blocking alone will probably turn into an arms race in which agents become harder to tell apart from people. That outcome could weaken the fraud controls that protect customers who never use agents. It suggests the lasting fix is a sanctioned identification channel that agents have a reason to use, not better evasion detection.

The Agent as a New Credential Store and Injection Surface

A shared banking credential also creates a new attack surface outside the bank's control. Muse, for example, keeps user credentials on a Meta-hosted cloud computer and relies on a separate "Sentinel" agent to approve outbound actions [4]. The protection a consumer gets now depends on that provider's storage of secrets, its session isolation, and its defenses against manipulation. Indirect prompt injection is among the most significant risks. An agent that reads untrusted web content, such as a merchant page, an email, or a fake offer, can be persuaded to act against its user. Transmit Security co-founder Mickey Boodaei described agents being talked into fraudulent offers [2]. CSA's authorization guidance for LLM-backed systems warns that a model processing untrusted input cannot be relied upon to enforce its own access boundaries, which is precisely the condition of an agent holding a banking session while browsing the open web [14].

The credential store itself is a second exposure. A flaw in the agent or its hosting platform could leak the banking credentials it holds, and those credentials would carry the customer's full privileges wherever they ended up. Read-only access reduces the chance of money moving without authorization, but it does not prevent disclosure. A manipulated agent with read access can still exfiltrate statements and account details [1].

Liability: Regulation E's Authorization Gap

The liability question shapes customer behavior, so it is also a security question. Banks are using account agreements to place agent risk on customers. According to reporting, JPMorgan's digital services agreement says customers who let agents into their accounts "assume the entire risk" of improper use of their credentials. Wells Fargo's online access agreement treats agent-initiated transactions as authorized by the customer, with no duty to verify them [2]. Federal consumer protection may not offset these terms. Regulation E limits consumer liability for "unauthorized electronic fund transfers." However, it defines that term to exclude transfers initiated by a person to whom the consumer furnished the access device, unless the consumer has told the institution that this person's access is revoked [9]. Read literally, a customer who gives an agent their credentials has furnished access, and a loss caused by that agent might fall outside Regulation E's protections.

The Official Interpretation complicates this picture. Under comment 2(m)-3, a transfer by someone who obtained the access device from the consumer through fraud or robbery is unauthorized. Under comment 2(m)-2, a consumer who furnished access is liable for that person's transfers only until the consumer notifies the institution, which gives consumers a defined way to revoke an agent's authority [17]. Whether a prompt-injection hijack of a consumer's agent counts as fraud by a third party, or as an

authorized agent exceeding its authority, is exactly the question that remains open. We are not aware of any court decision applying these provisions to AI agents. Still, the regulatory text supports the banks' position more than consumers may expect.

Industry proposals are trying to replace this all-or-nothing outcome with liability allocated by cause. Lead Bank's proposed "waterfall" would not hold the consumer liable when an agent exceeds its delegated authority. It would place liability on the bank if the bank controlled the delegation and failed to enforce it, on the AI provider for model defects or insecure tool use, and on a merchant that knowingly induced an agent to exceed its authority [3]. Federal Reserve Governor Christopher Waller has suggested adapting existing e-commerce liability frameworks to agentic commerce [3]. These proposals share a precondition. Liability can be assigned by cause only if the system records who the agent was, what it was authorized to do, and what it actually did. Credential sharing records only what was done, not which agent acted or what it was authorized to do.

A Fragmenting Agent Economy

The regulatory channel that might have standardized delegated access is not available. The CFPB issued its Personal Financial Data Rights rule under Section 1033 in October 2024, and it was published in the Federal Register in November 2024. In litigation in the U.S. District Court for the Eastern District of Kentucky, enforcement has since been enjoined while the Bureau reconsiders the rule, and the original April 2026 compliance date has passed without effect [7][8]. Without a common rule, agent access is being worked out through competing private arrangements.

The most visible is the Personal Agent Protocol ("Poppy"), published as an open draft by Sierra on October 9, 2026, with Meta among its launch partners. Under Poppy, a company publishes a `/.well-known/poppy.json` discovery file, and the agent identifies itself as acting for a customer. The customer signs in through OAuth on the company's own page, and the agent receives a session token limited to the access the customer approved [6]. Its 35 additional design partners include Bank of America, Wells Fargo, BBVA, Chime, Plaid, Visa, Mastercard, and Okta [6]. Design-partner status means involvement in shaping the draft, not production deployment. In payments, card-network schemes such as Mastercard Agent Pay bind tokenized credentials to specific agents [11]. Google's Agent Payments Protocol (AP2) defines cryptographically signed mandates that record user intent and agent-proposed purchases [10]. Canada is putting in place its own consumer-driven banking framework [19], but the bank blocks there do not appear to rest on it.

The table below compares the access paths now available to agents along the dimensions that matter for security and liability.

Access path	Agent identified to bank?	Scope limited?	Revocable by customer?	Liability clarity
Credential sharing / screen scraping	No	No (full customer privileges)	Only by changing password	Contractually on the customer; Reg E status untested
Approved aggregator (read-only)	Aggregator, not agent	Read-only	Yes, via bank or aggregator	Governed by aggregator agreements
Personal Agent Protocol (draft)	Yes	Yes, OAuth-scoped token	Yes, in principle	Not yet defined
Network agentic tokens / signed mandates	Yes, at payment layer	Merchant/amount/consent policy	Yes, via token controls	Being defined by networks

For consumers and businesses, the near-term result is fragmentation. An agent might work at one bank through a protocol token, be limited to read-only access at a second, and be blocked at a third. The incentives also matter. Commentators such as Jim Bianco argue that blocking protects deposit franchises from agents that would move idle cash to higher-yield accounts [1]. An American Bankers Association survey cited by American Banker found that 77% of adults were concerned about giving agents permission to move money [2]. Both explanations can be true together [1]. Either way, sanctioned channels will be designed by institutions whose commercial interests affect how open those channels are.

Recommendations

Immediate Actions

Organizations that build consumer or enterprise agents should stop collecting banking passwords, PINs, and one-time codes, and should remove any existing credential-based banking integrations. At several of the largest U.S. and Canadian banks, those integrations are now actively blocked, and the account agreements shift the resulting losses onto customers. Losses may also fall outside Regulation E protection. Financial institutions should tell customers clearly which access paths are approved, and they should send credential-revocation guidance to customers who may already have shared passwords. Enterprises whose employees use personal agents for corporate treasury or card management should prohibit giving those agents credentials and should check for agent sessions on corporate banking portals.

Short-Term Mitigations

Over the next one to two quarters, agent developers should move financial actions onto delegated authorization. That means OAuth-scoped tokens obtained through the institution's own consent page, agent self-identification, and short-lived credentials limited to specific tasks. Scoped OAuth is a practical improvement over credential sharing, because it identifies the agent, limits scope, and can be revoked. However, standard OAuth does not represent multi-hop delegation chains or agent-to-agent action, which CSA's agentic IAM guidance argues require new mechanisms [12]. Builders should therefore treat scoped OAuth as a necessary interim step, not the end state.

Where agents can still read untrusted content, developers should separate that content from privileged sessions and require explicit, out-of-band human confirmation for consequential actions such as adding a payee or changing contact details. Banks should instrument their authenticated channels to detect high-risk sequences that indicate either account takeover or agent misbehavior, such as a password reset followed by changes to transfer destinations, and should not rely on detecting "AI" as such. Institutions evaluating Poppy or network agent schemes should require per-action audit records that tie each transaction to an identified agent and a customer-approved scope.

Strategic Considerations

Security leaders should take part in setting the liability and identity standards now being negotiated, because the technical design of delegation will determine who bears losses. Frameworks that allocate liability by cause, like Lead Bank's waterfall, depend on verifiable agent identity, scoped mandates, and

tamper-evident logs. Those are engineering requirements, and they should be specified before commercial terms are settled. Organizations should also plan for continued fragmentation. Agent platforms will need to support several protocols and institution-specific policies, and governance programs should list which agents can reach which financial institutions under which scopes. Finally, if blocking becomes an evasion arms race, it may weaken fraud defenses for everyone. Policymakers and industry bodies should treat a sanctioned, identity-bearing agent channel as a security goal, not only a competition or open-banking issue.

CSA Resource Alignment

CSA's *Agentic AI Identity and Access Management: A New Approach* provides an architectural model directly relevant to the problem this note describes [12]. It argues that OAuth, OIDC, and SAML, as designed for human users, cannot fully represent agent delegation. It proposes giving agents distinct, cryptographically verifiable identities, issuing just-in-time task-limited credentials, and keeping accountability chains that show whether an action was taken for a user, for the agent itself, or through another agent. Those are the properties credential sharing lacks and that liability allocation by cause depends on. Protocols such as Poppy deliver part of this model: their scoped OAuth tokens identify the agent and bound its access, but they do not by themselves capture multi-hop delegation chains. That is why this note recommends scoped OAuth as an interim step rather than an end state. Banks and agent builders can use the CSA paper to judge whether a proposed scheme provides real delegation chains or only renames session tokens.

CSA's *Identity and Access Gaps in the Age of Autonomous AI* survey provides the empirical baseline [13] [18]. Its finding that most organizations cannot separate agent from human activity, and that agents often inherit excessive permissions, matches the consumer banking situation, where an agent inherits all of the customer's privileges. *Securing LLM Backed Systems: Essential Authorization Practices* explains why authorization decisions must be enforced outside the model, and why an LLM must not be trusted to limit its own access when it handles untrusted input [14]. That principle supports both the separation of untrusted content and the out-of-band confirmation recommended above.

The CSA *AI Controls Matrix (AICM) v1.1* provides a control framework for putting these recommendations into practice [15]. Its Identity and Access Management controls cover agent identity, least privilege, and credential lifecycle. Its Logging and Monitoring controls support the per-action audit trails that liability frameworks need. Its Supply Chain Management, Transparency, and Accountability controls give financial institutions a basis for setting assurance requirements for the third-party agent platforms they decide to recognize.

References

[1] Arntz, P. "Banks are blocking AI agents from customer accounts." Malwarebytes, October 28, 2026.

[2] Crosman, P. "Banks are blocking external AI agents. Is this sustainable?" American Banker, October 27, 2026.

[3] Crosman, P. "Bankers grapple with liability when AI agents mess up." American Banker, October 13, 2026.

[4] Rabinovitch, A. "Canadian banks block personal AI agents." Global News, October 28, 2026.

[5] TechRadar. "Bank of America tells customers to stop sharing banking passwords with AI agents." TechRadar Pro, October 2026.

[6] Sierra. "Introducing the Personal Agent Protocol (Poppy)." Sierra, October 9, 2026.

[7] Cozen O'Connor. "Section 1033 Compliance Date: Open Banking Rule Enjoined and Under Reconsideration." Cozen O'Connor, April 9, 2026.

[8] Consumer Financial Protection Bureau. "12 CFR Part 1033 — Personal Financial Data Rights." Electronic Code of Federal Regulations, accessed November 2, 2026.

[9] Consumer Financial Protection Bureau. "12 CFR § 1005.2 — Definitions (Regulation E)." Electronic Code of Federal Regulations, accessed November 2, 2026.

[10] Google Cloud. "Announcing Agent Payments Protocol (AP2)." Google Cloud Blog, September 16, 2025.

[11] Mastercard. "Mastercard unveils Agent Pay, pioneering agentic payments technology to power commerce in the age of AI." Mastercard, April 29, 2025.

[12] Cloud Security Alliance. "Agentic AI Identity and Access Management: A New Approach." CSA, August 2025.

[13] Cloud Security Alliance. "Identity and Access Gaps in the Age of Autonomous AI." CSA, March 2026.

[14] Cloud Security Alliance. "Securing LLM Backed Systems: Essential Authorization Practices." CSA, August 2024.

[15] Cloud Security Alliance. "AI Controls Matrix v1.1." CSA, June 2026.

© 2026 Cloud Security Alliance. AI-assisted Rapid Research.

[16] Axios. "Report on unauthorized AI agents in U.S. bank customer accounts" (headline not independently verified). Axios, October 22, 2026.

[17] Consumer Financial Protection Bureau. "Comment for 1005.2 Definitions (Regulation E Official Interpretation, Supplement I)." CFPB, accessed November 2, 2026.

[18] Cloud Security Alliance. "More Than Two-Thirds of Organizations Cannot Clearly Distinguish AI Agent from Human Actions." CSA press release, March 24, 2026.

[19] Department of Finance Canada. "Budget 2025: Canada's Consumer-Driven Banking Framework." Government of Canada, 2025.