



2026

KNOW YOUR AGENT

Why Banks Must Govern AI Agents at the Front Door

This field guide explains how banks can identify, govern and audit customer-owned AI agents before they act.



www.ovationcxm.ai

The Know Your Agent imperative

Customer-owned AI agents are moving from research to execution at a dizzying pace. What started as chat-based customer queries is evolving into autonomous agents that can execute banking tasks, ranging from balance inquiries and account servicing requests to more consequential actions, including initiating payments, moving treasury funds and executing lending requests on behalf of corporate clients.

Consumer adoption of personal AI agents like Meta Muse and agent platforms like Instinct is rapidly surging.¹ These agents are already disrupting retail finance by comparing rates, products and services, and automating real-world purchases. Business owners and corporate operators naturally expect the same seamless, agent-driven execution across their commercial banking workflows. However, with verification challenges mounting, allowing external agents to connect cannot mean granting unmonitored access. This shift creates significant commercial opportunity alongside operational and compliance risks, requiring banks to deploy a robust Know Your Agent (KYA) governance framework to authenticate, permission, govern and audit all machine interactions.

OvationCXM's 2026 *When AI Agents Become the Customer* report of 522 U.S. finance leaders found that 88% would use agentic AI in banking if it were secure and compliant, and 69% would pay for bank connectivity that enables AI-initiated workflows.² Crucially, in another OvationCXM survey, 43% said they would seek workarounds if their current institution could not support agent-initiated activity.³

That creates a new challenge for banks: Simply blocking agent-initiated activity may not be a viable long-term strategy, but allowing an agent to connect cannot mean allowing it to act without oversight.

This rapid evolution of agentic banking creates both risk and opportunity. Banks that fail to establish clear controls risk losing visibility into agent activity and control over customer relationships. Those that act now can turn agentic activity into a governed, auditable and potentially revenue-generating channel.

This Guide Covers

- The rise of agentic banking and the urgency to govern the front door
- The KYA model and its foundational requirements
- How the KYA lifecycle works
- How to operationalize KYA governance at scale

Three Principles In This Guide

- **Non-human identities must have accountable owners.** Every agent action must trace back to an authorized human or corporate entity and the customer or corporate entity the agent represents.
- **Runtime controls are required.** Real-time validation of identity, authority, purpose and scope is essential to manage agent-initiated activity safely.
- **Auditability must be built in.** Complete records are necessary for compliance, investigation and regulatory scrutiny.



KEY TAKEAWAY

Banks need more than agent identity. They need a way to establish ownership, evaluate authority in real time and preserve a complete record of every agent-initiated action.

The rise of agentic banking and the urgency to govern the front door

Customer-owned AI agents are no longer a hypothetical concept. These autonomous agents are starting to initiate banking activity on behalf of corporate clients. OvationCXM's 2026 agentic AI readiness survey found that nearly two-thirds (64%) of U.S. finance leaders plan to use agentic AI in at least two finance workflow areas within two years.² Another study showed that 75% of organizations are already testing or using AI agents for finance and banking activities.³

For banks, this creates a fundamentally different type of customer interaction. A human typically arrives through a bank-controlled channel with an established identity, role and set of permissions. A customer-owned AI agent may originate outside the bank, operate across multiple systems and make requests at machine speed.

More broadly, emerging industry research is converging on the need to establish identity, authority, oversight and controls for non-human actors. But written policies alone do not determine whether an individual request should proceed.^{4,5} Those controls must be enforceable when the agent interacts with the bank.

The question for banks therefore shifts from "Do we allow AI agents?" to "Under what conditions can this agent perform this action for this customer?"



ACTIONS TO GET STARTED

- Identify gaps between existing identity frameworks and the requirements for governing non-human actors.
- Assess current exposure by mapping which customer segments are most likely to deploy AI agents in the next 12 to 18 months.
- Engage with risk and compliance teams early to define acceptable risk thresholds for agent-initiated activity.
- Map how the bank currently determines identity, authority and permissions when a request originates outside a traditional human-controlled channel.



KEY TAKEAWAY

Banks cannot treat agent access as a simple yes-or-no decision. Governance must operate at the level of the agent, the customer it represents and the action it is requesting.

The KYA model

Just as Know Your Customer (KYC) rules are required to establish a commercial banking relationship, KYA governance is required before letting autonomous AI agents transact on a client's behalf. Traditional identity and access management frameworks were largely designed around human users and established machine-to-machine connections. They lack the real-time controls, ownership linkage and auditability required when external agents initiate financial transactions on behalf of corporate clients.

Identity alone therefore does not establish authority. A bank may know what an agent is and still need to establish who it represents, what it has been authorized to do and whether that authority applies to the specific request being made.

This distinction is central to KYA.

Connectivity answers whether an agent can reach the bank. KYA governance helps determine whether the agent can act once it arrives.

Creating a broadly accepted standard for agents to identify themselves is a critical industry imperative. According to technologist Nikita Bier, it is "the most important technology imperative of the next 5 years."⁶ Institutions face a "cat-and-mouse game" as agents try to circumvent detection to maintain their utility.

KYA governance allows organizations to move beyond playing defense and address these gaps by treating every AI agent as an accountable non-human actor. It establishes foundational requirements that together create a practical governance framework.

When a customer-owned agent arrives, the bank should be able to answer a core set of questions before acting on its request:⁷

- **Identity:** Which agent is making the request?
- **Representation:** Which customer or corporate entity does it represent?
- **Ownership:** Who is accountable for the agent?
- **Authority:** What has the agent been authorized to do?
- **Context:** Does that authority apply to this specific request, amount, account and circumstance?
- **Control:** Should the request be allowed, routed for approval, escalated or blocked?
- **Audit:** Can the bank reconstruct what happened, why and under whose authority?

These questions move KYA beyond agent identification. They connect identity to delegated authority, then apply that authority to the requested action. This approach moves banks beyond static policies toward enforceable, real-time controls that can keep pace with the speed and autonomy of agentic banking.

Minimum KYA registration record

Before a customer-owned AI agent can act, the bank should maintain a registered record that establishes who the agent is, who it represents and the boundaries of its authority. At a minimum, that record should include:

- Agent identifier
- Named owner or accountable party
- Customer or corporate entity represented
- Authorized purpose
- Credential or authentication method
- Permitted tools, accounts and action scope
- Operating environment
- Agent fingerprint
- Current status, including active, suspended or revoked

Higher authority and greater operational dependence should trigger stronger controls, evidence requirements and oversight.

Fingerprint the agent, not just the credential

Credentials establish how an agent authenticates, but they do not necessarily distinguish one agent from another. The bank should maintain a persistent fingerprint that allows it to recognize the specific registered agent making a request, even when multiple agents operate for the same customer or use similar authentication methods.

This gives the bank a durable identity reference that can be evaluated alongside ownership, credentials, scope and request context at runtime.



ACTIONS TO GET STARTED

- Define ownership requirements for all customer-owned AI agents.
- Establish real-time validation standards for identity, representation, authority, purpose, credentials and scope.
- Define which requests can proceed automatically and which require approval, escalation or denial.
- Implement revocation capabilities that do not require manual engineering intervention.
- Create audit records that risk and compliance teams can access directly.



KEY TAKEAWAY

Knowing the agent is not enough. Banks must connect agent identity to customer authority and evaluate that authority against each requested action in real time.

How the KYA lifecycle works

KYA governance operates as a continuous lifecycle, beginning with initial agent registration and approval as the first operational phase before moving on to runtime transaction enforcement.⁸

Separating these two phases ensures that the autonomous agent is unable to attempt live transactions until its owner is verified and its technical credentials are validated.

- 1 **Agent arrives:** A customer-owned AI agent initiates a request.
- 2 **Identify + fingerprint:** Confirm which specific agent is making the request and match it to its registered identity.
- 3 **Resolve owner + customer:** Determine who is accountable for the agent and which customer or corporate entity it represents.
- 4 **Validate credential + scope:** Check that the credential is valid and that the agent's registered permissions cover the requested capability.
- 5 **Evaluate the requested action:** Apply policy to the actual request, including account, amount, purpose, workflow state and any required approvals.
- 6 **Allow, route, escalate or block:** The bank makes a deterministic execution decision outside the model.
- 7 **Write the audit record:** Record the agent, owner, customer, request, policy decision, approvals, downstream result and timestamp.

Revocation must be immediate and granular:

- Revoke one agent
- Suspend or revoke a credential or permission scope
- Revoke all agents associated with a customer or corporate entity
- Prevent revoked agents from executing future requests immediately



KEY TAKEAWAY

KYA continuously validates agent identity, ownership, credentials and scope - backed by instant revocation and an examiner-ready audit trail.

How to operationalize KYA governance at scale

Implementing KYA at scale requires more than policy updates. Banks must embed governance into existing processes while building the technical capability to authenticate agents, enforce boundaries, revoke access when necessary and produce audit-ready records without burdening engineering teams.

Operationally, that means treating agent requests as governed work, not simply API traffic. Successful operationalization typically involves three practical steps:

Step one: Establish the agent and its authority. Banks must integrate KYA requirements into customer onboarding and identity management processes so agent governance becomes part of standard workflows, not a separate initiative. The agent should be linked to the customer it represents, an accountable owner and clearly defined permissions.

Step two: Evaluate each request at runtime. Banks need real-time policy enforcement capabilities that can validate identity, authority, purpose and scope at the moment of interaction. A balance inquiry, an account information change and a wire initiation should not necessarily follow the same decision path simply because they originate from the same agent.

Step three: Orchestrate the appropriate response and preserve the record. Banks need the ability to allow an authorized request, route activity requiring human approval, escalate exceptions or block prohibited activity while maintaining complete audit and reporting records.

KYA does not end when agents pass registration and validation.

Agent governance should continue dynamically after access is granted. Banks should monitor how an agent is actually being used, including request patterns, overrides, exceptions, customer impact, failure modes and changes in practical authority. Changes to an agent's owner, purpose, credential, permissions, behavior or operating environment should trigger re-review. The same should apply when an agent begins performing activity materially different from what was originally approved.

Internal AI and inference engines can be deployed to analyze incoming requests and interactions from external agents in real time. These models can detect subtle anomalies, such as scope creep or unauthorized automation shifts, based on their broader evaluation of reasoning patterns and behavioral baselines - and they can do so before the agents impact internal systems. Any material changes can automatically trigger a formal re-review.

The objective is to prevent trust and authority from expanding informally simply because an agent has operated successfully over time. This is where governance becomes operational. The bank is not merely authenticating the agent at the front door. It is governing what that agent can do throughout the interaction.

Banks that take these steps will be better positioned to support corporate clients that want to delegate financial activity to their own agents while maintaining control over the relationship and meeting regulatory expectations.

What risk and compliance should be able to see

For every material agent request, the bank should preserve a record that includes:

- Agent ID and fingerprint
- Accountable owner
- Customer or corporate entity represented
- Requested action
- Account or resource involved
- Permission or scope evaluated
- Policy decision
- Approval or escalation, if required
- Downstream execution result
- Timestamp
- Agent and credential status at the time of execution

Risk, Compliance or Audit should be able to reconstruct an agent-initiated action without submitting an engineering request.



ACTIONS TO GET STARTED

- Pilot KYA controls with a limited set of high-value corporate clients to test workflows and surface operational gaps.
- Integrate KYA requirements into existing customer onboarding and identity processes.
- Establish a cross-functional working group with internal risk, compliance, technology and product leaders to define ownership, authority, approval, revocation and audit requirements.
- Evaluate current technology capabilities against the ability to deliver real-time validation and revocation at scale.
- Test multiple outcomes, including allowing, routing, escalating and blocking requests, rather than designing only for successful agent requests.



KEY TAKEAWAY

KYA becomes meaningful when banks can translate identity and delegated authority into an enforceable decision for every agent request.

Conclusion

As agentic banking accelerates, the institutions that establish clear ownership, real-time controls and auditable records will be best positioned to lead.

But the defining challenge is not simply recognizing an AI agent. It is determining what that agent is authorized to do once it reaches the bank.

Customer-owned AI agents will arrive from ERPs, treasury platforms, accounts payable and accounts receivable systems, general-purpose AI platforms and other environments the bank does not control. Banks will need infrastructure that can connect agents to the customer it represents, evaluate the authority behind each request and coordinate the appropriate response across systems, teams and approval processes.

Protocols and connectivity standards can help agents reach the bank. Governance determines whether they can act.

OvationCXM helps banks operationalize KYA governance through its Agentic Transaction Control Layer (ATCL), which provides a control and orchestration layer for agent-initiated activity. ATCL enables banks to evaluate agent identity, customer representation, authority, permissions and request context, then allow, route, escalate or block the request while preserving a complete audit trail.

For banks, the opportunity is not to open or close the front door to AI agents. It is to control what happens after they arrive.



Know the agent.
Govern the action.



THE OPPORTUNITY

For banks, the opportunity is not to open or close the front door to AI agents. It is to control what happens after they arrive.

References

Sources published in 2025-2026 were used wherever possible. Older sources are retained only when they remain the best or primary authority for a specific concept.

1. Forbes. (2026, September 30). Meta's Muse AI Assistant Hits 5 Million Downloads, Outpacing Growth of ChatGPT, Grok and Claude. <https://www.forbes.com/sites/maryroeloffs/2026/09/30/metas-muse-ai-assistant-hits-5-million-downloads-outpacing-growth-of-chatgpt-grok-and-claude/>
2. OvationCXM. (2026, June). When AI Agents Become the Customer. <https://www.ovationcxm.com/agentic-banking-report>
3. OvationCXM. (2026, September). AI Agents: From Experiment to Execution. <https://www.ovationcxm.com/agentic-banking-research-ai-agents-from-experimentation-to-execution>
4. Gartner. (2026, September 24). Agentic AI Fails Where Governance Stops. <https://www.gartner.com/en/articles/agentic-ai-infrastructure-governance>
5. Forrester. (2026, February 24). Agentic AI Is on the Cusp of Transforming Financial Services. <https://www.forrester.com/blogs/agentic-ai-is-on-the-cusp-of-transforming-financial-services/>
6. Bier, N. (2026, October 7). Post on X. <https://x.com/nikitabier/status/2107157904168239416>
7. OvationCXM. (2026, September). Prepare for the Agentic Customer: A Field Guide to Governing Customer AI Agents Across Banking and Payments. <https://www.ovationcxm.com/ebooks/prepare-for-the-agent-customer>
8. Know Your Agent. (2026). Know Your Agent: The Governance Standard for AI in Financial Services. <https://www.kyabook.com/>



The orchestration system for agentic banking

Schedule an ATCL briefing at
www.ovationcxm.ai/atcl.

OvationCXM provides the orchestration operating system for financial services. Its CXMEngine platform connects systems, data, teams and ecosystem partners to coordinate complex customer and operational workflows without requiring institutions to replace existing technology. OvationCXM's Agentic Transaction Control Layer (ATCL) extends that orchestration to client-facing AI agents, enabling institutions to register, authenticate and permission incoming agents and coordinate their approved activity across systems, workflows, partners and human approvals. Learn more about OvationCXM at www.ovationcxm.ai.