OvationCXM

2026

PREPARE FOR THE AGENTIC CUSTOMER

A field guide to governing customer AI agents across banking and payments

www.ovationcxm.ai

Get your bank ready for customer agents

How to stand up a governed front door for client-owned AI — using the portal, APIs and banking core you already have.

This field guide contains nine actionable steps to take you from start to finish in getting banks and enterprises ready to work with client-owned AI agents. This guide will cover:

- What "customer agent" means on the bank side (and what it does not)
- How to reuse treasury entitlements, APIs and approval paths
- How to scope a Phase 1 deployment and prove control
- What compliance will ask for before money moves

Who Is This For?

- Bank Leadership
- Product Teams
- IT Teams
- Compliance Teams
- AI Oversight and Risk Management Teams
- Digital Transformation Teams

Table of Contents

02 Align on what you are preparing for

03 Treat this as a new channel, not a science project

04 Inventory what you already have

05 Specify the governed front door

06 Start safe. Expand as controls are proven

08 Reuse human controls; extend them to non-humans

09 Get compliance in the design, not at the end

10 Measure the first deployment

11 How OvationCXM can help

Align on what you are preparing for

Clients are already using AI in financial workflows. In the OvationCXM 2026 Agentic AI Readiness Survey (n=522 U.S. finance leaders), 76% already use automation or AI for banking-adjacent work — 62% automation, 14% AI. This is happening right now. It is a present-day reality, not a future concern.

These agents are already embedded in the systems finance teams rely on daily, including QuickBooks/Intuit, SAP Joule, Oracle Fusion ERP, Workday, Salesforce Agentforce, ServiceNow and others. As commercial clients often run several of these platforms at once, banks end up dealing with a fragmented patchwork of vendor agents, each bringing its own identity model, permissions and definition of what it means to "talk to the bank."

Finance teams rely on much more than traditional ERP suites. General-purpose AI tools currently lead usage — led by ChatGPT at 39.4%, alongside Gemini and Microsoft Copilot — followed by embedded product agents like QuickBooks (24.5%), Workday (19.3%), Oracle (16%) and Microsoft Dynamics (13.9%). In the near term, most inbound banking requests will come through ChatGPT and Copilot performing simple read-only tasks like checking balances, history and status, rather than through purpose-built banking agents. Waiting to build a dedicated "banking agent" means missing client traffic already arriving today.

Most clients want to use agentic AI for their banking, provided it is secure and compliant. In fact, 88% would allow an AI agent to prepare, recommend or execute tasks, and 54% would let it run full workflows automatically. This is already driving business: 71% consider AI capabilities when choosing or expanding a bank, 51% view AI support as critical and 69% are willing to pay for direct connectivity.

However, clients do not want total autonomy — they demand strong controls and oversight across every task. Because clients expect AI tools to connect seamlessly with their bank, financial institutions have a key opportunity in the next planning cycle to lead this channel while expectations are high.

Customer agents are client-owned software applications — running in ERP, treasury systems, chat and AI workspaces — that initiate banking tasks under an authorized user's authority. This is fundamentally different from a bank-provided chatbot, a new mobile app or an agent directory. Instead, the bank's role is to establish a secure, governed entry point to receive, authenticate, monitor and execute these incoming requests without granting direct access to core backend systems. If banks fail to establish and control this channel, external platforms like ERPs, treasury management systems or fintechs will take over the client relationship instead.

KEY TAKEAWAY

88%

of finance leaders would allow an AI agent to prepare, recommend or execute banking tasks

Treat this as a new channel, not a science project

Treasury portals, mobile apps, file transfers, APIs and SWIFT represent existing customer channels; client-owned AI agents are simply the next evolutionary interface channel. To prepare, Product, IT, digital and AI leaders must appoint a dedicated owner for Agent Experience (AX), mirroring how they manage digital Customer Experience (CX). Leadership must also align on prioritizing controlled access as the immediate outcome, rather than fully autonomous transactions.

Treat AI agents like a core banking channel — such as online, mobile or host-to-host banking — rather than a lab experiment. Just as those legacy channels required dedicated ownership, entitlement models, approval workflows, auditing and kill switches, agent traffic needs these same five operational pillars before accessing account balances.

Assign clear ownership this quarter. Designate one leader with overall accountability for the channel and define specific roles for the rest. For instance, Product should manage the agent experience (AX) — determining visible tasks, completion criteria and success metrics. IT should control identity, credentials and connectivity (e.g., via API, MCP or agent portal), deciding if requests reach core systems. Digital transformation must oversee the operating model across treasury, digital and servicing to ensure AX is fully integrated, not a CX afterthought. Finally, AI leaders must manage model and tool risk, defining allowable agent calls and protocols for handling model drift.

Define your initial goal in a single sentence: establish controlled access for a specific agent running a specific workflow. If you cannot identify the system owner, the targeted workflow, the designated human with authority to reject the request and the audit log to document it, you are not ready to open the channel.

If the bank fails to build this governed entry point, clients will send AI agents anyway. They will bypass official channels using shared portal credentials, fragile RPA scripts on the treasury interface or third-party fintechs embedded in their ERP. Because this automated traffic will look like a standard user, you won't know which agent acted, under whose authority or how to revoke its access. As a result, the ERP, treasury management system or software platform becomes the primary client layer — leaving the bank with all of the risk and none of the relationship.

ACTIONS TO GET STARTED

This week, designate an Agent Experience (AX) owner, inventory every existing entry point a machine can access — such as portal logins, API keys or SFTP — and specify which ones agents will be permitted to use during Phase 1. Later this month, elevate agent-channel readiness to the same steering committee package as digital Customer Experience (CX), ensuring Compliance is actively involved. Remember, your first shippable deliverable isn't an AI agent, but a one-page channel charter that explicitly defines the owner, permitted tasks, prohibited tasks and emergency shutdown procedures.

Inventory what you already have

You do not need to replace any core systems or rebuild rules to support incoming client-owned agents. Take stock of what your bank has already spent decades building and extend these secure frameworks to include machine callers. List treasury entitlements, dual control, OTP, dollar limits and audit logs. Split every commercial function into API, file/SWIFT or portal-only. Phase 1 should be simple and low risk. For example, reporting APIs that are already exposed. The core remains the system of record. The Agentic Transaction Control Layer (ATCL) fronts it; it does not replace FIS, Fiserv, Jack Henry or Hogan.

Treat the agent interface as fundamentally distinct from human interfaces. While today's banking APIs serve as system pipes to transmit calls, an agent interface acts as a controlled boundary that defines which specific tasks the bank permits. This usually takes the form of an MCP wrapper over those existing APIs, built with structured Agent Journey steps inside an agent-specific portal. In short, the API provides the connection, while the MCP and journey steps enforce essential guardrails — defining available tools, execution order and required verification evidence. Banks must inventory both and avoid assuming an existing API automatically means an AI agent is authorized to call it.

Run the inventory exercise as an active working session, not an extended research project. Product supplies the commercial task list detailing what clients currently execute in treasury and servicing. At the same time, IT provides the interface map covering every active API, file drop, SWIFT gate and portal URL those tasks use.

Digital transformation contributes to the admin entitlement and approval matrix, and the AI team identifies the candidate tools an inbound agent might attempt to call. Consolidate everything into a single spreadsheet with just four columns: the task, its current interface, the existing human controls and whether it is allowed in Phase 1 (yes, no or later).

Be explicit when defining functions: an "API" strictly means a documented call a machine executes using a bank-issued credential; "File/SWIFT" refers to batch processing rather than conversational tasks; and "Portal-only" signifies that manual human clicks are required. Phase 1 must focus exclusively on API-enabled, read-only functions — such as checking balances, transaction history, statements and reporting. If a client can currently access their balance only through the portal, that task cannot enter Phase 1 until it is formally exposed as an API or dedicated agent task, as screen-scraping should never be treated as an acceptable interface for client-owned agents.

Leverage your existing trusted controls rather than redesigning them for AI agents. Security mechanisms like dual authorization, OTPs, transaction limits and audit logs remain the same — you simply assign a machine entity (the registered agent, linked to a human owner) to them. If a workflow currently demands two human approvers, it still requires two human approvers when an agent initiates it. Finally, if a task lacks a reliable audit trail today, it is not ready for agent automation.

ACTIONS TO GET STARTED

This week, export the treasury entitlement report and flag every permission a non-human identity could potentially inherit. Later this month, share the four-column inventory with the Agent Experience (AX) owner and Compliance team. Exclude from Phase 1 any task that cannot be clearly mapped to both a secure connection and an established control.

Specify the governed front door

Finance leaders have already defined the necessary requirements: fraud and security, configurable permissions, confirmation and reconciliation, audit trails and human override. Adapting these controls for client-owned AI agents requires more than just creating another REST endpoint; it demands a four-part front door:

- **Know Your Agent (KYA):** Register each agent as an accountable non-human user by binding it to a specific legal entity, a named human owner and captured credentials.
- **Policy:** Restrict the agent strictly to published actions using least-privilege tools, read vs. write rules, dollar limits, eligible accounts and time windows.
- **Human-in-the-Loop (HITL):** Ensure agents can only *request* high-impact actions — never silently complete them — by enforcing OTPs, dual approvals or bank reviews whenever policy limits are exceeded.
- **Revoke:** Maintain an immediate kill switch to freeze credentials, halt in-flight activity and pause or shut down a single agent (or all agents for a customer) in seconds.

APIs merely transmit data; banks need a governed front door that evaluates whether a call can proceed, verifies whose authority it is acting under and captures supporting evidence.

Document all four components before engineering designs anything. To execute this, Product first defines the allowable tasks and the verification evidence required for each. IT then approves how KYA data (agent ID, owner, legal entity, credentials and digital fingerprint) is stored and how revocation immediately terminates active sessions.

Next, digital transformation aligns existing treasury approval workflows with agent journeys so Human-in-the-Loop (HITL) actions use current approval paths rather than requiring a separate queue. Finally, inventory every tool an agent could invoke and explicitly classify each as read, write or blocked.

KYA is a living record, not a static spreadsheet — tracking agent ownership, purpose, permissions, activity timestamps and fingerprint verification. If you cannot revoke access from that record in seconds, true KYA is not in place. Furthermore, policy must function as a published menu rather than a prompt; if an unlisted tool is requested, the system automatically denies and logs it. Human-in-the-Loop (HITL) triggers should mirror existing human thresholds, including dollar limits, payment types, new beneficiaries and after-hours activity. Finally, revocation requires three distinct switches: one for an individual agent, another for all agents associated with a customer and a global switch for all agents across your organization. All must be thoroughly tested before Phase 1 launches.

The front door also must see the calls that fail. A refused registration, a scope violation or a replayed credential: Each is a piece of evidence, not noise. If the only records you keep are successful API responses, Compliance should not sign off on the path.

ACTIONS TO GET STARTED

This week, review your inventory's Phase 1 task list and identify which of the four governance components each task is missing. You are officially ready to design the system entry point once a controller can answer five core questions for that task: identity (who is the agent), ownership (who owns it), authorization (what it may do), auditing (what it just did) and control (how to stop it).

Start safe. Expand as controls are proven

Avoid launching with fully autonomous money movement. Instead, target a single mid-market commercial segment — or more importantly, a single workflow — and clearly define what success looks like before you begin building. Start with a read-only use case to prove the concept before moving to higher-impact actions.

1

Phase 1: Read-Only Operations

Focus on read-only tasks like checking balances, viewing transaction history, pulling statements and generating reports. Timestamp and log every request with the agent's identity, the requested task and whether access was granted or denied. Even rejected calls create an audit record so agents blocked at the door remain fully trackable and reportable.

2

Phase 2: Instructed Workflow Assistance

Expand to assisted servicing workflows, including checking task status, managing documents, handling servicing and reconciliation support and initiating customer service requests. Auditing in this phase captures the incoming request, the specific agent journey or task touched, any human who approved an exception and the final outcome of the work item.

3

Phase 3: Controlled Transaction Initiation

Enable high-value operations — such as ACH payments, wire transfers, bill pay, payroll, commercial card management, FX and beneficiary changes — governed strictly by OTPs, dollar limits and dual approvals. Audit records must compile a full Human-in-the-Loop (HITL) package, documenting the OTP delivery channel, the approving individual, the linked payment details and whether the agent's credentials were subsequently revoked.

If a call is refused at the front door, record the agent and the refusal. There can be no silent drops. Compliance should be able to report “blocked at the door” the same way it reports “completed.”

Select a single customer segment to focus on from start to finish. Many banks begin with the mid-market commercial sector because it offers meaningful transaction volume, well-understood entitlement structures and finance leaders willing to pay for connectivity. Focus on deploying a single workflow rather than an entire platform; simple use cases like balance inquiries or statement retrieval are ideal for Phase 1.

Define completion by establishing four explicit criteria across your rollout phases. Phase 1 is complete when a registered agent can successfully pull an authorized balance, unauthorized requests are recorded, a controller can export logs for both events and the revocation function immediately blocks subsequent attempts.

Phase 2 is complete when an agent can advance or resolve an assigned work item using the same exception workflow as a human user. Finally, Phase 3 is complete when transaction requests are strictly blocked from settling without an OTP or dual authorization, and a complete Human-in-the-Loop (HITL) audit record is attached to the payment.

Assign responsibilities clearly across teams to maintain accountability and momentum: Product defines the target workflows and completion criteria, while IT builds the Phase 1 connections and manages the "deny-and-log" pathways. Digital transformation oversees the pilot mid-market customer segment and coordinates the treasury approvers handling the Human-in-the-Loop (HITL) queue. The collective working group retains joint ownership of the approved tool inventory to prevent unauthorized write permissions from creeping into Phase 1. At the same time, Compliance reviews audit log exports immediately rather than waiting until the end of Phase 3.

ACTIONS TO GET STARTED

This week, identify the target customer segment, select the single initial workflow and assign the executive responsible for signing off on completion. Later this month, execute 10 Phase 1 test calls in a sandbox environment — five permitted and five blocked — and generate an audit log extract that meets controller standards. If the five rejected attempts fail to appear in that audit report, the system is not secure enough to expand.

Reuse human controls; extend them to non-humans

Existing governance mechanisms — such as company admin configurations, maker-checker dual controls, transaction limits and OTP verification — already function effectively for human users. Instead of creating new frameworks, register client-owned AI agents as accountable non-human users tied to a specific owner, explicit purpose and defined mandate. Over time, repurpose existing bank portals as the administrative and exception-management interface for human users, while the AI agent operates directly from the client's environment using machine-readable tasks within a dedicated agent portal and MCP proxies. Avoid creating a parallel control framework. Instead, bind the agent directly to an existing entitlements record — linking its owner, legal entity, accessible accounts, dollar limits and approval chains. Crucially, an agent can never perform a task that its human owner is not authorized to execute.

Product defines this mandate in a concise, one-line summary, while IT issues credentials against that specific entitlement record rather than a generic shared service account. Digital transformation maintains company admin portals as the central location for human managers to modify access scope, ensuring AI agents never receive tools outside their defined mandate.

Bank portals remain dedicated to human interactions. Critical governance tasks — including approvals, exception handling, beneficiary updates and access revocation — must occur within these existing portals or through established treasury OTP channels like SMS.

The AI agent must never "log in as" its human owner; rather, it authenticates under its own distinct identity, operating strictly under the owner's explicit authority and through published, governed tasks.

ACTIONS TO GET STARTED

This week, review your inventory's Phase 1 task list and identify which of the four governance components each task is missing. You are officially ready to design the system entry point once a controller can answer five core questions for that task: identity (who is the agent), ownership (who owns it), authorization (what it may do), auditing (what it just did) and control (how to stop it).

Get compliance in the design, not at the end

Before any machine-initiated bank interaction, Risk and Compliance will require definitive answers to five questions:

01 **Who is the agent?**

02 **Who authorized it?**

03 **What is it allowed to do?**

04 **What did it actually do?**

05 **How do we stop it instantly if it goes out of scope?**

Integrate these fundamental governance answers directly into Phase 1. KYA verifies identity and authorization by defining the agent's identity, owner, legal entity, intent, credentials and digital fingerprint. Organizational policies combined with a published action menu restrict what the agent can execute, while trace IDs and session logs capture all activity — including rejected calls. Finally, agent-specific revocation and customer-wide access pauses ensure immediate containment.

As an optional future enhancement on core IBM Z systems, banks can enable continuous dynamic policy enforcement at the velocity required for high-throughput commercial channels by leveraging the on-chip AI accelerator in Telum II processors, with the option to add Spyre inference accelerator cards for even higher capacity workloads.

ACTIONS TO GET STARTED

This week, document the five core governance answers for your single Phase 1 workflow on a one-page summary and review it directly with Compliance. Later this month, generate a single log extract demonstrating an allowed call, a blocked call and an access revocation. If Compliance cannot independently interpret and utilize that log extract without filing an engineering ticket, the system design is incomplete.

Measure the first deployment

Evaluate the initial deployment using only four metrics: fully registered agents bound to a named owner, incoming requests that cannot reach core systems without passing through governance checkpoints, direct reuse of existing treasury approval workflows and an exportable audit log for a single workflow. True success is not simply deploying AI to production, but establishing a governed framework that a controller will formally approve.

Evaluate each metric using a strict binary standard:

ACTIONS TO GET STARTED

This week, identify the four evidence artifacts and assign ownership for generating each. Later this month, execute the 10-call sandbox simulation from Section 5 and attach the resulting audit extract directly to the channel charter. Do not expand the list of allowed workflows until a controller formally signs off on this complete documentation package.

Agent Registration

You either have a complete KYA record linked to a named owner, or you do not.

Checkpoint Enforcement

You can either prove a denied request was intercepted before reaching core functions, or you cannot.

Control Integration

You either reused existing maker-checker and OTP pathways, or you created a parallel queue.

Audit Accessibility

You provided Compliance with an exportable file they can retain, or you walked them through a live screen demonstration.

Phase 1 requires achieving all four positive outcomes — securing three out of four is an incomplete launch.

How OvationCXM can help

The work detailed in this guide centers entirely on building the front door: KYA, policy enforcement, human controls, instant revocation and an audit trail robust enough for a controller's formal sign-off. You can assemble every piece of this framework using your institution's existing assets and entitlement structures. Most banks stall in the assembly itself — linking an agent's identity to a human owner, publishing tasks through a dedicated agent interface, logging rejected attempts and ensuring the emergency kill switch executes in seconds.

Client Authorizes. Agent Executes. Bank Governs.

The Agentic Transaction Control Layer (ATCL) serves as that essential assembly. Positioned directly between client-authorized agents and bank-approved workflows, it handles authentication, permissioning, routing, approvals, logging and auditing. This approach lets the bank continue relying on its existing infrastructure, including HITL, OTP, monitoring and revocation mechanisms — without replacing core banking systems. By leveraging the inventory defined in Section 3 rather than waiting to build an entirely new channel stack, it offers the fastest path to a defensible Phase 1 deployment.

When the inventory highlights APIs that are not yet agent-ready, OvationCXM has experience helping banks wrap existing endpoints with MCP and expose them as bounded tasks, including structured Agent Journey steps within a dedicated agent portal.

This capabilities package can be deployed alongside ATCL whenever an institution requires the complete end-to-end integration path, not just the control layer.

ACTIONS TO GET STARTED

While the base ATCL offering runs as a fully managed service with OvationCXM, banks can deploy the control plane in a hybrid pattern next to their core. They can do this in the cloud or on IBM Z. The IBM Z option lets banks run the control plane in the same secure, high-performance environment already used for core workloads, keeping agent traffic, policy and audit inside that boundary. Phase 1 outcomes and the client experience stay the same. Our Forward Deployed Engineers are available to support every stage of preparing your enterprise for customer-facing AI agents.

OvationCXM

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.