# Banking AI Governance Readiness Assessment

*The 14 items every regional and mid-market bank using AI needs to address — mapped to OCC, FDIC, Federal Reserve, SEC, and SOC 2 expectations. Use this for audit committee agendas, third-party-risk reviews, and regulatory-exam preparation.*

**From HitLai Institute — Module B1: "Governed AI for Banking — From Shadow AI to Auditable AI"**

---

## WHAT THIS IS

Banks are using AI faster than they are governing it. Tellers, analysts, and operations staff use ChatGPT and Copilot daily — often on data the bank's information-security policy does not permit to leave the network. Compliance and audit teams find out about it after a vendor demo, a regulatory inquiry, or a customer complaint.

This is not a problem unique to banking, but it is far more consequential in banking. Customer data, transaction records, KYC files, and AML investigations are subject to specific regulatory expectations under the OCC, FDIC, Federal Reserve, SEC, FinCEN, GLBA, and the bank's SOC 2 obligations. Shadow AI inside a bank is not a productivity convenience — it is a regulatory exposure with a quarterly clock attached.

This assessment is a starting point. It is not legal, regulatory, or audit advice. Use it to ground your AI governance conversation with internal audit, compliance, and your primary regulator.

---

## THE 14-ITEM ASSESSMENT

For each item, mark Done / In Progress / Not Yet. Anything below "Done" is a gap — assign an owner and a target close date before the next audit committee meeting.

---

### GOVERNANCE STRUCTURE

☐ **1. The bank has a written AI governance policy approved by the board (or risk committee).** - Defines approved use cases, prohibited uses, owner roles, escalation paths. - Updated at least annually and after any material AI deployment. - Aligned with the bank's third-party risk management program and model risk management framework (SR 11-7).

☐ **2. An accountable AI governance owner is named.** - Single point of accountability — typically CRO, CIO, Chief Compliance Officer, or a designated AI Risk Officer. - Has authority to approve, restrict, or shut down AI use across business lines. - Reports AI risk posture to the audit committee on a defined cadence.

☐ **3. An inventory of every AI tool in use exists.** - Tool, vendor, business owner, data classification handled, supervision protocol, contractual data-handling terms. - Refreshed quarterly. Anything not on the inventory is, by policy, not approved. - This includes Copilot, ChatGPT Enterprise, embedded AI features in vendor products, and any internal model deployments.

---

## REGULATORY MAPPING

☐ **4. AI use cases are mapped to the relevant regulatory frameworks.** - OCC heightened standards and model risk management (SR 11-7 / OCC 2011-12) for any model used in credit, capital, or operational decisioning. - FDIC guidance on third-party AI and consumer protection. - Federal Reserve SR 11-7 model risk management — including non-quantitative AI models in operations. - SEC requirements for any AI use in investment advice, broker-dealer activity, or disclosures (Reg BI, ATS, marketing rule). - FinCEN expectations for AI in BSA/AML programs. - GLBA Safeguards Rule and the FFIEC IT Examination Handbook on information security. - SOC 2 Trust Services Criteria — security, availability, confidentiality, processing integrity, privacy. - State banking regulators and the NYDFS Part 500 (for New York–chartered or supervised institutions).

☐ **5. Model risk management coverage extends to AI.** - AI models — including LLM-driven workflows — are inventoried in the bank's model inventory under SR 11-7. - Each material AI use has a model risk tier, validation evidence, performance monitoring plan, and a periodic review schedule. - Generative AI use cases have a defensible position on whether they constitute "models" under your framework — get this in writing.

☐ **6. Fair lending and UDAAP risks are evaluated for any customer-facing AI.** - Any AI that touches loan decisions, account opening, credit-line management, collections messaging, or marketing has documented fair-lending testing. - Disparate-impact testing on a representative sample. - UDAAP review of AI-generated customer communications.

---

## DATA HANDLING & CONFIDENTIALITY

☐ **7. Customer data does not flow to public or unapproved AI providers.** - No employee may paste customer PII, account numbers, transaction history, KYC files, or AML investigation material into free ChatGPT, free Claude, free Gemini, or any tool not on the approved inventory. - Either approved enterprise tiers with contractual no-training, no-retention, and tenant-isolation terms OR self-hosted AI on the bank's network. - DLP rules block paste-to-known-AI-domains for anything classified as customer data.

☐ **8. Self-hosted deployment is available for the most sensitive workloads.** - AML investigations, suspicious activity narratives, internal fraud investigations, and any matter touching a regulator request run on AI infrastructure inside the bank's network. - Local AI via Ollama, vLLM, or equivalent — air-gapped where the matter requires. - Data never leaves the bank's network for these workloads.

☐ **9. Vendor agreements with AI providers meet the bank's third-party risk standards.** - No-training clauses in writing. - Data residency commitments that match your regulatory obligations. - Sub-processor transparency. - Breach notification and audit-rights clauses consistent with FFIEC guidance. - Annual reassessment.

---

## CONTROLS, SUPERVISION & AUDIT

☐ **10. Every AI interaction is logged.** - User identity, business unit, prompt, response, model used, timestamp, customer/account reference if applicable. - Audit log is tamper-evident and retained per the bank's records-retention policy (typically 5-7 years). - Audit log is searchable by examiner / internal auditor within 60 seconds.

☐ **11. AI output that affects customers or filings is verified before external use.** - Customer-facing communications, loan decisions, AML/BSA filings, regulatory responses — every AI-assisted output is reviewed by a qualified bank employee before transmission. - The verifier's name and the verification step are logged.

☐ **12. The trust journey has explicit ceilings for high-risk activities.** - KYC dispositioning, AML/BSA alert clearing, SAR filing decisions, fraud loss writeoffs, credit decisions, and any regulatory filing are NEVER fully autonomous. - These remain at "AI suggests" or "AI acts, you approve" — never at "AI handles routine." - The ceiling is enforced in the workflow configuration, not by policy alone.

---

### TRAINING & CULTURE

☐ **13. Every employee with AI tool access has completed AI-use training and signed an acknowledgment.** - Covers what's approved, what's prohibited, escalation paths, data classification, and the bank's incident reporting expectations. - Refreshed annually, with role-specific add-ons for AML, lending, treasury, and customer service. - Acknowledgment tracked alongside other compliance attestations.

☐ **14. There is a defined intake and review process for new AI tools.** - Business unit proposes; AI governance owner, IT security, compliance, and third-party risk evaluate. - Documented decision (approve / restrict / decline) with conditions. - Shadow-AI amnesty: an active channel for employees to disclose unsanctioned tool usage without punitive consequences, paired with rapid remediation.

---

## DOCUMENTATION YOUR BANK SHOULD BE ABLE TO PRODUCE WITHIN 48 HOURS

When an examiner asks, or your internal audit team requests, or a major customer's vendor risk team requests:

1. **AI Governance Policy** (current, board-approved)
2. **AI Use Inventory** (every approved tool, every business owner)
3. **Regulatory Mapping Matrix** — your AI use cases mapped to OCC, FDIC, Fed, SEC, FinCEN, GLBA, NYDFS Part 500, and SOC 2 control families
4. **Model risk inventory** including AI / generative AI entries with tier, validation evidence, and review schedule
5. **Vendor risk files** for each AI vendor — including data handling, sub-processor list, and breach notification terms
6. **Training log** — every employee, every acknowledgment, every annual refresh
7. **Audit log sample** showing user, prompt, response, verification, and outcome for any AI-assisted decision in the last 90 days
8. **Incident log** — any shadow-AI disclosure, any data-loss event, any AI-related customer complaint

If you cannot produce these in 48 hours, your bank is not yet exam-ready on AI.

---

## THE FOUR MISTAKES THAT CREATE REGULATORY EXPOSURE

1. **Treating Copilot or ChatGPT Enterprise as "already governed" because the vendor said so.** Contractual no-training is necessary but not sufficient. Without the bank's own inventory, training, supervision, and audit log, you cannot defend the use under SR 11-7 or to a primary regulator.
2. **Allowing AI inside KYC, AML, or fraud workflows without an explicit ceiling.** Even when AI is genuinely better at the first-pass triage, the final disposition belongs to a qualified analyst. A SAR filed, a customer offboarded, or an alert cleared by AI alone is a finding waiting to happen.
3. **No audit log.** Regulators no longer accept "we use AI in operations" without the ability to reconstruct what AI did, on what data, for what decision, reviewed by whom. The audit log is not optional.
4. **Shadow AI without a disclosure channel.** Punishing employees who admit to using ChatGPT on bank data drives the behavior underground. An amnesty channel paired with rapid remediation is the only practical way to surface the real footprint.

## THE TRUST JOURNEY FOR BANKING

| Setting | Use case in banking | Notes |
|---|---|---|
| **AI watches** | New AI tool, new business line, brand-new use case | Always start here. Observe before you trust. |
| **AI suggests** | KYC document extraction with analyst review, AML alert summarization, vendor-invoice categorization | Default for any operational AI work. The qualified employee decides. |
| **AI acts, you approve** | Customer service first-draft responses, internal report generation, vendor onboarding paperwork | Only after the trust setting earns it on logged data. |
| **AI handles routine** | Internal calendar entries from extracted contract dates, low-risk paperwork triage, internal knowledge-base routing | NEVER for KYC dispositioning, AML clearing, SAR decisioning, credit decisions, regulatory filings, or customer-impact decisions. |

The ceilings above are not configurable in your workflow platform if it is set up correctly for banking. They are enforced at the system level — and confirmed during the bank's annual governance review.

## THE HITLAI / AICTRLNET ANSWER, IN ONE LINE

> AI tools, your team, your systems — running together, safely.

For a bank: AI extracts, summarizes, drafts, and pre-screens. Your AML analysts, KYC dispositioners, credit officers, and operations teams do the work AI cannot. Supervisors review and approve. Customer data and investigation material stay inside your bank's network — including with self-hosted, air-gapped local AI via Ollama or vLLM. Your core, AML platform, KYC vendor, and case management systems take what's been verified. Every step auditable for OCC, FDIC, Fed, SEC, and SOC 2.

**Want a board-ready walkthrough of this assessment for YOUR bank?** That's Module B1 — bring your current AI inventory (or your "we have no inventory yet"). → **hitlai.net/institute/banking**

---

*This assessment is not legal, regulatory, or audit advice. It is an operational framework intended for internal bank use under the supervision of the bank's compliance, audit, and risk functions, and in coordination with the bank's primary federal and state regulators.*