# Banking AI Regulatory Mapping Matrix

*Every AI use case in your bank, mapped to OCC, FDIC, Federal Reserve, SEC, FinCEN, GLBA, NYDFS Part 500, and SOC 2. Use this for the audit committee, examiner pre-meetings, and the third-party-risk review of any new AI tool.*

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

---

## WHY THIS EXISTS

Banks live under a regulatory stack that did not anticipate generative AI. The good news: most of the framework already applies. The bad news: it applies in ways that are not obvious until an examiner asks the question.

A bank that can produce a single page mapping every AI use case to the relevant rules, expectations, and control families is a bank that:

- Walks into exams with a defensible position
- Closes third-party-risk reviews in days, not quarters
- Reduces the chance that a shadow-AI incident becomes a Matter Requiring Attention (MRA)
- Gives the board what it actually wants — a clear answer to "are we OK on AI?"

This matrix is the working document for that answer. Adapt to your bank's charter, business mix, and primary regulator.

---

## HOW TO USE THIS MATRIX

For every AI use case in your bank's inventory:

1. Identify the use case (e.g., "Copilot drafting customer-service emails," "KYC document extraction," "AML alert summarization").
2. For each regulatory framework below, mark whether the use case is in scope.
3. If in scope, name the specific obligation, the bank's control, and the evidence the control is operating.
4. If a row is empty for an in-scope use case, that is a gap. Assign an owner and a close date.

A single use case typically maps to 4-7 frameworks. If yours maps to fewer, double-check.

---

## THE FRAMEWORKS

### 1. OCC — Heightened Standards & Model Risk

**Source documents:** OCC Bulletin 2011-12 (Model Risk Management), OCC Heightened Standards, OCC's third-party risk management guidance (OCC 2013-29 and successors).

**In scope when:** the AI use case influences a credit decision, capital calculation, operational decision affecting risk, fair lending, or a third-party arrangement.

**Key obligations to map:** - Model identification, validation, ongoing monitoring, and periodic review (treat generative-AI workflows that drive decisions as models unless legal and risk concur otherwise — get the position in writing either way). - Third-party risk lifecycle: due diligence, contract terms, ongoing monitoring, termination. - Heightened standards for large banks: risk governance framework includes AI risk explicitly.

**Evidence to maintain:** model inventory entry, validation report, monitoring dashboard, vendor due diligence file.

---

### 2. FDIC

**Source documents:** FDIC FIL guidance on third-party arrangements, FDIC Consumer Compliance Examination Manual, FDIC interpretive letters on AI/ML use.

**In scope when:** any AI use that touches consumer products, deposit operations, or third-party arrangements at FDIC-supervised institutions.

**Key obligations to map:** - Consumer protection — fair lending, UDAAP, advertising and disclosures generated or screened by AI. - Third-party arrangement controls aligned with the Interagency Guidance on Third-Party Relationships (2023). - Information security under the Safeguards Rule.

**Evidence:** fair-lending testing report, UDAAP review, vendor contract clauses, disclosure templates.

---

### 3. Federal Reserve — SR 11-7 & SR Letters

**Source documents:** SR 11-7 Model Risk Management, SR 13-19 Third-Party Risk, recent SR letters on AI / generative AI.

**In scope when:** any AI used by a Fed-supervised institution that affects risk, capital, financial reporting, or material business decisions.

**Key obligations to map:** - Effective challenge: validation independent from the model owner. - Documentation: model purpose, methodology, limitations, monitoring plan. - Compensating controls where validation is limited (often the case for proprietary LLMs).

**Evidence:** validation memo, effective-challenge documentation, compensating-control rationale.

---

### 4. SEC — Where Applicable

**Source documents:** Regulation Best Interest, Investment Advisers Act, Marketing Rule (Rule 206(4)-1), Reg ATS, recent SEC guidance on AI use by registrants.

**In scope when:** the bank or affiliate has a broker-dealer, investment adviser, or municipal advisor registration and AI is used in advice, recommendations, marketing, or trading.

**Key obligations to map:** - Conflicts of interest from AI-driven recommendations (the “predictive data analytics” framework). - Marketing-rule compliance for any AI-generated client-facing materials. - Books and records — AI interactions affecting advice or trades are records.

**Evidence:** disclosure language, supervisory review records, retention policy that includes AI logs.

---

### 5. FinCEN — BSA / AML

**Source documents:** BSA, FFIEC BSA/AML Examination Manual, FinCEN guidance on innovative technologies (joint statement with the agencies, 2018 and successors).

**In scope when:** AI is used in transaction monitoring, KYC, customer risk rating, sanctions screening, SAR drafting, or 314(a)/(b) response.

**Key obligations to map:** - Independent testing of AI-driven monitoring effectiveness. - Risk-based assessment of model performance — no degradation of detection. - SAR quality and timeliness — AI-drafted narratives still require a qualified BSA officer’s review and sign-off. - Documentation of model changes and tuning decisions.

**Evidence:** independent testing report, model change log, SAR sample review, AML risk assessment update.

---

### 6. GLBA — Information Security

**Source documents:** Gramm-Leach-Bliley Safeguards Rule, FFIEC IT Examination Handbook, NIST cybersecurity standards adopted by the bank.

**In scope when:** customer NPI (non-public personal information) is involved in any AI workflow.

**Key obligations to map:** - Encryption in transit and at rest for any AI infrastructure handling NPI. - Access controls and identity management for AI tool use. - Incident response procedures that include AI vendor scenarios. - Annual risk assessment that addresses AI.

**Evidence:** safeguards risk assessment, encryption inventory, access review, incident playbook.

---

### 7. NYDFS Part 500 (and equivalent state regulators)

**Source documents:** 23 NYCRR 500 (NYDFS Cybersecurity Regulation, amended 2023), and equivalent state guidance.

**In scope when:** the bank is supervised by NYDFS or operates in states with parallel cybersecurity regulations (CA, MA, NJ, etc.).

**Key obligations to map:** - CISO oversight of AI cybersecurity risk. - Multi-factor authentication for AI tool access where applicable. - 72-hour incident reporting that includes AI-related cybersecurity events. - Third-party service-provider security requirements.

**Evidence:** CISO board report on AI risk, MFA enforcement evidence, incident-reporting drills.

---

### 8. SOC 2 (Trust Services Criteria)

**Source documents:** AICPA Trust Services Criteria (TSC) — Security, Availability, Confidentiality, Processing Integrity, Privacy.

**In scope when:** the bank operates a service organization, or vendors handling bank data are evaluated under SOC 2.

**Key obligations to map:** - Security: logical access to AI infrastructure and data. - Availability: uptime and recovery for AI services in critical workflows. - Confidentiality: classified data handling in AI prompts and outputs. - Processing Integrity: AI-driven processes produce accurate, complete, valid results. - Privacy: notice, choice, consent for AI use on customer data.

**Evidence:** SOC 2 report (own and vendors'), control test results, exceptions log.

---

### 9. State Banking Regulators

**Source documents:** State banking department guidance (CSBS coordination, state AI/data-protection laws).

**In scope when:** the bank has a state charter or operates in states with specific AI or data-protection regulation (CA, IL, TX, VA, CO, etc.).

**Key obligations to map:** - State-specific consumer notice requirements. - State data-protection laws affecting AI use on residents' data. - State-specific advertising / marketing AI restrictions.

**Evidence:** state-specific disclosure templates, data-residency confirmation.

---

## SAMPLE MAPPING — KYC DOCUMENT EXTRACTION

| Framework | In Scope? | Obligation | Control | Evidence |
|---|---|---|---|---|
| OCC SR 11-7 / 2011-12 | Yes | Material AI in KYC workflow; treat as model | Model inventory entry, annual validation | Model file in MRM repository |
| FDIC Third-Party | Yes | AI vendor evaluated as third party | Vendor due diligence, contract clauses | Vendor risk file |
| Fed SR 11-7 | Yes | Same as OCC | Effective challenge by independent reviewer | Validation memo |
| FinCEN BSA/AML | Yes | KYC is part of CIP; AI affects customer risk rating | Analyst review and sign-off on AI output | Audit log + SAR sample review |
| GLBA | Yes | Customer NPI in flow | Self-hosted deployment, encryption, access control | Safeguards risk assessment |
| NYDFS 500 | If supervised | CISO oversight, MFA, incident response | CISO board report, MFA logs | Quarterly CISO report |
| SOC 2 | Yes | Security, confidentiality, processing integrity | Test results from latest SOC 2 audit | SOC 2 report |
| State regulators | Depends | State data laws | State-specific disclosure | Disclosure templates |
| SEC | Generally no | N/A unless securities advice in scope | N/A | N/A |

## SAMPLE MAPPING — AI-DRAFTED CUSTOMER SERVICE EMAILS

| Framework | In Scope? | Obligation | Control | Evidence |
|---|---|---|---|---|
| OCC | Yes | Operational risk; UDAAP review | Pre-send review by qualified employee | Audit log + sample reviews |
| FDIC | Yes | UDAAP, consumer protection, fair lending in messaging | Sampling, exception reporting | Compliance review log |
| Fed | If material | Operational AI | Documentation of supervision | Process documentation |
| FinCEN | If touches AML matters | Customer communication during investigations | Reviewed by BSA officer | Sample log |
| GLBA | Yes | NPI handling | DLP rules, approved tools list | DLP exception log |
| NYDFS 500 | If supervised | Incident response for breaches | Drill records | Incident drills |
| SOC 2 | Yes | Processing integrity, confidentiality | Control tests | SOC 2 report |
| State regulators | Yes | State consumer protection | Disclosures and complaint procedures | Disclosure files |
| SEC | Generally no | N/A | N/A | N/A |

## SAMPLE MAPPING — AML ALERT SUMMARIZATION

| Framework | In Scope? | Obligation | Control | Evidence |
|---|---|---|---|---|
| OCC | Yes | BSA/AML program adequacy | Independent testing | Audit report |
| FDIC | Yes | Same | Independent testing | Audit report |
| Fed | Yes | Operational AI in core compliance | Effective challenge | Validation memo |
| FinCEN | Yes (critical) | AML program effectiveness, SAR quality | Analyst review of every AI summary, sign-off log | Audit log, SAR sample |
| GLBA | Yes | NPI throughout investigation | Self-hosted only | Architecture diagram, encryption inventory |
| NYDFS 500 | If supervised | Cyber controls around investigation data | Access logs | Quarterly review |
| SOC 2 | Yes | Confidentiality (investigation material) | Tenant isolation, encryption | SOC 2 report |
| State regulators | Depends | State AML/financial-crime law | State-specific procedures | Procedures file |
| SEC | If affiliate touches | If broker-dealer involved | Coordinated supervision | Cross-functional log |

## WHEN THE MATRIX BREAKS DOWN — AND WHAT TO DO

You will find use cases where a column is genuinely ambiguous (e.g., “is this a model under SR 11-7?”). When that happens:

1. **Write the position down.** Have legal, risk, and compliance sign the position memo.
2. **Pick the more conservative path operationally.** If in doubt about model designation, run it through model risk management anyway. Cheaper than the finding.
3. **Bring the position to your primary regulator early.** Examiners prefer to weigh in before the deployment, not after.
4. **Re-evaluate on a cycle.** Regulatory guidance is evolving. A position taken in 2026 may not survive 2027.

## THE COMPETITIVE EDGE

Banks that walk into 2026-2027 exams with a clean matrix:

- Spend less time defending AI use to examiners; more time on substantive risk topics.
- Close vendor risk reviews from large customers faster.
- Make better build-vs.-buy decisions on AI vendors (the matrix reveals where vendor risk is highest).
- Reduce the probability of an MRA on AI-related controls.

Banks that don't:

- Find out about AI gaps during the exam — the worst time.
- Get differential treatment from cautious customers in vendor risk reviews.
- Pay for AI tooling that cannot survive a meaningful audit.

---

## THE HITLAI / AICTRLNET ANSWER, IN ONE LINE

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

For a bank: AI extracts and pre-screens. AML, KYC, fraud, and credit specialists do the work AI cannot. Supervisors review. Customer data stays inside your bank's network — including with self-hosted, air-gapped local AI via Ollama or vLLM, mapped to OCC, FDIC, Fed, SEC, FinCEN, GLBA, NYDFS, and SOC 2. Your core, AML platform, KYC vendor, and case management systems take what's been verified. Every step auditable.

**Want a working session on filling this matrix in for YOUR bank?** That's part of Module B1 — bring your current AI inventory and one regulatory-mapping question you can't answer today. → **hitlai.net/institute/banking**

---

*This matrix is not legal, regulatory, or audit advice. It is operational guidance for 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.*