

# Banking AI Business Case – Template

*The board-ready and risk-committee-ready business case for governed AI orchestration at a regional or mid-market bank. Three sections, one ROI worksheet, and a defensible answer to the question “why now and why this approach?”*

**From HitLai Institute – Module B3: “AI Strategy for Banking Leadership”**

---

## WHY THIS TEMPLATE EXISTS

---

Most banking AI proposals fail because they answer the wrong question. The board does not need a deck about how transformers work. The risk committee does not need a tour of vendor logos. The CFO does not need an AI-feature comparison chart.

They need three things:

1. **What does this do for the bank** – in dollars saved, capacity recovered, exam-readiness improved, or revenue defended?
2. **What does this cost and what could go wrong** – including the failure modes that get banks into trouble?
3. **Why this approach now** – and what happens if we wait?

This template answers those three questions for governed AI orchestration specifically. Adapt to your bank’s strategic-planning format.

---

## SECTION A – EXECUTIVE SUMMARY

---

(One page. Read top to bottom. This is what the board sees first.)

### A.1 Recommendation

[Bank Name] should deploy governed AI orchestration across [3-5 named workflows in priority order – typically a subset of B2-01] over [N] months, on a [self-hosted / private-tenant / hybrid] deployment architecture, with [CRO / CIO / CCO / AI Risk Officer] as the named accountable owner.

### A.2 Three-line problem statement

- Employees are already using public AI tools on bank data in unsanctioned ways. The risk is real and unmeasured.
- KYC, AML, customer service, and back-office paperwork volumes are rising while headcount budgets are not.
- Examiners and large customers’ vendor risk teams are starting to ask AI-specific questions the bank cannot fully answer today.

### **A.3 What we are proposing – in one sentence**

Deploy a governed AI orchestration layer that sits between approved AI models and the bank's existing systems (core, AML platform, KYC vendor, case management, payments) – with policy enforcement, audit logging, and supervision built in, deployable inside the bank's network for sensitive workloads.

### **A.4 Year-1 expected outcomes (illustrative – fill in with your numbers)**

- Hours recovered: \_\_\_ FTE-equivalent / year (across KYC, AML triage, customer service, AP)
- Shadow-AI exposure: Reduced from “unmeasured” to “inventoried, monitored, and bounded”
- Exam-readiness: Documented governance pack ready in 48 hours; AI-specific finding probability materially lowered
- Customer experience: Faster response times on routine inquiries with no decrease in complaint quality
- Strategic optionality: Position to offer governed-AI features in commercial and treasury products

### **A.5 Investment range**

- Year 1: \$\_\_\_\_\_ (platform + deployment + governance setup + training)
- Year 2-3 steady state: \$\_\_\_\_\_ / year (subscription + ongoing governance)
- Internal effort: \_\_\_ FTE for setup; \_\_\_ FTE steady-state across IT, compliance, business owners

### **A.6 What we are NOT proposing**

- Replacing AML analysts, KYC dispositioners, credit officers, sanctions reviewers, or any decision-maker with AI
- Sending customer data to public AI services
- Adopting a single AI vendor without a fallback or multi-model strategy
- Bypassing model risk management

---

## **SECTION B – PROBLEM AND OPPORTUNITY**

### **B.1 The shadow-AI baseline**

Estimate the bank's actual AI footprint today (be honest, not aspirational):

Tool / vendor	Approved?	Estimated bank-wide use (hrs/ week)	Data classifications likely touched
ChatGPT (free)	<input type="checkbox"/>		
ChatGPT Enterprise	<input type="checkbox"/>		
Microsoft Copilot	<input type="checkbox"/>		
Claude (free)	<input type="checkbox"/>		
Claude for Work	<input type="checkbox"/>		
Gemini (consumer)	<input type="checkbox"/>		
Gemini for Workspace	<input type="checkbox"/>		
Embedded AI in vendor products (KYC, AML, CRM, etc.)	<input type="checkbox"/>		
Other: _____	<input type="checkbox"/>		

The point of this table is not blame. The point is the realistic baseline. Banks that pretend the baseline is zero are banks that get surprised.

## B.2 The pressure curve

- **Internal pressure** — business units want faster turnaround; ops want headcount relief; tech wants modern tooling.
- **External pressure** — examiners are asking; large customers' vendor risk teams are asking; competitors are advertising AI features.
- **Risk pressure** — every quarter without governance is another quarter of compounding shadow-AI exposure.

## B.3 The opportunity

- Reclaim ops capacity from rising work volumes without adding headcount
- Reduce shadow-AI exposure from “uncapped” to “inventoried + bounded”
- Walk into exams with a defensible AI posture and the documentation to back it
- Enable governed-AI features in commercial banking products (treasury, lending, deposit operations) where larger competitors are already moving

## SECTION C – APPROACH

---

### C.1 What governed AI orchestration is

A platform that sits between approved AI models and the bank’s existing systems, with these layers:

1. **AI model layer** – multiple approved providers (frontier-cloud with no-training contracts, self-hosted local models for sensitive workloads). Vendor-independent.
2. **Orchestration layer** – workflow definitions, trust settings, supervision rules, and audit logging.
3. **Integration layer** – connectors to the bank’s core, AML platform, KYC vendor, loan origination, case management, CRM, payments, and document management. What’s missing, the platform builds, with the bank’s IT reviewing before it goes live.
4. **Governance layer** – policy enforcement, role-based access, audit log, model risk management hooks, and incident response.

### C.2 The three deployment options

Option	When	Trade-offs
<b>Self-hosted on bank infrastructure</b>	KYC, AML, fraud, regulator-related work, any matter with elevated sensitivity	Highest assurance; data never leaves bank network; requires bank IT capability; local models slightly behind frontier on absolute capability
<b>Private-tenant cloud</b>	General operations, customer service starters, non-NPI workflows	Frontier model quality; lower IT burden; requires careful contract review
<b>Hybrid</b>	Most realistic for regional and mid-market banks	Right tool per workflow; cost-balanced; requires clear routing policy

### C.3 What we are buying vs. what we are doing ourselves

Layer	Vendor	Bank-owned
Core platform	HitLai / AICtrlNet	—
Local model inference (if self-hosted)	Llama / Qwen / Mistral via Ollama or vLLM (open weights)	Operated by bank
Cloud model providers (if private-tenant)	Anthropic / OpenAI / Google enterprise tiers	Contracts owned by bank
Integrations to bank systems	Platform connectors + auto-build	Reviewed and approved by bank IT
Governance policy and supervision	—	Bank-owned, mapped to existing frameworks
Audit log retention and search	Platform + bank archive	Retention policy owned by bank
Model risk management	—	Bank-owned under SR 11-7

### C.4 Why now

- **Earlier is cheaper.** Shadow-AI cleanup costs compound. Every quarter of uncatalogued use grows the cleanup effort.
- **Examiners are calibrating.** Banks that walk into 2026-2027 exams with a documented posture get a different conversation than banks that don't.
- **Vendor-independence has a window.** The longer the bank waits, the more pressure there is to standardize on a single AI vendor — losing the flexibility to switch as the technology shifts.
- **Customer expectations are moving.** Commercial customers in regulated industries (their own legal teams and procurement) are starting to score banks on AI policy as part of vendor risk reviews.

### C.5 Why this approach (vs. alternatives)

Alternative	What it gives us	What it doesn't give us
Do nothing	No cost, no change	Compounding shadow AI; no defensible posture for exams; left behind on customer experience
Single-vendor AI (e.g., Microsoft Copilot only)	Familiar vendor relationship	Vendor lock-in; no self-hosted path; limited control over data residency and model choice
Custom-built internal AI platform	Maximum control	18-24 months to build; ongoing engineering cost; lack of regulatory deployment patterns
Embedded AI features inside each vendor's product (siloed)	Quick wins	No central governance; each vendor's policy is different; no unified audit log; high risk of finding
<b>Governed AI orchestration platform (proposed)</b>	Central governance + audit log; multi-model; self-hosted option; vendor-independent	New vendor relationship to manage; need to commit to a deployment architecture

## SECTION D – ROI WORKSHEET

Fill in the gray cells with the bank's own numbers; use the catalog ranges in [[B2-01-Banking-Workflow-Catalog]] as a starting point.

### D.1 Hours recovered per year

Workflow	Starting volume / month	Hours/unit (current)	Hours/unit (with AI assist)	Hours saved / month	Hours / year
KYC document extraction					
AML alert summarization					
Vendor invoice processing					
Customer service first-drafts					
Loan document pre-review					
Reconciliation triage					
<b>TOTAL</b>					

### D.2 Equivalent FTE recovered

Total hours / year ÷ 2,000 hours/FTE = \_\_\_ FTE-equivalent

### D.3 Dollar value of recovered capacity

Avoided hire(s) at fully-loaded \$\_\_\_\_\_ × \_\_\_ FTE = \$\_\_\_\_\_

Reassigned to revenue-generating activity at \$\_\_\_\_\_ contribution × \_\_\_ FTE = \$\_\_\_\_\_

Reduced overtime / contractor spend \$\_\_\_\_\_

Other: \$\_\_\_\_\_

**Subtotal capacity value: \$\_\_\_\_\_ / year**

### D.4 Risk and exam value (best-estimate, not booked)

- Avoided MRA on AI controls: estimated cost (legal, consulting, remediation, possible CMP) \$\_\_\_\_\_
- Avoided customer-impact incident (UDAAP, breach, SAR failing, fair-lending) \$\_\_\_\_\_
- Avoided shadow-AI breach with regulatory reporting: \$\_\_\_\_\_
- Avoided lost commercial deal due to AI-policy vendor risk review failure: \$\_\_\_\_\_

**Subtotal risk-avoidance value: \$\_\_\_\_\_ / year**

**D.5 Strategic / revenue defense**

- Avoided customer-experience erosion vs. competitors with faster responses: \$\_\_\_\_\_
- Pipeline of governed-AI features for commercial customers in regulated industries: \$\_\_\_\_\_
- Stronger position for M&A or partnership discussions where AI maturity matters: \$\_\_\_\_\_

**Subtotal strategic value: \$\_\_\_\_\_ / year**

**D.6 Total benefit (illustrative)**

A + B + C above = \$\_\_\_\_\_ / year

**D.7 Cost**

Cost category	Year 1	Year 2-3 / yr
Platform subscription		
Implementation services		
Infrastructure (self-hosted if applicable)		
Internal effort (loaded FTE)		
Training		
Governance refresh (annual)		
<b>TOTAL</b>		

**D.8 Payback period and ROI**

- Net benefit year 1: \$\_\_\_\_\_
- Payback months: \_\_\_\_
- 3-year ROI: \_\_\_\_ %

(Honest framing for the board: the hours-recovered case alone is rarely the entire ROI. The risk-avoidance and strategic-defense legs of the case are real but harder to quantify — present them in ranges, not point estimates.)

---

**SECTION E — RISKS AND COMPENSATING CONTROLS**

---

(Most boards expect this section. Be candid; vague risk language signals weakness, not strength.)

Risk	Likelihood	Impact	Compensating control	Residual
Vendor disruption or business-model change			Multi-model architecture; data portability; exit plan	
Model performance drift			Monitoring, periodic validation, fallback	
Shadow-AI gap persists despite deployment			Mandatory amnesty channel + remediation; DLP rules	
Examiner views the deployment unfavorably			Pre-meeting briefing; documented governance pack	
Customer or staff resistance			Change management; tier the rollout; communicate guardrails	
Implementation delay			Phased rollout; pilot before scale; quarterly checkpoints	
Data-loss incident through AI			Self-hosted for sensitive workloads; encryption; audit; incident response drill	
Inadequate model risk treatment			MRM owner sign-off pre-deployment; periodic review	
Fair-lending or UDAAP exposure			Sampling, testing, supervisor review, complaint correlation	

## SECTION F – TIMELINE AND MILESTONES

---

### Quarter 1 – Governance foundation

- AI governance policy approved by board / risk committee
- Named accountable owner appointed
- Tool inventory completed
- Vendor selected; contract signed
- Deployment architecture decided

### Quarter 2 – Deployment and pilot

- Self-hosted / private-tenant / hybrid deployed

- 1-2 lowest-risk workflows go live with single business unit
- Training launched
- Audit log operational

**Quarter 3 — Scale to priority workflows**

- 3-5 workflows live across business units
- Supervision rhythm operating
- First quarterly model risk review

**Quarter 4 — Exam-ready posture**

- Governance pack documented for every live workflow
- Internal audit walkthrough completed
- Year-1 outcomes measured against business case
- Year-2 plan to board

**SECTION G — DECISION REQUESTED**

The board / risk committee is asked to:

- Approve the deployment of governed AI orchestration as described
- Approve the year-1 budget of \$\_\_\_\_\_
- Confirm the named accountable owner
- Authorize policy and procedure refresh as required
- Receive quarterly updates on shadow-AI reduction, exam-readiness posture, and ROI tracking

**THE HITLAI / AICTRLNET ANSWER, IN ONE LINE**

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

For a bank: AI extracts, pre-screens, and drafts; your 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. Every workflow audit-ready for OCC, FDIC, Fed, SEC, FinCEN, GLBA, NYDFS, and SOC 2. ROI in the hours-recovered case; durable upside in risk avoidance and strategic defense.

**Want a working session on filling in this business case for YOUR bank’s specific operating model and regulator?** That’s Module B3 — the 60-minute leadership-ready strategy session. → [hitlai.net/institute/banking](https://hitlai.net/institute/banking)

*This template is not legal, regulatory, audit, or investment advice. It is a planning aid for internal bank use under the supervision of compliance, audit, risk, and finance.*

© 2026 Bodaty LLC. All rights reserved. AICtrlNet™, HitLai™, and “Governed AI Orchestration”™ are trademarks of Bodaty LLC.