BANKING EXECUTIVE AI GOVERNANCE & DEFENSE REVIEW

Govern authority. Preserve human judgment. Prove recoverability.

A decision template for one consequential AI-enabled banking workflow

Built from November 2026 TRAIT, the proprietary Organizational Field Manual, and Employee Field Guide v1.2

myWorkforceAgents.ai | Dr. Timothy J. Giardino

The executive decision

For one defined workflow, establish whether consequential action is:

AUTHORIZED

Who granted the mandate and under what limits?

ENFORCEABLE

Do technical controls constrain execution?

OBSERVABLE

Can the bank detect cumulative consequences?

INTERRUPTIBLE

Can accountable people restrict harmful action?

RECOVERABLE

Can outcomes be reconstructed and corrected?

This complements existing risk, cyber, fraud, model, vendor, and incident programs. It does not presume a gap.

Select and bound the workflow

- Trusted customer or business outcome
- Work sequence, execution points, channels, systems, and vendors
- Principal, agent or human actor, mandate, expiration and revocation
- Named business owner, risk/security owner, decision authority
- Current operating state and next review date

Record scope and evidence identifiers only. Keep confidential bank data in approved internal systems.

Do not combine different threat pathways

Malicious intrusion or API abuse

Test payment-level enforcement when normal application checks are bypassed.

Customer-authorized external agent

Verify the mandate independently from login credentials and provider assertions.

Bank- or vendor-operated agent

Bound and monitor assigned digital work; restrict and revoke authority when warranted.

Human deception or manipulated content

Establish verification and escalation for unexpected instructions or suspicious source material.

Identify the threat pathway most relevant to your institution. Different threats require different controls.

Five connected evidence proofs

01 Identity	Agent/provider, represented party, access channel, delegation chain and accountable owner
02 Mandate	Purpose, permitted data and actions, duration, financial limits, confirmation, revocation
03 Operating Enforcement	Server-side and other external controls; negative, expired and revoked permission testing
04 Trajectory	Cumulative activity across sessions, retries, connected tools and delegated agents
05 Recourse and Recovery	Reconstruction, customer challenge, pending-work reconciliation and reauthorization

For every proof: record owner, evidence reference/date, status, limitation, corrective action, and due date.

Evidence status is not authorization

01

Not applicable, with reason

02

Not examined

03

Documented policy only

04

Implemented, not tested

05

Tested within defined scope

06

Failed or evidence unavailable

Policy, implementation, test results and the authorization to operate are four different claims.

Frontline human readiness

Decision clarity

What is the employee approving: wording, evidence, exception or action?

Agent-originated requests

Can staff verify external agent mandates and route disputed authority?

Handoff acceptance

Are work state, uncertainty, decision, owner and deadline transferred and accepted?

Intervention rights

Are reporting routes, permitted pause authority and reachable decision makers clear?

Qualified capacity

Can trained reviewers meet actual deadlines across shifts and shared queues?

Learning and trust

Can staff challenge outputs, identify hidden work and practice independent judgment?

The Employee Guide is educational. It does not transfer system design or security responsibility to employees.

Adverse scenario and controlled response

TRIGGER	Expired mandate; unauthorized tool chain; API bypass; suspicious payment pattern; reviewer unavailable; vendor outage
CONTAIN	Identify the restriction authority and enforcement point. Confirm downstream work that persists after initiating activity stops.
RECONCILE	Separate proposed, authorized, attempted, completed, pending, uncertain and disputed actions. Assign customer-impact ownership.
ABSORB	Prove trained human coverage, deadlines, and feasible reduced service. Do not assume people can absorb machine-scale work.
REAUTHORIZE	Confirm valid mandate, retested controls, reconciled work, coverage, residual-risk approval and next review.

Executive authorization decision

- Operating state: authorized with conditions / restricted / paused / undecided
- Precise permitted and prohibited actions
- Required evidence and unresolved limitations
- Named human coverage and backup
- Customer remediation and external consequence owner
- Approving authority, effective date, expiry and restriction triggers

No required safety evidence, no consequential autonomy.

A proportionate 30-day adoption path

DAYS 1–5

Select workflow, owners, business outcome, evidence register and existing baseline

DAYS 6–12

Map actual mandates, control locations, handoffs and required human judgment

DAYS 13–20

Run one authorized adverse-scenario tabletop or controlled test with specialists

DAYS 21–30

Record approved changes, employee support needs, missing evidence and retest dates

Adapt the cadence to the bank’s risk process, legal requirements, and operational availability.

Doctrine crosswalk

TRAIT Framework | Public

Governance sequence, operating envelopes, external mandates, action/session/trajectory/outcome controls, conditional authority.

Organizational Field Manual | Proprietary

Deeper operating methods for authorized advisory application. Not a free distributed implementation manual.

Employee Field Guide v1.2 | Public

Role clarity, approved AI use, handoffs, challenge, oversight capacity, hidden work, learning, and responsible fallback.

myworkforceagents.ai/frameworks-resources

Devil’s Advocate: test before claiming

- A mature bank may already have adequate controls. Establish incremental value first.
- An external customer agent is not the same threat actor as a malicious intruder.
- A human approval step cannot compensate for bypassable server-side payment controls.
- A shutdown can create unsafe backlogs if qualified manual capacity is insufficient.
- A completed assessment is not certification, regulatory compliance, or proven attack prevention.

Source basis: November 2026 TRAIT, Organizational Field Manual, Employee Field Guide v1.2; CERT-In advisories (October 9 and November 4, 2026), Reuters (October 22, 2026), and US Treasury financial-services AI risk guidance (February 2026).

Start strengthening your AI governance today.

Two free resources. Your approved AI assistant. No new software purchases or consulting fees required.

01 TRAIT FRAMEWORK

Your organizational starting point

Map one consequential workflow. Clarify authority, accountable owners, decision gates, intervention, and recovery.

02 EMPLOYEE FIELD GUIDE

Your people-centered starting point

Help staff understand approved AI use, verify outputs, question unclear mandates, escalate safely, and practice fallback.

START TODAY

Give both guides to your approved AI assistant. Use a fictional or authorized non-sensitive banking scenario to identify one readiness gap and assign a human owner.

Use your own prompt or master prompt within the employee field guide to get started.

FREE PUBLIC RESOURCES myworkforceagents.ai/frameworks-resources

Educational readiness and transition planning, not attack prevention, security certification, or compliance assurance. Use approved tools; never upload sensitive banking data without authorization.