SAI360

Preparing for the EU AI Act

AI GOVERNANCE & MODEL RISK IN BANKING

A practical guide for banking risk, compliance, and governance leaders preparing to govern AI with the rigor, traceability, and accountability regulators increasingly expect.

The challenge is no longer whether banks use AI. The challenge is whether they can prove control over it. The EU AI Act turns AI governance from an internal best practice into a regulatory obligation with evidence requirements, named accountability, and ongoing oversight expectations.

Table of Contents

THE CORE CHALLENGE 3

EU AI ACT IMPACT 4

THE GOVERNANCE GAP 5

SIX PILLARS OF AI GOVERNANCE 6

GEN AI GOVERNANCE 7

READINESS DIAGNOSTIC 8

LEGACY VS. MODERN GOVERNANCE 9

OPERATIONAL GOVERNANCE 10

MRM INTEGRATION 11

READINESS ROADMAP 12

OPERATIONAL ADVANTAGE 13

HOW SAI360 CAN HELP 14

01 / THE CORE CHALLENGE

AI adoption has outpaced governance in banking

AI adoption in banking expanded rapidly from targeted use cases to enterprise-wide deployment. Banks moved quickly because AI delivered clear gains in efficiency. But governance has not kept pace. Generative AI, embedded vendor capabilities, intelligent automation, and business-led experimentation are expanding AI exposure beyond traditional model risk processes.

At the same time, regulators are increasing expectations around traceability, accountability, human oversight, and operational resilience. The challenge for many banks is no longer whether AI is being used — it is whether governance teams can maintain visibility and evidence across a rapidly evolving AI landscape.

For many institutions, the issue is not whether AI governance policies exist. The issue is whether governance mechanisms operate consistently across a rapidly expanding AI environment. AI systems now emerge from multiple sources simultaneously: internally developed models, vendor platforms, embedded copilots, low-code automation tools, and employee-led experimentation. Governance teams often discover exposure only after AI capabilities are already operational.

Most banks do not have an AI adoption problem. They have an AI governance evidence problem.

INCOMPLETE INVENTORY

Shadow AI, embedded vendor capabilities, business-led pilots, and informal GenAI use can sit outside formal model inventories. Without visibility, risk teams cannot classify, govern, or monitor AI exposure.

FRAGMENTED ACCOUNTABILITY

AI governance spans business, data science, technology, compliance, legal, operational risk, model risk, third-party risk, and audit. Without clear ownership, accountability becomes diffused.

INSUFFICIENT EVIDENCE

Regulators will not only ask whether controls exist. They will expect evidence that controls are working, decisions are traceable, and oversight is active across the AI lifecycle.

This shift has major implications for banking leadership teams. AI governance can no longer operate as an isolated model risk exercise. It increasingly requires coordination across compliance, operational risk, technology, legal, procurement, privacy, internal audit, and executive oversight functions. Institutions that fail to operationalize governance across these areas may struggle to demonstrate consistent control environments under regulatory scrutiny.

Regulators increasingly view unmanaged AI risk not as a technology issue, but as a financial stability issue. Recent IMF warnings around AI-driven operational disruption, cyber exposure, and correlated failures are accelerating board-level scrutiny across the banking sector.

What the EU AI Act changes for banks

The EU AI Act moves AI governance from broad supervisory guidance to a formal operational accountability framework. Regulators will increasingly expect banks to identify AI systems, classify risk, document controls, monitor performance, and demonstrate human oversight with defensible evidence.

Historically, governance expectations focused primarily on model development and validation standards. The EU AI Act expands those expectations across the full AI lifecycle — including deployment, monitoring, vendor oversight, operational use, and ongoing accountability.

Governance is no longer evaluated primarily through policy documentation. It is evaluated through evidence that controls are actively operating in practice.

Risk Classification

AI SYSTEMS MUST BE ASSESSED BY USE CASE AND IMPACT

Banks need visibility into where AI is being used, what decisions it supports, and whether the use case creates elevated customer, operational, or regulatory risk.

Human Oversight

HUMAN REVIEW MUST BE STRUCTURED AND DEMONSTRABLE

Institutions need defined review procedures, escalation paths, override standards, and evidence showing how human judgment is applied in practice..

Technical Documentation

GOVERNANCE EVIDENCE MUST BE AUDIT-READY

Validation records, testing results, monitoring activities, approvals, model changes, and remediation evidence must withstand supervisory review.

Third-Party Accountability

VENDOR AI DOES NOT REMOVE INSTITUTIONAL RESPONSIBILITY

Banks remain accountable for AI risk even when capabilities originate from third-party vendors or embedded platforms.

WHAT THIS MEANS FOR LEADERSHIP

AI governance is becoming an enterprise risk discipline rather than a standalone model governance activity. Boards, regulators, and audit functions will increasingly expect:

- Clear accountability for AI systems
- Centralized governance evidence
- Continuous monitoring capabilities
- Effective third-party oversight
- Operational readiness for regulatory review

For many institutions, the challenge is no longer creating policies. It is operationalizing governance consistently across a rapidly expanding AI environment.

Why traditional governance models struggle with AI

Most banks already maintain mature risk, compliance, audit, and model risk functions. The challenge is that AI risk cuts across all of them simultaneously.

Traditional governance models were designed for relatively stable systems with defined ownership, periodic validation cycles, and centralized oversight. AI environments increasingly operate differently. Models may evolve continuously, rely on dynamic data sources, appear inside third-party platforms, or spread through decentralized business experimentation before formal governance processes engage.

As AI adoption accelerates, governance gaps often emerge gradually rather than through a single failure point. Small gaps in visibility, accountability, monitoring, or documentation can accumulate into material regulatory and operational risk over time.

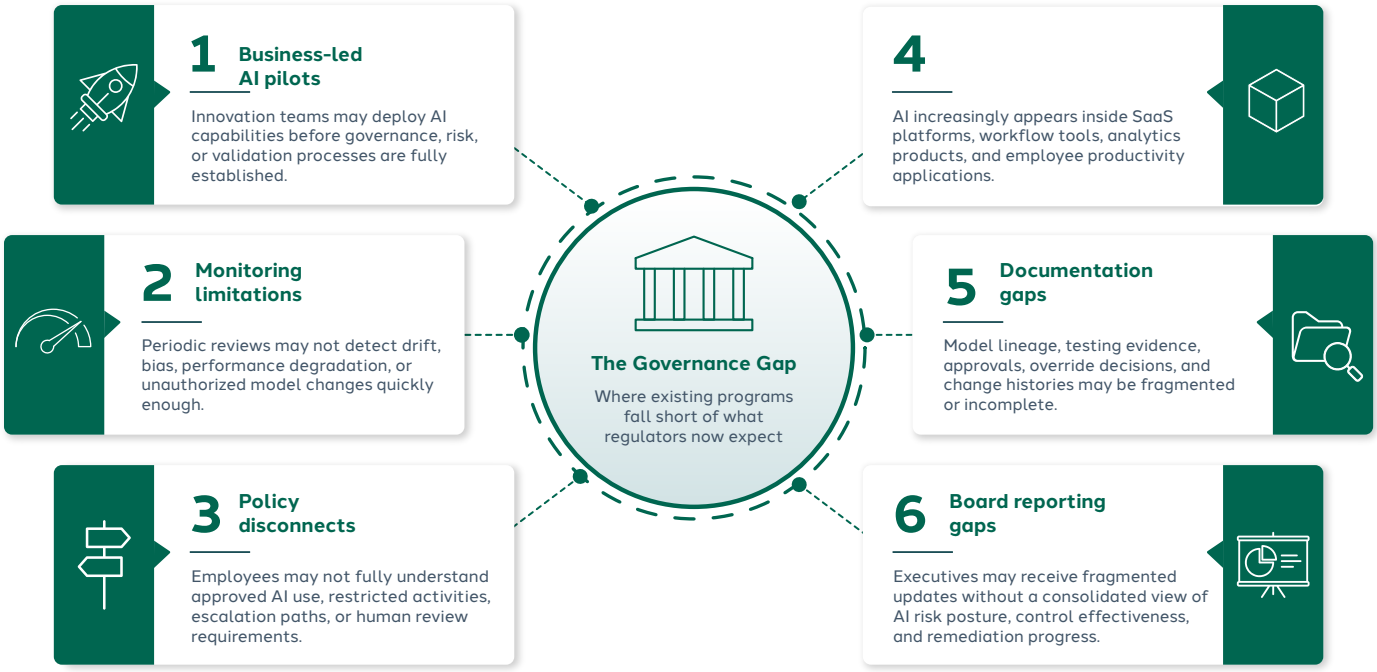

Most AI governance failures will not emerge from a single catastrophic event. They will emerge from accumulated control fragmentation across systems, vendors, workflows, and business teams.

The institutions best prepared for evolving regulation will be those that operationalize governance continuously, not those relying solely on periodic review and static documentation.

Six pillars of AI governance in banking

AI governance should be embedded into the bank's existing risk architecture — connecting AI oversight into model risk, operational risk, compliance, third-party risk, policy, training, and audit.

Effective AI governance requires more than isolated controls. Banks need an operating model that connects governance activities across the enterprise while aligning oversight responsibilities to the AI lifecycle itself. The objective is not to create a separate governance function for AI, but to integrate AI oversight into existing risk, compliance, and operational frameworks in a scalable and defensible way.

1 Inventory & classification

A complete inventory of internal models, vendor AI, GenAI tools, and AI-enabled processes. Classify by use case, impact, and risk level. Flag high-risk customer-impacting applications.

2 Governance & accountability

Every AI system needs named ownership, defined oversight, approval workflows, escalation rules, and clear accountability across the three lines of defense.

3 Risk controls

AI performance changes as data and conditions evolve. Controls must detect issues before they become regulatory, customer, or operational failures. Monitor drift, bias, and performance degradation.

4 Documentation & audit

If it is not documented, it did not happen. Maintain model change logs, human override rationale, data lineage, and a defensible record of AI decisions and controls.

5 Vendor oversight

Third-party AI introduces concentration risk, transparency limitations, and accountability challenges. Require AI risk disclosures, review model transparency, and monitor vendor obligations over time.

6 Policy & training

Policies only work when employees understand them. Role-based training on approved AI use, prohibited practices, data handling, escalation, and human review requirements.

While organizational structures vary across institutions, regulators increasingly expect these core governance capabilities to exist in some form. The maturity of the control environment may differ by institution size and AI adoption level, but visibility, accountability, monitoring, and evidence management are rapidly becoming baseline expectations.

Generative AI expands governance from models to enterprise behavior

Generative AI creates governance challenges that differ materially from traditional statistical and predictive models. Instead of producing narrow, repeatable outputs, GenAI systems can generate open-ended responses, interact dynamically with users, and draw on large volumes of structured and unstructured information.

For banks, the risk is not only the model itself. It is how employees, vendors, copilots, and business workflows use GenAI in daily operations. Governance must therefore extend beyond model validation to address prompts, outputs, data handling, human review, auditability, and policy compliance.

EXECUTIVE TAKEAWAY

The governance challenge with GenAI is not just model behavior. It is the speed at which unmanaged AI usage can spread across the enterprise.

Banks want to...	But must also...
Enable employee productivity with copilots and assistants	Prevent sensitive customer, employee, and business data from entering unapproved tools
Accelerate analysis, drafting, and decision support	Control hallucinations, unsupported conclusions, and over-reliance on AI-generated outputs
Deploy GenAI capabilities across business workflows	Maintain clear human accountability, escalation rules, and review standards
Use third-party and embedded AI at scale	Preserve vendor transparency, audit trails, documentation, and regulatory defensibility

Why traditional controls may miss the risk

Traditional model governance focuses on known systems, defined inputs, and periodic review. GenAI risk often emerges through decentralized usage, variable prompts, changing outputs, embedded vendor capabilities, and informal employee workflows.

What governance must now address

Banks need policies, training, monitoring, and evidence capture that define approved use, restricted use, human review requirements, data boundaries, prompt/output logging, and escalation paths for high-impact use cases.

06 / READINESS DIAGNOSTIC

What regulators will expect banks to prove

Regulatory defensibility depends on evidence. Banks should be prepared to demonstrate not only that AI governance policies exist, but that oversight activities are actively operating across the enterprise. The following questions reflect the types of evidence regulators, auditors, and boards are increasingly expected to review.

1. AI Visibility

- Can you identify every AI system in use?
- Can you classify AI systems by use case and risk level?
- Can you prove who owns each AI system?
- Can you identify embedded AI inside vendor platforms?
- Can you detect shadow AI and employee-led experimentation?

2. Control Evidence

- Can you show how high-risk systems are approved?
- Can you document drift monitoring and threshold breaches?
- Can you produce fairness and bias testing evidence?
- Can you demonstrate validation testing before production deployment?
- Can you show control testing results and remediation status?

3. Human & Third-Party Oversight

- Can you explain high-impact AI-supported decisions?
- Can you show human oversight and override records?
- Can you prove third-party AI controls are adequate?
- Can you verify vendor compliance with your governance standards?
- Can you document escalation paths for AI-related concerns?

4. Audit & Board Readiness

- Can you show employee training and attestations?
- Can you evidence incident response and remediation?
- Can you report AI risk posture to executives and the board?
- Can you produce a consolidated AI risk dashboard on demand?
- Can you produce audit-ready evidence without manual reconstruction?

WHERE READINESS BREAKS DOWN

If the answer to these questions lives across spreadsheets, shared drives, email threads, vendor portals, and disconnected systems, the governance process itself becomes a risk. Readiness requires a defensible system of record — not a last-minute evidence collection exercise.

From periodic review to continuous oversight

AI governance cannot rely solely on static documentation and annual review cycles. Banks need a connected, evidence-based approach that supports continuous monitoring.

LEGACY APPROACH

- Model inventories updated periodically
- Validation occurs at fixed intervals
- AI risk evidence is scattered across teams
- Vendor AI reviewed through generic vendor controls
- Policy training disconnected from actual AI use cases
- Board reporting focuses on status rather than risk posture

MODERN APPROACH

- AI systems are continuously inventoried and classified
- Controls monitor drift, fairness, performance, and use
- Evidence is centralized and audit-ready
- Vendor AI assessed through AI-specific oversight
- Training is role-based and tied to approved AI use
- Executives see AI risk exposure, trends, and remediation progress

Many existing governance frameworks were designed around relatively stable model environments with periodic validation cycles and limited operational change. AI systems increasingly operate differently. Models may update more frequently, interact with dynamic data sources, rely on third-party infrastructure, or generate outputs that vary across contexts. Governance approaches must evolve accordingly.

GOVERNANCE AREA	COMMON GAP	MODERN REQUIREMENT
Inventory	Incomplete view of AI across vendors, business tools, and GenAI use.	Enterprise AI inventory with classification by use case, risk, owner, and control status.
Model risk	Validation cycles may not keep pace with changing AI behavior.	Continuous monitoring of drift, performance, fairness, and retraining triggers.
Third-party risk	Vendor reviews may not capture AI-specific transparency needs.	AI-specific vendor assessments, documentation requests, contractual controls, and ongoing review.
Compliance	Policies not clearly defined, approved, restricted, and prohibited AI uses.	Role-based AI policies, training, attestations,
Audit	Evidence is fragmented and difficult to reconstruct.	Centralized logs, decision evidence, change history, review records, and control testing results.

From periodic review to continuous oversight

Traditional governance models were designed for slower-moving systems with periodic validation cycles and centralized oversight. AI systems operate differently. They may update more frequently, rely on dynamic data sources, generate variable outputs, and introduce risk across decentralized workflows and third-party platforms. Governance must evolve from static documentation to continuous operational oversight.

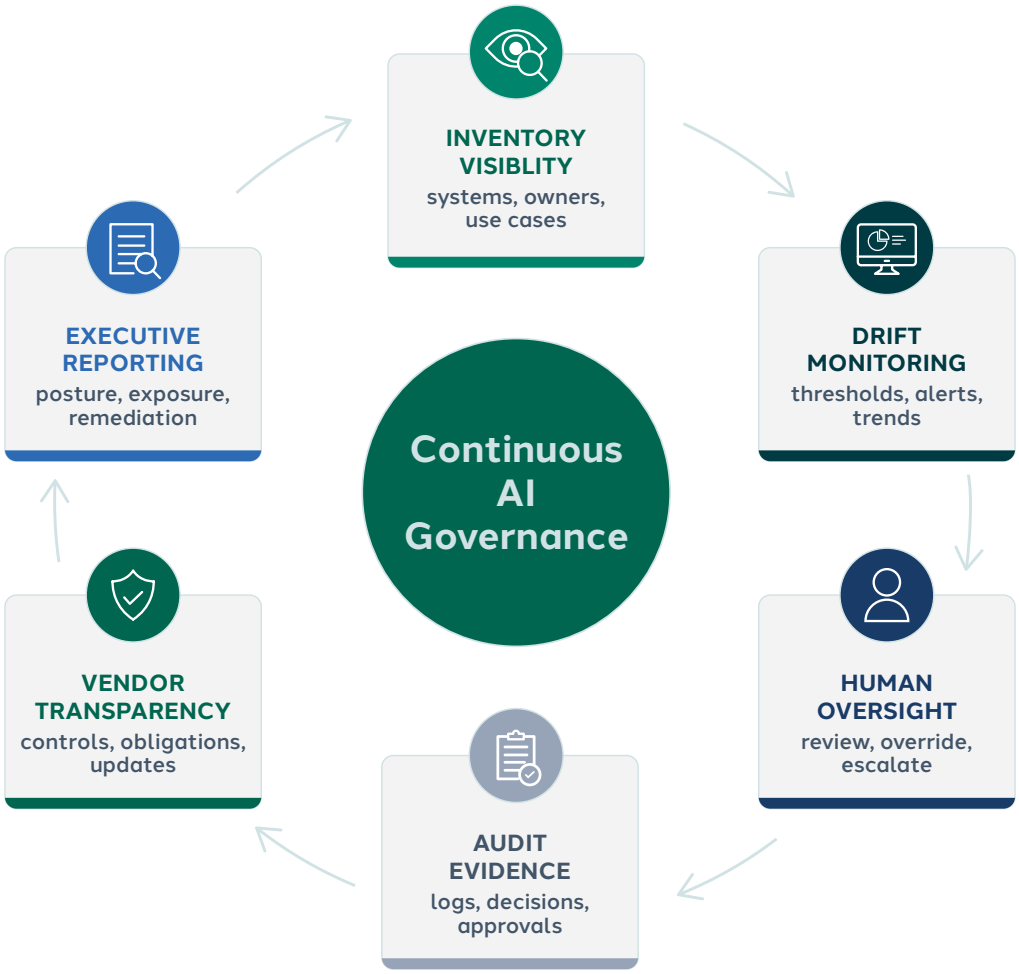

Operational AI governance is not a periodic review exercise. It is a continuous oversight capability that combines visibility, monitoring, accountability, evidence management, and executive reporting into a connected governance operating model.

CONTINUOUS VISIBILITY
AI systems, vendor tools, and GenAI usage must remain continuously inventoried, classified, and monitored as environments evolve.

ACTIVE OVERSIGHT
Governance teams need defined escalation paths, human review standards, and monitoring thresholds that operate in real time.

DEFENSIBLE EVIDENCE
Regulators increasingly expect governance evidence to be centralized, traceable, and available without manual reconstruction.

09 / MRM INTEGRATION

AI governance should extend existing risk infrastructure

Banks are not starting from zero. Most institutions already have mature model risk, operational risk, compliance, vendor management, and audit functions. The challenge is that AI risk does not fit neatly into any one of them. Effective AI governance depends on integrating oversight across existing governance structures rather than creating isolated AI governance programs.

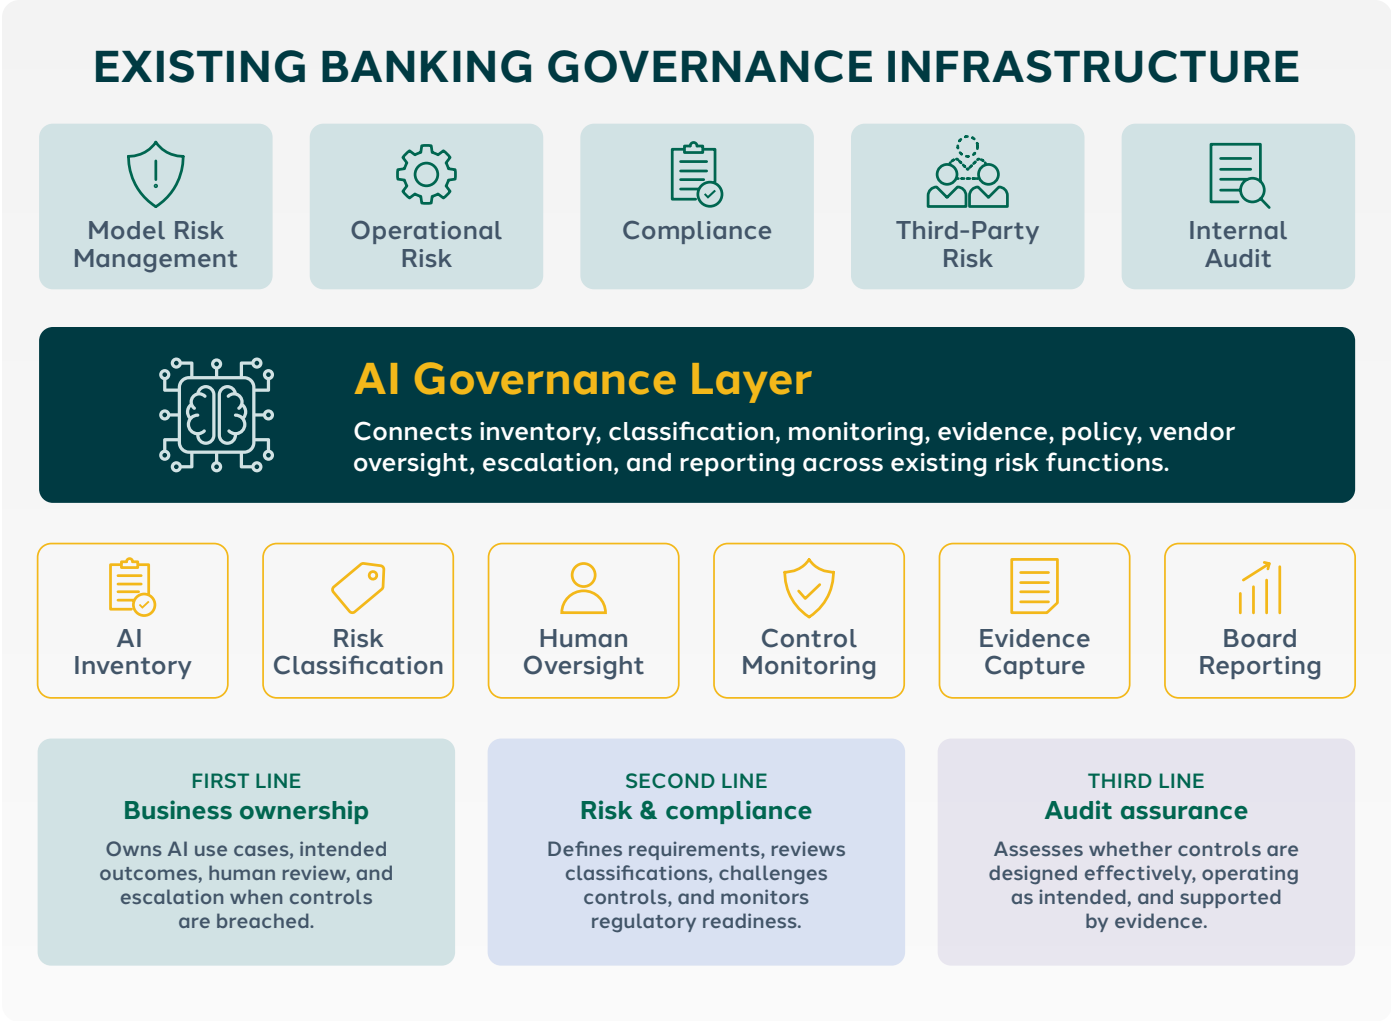

The objective is not to replace MRM. It is to extend MRM with the visibility, workflow, evidence capture, and cross-functional accountability required for AI systems that evolve continuously.

The institutions most prepared for AI regulation will strengthen existing governance infrastructure with clearer ownership, continuous monitoring, centralized evidence, and coordinated oversight across the AI lifecycle.

From AI visibility to regulatory readiness

Banks do not need to operationalize every aspect of AI governance at once. Most institutions progress through a staged maturity journey that begins with visibility and evolves toward control, continuous oversight, centralized evidence, and regulator-ready reporting.

Traditional governance programs were not designed for decentralized AI adoption, embedded vendor intelligence, or rapidly evolving GenAI use cases. As regulatory expectations mature, banks need governance models that can scale from initial visibility to continuous oversight and evidence-based accountability.

VISIBLITY → CLASSIFICATION → ACCOUNTABILITY → CONTINUOUS OVERSIGHT → READINESS

1	2	3	4	5
0-30 Days	30-60 Days	60-90 Days	90-120 Days	Ongoing
Discover AI exposure • Identify models, vendor AI, embedded AI, and GenAI tools • Map business owners and high-impact workflows • Review existing policies and approval paths	**Classify risk** • Define AI risk tiers and classification criteria • Assess use case, customer impact, and data sensitivity • Map systems to regulatory and control requirements	**Assign ownership** • Define business owners and risk reviewers • Establish escalation thresholds and committee approvals • Clarify 1st , 2nd, and 3rd line responsibilities	**Operationalize controls** • Monitor performance, drift, fairness and vendor obligations • Implement issue escalation and remediation • Capture evidence through normal governance workflows	**Prove readiness** • Report AI risk posture to executives and the board • Test controls and close governance gaps • Prepare regulator, audit, and board-ready evidence
OUTPUT Initial AI exposure map and governance gap list	**OUTPUT** AI governance operating model and control requirements	**OUTPUT** Documented accountability model and review workflows.	**OUTPUT** Operationalized AI governance process.	**OUTPUT** Continuous AI governance and readiness reporting.

Readiness is cumulative. Each phase strengthens the bank's ability to identify AI exposure, control high-risk systems, and demonstrate governance with evidence when regulators, auditors, or the board ask.

Turning AI governance into operational advantage

Regulatory readiness is only part of the value. Banks that operationalize AI governance can move faster with confidence, reduce unmanaged exposure, and build greater trust in AI-supported decisions. The advantage comes from connecting governance activities into a repeatable operating model — not managing AI risk through disconnected processes.

CONNECTED AI GOVERNANCE OPERATING MODEL

Centralize AI visibility

Maintain an inventory of AI systems, owners, vendors, use cases, risk ratings, control status, and approval history.

Connect governance workflows

Route reviews through risk assessment, validation, policy approval, third-party review, issue management, and escalation.

Monitor AI risk continuously

Track drift, performance, threshold breaches, ownership gaps, exceptions, vendor obligations, and emerging AI risk indicators.

Capture defensible evidence

Document decisions, approvals, model changes, testing, training, attestations, incidents, vendor reviews, and audit actions.

Report with confidence

Give executives, committees, and boards a clear view of AI risk posture, readiness gaps, remediation progress, and control maturity.

Why it matters beyond compliance

Strong AI governance is not only a defensive control. It becomes the infrastructure that allows responsible AI adoption to scale.

Move faster with confidence

Clear paths for approval, testing, monitoring, and escalation reduce friction for responsible innovation.

Reduce unmanaged exposure

Centralized oversight helps identify shadow AI, vendor opacity, control gaps, and high-risk use cases earlier.

Build trust in AI outcomes

Traceability, explainability, human oversight, and evidence strengthen confidence in AI-supported decisions.

For risk and compliance leaders, the strategic question is no longer whether AI can create value. It is whether the institution can govern AI well enough to capture that value responsibly.

AI governance is becoming a regulatory expectation

AI adoption is accelerating faster than governance maturity. Regulators increasingly expect traceable, evidence-based oversight across the AI lifecycle — including visibility, accountability, human oversight, monitoring, and defensible evidence.

OPERATIONALIZE GOVERNANCE EARLY
Institutions that establish governance foundations now will scale AI adoption more safely and efficiently.

STRENGTHEN REGULATOR CONFIDENCE
Centralized evidence, monitoring, and accountability improve defensibility with regulators, auditors, and boards.

BUILD LONG-TERM TRUST
Responsible AI governance supports operational resilience, customer trust, and sustainable innovation.

Banks that operationalize AI governance now will be better positioned to move faster, reduce unmanaged exposure, and adapt as regulatory expectations continue to evolve.

Learn how SAI360 helps banks operationalize AI governance