The A.I. Framework: A Layered Mosaic for Community Banks

Joe McMann: Co-Founder & CRO
Alec Crawford: Founder & CEO
Frank Fitzgerald: Co-Founder & CTO

Executive Summary

Community banks operate inside fixed constraints of risk, regulation, and trust. Artificial intelligence is already active inside those constraints, whether governed or not. Employees are using A.I. today without uniform oversight, consistent controls, or institutional accountability. That condition does not call for experimentation. It requires structure. Community banks cannot wait for perfect data or external mandates. They must define how intelligence is governed, secured, and applied inside the institution. Confidence emerges when intelligence operates with cohesion, not improvisation.

The A.I. Framework gives community banks that model. It establishes a single architecture for how intelligence is introduced, supervised, validated, and scaled. It replaces fragmented tools with an integrated operating system that unites governance, risk management, compliance, cybersecurity, and human capital. Every A.I.-influenced decision becomes traceable, explainable, and defensible by design.

This Framework advances in four phases:

- A.I. 1.0: establishes boundaries, oversight, and readiness
- A.I. 2.0: introduces governed capabilities through defined use cases that improve accuracy and consistency without introducing unmanaged risk
- A.I. 3.0: aligns the institution by integrating intelligence vertically and horizontally across business lines, replacing silos with a unified platform
- A.I. 4.0: defines the horizon for what enterprise intelligence could become

Most community banks currently exist in A.I. 0 or A.I. 1.0. Some have experimented with isolated tools but lack the infrastructure required to scale safely. This Framework establishes the required direction for institutional progression and disciplined execution. No matter the starting point, the Framework demonstrates how to build structure before scale, capability before complexity, and cohesion before expansion. Hesitation from uncertainty only creates operational, compliance, and reputational risk.

Governed intelligence strengthens community banking by reinforcing institutional authority over risk, decisions, and accountability. It improves accuracy, reduces operational strain, and strengthens compliance. Human capital benefits follow structure. When repetitive work is removed under clear controls, employees spend more time applying judgment where it matters most, and institutions retain talent because tools operate at the same standard as fiduciary responsibility.

This Framework gives community banks a governed path forward. It replaces uncertainty with structure, fragmentation with cohesion, and experimentation with disciplined execution. It prepares institutions to operate in an environment where intelligence is embedded in decisions and workflows without eroding control. Institutions that govern intelligence as an institutional asset will set the standard for their peer group. This Framework exists to establish that control and to enable disciplined advancement under regulatory, fiduciary, and operational accountability.

Table of Contents

INTRODUCTION

Community banks operate under standards that do not bend. Risk, capital, compliance, and trust define these institutions. Regulators expect structure. Customers expect stability. Artificial intelligence will not change those expectations. A.I. is already inside every bank, and its influence continues to grow as intelligence matures. The question is no longer whether A.I. will shape bank operations, but how the institution will shape A.I. through disciplined governance. Community banks cannot wait for external direction or perfect data. They must decide how intelligence will operate, how it will be supervised, and how it will shape decisions across the institution.

Institutions begin this work from different levels of maturity. Some operate without a governing framework. Others rely on isolated tools that solve local problems while expanding institutional exposure. In many cases, employees turn to external systems because internal capability has not been provided. These conditions do not prevent progress, but they do increase risk. Readiness is not optional. Boards and regulators will expect structure before scale. Community banks must protect employees from the exposure created by ungoverned A.I. They require a model that directs intelligence deliberately, protects the institution by design, and moves activity from scattered usage to supervised practice.

This Framework establishes that model. It defines a single operating standard for how intelligence enters the institution, how it is governed, how it is validated, and how its use is documented. It gives Boards a structure they can oversee and executives a system they can execute. It unifies governance, risk, compliance, and cybersecurity under one controlled operating approach (AI GRCC). It replaces improvisation with discipline. It moves intelligence from disconnected tools into a coordinated system built for cohesion, accountability, and control.

The A.I. Framework Defined:

- A.I. 1.0: establishes boundaries, oversight, and readiness
- A.I. 2.0: introduces controlled capabilities through defined use cases that improve speed and quality without exposing the bank to unmanaged risk
- A.I. 3.0: aligns the institution by integrating intelligence vertically and horizontally across business lines, replacing silos with a unified controlled platform
- A.I. 4.0: defines the horizon for what enterprise intelligence could become

Community banks need intelligence that strengthens the institution holistically. Fragmentation erodes oversight. It creates blind spots regulators will not excuse and risks the institution cannot measure. Cohesion strengthens control. It aligns data, policy, and judgment under one governed model. When intelligence shifts from pushed reports to pulled insight, direction is restored. When data becomes a governed asset instead of a distributed dependency, decisions gain precision. Institutions advance when A.I. follows one architecture, one standard, and one source of truth across every workflow it influences.

The objective is disciplined clarity. A.I. does not replace the judgment that defines community banking. It strengthens it by improving accuracy, transparency, and alignment with institutional standards. It strengthens compliance by making boundaries

explicit and risk visible earlier in the process. When intelligence operates within one governed model, decisions become faster to reach and easier to defend. Institutions that succeed will not rely on disconnected tools. They adopt operating discipline. They govern intelligence as an institutional asset. This Framework exists to establish that discipline before scale, complexity, or speed introduce irreversible risk.

A.I. 1.0: Governance, Boundaries, and Institutional Readiness

A.I. 1.0 defines the requirements a community bank must meet before any employee, vendor, or system uses artificial intelligence on behalf of the institution. It is not experimentation. It is preparation. It is the discipline that allows every later phase to move deliberately and without losing control.

The first requirement is an enterprise A.I. policy. This is the bank's authoritative rulebook for intelligence. It sets purpose, scope, and definitions. It identifies which activities qualify as A.I. It establishes boundaries for what A.I. may and may not influence. It defines which data classes are in scope, which are restricted, and which are prohibited. It sets expectations for human oversight, documentation, and escalation. It mandates that no A.I. use case is allowed without approval. It is written so every business line, branch, and employee can understand the rules.

The second requirement is formal institutional governance. Community banks need a standing A.I. governance committee that sits at the same level as other risk and technology forums. It includes senior leaders from risk, compliance, information security, operations, technology, and business. The committee owns policy. It reviews and approves use cases. It classifies them by risk and readiness. It sets the bank's A.I. priorities. It records decisions and assigns accountable owners. It meets on a fixed cadence and produces minutes that can be shown to auditors and regulators.

A.I. 1.0 requires a structured intake and approval process. The bank needs a single front door for A.I. Every idea, from any department, enters through that door. The intake documents the business objective, data involved, proposed model or vendor, affected users, and potential risks. The governance committee uses that information to classify each use case, approve, reject, or send back for refinement. Nothing skips this step. This prevents shadow projects, vendor sprawl, and uncontrolled experimentation.

Risk classification must be defined before execution. Not all A.I. carries the same risk. The bank creates clear tiers such as informational, advisory, decision-supporting, and decision-impacting. Each tier maps to specific requirements for testing, validation, monitoring, and human oversight. A chatbot that answers internal HR questions does not require the same controls as an A.I. agent that drafts credit memos or supports underwriting. A.I. 1.0 requires those distinctions upfront so the bank does not treat every use case as harmless or catastrophic by default.

Vendor and model oversight must be integrated into the program. Community banks rely on external platforms. A.I. 1.0 requires every vendor to be evaluated through the existing third-party risk process, including security, privacy, model behavior, data handling, incident response, and regulatory alignment. The bank maintains an up-to-date inventory of all A.I. vendors, the models in use, and their approved purposes. No department may independently purchase or deploy a new A.I. tool.

Data governance is non-negotiable. Before any A.I. activity, the bank determines which data sources may be used, the conditions for use, and the protections required. That includes customer data, employee data, transaction data, and confidential information. Sensitive fields are identified, and potentially masked, tokenized, or excluded. The policy states where data may flow, where it may not, and whether outputs may persist or be reused. A.I. 1.0 reduces exposure by design, not by reaction.

Cybersecurity must expand to cover A.I. specific risk. The security team defines explicit A.I. controls for access, logging, prompt and output monitoring, anomaly detection, and incident response for misuse or unexpected behavior. The bank determines which users can access which tools, from which devices, and under which conditions. A.I. 1.0 integrates intelligence into the security architecture rather than treating it as an external channel.

Employees must understand the rules before using A.I. A.I. 1.0 requires training and communication across the institution. Staff learn the policy, the risks, the boundaries, and the approved tools. They learn what is prohibited, training is mandatory and maintained as the program evolves.

Monitoring and auditability complete pre-execution readiness. The bank determines how it will evidence control. That includes logging A.I. activity, tracking use cases, recording approvals, and maintaining documentation of risk assessments and testing. The bank defines key risk indicators and controls for A.I. usage. A regulator should be able to ask how A.I. is used and receive a complete, documented answer. A.I. 1.0 ensures intelligence can be supervised before it scales.

A.I. 1.0 concludes with a readiness certification. Before the bank advances to A.I. 2.0, it must answer yes to questions such as:

- Do we have an approved A.I. policy?
- Do we have a functioning governance committee?
- Do we have a single intake process?
- Do we classify A.I. risks?
- Do we control vendor and tool access?
- Do we define which data A.I. may see?
- Do we have AI-specific security controls?
- Do we train employees before use?
- Do we log and document A.I. activity?

If any answer is no, the bank remains in A.I. 1.0. and is not ready to proceed.

A.I. 1.0 is where discipline is established. It is where the bank decides how intelligence will be governed, not just where it will be used. It front-loads work so later phases move faster and with less risk. Community banks that invest in A.I. 1.0 find adoption cleaner, safer, and easier to defend. Those that skip A.I. 1.0 experience A.I. as a series of disconnected tools, compliance surprises, and avoidable stress. No bank should advance to A.I. 2.0 until every requirement in A.I. 1.0 is complete. The Framework begins here because nothing that follows can stand without this foundation. In A.I. 1.0, the bank does not pursue capability. It pursues control. Discipline becomes architecture. Readiness becomes the institution’s advantage.

A.I. 1.0: Operational Execution Requirements

A.I. 1.0 requires more than principles. It requires evidence. Community banks must show the steps they follow, the artifacts they produce, and the controls they put in place before any A.I. use case moves into production. These requirements form the operational core of A.I. 1.0. They define readiness. They enable the institution to defend its decisions to its Board, its regulators, and its customers.

1. Required Governance Artifacts

The bank completes, approves, and stores the following documents:

- A.I. Policy Version 1.0
- A.I. Governance Committee Charter
- A.I. Risk Appetite Statement
- A.I. Use-Case Intake Form and workflow
- A.I. Use-Case Approval Log
- A.I. Risk Classification Matrix
- A.I. Vendor Oversight Checklist
- A.I. Model Inventory
- A.I. Data Governance Map
- A.I. Access Control Matrix
- A.I. Incident Response Addendum
- A.I. Monitoring and Audit Protocol

All these artifacts must exist before any A.I. tool is deployed.

2. Required Cybersecurity Controls

- Logging of all A.I. interactions
- Restriction to approved users, devices, and environments
- Blocking uploads to public or unmanaged systems
- Defined anomaly-detection rules
- Reviewable audit trails
- Secure storage of all A.I. outputs
- Encryption or masking of sensitive fields
- Integration of A.I. activity into the SIEM

These controls must be tested before use.

3. Required Vendor and Model Oversight

- Full third-party due diligence
- Review of data handling, residency, and subcontractors
- Verification of audit logs and access controls
- Confirmation that bank data is not used for model training
- Updated vendor risk rating

No vendor may be used without these steps.

4. Required Workforce Preparation

- Mandatory training on policy, boundaries, and risk
- Clear guide to acceptable and prohibited use
- Examples of acceptable and unacceptable prompts
- Identification of approved tools
- Recorded evidence of training completion

Training precedes execution.

5. Required Communication Across the Institution

- Announcement of policy, governance structure, and expectations
- Publication of the intake process
- Explicit prohibition of shadow A.I.
- Clear escalation paths
- Periodic updates as the program evolves

Communication establishes authority and removes ambiguity across the institution.

6. Required Documentation and Evidence

- Intake records and approval decisions
- Risk assessments and testing documentation
- Governance committee minutes
- Training evidence
- Logs of A.I. activity
- Control-testing evidence

Documented evidence converts governance from intent into defensible control.

7. Required Pre-Execution Sequencing

- Approve A.I. Policy
- Convene A.I. Governance Committee
- Approve Committee Charter
- Publish Risk Appetite Statement
- Establish intake process
- Finalize Risk Classification Matrix
- Complete vendor due diligence
- Map data categories
- Approve cybersecurity controls
- Train employees
- Test monitoring and logging

- Validate documentation processes
- Certify readiness
- Begin A.I. 2.0

Each step must be completed in sequence.

8. Required Readiness Certification:

“The institution has completed all governance, risk, compliance, cybersecurity, data protection, vendor oversight, monitoring, and training requirements for A.I. 1.0. The bank is ready for supervised execution under A.I. 2.0.

A.I. 2.0: Controlled Introduction and Early Execution

A.I. 2.0 begins when the bank moves from preparation into supervised execution. It introduces intelligence into real workflows with defined objectives and disciplined boundaries. It is not scale. It is controlled execution. A.I. 2.0 demonstrates how intelligence behaves inside an institution and how employees respond when direction replaces improvisation. It gives the bank momentum while minimizing risk.

Most banks do not enter A.I. 2.0 from a clean slate. Employees already use public tools in quiet ways. They copy text into external systems, test ideas in uncontrolled environments, and expose the institution to risk without awareness. This elusive use of artificial intelligence is commonly referred to as shadow A.I. Research indicates that approximately 70% of employees report using shadow A.I. in the workplace and 45% of folks who use A.I. for work admit to putting customer or financial data into public large language models. A.I. 2.0 replaces shadow A.I. with supervised A.I., protecting from unintended consequences by offering governed alternatives.

The defining feature of A.I. 2.0 is contained capability within vertical workflows. Each tool improves a single process. Each deployment stands alone. Intelligence remains within the boundaries of a specific task. This produces value without requiring the bank to rebuild its technology stack. It also reveals the limits of fragmentation. A.I. 2.0 improves the work. It does not integrate the institution.

A.I. 2.0 also serves a critical diagnostic function. It exposes how well the institution's governance, risk, compliance, and cybersecurity controls perform under real operating conditions. Gaps become visible. Intake friction appears. Approval delays surface. Logging weaknesses are revealed. These signals are not failures. They are the evidence the institution needs to refine controls before intelligence expands further

Execution in A.I. 2.0 requires discipline. Every tool must satisfy the controls established in A.I. 1.0: governance, intake, risk classification, data restrictions, and security review. No department deploys A.I. because it appears useful. Each deployment requires a defined use case, a documented test, a risk tier, and a clear group of approved users. Execution follows a disciplined pattern: begin narrow, restrict data access, log activity, review outputs, enforce guardrails, and retire tools that cannot be supervised. A.I. 2.0 works when every deployment stays narrow, supervised, documented, and aligned to the controls established in A.I. 1.0.

The most defensible A.I. 2.0 use cases share three traits. They carry negligible risk. They operate at high volume. They remove friction from work that already strains capacity. These early deployments deliver measurable value and give Boards and regulators the evidence they expect before intelligence expands into more complex functions.

Practical Use Cases: Bank-Specific ROI

Below are five disciplined, regulator-aligned use cases. Each demonstrates value, boundaries, and measurable return.

1. **Meeting Intelligence (Transcription, summarization, action extraction)**
   This tool records meetings, produces transcripts, extracts decisions, and generates action items. It reduces administrative burden without touching customer data or regulated analysis.
   Value: faster follow-up, clearer documentation, and more accurate records.
   Risk: exposes internal strategy or personnel discussions if storage rules are weak.
   Oversight: approved users, required redaction, controlled storage, fixed retention periods.

   Bank ROI:
   Meeting intelligence typically saves 30–60 minutes per meeting. A bank running 40–60 recurring meetings weekly regains 25–50 staff hours per week, or $70,000–$140,000 annually at standard labor costs. Documentation accuracy improves 20–40 percent, strengthening audit and examination readiness.

2. **Knowledge Center (Policy summaries, internal research)**
   This tool accelerates access to internal policies, procedures, product rules, supervisory guidance, and training material. It strengthens consistency and reduces reliance on subject-matter experts.
   Value: faster answers, fewer bottlenecks, consistent explanations across departments.
   Risk: inaccurate or outdated outputs that create compliance exposure through inconsistent policy interpretation or supervisory misalignment.
   Oversight: a curated corpus, human verification, and rules that prohibit using A.I. for final regulatory interpretation.

   Bank ROI:
   Research time decreases 40–70 percent. A compliance officer typically regains 6–10 hours per week, equal to $20,000–$35,000 annually. Better first-pass interpretation reduces remediation events costing $10,000–$50,000 each.

3. **Credit Support (Memo drafting, structuring, outline generation)**
   This tool drafts credit outlines, deposit analysis structures, summaries, and risk frames. It accelerates the first-draft process but never replaces underwriting judgment.
   Value: faster starts, standard structure, earlier detection of missing data.
   Risk: incorrect interpretation or unintended inclusion of restricted fields.
   Oversight: strict data rules, template alignment with credit policy, and mandatory human ownership of conclusions.

Bank ROI:
Drafting time decreases 25–40 percent. Lenders save 3–5 hours per memo, equal to 150–350 hours annually per lender. At $75 per hour, this returns $11,000–$26,000 per lender. Rework declines 15–25 percent.

4. **Customer Support / Call Center Enhancement**
   A supervised A.I. agent retrieves answers from internal systems during live calls. It does not speak to customers. It improves accuracy and reduces handling time.
   Value: shorter calls, higher first-contact resolution, and consistent responses.
   Risk: inaccurate guidance or unintended exposure of internal-only information that could create compliance, conduct, or reputational risk during customer interactions.
   Oversight: mandatory human confirmation before any guidance is used, continuous monitoring of agent outputs, documented escalation procedures, and periodic quality review by compliance and operations.

   Bank ROI:
   Handle time decreases 20–35 percent. A ten-agent call center regains 3,000–4,800 staff hours per year, equal to $135,000–$215,000 in capacity. Consistency reduces compliance exposure tied to misstatements.

5. **Internal Drafting & Correspondence**
   A.I. assists with emails, letters, summaries, and internal documentation. All output is reviewed by a human. The tool accelerates work and improves consistency.
   Value: less drafting time, stronger language, better documentation.
   Risk: inadvertent inclusion of sensitive or restricted information, unauthorized tone or content deviations, or use outside approved templates and purposes.
   Oversight: restricted access to approved content sources, enforced use of controlled language libraries, mandatory human review before distribution, and retention of reviewed outputs for audit and supervision.

   Bank ROI:
   Drafting effort drops 50–70 percent for routine tasks. Banks producing 200+ communications annually regain 600+ hours per year, equal to $35,000–$40,000 in capacity. Language quality becomes more uniform.

Consolidated ROI Summary

A.I. 2.0 produces measurable gains across the institution. These gains are real, repeatable, and momentum-creating, but remain only local. Each improvement increases efficiency inside a single workflow. The institution becomes faster but not more integrated. However, efficiency is not intelligence. Local gains do not create governed capability. That requires a more unified approach.

The Cost of Fragmentation

- A.I. 2.0 exposes the cost of fragmented intelligence
- Each tool demands onboarding
- Each tool demands access control
- Each tool demands monitoring
- Each tool demands auditability
- Each tool expands the bank's risk surface
- Financial cost accumulates across vendors and workflows
- Compliance cost accelerates with each additional system
- Oversight cost becomes unsustainable at scale
- Control becomes reactive
- Policies strain
- Human oversight becomes the limit
- Local efficiency creates institutional fragility

Fragmented tools force the bank to supervise A.I. one system at a time while employees operate across all systems simultaneously.

Why A.I. 2.0 Cannot Scale

- A.I. 2.0 is not the destination
- A.I. 2.0 is the constraint
- As tools multiply, data scatters
- As data scatters, oversight weakens
- As oversight weakens, regulatory exposure increases

More tools do not create more intelligence; they create noise; they create risk; they create hidden costs.

A.I. 3.0 provides the alternative. A.I. 3.0 aligns local gains under one architecture that governs intelligence as a system. That is when intelligence becomes institutional and return on investment becomes return on intelligence.

A.I. 3.0: Platformization, Cohesion, and Institutional Intelligence

A community bank enters A.I. 3.0 when intelligence stops improving individual workflows and starts strengthening the institution holistically. This phase replaces fragmentation with cohesion by aligning data, policy, governance, and judgment under one architecture instead of scattered tools. A.I. 3.0 is where the bank establishes a governed environment that integrates intelligence vertically and horizontally. It is the phase most banks aspire to achieve.

A.I. 3.0 begins with platformization: the shift from independent tools to a governed operating system for intelligence. In A.I. 2.0, departments optimized their own work using narrow tools that live inside silos. In A.I. 3.0, the bank replaces that model with a platform that governs all intelligence in one place. It sets rules once and enforces them everywhere. It treats data as shared infrastructure rather than departmental property. It builds the foundation needed for intelligence to operate consistently across the institution.

The platform enables interoperability, not just integration. Interoperability unifies systems. Governed data reinforces other workflows. Insights strengthen judgment across functions. Credit, fraud, risk, compliance, operations, and customer service stop operating on disconnected information. They work from a shared source of truth with a common set of controls. This is the moment when intelligence begins to operate like infrastructure rather than software.

A.I. 3.0 also reflects a regulatory reality: examiners evaluate the institution, not the tool. Fragmented A.I. obscures oversight. It creates blind spots regulators cannot trace and risks the bank cannot measure. A unified platform eliminates those blind spots. It provides one audit trail, one policy engine, one set of guardrails, and one record of decisions. It gives risk, compliance, and cybersecurity full visibility into every model, every agent, every prompt, every output, and every workflow. It gives the Board evidence that governance is not conceptual but operational.

A governed platform becomes the bank's center of data gravity. Intelligence accumulates where it is supervised. Workflows accumulate where they are standardized. Data accumulates where it is protected. As the platform becomes the strategic core of A.I. activity, the value of each additional use case compounds. Every new agent becomes more capable because it draws from an expanding universe of governed intelligence. Every decision becomes more precise because the same standards shape the information behind it. This compounding effect does not occur with distributed tools.

Vertical and horizontal intelligence define A.I. 3.0 as a system:

- Vertical intelligence brings depth: precision inside a department with refined data, structured learning, and consistent execution.
- Horizontal intelligence brings perspective: patterns that span departments and reveal relationships no silo can see alone.

When both dimensions operate together, the bank moves from pushed reporting to pulled insight. Executives no longer wait for information. They extract insight from governed data in real time.

Agentic workflows define the next layer of capability. In A.I. 2.0, tools perform tasks. In A.I. 3.0, agents perform sequences. They coordinate multi-step processes across systems, enforce policy by design, and operate with narrow permissions aligned to the bank's risk appetite. Agents retrieve, validate, enrich, compare, summarize, classify, and recommend. Their actions are deterministic, traceable, and governed. Agents do not replace people. They extend human capability, remove friction from work and elevate judgment.

In this manner, human capital becomes a strategic advantage in A.I. 3.0. A governed platform gives employees the ability to perform meaningful work with less fatigue. It removes administrative weight. It restores time. It improves the work experience by giving people more space to apply expertise and less obligation to complete repetitive tasks. Banks that adopt governed intelligence attract talent. Banks that resist will lose it. People will choose institutions that reduce friction, expand capability, and respect their time. A.I. 3.0 turns the bank into a destination for talent, not a source of attrition.

As part of this, change management becomes an outcome, not an obstacle. People adopt systems that feel consistent, safe, and useful. A governed platform delivers that environment. It reduces fear by establishing boundaries. It builds confidence by producing reliable results. It strengthens adoption because it makes daily work easier, faster, and more accurate. When intelligence operates inside guardrails, people use it with trust.

Ultimately A.I. 3.0 increases decision velocity across every function: credit, fraud, service, onboarding, operations, and risk review. Cycle times fall because information becomes more accurate, more complete, and more aligned. Decisions take less time because judgment is supported rather than strained. Institutions that achieve A.I. 3.0 move faster than peers not by taking more risk but by removing friction from insight.

This is the phase where a community bank gains scale without adding headcount. It is where intelligence becomes coherent. It is where the institution stops experimenting and starts advancing. A.I. 3.0 is what large institutions spend billions to attempt to build. Community banks cannot replicate that investment, but they can reach the same destination through a governed platform built for high-risk environments, one that unifies data, models, agents, controls, and oversight and turns intelligence into infrastructure.

A.I. 3.0 is the community bank objective, where the institution gains leverage, not just efficiency, where A.I. becomes safe, scalable, supervised, and strategic, where community banks regain control over their data, their workflows, and their future. A.I. 3.0 is alignment. A.I. 3.0 is acceleration. A.I. 3.0 is cohesion. It prepares the institution for enterprise intelligence, where intelligence becomes the bank's operating system and every workflow it touches is governed, measured, and strengthened by design.

Director's Checklist

Before a community bank adopts integrated intelligence, its Board must confirm that the institution governs A.I. with discipline. This Director's Checklist establishes the standards that determine whether the bank is ready to function as an A.I. 3.0 institution.

The Checklist provides the Board with a disciplined method to evaluate whether the bank is prepared to move from isolated tools to cohesive, institution-wide intelligence operating under defined governance.

Almost no community bank will advance beyond A.I. 3.0 in the near term. For that reason, this Checklist sets the threshold that must be met before integrated intelligence can be adopted safely, scaled responsibly, and defended to regulators. The institution is ready to proceed only when every answer is yes.

Governance & Oversight

1. Do we have a Board-approved A.I. policy that defines scope, boundaries, and controls?
2. Is there a standing A.I. governance committee with documented minutes and accountable owners?
3. Does every A.I. activity, without exception, enter through a single intake and approval process?

Risk, Compliance & Cybersecurity

1. Have all A.I. use cases been classified into defined risk tiers with associated oversight requirements?
2. Do our cybersecurity controls explicitly cover A.I. access, logging, continuous monitoring, anomaly detection, and incident response across all A.I.-influenced workflows?
3. Does compliance have full visibility into all A.I. activity, including vendor tools and internal workflows?

Vendors, Tools & Architecture

1. Are all A.I. tools across the institution inventoried, approved, governed, and monitored under a single enterprise standard?
2. Do we understand and actively measure the financial, operational, compliance, and control cost of maintaining multiple independent A.I. tools?
3. Have we evaluated whether a platform model would reduce complexity, strengthen oversight, and improve resiliency?

People & Capability

1. Do employees have approved, governed A.I. tools that protect them from the risks of unmonitored use?
2. Are we improving human capital by removing repetitive work and increasing the quality and consistency of output?

Readiness & Maturity

1. Can management articulate which phase of the A.I. Framework we occupy today and why?
2. Do we have defined criteria for advancing from A.I. 1.0 to A.I. 2.0 to A.I. 3.0?
3. Can we provide examiners with unmistakable evidence of control approvals, testing, documentation, and monitoring?

Strategic Direction

1. Are we building toward cohesion, or are we accumulating tools faster than we can govern them?
2. Is A.I. improving decision quality, strengthening trust, and reinforcing control, not just increasing speed?

Catalyst Scorecard: Readiness for Enterprise Intelligence

Community banks advance unevenly through A.I. 1.0, A.I. 2.0, and A.I. 3.0. Some establish governance before deploying a single tool. Others deploy tools before understanding the risk they introduce. Some build pockets of excellence that never extend beyond a single department. A.I. 3.0 creates institutional cohesion. Enterprise Intelligence refers to A.I. 4.0: the phase where governed intelligence operates as an institution-wide system rather than a collection of tools.

The Catalyst Scorecard tests whether that cohesion is strong enough to support enterprise intelligence, but enterprise intelligence is not the next step for every bank. It is the next step for the bank that has earned it. The Catalyst Scorecard gives Boards and executives a disciplined method to determine whether the conditions for scale exist. It clarifies what must already be in place, what must be strengthened, and what must be corrected before intelligence operates across the institution without increasing risk.

This is an institutional assessment that measures governance, data discipline, workflow consistency, oversight, and cultural readiness. It prevents the bank from moving faster than its controls or slower than its opportunity. Each block represents a discrete readiness dimension that links aspiration to execution.

1. Governance Foundation

Required Standard:
• A.I. policy, committee, intake, and risk tiers operating in practice with documented enforcement, escalation, and evidence of use
Evidence of Readiness:
• Approved policy
• Committee minutes
• Intake records
• Risk tier definitions
Common Failure Modes:
• Irregular committee cadence
• Policy unenforced
• Intake bypassed
Consequence:
• A.I. grows faster than oversight; exam criticism

Score: Yes / Partial / No

2. Unified Data Controls

Required Standard:
• Clear, institution-wide definitions for data classes, permitted use, and prohibited exposure, enforced consistently across all A.I. workflows
Evidence of Readiness:
• Enterprise data dictionary mapped to A.I. use cases
• Sensitivity classifications enforced at ingestion, processing, and output

• Masking, tokenization, and field-level restrictions applied by architecture, not user discretion

Failure Modes:

• Rules differ by business line

• Unclear prohibited fields

Consequence:

• Data leakage; privacy violations

Score: Yes / Partial / No

3. Cross-Functional Coordination

Required Standard:

• Risk, Compliance, IT, Security, and Operations hold formally assigned roles, decision authority, and accountability within the A.I. governance structure, with cross-functional concurrence required for all A.I. approvals and deployments

Evidence of Readiness:

• Formal RACI or equivalent role-mapping for A.I. governance decisions

• Use-case approval records showing required sign-off from Risk, Compliance, IT, and Security

• Governance committee minutes evidencing cross-functional challenge, escalation, and resolution

Failure Modes:

• A.I. decisions made within technology or business units without enforceable risk and compliance veto authority

• Cross-functional review treated as advisory rather than mandatory, allowing deployments to proceed despite unresolved objections

Consequence:

• Examiner findings citing weak governance, unclear accountability, and inability to evidence enterprise-level oversight of A.I. decision-making

Score: Yes / Partial / No

4. Vendor & Model Oversight

Required Standard:

• The institution owns a complete and authoritative inventory of all A.I. vendors, models, and embedded A.I. capabilities, each approved for a defined purpose and risk tier before use

• Lifecycle governance is mandatory for every approved vendor and model, including change control, ongoing suitability review, and formal decommissioning

Evidence of Readiness:

• Centralized vendor and model inventory covering all direct, embedded, and inherited A.I. capabilities

• Completed third-party risk assessments that explicitly evaluate A.I. behavior, data handling, model change risk, and downstream use

• Documented approval records defining model purpose, permitted use cases,

limitations, and escalation thresholds
• Documented review cadence demonstrating continued oversight beyond initial approval
Failure Modes:
• Business units procure A.I.-enabled tools outside formal vendor risk processes
• Embedded or “hidden” A.I. capabilities are not disclosed, inventoried, or reviewed after deployment
Consequence:
• Unmanaged third-party exposure; inability to demonstrate institutional control over outsourced intelligence

Score: Yes / Partial / No

5. Visibility & Auditability

Required Standard:
• The institution can reconstruct any A.I.-influenced decision end-to-end, including prompts, data inputs, model or agent actions, outputs, and human approvals, without manual reassembly
• Auditability is continuous, centralized, and independent of individual tools or vendors
Evidence of Readiness:
• Centralized logging of prompts, outputs, actions, and system decisions across all A.I. usage
• Time-stamped records linking A.I. activity to users, use cases, and approval artifacts
• Version control for prompts, models, policies, and workflows
• Audit-ready evidence that can be produced on demand without vendor dependency
Failure Modes:
• Logs fragmented across tools, vendors, or business lines
• Reliance on screenshots, exports, or after-the-fact reconstruction
• Visibility limited to activity metrics rather than decision traceability
Consequence:
• Inability to evidence control during examination; supervisory criticism regardless of intent

Score: Yes / Partial / No

6. Consistency of Controls

Required Standard:
• A single, enforceable control framework governs all A.I. usage across the institution, regardless of business line, tool, vendor, or deployment model
• Guardrails are defined once, applied universally, and enforced by architecture rather than discretion, with overrides permitted only through documented approval, escalation, and time-bound exception
Evidence of Readiness:
• Centralized access control rules applied consistently across all A.I. tools and

workflows
• Uniform data restrictions, usage boundaries, and approval requirements enforced by architecture, not policy alone
• Central control plane or equivalent mechanism demonstrating consistent enforcement across vendors and internal systems
• Documented exception process with approvals, rationale, and expiration dates

Failure Modes:
• Controls defined at a policy level but implemented differently by tool or department
• Vendors enforcing their own rules instead of institutional standards
• Exceptions granted informally or left in place indefinitely

Consequence:
• Uneven risk posture; inability to defend why similar A.I. activities are governed differently across the institution

Score: Yes / Partial / No

7. Horizontal Integration Capacity

Required Standard:
• The institution enables cross-functional data use **only** through governed, approved integrations that preserve data classification, access restrictions, **purpose limitation**, and auditability
• Horizontal data movement is intentional, documented, **use-case specific**, and limited to purposes explicitly approved through A.I. governance

Evidence of Readiness:
• Approved data-sharing rules that define which data may move across functions and under what conditions
• Governed integration mechanisms that enforce access controls, logging, and usage restrictions across domains
• Demonstrated cross-domain insights produced from governed data sharing that can be traced to approved use cases, access rules, and audit logs, without violating data ownership or privacy boundaries

Failure Modes:
• Data technically shareable but not governed
• Informal data movement between departments without documented approval
• Horizontal insight dependent on manual exports or workarounds

Consequence:
• Intelligence remains trapped in verticals; horizontal insight exists without defensible controls

Score: Yes / Partial / No

8. Agentic Workflow Readiness

Required Standard:
• The institution has fully documented, standardized, and repeatable workflows suitable

for supervised agent execution
• Each workflow has defined inputs, steps, decision points, outputs, and accountable human ownership
• Agent execution is prohibited in any workflow where variability, discretionary judgment, or exception handling is not formally defined and governed

Evidence of Readiness:
• Approved process maps identifying agent-eligible steps versus human-only judgment points
• Step-by-step SOPs aligned to policy and risk tier
• Documented workflow owners responsible for accuracy, escalation, and outcome quality
• Clear separation between advisory agent output and final human decision authority

Failure Modes:
• Agents deployed on undocumented or inconsistent processes
• Employees executing the “same” task differently across roles or locations
• Agents introduced to compensate for broken, undefined, or inconsistently executed workflows

Consequence:
• Agent behavior becomes unpredictable; errors cannot be traced to process or control failure
• Supervisory findings tied to unsafe automation and lack of accountable ownership

Score: Yes / Partial / No

9. Human Capital Readiness

Required Standard:
• Employees are trained, permissioned, and governed based on role-specific A.I. authority
• Use of A.I. is explicitly tied to policy, approved tools, and defined responsibilities
• System access to A.I. tools is conditional upon completion of required training, documented acknowledgement of standards, and ongoing compliance with usage requirements
• No employee may use A.I. on behalf of the institution without documented training and acknowledgement of standards

Evidence of Readiness:
• Role-based training programs aligned to risk tiers and permitted use cases
• Training completion records programmatically linked to system access, tool permissions, and continued eligibility for A.I. use
• Clear employee attestations acknowledging A.I. policy, boundaries, and prohibited behavior
• Adoption metrics demonstrating use of approved tools and decline of unapproved alternatives

Failure Modes:
• Training delivered generically without role differentiation
• Employees using unapproved tools due to unavailable, unclear, or weakly enforced

governed alternatives
• Cultural normalization of policy exceptions or informal workarounds

Consequence:
• Shadow A.I. persists; employee behavior becomes the institution's largest unmanaged risk
• Loss of examiner confidence due to inability to demonstrate behavioral control

Score: Yes / Partial / No

10. Risk Appetite Alignment

Required Standard:
• The institution has explicitly defined the permissible role of A.I. in recommendations, decision support, and execution, with final decision authority always retained by a designated human role
• Risk appetite boundaries specify where A.I. may inform judgment, where it may recommend actions, and where it is prohibited from influencing outcomes
• All A.I. use cases are mapped to materiality thresholds and escalation requirements consistent with the institution's Risk Appetite Statement, with escalation triggered automatically when defined thresholds are met or approached

Evidence of Readiness:
• Risk Appetite Statement that explicitly addresses A.I. usage and decision influence
• Documented materiality thresholds defining when A.I. output requires human review, approval, or override
• Use-case documentation showing alignment between A.I. function and approved decision authority
• Escalation protocols embedded in workflows and systems, triggered automatically when A.I. activity approaches or exceeds defined risk limits, and recorded as auditable events

Failure Modes:
• A.I. influence expands beyond documented authority through informal practice, undocumented exceptions, or human deference to model output
• Risk appetite defined at a conceptual level but not operationalized in workflows
• Human reviewers defer to A.I. outputs without understanding limits or confidence boundaries

Consequence:
• A.I. participates in decisions beyond board-approved tolerance
• Regulatory findings tied to unmanaged decision influence and unclear accountability

Score: Yes / Partial / No

11. Decision Velocity Infrastructure

Required Standard:
• The institution's workflows, data pipelines, and approval structures are deliberately designed to increase decision velocity through pre-approved paths, without bypassing or weakening controls

• Decision speed increases are achieved only through standardization, automation, and pre-approved pathways, with no reliance on ad hoc judgment, discretionary shortcuts, or informal escalation
• A.I. acceleration operates within defined intake, review, and escalation mechanisms

Evidence of Readiness:
• Standardized intake workflows with defined service-level expectations by risk tier
• Clean, validated datasets consistently available to approved A.I. use cases
• Documented approval paths that scale with volume without adding manual bottlenecks
• Metrics demonstrating reduced cycle time at increased volume without growth in exceptions, control breaks, or manual intervention

Failure Modes:
• Faster outcomes dependent on individual intervention, informal prioritization, or reviewer discretion rather than system design
• Inconsistent turnaround times across similar use cases
• Control steps skipped to maintain perceived speed

Consequence:
• A.I. fails to deliver sustainable acceleration
• Speed gains collapse under scale or trigger supervisory concern

Score: Yes / Partial / No

12. Cybersecurity Coverage

Required Standard:
• The institution's cybersecurity program explicitly incorporates A.I. specific threats, misuse scenarios, and abnormal behavior patterns into its core detection, monitoring, and response framework
• Security controls owned by the cybersecurity function extend to prompts, outputs, agent actions, model access, and data movement associated with all A.I. usage
• A.I. activity is monitored continuously and integrated into the bank's incident detection and response processes

Evidence of Readiness:
• Centralized logging of prompts, outputs, agent actions, and access events
• Defined alerts for anomalous behavior, misuse, policy violations, and unexpected model actions
• Incident response procedures that explicitly include A.I.-related events
• Continuous monitoring of A.I. activity integrated into the SOC or equivalent monitoring function, with defined alert thresholds, triage procedures, and response ownership

Failure Modes:
• Reliance on traditional perimeter or endpoint controls without A.I.-specific visibility
• A.I. activity excluded from security monitoring, treated as application noise, or assumed to be governed solely through policy rather than detection and response
• Security teams unable to interpret or respond to A.I.-driven incidents

Consequence:
• Misuse or compromise goes undetected
• Delayed response increases operational, regulatory, and reputational impact

Score: Yes / Partial / No

13. Enterprise Consistency

Required Standard:

• A single, authoritative A.I. standard governs all intelligence usage across the institution and supersedes all local interpretations, practices, or tool-specific rules regardless of business line, department, vendor, or deployment model
• Interpretations, controls, and enforcement are consistent enterprise-wide and cannot be redefined locally without formal approval
• Exceptions are rare, formally documented, time-bound, and explicitly approved through governance, with defined expiration and mandatory re-evaluation

Evidence of Readiness:

• One unified A.I. policy applied consistently across all lines of business
• Standardized interpretations and implementation guidance issued centrally
• Documented exception register with rationale, approving authority, and expiration date
• Evidence that controls and standards are enforced uniformly in practice

Failure Modes:

• Lines of business interpret, modify, or operationalize A.I. policy independently without formal governance approval
• Controls vary by department, vendor, or use case without formal approval
• Exceptions accumulate without review or sunset

Consequence:

• Fragmentation re-emerges
• Governance weakens as standards lose authority
• The institution cannot demonstrate consistent control to regulators

Score: Yes / Partial / No

14. Scalability Without Friction

Required Standard:

• The institution can deploy new A.I. use cases quickly because governance, controls, and documentation are pre-defined, standardized, and enforced
• Scale is achieved through repeatable, risk-tiered processes with predefined approval paths, not ad hoc approvals or manual workarounds
• Speed increases because structure already exists

Evidence of Readiness:

• Defined intake-to-approval timelines that are consistently met
• Standardized onboarding, monitoring, and control patterns reused across use cases
• Automation supporting approvals, monitoring, logging, and reporting
• Clear ownership for each stage of deployment, from intake through ongoing oversight

Failure Modes:

• New use cases require bespoke approvals each time
• Scaling depends on individual reviewers or informal coordination

• Manual processes create bottlenecks as volume increases
• Governance slows innovation instead of enabling it
Consequence:
• A.I. momentum stalls
• Business units bypass controls to maintain speed
• The institution cannot scale intelligence without increasing risk

Score: Yes / Partial / No

15. Institutional Ownership of Data

Required Standard:
• The bank exercises direct operational control over its data, models, prompts, guardrails, and audit evidence at all times
• Intelligence operates inside infrastructure and architectures where the institution defines and enforces access, retention, monitoring, security, and data residency
• No critical intelligence function depends on vendor-controlled data custody or opaque processing
Evidence of Readiness:
• A.I. workloads operate within bank-controlled or contractually protected private environments
• Clear data ownership and custody provisions documented for every A.I. vendor and model
• The bank can retain, export, reconstruct, and audit all prompts, outputs, logs, models, and decision artifacts independently of any vendor system or permission
• Exit and transition plans exist to prevent loss of data, models, or audit history if a vendor relationship ends
Failure Modes:
• Reliance on SaaS tools that store or process data outside the bank's control
• Vendor terms limit visibility, retention, or audit access
• Inability to migrate intelligence assets without disruption or loss of evidence
Consequence:
• Loss of institutional control over intelligence
• Increased regulatory and compliance exposure
• Long-term dependency that constrains governance, scale, and resilience

Score: Yes / Partial / No

Interpreting A Bank Score:

12–15 "Yes" Responses: Platform-Ready (Very, very few will be here)

The institution has established the governance, data control, and operational discipline required for enterprise intelligence. Intelligence can operate as a system rather than a collection of tools. A governed platform will scale safely, compound value, and accelerate decision quality. The bank is positioned to advance toward A.I. 4.0.

8–11 “Yes” Responses: Strength with Gaps (A few will be here)

The institution has made meaningful progress but has not yet achieved cohesion. Governance exists but is uneven. Controls function but are not yet universal. A platform model will resolve fragmentation and enforce consistency, but only if identified gaps are addressed first. Advancement without remediation increases risk.

5–7 “Yes” Responses: Early Maturity (Many will be here)

The foundation remains incomplete. Governance, data discipline, and oversight are present in pockets but not institutionalized. Intelligence operates locally without systemic control. The bank must strengthen structure, clarify authority, and close material gaps before enterprise intelligence is viable.

0–4 “Yes” Responses: Structural Risk (Most will be here)

A.I. usage has outpaced governance. Oversight is reactive or absent. Risk, compliance, and control gaps are measurable and defensible only by chance. The institution should return to A.I. 1.0 and A.I. 2.0 to establish foundational discipline before pursuing scale.

Board-Level Guidance

Community bank Boards should evaluate A.I. through structure, not ambition. A high score reflects readiness. A low score reflects risk. Both outcomes provide clear direction for oversight. The Catalyst Scorecard defines what must exist before intelligence can operate at scale. Advancement is earned through discipline. Accelerated intelligence is not a leap. It is the result of control, cohesion, and deliberate execution.

A.I. 4.0: Enterprise Intelligence

A.I. 4.0 remains on the horizon. Few community banks will ever reach it, and none should feel pressure to pursue it. This phase defines what fully governed intelligence becomes when it operates as institutional infrastructure rather than a collection of workflows. It represents a long-term direction that informs present architectural decisions rather than a near-term objective to be pursued.

In A.I. 4.0, intelligence operates under a single architecture. Models, agents, data, policies, and controls follow one standard. Every workflow that uses intelligence operates within the same boundaries, audit trail, and security framework. Decisions become faster to reach, easier to explain, and simpler to defend because each component adheres to a unified operating model. A.I. 4.0 builds intelligence above the core rather than forcing the institution to rebuild it.

This phase introduces digital workers and autonomous agents as supervised participants in institutional workflows. These agents do not operate independently or replace judgment. They execute narrowly defined tasks, follow prescribed sequences, and operate with explicit authority and constraints set by policy. Their value comes from consistency, speed, and traceability, not autonomy. Human ownership remains explicit at every decision point that carries fiduciary, regulatory, or reputational consequence.

A.I. 4.0 changes how institutions scale by introducing supervised digital workers and autonomous agents as institutional labor, not experimental automation. These agents operate inside predefined authority, execute repeatable work, and escalate judgment rather than replace it. They support underwriting preparation, operational execution, fraud triage, and service workflows while remaining continuously monitored, logged, and governed. Their value is additive. They extend capacity without altering accountability. Human judgment remains the decision-maker. Agents remain the instrument.

This shift redefines human capital. When employees spend more time applying expertise and less time managing process, performance improves. Institutions that modernize workflows through governed intelligence attract and retain talent. Efficiency is not the objective; leverage is. A governed platform creates leverage that individual tools cannot provide.

Enterprise maturity does not eliminate innovation. New capabilities will continue to emerge. A.I. 4.0 does not restrict adoption. It disciplines it. Every new capability enters through the same architecture, operates under the same controls, and produces consistent audit evidence. Innovation strengthens the institution without reintroducing fragmentation or unmanaged risk.

Understanding A.I. 4.0 matters even for banks that never reach it. It establishes direction. It aligns present decisions with future architecture. It prevents short-term gains from creating long-term fragility. The challenge of A.I. 4.0 is not construction. It is sustained governance, continuous investment, and permanent operational discipline. A.I. 4.0 is aspirational by design. It provides clarity without pressure and direction without deadlines.

Conclusion: Direction, Discipline, and the Path Forward

Community banks are now stewards of intelligence by default. Artificial intelligence has already entered the institution. The only remaining question is whether it will operate under bank-defined governance or outside it. Power without structure creates risk. Intelligence without discipline creates exposure. Institutions that endure will not be defined by speed. They will be defined by cohesion, confidence, and control.

The A.I. Framework exists because intelligence does not mature evenly. It must be directed. Each phase establishes a necessary condition for safe advancement.

- A.I. 1.0 establishes governance and boundaries
- A.I. 2.0 introduces capability without surrendering oversight
- A.I. 3.0 aligns intelligence across the institution as the operating system
- A.I. 4.0 defines the horizon that informs every architectural decision made today

These are not optional stages. They are the structural requirements for intelligence to operate inside a regulated institution without compromising trust.

Artificial intelligence does not replace the relationships that define community banking. It intensifies their responsibility. Every model, agent, output, and workflow now carries fiduciary, regulatory, and reputational consequence. Governed intelligence strengthens judgment by making decisions explainable, traceable, and defensible. It strengthens compliance by enforcing boundaries before risk materializes. It strengthens human capital by removing friction and returning time to work that requires expertise, accountability, and leadership.

This Framework exists because trial and error is no longer survivable. Regulators will not tolerate unmanaged intelligence. Customers will not forgive it. Employees will not remain in institutions that ask them to operate without protection, clarity, and standards. Competitive advantage will not emerge from fragmented tools or borrowed intelligence. It will emerge from ownership, structure, and disciplined execution.

Community banking has always been built on trust. Artificial intelligence raises the cost of maintaining it. Institutions that govern intelligence as an institutional asset will retain talent, sustain relevance, and compound advantage. Institutions that do not will discover too late that intelligence has already reshaped their risk profile. The future will not reward ambition. It will reward cohesion that builds confidence and sustains control under disciplined leadership.

About the Author:

Joe McMann
Co-Founder and Chief Revenue Officer - Artificial Intelligence Risk (AIR)

Joe McMann is a lifelong entrepreneur and former investment banker whose work is focused on making artificial intelligence safe, secure and compliant for financial institutions. At AIR he leads growth and partnerships across community banks, credit unions, wealth and asset management firms. His approach helps organizations harness agentic artificial intelligence through AIR's framework for AI GRCC: Governance, Risk Management, Regulatory Compliance and Cybersecurity, helping to restore trust in data-driven decision making and innovation accountability at every strategic level.

LinkedIn: linkedin.com/in/joemcmann
Email: joemcmann@aicrisk.com