

START SMALL, BUILD RIGHT

A Practical AI Adoption Path for Community Banks


Yulia Boozer

Founder, Iskratek Consulting




Where is your bank today?

AI adoption looks very different from one organization to another.



The point is not maturity. The point is that banks are entering AI from very different places.

AI today and where it may be headed



What is different is that AI, especially once it starts participating in work rather than merely generating text, does not behave like the traditional deterministic software banks are used to.

So how do we approach this successfully?

AI does not erase what we know about technology adoption. It changes where our familiar instincts need to be adjusted.

01

Recognize that the mental model needs to shift

Keep what still works. Reconsider the assumptions AI changes.

02

Understand what has changed

The familiar lifecycle remains, but the mechanics inside it are different.

03

Understand what AI introduces that is new

New design and operating questions appear once intelligence enters the workflow.

Then make practical choices about where and how to start.

What has changed?

AI pilots follow the same high-level stages as traditional software, but the core mechanics change dramatically.

Dimension	Traditional software	AI
Success / testing	- Defined behavior and expected outputs - Functional pass / fail testing - If it works only 90% of the time, it is generally considered broken	- Outputs are not always identical or perfectly predictable - Accuracy is measured across many examples 90% accuracy may be very strong for some tasks - Success depends on task, risk, and consequence of error
Data	- Data feeds predefined logic - Schemas, mappings, and APIs - Valid inputs drive expected behavior - Business rules encoded in software	- Context and domain meaning directly shape the answer - Meaning matters, not just structure - Different evidence can change the result - Grounding affects accuracy and hallucination risk
Security / risk	- Role-based access and permissions - Encryption and infrastructure security - Application and network controls - Known system boundaries	- Model/data exposure and retention terms - Sensitive data entered into prompts/context - AI authority and action boundaries - Bias, unsafe, or unauthorized behavior
Cost	- Predictable licensing or subscription costs - Infrastructure sized in advance - Cost changes relatively slowly - Easier to forecast at scale	- Usage-driven model and API costs - Cost varies by model, context, and volume - Tool calls and retries can add cost - Pilot economics may change significantly at scale
User role	- Learn the designed process - Follow defined steps - Execute consistently - Report system errors	- Verify and critique outputs - Judge when to trust or challenge - Correct or escalate when needed - Avoid over-trust and blind skepticism

What is new?

01

Data foundations

Information designed for reporting may not preserve what AI needs for reasoning.

AI needs more than clean data. It needs the context and meaning required for the work.

02

Human–AI operating model

Where may AI interpret, recommend, or act, and where does a person decide?

Decision boundaries must be designed, not discovered after an incident.

03

Operational reality

AI behavior, knowledge, and context require ongoing ownership.

Deploying intelligence is not a project. It is an operating capability.

04

Building beyond the first use case

One successful implementation does not automatically create repeatable AI capability.

A pilot proves the concept. It does not prove the foundation.

Don't get derailed at the starting line

This is a common mistake, not the right path.

Existing gaps and workarounds

Missing software capabilities

Vendor limitations

Enhancement requests

Integration gaps

Manual workarounds

Technology wish list

becomes

Wrong starting point


“Could AI fix this?”

Old gaps get relabeled as AI use cases, even when ordinary software or deterministic automation may be the better answer.

Don't use AI simply because AI can perform the task. Use it where intelligence adds something ordinary software cannot.

Start by deciding what actually requires intelligence


Don't force-fit AI.



Where are people spending significant time gathering information, reconciling systems, interpreting what matters, following up and making sure work gets completed?


Five gates for a bank-ready AI use case

These are not stages of a pilot. They answer one question: does this deserve to become our first AI use case?



Let's put it into practice.

Which AI opportunity would you choose, and does your answer change once you apply the filter?



What “Build Right” looks like from the start

Built into the AI solution



Data, evidence, and fallback behavior are explicitly defined

- Authoritative sources
- Required context
- Data freshness
- Missing/conflicting evidence
- Safe fallback / escalation
- Documented in solution config



Guardrails, privileges, and human oversight are designed in

- User permissions
- Data access limits
- Action boundaries
- Human approvals
- Policy guardrails
- Stop / escalate / refuse



Traceability and monitoring are built in

- Source traceability
- Tool / action history
- Human approvals & overrides
- Decision lineage
- Usage monitoring
- Cost controls
- Usage caps / limits

Established by the organization



The baseline, success measures, and evaluation approach are established before launch

- Current-state baseline
- Success thresholds
- Agreed test scenarios
- Automated evals
- Regression testing
- Review cadence
- Continue / correct / stop criteria



Ongoing ownership and monitoring are established


- Named owner
- Exception handling
- Behavior review
- Model output quality monitoring
- Deterioration detection
- Defined human response
- Automated rollback triggers

How do you know it actually works?

The conversion decision rests on evidence, not impression.

Evaluation dimension	Core focus	Key metrics & signals
Output Quality & Task Success	Does the AI complete the intended work correctly, completely, and usefully?	Task success rate; accuracy by task; completeness; correct tool/action execution; workflow/action integrity; behavioral consistency; error or unsupported-output rate
Grounding & Traceability	Is the output supported by authoritative evidence, and can we trace it back to the sources, decisions, and actions that produced it?	Groundedness; source attribution; evidence coverage; authoritative-source use; source freshness; unsupported claims; traceable source/action path; decision lineage
Human Impact & Adoption	Does the AI improve the workflow for the people using it?	Net time saved after review and correction time; edit/override rate; adoption and repeat usage; user feedback; workflow friction
Unit Economics & TCO	Does the value justify the ongoing cost of operating the AI?	Cost per successful task/workflow; cost distribution and percentiles (median, p95); human-review cost; model/API usage; support and maintenance effort
Technical Performance & Reliability	Does the system perform dependably under real operating conditions?	Availability; latency distribution / p95; failure rate; retrieval/tool execution reliability; fallback and recovery behavior
Risk, Safety & Appropriate Behavior	Does the AI stay within the evidence, authority, policies, and boundaries the bank defined?	Correct escalation; correct refusal; policy and permission adherence; sensitive-data handling; unsupported or unauthorized action rate; relevant bias/compliance tests

START SMALL, BUILD RIGHT



**The first use case isn't just about proving AI can work.
It's about building the bank's confidence and capability
to use it well.**

Yulia Boozer | Iskratek Consulting