SmarterFi.ai
AI for Community Banking



WHITE PAPER · NOVEMBER 2026

AI That Compounds

Why financial institutions need a governed learning foundation, not more point solutions

> How SmarterFI built one shared, governed foundation for community banking that every application plugs into, and that gets smarter with use

By Jamie Donahue, Founder, President & CAIO, SmarterFI · Palm Beach, Florida

Cite as: Donahue, J. (2026). AI That Compounds: Why financial institutions need a governed learning foundation, not more point solutions. SmarterFI. Version 1.0. CC BY 4.0

One foundation
every app uses the same governed data, models and audit trail

Learns with use
outcomes improve each institution's own models

Numbers from engines
words from the language model

Examinable
every result traced, recorded and reviewable

For community bank and credit union CEOs, CIOs, CROs, boards, and risk, compliance and audit leaders

© 2026 SmarterFI. Licensed under CC BY 4.0.

1 Executive Summary

Most artificial intelligence in financial services today arrives one tool at a time: a chatbot for the contact center, a document extractor for loan operations, a fraud score from one vendor, a reporting assistant from another. Each is sold on an efficiency case, and some deliver it. But a collection of point solutions does not add up to an AI capability. Each tool starts from zero, copies bank data into its own store, brings its own model to validate and its own contract to manage, and leaves its own audit trail. The institution ends up renting intelligence it never owns, and the tools get no smarter for having been used.

This paper argues for a different model: a **governed learning foundation**. One shared layer holds the institution's mapped data, its models, its controls and its audit trail. Every application plugs into it, and each new application starts from what the last one built. As staff use the applications, their decisions and outcomes, recorded under the institution's control, become the material for better institution-specific models, released only after testing, independent challenge and human approval.

> **Our position**
> The question for a board is not “which AI tool should we buy next?” but “what are we building that gets better every time we use it?” Efficiency is a result of a good foundation, not a substitute for one.

SmarterFI was built around that idea, and it is connected to 13 different banking cores today, so the foundation starts from live institution data. Our foundation is designed around three model layers behind one governed gateway:

- **LQM (large quantitative model), running in production use today:** tested, deterministic calculators for every figure, plus learned quantitative models for estimates such as loss rates and deposit behavior.
- **LBM (large behavioral model):** the behavioral layer SmarterFI is building, trained per institution to anticipate customer, account and portfolio behavior for staff to act on.
- **LLM (large language model):** a licensed, self-hosted open-weight model that reads, retrieves and explains, and never produces the numbers.

In the design, each institution gets its own dedicated environment, with its own subscription, network, keys and data stores, in SmarterFI's private cloud. Every result is traced to the data, model version and approver behind it, and written to a write-once audit store. The design is built to fit the way banks already govern models: the agencies' revised model risk guidance (SR 26-2) and the interagency third-party guidance give boards and examiners a familiar yardstick for the quantitative layers, while generative AI is governed under the institution's own risk framework.[1][2][3]

2 AI in Financial Services Today

AI is no longer new to financial services. Banks and credit unions use it for customer service, fraud detection, credit decisions and back-office work, and the regulators themselves increasingly use it in supervision.[4] The question is no longer whether institutions will use AI. It is what they are accumulating as they do.

A pattern of point solutions

For most institutions, and especially for community banks and credit unions, adoption has followed the path of least resistance: buy a tool for a specific task, connect it to the data it needs and measure the time saved.

© 2026 SmarterFI. Licensed under CC BY 4.0. · www.smarterfi.ai

A focused tool is easier to explain to a budget committee than a platform. But repeated across a dozen use cases, that path produces three structural problems.

Problem 1: Point solutions do not compound

A document extractor that has read ten thousand loan files does not make the fraud model better. The chatbot does not learn from what the credit team corrected. Each tool's value is capped at its own task, and the institution's second, fifth and tenth AI projects cost as much effort as the first, because nothing underneath them is shared. The knowledge an institution creates by using AI, which outputs staff accepted, which they overrode and what actually happened next, is scattered across vendor systems or not captured at all.

Problem 2: Point solutions fragment data and governance

Every tool needs data, so every tool gets its own copy, mapped its own way. Customer and account records drift apart across systems. Every vendor model is a separate entry in the model inventory, a separate validation exercise and a separate third-party relationship. The revised interagency model risk guidance asks banks to consider aggregate model risk, including reliance on common assumptions, data or methodologies across models.[2] That is hard to assess when each model's data, assumptions and logs sit with a different vendor in a different format. The interagency third-party guidance similarly expects due diligence on each third party, including its use of subcontractors.[3][5] Ten AI vendors mean ten sets of that work.

Problem 3: Point solutions rent intelligence

Most AI tools are built on a small number of large foundation models owned by a handful of providers. In its December 2024 report on AI in financial services, the U.S. Treasury summarized public comments warning that this concentration increases smaller institutions' dependency on a few large companies, and that an interruption at a single AI provider could disrupt many institutions at once.[6] When the intelligence lives in a vendor's model and the vendor's logs, the institution owns neither the capability nor the record of how it was used. If the vendor changes the model, the bank's results change with it.

> **Efficiency is a result, not a strategy**
> Operational improvement is a fine outcome of AI. It is a poor organizing principle. An institution that buys AI only to save hours will own a stack of tools, each one capped at the hours it saves. An institution that builds a foundation will save those hours too, and will also own something that improves.

3 Why a Learning Foundation

A learning foundation is the shared layer that sits between an institution's systems of record and its AI applications. It holds four things that point solutions keep separately: the institution's **mapped data** in one canonical schema; its **models**, numerical, behavioral and language; its **controls**, identity, scope, checks and approvals; and its **record**, a single audit trail of what each model was asked, what it returned and who acted on it. Applications are thin by comparison: screens and workflows that call the foundation for everything else.

© 2026 SmarterFI. Licensed under CC BY 4.0. · www.smarterfi.ai

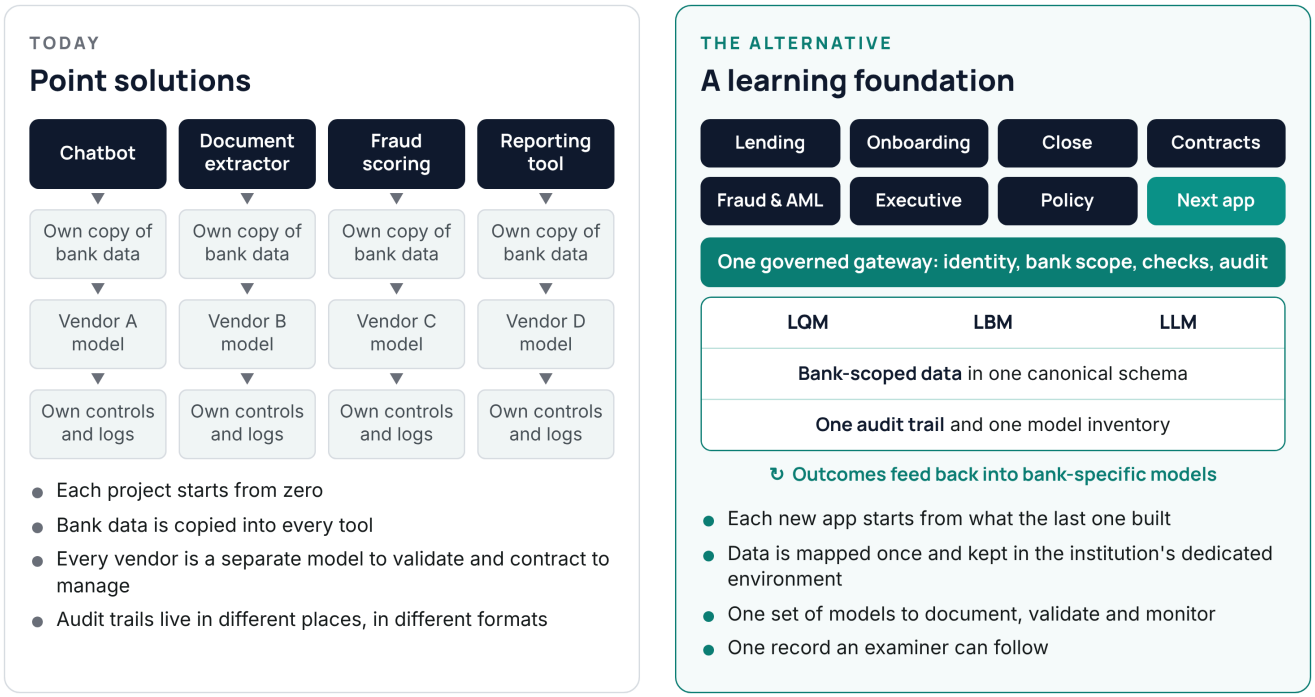


Figure 1. Point solutions versus a learning foundation. Each point solution brings its own data copy, model and logs; a foundation gives every application the same governed data, models and audit trail, and feeds outcomes back into the institution's own models.

Five properties that matter to a board

1. **It compounds.** The data mapping, controls and models built for the first application are reused by the second. Each new application is faster and cheaper to add, and the foundation's models improve as outcomes accumulate.
2. **It keeps one version of the data.** Data is mapped once, checked against data contracts and kept in the institution's dedicated environment. Applications read the same records instead of each holding a copy.
3. **It governs once.** Identity, scope checks, approvals and logging are enforced in one gateway for every application. Policy changes are made in one place.
4. **It explains by construction.** Because every application goes through the same gateway, every result carries the same trace: input references, model and calculator versions, retrieved passages and the person who acted on it.
5. **It fits model risk management.** One set of models, documented and validated once, monitored continuously, is easier to govern than a dozen opaque vendor models. Section 8 explains how this maps to current guidance.

The honest trade-off

A shared foundation concentrates dependency. If the foundation fails or a shared model is wrong, more than one application is affected. That is exactly the aggregate risk the model risk guidance asks banks to consider.[2] We think the answer is to design for it rather than to avoid foundations. Institution-specific models are trained and approved separately; every model version can be rolled back to the prior approved version; serving is stateless, so one request cannot affect another; and an institution's data and audit records stay in its own environment and can be exported if it leaves. A foundation should make dependency visible and controllable. A drawer full of point solutions only hides it.

© 2026 SmarterFI. Licensed under CC BY 4.0. · www.smarterfi.ai

4 How I Built SmarterFI

> **FROM THE AUTHOR**
>
> *I spent my career on the builder side of banking technology: as President and CTO of BM Technologies, as Global Head of Cloud at Finastra, as SVP of Cloud Engineering at First Republic, and in CIO and CISO roles before that. In every one of those roles, my first question for any vendor was simple: can it show its work? When I started SmarterFI, I wanted to build AI that community banks could answer "yes" to, every time, without an enterprise-sized team.*

SmarterFI did not start with a model. It started with the problem every community bank AI conversation runs into: the data. In the early days, before any live connections, we built a mock sandbox that mimicked the API styles of the major core providers so we could test our own applications against each one. That work taught us where AI projects actually get stuck, and it shaped how we built the live core connections SmarterFI runs today. A clean API to a core still returns whatever is in the core: email addresses that don't match between the loan system and the core, customers missing from the CRM, journal entries with no approver. None of it is dramatic, and every bit of it quietly breaks an AI application built on top.[7]

So the first thing we built was not an AI feature. It was **Data Mapping & Discovery**: a guided interview and a read-only scan that documents an institution's systems, how they connect and where the gaps are, before anyone spends money on AI. The Data Mapper reads metadata by default, not customer or account records, and it became the ingestion path for everything else.[7] Once the data is mapped into one canonical schema, every application can use it. That was the moment the foundation idea became concrete: build the hard, shared parts once, govern them properly and let every application stand on them.

From there we set the design principles that every part of SmarterFI follows:

1. **Own what differentiates.** SmarterFI writes its own ingestion and mapping code, canonical schema, calculators, model training, prompts, evaluations and orchestration. We do not resell someone else's model as our intelligence.
2. **Bank data is bank-scoped by default.** Each institution has its own environment, keys, data stores and index. Pooling data across institutions requires written permission from every institution involved.
3. **Numbers come from engines; words come from the language model.** Every figure a user sees comes from a tested calculator or a validated model. The language model explains and cites; it never computes.
4. **People decide.** Credit, risk, compliance and customer decisions stay with the institution's staff. Models advise, and the system records who acted on what.
5. **Every result is traceable** to its data, model, code and prompt versions, and to its approver.
6. **Private and minimal by default.** Data services sit behind private endpoints. Expensive services are added only when there is a demonstrated need.
7. **Say what is built and what is planned.** We separate what runs today from what we are still building, and we do not quote results we have not measured.

We also built applications, because a foundation nobody uses teaches nothing. Each application was designed from the start to be a client of the shared layer rather than a self-contained product. Section 6 shows how they connect.

© 2026 SmarterFI. Licensed under CC BY 4.0. · www.smarterfi.ai

5 The Architecture in Plain English

Figure 2 shows the SmarterFI foundation from top to bottom. Applications sit at the top. Every request passes through one governed gateway. Each institution's data lives in its own dedicated environment. The model layers sit in a separate SmarterFI model plane. A platform layer underneath provides identity, networking, keys and monitoring.

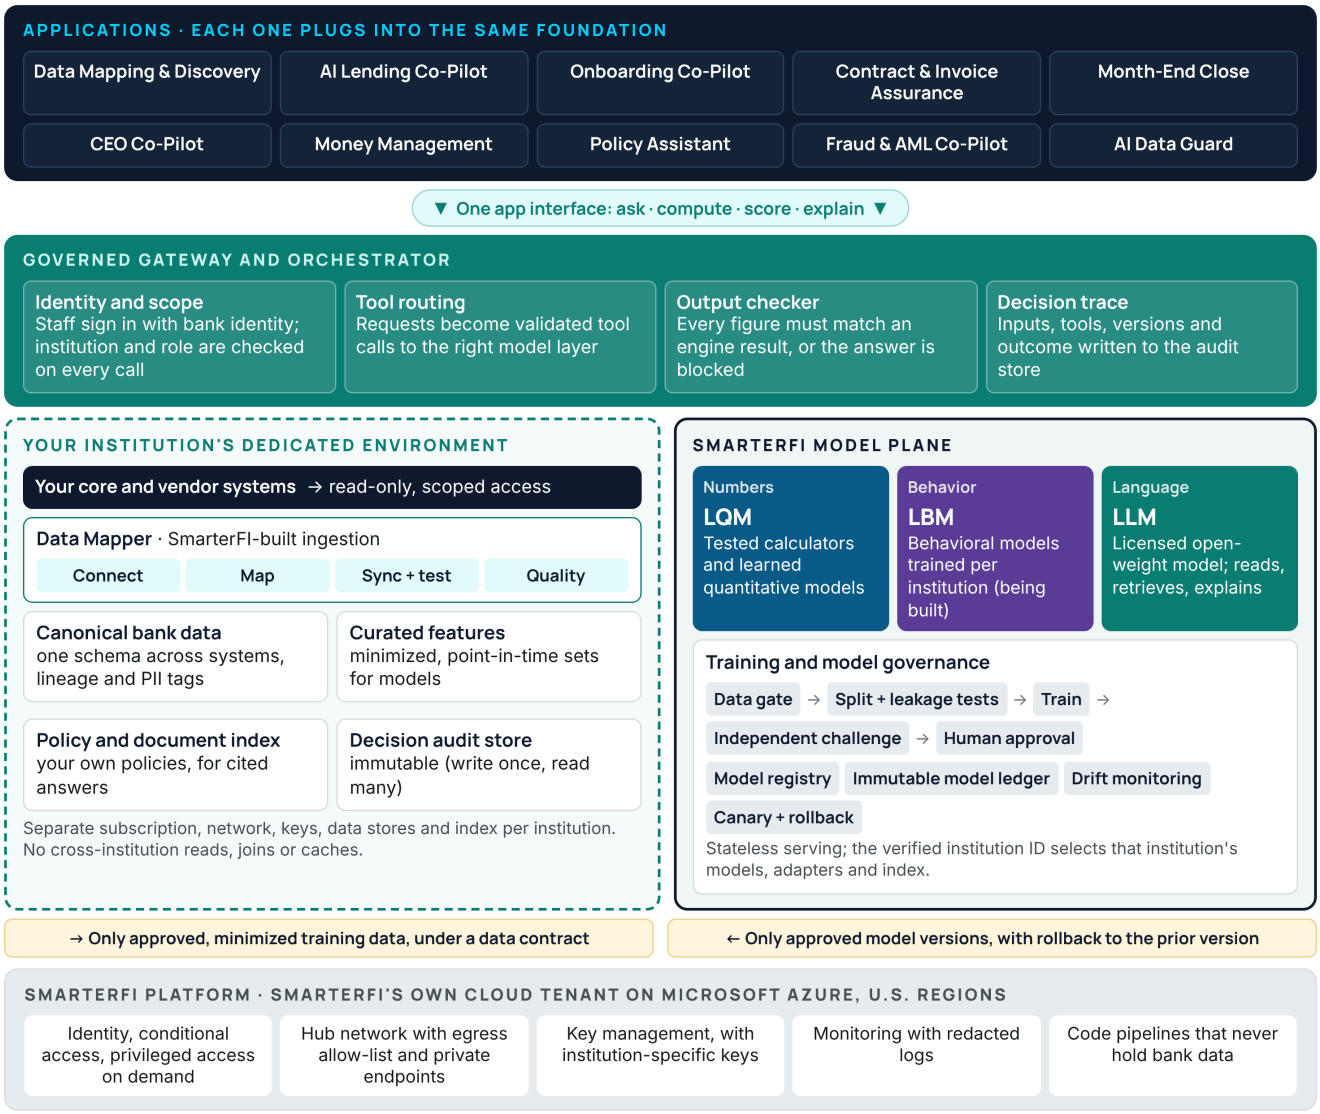


Figure 2. The SmarterFI learning foundation (target architecture). One gateway serves every application; each institution's data, index and audit store sit in its own dedicated environment; models are trained, approved and served in the SmarterFI model plane.

Walking through a request

1. **Sign in and scope.** Staff sign in with the institution's own identity. The gateway sets institution and role from verified claims, never from anything the caller sends.
2. **Plan.** The language model turns the request into structured tool calls from a registered catalog (compute this ratio, score this account, retrieve this policy). Each call is checked against its schema, the user's permissions and the institution's scope; invalid calls are rejected, not quietly repaired.

© 2026 SmarterFI. Licensed under CC BY 4.0. · www.smarterfi.ai

3. **Compute and retrieve.** The LQM computes figures with decimal arithmetic and versioned rules; the LBM returns scores with reason codes; retrieval returns passages from the institution's own policy index.
4. **Explain and check.** The language model writes only from those results and passages, citing each one. The output checker matches every number, rate and date to a result or passage; if one does not match, the answer is blocked or replaced with a plain "I can't answer that from the available data."
5. **Record.** The full trace (input references, tool calls, results, model and prompt versions, checker outcome) is written to the institution's write-once audit store.

Separation between institutions

SmarterFI runs in its own cloud tenant, and each institution gets its own subscription with its own network, keys, databases, document index and audit store. A subscription is strong separation but not, on its own, complete isolation, so separation is enforced in layers: per-institution identity groups and role checks, private networking with an egress allow-list, institution-specific encryption keys, separate data stores, and logs redacted before central monitoring. In the shared model plane, every call carries the verified institution identifier set by the gateway, serving keeps no shared caches or conversation memory, and that identifier selects the institution's own models and adapters.

> **Where SmarterFI is today, and starting small**
> SmarterFI is connected to 13 different banking cores today. Its applications work from those live core connections, and every new core connection is reused by every application on the foundation. The LQM is running in production use in SmarterFI's private cloud on Microsoft Azure, in U.S. regions. The other owned model layers in Figure 2 are being built in stages there, and the LBM's definition is still evolving. New capacity starts small: CPU-based training by default, with specialized compute added only when an evaluation shows the need. This paper does not report accuracy, validation or customer results.

6 Three Model Layers and How Apps Connect

SmarterFI uses three model layers because a bank's AI has to do three different kinds of work, and each needs different controls. LQM and LBM are SmarterFI product terms, not regulatory or industry categories; LLM is the general industry term (Table 1).

	LQM: numbers	LBM: behavior	LLM: language
What it does	Computes figures and quantitative estimates	Anticipates customer, account and portfolio behavior for staff to act on	Understands requests, retrieves policy text, calls tools, explains verified results
What it is	(a) Deterministic calculators; (b) statistical and machine-learning models	Institution-specific behavioral models: interpretable baselines first, sequence models only if they prove their value	A licensed open-weight model, self-hosted, with retrieval and optional fine-tuned adapters
Examples	Debt service coverage, amortization, interest-rate sensitivity; loss-rate, prepayment and deposit-decay estimates	Deposit run-off and attrition risk, anomalies routed for human review, product propensity	Plain-English answers that cite tool results and the institution's own policies
Trained on	Calculators: none. Models: the institution's portfolio history, under contract	The institution's own behavioral data; pooled only with written permission	Base model per its license; adapters only on licensed, consented, de-identified text

© 2026 SmarterFI. Licensed under CC BY 4.0. · www.smarterfi.ai

	LQM: numbers	LBM: behavior	LLM: language
Key control	Decimal arithmetic, declared rounding, golden tests reconciled to core outputs	Out-of-time testing, calibration, reason codes, fair-lending review where relevant	Output checker: no unverified figure reaches a user
Status	Running in production use	Being built; definition still evolving	Self-hosted model being built in stages; managed model access is an option, not owned weights

Table 1. The three model layers in SmarterFI's design, with the current status of each.

The LQM: every figure from a tested engine

The LQM is running in production use today. Calculations such as debt service coverage, amortization and covenant tests are versioned functions with no learned parameters. They use decimal or fixed-point arithmetic rather than binary floating point, which cannot represent most decimal fractions exactly;[8] rounding rules are declared per calculation and reconciled to the institution's core outputs. Estimates such as probability of default, loss rates and deposit behavior are statistical models, and they are treated as models: documented, validated out-of-sample and out-of-time, and monitored, with outputs stated as estimates with uncertainty.[9]

The LBM: the behavioral layer we are building

Community banks know their customers, but much of that knowledge lives in people's heads and in transaction histories nobody has time to read. The LBM is the behavioral layer SmarterFI is building to make that knowledge usable: which deposit relationships are at risk of leaving, which accounts show patterns worth a second look, which small-business customers may need a line of credit before they ask. It is designed to be trained per institution, from that institution's own history, with features defined so that no prediction uses information from after the prediction date, and every score shown to staff will carry reason codes. It supports staff judgment; it does not make credit, fraud or customer decisions.

The LLM: a language interface, not an oracle

Our language layer is designed around a licensed open-weight model that SmarterFI hosts itself, so prompts and documents stay inside SmarterFI's environment. SmarterFI does not own the base model's weights; it owns the adapters, prompts, evaluations and orchestration around them. Managed models such as Azure OpenAI remain an option, and those weights are not ours either. Adaptation is staged: retrieval over the institution's own documents first, light fine-tuning only where evaluation shows a gap, and no from-scratch pretraining.

How applications plug in

Every SmarterFI application uses the same client interface to the foundation, with four verbs: **ask**, **compute**, **score** and **explain**. Feature flags switch each application, per institution and per feature, onto the foundation, and each runs in shadow mode on approved data before users see its results. Because SmarterFI is connected to 13 different banking cores today, an application written once can serve institutions on any of those cores. Table 2 shows how the applications in SmarterFI's catalog draw on the shared layers.[7]

Application	What it does	Draws on
Data Mapping & Discovery	Documents systems, connections and data gaps; becomes the ingestion path for the foundation	Canonical data, data contracts, quality checks

© 2026 SmarterFI. Licensed under CC BY 4.0. · www.smarterfi.ai

Application	What it does	Draws on
AI Lending Co-Pilot	Checks loan packages, spreads tax returns and drafts credit memos with a citation for every number	LLM to read and draft; LQM for spreads and ratios; audit trail
Onboarding Co-Pilot	AI-assisted account opening with the institution's own BSA/AML rules and reviewer notes	Deterministic rules; LBM anomaly signals for review; LLM-drafted notes
Contract & Invoice Assurance	Turns vendor contract terms into data and checks invoices against them	LLM extraction with citations; LQM invoice checks
Month-End Close	Structured close workflows, reconciliations and exception routing	LQM reconciliations; LBM exception patterns; audit trail
CEO Co-Pilot	Daily briefs, scenario models and board-ready reporting from the bank's own figures	LQM scenarios; LBM deposit behavior; LLM briefings
Money Management	Customer-facing guidance under the institution's brand	LBM spending patterns; LQM projections; LLM plain-language guidance
Policy Assistant	Answers staff questions only from the institution's own policies, with citations	LLM retrieval over the institution's policy index
Fraud & AML Co-Pilot	Supports fraud and AML review across onboarding, logins and transactions	LBM anomaly scores with reason codes; LLM case narratives; people decide
AI Data Guard	Detects and controls sensitive data before it reaches public AI tools, with policy and audit reporting	Shared policy, identity and audit controls at the edge

Table 2. How SmarterFI applications draw on the shared layers. Model layers marked as being built in Table 1 are added to each application as they are released.

7 How the Foundation Learns

"Gets smarter with use" is a claim that deserves precision, because in banking an AI system that changes itself without control is a liability, not a feature. In the SmarterFI foundation, learning is a governed release process, not a side effect.

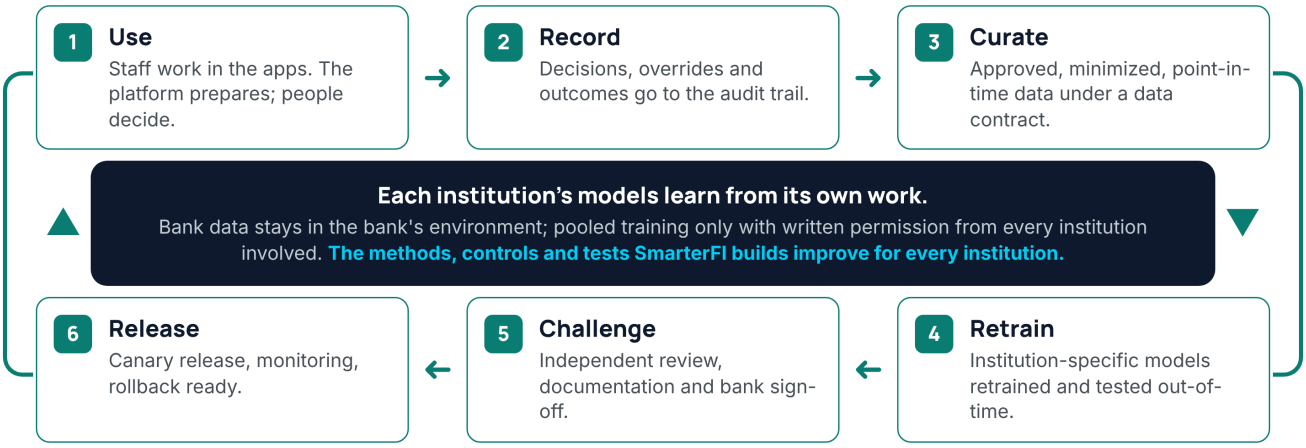


Figure 3. The governed learning loop. Outcomes recorded in daily use become curated training data; new model versions are released only after testing, independent challenge and human approval.

© 2026 SmarterFI. Licensed under CC BY 4.0. · www.smarterfi.ai

The loop runs in six steps. Staff use the applications, and the platform records what they accepted, what they overrode and, later, what actually happened: whether a flagged anomaly was confirmed, whether a deposit relationship left, whether a memo figure was corrected. Approved, minimized data, defined by a data contract and fixed at the point in time of each prediction, flows to the training pipeline. Institution-specific models are retrained and tested on later periods than they were trained on. Developers do not validate their own models: an independent reviewer challenges the results, and the institution signs off on models trained on its data. Approved versions are released gradually, monitored for drift and kept ready to roll back.

What compounds, and for whom

Two different things improve. **Each institution's models** improve from that institution's own outcomes, and that learning stays in its environment. **SmarterFI's methods** improve across every institution without any data moving: better calculators and golden tests, mappings for more core and vendor systems, sharper evaluation sets, stronger output checks. Pooled, cross-institution training is possible only with written permission from every institution involved and legal review; it is never the default.

Every change leaves a record

A model registry catalogs which versions exist and who may promote them,[10] but it is not an audit trail. SmarterFI's design keeps two write-once records in immutable storage:[11] a **model ledger** recording the dataset snapshot, data contract, code commit, environment, evaluation results, validation findings and approver for every model version; and a **decision audit store** in each institution's environment recording each request, the tools called, the versions used and the outcome. Calculators reproduce exactly from the same version and inputs; statistical models reproduce within stated tolerances, because training on modern hardware is not bit-for-bit reproducible, and we say so.

8 Governance, Auditability and Regulator Readiness

A foundation should make an institution's conversation with its examiners easier. Here is how the SmarterFI design maps to current guidance.

Model risk management under SR 26-2

On April 17, 2026, the Federal Reserve, OCC and FDIC issued revised guidance on model risk management. The Federal Reserve's SR 26-2 and the OCC's Bulletin 2026-13 supersede SR 11-7 and its OCC counterpart, Bulletin 2011-12.[1][12] Three points matter here. First, the guidance is expected to be most relevant to banking organizations with more than $30 billion in assets; smaller institutions' models are typically governed by internal practices suited to their size and risk, unless their model risk is significant.[2] Second, it defines a model as a complex quantitative method that turns input data into quantitative estimates, and it excludes simple arithmetic and deterministic rule-based processes.[2] Third, it places generative and agentic AI outside its scope, while saying the institution's own risk management should guide controls for anything it does not cover.[2]

The three-layer design follows those lines. The LQM's statistical models and the LBM are models in the guidance's sense, so they are built to its principles: clear purpose, conceptual soundness, out-of-time testing, outcomes analysis, ongoing monitoring and effective challenge by objective, qualified reviewers.[2] The LQM's calculators fall outside the definition, but are still tested against reference cases. The LLM is governed under the institution's risk framework, supported by the NIST AI Risk Management Framework and its Generative AI Profile.[13][14] For credit unions, GAO has recommended that NCUA expand its model risk guidance to cover a broader range of models, so a vendor that documents to these principles is easier to govern either way.[4]

© 2026 SmarterFI. Licensed under CC BY 4.0. · www.smarterfi.ai

Vendor models and third parties

The guidance asks institutions to understand vendor models, including their design, development data and performance, and to monitor them over time.[2] The guidance itself acknowledges that proprietary vendor components can make this hard.[2] A foundation can be designed for it: SmarterFI's design includes a documentation pack for each model so the institution, or a validator it engages, can perform its own validation. Under the interagency third-party guidance, which replaced OCC Bulletin 2013-29 in 2023, the institution still performs its own due diligence on SmarterFI, including SmarterFI's reliance on subcontractors.[3][5] A foundation reduces the number of AI vendors to assess, but it does not remove the chain: the cloud provider, model licensors and core providers remain, and SmarterFI documents them. A cloud provider's audit reports evidence that provider's controls, not SmarterFI's or the institution's.[15]

What an examiner or board may ask	How the foundation answers
Where did this number come from?	Every figure traces to a calculator or model result with its version, inputs and data references, and the output checker blocks any figure that does not.
Which models do you use, and who validated them?	One inventory for every application, with a model ledger recording data, code, evaluation, independent review and approver for each version.
Who made the decision?	A named person at the institution. Models advise; the audit store records who acted on which output.
Is our data used for anyone else?	No, by default. Institution data stays in its dedicated environment. Pooling requires written permission from every institution involved.
What happens if a model goes wrong?	Monitoring thresholds set at validation, canary releases and rollback to the prior approved version without retraining.
What if we leave?	Data and audit records are the institution's; retention, deletion and return, including of institution-specific models, follow the agreement.

9 What This Means for Your Institution

Building on a foundation does not mean a multi-year program before anyone sees value. It means sequencing the work so every step leaves something reusable behind.

Step	What you do	What you keep
1. Map your data	Run a read-only discovery of your core, vendor systems and how they connect. Score readiness and find the gaps to close first.	A systems map, a data readiness baseline and a plan your board can read
2. Set the rules	Agree on AI use policy, model inventory, data rights and who approves what. Govern staff use of public AI tools rather than banning it.	One governance model every future application inherits
3. Stand up the foundation	Provision your dedicated environment with identity, keys, canonical data and audit store. Connect one high-value application in shadow mode.	The shared layer, proven on a real workflow without customer impact
4. Add applications	Each new application plugs into the same data, gateway and audit trail, behind feature flags, one at a time.	Lower effort per application; one record across all of them
5. Close the loop	Capture outcomes, retrain institution-specific models, review them independently and release them under approval.	Models that improve with your own experience

© 2026 SmarterFI. Licensed under CC BY 4.0. · www.smarterfi.ai

Questions for your next AI decision

- [ ] Will this tool use our data where it already lives, or make another copy?
- [ ] Can every figure it shows be traced to a tested calculation or a validated model?
- [ ] Does it get better as our staff use it, and do we keep that improvement if we leave?
- [ ] Can the vendor give our validator the documentation the model risk guidance describes?
- [ ] Who are the vendor's subcontractors, including model providers, and what happens if one changes?
- [ ] Does it share identity, controls and audit with our other AI applications, or bring its own?

10 How SmarterFI Is Different

	Point solutions	A learning foundation (SmarterFI)
Unit of investment	One tool per task	One foundation; applications as thin clients
Second, fifth, tenth project	Each starts from zero	Each reuses data, controls, models and audit
Core connectivity	Each vendor integrates with the core separately, if at all	Connected to 13 different banking cores today; each connection serves every application
Data	Copied into each vendor's store and mapped separately	Mapped once into one canonical schema in the institution's dedicated environment
Learning	Capped at each tool; usage knowledge stays with vendors	Outcomes improve the institution's own models under governed release
Numbers	Often generated by the same model that writes the text	Computed by tested engines; the language model explains and cites
Governance	Separate controls, logs and formats per tool	One gateway, one model inventory, one write-once audit trail
Model risk	Many opaque vendor models to validate	Documented models built to SR 26-2 principles, with a documentation pack for validators
Third parties	One vendor relationship per tool, each with its own fourth parties	Fewer AI vendors to assess; subcontractors documented
Ownership	Intelligence rented; results change when the vendor changes the model	Institution owns its data and records; SmarterFI owns the methods; model versions change only on approval
Exit	Knowledge and logs stay behind	Data and audit records returned under the agreement

None of this makes point solutions useless. A focused tool can be the right answer for a narrow problem. The difference is what sits underneath: when the next application arrives, does the institution start again, or start from what it already built?

11 Conclusion

Financial services is moving from asking whether to use AI to asking what kind of AI capability to own. Community banks and credit unions have too often been offered a growing pile of tools, each solving one

© 2026 SmarterFI. Licensed under CC BY 4.0. · www.smarterfi.ai

problem, none of them learning, and every one adding data copies, models to validate and contracts to manage.

We think the industry should hold AI to the standard banks already hold everything else to: build on a sound foundation, govern it properly and keep a record anyone can follow. A governed learning foundation, with numbers from tested engines, behavior from models trained on each institution's own history, words from a language model that cites its sources, and people making the decisions, is how AI becomes an asset that compounds rather than an expense that repeats.

Let's talk about your foundation

I would welcome a 30-minute working session with your leadership team: where your data stands, what a foundation would look like for your institution, and what your risk committee and examiners would want to see.

BOOK A SESSION WITH JAMIE
calendly.com/jamie-donahue-smarterfi/30min

LEARN MORE
www.smarterfi.ai

Jamie Donahue, Founder, President & CAIO, SmarterFI

R References

Web sources accessed November 9, 2026. Public sources support general statements about the industry and regulation and do not refer to SmarterFI.

1. Board of Governors of the Federal Reserve System, SR 26-2, *Revised Guidance on Model Risk Management* (April 17, 2026); supersedes SR 11-7 and SR 21-8. https://www.federalreserve.gov/supervisionreg/srletters/SR2602.htm
2. Federal Reserve, OCC and FDIC, *Supervisory Guidance on Model Risk Management* (attachment to SR 26-2): scope, definition of "model", aggregate risk, effective challenge, vendor products, and the note on generative and agentic AI. https://www.federalreserve.gov/frrs/guidance/supervisory-guidance-on-model-risk-management.htm
3. Office of the Comptroller of the Currency, Bulletin 2023-17, *Third-Party Relationships: Interagency Guidance on Risk Management* (June 6, 2023); rescinds OCC Bulletins 2013-29 and 2020-10. https://www.occ.gov/news-issuances/bulletins/2023/bulletin-2023-17.html
4. U.S. Government Accountability Office, GAO-25-107197, *Artificial Intelligence: Use and Oversight in Financial Services* (May 19, 2025). https://www.gao.gov/products/gao-25-107197
5. Federal Reserve, OCC and FDIC, *Interagency Guidance on Third-Party Relationships: Risk Management* (2023), including due diligence on a third party's use of subcontractors. https://www.federalreserve.gov/frrs/guidance/interagency-guidance-on-third-party-relationships.htm
6. U.S. Department of the Treasury, *Artificial Intelligence in Financial Services* (December 2024) (executive summary; concentration-related risks). https://home.treasury.gov/system/files/216/Artificial-Intelligence-in-Financial-Services.pdf
7. SmarterFI, *Platform & Solutions* and *Technology & Architecture* pages, and the SmarterFI App Catalog. https://smarterfi.ai/platform
8. Python Software Foundation, *Floating-Point Arithmetic: Issues and Limitations*. https://docs.python.org/3/tutorial/floatingpoint.html
9. Donahue, J. (2026). *Why LLMs Are Not Enough for Community Banks: The case for a private large quantitative model*. SmarterFI. Version 1.1. https://doi.org/10.5281/zenodo.23884876
10. Microsoft Learn, *Manage models registry in Azure Machine Learning with MLflow*. https://learn.microsoft.com/en-us/azure/machine-learning/how-to-manage-models-mlflow
11. Microsoft Learn, *Immutable storage for Azure Blob Storage* (write once, read many). https://learn.microsoft.com/en-us/azure/storage/blobs/immutable-storage-overview
12. Office of the Comptroller of the Currency, Bulletin 2026-13, *Model Risk Management: Revised Guidance* (April 17, 2026); rescinds OCC Bulletin 2011-12. https://www.occ.gov/news-issuances/bulletins/2026/bulletin-2026-13.html
13. National Institute of Standards and Technology, NIST AI 100-1, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (January 2023). https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf
14. National Institute of Standards and Technology, NIST AI 600-1, *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile* (July 2024). https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf
15. Microsoft Learn, *Shared responsibility in the cloud*. https://learn.microsoft.com/en-us/azure/security/fundamentals/shared-responsibility

© 2026 SmarterFI. Licensed under CC BY 4.0. · www.smarterfi.ai

SmarterFi.ai
AI for Community Banking

ABOUT SMARTERFI

SmarterFI (SmarterFi.ai, “AI for Community Banking”) builds a governed AI foundation and a catalog of applications for community banks and credit unions, including Data Mapping & Discovery, AI Lending Co-Pilot, Onboarding Co-Pilot, Contract & Invoice Assurance, Month-End Close, CEO Co-Pilot, Money Management, Policy Assistant, Fraud & AML Co-Pilot and AI Data Guard. SmarterFI is connected to 13 different banking cores today, and each institution runs in a dedicated environment in SmarterFI's private cloud, with its own data, keys and audit trail. SmarterFI also convenes the AI Consortium for Community Banks, a forum for executives at institutions adopting AI together.

ABOUT THE AUTHOR

Jamie Donahue is Founder, President & Chief AI Officer (CAIO) of SmarterFI, where he leads the design and build of SmarterFI's AI foundation for community banks and credit unions. His career spans banking, fintech and financial technology providers. He was previously President & Chief Digital Officer of First Carolina Bank's BM Technologies Division and President & Chief Technology Officer of BM Technologies. Earlier, he served as Global Head of Cloud at Finastra and SVP of Cloud Engineering at First Republic, and held CIO, CISO and SVP technology roles. A US Navy veteran, he is based in Palm Beach, Florida. LinkedIn: linkedin.com/in/jamesmdonahue

WEB
www.smarterfi.ai

EMAIL
jamie.donahue@smarterfi.ai

CITE AS

Donahue, J. (2026). AI That Compounds: Why financial institutions need a governed learning foundation, not more point solutions. SmarterFI. Version 1.0. CC BY 4.0

© 2026 SmarterFI. Licensed under Creative Commons Attribution 4.0 International (CC BY 4.0): you may share and adapt this paper with attribution to SmarterFI and the author (https://creativecommons.org/licenses/by/4.0/). Describes SmarterFI's architecture and design as of November 2026: SmarterFI is connected to 13 different banking cores today and the LQM is running in production use in SmarterFI's private cloud; the LBM and other model layers described as being built are not yet in production. It does not report accuracy, validation or customer results. Provided as is, without warranties of any kind, as set out in the CC BY 4.0 license. Not legal or regulatory advice; does not modify any agreement or imply endorsement by any organization cited. Each institution remains responsible for its own model risk, third-party risk, compliance and fair-lending programs. The license does not grant rights in the SmarterFI name or logo. Microsoft and Azure are trademarks of the Microsoft group of companies; other names are the property of their owners.