SmarterFi.ai
AI for Community Banking

TECHNICAL ARCHITECTURE BRIEF · OCTOBER 2026

Eliminating Frontier Model Risk

SmarterFI Technical Architecture Brief

> How the SmarterFI AIOS keeps generative AI inside a governed, isolated and examinable control path for community banks and credit unions

By James (Jamie) Donahue, Founder, President & Chief AI Officer, SmarterFI.ai

Cite as: Donahue, J. (2026). Eliminating Frontier Model Risk: SmarterFI Technical Architecture Brief. SmarterFI.ai. Version 1.1. https://doi.org/10.5281/zenodo.23884799. CC BY 4.0

Identifiers masked
before any model call

No number
from a language model

A person decides
every consequential action

Append-only
examiner-ready audit log

For bank and credit union technology buyers, risk committees, compliance and internal audit

1 Executive Summary

Frontier models, the large general-purpose language models behind today's generative AI, give community banks and credit unions capabilities that were out of reach a few years ago. They also raise questions that risk committees are right to ask: Can the model invent a number? Can it see a customer's Social Security number? Can a document hide instructions that the model follows? Will the provider keep or train on our data? Can we show an examiner exactly what happened?

SmarterFI's answer is architectural. No design removes model risk entirely, so the SmarterFI AIOS takes the frontier model **out of the decision path** and confines it to a narrow, governed role: it reads, extracts and drafts, and it must cite its source. Everything that carries institutional risk happens around it, in controls the bank can inspect:

- **Identifiers never reach a model unprotected.** AI Data Guard masks or tokenizes sensitive data before any model call.
- **No number comes from a language model.** Every calculation runs in versioned, tested, deterministic code. A value whose citation does not check out is discarded, and “not found” is a valid answer.
- **People decide.** Agents prepare the work; a person makes every decision that affects a customer, financial reporting or a regulatory filing, and that decision is recorded.
- **Your data stays yours.** Each institution has its own isolated environment, with U.S.-only storage and inference, enterprise model endpoints, and a contractual commitment that bank data is not used to train models.
- **Everything is logged.** An append-only audit log captures prompts, outputs, model versions and user actions, and the bank can export it for examiners.

This brief documents each control, maps it to the revised interagency model risk guidance (SR 26-2 / OCC Bulletin 2026-13, which superseded SR 11-7), the FFIEC IT Examination Handbook, GLBA, interagency third-party risk management guidance, SOC 2 and business continuity expectations, and crosswalks it to three 2026 community-bank AI frameworks. It closes with a checklist and FAQ for risk committees.

> **How to read the status labels**
> Every control carries a status so that reviewers can weigh the evidence: LIVE deployed in SmarterFI's production environment on Microsoft Azure or published as a platform commitment on smarterfi.ai; BUILT implemented in SmarterFI's application code and demonstrated in its public prototypes, which run on synthetic data only; SPECIFIED required by SmarterFI's product requirements or defined in the customer agreement. SmarterFI's controls are designed to SOC 2, and its landing zone runs on Microsoft Azure, whose infrastructure carries Microsoft's SOC 2 Type II attestation (Section 8).

2 Frontier Model Risk, Defined

Risk committees evaluating generative AI typically raise the same nine exposures. The table summarizes how the SmarterFI architecture answers each one; the sections that follow give the control-level detail.

Exposure	What can go wrong	SmarterFI architectural response	See
Fabrication	A model invents a figure, a clause or a citation	Model extracts and cites only; citations are verified or the value is discarded; all arithmetic in deterministic code; “not found” preferred to guessing	4
NPI leakage	Customer data reaches a model provider or a public AI tool	AI Data Guard masks identifiers before any model call; Data Guard application monitors staff use of public AI tools; production egress allow-list	4, 5

Exposure	What can go wrong	SmarterFI architectural response	See
Prompt injection	Instructions hidden in a document or email are executed	Document text passed as data, never as instructions; Prompt Shields; assistants cannot take actions	4
Provider reuse	Prompts retained or used for training	Azure-sold models only; no-training terms from Microsoft and in SmarterFI's agreement	5
Cross-institution exposure	One bank's data visible to another	Dedicated environment and keys per institution; no shared models trained on other banks' data	5
Silent change and drift	Outputs shift when a model is updated or degrades	Pinned model versions, advance notice of material changes, ongoing monitoring	7
Automated decisions	AI approves, declines or sends without review	Human approval gates; no automated credit or account-opening denials by default	4, 7
Unexaminable activity	No reliable record of what the AI did	Append-only audit log of prompts, outputs, versions and actions; examiner export	6
Vendor dependence	Lock-in or loss of data on exit	Bank owns data and outputs; export with data dictionary and audit logs; certified deletion	8

> **Design principle**
> The bank governs the AI; the AI does not govern the bank. Governance sits between every agent and every action, and the bank's core systems, data and decisions stay on the bank's side of the boundary. LIVE

3 Architecture Overview

The SmarterFI AIOS is a multi-agent, policy-governed platform that runs in SmarterFI's private cloud, with each institution in an isolated environment, U.S. regions only LIVE. It enhances the core and vendor systems a bank already runs rather than replacing them. Figure 1 shows the production architecture for each institution LIVE.

Architecture principle	How it is enforced	Status
Isolation by subscription	Each institution's data lives in its own Azure subscription: the unit of access, policy, cost and exit. Spokes peer only to a hub, never to each other.	LIVE
Shared logic, not shared data	Shared application logic does not store, cache or retain bank data outside the bank's environment except transiently to process a request.	SPECIFIED
Private by default	No public endpoints on data services; inbound only through Front Door Premium with WAF; outbound only through Azure Firewall with an allow-list.	LIVE
U.S. only, enforced by policy	Azure Policy allows only East US 2 (primary) and Central US (recovery); model deployments limited to Data Zone Standard (US); Global deployment types denied.	LIVE
Least privilege, just in time	No standing Owner rights; Privileged Identity Management with approval; phishing-resistant MFA for administrators; workload identities instead of secrets.	LIVE
Everything as code	Infrastructure deployed by pipeline with pull-request review, required checks and a production approval gate; portal changes in production are break-glass only.	LIVE

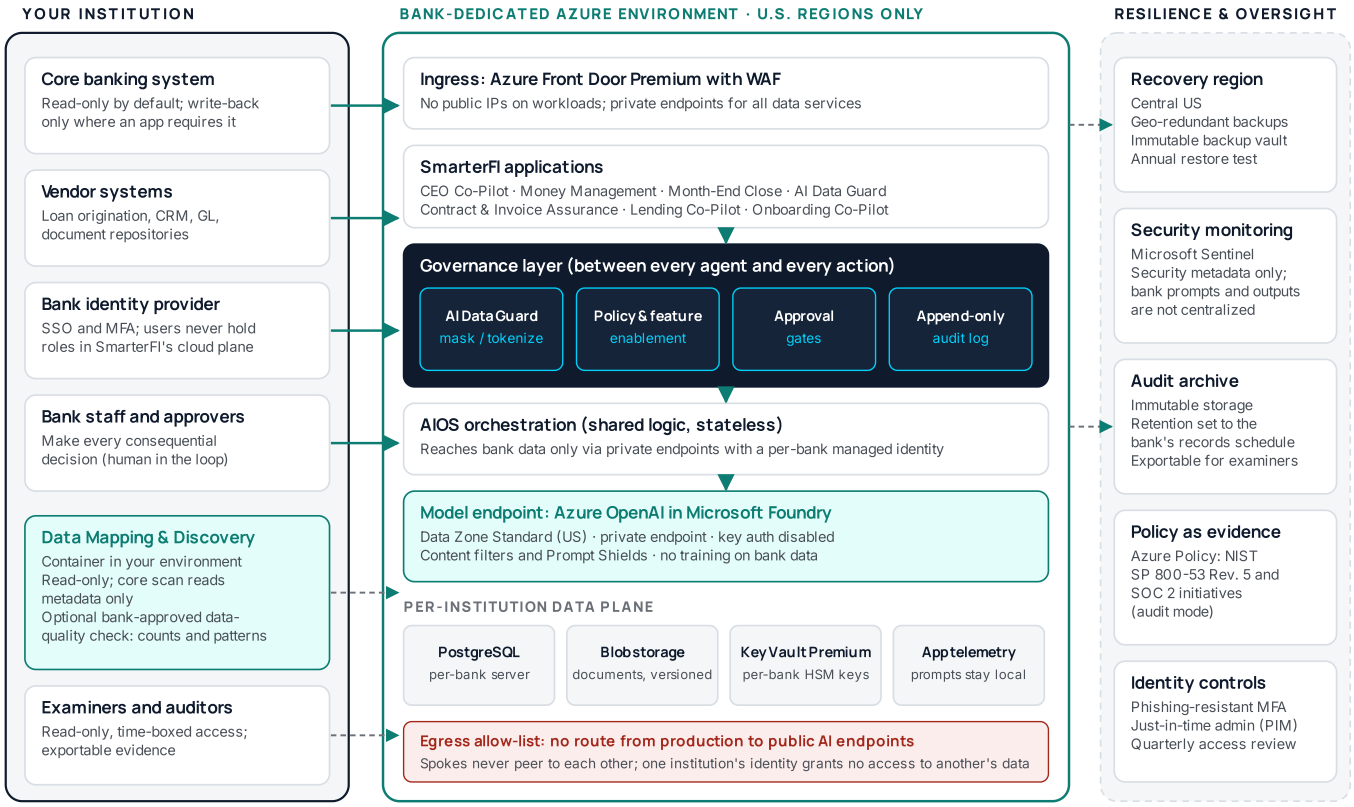


Figure 1. Production architecture. Each institution receives its own Azure environment in U.S. regions; shared AIOS logic is stateless and reaches bank data only through private endpoints; the governance layer sits between every agent and every action.

Hosting. The platform runs only in SmarterFI's private cloud, in an environment dedicated to the institution. SmarterFI does not offer an in-house version; only Data Mapping & Discovery can run in a bank's environment.

4 Guardrails: The Governed Path of Every Model Call

Every AI interaction follows the same six-stage path. Two stages, AI Data Guard before the model and output controls after it, exist specifically to neutralize frontier-model behavior: one prevents the model from seeing what it should not, the other prevents the institution from relying on what the model should not have produced.

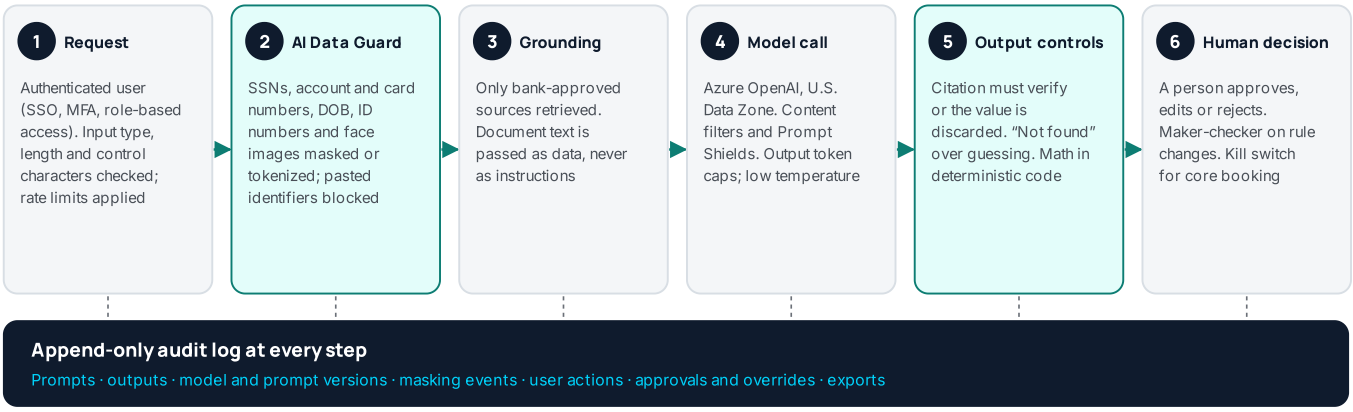


Figure 2. The path a request takes. The model is one stage of six; it is surrounded by input controls, identifier masking, source grounding, output verification and a human decision, all written to the audit log.

Control area	Control detail	Status
Input filtering	Server-side validation rejects malformed requests, strips control characters and caps question length (500–600 characters in current applications) and response tokens. Per-user and global rate limits. Errors return generic messages, so provider or system details never reach the browser.	BUILT
PII redaction	AI Data Guard masks, tokenizes or excludes SSNs and TINs, full account and card numbers, dates of birth, ID numbers and face images before any model call; the reviewer sees real values, the model does not. Images are cropped to the fields needed. Identifiers pasted into staff chat are blocked and the block is logged. A model gateway applies PII detection and redaction before any request leaves the bank's environment and records that redaction occurred.	SPECIFIED BUILT
Prompt-injection defenses	Retrieved document and record text is passed to the model as data; instructions inside documents are not executed; tool use is limited to retrieval. Assistants cannot approve, dispute, send, change terms or give notice. Azure content filters with Prompt Shields for jailbreak and indirect prompt injection are part of the deployed Azure model stack. Evaluation sets test prompt-injection resistance and refusal of out-of-scope requests.	LIVE SPECIFIED
Output filtering	Answers are restricted to supplied, numbered sources and must carry citation markers; citations are validated against the retrieved set; when sources do not answer, the system returns "not found" rather than guessing. Extraction is schema-constrained with page and clause citations and confidence scores. Output is rendered as sanitized text with internal links only.	BUILT SPECIFIED
Deterministic computation	Dollar figures, ratios and variances are computed in versioned, tested code; totals are precomputed server-side before any model sees them; drafted letters receive numbers injected from data.	LIVE BUILT
Human in the loop	Each consequential decision point requires an affirmative action by an authorized bank user and is logged. Examples: contract terms become active only after verification; invoice findings need a disposition, with a reason for dismissals; dispute letters require approval above the bank's threshold and a person sends them; onboarding reviewers, not the AI, approve or decline; BSA rule changes require a second approver (maker-checker); a kill switch pauses core booking or any vendor check.	BUILT SPECIFIED
Feature enablement	Each AI capability is separately enabled and off by default; nothing runs until the bank approves it. Each carries a risk tier (informational, advisory, decision-supporting) and an approval record.	BUILT
Application security	Strict Content Security Policy, HSTS, frame and MIME protections; signed HttpOnly, Secure, SameSite=Strict sessions with timing-safe verification; same-origin checks; failed-login lockout; no third-party scripts or fonts; no source maps. Verified in an October 2026 hardening review of SmarterFI applications.	BUILT

5 Data Isolation, Residency and Data Minimization

SmarterFI does not describe its environment as physically air-gapped. Instead it applies a **logical air gap**: each institution's data plane has no public endpoints, no network path to other institutions, and no route from production to public AI services, and the Data Mapper is built to run inside the bank's network with limited or no outbound internet access.

Control	Detail	Status
Tenant isolation	Bank data hosted in an Azure environment dedicated to the institution and isolated from other customers; personal information never commingled with or accessible to another customer, consortium member or reseller. Each institution has its own PostgreSQL server, storage, Key Vault with institution-specific keys and model resource, all behind private endpoints.	LIVE SPECIFIED
U.S. hosting	Storage and inference in the United States only (East US 2 primary, Central US recovery); administrative sign-in from outside the United States is blocked by Conditional Access policy.	LIVE SPECIFIED
No training on customer data	Bank data is not used to train, fine-tune or improve any model used for another customer. Foundation model services are configured so the provider does not train on bank data. Microsoft states that prompts and completions for Azure-sold models are not available to other customers or to OpenAI and are not used to train foundation models without permission.	LIVE SPECIFIED
Provider retention	Enterprise model endpoints with zero data retention.[1] Application telemetry, including prompts and outputs, stays in the bank's dedicated environment; only security metadata is centralized.	LIVE
Model provider allow-list	Bank data goes only to Azure-sold models through the bank's private endpoint. Other AI tools are restricted to development with synthetic or public data; production has no network route to public AI APIs.	LIVE
Encryption and keys	TLS 1.2 or higher in transit; AES-256 at rest; HSM-protected keys held in a per-institution Key Vault Premium.	LIVE SPECIFIED
Non-production data	Personal information masked in non-production environments; no bank data in development or test without the bank's approval. Public prototypes use fictional institutions and synthetic data only.	SPECIFIED BUILT

[1] Zero data retention uses Azure OpenAI's modified abuse monitoring configuration: prompts and completions are not stored by the model provider for abuse-monitoring review, and no customer content is retained by the model provider.

5.1 Data Mapping & Discovery: metadata-only by design

Every SmarterFI engagement starts with Data Mapping & Discovery, a small containerized tool that runs inside the bank's environment (or in SmarterFI's private cloud) and documents the core, vendor systems and how they connect LIVE.

- **Read-only, metadata-only core scan.** The scan reads system names, schemas, volumes and integration methods. It never reads customer, account or transaction records unless the bank approves. LIVE
- **Optional, bank-approved data-quality check.** If the bank turns it on, it runs inside the bank's environment, compares record values across systems and reports only counts and patterns (for example, the share of loan-system emails that do not match the core). LIVE
- **Least privilege and full logging.** The bank supplies least-privilege credentials; every scan action is logged with user, timestamp, target and scope; the bank can review and redact results before they influence scoring or exports. SPECIFIED
- **Results stay with the bank.** Outputs are a systems map, a readiness score across seven areas and a deployment plan; the bank retains the full package and shares only an approved summary. LIVE SPECIFIED

5.2 Core connections: read-only by default

Analytical applications connect to the core read-only; for example, Lending Co-Pilot core connectors are read-only in pilot, and Contract & Invoice Assurance reads GL and AP data SPECIFIED. Write-back is limited to applications whose purpose requires it. Onboarding Co-Pilot books approved accounts to the core only after a person approves them, uses idempotency keys so a retry cannot create a duplicate account, logs every core call and can be paused by a kill switch BUILT.

6 Audit Logging and Examiner Evidence

A regulator should be able to ask how AI was used and receive a complete, traceable answer. SmarterFI treats the audit log as a product feature, not a by-product.

Attribute	SmarterFI control	Status
What is logged	Prompts, outputs, model versions and user actions LIVE. In applications: every applicant step, vendor call, AI call (feature, masked input summary, output, model and prompt version), reviewer action, rule change, core call, approval, override and export, with who, what, when and before/after values.	LIVE SPECIFIED
Immutability	Append-only by design; tamper evidence through a hash chain that links each event to the one before it (demonstrated in the Onboarding Co-Pilot prototype); immutable storage policies on audit containers in production.	LIVE BUILT
Location	Application logs stay in the bank's dedicated environment, so prompts and outputs are not centralized; only security metadata is sent to SmarterFI's security monitoring.	LIVE
Retention	Security logs and audit records retained for the periods defined in the customer agreement, with audit records archived to immutable storage and retention configurable to the bank's records schedule. Where a rule sets a minimum (e.g., BSA CIP records five years after account closure), retention cannot be configured shorter; legal holds supported; deletions logged.	LIVE SPECIFIED
Exportability	Filters by application, customer, user, event type and date; CSV and JSON export; SIEM export (syslog/JSON); examiner packs by date range, user, topic or sign-off status, including the source snapshots referenced; exam sampling tools. On exit, audit logs are exported with the bank's data.	SPECIFIED BUILT
Reconstruction	Any finding or answer can be reconstructed: inputs, citations to document, page and line, rule and model versions, and every human decision.	BUILT SPECIFIED
Examiner and auditor access	Read-only auditor role in each application; time-boxed read-only access for auditors; SmarterFI acknowledges examination by the bank's regulators under the Bank Service Company Act.	BUILT SPECIFIED

> **Why this matters for examiners**
> Because the model never calculates and never decides, the audit question reduces to three verifiable facts: which source the value came from, which deterministic rule used it, and which named person made the decision.

7 Model Governance Mapped to Model Risk Guidance

On April 17, 2026 the agencies issued Revised Guidance on Model Risk Management (Federal Reserve SR 26-2; OCC Bulletin 2026-13), which supersedes SR 11-7 and OCC Bulletin 2011-12. The revised guidance is expected to be most relevant to banks over $30 billion in assets, and it places generative and agentic AI outside its formal scope. The agencies nevertheless expect a bank's risk-management and governance practices to guide appropriate controls for

tools outside the guidance. SmarterFI therefore applies the familiar SR 11-7 disciplines, carried forward in SR 26-2, to its generative components.

Model risk principle	SmarterFI control	Status
Model inventory	Per-application inventory: component, type, purpose, version, risk tier, SmarterFI and bank owners, evaluation and status. Deterministic rules engines are identified as software, not models. Vendors' own models (e.g., identity and fraud scores) are recorded as third-party entries.	BUILT SPECIFIED
Development and documentation	For each AI feature: intended use, scope and known limitations; design and methodology; data categories, including the foundation model; controls against inaccurate or fabricated outputs; testing and results; monitoring metrics; change log.	SPECIFIED
Validation and effective challenge	Cooperation with the bank's own or independent validator, including test environments and results. Regression and gold-set evaluation per release. Citations let any reviewer challenge any value.	SPECIFIED
Ongoing monitoring	Monitoring for performance degradation, drift and error patterns; prompt notice to the bank of issues likely to cause materially inaccurate outputs; acceptance, override and reviewer-edit metrics.	SPECIFIED BUILT
Change management	Advance notice, as defined in the customer agreement, before a material change, such as a new foundation model, with release notes, updated documentation and, where feasible, a test environment first. Model versions pinned; provider retirements planned as a standing activity.	LIVE SPECIFIED
Use limits	No automated credit decisions, account-opening denials or other adverse decisions without human review unless the bank requests and approves them in writing under its own model risk process.	SPECIFIED
Explainability and fair lending	Where an output supports a credit decision, principal factors and source data are shown to support specific adverse-action reasons under ECOA and Regulation B. No prohibited-basis inputs or known proxies; disparate-impact testing before release and after material changes, with results shared.	SPECIFIED
Accountability	The bank owns its AI program, risk classifications and decisions; SmarterFI supplies the evidence and controls. Detailed materials are available for review under confidentiality.	SPECIFIED

8 Regulatory, Third-Party and Assurance Mapping

Framework	Expectation	SmarterFI control	Status
FFIEC IT Examination Handbook	Sound information security, architecture, operations, business continuity and outsourcing oversight	Written information security program aligned to NIST CSF 2.0 or ISO/IEC 27001, taking the FFIEC Handbook into account, reviewed annually. Azure Policy NIST SP 800-53 Rev. 5 initiative as continuous control evidence. (The FFIEC CAT was sunset August 31, 2025.)	LIVE SPECIFIED
GLBA and Interagency Information Security Standards	Safeguard customer information; protect against anticipated threats and unauthorized access; oversee service providers	Program objectives mirror the Interagency Guidelines; Regulation P reuse and redisclosure limits; no cross-bank sharing; Data Processing and Information Security Addendum with minimum controls.	SPECIFIED

Framework	Expectation	SmarterFI control	Status
Interagency TPRM guidance (2023; Sept. 11, 2026 proposal)	Due diligence, contract provisions, ongoing monitoring and termination	Due-diligence package; agreement mapped to TPRM contract topics; audit rights; subcontractor list with advance notice of changes; ongoing monitoring information (financial condition, BC/DR results, subcontractors, assurance reports, model documentation); notice of material changes as defined in the customer agreement.	SPECIFIED
Bank Service Company Act, 12 U.S.C. 1867(c)	Services subject to examination as if performed by the bank	Contractual acknowledgment; regulator access to personnel, records, systems and subcontractors; support for the bank's service-relationship notice.	LIVE SPECIFIED
Computer-Security Incident Notification Rule	Bank service providers notify customer banks of incidents disrupting services for four or more hours	Notice to the bank of security incidents involving bank data within the timeframe defined in the customer agreement; notice as soon as possible for disruptions of four hours or more, consistent with the rule; written root-cause analysis; incident response plan tested annually.	SPECIFIED
SOC 2 (Security, Availability, Confidentiality)	Independent assurance over vendor controls	SmarterFI's controls are designed to SOC 2 from day one LIVE. The SmarterFI landing zone runs on Microsoft Azure, whose infrastructure carries Microsoft's SOC 2 Type II attestation; Azure SOC reports are available to banks for their third-party file. SmarterFI also provides a standardized security questionnaire (SIG Lite or CAIQ) and its information security policies on request. SmarterFI does not represent that it holds its own SOC 2 report.	LIVE SPECIFIED
Business continuity and disaster recovery	Resilience to region failure, cyber events and vendor failure	Plan covers region failure, key personnel, subcontractor failure, ransomware and facility disruption. Recovery time and recovery point objectives are defined in the customer agreement; encrypted, tamper-protected backups to a separate U.S. region (Central US); stateless compute redeployed from code; annual restore test with results available on request.	LIVE SPECIFIED
Identity and access	Least privilege, MFA, access review	SSO through the bank's identity provider; MFA for all administrative and remote access; quarterly privileged-access review; prompt removal of access on a change of role or departure.	LIVE SPECIFIED
Exit	Orderly termination and data return	Bank owns data and outputs; export in CSV, JSON or Parquet with a data dictionary and audit logs; NIST SP 800-88 deletion with an officer's certificate; transition assistance as defined in the customer agreement.	SPECIFIED

9 Aligned with 2026 Regulatory Frameworks

In 2026 three community-bank resources set out practical expectations for AI security and governance. The crosswalks below map each recommendation to the SmarterFI control that meets it, and note where SmarterFI's control goes further. These resources are industry guidance, not regulation, and reference to them does not imply endorsement of SmarterFI by ICBA, Compliance Alliance, the American Bankers Association or Verapath.

9.1 ICBA Community Banker AI Security Readiness Guide (June 3, 2026)

The crosswalk below is based on ICBA's published summary of the guide (its June 3, 2026 announcement) and on Compliance Alliance's published description of the topics it covers.

ICBA topic (as publicly described)	SmarterFI control that meets it	Where SmarterFI exceeds
"The AI threat shift", including AI-enabled fraud and other emerging threats	Onboarding Co-Pilot identity, document and selfie checks with OFAC and fraud screening; Prompt Shields; WAF ingress; Defender for Cloud, including AI threat protection	AI Data Guard also detects and controls sensitive data that staff share with public AI tools, closing the most common internal leakage path
"Strengthening third-party due diligence"	Due-diligence package: security overview, BC/DR summary, Azure SOC reports, security questionnaire, subcontractor list	Model documentation per AI feature and a model inventory delivered alongside standard TPRM artifacts; time-boxed auditor access to the environment
Vendor dependence	Bank owns data and outputs; full export with data dictionary and audit logs; certified deletion	Per-institution environment supports a clean exit with export and transition assistance; modular apps can be adopted or exited individually
Contract language	Customer agreement structured on the interagency TPRM contract topics, with regulator access and Bank Service Company Act acknowledgment	AI-specific terms: human review, no automated adverse decisions by default, no training on bank data, advance model-change notice, fair-lending testing
"Updating and testing incident response plans"	Incident response plan tested annually; defined notice timelines and content; cooperation with regulators and forensics	Notice obligations aligned to the computer-security incident notification rule; regular updates during active response; written root-cause analysis
Board oversight	Governance evidence for the bank's AI committee and board: feature approvals, risk tiers, evaluation and override metrics	Every AI decision point is reconstructable from the audit log, giving the board evidence rather than attestations
Cyber insurance considerations	Insurance obligations defined in the customer agreement	Breach-cost allocation for incidents caused by SmarterFI or its subcontractors addressed in the customer agreement

9.2 Compliance Alliance: "AI Governance: A Practical Priority for Community Bank Compliance" (August 2026)

Recommendation (quoted or closely paraphrased)	SmarterFI control that meets it	Where SmarterFI exceeds
"Start with a complete inventory" of every AI system, model and vendor product that uses AI	Model inventory per application, including third-party vendor models	Data Mapping & Discovery maps the bank's wider system landscape, surfacing where AI and data dependencies sit across the core and vendors

Recommendation (quoted or closely paraphrased)	SmarterFI control that meets it	Where SmarterFI exceeds
Assess each system for fairness, transparency, explainability and consumer impact	Fair-lending and disparate-impact testing; principal factors and source data shown for credit outputs	Every extracted value cites document, page and line; values that fail citation are discarded
"Prohibit the entry of nonpublic personal information … into external AI tools unless the arrangement has been reviewed and approved"	AI Data Guard masks identifiers before any model call; pasted identifiers are blocked and logged	Production has no network route to public AI APIs; the Data Guard application monitors staff use of public AI tools
Apply model risk principles; for generative tools, "document the purpose, data-handling rules, human-review requirements, and escalation paths"	Per-feature documentation of intended use, limits, data, testing, monitoring and change log	SR 26-2 disciplines applied to generative components even though they are outside its formal scope
Vendors should disclose AI use and provide information on data sources, decisioning, testing, controls and documentation	Foundation model and configuration disclosed on request; subcontractor list; validation support	Advance notice before a material model change, with a test environment first where feasible
Written AI policy addressing inventory, human oversight, data protections, vendor expectations and board reporting	Features off by default and enabled only after bank approval, giving the bank's policy an enforcement point	Risk tier and approval record attached to each feature in the application itself
Human oversight for material or higher-risk uses; "Document where that review occurs and who is responsible"	Human approval gates, each recorded in the audit log with the named user	Maker-checker on rule changes; a person, not the AI, sends any external communication
Lifecycle monitoring built into risk and compliance calendars	Drift and degradation monitoring with notice to the bank	Acceptance, override and reviewer-edit metrics available as ongoing performance evidence

9.3 The AI Framework for Community Banks (Verapath; published on ABA.com, June 26, 2026)

The framework's A.I. 1.0 stage lists required cybersecurity controls and vendor oversight steps that must be in place before any AI tool is deployed. Requirements are quoted verbatim.

Framework requirement (verbatim)	SmarterFI control	Status
"Logging of all A.I. interactions and outputs"	Append-only log of prompts, outputs, model versions and user actions	LIVE BUILT
"Restriction to approved users, devices, and environments"	SSO via the bank's identity provider, role-based access, U.S.-only sign-in, per-feature enablement	LIVE BUILT
"Prohibition of uploads to public or unmanaged systems"	No route from production to public AI APIs; Data Guard controls on public AI tools	LIVE
"Defined anomaly and misuse detection rules"	Rate limits and lockouts; Sentinel analytics; Defender for AI services	LIVE BUILT
"Reviewable audit trails tied to users and use cases"	Every event records user, role, feature and record reference	BUILT
"Secure storage of all A.I. outputs"	Outputs stored in the bank's dedicated environment, encrypted, with immutable audit archive	LIVE
"Encryption or masking of sensitive customer and institutional data"	AES-256 at rest, TLS 1.2+, AI Data Guard masking before model calls	SPECIFIED

Framework requirement (verbatim)	SmarterFI control	Status
“Integration of A.I. activity into existing security monitoring processes”	Security metadata to SIEM; SIEM export (syslog/JSON) to the bank's tools	LIVE SPECIFIED
“Full third-party due diligence”	Due-diligence package and questionnaire support	SPECIFIED
“Review of data handling, residency, and subcontractors”	U.S.-only processing; subcontractor list including model providers	LIVE SPECIFIED
“Verification of access controls, logging, and supervision capabilities”	Read-only auditor role; walkthrough of controls; audit rights	BUILT SPECIFIED
“Confirmation that bank data is not used for external model training”	Contractual no-training commitment and provider configuration	LIVE SPECIFIED
“Assignment and documentation of vendor risk classification”	Owned by the bank; SmarterFI supplies the evidence to support it	Bank

10 Risk Committee Checklist and FAQ

10.1 Pre-approval checklist

- [ ] The use case is approved and risk-tiered by our AI governance committee; only approved SmarterFI features are enabled.
- [ ] We have confirmed the hosting model (a dedicated environment in SmarterFI's private cloud).
- [ ] Data classes in scope are defined, and AI Data Guard masking is confirmed for our identifiers.
- [ ] The agreement contains U.S.-only processing, no-training and no cross-bank sharing terms; the foundation model and its configuration are disclosed.
- [ ] We have reviewed the subcontractor list and the change-notice process.
- [ ] We have received model documentation: intended use, limitations, testing results, monitoring metrics and change log.
- [ ] Every human checkpoint is mapped to a named role in our organization.
- [ ] We have tested an audit-log export and set retention to our records schedule.
- [ ] Incident contacts are designated and notice timelines are agreed.
- [ ] We have the BC/DR summary and latest restore-test results; RTO and RPO are defined in our agreement.
- [ ] We have SmarterFI's security evidence (SIG Lite or CAIQ, policies, Microsoft Azure SOC reports).
- [ ] The exit plan, export formats and deletion certificate are confirmed.
- [ ] Staff complete training before they receive access.

10.2 Frequently asked questions

Does SmarterFI train models on our data?
No. Bank data is not used to train, fine-tune or improve any model used for another customer, and the foundation model services are configured so the provider does not train on it.

Can the model see our customers' Social Security or account numbers?
No. AI Data Guard masks or tokenizes identifiers before any model call. Reviewers see real values in the application; the model does not.

What happens if the model is wrong?
Its output cannot become a number of record: values must cite a verifiable source or are discarded, calculations run in deterministic code, and a person decides. Errors are visible in the audit log and in override metrics.

Where is our data, and who can reach it?
In an Azure environment dedicated to your institution, in U.S. regions only. Administrative access is just-in-time, MFA-protected and reviewed quarterly; administrative sign-in from outside the United States is blocked by policy.

Does SmarterFI have SOC 2?
SmarterFI does not claim its own SOC 2 report. Its controls are designed to SOC 2 from day one, and its landing zone runs on Microsoft Azure, whose infrastructure carries Microsoft's SOC 2 Type II attestation. SmarterFI provides a standardized security questionnaire, its policies and Microsoft Azure's SOC reports on request.

Does SmarterFI make credit or account decisions?
No. The platform prepares the work and drafts recommendations; it does not make automated credit decisions or account-opening denials unless your institution requests and approves such a capability in writing.

Do the public demonstrations use real bank data?
No. SmarterFI's prototypes use fictional institutions and synthetic data running against a simulated core.

Can we leave?
Yes. You own your data and outputs. SmarterFI exports them with a data dictionary and audit logs, deletes them to NIST SP 800-88 with a certificate, and provides transition assistance as defined in your agreement.

A Sources

1. SmarterFI.ai, Platform & Solutions and Technology & Architecture pages (accessed October 21, 2026). https://smarterfi.ai/platform ; https://smarterfi.ai/platform/technology
2. Federal Reserve SR 26-2 and OCC Bulletin 2026-13, Revised Guidance on Model Risk Management (April 17, 2026). https://www.federalreserve.gov/supervisionreg/srletters/SR2602.pdf ; https://www.occ.gov/news-issuances/bulletins/2026/bulletin-2026-13.html
3. OCC Bulletin 2023-17, Third-Party Relationships: Interagency Guidance on Risk Management (June 6, 2023). https://www.occ.gov/news-issuances/bulletins/2023/bulletin-2023-17.html
4. FDIC, Proposed Interagency Third-Party Risk Management Guidance and Joint Statement on Community Banks' Engagement with Core Service Providers (September 11, 2026). https://www.fdic.gov/news/financial-institution-letters/2026/proposed-interagency-third-party-risk-management-guidance
5. OCC Bulletin 2021-55, Computer-Security Incident Notification: Final Rule. https://www.occ.gov/news-issuances/bulletins/2021/bulletin-2021-55.html
6. 12 U.S.C. 1867(c) (Bank Service Company Act); 15 U.S.C. 6801 (GLBA); 12 C.F.R. Part 364, Appendix B (Interagency Guidelines Establishing Information Security Standards).
7. FFIEC, Cybersecurity Assessment Tool sunset (August 31, 2025). https://www.ffiec.gov/resources/cat
8. ICBA, “ICBA Releases Community Banker Guide on Artificial Intelligence Security Readiness” (June 3, 2026). https://www.icba.org/w/icba-releases-community-banker-guide-on-artificial-intelligence-security-readiness
9. Compliance Alliance, “AI Governance: A Practical Priority for Community Bank Compliance” (August 2026). https://compliancealliance.com/news-events/newsletter/august-2026-newsletters/ai-governance-a-practical-priority-for-community-bank-compliance/
10. Verapath, The AI Framework for Community Banks (published on ABA.com, June 26, 2026). https://www.aba.com/news-research/analysis-guides/the-ai-framework-for-community-banks
11. Microsoft Learn, Data, privacy, and security for Foundry Models sold by Azure; Abuse monitoring; Content filtering; Deployment types. https://learn.microsoft.com/en-us/azure/foundry/responsible-ai/openai/data-privacy

SmarterFi.ai
AI for Community Banking

Request a technical briefing

See the governed path in a live walkthrough, review the model documentation pack, or run Data Mapping & Discovery against your own environment.

WEB
www.smarterfi.ai

EMAIL
jamie.donahue@smarterfi.ai

BOOK 30 MINUTES
calendly.com/jamie-donahue-smarterfi/30min

ABOUT THE AUTHOR

James (Jamie) Donahue is Founder, President & Chief AI Officer of SmarterFI.ai (2026–present), where he is building AI for community banks and credit unions. He was previously President & Chief Digital Officer of First Carolina Bank's BM Technologies Division (2025–2026) and President & Chief Technology Officer of BM Technologies (2020–2025). 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. He is based in Palm Beach, Florida. LinkedIn: linkedin.com/in/jamesmdonahue

CITE AS
Donahue, J. (2026). Eliminating Frontier Model Risk: SmarterFI Technical Architecture Brief. SmarterFI.ai. Version 1.1. https://doi.org/10.5281/zenodo.23884799. CC BY 4.0

REVISION
Version 1.1 (November 3, 2026): added citation and DOI line; clarified that the SmarterFI platform runs only in SmarterFI's private cloud and that only Data Mapping & Discovery can run in a bank's environment; licence notice aligned to CC BY 4.0. Version 1.0 published October 27, 2026.

© 2026 SmarterFI.ai. Licensed CC BY 4.0.This brief describes SmarterFI's controls and architecture as of October 2026; status labels indicate the evidence for each control. It is not legal advice, does not modify any agreement, and does not imply endorsement by any organization cited. Each institution remains responsible for its own model risk, third-party risk, compliance and fair-lending programs.