

Seekr for Financial Services

Close the gap between deploying AI and defending it

Financial institutions are deploying AI into credit decisioning, fraud detection, compliance screening, portfolio analysis, and regulatory reporting. The pressure to move faster is real. So is the pressure to answer for it.

Every one of these workflows produces decisions that someone will eventually examine. A regulator, an auditor, a risk committee, a counterparty in litigation. When that moment arrives, the institution needs more than confidence in its technology. It needs evidence.

Seekr® is trusted AI for the decisions you have to defend. Certify a model, take on a workload your team dreads, or bring AI to data you can't move. Every output is explainable and holds up when someone asks how you got there.

The AI accountability gap in financial services

AI is advancing faster than your ability to trust it with decisions that carry real weight. Under pressure to deploy, the gap between what AI can do and what you can answer for leaves you without good options: stall progress, or move ahead without knowing you can defend the results.



Model selection without evidence.

Teams often select AI models using generic benchmarks or vendor claims that don't reflect your institution's data, risk thresholds, or regulatory requirements. When an examiner challenges a model choice, there's no defensible record of how or why it was selected.



No visibility into what changed.

Fine-tuning and prompt changes alter model behavior, but teams rarely know exactly where performance improved, where it regressed, or how risk shifted. Every change introduces uncertainty without a before-and-after evidence trail.



Observability is not explainability.

Existing monitoring tools log what happened: which documents were retrieved, which tools were called, how long each step took. That's useful for debugging. It's not sufficient for defending a decision. Logs show process. They don't show influence.



Governance treated as a one-time approval.

Models drift as markets move, customer behavior changes, and regulations evolve. A model approved last quarter may no longer meet the criteria it was approved against, and most institutions have no mechanism to detect or document that shift.



Institutional data left on the table.

Financial institutions hold decades of proprietary data: loan performance histories, transaction patterns, customer behavior, internal decisioning records. Most of this data sits unused because teams lack a safe, governed way to build AI on it without exposing it to third-party model providers or losing control of how it's used.

How Seekr works for financial services

Seekr closes the gap between deploying AI and defending its decisions. The platform spans the full AI lifecycle, from data preparation through production monitoring, and produces the evidence trail financial services teams need at every stage.

Prove why every decision holds up

You can't use AI for decisions you'll have to defend if you can't show your work. Explainability starts before the model runs. You set how much latitude it has, how deterministic each step is, and where a human stays in the loop. From there, you can watch every step, tool call, and output as it happens, and trace the reasoning and provenance behind any output across the model, its tools, and its data. When an output is wrong, you can challenge it, correct it, and confirm the correction. And before any model ships, you can test, score, and certify it for bias, accuracy, and risk.

Keep your data, your models, and your control

Your data is your competitive advantage, and it can't leave your environment to train someone else's model. But keeping it safe isn't enough. Seekr lets you activate that data by building domain-specific models on your own data like loan histories, transaction patterns, and other institutional knowledge, and running them where you need them: in the cloud, on premises, or air gapped. You own the model, your knowledge stays yours, and the intelligence you build compounds over time, rather than subsidizing a frontier lab's next release.

Get decision-ready results without runaway cost

Seekr helps you right-size continuously, matching the right model to each task on your own data and criteria, and providing evidence to defend every choice on cost and performance. You use the least resource-intensive model proven to meet the requirement and escalate to a more capable model when the work demands it, so cost stays predictable as usage scales.



The products

seekrFlow™

SeekrFlow™ is the platform. It handles data preparation, model fine-tuning, agentic workflow orchestration, and production monitoring in a single environment. For financial services, this means you prepare training data from unstructured loan files, regulatory filings, and transaction records with full lineage. You fine-tune domain-specific models on your institutional data. You deploy agents that execute multi-step compliance workflows under your oversight. And you monitor everything from one control point.

seekrGuard

SeekrGuard™ lets risk and compliance teams evaluate AI models using their own data, scenarios, and risk criteria as the benchmark. You compare candidate models side by side, define institution-specific risk profiles, and score for bias, accuracy, and failure modes before any model reaches a production decision. As models evolve, SeekrGuard supports repeatable testing against the same criteria, creating a documented record of performance, risk, and change over time.

What your institution gets

- ☒ **Explainability at the architecture level.**
 Every output carries an influence score and a source trail that traces the answer back to the data, the reasoning, and the model version that produced it. The explainability is built into the architecture at every stage of the workflow.
- ☒ **Evidence that survives examination.**
 When an examiner, an auditor, or a risk committee asks how a decision was reached, the evidence is already assembled. No reconstruction, no reverse engineering, no plausible narratives.
- ☒ **Continuous governance.**
 Drift detection, ongoing re-evaluation, and a documented performance record that reflects what the model is doing now, not what it did when you first approved it.
- ☒ **Sovereignty over your data and models.**
 Own your models, and bring them to your data instead of sending your data to a third-party. No vendor lock-in, no black-box constraints, no dependency on a frontier lab's roadmap or pricing changes.
- ☒ **Deployment where your data lives.**
 Cloud, on premises, or air gapped. Your deployment meets the data residency and cross-border transfer requirements your regulators expect.
- ☒ **More value from your institutional data.**
 Turn decades of proprietary loan histories, transaction patterns, and decisioning records into domain-specific AI models that compound your competitive advantage rather than sitting unused.
- ☒ **AI economics that scale.**
 Match the right model and the right compute to each workload, with evidence to defend every choice on cost and performance. Cost stays predictable as usage grows because you use the least resource-intensive model proven to meet the requirement.

Use cases

Credit risk model governance.

Compare candidate credit models against your institution's historical downturn and stress scenarios, score for bias across protected classes, and produce audit-ready evidence for model risk management reviews.

Fraud detection and compliance.

Deploy AI agents that clear routine low-risk transactions autonomously while escalating edge cases for human review, reducing false positives and incorporating feedback to align with evolving regulations.

Regulatory document processing.

Automate high-volume compliance workflows such as filings, disclosures, and audit evidence assembly with full traceability at every step.

Anti-money laundering (AML) screening.

Apply domain-specific models to transaction monitoring and suspicious activity detection with explainable outputs that satisfy examiner scrutiny.

Know Your Customer (KYC) document review.

Accelerate identity verification and due diligence document review while maintaining the audit trail required for regulatory examination.

Model drift detection and re-evaluation.

Continuously monitor production models for shifts in accuracy, bias, or risk posture as market conditions change, and retest against original approval criteria.

Portfolio analysis and reporting.

Generate explainable, source-attributed analysis across portfolio holdings, risk exposures, and performance metrics for internal review and client reporting.

Regulatory stress testing.


Test AI-assisted models against scenario-specific stress conditions and produce structured evidence of model behavior under adverse conditions.

Claims and dispute evidence assembly.

Assemble traceable evidence trails for customer disputes, counterparty challenges, or litigation where AI-assisted decisions are at issue.

Internal audit and examination preparation.

Produce structured, reproducible records of model selection, evaluation, deployment, and ongoing performance for internal audit cycles and regulatory examinations. every step.






Why Seekr

Observability tells you what happened. Seekr tells you what influenced the outcome. In a regulatory examination or a dispute, the difference between showing your workflow and showing your evidence is the difference between describing a process and defending a decision.

Financial institutions need AI they can stand behind when it matters. Seekr gives your risk, compliance, and audit teams the evidence to do exactly that.



Next steps

Book a 30-minute technical briefing to see how Seekr works for financial services risk, compliance, and credit decisioning. No commitment required.