

Seekr for Retail

Close the gap between automating decisions and standing behind them

Retailers are deploying AI into demand forecasting, pricing, personalization, inventory management, loss prevention, and customer experience. The pressure to protect margins, reduce stockouts, and meet customer expectations is real. So is the pressure to explain how those decisions are made when a customer, a regulator, or a brand partner asks.

Every one of these workflows produces decisions that someone will eventually question. An FTC enforcement action on pricing practices, a state attorney general investigating personalization bias, a brand partner demanding transparency on how their products are merchandised, a customer challenging a fraud flag on their account. When that moment arrives, the organization 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 retail

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.



Pricing and personalization under scrutiny.

AI models set prices, personalize offers, and target promotions across millions of customers, but when a regulator investigates discriminatory pricing or a consumer advocacy group challenges your personalization practices, there's often no structured record of what influenced the decision.



Demand and inventory decisions without traceability.

Teams adopt AI for demand forecasting and inventory allocation, but the models produce recommendations without traceable reasoning. When a stockout costs millions or an overstock forces markdowns, the planning rationale is difficult to reconstruct and defend.



Customer data across channels.

Retail AI touches loyalty data, purchase histories, browsing behavior, and payment information across online, mobile, and in-store channels. Models running on this data require deployment options that keep customer information within controlled environments and comply with privacy regulations across jurisdictions.



Governance treated as a one-time approval.

Consumer behavior shifts seasonally, competitive landscapes change, supply chains disrupt, and regulations evolve. A model approved last quarter may no longer meet the criteria it was approved against, and most retailers have no mechanism to detect or document that shift.

How Seekr works for retail

Seekr is designed to close the gap between automating decisions and standing behind them. The platform spans the full AI lifecycle, from data preparation through production monitoring, and produces the evidence trail retail teams need at every stage.

Prove why every decision holds up

You can't automate decisions your customers and regulators will scrutinize 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.

Keep your data, your models, and your control

Customer data is your responsibility, and it can't leave your environment to train someone else's model. Seekr lets you build domain-specific models on your own data and run them where you need them: in the cloud, on premises, or air gapped. You own the model, your knowledge stays yours, and nothing rides on a frontier lab's roadmap or fine print.

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 retail, this means you prepare training data from POS systems, loyalty programs, e-commerce platforms, and supply chain records with full lineage. You fine-tune domain-specific models on your proprietary data. You deploy agents that execute multi-step merchandising, pricing, and customer service workflows under your oversight. And you monitor everything from one control point.

seekrGuard


SeekrGuard™ is designed to help merchandising, operations, 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 category-specific and channel-specific risk profiles, and score for accuracy, bias, and failure modes before any model reaches production. As models evolve, SeekrGuard supports repeatable testing against the same criteria, creating a documented record of performance, risk, and change over time.

What your organization 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 scrutiny.**
 When a regulator, a brand partner, or a customer asks how a pricing, personalization, or fraud 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 security and privacy requirements your customers and regulators expect.

Use cases

- Demand forecasting and inventory optimization.**
 Apply AI to demand prediction, allocation, and replenishment with traceable, defensible outputs that document the reasoning behind every inventory decision.
- Dynamic pricing and markdown optimization.**
 Deploy AI-assisted pricing that produces the evidence trail your legal and compliance teams require, with explainable outputs that document what influenced every price change. For surge and event-driven pricing, prove that adjustments are driven by actual supply chain dynamics rather than opportunistic markups.
- Personalization and recommendation.**
 Build personalization models on your own first-party data, running in your environment, with explainable targeting outputs that satisfy privacy regulations and fairness requirements.
- Loss prevention and fraud detection.**
 Deploy AI agents that flag suspicious transactions and returns while producing explainable outputs that support investigation, customer communication, and dispute resolution.
- Supply chain visibility and vendor assessment.**
 Apply domain-specific models to supplier evaluation, lead time prediction, and disruption risk assessment with traceable outputs that support procurement decisions.
- Trade agreement reconciliation.**
 Reconcile supplier trade agreements, volume rebates, co-op marketing funds, and early payment discounts against actual purchasing and payment data at scale. Recover negotiated savings that quietly disappear when contract terms can't be checked against transaction records manually.
- Customer experience and service automation.**
 Deploy governed AI agents for customer service, order management, and issue resolution with audit-ready documentation that traces what influenced each customer interaction.
- Assortment planning and merchandising.**
 Optimize category assortment, planogram placement, and seasonal planning with AI that documents the reasoning behind every merchandising recommendation.
- Store operations and workforce optimization.**
 Apply AI to labor scheduling, task allocation, and store performance analysis with explainable outputs that support operational reviews and workforce planning.
- Marketing attribution and campaign optimization.**
 Apply AI to media mix modeling, attribution analysis, and campaign targeting with traceable outputs that satisfy brand partners and internal stakeholders.
- Model drift detection and re-evaluation.**
 Continuously monitor production models for shifts in accuracy or risk posture as consumer behavior, competitive dynamics, and seasonal patterns change, and retest against original approval criteria.






Why Seekr

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

Retailers need AI they can stand behind when it matters. Seekr gives your merchandising, operations, and compliance teams the evidence to do exactly that.



Next steps

Book a 30-minute technical briefing to see how Seekr works for retail operations, merchandising, and customer experience. No commitment required.