

Seekr for Healthcare

Close the gap between deploying AI in clinical and administrative workflows and proving it performs safely

Health systems, payers, and life sciences organizations are deploying AI into clinical decision support, medical coding, patient risk stratification, claims processing, and administrative operations. The pressure to reduce clinician burden, improve outcomes, and control costs is real. So is the pressure to prove that every AI-assisted output meets the clinical, regulatory, and privacy standards healthcare demands.

Every one of these workflows produces decisions that someone will eventually examine. A Food and Drug Administration (FDA) reviewer evaluating software as a medical device, a Health Insurance Portability and Accountability Act (HIPAA) auditor assessing Business Associate Agreement (BAA) compliance and data handling, a Centers for Medicare & Medicaid Services (CMS) auditor reviewing coding accuracy, a malpractice attorney questioning a clinical recommendation, a patient asking why a treatment path was suggested. 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 healthcare

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.



Clinical outputs without defensible reasoning.

AI models support diagnoses, flag patient risks, and recommend treatment pathways, but when a clinician needs to explain the basis for a recommendation to a patient, or a malpractice attorney challenges a clinical decision, there's often no structured record of what influenced the output.



PHI that cannot leave the environment.

Healthcare AI touches protected health information across EHRs, imaging systems, lab results, and claims data. HIPAA permits PHI processing by BAAs under executed agreements, but organizations need deployment options that enforce contractual boundaries, satisfy the BAA chain from the Electronic Health Record (EHR) vendor through the AI platform, and prevent data from reaching third-party model providers outside that chain.



Regulatory complexity across clinical and administrative AI.

Clinical AI may fall under FDA oversight as software as a medical device. Administrative AI must satisfy CMS coding accuracy requirements and payer audit standards. Most organizations have no systematic way to evaluate, document, and maintain compliance across both domains simultaneously.



Governance treated as a one-time approval.

Patient populations change, clinical guidelines evolve, coding rules update, and disease prevalence shifts. A model approved last quarter may no longer meet the criteria it was approved against, and most health systems have no mechanism to detect or document that shift.

How Seekr works for healthcare

Seekr is designed to close the gap between deploying AI and proving it performs safely. The platform spans the full AI lifecycle, from data preparation through production monitoring, and produces the evidence trail healthcare teams need at

Prove why every decision holds up


You can't use AI for decisions clinicians, patients, 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

Patient data is your most sensitive asset. HIPAA permits processing by Business Associates under executed BAAs, but your data should never train another organization's models or leave the BAA chain you've approved. 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 healthcare, this means Seekr can connect to clinical data sources through standard interoperability protocols like Fast Health Interoperability Resources (FHIR) and Health Level 7 (HL7), subject to the BAA chain and access approvals each system requires. Data preparation can handles claims databases, clinical registries, and structured EHR reports with full lineage. You can fine-tune domain-specific models on your institutional data. You

seekrGuard

SeekrGuard™ is designed to help clinical informatics, compliance, and operations teams evaluate AI models using their own data, patient populations, and clinical criteria as the benchmark. For imaging workflows, Seekr can support de-identification workflows for PHI embedded in Digital Imaging and Communications in Medicine (DICOM) metadata and pixel data before any training use. You compare candidate models side by side, define specialty-specific and population-specific risk profiles, and score for accuracy, safety, and bias before any model reaches clinical or administrative production. As models evolve, SeekrGuard supports repeatable testing against the same criteria, creating a documented record of performance, risk, and change over time.

Use cases

Clinical decision support governance.

Evaluate and monitor AI models that support diagnoses, risk scoring, and treatment recommendations, helping produce the evidence trail clinicians, patients, and regulators require.

Medical coding and billing automation.

Deploy AI agents for ICD, CPT, and DRG coding with explainable outputs that document what influenced each code assignment and satisfy CMS audit requirements.

Patient risk stratification.

Apply domain-specific models to patient populations for readmission risk, deterioration prediction, and care management prioritization with traceable, defensible outputs.

Claims processing and denial management.

Automate claims adjudication, prior authorization, and denial appeals with full traceability at every step and explainable outputs that support payer and provider review.

Clinical documentation and note generation.

Deploy governed AI agents for clinical note generation, discharge summary creation, and documentation workflows with audit-ready evidence of what influenced each output.

Radiology and pathology AI governance.

Evaluate and monitor imaging AI models for diagnostic accuracy, bias across patient demographics, and performance drift, helping produce documented evidence that can support FDA and institutional oversight.

Drug discovery and research support.

Apply AI to literature review, compound screening, and biomarker analysis with traceable outputs that document the reasoning behind every research recommendation.

Clinical trial matching and eligibility.

Automate patient screening and trial eligibility determination with explainable outputs that document what criteria each patient met or missed.

Administrative workflow automation.

Deploy governed AI agents across scheduling, referral management, prior authorization, and revenue cycle workflows with audit-ready documentation at every step.

Model drift detection and re-evaluation.

Continuously monitor production models for shifts in accuracy or risk posture as patient populations change, clinical guidelines evolve, or coding rules update, and retest against original approval criteria.




Why Seekr

Observability tells you what happened. Seekr tells you what influenced the outcome.

In an FDA review, a HIPAA audit, or a malpractice proceeding, the difference between showing your workflow and showing your evidence is the difference between describing a process and defending a decision.

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



Next steps

Book a 30-minute technical briefing to see how Seekr works for healthcare clinical and administrative AI. No commitment required.