

The Enterprise AI and **Agentic Control Plane** for Healthcare

Make Governance Enforceable, Prove Controls Work, and Protect PHI in Every AI Interaction

Your healthcare organization is already running AI across three rapidly expanding vectors: clinical AI models and agents built in-house or integrated into EHRs, public AI applications like ChatGPT and Cursor used by clinical and administrative staff, and embedded AI features inside everyday SaaS platforms like Epic, Cerner, telehealth systems, and medical coding tools.

Agentic AI now handles patient interactions, documentation, scheduling, and care coordination, often accessing PHI autonomously and invoking tools across connected systems. Policies exist on paper. Most organizations cannot prove those controls are consistently enforced across live AI and agentic workflows. That is the control gap.

Singlr closes it. We transform AI governance from static policy into enforceable, measurable control across every AI interaction in your environment, so you can balance the speed of AI adoption with the HIPAA compliance, patient safety, and PHI protection your organization requires.

Every AI interaction must be:

Transparent

Maintain full visibility into which AI service is in use, by whom, and for what purpose, across clinicians, administrative staff, and AI agents operating across your environment.

Secure

Prevent PHI exposure, model training on patient data, prompt injection, and misuse that could impact patient safety across user-facing AI tools and autonomous agent interactions alike.

Compliant

Align with HIPAA, HITECH Act, state privacy laws, and FDA guidelines. Enforce security policies in real time and ensure that Business Associate Agreements (BAAs) are in place across all AI services and agentic workflows.

Auditable

Generate longitudinal, tamper-evident proof of AI control behavior to satisfy internal governance teams, OCR auditors, and breach notification requirements. Not reporting. Verifiable evidence.

Healthcare Leaders are Facing Urgent Questions:

- Are clinicians using AI tools that aren't covered by BAAs?
- Is protected health information being used to train third-party AI models?
- Which AI tools and agents are accessing or processing PHI for documentation or decision support without vetting?
- Are employees using personal AI accounts to handle patient data?
- How do we prove autonomous AI agents stay within HIPAA and patient safety boundaries, not just assume they do?
- Are embedded AI features in our EHR and SaaS platforms creating control gaps we can't see?

"Singlr has given us the confidence to adopt AI at scale without compromising security, compliance, or governance. Their platform provides visibility and control over how AI is used, ensuring alignment with regulations and our internal standards. What stood out was the speed to value. We moved from evaluation to impact in weeks, not months. Singlr has strengthened our security posture while accelerating our ability to deliver AI-driven solutions that advance our mission. They are more than a technology provider; they are a trusted partner helping us responsibly harness the power of AI."

—Lonnie Johnson,
Chief Information Officer, KVC Health Systems

We Put You In Control Of AI™

 Contextual AI Discovery

Singlr uncovers every AI model, agent, and service in your environment, including default-on AI features hidden inside healthcare SaaS, and maps usage by department, user, and data flow. From third-party apps and in-house clinical applications to MCP servers, browser extensions, and fourth-party processors used by your vendors.

 Global AI Intelligence with Singlr Pulse™

Singlr Pulse™ is the real-time risk intelligence engine powering the control plane. It maintains a continuously updated global directory of AI services, agents, and datasets, with in-depth metadata on capabilities, compliance posture, subprocessors, and risk factors, feeding every enforcement and governance decision.

 Accelerated Vetting & Onboarding

Singlr combines Pulse™ intelligence with automated research, continuous vetting, and context-rich reviews to shorten decision cycles and get clinical AI tools into production faster, with confidence.

 AI-Powered Policy Controls

Singlr turns governance intent into live, context-aware controls that allow, guide, restrict, redact, or block AI interactions based on risk, user role, and data sensitivity, including BAA-aware enforcement for PHI. Controls extend to agentic workflows, governing how agents access data and invoke tools across systems, and can be pushed to SASE, cloud security platforms, API gateways, and MDM.

 Pre-Deployment AI Red-Teaming

Simulate targeted attacks on in-house clinical applications, models, and agents using customizable scenarios aligned to responsible AI risks, NIST AI RMF, OWASP Top 10 for LLMs, and healthcare-specific regulatory standards. Validate governance intent before it reaches production.

 Fine-Grained Runtime Detection & Control

Enforce governance at execution time across every AI interaction. Singlr provides enforcement across user activity, PHI/ePHI exposure, file uploads, prompt content, copy-paste actions, account types, and MCP-based services. Critically, Singlr measures whether those controls are actually working, not just whether they are configured.

 Reporting, Compliance, & Tracking

Singlr delivers independent, longitudinal proof that controls are operating as intended across heterogeneous healthcare environments. This is tamper-evident evidence supporting audit readiness against HIPAA, HITECH, NIST AI RMF, SOC 2, and GDPR, with clear visibility into AI adoption, enforcement outcomes, and risk reduction over time.

 Singlr is Integration First by Design:

- **Extensible:** works across diverse environments, no vendor lock-in
- **Accelerates time-to-value:** cuts onboarding from weeks to hours
- **Maximizes ROI:** connects with your existing identity, security, compliance, and DevSecOps tools

The Singlr Unified Control Plane

Singlr operates through three tightly integrated runtime pillars, encircled by the **Singlr Assurance™ Layer**, which provides continuous intelligence and independent proof across all three.



Singlr Assurance™ Layer

The persistent intelligence and proof system that surrounds all three runtime pillars. Not just reporting. Longitudinal, tamper-evident evidence that controls are operating as intended, continuously, across your entire environment.



Singlr Runtime Governance™

Define enforceable intent and risk thresholds across AI systems, agents, and cross-system interactions.



Singlr Runtime Control™

Translate governance into real-time enforceable boundaries across AI interactions and agentic execution paths.



Singlr Runtime Security™

Focus security on true adversarial behavior across AI systems and agents, reducing preventable escalation.

© Singlr AI. All rights reserved.

www.singlr.ai







