

The Enterprise AI and Agentic Control Plane

Make Governance Enforceable. Prove Controls Work. Protect Data in Every AI Interaction

Your organization is already running AI across three rapidly expanding vectors: in-house models and agents, public tools like ChatGPT and Cursor, and embedded features within SaaS platforms such as Grammarly, Notion, and Slack.

Many of these services operate without formal oversight, making it hard to balance rapid AI adoption with the security, compliance, and cost controls your business requires. 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, connecting **governance intent** to **runtime enforcement** for safe, governed innovation at scale.

Every AI interaction must be:

Transparent

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

Secure

Prevent data leakage, model training on sensitive information, prompt injection, and adversarial misuse across user-facing tools and autonomous agentic workflows.

Compliant

Align with evolving regulations including EU AI Act, NIST AI RMF, GDPR, HIPAA, and ISO. Enforce security policies in real time across every AI interaction.

Auditable

Longitudinal, tamper-evident proof of AI control behavior to satisfy internal governance teams, regulators, and auditors. Not context-rich logs. Verifiable evidence.

AI initiatives are raising urgent questions for enterprise leaders:

- How efficiently are new AI tools and projects vetted and onboarded?
- Where is sensitive data being accessed or processed by AI?
- Are employees or agentic workflows exposing confidential data?
- What unsanctioned AI interactions are expanding our risk surface?
- Are AI agents or plugins triggering actions across systems without governance controls?
- How do embedded AI features affect compliance with GDPR, HIPAA, and the EU AI Act?

The Singlr Unified Control Plane

Singlr operates through three tightly integrated runtime pillars, each encircled by the **Singlr Assurance™ Layer**, providing continuous intelligence and independent proof across your entire AI environment.



Powered by the **Singlr Assurance™ Layer**



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.

Contextual AI Discovery

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

Global AI Intelligence with Singulr Pulse™

Singulr 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 in real time.

AI-Powered Policy Controls

Singulr turns governance intent into live, context-aware controls that allow, guide, restrict, redact, or block AI interactions based on risk level, user role, and data sensitivity. 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, working alongside Singulr's native runtime enforcement.

Pre-Deployment AI Red-Teaming

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

Fine-Grained Runtime Detection & Control

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

Reporting, Compliance, & Tracking

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

Accelerated Vetting & Onboarding

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

Singulr is Integration First by Design:

Singulr is built to work across diverse environments with no vendor lock-in. It cuts AI tool onboarding from weeks to hours and connects with your existing identity, security, compliance, and DevSecOps tools to maximize ROI on your current investments.

- **Sample Integrations:**

- Identity and Access: Okta, Azure AD
- Cloud and Infrastructure: AWS, Azure, GCP
- Security Stack: SIEMs, DLPs, endpoint security
- AI Platforms and MCP-based services



List above is not exhaustive

About Singulr AI:

Founded in Palo Alto, California, Singulr AI is the Unified AI Control Plane for the agentic enterprise. The platform transforms AI governance from static documentation into enforceable, measurable runtime control, enabling organizations to scale AI safely across cloud, SaaS, and enterprise.

Singulr is an AI governance and security platform that helps streamline and secure enterprise AI use at scale. Enterprise IT, Security, Privacy, and Risk teams can now accelerate AI-driven innovation while managing the end-to-end AI adoption lifecycle, mitigating risks and costs, all powered by Singulr Pulse™, the industry's leading AI risk intelligence system.

For more information, visit <https://www.singulr.ai>