Kiteworks

TOP 5

Top 5 Reasons Organizations Deploying OpenClaw Need Kiteworks

Every organization moving toward OpenClaw is securing the runtime layer—sandboxing, tool controls, and network guardrails. But runtime security does not answer the questions auditors, regulators, and boards will ask: what data was accessed, under what authorization, with what encryption, and where is the audit trail? Kiteworks Compliant AI, delivered through the Kiteworks Private Data Network, provides the governed data layer—the missing third layer of a complete OpenClaw strategy—with zero-trust AI data access, ABAC policy enforcement, FIPS 140-3 encryption, and tamper-evident audit trails for every interaction with sensitive data.

1. Your Auditor Won't Ask About the Runtime—They'll Ask About the Data

Runtime controls govern how agents execute, but compliance requirements are enforced at the data layer. Organizations consistently struggle to enforce purpose limitations, terminate misbehaving agents, and produce audit evidence on demand. Kiteworks Compliant AI sits between AI agents and the data they access—authenticating every request, enforcing ABAC policies per operation, applying FIPS 140-3 encryption, and capturing tamper-evident audit logs mapped to regulatory controls. When auditors ask for proof, the answer is immediate and defensible.

2. OpenClaw Agents Introduce Structural Risk—Kiteworks Controls the Data Instead

AI agents can be manipulated through social engineering, can take actions beyond their intended scope, and can expose data across channels. These risks are inherent to how agents operate and cannot be fully eliminated at the runtime layer. Kiteworks Compliant AI removes the dependency on trusting the agent by governing access at the data layer—ensuring that every interaction is authenticated, authorized, encrypted, and logged, regardless of agent behavior.

3. Local Execution Expands the Blast Radius Without Central Governance

Running models locally improves data sovereignty, but it also means agents inherit the full privileges of the environment—email, file systems, messaging platforms, and developer tools. Without a centralized control point, risk scales with access. Kiteworks Compliant AI restores that control by enforcing consistent data governance across environments, ensuring that sensitive data remains protected even when agents run outside centralized infrastructure.

4. Shadow AI Agents Are Already Operating Outside IT Visibility

AI agents can be deployed on endpoints without formal approval and are already spreading across enterprise environments. Attempting to block them is ineffective; they will operate wherever users have access to data. Kiteworks Compliant AI addresses this reality by governing the data itself—requiring authentication, enforcing policy, and logging every interaction so that even unmanaged or unauthorized agents cannot access sensitive content without control and visibility.

5. Governance Is What Enables AI to Scale

AI adoption slows when organizations cannot confidently control data access or demonstrate compliance. Manual review processes do not scale with agent-driven workflows. Kiteworks Compliant AI embeds governance directly into the data layer—automating authentication, authorization, encryption, and audit logging for every interaction. This allows organizations to deploy AI faster while maintaining the control, visibility, and auditability required to operate at scale.