

AGENTIC AI SECURITY SERVICES

Lifecycle Management for an Autonomous Workforce

AI agents are becoming autonomous actors in the enterprise, accessing systems, using data, and executing actions often without human oversight. Yet most organizations have not included them in their identity lifecycle, or set trust boundaries or governance controls to manage those agents.

Enable autonomous AI to move faster without sacrificing visibility, security, or control. GuidePoint helps organizations establish the identity and governance foundation needed to securely scale AI agents across the business. Our Agentic AI Security Services provide a practical, prioritized approach to discovering, governing, and managing agents across your environment. From visibility and access governance to lifecycle management and continuous oversight, GuidePoint AI Security Services help teams extend cybersecurity disciplines to a new class of enterprise actors.

Better Manage AI Agents at Scale


As organizations deploy AI agents alongside service accounts, APIs, cloud workloads, CI/CD pipelines, and other non-human identities (NHIs), existing security models and processes built primarily around human users can't keep up. AI agents need credentials, permissions and access like human users, but they operate at machine speed. Agents are emerging as a new class of enterprise identity, and their adoption velocity demands dedicated governance and controls.

Agentic AI security spans three control planes: identity (who or what may act), data (what an agent may access), and runtime (what an agent may do in the moment). Identity is the foundational control plane upon which both data and runtime controls depend.

GuidePoint helps you modernize your cybersecurity, identity, and governance strategies to secure, manage, and scale autonomous AI across your environment.

Our approach helps organizations:

- ✓ Discover and inventory AI agents across the environment
- ✓ Establish trust boundaries and kill-switch controls for autonomous operations
- ✓ Strengthen governance and accountability as agent adoption scales
- ✓ Reduce risk exposure without adding operational friction
- ✓ Extend existing security programs to address agentic AI



Stronger Together. Protecting What's Next.

Working with GuidePoint, your organization will be backed by our elite team of highly trained cybersecurity engineers, architects and consultants who come from organizations of all sizes, including Fortune 100 companies, the Department of War, and U.S. Intelligence Community.



Comprehensive Security Services for Agentic AI

GuidePoint helps organizations build the identity, governance, and operational foundation needed to securely adopt and scale autonomous AI agents securely — from strategy and assessment through implementation and ongoing management. Whether you're just getting started or looking to mature our existing program, our tailored services meet you where you are:

Strategy, Architecture, and Advisory – A full readiness assessment for AI agent adoption, spanning identity governance, trust boundaries, NHI alignment, and operational risk, paired with architectural guidance and a prioritized roadmap aligned to your security posture and business objectives.

Implementation and Deployment – Deploy identity lifecycle controls and AI security tooling across your environment, including discovery, classification provisioning, access governance, kill-switch mechanisms, and deprovisioning, integrated into existing architectures or built from the ground up.

AI Forward Deployed Engineering – Embedded engineering resources to deploy AI and agentic solutions, force-multiply adoption of security architectures and tooling, and accelerate operationalization of identity and governance controls.

Red Teaming and Adversarial Testing – Proactive offensive testing of AI agent deployments targeting the your unique attack surface introduced by agentic systems, including prompt injection, jailbreak techniques, tool-use abuse, privilege escalation, data exfiltration through agent channels, and multi-agent collusion scenarios, with actionable findings and risk-rated remediation guidance to help your team understand how adversaries will target your agents before they do.

Compliance and Regulatory Alignment – Map AI agent deployments against emerging regulatory frameworks, including the EU AI Act, NIST AI RMF, ISO 42001, and sector-specific requirements, with audit preparation, evidence collection, gap analysis, and ongoing compliance monitoring to stay ahead of a rapidly evolving regulatory landscape without slowing agent adoption.

Security Tooling and Evaluation – Structured comparison of AI security and identity governance solutions against your organization's functional requirements, integration constraints, and long-term strategic direction, including lab-based proof of value testing and best practice guidance to accelerate tool selection.

Managed Services and Operations – Ongoing monitoring, policy tuning, and platform operations that evolve alongside your agent population, business needs, and regulatory landscape, from lifecycle management through continuous governance enforcement.

AI CISO – Comprehensive security team enablement for agentic AI, including SOC agentification, AI-driven exposure management, and detection and remediation capabilities using vendor and custom-built tooling.

Incident Response and Forensics – Development and operationalization of AI-agent-specific incident response capabilities and playbooks, including containment procedures for compromised agents, forensic analysis of agent decision chains and tool invocations, evidence preservation workflows, and structured tabletop exercises simulating agent-related breaches. Services close the gap between traditional IR playbooks and the novel failure modes introduced by autonomous systems.

Supply Chain and Third-Party Agent Risk – DAssesses risk introduced by third-party agents, plugins, tools, and model providers across the environment. Covers model provenance validation, plugin and integration security assessments, agent marketplace ecosystem evaluation, and vendor risk framework development tailored to the unique trust relationships and data flows inherent in agentic architectures.

Data Protection and Privacy – Design and implement data loss prevention controls purpose-built for agent interactions, covering PII and PHI handling in prompts and responses, data classification for agent-accessible resources, prompt and response logging policies, and retention and redaction strategies. Addresses the novel data flow paths created by autonomous agents that traditional DLP architectures were not designed to govern.

Secure Agent Development Lifecycle – Embed security into the full agent development lifecycle with secure-by-design guidance for engineering teams building agentic systems. Covers secure coding patterns, guardrail implementation, output validation, CI/CD pipeline hardening for agent deployments, and pre-production safety testing across hallucination, toxicity, and tool-misuse boundaries.

Agent Observability and Behavioral Auditing – Design and deploy purpose-built observability infrastructure for AI agent environments, including structured logging of agent reasoning chains, tool invocations, data access patterns, and decision outputs across the full agent lifecycle. Implement real-time telemetry pipelines, behavioral baselining, anomaly detection, and audit-ready dashboards that provide security, compliance, and engineering teams with full traceability into what agents are doing, why, and with what impact. Covers integration with existing SIEM and SOAR platforms, custom alerting on policy-violating agent behaviors, and long-term retention architectures aligned to regulatory and forensic requirements.

Training and Organizational Enablement – Broad-based enablement programs spanning developer training on secure agent design, end-user awareness covering social engineering via agents and trust calibration, and executive and board-level AI risk literacy. Goes beyond security team enablement to ensure your entire organization understands how to build, use, and govern AI agents responsibly and with appropriate skepticism.

Responsible AI and Ethical Risk Assessment – Evaluate AI agent deployments against responsible AI principles, including fairness, bias, transparency, explainability, and accountability, mapped to your organization's values, industry norms, and emerging regulatory expectations such as the EU AI Act and NIST AI RMF. Deliver risk-tiered assessments of agent use cases, recommend guardrails and human-in-the-loop controls proportionate to impact, and establish governance frameworks that ensure ethical considerations are embedded into agent design, deployment, and ongoing operation rather than bolted on after the fact.

Why GuidePoint Security

When cyber threats evolve, organizations need confidence that every identity, human and non-human, is secure. GuidePoint Security combines cybersecurity expertise, identity strategy, and technology guidance to help organizations build resilient AI Identity programs. Our vendor-objective approach enables clients to assess AI risks, strengthen governance and lifecycle controls, and implement scalable strategies that reduce identity risk and secure the automation, applications, and AI capabilities that power the business.

About Us

GuidePoint Security brings together proven expertise, great relationships, and leading technologies to solve our client's most complex cybersecurity challenges. As a trusted cybersecurity advisor and partner, GuidePoint keeps people, data, and operations safe. We deliver tailored cybersecurity services that adapt to safeguard the nation's leading organizations today and provide complete confidence in their cybersecurity tomorrow. Stronger Together. Protecting What's Next.



1900 Reston Metro Plaza, Suite 701 • Reston, VA 20190
[guidepointsecurity.com](https://www.guidepointsecurity.com) • info@guidepointsecurity.com • (877) 889-0132
DS.AAI.2609