cyera

Solution Brief

Cyera Agent Guardian

Securing the Agentic Enterprise

Enterprises are required to deploy an autonomous workforce to keep up with the speed of business in the AI era. Security professionals are expected to find ways to enable the rollout of autonomous agents without compromising on security. But agentic adoption introduces new challenges, as agents operate independently, call application programming interfaces, execute system command lines, query structured databases, and write directly to production environments. This active execution occurs with minimal human oversight and often operates entirely outside of real-time security validation.

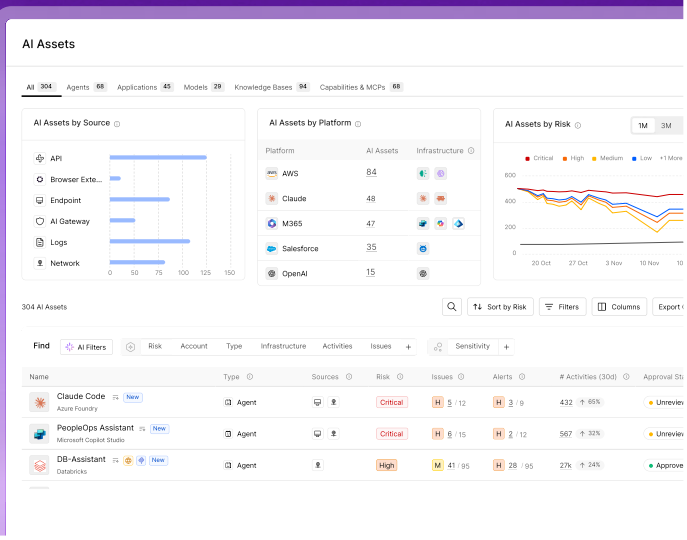

Cyera Agent Guardian establishes an active control layer across the enterprise agentic ecosystem. Discovering agent activity across platforms and locations, enforcing intent-based policy boundaries, assessing operational risk profiles, monitoring intermediate execution steps, and blocking destructive actions in real time. This unified solution allows organizations to safely deploy autonomous agents, enabling the AI adoption the business requires.

The Core Enterprise Challenges

Active Enablement Friction	High business demand for autonomous agents forces security teams to balance rapid innovation with strict risk and compliance controls.
Shadow Agent Proliferation	Decentralized creation across low-code tools, model suites, and dev workstations makes tracking agent inventory and access impossible.
Toxic Access & Capability Combinations	Unchecked convergence of privileged identities, execution capabilities, and broad data access risks database deletion, misconfigurations, or leaks.
Intent Drift	Agents take technically permitted actions outside their intended business function or fall into costly infinite tool-execution loops.
The Intermediate Gap	Standard logs capture only initial prompts and final responses, leaving intermediate tool calls, prompt translations, and DB queries unmonitored.
The Prevention & Remediation Void	Static blocking disrupts critical automation, while passive alerts fail to stop non-deterministic agent data loss in real time.

The Core Security Lifecycle Pillars

Securing the agent workforce requires a systematic operational framework that spans from initial discovery to real-time execution defense. Cyera structures this protection lifecycle into four sequential phases to establish complete visibility, continuous policy alignment, and automated safety verification.

Discover

Find and map all agents

Govern

Control your agentic enterprise

Protect

Real-time protection for agent actions

Validate

Continuous adversarial verification

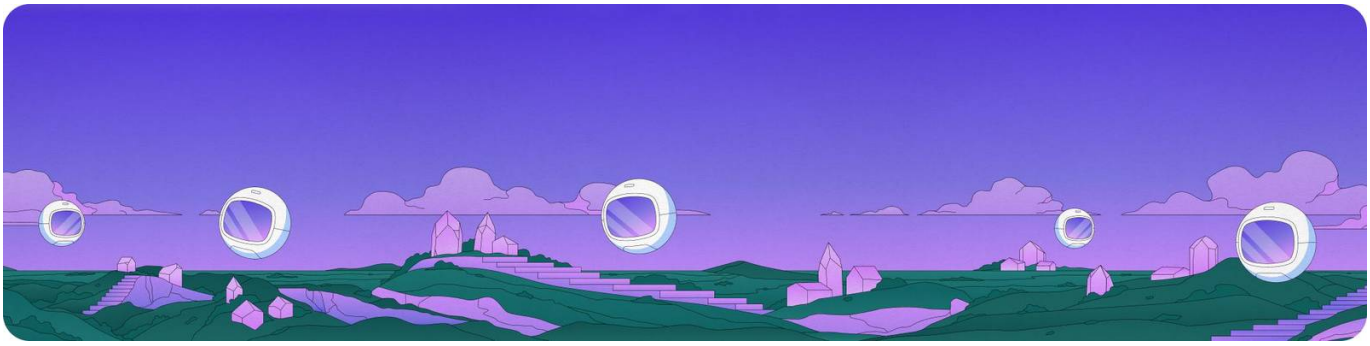

Discover: Find and map all agents

Cyera automatically inventories the corporate agentic footprint across all cloud environments, SaaS platforms, and workstations. The discovery engine maps connections to build a visual Agent Graph and an AI inventory, revealing what exists, what data it can access, and its operational scope.

- **Unified AI inventory:** Inventory all AI assets, including platform and custom agents, to provide visibility into the organizational agent environment.
- **End-to-end action visibility:** Monitor the full execution chain, prompts to tool calls, token monitoring, and track agent behavior and intent in real time.
- **Unified agent mapping:** Standardize fragmented architectures into a clear, visual map of every agent's technical lineage.

Govern: Control your agentic enterprise

Cyera Agent Guardian provides active risk and posture management, establishing clear behavioral boundaries for autonomous agents. By analyzing configurations, access rights, and intent, Cyera turns unmanaged agent deployments into observable, risk-managed business drivers.

- **Policy guardrails:** Define rules spanning both static posture configurations and live runtime execution flows to prevent unauthorized tool usage, enforce access controls, and block dangerous behavior.
- **Agent risk model:** Assess each agent's risk by weighing data sensitivity, existing security controls, and the agent's level of exposure.
- **Continuous drift monitoring:** Track configuration and behavioral drift between approved states and actual runtime states by comparing authorized author intent with live permissions and execution.

Protect: Real-time protection for agent actions

Cyera operates inline to evaluate user interactions, tool invocations, and data retrieval processes. The intermediate execution layer intercepts dangerous operations before they can propagate, allowing organizations to maintain full security control without degrading the user experience.

- **Human-in-the-loop enforcement:** Surface high-risk or sensitive actions for immediate human review and sign-off, ensuring operational safety without stalling critical automated workflows.
- **Inline action blocking:** Intercept and block destructive tool calls, prompt injections, and other unauthorized operations in real-time, before they can execute or cause downstream harm.
- **Agent kill switch:** Instantly isolate or disable compromised agents based on real-time risk levels, providing immediate control to revoke access and block tool routes when Cyera detects repeated attempts to perform dangerous actions.

Validate: Continuous adversarial verification

Cyera implements continuous security testing to ensure that production agents remain resilient against sophisticated exploitation. As AI systems are probabilistic by design, existing guardrails can lose efficacy even without an apparent change in configuration, making continuous validation essential to verify the effectiveness of protections.

- **Continuous red teaming:** Launch AI-driven attack scenarios against active agents to simulate advanced adversarial methodologies.
- **Automated policy validation:** Continuously verify that the governance, data protection, and execution policies established in earlier lifecycle stages are actively and correctly enforced at runtime.
- **Policy review audit trails:** Provide verified timestamping of the most recent validation run for each policy, delivering a continuous compliance review trail that shows security guardrails are actively tested and up to date.

Architecture and Deployment Control Points

Cyera offers multiple integration points to balance coverage depth, latency constraints, and implementation complexity. These controls can operate concurrently depending on your architecture, meeting our customers exactly where they are on their AI adoption journey.

Cloud Integrations	Direct integrations to managed agent builders, including Microsoft Copilot Studio, AWS Bedrock, and Azure AI Foundry, to analyze active configurations and monitor interactions with negligible operational impact.
AI Gateway Integration	A centralized inline plugin for AI gateways such as LiteLLM, Kong and Apigee, that intercepts all interactions between the agent and the LLM provider, enabling real-time content blocking, prompt injection defense, and active blocking.
Cyera Endpoint Agent	A local agent that prevents source code leaks and unauthorize d access. In ephemeral workspaces, MDM or EDR systems (e.g., CrowdStrike, SentinelOne) trigger scripts to scan local command-line tools, directories, and model context protocol servers.

Browser Shield	A lightweight browser extension that monitors and secures web-based AI, discovering AI usage, and blocking client-side leaks and injections in real time.
Agent Framework Hooks	Distributed through plugin marketplaces. Enforces inline protection for prompts, responses and tool calls through hooks.

Unified Deployment Scope

Local Agents - Claude Code, Claude Cowork

Platform Agents - Copilot Studio, Microsoft Foundry Agents, Bedrock Agents, Snowflake Cortex AI

Agent Platforms - Microsoft Foundry, AWS Bedrock

The Cyera Advantage: A Unified Platform for agent security

The Posture and Runtime Convergence	Cyera unifies static configuration auditing with dynamic execution defense. This single-platform approach combines data security posture management, DLP orchestration, identity understanding, and agent security into a single operational framework. Providing robust agent security management from both posture and real-time runtime perspectives.
Data-Centric Context	By integrating directly with Cyera's industry-leading data security posture management capabilities, Cyera understands the exact sensitivity of the data accessed by each agent. This deep context ensures high-precision alerting, separating routine, authorized queries from anomalous attempts to access or exfiltrate sensitive corporate data assets.
Centralized Data Protection Policies	Smart security policies are configured once and applied universally. Whether data is accessed by a human employee or an autonomous agent calling a database API, identical data protection, compliance, and privacy rules govern the interaction. This unified approach eliminates policy fragmentation and ensures consistent risk management.

Summary and Next Steps

The rapid expansion of autonomous agents introduces unprecedented operational and data protection challenges to the enterprise. Cyera Agent Guardian combines comprehensive discovery, intent-based governance, real-time runtime intercept, and continuous adversarial red teaming to enable the business to deploy agents with confidence.

For more information visit cyera.com or book a demo

Book a demo →