


Cyera Endpoint for Agent Security

Secure On-Device Agentic Work Grounded in Data and Identity Context



The Agentic Security Paradigm

Employees are using AI agents to execute more of their tasks, freeing them to focus on high-value work that requires deep expertise. Teams are building multi-step agentic workflows on employee endpoints through native desktop AI applications, browser workflows, command-line tools, and developer environments such as Claude Code and Codex. These agents operate autonomously on the endpoint, often with broad employee permissions and credentials. That autonomy and inherited access mean every new agent, tool, or connection becomes another potential path for data to leave the device - often outside the visibility of traditional security controls, including network proxies that operate downstream from the application runtime. The result is shadow AI: a steady, largely invisible flow of sensitive corporate data and secrets into unapproved AI tools and external services - leaving security teams unable to answer the basic question of what employees are actually doing.

This creates four problems:

Shadow Agent Sprawl

Employees routinely deploy local AI desktop apps, coding assistants, and unvetted MCP servers on workstations, bypassing central administrative oversight and security policies.

Inherited Privileged Access

Autonomous agents inherit employees' broad entitlements and credentials, enabling them to query confidential local files, run terminal scripts, and access enterprise data stores without continuous re-authorization.

Intermediate Execution Gap

The highest-risk activity may happen before an agent produces a final response: reading a file, assembling context, executing code, calling a tool, and changing local or remote systems. Controls positioned only at the browser or gateway see fragments of this chain rather than the on-device activity that produced the request.

Agent Execution Drift

Agents can act unpredictably because of their probabilistic nature. Without strict guardrails, an agent might deviate from its intended business goal, potentially running unauthorized commands, altering critical system settings, or wasting resources in infinite loops.

Agent Security Controls Offered by the Cyera Platform

Cyera secures the agentic ecosystem through multiple control points: agent and framework hooks, an AI gateway plugin, cloud/AI platform integrations, Cyera Browser Shield, Cyera Endpoint, and Cyera Endpoint Analyzer. Each feeds continuous telemetry and context to the Cyera platform's policy engine which is grounded in shared data classification and identity context - so every control point acts on the same understanding of what's sensitive and who, or what, is acting on it. Agent Guardian supports modular deployment across one or more control points. Each organization will use a different set of control points based on its tech stack, required use cases, and maturity of their agentic program.

Introducing Cyera Endpoint

Cyera Endpoint serves as the primary on-device control point for Cyera Agent Guardian. It provides runtime visibility and protection to secure agentic applications, such as Claude Code and Cowork, as they operate autonomously on employee and developer machines. Cyera Endpoint continuously evaluates a local agent's actions while it is happening - calling tools, reading files, running shell commands - against its intended goals and organizational policy. It can allow, block, or coach users on risky agent behavior such as unauthorized data movement and secret credential access. Cyera combines application runtime context and controls with the Cyera platform's data and identity intelligence, helping security teams discover local agents, govern human and non-human activity, and enforce controls in real time before data moves.

Core Capabilities

Agent Discovery

AI Asset and Agent Discovery: Automatically inventories local AI assets and agents on enterprise workstations, mapping installed AI desktop applications, tools, models, and MCP servers, along with all their configurations, to provide unified visual lineage within the Cyera platform.

Runtime Visibility and Enforcement

Operates inline within local execution paths to intercept file reads, script executions, and tool invocations. Intervenes and gates which agent actions are allowed before execution. Inline block/alert/coach enforcement is executed locally on the device as each agent action is invoked.

File System Monitoring

Continuously inspects raw file reads and writes on local drives, evaluating file contents against the data classification engine of the Cyera platform to detect sensitive data exposure before prompts transmit across network boundaries.

Real-Time User Coaching

Delivers contextual coaching and data sensitivity guidance to users at the moment of risky agent interactions, reducing friction for legitimate business use while intercepting critical policy violations.



Core Use Cases

Safe Agentic Workflows and Rollouts

Monitors local coding and workspace agent activity and enforces runtime security guardrails for teams, detecting prompt injection, secret leakage and policy violations within development and business workflows.

Shadow AI Usage

Continuously discovers and inventories AI assets running on endpoints - local AI agents, CLI tools, MCPs, and browser extensions - mapping each to the device, identity, and data.

Personal-Account Control

Enforces corporate data protection policy regardless of whether an employee is logged into a corporate or personal AI account, without relying on corporate SSO or sanctioned browser session.

AI Application Data Leakage Prevention

Prevents source code, proprietary files, secrets, or credentials from leaving the device through prompts, file uploads, or tool interactions that bypass network-layer DLP.

Cyera Endpoint Architecture

Cyera Endpoint runs as a lightweight software application on employee Windows and macOS devices. It is managed through the Cyera platform and distributed through standard enterprise software-management tools, such as Microsoft Intune and Jamf. Its on-device interception layer observes application and operating-system activity, including native AI desktop interactions, local-agent actions, and file operations. Each action is evaluated in real time and enforced inline - before the agent's action or data movement completes. Securing agentic workflows takes more than one control. None goes deeper than Cyera Endpoint at the point of user interaction.

The Cyera Advantage: Data and Identity Context

Point security controls often evaluate isolated signals - such as a process, prompt, tool call, or file transfer - without the broader context needed to understand the actual risk. Cyera Endpoint combines on-device visibility with Cyera's data classification, identity, access, and business context, helping security teams understand what data is involved, who or what is acting on it, and whether the action should be allowed. The result is more precise protection for local-agent workflows with less unnecessary disruption.

Data and Identity Understanding and Context

Deep data classification evaluates file sensitivity while identity context verifies exact user entitlements during local agent interactions.

Unified Policy Engine

A single administration console within the Cyera platform applies consistent data protection rules whether data is accessed by human employees or autonomous agents.

Contextual Execution Control

Real-time evaluation stops destructive intermediate actions and unauthorized tool invocations before terminal commands affect production environments.



The Agent Security Lifecycle

Discover

Automatically inventories local personal and enterprise agent footprints across workstations, mapping installed tools, model files, and MCP servers to provide unified visual lineage within the Cyera platform.

Govern

Manages operational risk by mapping every local agent to operator identity entitlements and corporate access policies, continuously auditing configurations to uncover over-privileged access and unauthorized drift.

Protect


Operates inline within local execution paths to intercept file reads, script executions, and tool invocations, dynamically redacting sensitive content and isolating compromised sessions using an agent kill switch.

Validate

Executes validation scenarios, including simulated prompt injections and privilege escalations, to continuously verify that local guardrails remain resilient against adversarial manipulation.

Summary

Local agents can accelerate development and knowledge work, but organizations need visibility and control where those agents operate. Cyera Endpoint gives security teams an on-device control point grounded in Cyera's data and identity intelligence, helping them discover local AI, apply policy to agent actions, and protect sensitive data at the point of use.



Explore how Cyera Endpoint secures the AI apps and agents running on your devices.

Book a demo

