AppViewX

Agent Identity Security | Solution Overview

AI-Native Identity Security for Every Agent at Enterprise Scale

Challenge

Your workforce is no longer just human, but your identity security still treats it as if it is. Existing Identity and Access Management (IAM) tools were built for people, with no way to discover, govern, or secure what agents do.

Shadow agents proliferate without oversight

Developers install coding assistants. Business teams spin up agents inside SaaS platforms. Sanctioned deployments spawn sub-agents no one provisioned, and none of them appear in any inventory.

Agents are over-permissioned and under-governed

IAM tools apply the human access model to agents, with no insight into what an agent can reach once it gets there. Agents carry shared credentials and hold standing access they were never meant to use.

Accountability gaps make every audit a scramble

Even when an agent is known, context is thin: who owns it, what it touches, and whether its behavior maps to any compliance framework. Every audit is a scramble. Every incident is an investigation from scratch.

Rogue behavior is detected after damage is done

When an agent misuses its access, security teams find out after the fact. There is no runtime control, nothing standing between an agent and the action it takes.

70%
of organizations give AI systems more access than they'd give a human doing the same job[1].

97%
had an identity-related security incident last year, 70% of those tied to AI activity[2].

68%
of security teams can't distinguish a legitimate AI agent from a malicious actor posing as one[3].

AppViewX Agent Identity Security

Secure the Identity Layer for AI Agents

AppViewX Agent Identity Security is purpose-built at the identity layer for AI agents. It discovers every agent, sanctioned or shadow, so nothing operates outside your line of sight. It governs agent privileges and posture, maps them to compliance frameworks, and keeps audit-ready activity logs, so exposure shrinks and audits stop being a scramble. And it secures every agent action at runtime, stopping risky actions before they execute.

Discovery

See every agent in your environment within minutes, sanctioned or shadow. Know what each one is built from, who owns it, and what it can reach

Governance

Keep every agent assigned to an owner, least privileged, and audit-ready with AI spend held inside of limits that you set.

Runtime Security

Every agent action is checked before it executes. Risky actions are blocked and any agent can be stopped mid-session.

Source: [1] 2026 Infrastructure Identity Survey (NHIMG), [2] GitGuardian, "Non-human identity sprawl and AI agents in enterprise security," 2026
[3] CSA "Identity and Access Gaps in the Age of Autonomous AI"

What It Delivers

Discovery, Governance, and Runtime Security in One Platform

Multi-Method Discovery

Find every agent through API integrations with EDR, SaaS, Cloud platforms, endpoint software, an AI coding assistant plugin, and an SDK for custom agents.

Lifecycle and Cost Governance

No-code workflows register new agents, assign owners, track posture drift, and revoke credentials at decommission. Set token usage thresholds to control spending.

Secure MCP Gateway

Discover every MCP server and assess their risk. Agents get just-in-time access to servers and tools. Sensitive data is redacted in both directions with built-in controls or Microsoft Purview.

AIBOM and Skills Discovery

Each agent's AIBOM lists its models, MCP servers, credentials, identities, and skills and lives in an access graph showing how they connect.

Continuous Compliance

Demonstrate alignment with OWASP Top 10 for LLM Agentic Applications and Agentic Skills, MITRE ATLAS, GDPR, HIPAA, ISO 27001/42001, and NIST SP 800-53 Rev. 5, alongside NIST AI RMF, the EU AI Act, SOC 2, and SEC Cyber Disclosure.

Agent Kill Switch

Terminate an agent and all its active sessions by runtime policy, based on resource, agent, intent, or automatically when a no-code workflow fires on a detected change or signal from another security tool.

What Security Leaders Get

Accelerate AI adoption without scaling risk	Approve and scale AI initiatives instead of blocking them, with visibility, governance, and runtime controls that work alongside the IAM, PAM, and SIEM tools already in place.
Know every agent in your environment	Give leadership a complete current inventory of every agent, including the ones no one approved, and keep it current as new agents appear.
Be ready before the auditor asks	Prove control of your AI agent estate to regulators, auditors, and cyber insurance carriers with reports pulled from the platform, not evidence assembled through manual projects.
Stop incidents instead of investigating them	Contain a rogue agent while it's acting, not after the damage is done. Risky actions that are blocked never become incidents to investigate or disclose.

"AI agents are a new identity group, and their rapid proliferation creates new risks for the enterprise. To reduce risk, we must monitor, audit, and control their privileged access to sensitive data and systems. We must ensure their lifecycle is automated and governed. I've been collaborating with the AppViewX team, and their innovative Agent Identity Security solution is the right response for reducing the risk of agents in the enterprise."

John Barrow, JB Poindexter & Co. Chief Information Security Officer

Get a complete picture of your AI agent environment.

Schedule a discovery assessment, get a full inventory of your AI agents, their credentials, and their access.

AppViewX
info@appviewx.com
www.appviewx.com
+1 (212) 400-7541