




Agent Identity Security

See Every Agent. Stop Any Agent at Runtime.

AppViewX Agent Identity Security is purpose-built for the agent identity layer. It discovers every agent, whether sanctioned or shadow; governs their posture and maintains audit-ready activity logs; and monitors and controls, in real time, every action that the agent performs.

Traditional identity security is built for humans and assumes access is predictable, auditable, and owned.

AI agents break every one of those assumptions.

Sanctioned and shadow agents proliferate without centralized oversight as developers install coding assistants or business teams build them inside SaaS platforms. Each one brings models, credentials, privileges, skills, and connections to enterprise systems, and acts autonomously or on behalf of humans at a pace traditional access controls were never built to handle.

Agents also create other agents. A single sanctioned deployment spawns sub-agents and delegated tasks that were never provisioned, never reviewed, and never appear in any inventory.

Most enterprises lack a complete inventory of what's running, and the agents they can identify carry limited context: who owns them, what systems they access, and whether their behavior meets compliance requirements.

Existing identity and access tools apply the human model to agents, with no insight into what an agent can reach or what it does once it gets there. Agents act autonomously, carry shared credentials, and take actions no policy anticipated, even while operating within their assigned permissions. That creates a governance and enforcement gap at the identity layer.

Agent Identity is the Fastest Growing and Least Governed Attack Surface

The governance gap

92% agree governing AI agents is critical, but only **44%** have implemented any policy to do it.


AI access is over-permissioned

70% of organizations grant AI systems more access than they would give a human doing the same job.

Identity-related incidents are widespread

97% of organizations experienced an identity-related security incident last year; **70%** were tied to AI activity. **68%** of security teams say they can't distinguish a legitimate AI agent from a malicious actor posing as one.

AI-native identity security for every agent, at enterprise scale.



Discovery

Discover every agent, sanctioned or shadow, and map each to its users, credentials, endpoints, skills, MCP servers, and models in an AI Bill of Materials (AIBOM)

Governance

Map agent posture to compliance frameworks, maintain audit-ready logs, and review activity timelines in the live access graph

Runtime Security

Eliminate standing privileges with just-in-time access and terminate any agent with the Agent Kill Switch

What Agent Identity Security Delivers

Discovery

Get deep, contextual visibility across every agent in minutes

What it delivers

- Unified registry of every agent, sanctioned or shadow: coding, enterprise productivity, endpoint, SaaS, and custom-built
- Shadow AI discovery, including browser-based and short-lived, script-based agents
- Lightweight endpoint Guardian Agent plus API integrations across EDR, SaaS, and cloud platforms, so no shadow agent goes undetected
- AIBOM covering models, MCP tools, credentials, identities, and the skills agents can access, with unsafe skills flagged
- Discovery and continuous risk assessment of sanctioned and shadow MCP servers
- Access graph with real-time risk signals and timeline playback

Governance

Govern agent risk and prove compliance with audit-ready logs.

What it delivers

- Customizable enterprise AI risk profiles
- Agent lifecycle management workflows
- Compliance violation reporting, automated remediation
- Audit-ready activity logs
- AI token usage, trends, and costs, with threshold policies to prevent runaway spending
- Govern every MCP server, including shadow servers, through the MCP Gateway
- Assess and demonstrate alignment with OWASP Top 10 for Agentic Applications and Agentic Skills, MITRE ATLAS, GDPR, HIPAA, ISO 27001/42001, NIST SP 800-53 Rev. 5, NIST AI RMF, EU AI Act, SOC 2, and SEC Cyber Disclosure

Runtime Security

Stop agents from privilege misuse at runtime

What it delivers

- Agent Kill Switch: terminate any agent and its active sessions, sanctioned or shadow, automatically through event-driven workflows and Shared Signals Framework (SSF) signals, through runtime policies, or on demand
- Just-in-time access eliminates standing privileges, granting agents access to MCP servers and tools only when and for as long as needed, based on identity and agent context
- Intent-based access policies
- Sensitive data redaction with built-in controls or existing data security tools such as Microsoft Purview
- MCP Gateway with bi-directional policy enforcement
- Dynamic least agency enforcement: risk-scored actions evaluated against resource sensitivity in real time
- Behavior-based alerts
- Automated containment with forensic audit trail

What Security Leaders Get

Results	Business value
Accelerate AI adoption without scaling risk	Approve and scale AI initiatives instead of blocking them, using visibility, governance, and enforcement capabilities that help the organization remain secure.
Stop incidents instead of investigating them	Block risky actions before they execute and stop rogue agents with the Agent Kill Switch, helping teams avoid incident reports and reducing the business impact of disclosures.
Be ready before the auditor asks	Prove control to regulators, auditors, and cyber insurance carriers with reports pulled from the platform instead of evidence assembled through manual projects.
One solution for full agent lifecycle governance	Manage agent identity, access, posture, and threat response with a single platform that integrates with your existing security and IT ecosystem, including DLP, SIEM, ITDR, firewall, IdP, MDM, and EDR tools, SOC workflows, and collaboration platforms like Slack and Teams.

Where Agent Identity Security Applies

Secure enterprise productivity agents
AI assistants act on behalf of employees using delegated credentials, but without the behavioral baselines or access reviews applied to the humans they represent. Agent Identity Security inventories every AI agent and monitors each assistant's activity to enforce least agency at the point of action.

Secure coding agents
Coding agents operate with developer-level access to source code, build systems, and production environments. Agent Identity Security stops risky agentic actions while the code is being written, not just after it ships.

Govern SaaS agents
Business teams spin up agents inside SaaS platforms with inherited permissions that security never granted or reviewed. Agent Identity Security gives every agent an owner and governed lifecycle, so an agent never acts without an owner and a policy.

Secure endpoint agents
Shadow agents run in browsers and as short-lived scripts that never appear in a central inventory. A lightweight endpoint Guardian Agent discovers them, so they can be governed and stopped at runtime like any other agent.

Secure custom-built agents
Teams build their own agents with their own models, MCP tools, credentials, and skills. Agent Identity Security maps each one in the AIBOM, governs its posture, and controls its actions at runtime.

Why AppViewX

Purpose-built for agents.
Agent Identity Security delivers the industry's most comprehensive discovery, governance, and runtime security for every agent type.

No shadow agent goes undetected or unchecked. Alternative solutions rely on a single detection method. AppViewX pairs a lightweight endpoint Guardian Agent with API integrations across EDR, SaaS, and cloud platforms. The Agent Kill Switch covers sanctioned and shadow agents alike, with no dependence on an MCP gateway.

Trust that holds as cryptography evolves. OAuth governs what an agent may do on a user's behalf, but the agent must still prove who it is. AppViewX extends its core PKI and CLM platform to issue cryptographically verifiable, quantum-resilient agent identities, adding a layer of trust when agents authenticate through your enterprise IdP. Every identity chains to your AppViewX-managed PKI, so security teams get one trusted root, a tamper-evident audit trail, and a single place to build them.

Get a complete picture of your AI agent environment.

Talk to an AppViewX expert about Agent Identity Security.



165 Broadway, 23rd Floor,
New York, NY 10006

info@appviewx.com
www.appviewx.com

+1 (212) 400 7541
+1 (929) 955 0055