




Agent Identity Security

Defend Your Enterprise from Rogue Agents

AppViewX Agent Identity Security addresses the risk associated with AI agents. It sees, governs, and stops rogue agents at the identity layer, so security teams get runtime preventions with deep contextual visibility across every agent, not just detection.

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

AI agents break every one of those assumptions.

Known and unknown agents proliferate without centralized oversight as developers install coding assistants or business teams build them inside SaaS platforms. 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.

Built Around the Agent's Identity



What Agent Identity Security Delivers

Discovery	Governance	Runtime Security
Know what's running before it becomes a problem.	Govern agent risk and prove compliance.	Stop rogue behavior the moment it happens.
What it delivers - Unified registry of all agents — sanctioned and unsanctioned - Shadow AI detection, monitoring, and kill switch control - Full AI-BOM — MCPs, models, skills, users, credentials - Agent and MCP risk posture assessment - Access graph with real-time risk signals and timeline playback	What it delivers - Customizable enterprise AI risk profiles - Agent lifecycle management workflows - Compliance violation reporting, automated remediation - Audit-ready activity logs - Token usage insights and cost governance policies - Govern MCPs, models, skills, users and credentials - Continuous posture mapping to SOC 2, NIST AI RMF, EU AI Act, ISO 42001, OWASP Top 10 for LLM, MITRE ATLAS	What it delivers - Just-in-time privilege control for any agent action - Intent-based access policies - Enterprise DLP integrations support for run-time data security - MCP gateway with bi-directional policy enforcement - Dynamic least agency enforcement: risk-scored actions evaluated against resource sensitivity in real time - Event-triggered response workflows and behavior-based alerts - Risk-triggered agent kill switch - Integrated ITDR; 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 rogue agents mid-action with enforcement that decides before the action executes, 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 runs alongside the IAM, PAM, and SIEM tools already in place.

Where Agent Identity Security Applies

Secure enterprise use of AI assistants

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 in the development lifecycle

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 agents built by business teams in SaaS platforms

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.

Why AppViewX

Agent security is an identity problem.

Every agent acts through its identity, which makes the identity layer the chokepoint for agent-borne risk. AppViewX started with the identity layer and built outward. Agent Identity Security enforces controls by identifying every agent and requiring every agent to pass through it.

Agent access is controlled before and during the act. Agent Identity Security provides visibility into each agent's ownership, security posture, access, and activity. It uses that context to enforce governance, least privilege, least agency, and policy-based authorization both before an agent acts and at the moment of action.

Get a complete picture of your AI agent environment.

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



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

info@appviewx.com
www.appviewx.com

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