APONO

Datasheet

Securing Autonomous Agents with Intent-Driven Control

The Challenge of Deploying Agentic AI Securely

AI agents are already deployed across your organization, delivering real engineering velocity — but legacy privilege models can't handle the risks:

- Vulnerable to hallucinations, manipulation, and overreach
- Copilots and coding agents inherit their owner's credentials and permissions
- Standing privileges let agents carry out destructive actions at machine speed before anyone notices

At the center of the challenge are standing privileges. Without dynamic, Just-in-Time controls, organizations are forced to choose between the velocity benefits of autonomous agents and the security their environments demand.

Key Benefits

- Enable Secure Agent Productivity
- Fully Audited Across Environments
- Continuously Enforce Least Privilege

Introducing Apono Agent Privilege Guard

Apono is the partner enterprises need to unlock the full value of AI agents in their most sensitive environments. By understanding each agent's intent, Apono enforces dynamic guardrails that free agents to do more where it is safe and bring humans in only when it matters.

Every privilege is created at runtime, scoped to the task, and immediately eliminated on completion, returning every environment to Zero Standing Privileges by default.

Agent Privilege Guard continuously assesses and monitors every privilege request, adapting as risk and context change and validating that controls are enforced at every step. When an agent deviates from acceptable intent, the guardrail responds in real time, so the freedom to move fast never comes at the cost of security.

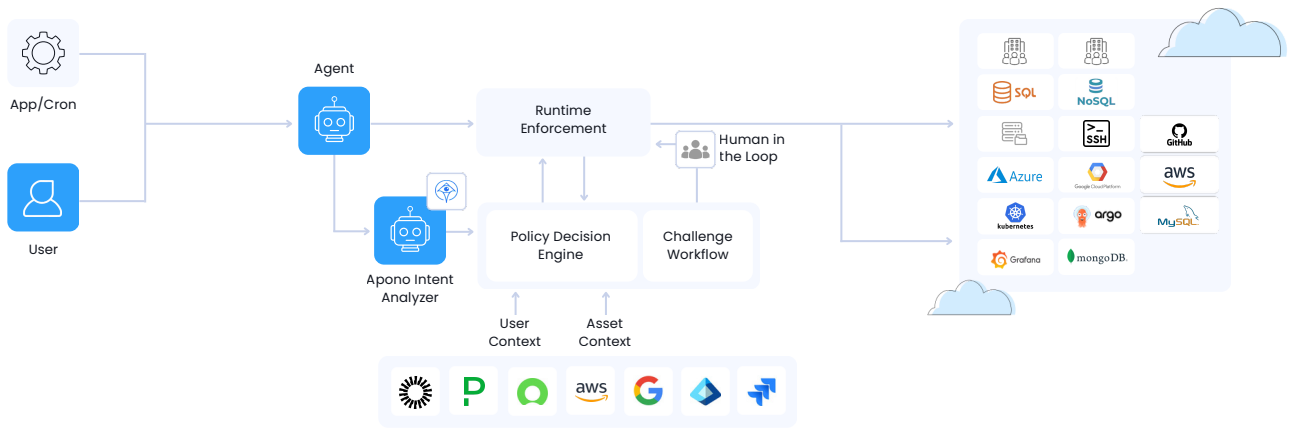

Placing Intent and Risk at the Center of the New Guardrails Approach

In the fast-paced world of Agentic AI, intent and risk are our primary indicators of if our agents are performing as expected or straying outside of the bounds of safe behavior. Apono's approach aims to strike the balance between operational needs and risks.

Apono assesses what an agent is attempting to do at runtime, making rapid assessments of risk that cause minimal disruptions to agent productivity. This process is based on a nuanced understanding of when a human is actually required to be brought in to approve an action, significantly enhancing velocity while simultaneously ensuring that security is never compromised.

The Agent Access Cycle

1

Declare Intent

The agent explicitly states **what** it intends to do and **why.**

2

Evaluate Context & Risk

Real-time assessment of environmental factors, data sensitivity, timing, and behavioral patterns.

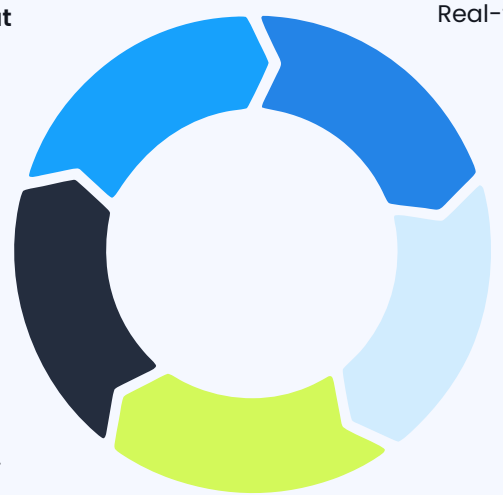

3

Grant Temporary Authority

Least privilege access is granted Just-in-Time, for a specific action, and expires immediately thereafter.

4

Enforce During Execution

Continuous monitoring and control of the agent's actions as they unfold, adapting to any deviation from intent.

5

Revoke & Learn

Access is immediately revoked post-action, with insights feeding on adaptive trust model for future decisions.

About Apono

Founded in 2022 by Rom Carmel and Ofir Stein, Apono delivers a Cloud Privileged Access Platform purpose-built for modern, fast-moving environments. With support for both human and non-human identities, Apono helps security and DevOps teams enforce least privilege at scale without slowing down delivery.

Trusted by Fortune 500 enterprise and recognized in Gartner's Magic Quadrant for Privileged Access Management two years running, Apono is shaping the future of secure cloud access.

APONO

www.apono.io

@Apono_official

@Apono