(Tuent) / Sentinel

RUNTIME BEHAVIORAL SECURITY
FOR AI AGENTS

A control layer between identity and execution.

Enterprises handed Claude Code and Cursor to hundreds of thousands of developers. Those agents read files, call APIs, touch databases. They execute, and they do not push back. Sentinel checks every tool call, file read, and intent **in-process, before it runs, deterministically, with no model in the monitoring path.** The question is no longer who has access. It is whether this actor should be allowed to do this, here, now.

TWO HOOK CHECKPOINTS · FOUR DETERMINISTIC DEFENSE LAYERS

PRE_EXECUTION › ROLE BOUNDARY | INTENT ALIGNMENT | BEHAVIORAL BASELINE | HOOK PREVENTION › POST_EXECUTION

On-prem. BYOK. Zero external model calls in the enforcement path. The thing protecting the agent does not inherit the prompt-injection surface of the agent it watches.

FOUR FAULT LINES IN AGENT SECURITY

DETERMINISTIC VS PROBABILISTIC

Deterministic. Same input, same verdict.

No model sits in the enforcement path, so nothing can be talked into a yes. An optional model tier recommends tightening only, never loosening, and never sees credentials.

ENDPOINT VS NETWORK

Neither. In-process.

Endpoint watches processes, network watches traffic. Both are blind to which file an agent read and whether it matched intent. Cisco AI Defense and Prisma AIRS sit at a proxy. Sentinel runs where the action happens.

INJECTION RESISTANCE

No model to fool.

A classifier defending an agent inherits the agent's own prompt-injection surface. Deterministic checks cannot be socially engineered into a yes.

DEFENSIBLE MOAT

Learned policy on proprietary behavioral data.

Policy recommends itself from your approvals and denials, paired with a two-persona operations console and a signed audit trail. The data compounds. Static rules do not.

In-process per-action enforcement is the least consolidated and least funded layer in the market. *We are building the layer the incumbents will want to bundle, not the one they already ship.*

TRACTION, PRESENT STATE

~1,200	**npm installs** of @tuent/sentinel. First real adoption signal, hardening in progress.
0.3.0	**Shipped Jun 28.** Signed audit trail live, false-positive recovery loop reachable end to end.
Weekly	**PANW advisor sessions** with a Sr. Director of AI Security PM, plus a 2.5 hour founding whiteboard. Technical validation, not a pilot.
2	**Design partnerships being fortified:** Omniloy (health AI), Janea Systems (ex-Microsoft). Not signed pilots yet.
Finalist	**BVA Board of Regents** interview complete.
2nd place	**Santa Clara pitch competition.**
Completed	**Bronco Ventures Prep School**, a founder fundraising accelerator.

TEAM

James Cunningham · Charlie Greenberg

Best friends since high school. Second venture together. Roles swapped. James on product, GTM, and the operator console. Charlie on the core enforcement engine. We have led a fundraise before. Not a bet on execution, a bet on the second rep.

ACCURACY NOTES · CLAUDE CODE IS THE ONLY LIVE ADAPTER. CROSS-VENDOR CORRELATION IS THE NEXT BUILD, NOT A PRESENT CLAIM.

"The thing protecting the agent cannot inherit the failure modes of the agent itself."

TUENT.AI
JAMES CUNNINGHAM
CO-FOUNDER