Your security stack cannot see them.
Your AI agents are running.

Every enterprise is deploying autonomous agents with privileged access to APIs, databases, and infrastructure. EDR sees processes. IAM sees users. Neither sees what an agent actually does — the tool calls, credential access, and decisions made at runtime. That gap is now your largest unmonitored attack surface.

The Blind Spot

Unmonitored privileged assets

Agents operate with high-privilege access — API keys, database write permissions, real operational power. Enterprises have 5× more agents running than they realize. Shadow AI is the new shadow IT, except agents execute code and access data.

Inherently vulnerable to logic manipulation

LLMs cannot distinguish between data and instructions. A malicious payload inside a CRM record, email, or RAG entry becomes a command. No malware. No signature to scan. The attack is the logic itself — what we call Agentjacking.

How Agentjacking Works?

1. **ATTACKER PLANTS PAYLOAD** — In email, CRM record, or RAG entry
2. **AGENT INGESTS DATA** — Payload treated as a legitimate instruction
3. **LOGIC HIJACKED** — Agent executes attacker's workflow autonomously
4. **RCE / EXFLIRTATION** — Credentials stolen, lateral movement, data breach
5. **WAF & EDR BLIND** — Traffic looks clean - authenticated agent, valid API calls

An enterprise can spend tens of millions on perimeter defense, yet a single ingested email can trigger a silent Remote Code Execution - leaving security teams without the telemetry to know it happened.

The Tenet Solution

Discovery

Automatically maps every agent running in your environment - including shadow AI deployed without security review. Full visibility before you can defend.

Detection & Response

A runtime endpoint sensor sits inside the agent process, correlating LLM calls, tool executions, syscalls, and filesystem operations in real time. Behavioral deviations are flagged with full session context - not noise.

Kill Switch

When an agent is about to exfiltrate credentials or move laterally, we simulate the outcome — then sever access in milliseconds with full justification. Prevention before it happens, not detection after.

24M+	1M+	1,000+	<30ms
Users protected in production	Agentic sessions analyzed	Unique agent profiles	p95 latency - invisible to developers

Cisco · Superior Court of California, County of Los Angeles · SHIFT4 · PatientPoint · gopuff

Integrates With

aws · splunk>

Maps to OWASP Top-10 for Agents · MITRE ATLAS · NIST AI RMF · EU AI Act

Built by Ex-Cisco AI Defense, Noname, SentinelOne & DeepMind Experts

See the kill switch in action

See a live attack demonstration.

Discover which of your agents are exposed

Contact us at: tenetsecurity.ai | hello@tenetsecurity.ai | +972-542613060