C1 SOLUTION BRIEF

C1 for AI Agent Security

Secure every agent. Govern every action.

TRUSTED BY IT & SECURITY

ramp · instacart · zscaler · DigitalOcean · Brex · klaviyo · qualtrics

AI agents created a new attack surface security teams haven't had to defend.

AI agents are already reaching into enterprise systems. A coding agent queries a Salesforce record, an assistant in Claude or Cursor pulls from an internal API, and a growing share run on their own with no person in the loop. Each one acts with credentials and calls tools, so every agent is a new identity to attack and a new way into the systems behind it. Fewer than 25% of organizations have documented policies for creating or removing AI identities,[1] so most agent credentials are issued fast, shared widely, and never retired.

The agent surface is different: a poisoned tool can hijack what an agent does, a scraped credential grants silent access, and a confused-deputy attack lets an agent reach data its owner never could. OWASP now ranks identity and privilege abuse among the top risks for agentic applications, and prescribes short-lived, task-scoped credentials as the defense.[2]

What C1 delivers

Adopt agents without adding risk

Every agent is governed the moment it appears, so the business can move on AI without opening new holes.

A stolen credential leads nowhere

Agent secrets are useless the instant they leave the agent, so theft stops turning into a breach.

Contain an attack in seconds

A compromised agent triggers its own alert and automatic containment, before the damage spreads.

C1 governs AI agents with the controls that already govern people.

C1 brings AI agents under the same governance that already protects the workforce, so security doesn't need a new playbook or a separate tool to trust them. Every agent gets an owner and a scope, its credentials stay vaulted and short-lived, and every tool call is checked against policy before it runs. An agent becomes one more identity the platform governs, held to the same standard as the people beside it.

Prevention comes first. Credentials are bound to the agent that holds them, so a stolen key is worthless off-host, and tool-call responses are filtered for sensitive data before it reaches the model. Because no defense is perfect, C1 also plants decoy credentials that turn an attacker's first use into an instant alert, and revokes a compromised agent's access in real time.

Security controls built for AI agents.

Every agent has someone accountable for it.

Each agent gets an owner, a scope, and a lifecycle, governed alongside people and service accounts. None runs in the dark, none outlives its purpose, and a name and history are always attached.

A stolen agent credential opens nothing.

Credentials stay vaulted, short-lived, and scoped to one task, and DPoP binds each token to the agent's own cryptographic key. A secret lifted from a config file or an agent's context is useless to whoever takes it.

Any MCP server becomes safe to use.

Every MCP server an agent touches sits behind C1's identity-aware gateway, which authenticates each call, enforces what the agent can do, and strips PII, secrets, and payment data from responses. A tool no one reviewed can't poison the agent or act as a confused deputy.

An attacker's first move becomes the first alert.

Decoy credentials planted in MCP configs and agent prompts turn a real intrusion into an instant signal: the first use fires a critical alert and auto-contains the source, rotating nearby secrets, suspending the identity, and quarantining its requests.

Prove what every agent did, and shut down what's compromised.

Every agent action is logged with full identity context and flows to the SIEM and SOAR already in place, as auditable as a person's. When a compromise signal arrives, C1 revokes the agent's access in real time.

Agents don't need a new security stack. They need the one that governs everyone else.

AI agents are the fastest-growing identities in the enterprise and the least understood. The teams that stay ahead won't bolt on a separate tool for agent threats. They'll govern agents inside the identity platform that already secures their people, where every agent has an owner, every credential is short-lived and bound to its holder, every tool call is checked, and every compromise is caught and revoked. C1 makes the agent the safest identity in the environment instead of the weakest.

GET STARTED

Talk to a product expert

Book a demo
c1.ai/request-demo

zscaler

"One of the key differentiators we found with C1 was their security-first approach to solving access."

Dheeraj Malik Director of Corporate Applications, Zscaler

About C1

C1

C1 empowers organizations to adopt AI securely and at speed by delivering the right access and context to every human, workload, and agent. Companies like Instacart, Ramp, Zscaler, and Brex trust C1 to accelerate AI adoption with confidence. Learn more at c1.ai.

Scale access at the speed of AI

[1] Cloud Security Alliance, State of NHI and AI Security Survey Report. [2] OWASP GenAI Security Project, Top 10 for Agentic Applications, December 2025.