C1 SOLUTION BRIEF

C1 for Agent Runtime Governance

Each agent scoped to its task. Every call checked. Every action owned.

TRUSTED BY IT & SECURITY | ramp | instacart | zscaler | DigitalOcean | Brex | klaviyo | qualtrics XM

Every action an agent takes looks the same, whether it's routine work or an attack.

In C1's 2026 survey of 508 IT and security leaders, 95% said their enterprises run AI agents autonomously, yet only 19% adjust access continuously based on policy and context.[1] Nearly everyone else hands an agent its access up front and hopes it stays on task. But an agent can use that access thousands of times an hour, steered by whatever content lands in front of it. Each action looks routine, right up until one isn't.

Security's answer to agent risk is fragmenting into a tool per surface: prompts, models, agent traffic, and whatever ships next quarter. But follow any agent incident to the end and you land on identity questions: whose agent, acting as what, with access to what, approved by whom. AI didn't create a new kind of risk. It handed the oldest one, access, to a whole new population of actors.

What C1 delivers

Risky calls stop before running

Every tool call is checked against its task and your policy before it runs.

No rewrites to your agents

Agent traffic routes through an identity-aware AI gateway. No SDK, no new framework to adopt.

Every action has an owner

Each call records the agent, tool, policy decision, outcome, and accountable owner.

C1 controls the call itself: security at the point where AI acts.

People use access. Agents call tools. That shift is why C1 controls the moment an agent tries to act. Agent traffic routes through an identity-aware AI gateway that registers your MCP (Model Context Protocol) servers and internal tools, with no rewrites to your agents required. Before a tool call runs, C1 checks for private data access, exfiltration attempts, and untrusted content that may be steering the agent through prompt injection.

Policy then decides what happens next: block the call, hold it for human approval, or redact sensitive fields from the tool response. Observe mode shows what a rule would have done before you enforce it. Each agent is scoped to only the tools its task requires, limiting the reach of risky calls. Every action is tied back to an accountable owner.

Runtime control for every agent.

Control every call without touching agent code.
The gateway sits in front of your MCP servers and internal tools, so every call passes through one point where policy runs. Shadow AI discovery surfaces the agents that aren't behind the gateway yet.

Catch dangerous calls before they run.
Each call is checked for what precedes damage: access to private data, attempts to move it out, and untrusted content that may be steering the agent through prompt injection. Policy stops the call when those signals combine.

Audit-ready, down to every agent action.
Each call records the agent, the tool, the policy decision, the outcome, and the accountable owner. An answer you can hand an auditor.

Shrink each agent's blast radius.
Every agent is scoped to the toolset its task needs, with each tool classified on a scale from read-only to dangerous. Auto-approve the reads, gate the calls that change or expose data.

Turn on enforcement with confidence.
Policy decides per rule: block, hold for approval, or redact tool responses. Observe mode shows what a rule would have done before you enforce it, and failed checks hold the call instead of waving it through.

Govern agents at runtime.

Point solutions will keep multiplying, because AI keeps shipping new surfaces to secure. The stack will keep changing, but the actor remains: something will still authenticate, hold access, and act on someone's behalf. Governing those actors doesn't take a new stack beside your identity program. It's the same discipline: owner, scope, policy, and audit, applied at the moment an agent acts.

And runtime governance is one motion in a larger loop on the control plane you already run. Shadow AI discovery sees what's running. Agentic vault secures the credentials your agents run on and plants decoys that catch theft. Agentic security & intelligence closes what goes wrong. Start by pointing traffic at the gateway and watching in observe mode; enforce when the rules prove themselves. Agents end up governed like everyone else in the company, which was always the point.

GET STARTED

Talk to a product expert

Book a demo
c1.ai/demo

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] C1, 2026 Future of Identity Report, a survey of 508 IT and security leaders at U.S. enterprises with 1,000+ employees.