

Agent SSO and Runtime Governance

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



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.ai'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 agents their access up front and hopes they stay 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.ai 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 C1 MCP 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.ai controls the call itself: security at the point where AI acts.

People use access. Agents call tools. That shift is why C1.ai controls the moment an agent tries to act. Agent traffic routes through C1 MCP Gateway, which registers your MCP (Model Context Protocol) servers and internal tools, with no rewrites to your agents required. Before a tool call runs, C1.ai 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.

Replace standing agent credentials with short-lived tokens and govern every tool call at runtime. Agents take one of two paths: through C1 MCP Gateway, in front of your MCP servers and internal tools, or, for MCP servers that support the standard, a scoped, C1.ai-signed token issued for the signed-in user. Shadow AI Discovery surfaces the agents that aren't governed yet.



Shrink each agent's blast radius.

Every agent is scoped to the toolset its task needs, with each tool classified as Read, Write, Destructive, Sensitive, or Dangerous. Auto-approve the reads, gate the calls that change or expose data.



Catch dangerous calls before they run.

Classifier rules gate on risk combinations, not just tool identity: private-data sensitivity, untrusted-content exposure, and exfiltration capability. When those signals combine, policy blocks the call or requires approval mid-call.



Turn on enforcement with confidence.

Policy decides per rule: block the call, require approval, or rewrite inputs, outputs, or the assistant's outgoing response. Run any rule in Observe mode to log what it would have done before it denies anything.



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.

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.

Runtime governance is one motion in a larger loop on the C1 Identity Platform. Shadow AI Discovery sees what's running. C1 Vault secures the credentials your agents run on, and decoy credentials, planted where an attacker would look, raise a critical finding the moment they're used. NHI & Identity Security closes what goes wrong. Start by pointing traffic at C1 MCP 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.ai



C1.ai is the identity platform for the AI era. Our platform connects humans and agents to enterprise systems with scoped access, enforced policies, human approval, and a record of every action. With C1.ai, the fastest path to AI adoption at scale is the governed path — giving organizations visibility and control without sacrificing security or slowing innovation. Companies such as Ramp, Zscaler, Qualtrics, and DoorDash trust C1.ai to govern identity, secure agents, and empower their teams to put AI to work.

Put AI to work.

