Enable the *agentic enterprise.* Scale productivity. Not exposure.

THE PROBLEM

AI agents are *already in your environment.*

They work alongside your teams across the business: writing and reviewing code, moving tickets, reading and drafting documents, querying data, and acting across customer and financial systems.

Agentic identities already outnumber people by roughly 100× and rising. Sub-agents and ephemeral workloads create nonstop authentication and authorization loops at machine speed.

Every one of them is acting under some credential. **Almost none are acting under policy.**

Authentication is only the *beginning.*

An agent can be perfectly authenticated and still inherit every permission available to the credential behind it. Traditional identity controls establish the session. They typically do not narrow authority at the moment of each tool call or preserve a link between the action and the accountable owner.

Any standing credential or direct connection outside the governed path becomes a bypass. A control that can be routed around is not a control.

For agentic workflows, least privilege must be enforced at runtime, on every call, before execution.

What goes wrong without *runtime governance.*

- **Long-lived credentials.** Agents may retain standing secrets to downstream systems beyond the task that required them.
- **Broad inherited access.** The agent operates with the reach of the connected account, not the narrow authority required for the task.
- **Consequential actions without approval.** Sensitive calls can execute without first reaching the accountable owner for review.
- **No single record. No single switch.** There is no central place to see which agent acted, under whose authority, and what policy allowed it, or to cut access at the agent or session level.

“Fully autonomous agents are not ready for most enterprise use cases, and human oversight remains essential. Semiautonomous deployments are what enterprises must plan for.”

— GARTNER® HYPE CYCLE™ FOR AGENTIC AI, 2026.

THE SOLUTION

Secure agentic workflows. *Enforce least privilege at runtime.*

Governance starts with identity. NewCore makes every agent known, accountable, and policy-scoped across systems, tools, and actions.

Authority is explicit and bounded. An agent drawing authority from a live user session cannot exceed that user's permissions. An agent operating independently runs under a scoped, time-bound policy ceiling. No agent can approve its own elevation.

Agent Guardian enforces that scope at runtime with short-lived credentials, authorization on every governed call, human approval when required, and a record of every decision.

One governance model for **what every agent can and cannot do.**

Five controls for *agents already at work.*

Control agent access

Define which systems and tools each agent can reach and what it can do in each one. Allow, ask, or deny per tool, assigned to people or groups. Anything not listed is denied.

Keep humans in the loop

When policy flags a sensitive action, the exact request reaches the accountable owner's phone with the agent, action, resource, scope, and risk context. They can approve, deny, or let it expire within a bounded window. Approval applies only to that request; lower-risk activity continues under policy.

Make the governed path the fastest path

Teams register a server and point an agent at it themselves, in minutes, which is easier than wiring around it. Ownership is captured at registration, so escalation paths exist before the first incident.

Enforce least privilege

Broker short-lived, task-scoped access on demand. The agent does not receive the downstream system's standing credential or inherit a broad service account.

Prove governance and retain control

One Event Log records the agent, the accountable person or team, the policy, the tool, the action, and the approval decision. Revocation is centralized at the agent or session level.

HOW IT WORKS

One governance model, *from request to audit trail.*

01 **Register**

Agents, people, applications, APIs, and MCP servers in one governed registry.

02 **Connect**

Authenticate the agent and bind its authority to an accountable person or team.

03 **Broker**

Issue short-lived, scoped access on demand.

04 **Authorize**

Evaluate each tool call at runtime. Allow, deny, or route to a responsible reviewer for mobile approval.

05 **Audit**

One record of every decision.

06 **Cut**

Terminate the relevant session, agent, or human access path.

Agents operate faster than human review can scale. **Governance must enforce routine decisions at machine speed, and bring people in only where policy requires.**

See agentic governance in action.

Watch an agent request a sensitive tool call and proceed only within the scope identity, policy, and approval permit. Scope a bounded pilot in an isolated environment, then expand once the model is proven.

Request a Demo →