The Agent Has No Manager

By Rahul Jindal

A finance agent at a mid-size GCC runs every night. It reconciles invoices, flags the odd ones, and pushes the clean batch to the payment queue. An engineer scoped it in an afternoon. It has read access to three systems and write access to one. Nobody can tell you whose headcount it sits under, who approves what it can touch, or what happens the night it decides a fraudulent invoice looks clean. It has a password. It does not have a manager.

Multiply that by every team that has quietly shipped one. This is where most enterprises are right now, and almost none of them have a name for the problem.

The clearest signal came from inside my own house. At Cloud Next this year, Google introduced Agent Identity, a new identity type built specifically for AI agents and kept deliberately separate from both human logins and the service accounts we have always handed to software. Look at what that decision admits. The company had two ways to say who someone is, decided neither one fit an agent, and built a third. Around it came Agent Gateway, which routes every agent's traffic so it cannot quietly reach a system it was never cleared for, and a Principal Access Boundary that sets hard limits on what an agent can never touch, no matter what it is told to do.

This is not one company's bet. The same season, outside Google, Deno shipped Clawpatrol, a firewall whose only job is to sit between an agent and the systems it reaches for, and Ory, a serious name in enterprise identity, shipped Talos, which issues narrow, time-boxed permissions instead of a borrowed master key. When the largest cloud and the independent toolmakers move the same way in the same window, the direction is the story.

The signal is loud. The smartest builders have stopped asking whether they can make an agent. They have started asking who the agent is allowed to be, and what stops it when it is wrong.

> ***"The enterprise question flipped from 'can we build it' to 'who do we let it act as, and what catches it when it fails.'"***

Why this lands on your desk, not the CISO's

The instinct is to file this under security and move on. That instinct is the trap. Security can build the firewall and mint the tokens. Security cannot answer the questions the firewall forces into the open, and every one of those questions is a People and Operations question.

When an agent acts inside your company, it inherits a role. Whose role? When it makes a call that costs money or touches a customer, accountability rolls up to a human. Which human? When it sits between two teams and does work that used to belong to a person on each side, the org chart now has a box that no leader owns. We spent a century building the discipline of deciding who is allowed to do what, who signs off, and who answers when it breaks. That discipline has a name. It is management. We are now hiring workers that the discipline was never written for.

> ***"A junior analyst with the access your agents already have would trigger a background check, a manager, and a review cycle. The agent got a service account and a Slack announcement."***

A diagnostic you can run on Monday

Pick the three agents already running in production somewhere in your business. For each one, force an answer to four questions. If any answer is a shrug, you have found the gap before it finds you.

1. Identity. Whose access does it borrow, and could that person do everything the agent can do? If the agent is more powerful than any single human it reports through, you have a permission you cannot govern.

2. Accountability. Name the one human who answers for its worst day. Not the team. The person. If three names come up, the real answer is zero.

3. Boundary. What is the action it is structurally unable to take, and what enforces that, code or hope? “We told it not to” is hope.

4. Review. Who looks at what it did, how often, and what gets it switched off? An agent with no review cycle is a hire with no performance management.

Most leaders expect to fail question two and pass the rest. In the rooms I have run, it is usually question three that goes quiet. Teams have decided what the agent should not do. Almost no one has made it impossible. The whole point of Google's Principal Access Boundary, and of tools like Clawpatrol and Talos, is to move that line from should not to cannot. The reason the industry built all of them at once is that intention was never enough.

What the People function owns here

This is the part that is genuinely new, and it is the part that belongs to you. Three shifts are coming whether or not anyone plans for them.

The org chart grows boxes with no person in them. Agents do real work that sits between human roles. Someone has to decide who owns that work, who its output flows to, and which budget it lands in. That is org design, and it is the People function's table.

Accountability needs a model before the incident, not after. When an agent causes a loss, “the AI did it” cannot be the answer your leadership gives a regulator, a customer, or a board. The chain from agent to human has to be designed in advance, the same way you design an approval matrix. The teams that write this down now will look composed in twelve months. The ones who wait will be writing it during the incident.

Capability becomes the unit of trust, not the login. The old model gave a trusted person broad access and hoped they used it well. The new model, the one Talos and Google's Principal Access Boundary are built around, gives any actor the narrowest possible capability for the shortest possible time. That is a cultural change as much as a technical one, and it reshapes how you think about trust for your human workforce too.

The principle that an agent should hold only what it needs for only as long as it needs it is a healthier model for people than the all-or-nothing access most companies still run.

The firms that win the next two years will not be the ones with the most agents. Everyone will have agents. They will be the ones who can answer, for any agent in their business, the simple question a good manager answers about any employee. Who are you, what are you allowed to do, and who answers for you. The technology to enforce that arrived this year. The management thinking is still missing. That gap is the work, and it is People work.

Build fast. But hire your agents like you hire your people. Give them a role, a boundary, an owner, and a review. The companies that skip that step are not moving faster. They are just finding out later.

Sources. Google Cloud, Agent Identity, Agent Gateway, Principal Access Boundary and Model Armor (new IAM for AI agents announced at Cloud Next 2026; Agent Identity for Agent Runtime now generally available, the rest in preview). Deno, Clawpatrol (agent security firewall). Ory, Talos (narrow, time-boxed permissions for users, services and AI agents).
