Thank you for downloading

Thank you for downloading this guide to roles in the AI era from House of Silver. We hope it helps you see how AI agents change the work of your people, and who should own what as they do.

When agents take on real work, three human roles appear: someone feeds the agents, someone directs them, and someone runs the system they work in. This guide describes each role, the skills it takes, how you know it is working, and how it changes as your organization grows. Most of the work grows out of jobs you already have, so read it with your own people in mind.

If you would value an experienced partner along the way, whether to map these roles onto your teams, run a leadership session, or help your people grow into them, we would be glad to hear from you.

Erik
Founder, House of Silver
hello@houseofsilver.ai | houseofsilver.ai

About House of Silver
House of Silver helps organizations adopt AI in ways that strengthen their people, keep their data safe, and deliver real impact. To see where your organization stands, take the free AI Readiness Assessment at houseofsilver.ai.

Roles in the AI Era

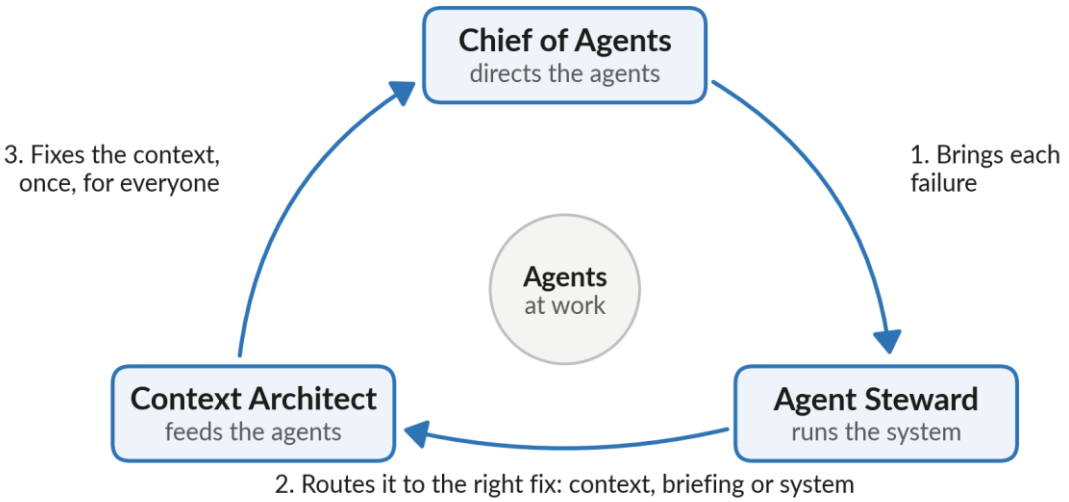
Contents

1. Context Architect 3
2. Chief of Agents 4
3. Agent Steward 5
How the roles change as you grow 6
Where to start 7
Sources 7

When agents take on real work in an organization, three human roles appear: someone feeds the agents, someone directs them, and someone runs the system they work in. We call them the **Context Architect**, the **Chief of Agents** and the **Agent Steward**.

None of them is a brand new person hired from outside. Each grows out of a role the organization already has: the subject matter expert, the operational leader and the adoption lead. What changes is where their time goes, and what they are measured on.

Microsoft, McKinsey, Harvard Business Review, KPMG and the legal press all describe the same work under different titles (see Sources). Each role also looks different at 30 people than at 1,000, so every description ends with how it scales.



How the three roles work together: a failure spotted once is fixed once, for everyone.

1. Context Architect

Purpose: make the organization's knowledge readable by agents, so every task starts from what the organization already knows instead of from zero.

Grows out of: the subject matter expert, the knowledge manager, the person everyone asks "how do we usually do this?"

Also called elsewhere: context engineer, orchestration engineer (KPMG), legal engineer.

Responsibilities

- Write down the organization's standards, methods and judgment calls as context an agent loads before it starts, not as documents someone has to remember to attach.
- Turn every mistake into a rule, and keep the reason next to the rule so people learn the history instead of just obeying it.
- Decide what an agent should know for each type of task, and what it should not see.
- Keep context current: retire outdated rules, merge duplicates, date everything.
- Review agent output in their field and trace each error back to a gap in the context.

Skills

- Deep expertise in at least one area of the business. You cannot write down what you do not know.
- Clear, structured writing. The reader is an agent that takes every word literally.
- Comfort saying "it depends" out loud, and then writing down what it depends on.
- Basic understanding of how agents load and use context. No coding required.

How you know it works

- The same mistake does not happen twice.
- A new hire, human or agent, produces work to your standard in days, not months.
- When an expert leaves, their judgment stays.

> ***The predictable fear:*** *"If I write it all down, I am replaceable." The answer is credit: the rule carries the author's name, and the expert becomes the person whose judgment runs in every project, not only the ones they touch.*

2. Chief of Agents

Purpose: set the agenda for how agents are used: decide what is done by a person and what by an agent, make sure the output is checked, and own the result in front of the client.

Grows out of: the leader who already decides how work is staffed and delivered: a founder, a COO, a department head. Everyone still directs their own agents; the Chief of Agents sets how.

Also called elsewhere: agent manager (HBR), hybrid manager (McKinsey), AI workforce manager (Microsoft), AI performance manager (KPMG), and at the top of a large organization, chief resources officer (Microsoft).

Responsibilities

- Decide which work is done by a person, which by an agent, and which by both, and keep that balance under review.
- Set how agents are briefed across their area: goal, constraints, what good looks like, the way you would brief a sharp new colleague.
- Make sure agent output is checked before it leaves the building. The client buys judgment, and judgment stays human.
- Spot where agents fail and bring each failure to the Agent Steward, so it is fixed once, for everyone.
- Keep people growing in mixed teams of people and agents, so they learn the work rather than only review it.

Skills

- Judgment of where agents add value and where they add risk, and the confidence to say no to the rest.
- Setting standards others follow: how agents are briefed, how their work is checked, and when a person takes over. Most agent failures are briefing failures.
- Skepticism without cynicism: trust the output enough to use it, not enough to skip the check.
- Coaching. Junior people now learn by reviewing agent work, and someone has to teach them how.

How you know it works

- Every team knows what its agents may do, who checks their work and who owns the result.
- The balance of people and agents is a decision, reviewed on a schedule, not something that happened by accident.
- Nothing reaches a client that no human has looked at.

> ***The predictable fear:*** *"I am losing the craft," usually heard from the team, and the Chief of Agents is the one who answers it. The craft does not disappear, it moves up a level: from doing the work to knowing what good work is, which was always the harder part.*

3. Agent Steward

Purpose: run the system the agents and people work in, so it stays reliable, safe, measured and actually used.

Grows out of: the AI adoption lead, the AI champion, the operations or IT lead who ended up owning "the AI stuff".

Also called elsewhere: AI coach (McKinsey), AI trainer and AI agent specialist (Microsoft), AI ops lead, AI governance and risk specialist, AI security specialist and AI ethicist (KPMG), evaluation engineer, forward deployed engineer, chief AI officer.

Responsibilities

- Own the reusable workflows: which tasks run through agents, in what steps, with what checks.
- Own the log. Every agent action is recorded with what changed and why, so the log becomes the organization's memory and its audit trail.
- Run the feedback loop: collect failures from the Chiefs of Agents, route them to the Context Architects, confirm the fix held.
- Set and enforce the guardrails: data access, client confidentiality, security, and regulation such as the EU AI Act.
- Measure what the agents deliver in hours, quality and risk, and stop what does not pay off.
- Help people adopt the system: training, office hours, a champions network.

Skills

- Systems thinking: seeing the workflow, not just the task.
- Change management. The tools are rarely the hard part, the habits are.
- Enough technical depth to configure agents and talk to IT, without having to be a developer.
- Working knowledge of data protection and AI regulation, or knowing who to ask.

How you know it works

- Agent use is steady and growing, not a pilot that faded after launch.
- Any agent action can be traced: who asked, what context it used, what it changed.
- Leadership can see what the agents return, in numbers, without asking for a report.

> ***The predictable fear:*** *"The log means I am being watched." The answer is framing and practice: the log is memory and credit, it shows who taught the system what, and it is never used as a surveillance tool.*

How the roles change as you grow

In a small organization the three roles are hats, worn part time by people who also do client work. In a large one each becomes a job, then a team, then a function with its own leader. The work stays the same; who carries it changes.

Role	About 30 people	About 100 people	About 500 people	1,000 people and more
Context Architect	• Every expert writes context for their own area • One person keeps the structure tidy • A few hours a week	• A named Context Architect per business area, often part time • Experts contribute • The architect edits, reviews and retires	• A small context team • A Context Architect in each business area • A lead who sets shared standards	• A central knowledge and context team • Led by a chief AI and knowledge officer • Context architects embedded in each business area
Chief of Agents	• One person sets how agents are used, often the founder or COO • Everyone directs their own agents within that agenda	• A named senior leader owns the agenda • Owns the balance of people and agents • Team leads apply it in their teams	• A Chief of Agents per major business area, such as marketing or operations • Each owns the balance of people and agents there	• A Chief of Agents in each function • An organization-wide Chief of Agents on the leadership team • Plans human and digital capacity together
Agent Steward	• Usually the founder or the most curious person • Part time • Covers workflows, the log and the basic rules	• A full time AI operations lead • A champions network in each team • A written AI policy	• A small AI operations team • Covers governance, evaluation and adoption • Stewards in the largest business areas	• An AI office, led by a Chief AI Officer • Governance, risk and security • Evaluation, reliability and adoption

The Chief of Agents changes the most. At 30 people it is one person, often the founder or COO, who sets how agents are used across the organization. At 100 it is a named senior leader, and team leads apply the agenda in their teams. At 500 each major business area has its own: a Chief of Agents for marketing, say, or for operations. At 1,000 and more they report to an organization-wide Chief of Agents on the leadership team, close to what Microsoft calls a chief resources officer: the person who decides how much of the organization's capacity is human and how much is digital, and plans hiring, training and agent investment as one budget.

Two things do not change with size: a human owns every result a client sees, and every rule carries the reason it exists.

Where to start

You do not need new hires or new titles to begin. Start with the people who already do this work informally, give them time for it, and make the work visible.

- Name a Context Architect for the one area where knowledge matters most, and give them a few hours a week to write down how good work is done there.
- Name one Chief of Agents to set how agents are used: what agents do, what a person checks, and who owns the result. In a small organization this is often the founder or COO.
- Choose one Agent Steward, often your most curious early adopter, to keep a record of which agents run, what they may do and what they deliver.
- Bring the three together once a month. Each brings one failure, and together they decide where it is fixed: in the context, in the briefing or in the system.
- Revisit the roles as you grow, using the table above. Turn a hat into a job when the work no longer fits in the hours.

Sources

- Microsoft Work Trend Index 2025: The year the Frontier Firm is born (agent boss, AI workforce manager, AI agent specialist, AI trainer, chief resources officer)
- McKinsey: Six shifts to build the agentic organization of the future (agent orchestrator, hybrid manager, AI coach)
- Harvard Business Review, Feb 2026: To thrive in the AI era, companies need agent managers (agent manager)
- KPMG: Agents of change, workforce (orchestration engineer, AI performance manager, governance, security, ethicist)
- Legaltech News via Morgan Lewis: Demand for legal engineers skyrockets (legal engineer, chief AI and knowledge officer)
- OfficeChai: Context engineering, Lutke and Karpathy (context engineering)