SoftServe

THE NEW CHAIN OF COMMAND

What it takes to keep institutional knowledge, governance, and accountability intact as agents join the org chart.



IN BRIEF

01 51% of software teams already use agentic AI in some form, but the org chart has yet to be updated to reflect it.

02 Productivity metrics capture what AI builds but miss what it depletes — institutional knowledge and human expertise.

03 Fiduciary exposure does not wait for regulatory clarity, and executives who treat governance as an IT problem will be the ones held accountable by their boards.

EXECUTIVE SUMMARY

Agentic AI is already in most workplaces in some form, and the efficiency numbers look good.



In fact, many new deployments keep delivering results, giving boards reason to continue funding the technology. Some are investing more in models, while others are capping their spending and turning to cheaper platforms for routine work. In either case, it's not the model that decides whether AI's value survives contact with the rest of the business. Organizations also have to redesign workflows, redefine roles, prepare their teams, and decide who's accountable when agents make decisions — and mistakes. Yet even when companies skip those steps, the productivity metrics still look great. The long-term value, on the other hand, does not.

Which raises challenges organizations have yet to address:

- Who owns the agents?
- Who bears liability when they fail?
- What happens to the institutional knowledge that disappears when AI takes?

These are not risks on the horizon. They are current conditions — ones that most leadership teams have yet to address, let alone identify. Every AI deployment without a management model, every workflow automated without assigned accountability, every junior role eliminated before it could capture what it knew — these are governance debts accruing now, invisible in the productivity metrics used to measure success.

The organization, of course, is no stranger to transformation. Peter Drucker, widely regarded as the father of modern management, identified three major restructurings over the last two centuries: the separation of management from ownership in the late 1800s, the rise of command-and-control hierarchy in the 1920s, and the move to the information-based organization in 1988.

Each shift remade how work was both done and governed. Yet for the better part of a century, the CORE (Clarity, Ownership, Relationships, and Execution) management model has remained intact. Change occurred only at the margins — flatter hierarchies by a degree, matrixed reporting lines in select departments, broader spans of control. But the modern enterprise has never been rebuilt from the ground up.

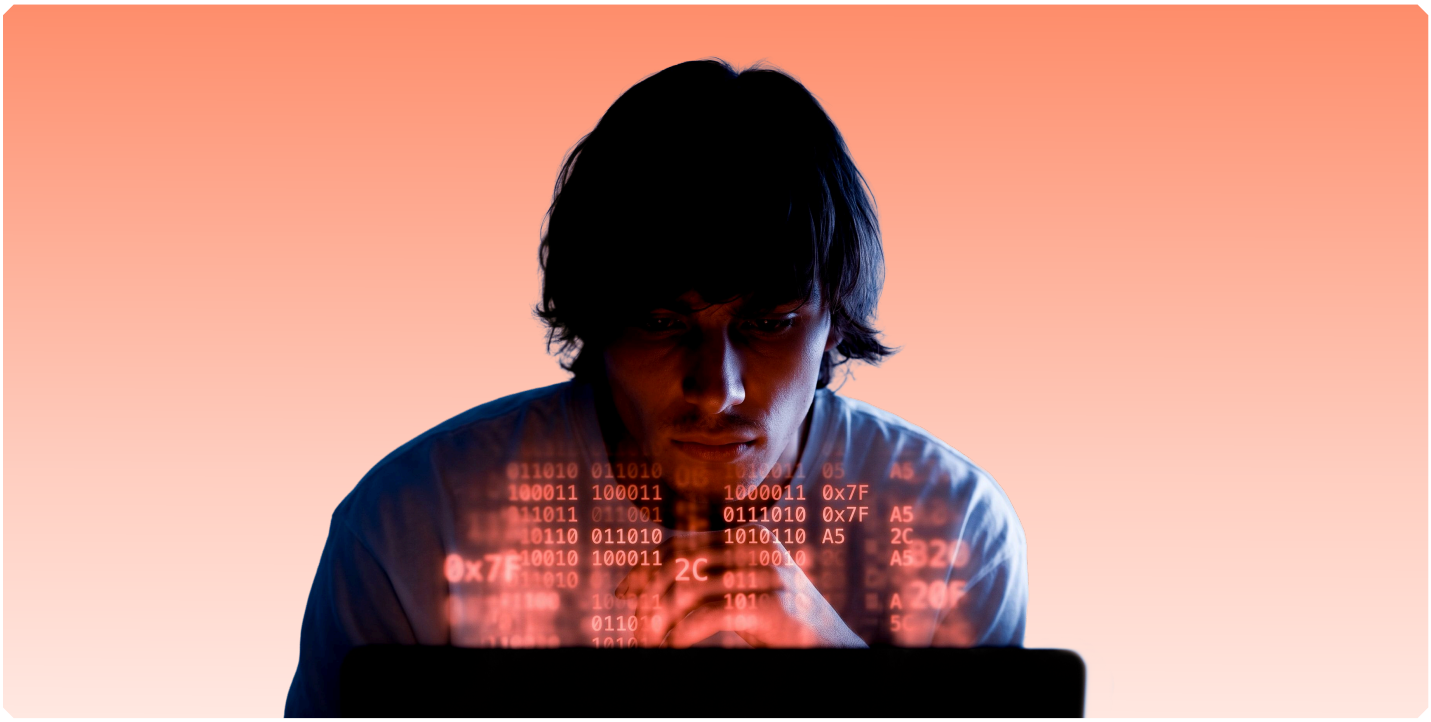

Then the agent arrived — and the org chart had no place for it.

Right now, AI agents are moving into material roles faster than organizational design can keep up with. Improvements in productivity are masking a knowledge crisis brewing beneath the surface. Organizations are automating away institutional expertise without transferring it to the systems that replace it.

And the executives signing off on those decisions lack the frameworks needed to govern those systems.

Now, they find themselves confronting a question Drucker never had to ask: Are we at a genuine inflection point — a radical remake of the organization and of management itself?

EVERY EFFICIENCY GAIN COMES FROM SOMEWHERE

Every AI deployment tells the same success story. Output rises, costs fall, and teams get more work done with the same — or fewer — people.

Our MIT Technology Review Insights global survey of 300 senior executives found that organizations expect agentic AI to accelerate software delivery by an average of 37%. But those efficiency metrics only measure what organizations would gain. They rarely capture what tends to disappear along the way — the institutional knowledge lost through attrition, the judgment that never transfers, the context that tells people when something isn't right.

Most organizations have seen some version of this before. A company streamlines a process, reduces headcount, and improves productivity. Years later, it discovers that the people who understood why the process worked are gone. And it often cannot find the right talent to replace them.

These companies keep repeating the pattern anyway — now at a scale experts believe dwarfs every previous round.

Adoption looks nearly universal on the surface. Optro, a GFC software solutions firm, found that 82% of more than 400 surveyed organizations are adopting AI across their functions. Relatedly, our MIT Technology Review Insights global study shows that 51% of software teams use agentic AI in some capacity, and another 45% expect to within the next year.

But deployment is not the same as scale. A recent LangChain survey found that 57.3% of organizations now have agents running in production environments — a 12% jump from the year prior. Another 30.4% are actively building agents and plan to launch them, suggesting that adoption continues to grow rather than stall. Yet the same report found that 67% of companies with more than 10,000 employees have agents in active use, compared with 50% of those with fewer than 100 employees. While many organizations have begun their agentic journeys, few have gotten past the deployment stage.



And that's expensive. Productivity is only a visible outcome, and much of the agentic AI spend right now is surface-level. Most organizations, after all, have focused on automating tasks rather than redesigning the processes underneath. And a pilot dropped into an existing workflow with no change to the work structure results in deployment without value — a deeper issue only discovered during an earnings call when a shareholder asks why the token bill doubled.

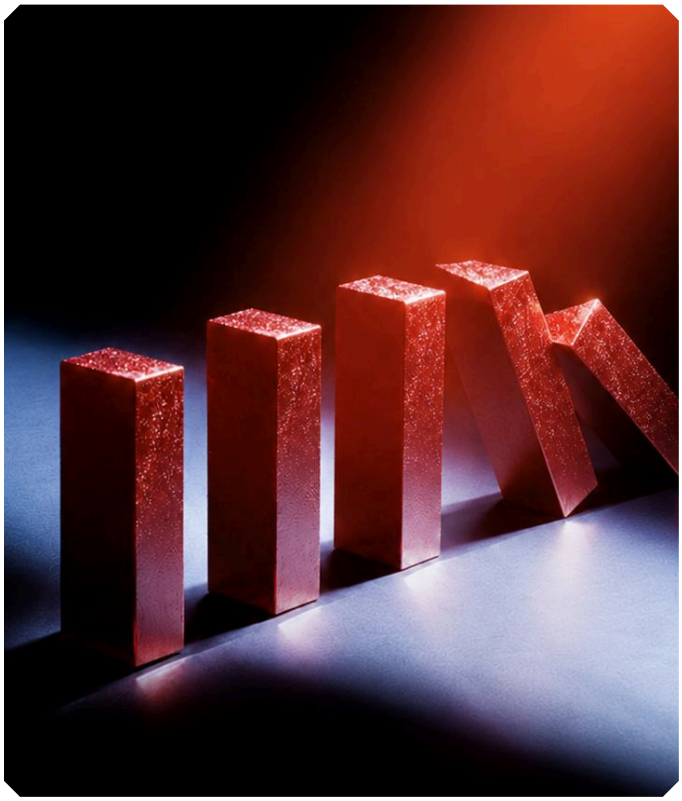

It's a question that's becoming harder to dodge. Gartner's latest research ties token overspending to weak controls around agent usage — too much autonomy, oversized context windows, no disciplined feedback loop. (This is a different problem than per-token pricing, which has fallen, according to industry trackers.) Cost controls can show how much an agent consumes. They cannot show why it consumed those resources or whether the decisions behind that consumption added value.

That requires visibility into how the agent operates. The cost shows up in the budget, but the failure often starts in governance. Without an inference audit of what an agent did, what informed its decision, and where the resources were spent, organizations cannot say whether rising spend is buying more value or simply more retries.

AI deploys quickly; organizational change doesn't. Teams often slow down before they speed up as they learn new ways of working. Some businesses expecting immediate or linear returns are now comparing the AI bill with the outsourcing models they replaced.

Companies that do redesign around AI, however, will likely experience improvements in both performance and financial profiles. In fact, organizations can reduce general and administrative expenses from as much as 8% of revenue to as low as 5% through AI-enabled redesign, according to Boston Consulting Group's research on corporate functions. But how long it takes companies to catch up when the work has already changed is an age-old question playing out again.

THE GOVERNANCE MODEL NEVER ANTICIPATED NON-HUMAN ACTORS

From paper to digital. Filing cabinets to cloud computing. In-person teams to remote offices. Every major — and inevitable — business transformation has produced a lag period where the work evolves well before the operating model does.

Agentic AI is creating that same lag again — only faster and with higher stakes.

In many organizations, AI agents already sit alongside employees — making decisions, completing tasks, and handling customer interactions that carry legal, financial, or regulatory consequences. SoftServe's conversations with enterprise leaders show just how those arrangements vary as organizations determine where agents belong and how much authority to give them. Some have issued employee IDs to their agents — granting them access to the same systems and controls as humans — on the premise that anything that acts like an employee should be treated like one. Others are building agents to replace full-time workers whose responsibilities they are absorbing, without an equivalent structure in place. Neither approach has been tested by a full budget cycle, a lawsuit, or a regulator.

Yet organizations still govern those agent activities as if people perform all of them. Businesses designed their chains of command, performance management, escalation paths, and liability for an entirely human workforce. As agents take on more responsibility, however, that assumption gets harder to sustain.

Consider what this looks like in practice. An agent managing customer escalations can issue refunds, rebook services, or flag accounts for review. No human identified to answer for its outputs. No defined limit on what it can commit to. No trigger to loop in a human before it makes a consequential decision.

This is not hypothetical. A separate Optro survey of more than 800 IT, audit, and GRC (Governance, Risk, and Compliance) professionals found that only 34% have an AI model inventory, and just 31% have AI incident response procedures. But organizations cannot govern what they have not cataloged — Grant Thornton's 2026 AI Impact Survey found that 78% of executives are not confident that they could pass an independent AI governance audit within 90 days.

A market is already forming to serve this uncertainty — governance-focused platforms and orchestration tools promising oversight that the current org chart doesn't yet provide. Neither category has a finished answer.

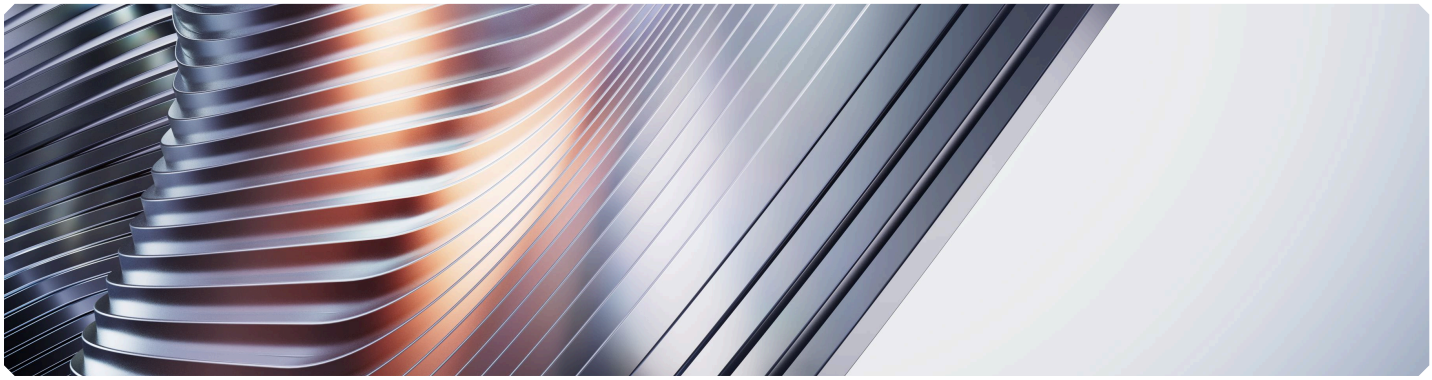

WHO OWNS THE DECISION?

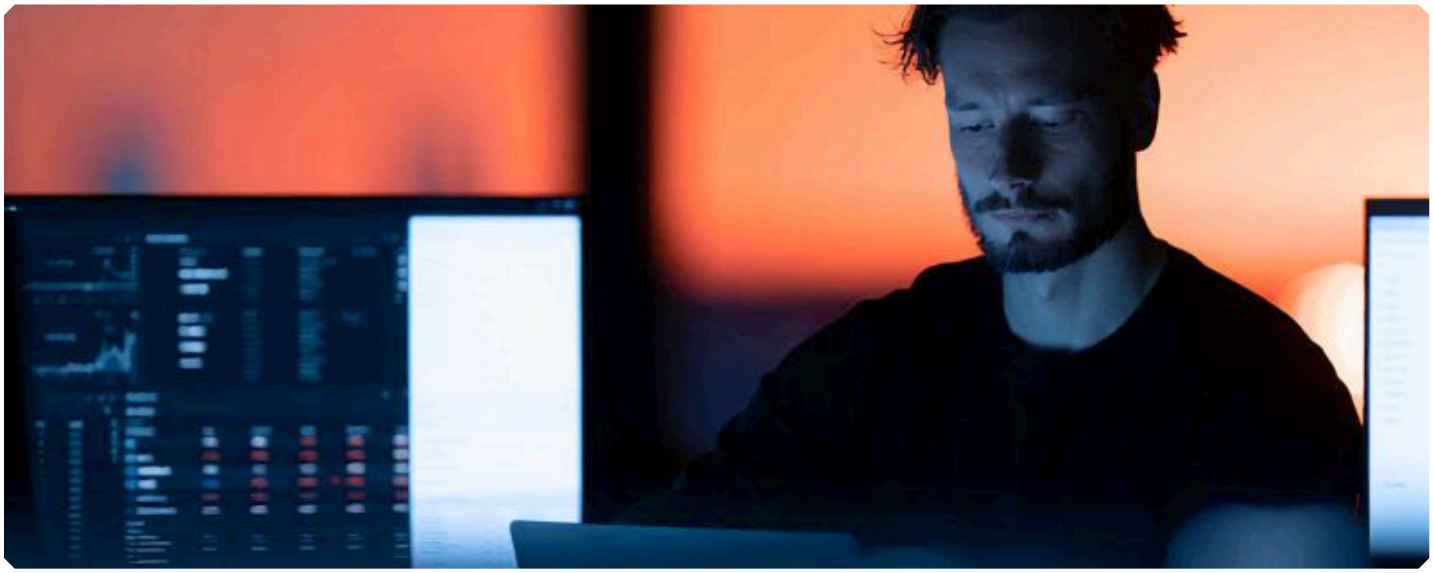

Most leadership teams can tell you who built their AI agents, who maintains them, and who manages the surrounding infrastructure. Ask who owns the outcome — the decision an agent influences, the commitment it extends, or the error it produces — and the answer is murkier.

Organizations have always linked authority to accountability; the more authority a person holds, the clearer the expectations for ownership, escalation, and oversight.

Agents now make consequential decisions within the same system, but few organizations have assigned human accountability for those actions. Gravitee's latest research on agentic AI security found that, as of April 2026, only 7.2% of organizations have named someone formally responsible for AI agent behavior — up roughly six points from four months earlier. That is still fewer than one in 10.

Most organizations, then, will have to settle questions they have yet to consider:

- What can an agent do on its own?
- When does a human need to get involved?
- What is the intent behind the decision we are asking the agent to make?
- Which actions require approval?
- Who owns the outcome when the stakes are financial, operational, or regulatory?

AUDITING THE AGENT, NOT JUST THE OUTCOME

Governing agentic AI demands more than assigning ownership. Organizations have to audit agents differently than most enterprises are built to do. Not because the tools are missing, but because they're auditing the wrong things.

Most governance approaches today focus on whether an agent performs reliably. But reliability alone cannot explain how an agent reached a decision. When one agent's output feeds another agent's actions, problems in the inputs, context, tool calls, or handoffs can travel downstream — even when every individual output looks fine at the time.

Tracing that chain back to its source, agent by agent, is a discipline in itself. Every time an organization retunes an agent's configuration, the inference chain resets, and the audit has to run again. That means operating two distinct processes instead of one: a reliability review, which most compliance functions already understand, and an inference audit — why an agent made its decision, what it dismissed, whether the same reasoning holds under new facts — which almost none have built yet. And organizations that ignore inference govern against the wrong standard.

The evidence shows this at scale. Our MIT Technology Review Insights global study found that governance, auditability, and reliability rank among the biggest obstacles organizations face in scaling agentic AI, with interviewees repeatedly stressing the need for ongoing human oversight as agents gain greater autonomy. Yet nearly three-quarters expect agents to manage most or all their development lifecycles within two years.

Optro's research, for its part, reveals where the strain is showing up: a third of organizations feel responsible for AI-related risks they can't control directly, and a third of CISOs report a mismatch between responsibility and authority over third-party risk. Even more striking, respondents identified employee behavior as the bigger risk — people using AI tools without oversight, and inside systems nobody vetted or approved. Put another way, authority over agents is expanding faster than any decision about who's supposed to hold it.

AI IS CHANGING HOW ORGANIZATIONS BUILD EXPERTISE

Most leaders learned the business by doing the work.

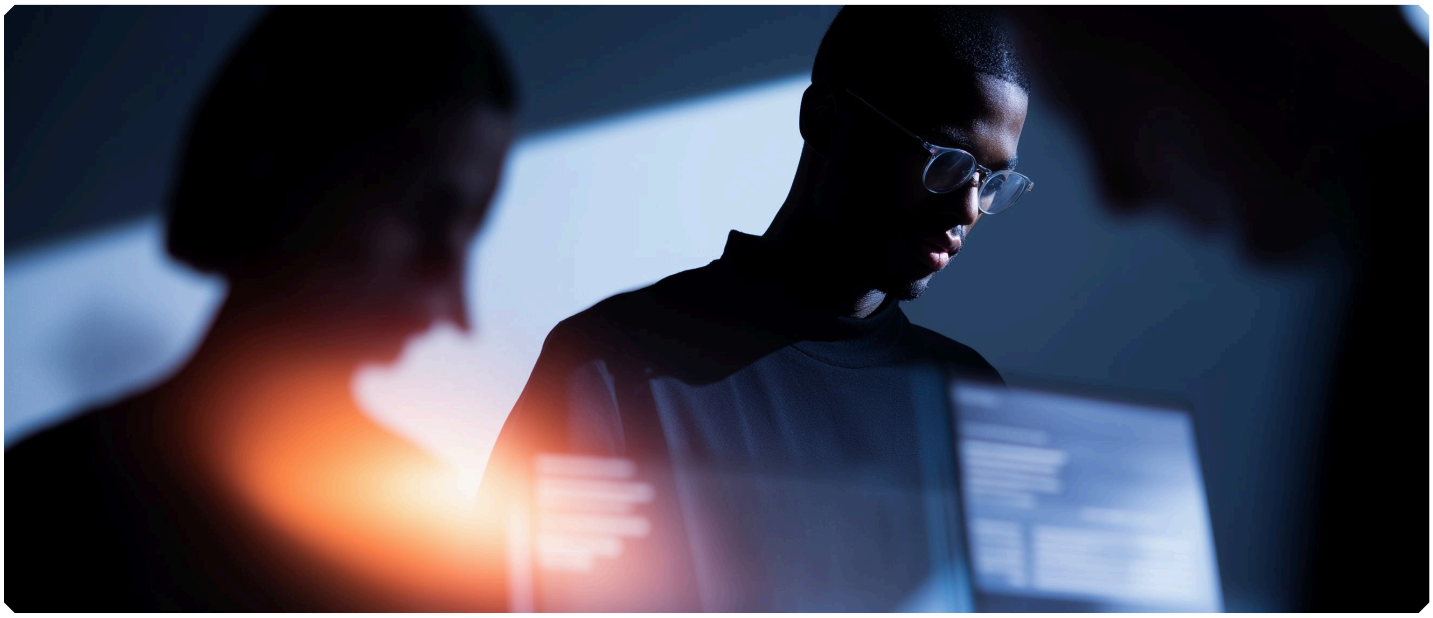

They started with smaller responsibilities. They saw what happened when things went wrong. They learned how decisions got made — and more importantly, why certain decisions should not be made. Over time, they developed a judgment that no documentation could capture fully.

Much of that expertise developed implicitly through the work itself, regardless of formal knowledge transfer programs. But AI is starting to change that.

Ask any executive what "agentic" means for their own workforce specifically, and many struggle to answer. That's not because they struggle to understand the technology. Rather, few organizations have defined what agentic AI means in the context of their people, processes, and operating model. The term arrived faster than the knowledge, and without a definition, companies do not have to decide which jobs to displace.

But research shows that the entry-level remits that once taught people how to work are the same ones AI is absorbing first. According to Korn Ferry Institute research, the routine tasks that define early-career roles — the ones that build experience before greater responsibility — face greater exposure to AI than the work done by more senior employees. And when those duties are automated, the learning that came with them vanishes, too.

Workers are feeling this change before the org chart may admit it. In a recent ADP global survey, only 16% of repetitive-task workers and 18% of skilled-task workers strongly agree their jobs were safe from elimination, compared to 30% of knowledge workers.

That divide is about to get harder — not easier — to manage. Our MIT Technology Review Insights global study found that more and more organizations expect agents to manage larger portions of the development lifecycle while creating new positions focused on configuring and evaluating those agents. More than half of respondents now identify AI engineering as a top hiring priority, even as automation may weaken the traditional experience pipeline that develops the talent those roles require.

At the same time, the people with the deepest institutional knowledge are reaching retirement age. Yet, most organizations still lack formal knowledge retention programs for employees nearing retirement. A recent American Productivity & Quality Center global survey found that although organizations worry about knowledge loss from retiring employees, only 8% consistently capture it before they leave, while 16% do not even attempt to do so. Knowledge transfer, when it happens, is ad hoc — rushed handovers, few shadowing sessions, exit interviews in someone's final week.

But the agents replacing entry-level positions do not inherit the knowledge those roles would have developed. Nor do they fill the information gaps left by retired employees. Instead, they inherit a configuration. Under real operating conditions, the difference can determine whether a system stays within expected parameters or fails in ways no one anticipated — because no one was left who remembered why certain guardrails existed.

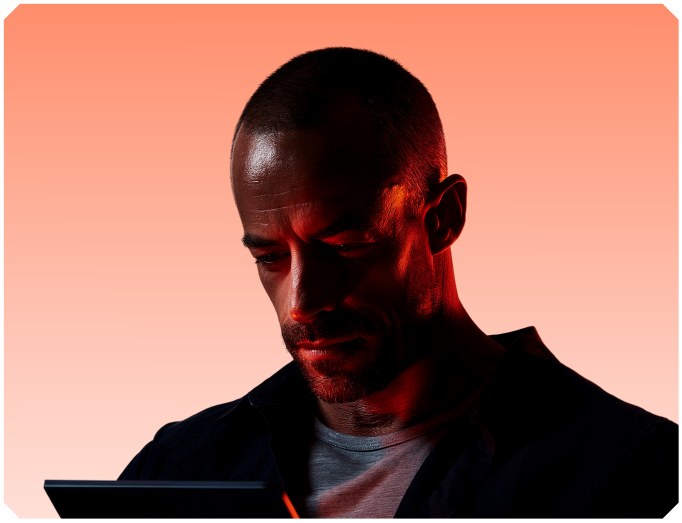

Knowledge at risk is not the kind that appears in transition plans. It's the reasoning cultivated over years of experience, developed by observing decisions play out. And the mechanism behind it is disappearing alongside the work itself.

That's because too many organizations treat data and knowledge as synonymous. Data — that is, a raw collection of facts and figures — feeds an agent. Knowledge, on the other hand, feeds a person. Yet, organizations are expanding agent authority while eliminating the expertise needed to supervise them. And while they can encode information, roles, processes, and documented knowledge into agent systems, tacit judgment is much harder to transfer.

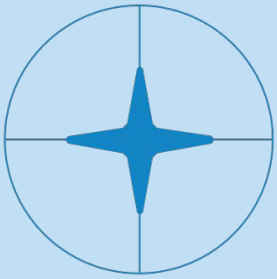

Human expertise is knowledge plus judgment. Agent expertise, for now, is data plus process. What disappears when a role gets automated first is the distance between the two.

AWARENESS IS NOT GOVERNANCE

Most conversations about AI governance still circle the technical layer — model accuracy, output quality, bias, explainability. But while these concerns are legitimate, they sit downstream of the governance questions with fiduciary weight.

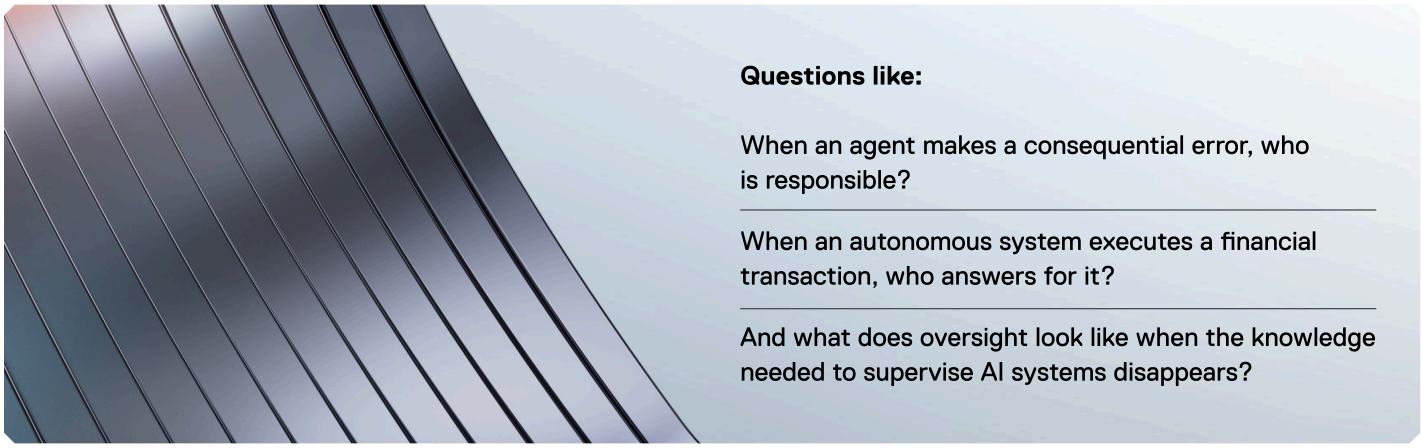

The answers, for most organizations, are incomplete. Nearly all respondents (96%) in Optro's survey report that boards or executives are involved in AI governance — yet only one in four has a fully implemented program to show for it. The rest have piecemeal policies in place or in development. But without clearly defined roles or coordinated processes between functions, those policies stay theoretical while AI agents keep running.

The obstacles, though, are not what most organizations assume. Gaps in clear ownership, thin internal expertise, and limited resources rank as the top barriers to AI governance, according to Optro's research. Fewer than 15% cited a lack of tools as the challenge. The biggest barriers are organizational rather than technological.

But many organizations are waiting for clearer rules before investing in AI governance. And that logic is understandable. Why build for requirements that haven't been fully defined?

The issue is that accountability arrives long before regulation does.

If an AI system contributes to a compliance failure, a flawed financial decision, or a customer issue, someone still has to explain what happened. Customers, boards, and auditors do not wait for regulatory guidance.



The EU Artificial Intelligence Act reflects this shift. Rather than focusing solely on model performance, it emphasizes human oversight, accountability, and auditability for high-risk systems — and carries fines of up to 7% of global revenue for violations. Yet Optro's research found that while 86% of organizations are aware of current or impending AI regulations, only 25% have a fully implemented governance program.

Part of why implementation lags behind awareness is structural. Agents don't absorb regulatory change independently. Organizations have to translate requirements into the rules, data, permissions, and workflows that govern them. That creates a new need for change management. And while companies already have processes to train their people as regulations evolve, they still need an equivalent discipline for agents.

Organizations making headway understand that awareness is not governance — and they're not waiting for guidance to be defined. According to Optro, 49% are already voluntarily aligning with the National Institute of Standards and Technology AI Risk Management Framework. Building for anticipated standards now will be less expensive than rebuilding under pressure later.

Rarely does an organization evaluate an employee once and assume performance will remain constant forever. After all, change happens — to markets, to customers, to business priorities. New situations can reveal strengths and weaknesses once hidden. The same applies to AI systems, where new data and evolving customer behavior can surface edge cases that testing never anticipated.

That is why governance cannot end at deployment.

FIVE STRATEGIES FOR GOVERNING AGENTIC AI

Governance works best as an operating discipline, with review processes, escalation paths, and accountability mechanisms that evolve alongside the systems themselves.

That starts with a deceptively simple question: Where are agents actually making decisions today?

For many organizations, the footprint often extends beyond deliberate deployments to AI embedded in vendor tools and other systems. This leaves some activity operating outside formal oversight.

Five strategies can bring that activity under control:

1

Build the inventory before anything else.

Catalog every agent, including third-party AI, operating in a material role — that is, where its output affects a financial transaction, a customer-facing decision, a regulatory obligation, or a judgment call a human used to make. Define what each one can act on, what decisions it influences, and who owns its outputs. Without this picture, every other governance effort rests on an assumption.

2

Assign named human accountability for every agent in a material role.

Ownership gaps rank as the leading barrier to AI governance, according to Optro's research. Every agent in a material role needs a named human owner — not a team, not a function, but a person — accountable for its outputs and empowered to act when something goes wrong.

3

Redesign workflows around accountability, not just efficiency.

Our MIT Technology Review Insights global study found that the hardest challenge in scaling agentic AI comes from organizational change. Companies should redesign their responsibilities and operating models rather than deploying new tools within existing structures.

4

Build escalation paths before deployment.

Define in advance the conditions under which human judgment overrides agent action. Review those conditions regularly as deployments mature. Governance that activates only after a failure has a different name — incident response.

5

Integrate AI governance across functions.

Optro found that organizations with high cross-functional integration of AI governance report faster risk identification (61% vs 54%), better regulatory readiness (56% vs 44%), and stronger executive reporting (58% vs 40%) than those managing AI governance in silos. Governance strengthens when shared across the organization.

REMODELING THE HUMAN-AI WORKFORCE

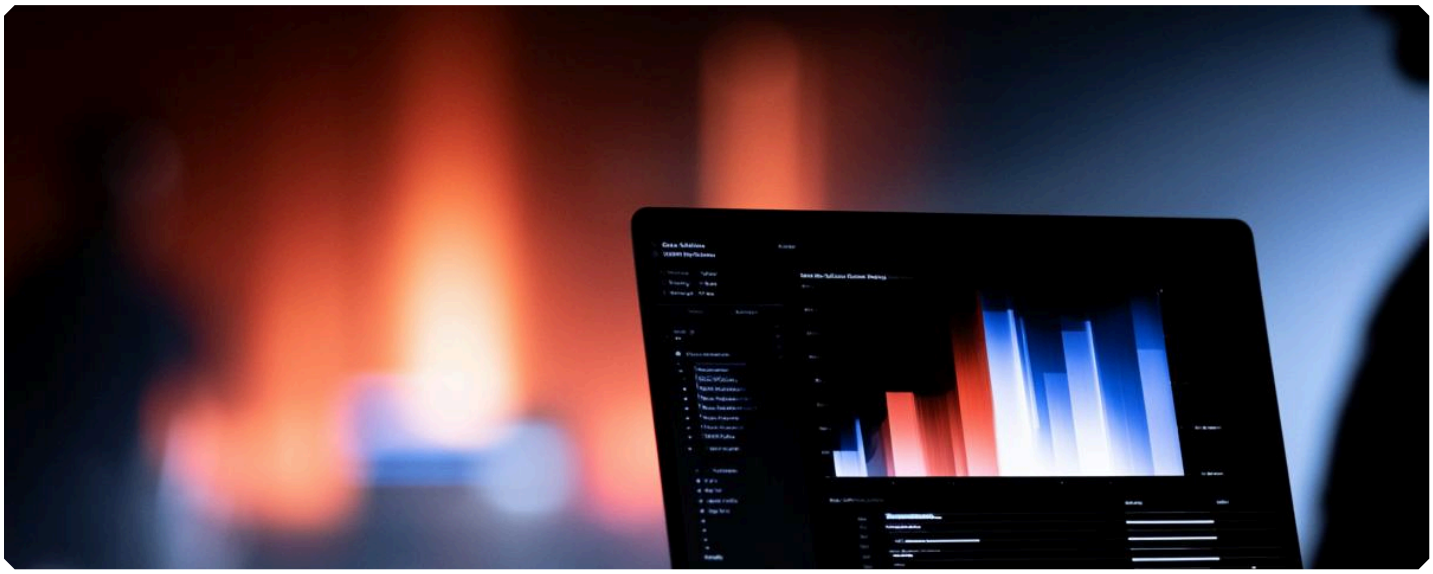

People and AI systems now work together in processes that affect customers, revenue, operations, and risk. Our MIT Technology Review Insights study traces the progression from AI assistants that improve individual productivity to agentic systems that increasingly manage workflows and lifecycles. As agents take on more of that work, organizations have to decide where their authority begins, where human judgment enters, and who ultimately owns the outcome.

That has implications for management itself. For decades, middle managers translated leadership priorities into team-level instructions while allocating resources, resolving expectations, and developing people. Agents are starting to take on some of the coordination and knowledge flow that once traveled through those layers, often without a manager in between. Propagation gets faster, but one question remains unanswered: What is middle management for once agents can execute the parts of the job that involve relaying information rather than exercising judgment?

Research suggests organizational structures are already becoming flatter. BCG's study on corporate functions finds that companies adopting AI-powered workflows are shedding layers and expanding spans of control. As a result, organizational charts have started to look less like a pyramid and more like an hourglass, with the middle layer continuing to shrink.

In fact, 41% of respondents in Korn Ferry's 2025 global workforce study say middle manager roles are disappearing from their company in an effort to cull labor costs.

What many organizations have yet to design is the management structure that reflects this reality. They know where the technology lives — but they are still working out what agents decide, when humans intervene, and who owns the outcome. They are investing in AI — but they are still determining how to preserve the experience and judgment that those systems depend on.

The evidence suggests that the distance between investment and governance readiness is wide and growing. Most organizations have policies, but few have implemented them. Most have board awareness, but few have operational ownership. And while many worry about what walks out the door when experienced employees retire, few consistently capture what those employees know before they leave.

The harder effort now is figuring out how organizations should operate around the technology. That means keeping oversight clear, holding onto institutional knowledge, and knowing where human judgment lives in the new chain of command. That foundation will determine how well the hybrid workforce works in practice.

The org chart for the future workforce needs an author. That job belongs to leadership.

AI GOVERNANCE & HYBRID WORKFORCE SELF-ASSESSMENT

Score each item 1–3. Note an owner for anything scoring 1 or 2. Revisit quarterly.

1 – NOT STARTED:	2 – IN PROGRESS:	3 – MATURE:
No meaningful action	Underway, but incomplete or inconsistent	Fully implemented, monitored, and effective

Example:

"Every agent in a material role has a named human owner." Score: 2. Notes: Customer-service agents covered; finance agents still default to the vendor. Owner: [name]. Next step: assign finance owners by the end of the quarter.

1 INVENTORY & VISIBILITY

Score:

We have a current inventory of all agents in material roles, including embedded third-party AI.	
We know what each agent can act on, what it influences, and its risks.	
We actively find and catalog shadow AI.	

2 ACCOUNTABILITY & OWNERSHIP

Score:

Every agent has a named human owner (a person, not a team).	
Policies define what agents can do alone, what needs approval, and what triggers escalation.	
Liability for AI-influenced decisions is explicitly assigned.	

3 KNOWLEDGE & EXPERTISE

Score:

We have real programs to capture institutional knowledge before people leave.	
Junior roles still build expertise, even as AI absorbs routine work.	
Agent knowledge is captured and reusable (logs, retrieval, and knowledge graphs).	
We're hiring for AI orchestration and oversight roles.	

4 GOVERNANCE & RISK

Score:

Our governance program is implemented, not just planned.	
Agents are monitored and audited after deployment, not just at launch.	
Escalation paths are tested before agents go live.	
Vendor and third-party AI risk has a clear owner.	

5 ORGANIZATION & CULTURE

Score:

Workflows were redesigned around agents, not just layered on top of them.	
Risk, legal, IT, and the business collaborate on governance.	
Leadership tracks the gap between deployment pace and governance maturity.	
Success is measured beyond productivity (knowledge, risk, judgment quality).	

SCORING

Total: ___ / 54

Below 27: Prioritize inventory and ownership now.

27–42: Progress made. Focus on the weakest category.

Above 42: Strong foundation. Refine and benchmark.

ACTION PLAN

- List your three lowest scores. Weight material-role items first.
- Assign one named person per priority. Set a deadline for each.
- Pick the fastest item to move from a 1 to a 2. Inventory and ownership items usually move fastest.
- List anything that needs a budget or new tooling. Keep it visible for the next planning cycle.
- Set your next assessment date now. Compare scores quarter over quarter.

Run this with a cross-functional leadership team. Use it to inform board updates, budget asks, and roadmap planning.

SoftServe

About Us

SoftServe is a global AI and technology services company specializing in autonomous enterprises with capabilities in AI-native applications, as well as physical and agentic AI, cloud, and data solutions. We believe technology should do more than work — it can elevate anything. We decode complexity to unlock its full potential, combining creative thinking, deep sector expertise, and precision engineering to imagine and build what's next. From strategy to execution, we turn technology into a force for growth and lasting impact.

Our global reputation is gained from more than 30 years of experience delivering superior digital solutions at exceptional speed by top-tier engineering talent to enterprise industries, including high tech, financial services, healthcare, life sciences, retail, energy, manufacturing, and more. Visit SoftServe's website, blog, and LinkedIn pages for more information.



NORTH AMERICAN HQ

201 W 5th Street, Suite 1550
Austin, TX 78701
+1 866 687 3588 (USA)
+1 647 948 7638 (Canada)

EMEA HQ

30 Cannon Street
London EC4 6XH
United Kingdom
+44 333 006 4341

info@softserveinc.com
www.softserveinc.com