




Authority without oversight

The 2026 agentic enterprise report



GRC

INTELLIGENCE

LAB



Table of contents

Executive summary

SECTION 1: The era of agentic authority

SECTION 2: How authority changes the risk profile

SECTION 3: The four agentic operating models

SECTION 4: When authority outpaces operating model

SECTION 5: The work is changing faster than the operating model

SECTION 6: 3 steps to a stronger agentic operating model

Conclusion

Appendix: Research methodology

03

04

05

06

07

09

10

11

12

Executive summary

Enterprise AI agents are now authorized to complete actions once reserved solely for humans.

This includes approving requests, triggering remediation, modifying controls and policies, changing access rights, and communicating externally. That means agents are now embedded directly into workflows with business, compliance, and regulatory consequences.


In a survey of over 400 GRC and audit leaders and qualitative interviews with members of Optro’s Customer Advisory Board (CAB), three findings emerged:

01 Agents are already authorized to do a lot. Agents are not only generating content or recommendations; 34% authorize agents to modify controls or policies, 27% permit changes to access rights, and 23% allow an agent to pause or shut down another agent.


02 But oversight of their actions remains uneven. Enterprises vary considerably in their ability to detect unexpected agent behavior, intervene effectively, and measure whether agents achieved the intended result. Among organizations with an agentic AI initiative that failed to deliver, 28% discovered it only after downstream impact.

03 And workflow redesign lags deployment. Ninety-six percent agree workflows should adapt to leverage AI capabilities, but only 9% have redesigned workflows from scratch around agents.


The collective risks posed by these trends are real and already underway. Among organizations deploying agents:




have experienced an inaccurate agent decision or output that was acted upon.



have seen an agent take an unintended action.



have experienced a control failure involving an agent.



have encountered a conflict between two agents.

External pressure is adding urgency. Three-quarters of GRC leaders say regulators or external auditors are already asking about AI agent use in their function. But authorization, oversight, measurement, and auditability are not developing at the same pace.

The organizational systems needed to govern agent authority are not advancing as fast as agents themselves. This report examines where those gaps are emerging and what GRC and audit leaders can do to adapt their operating models for the agentic reality.

SECTION 1

Agents have real authority

In the first wave of enterprise generative AI, the central governance question was whether humans could trust an AI-generated answer. Agentic AI raises new possibilities — and risks — entirely.

Surveyed organizations shared that:

38%

authorize agents to approve or reject transactions or requests.

34%

allow agents to modify controls or policies.

27%

permit changes to user permissions or access rights.

23%

allow one agent to pause or shut down another.

Organizations are granting agents permissions that could lead to steep consequences while still deciding where human judgment belongs. Using the example of identifying a regulatory change, one information security risk leader in a consumer goods organization described the potential:

“We’ve built a regulatory compliance agent, and the aim is that it does the automated landscape scanning for new regulations...flags the activity that’s needed, and does it. That’s the critical thing. Not just flag. Go and do it.”

That “go and do it” moment is where the governance challenge changes. Those boundaries should reflect consequence. Summarizing a policy does not require the same controls as changing it; identifying an access issue creates a different risk than revoking access.

Executive perspective

The C-suite findings reinforce how consequential these permissions have become:

41% report that their organizations authorize agents to modify controls or policies

32% permit changes to user permissions or access rights


28% allow agents to pause or shut down another agent

Executives report stronger containment of agentic-driven issues than the ability to reverse them. Organizations are most likely (**68%**) to say they can stop an agent or detect an issue quickly, while agent audit production (**55%**) and remediation (**51%**) are less common. High confidence in formal AI discipline sits alongside a meaningful operational recovery gap.

SECTION 2

Authority changes the risk profile

When agents can act, the stakes get higher. Events more serious than inaccurate outputs already take place in organizations using agents today:



Once an inaccurate output becomes an action, the consequences can affect enterprise systems, customers, controls, and evidence used elsewhere in the business. Delayed detection compounds the problem. Among organizations with an agentic AI initiative that failed to deliver, **28%** discovered it only after downstream impact. As one information security leader at a retail organization put it, “If the foundation’s not functioning well, and it’s still kind of broken, and you layer AI on top of it, you are not making things any better.”

When asked about the missing prerequisites likely to cause an agentic initiative to fail, survey respondents pointed to data quality and availability (**36%**) as the top issue. They also noted systems integration (**17%**), governance frameworks (**13%**), and team skills and training (**13%**).

The qualitative interviews backed this up:



Having data that is accurate and usable for AI is huge. I can’t oversell that enough based on what we’ve seen.



— An audit and risk executive at a health insurance organization



When I think about AI governance, it brings us together to make sure we’re not doing things that don’t make sense.



— An audit and risk executive at a health insurance organization



We also need to reinforce critical thinking, because if we don’t, information security could wither and die.



— An information security risk leader, consumer goods organization



As agents move into business processes, they need accurate procedures, decision rules, and operational context so that they don’t automate the gap rather than the process.

“If your procedures and processes aren’t documented to the letter, you’re training AI on nothing,” said an information security leader at a retail organization.

Four operating models emerge

Organizations are not approaching agent deployment in the same way.

Four operating models emerge from the research, differing in how businesses deploy and authorize agents, and the oversight and measurement surrounding them. It's important to note that these are archetypes, not a maturity ranking.

Measure	Moving fast	Deliberate	Flying blind	Disciplined
Share of total respondents	13%	23%	31%	33%
Agents can modify controls	20%	20%	47%	41%
Agents can approve transactions	20%	32%	42%	47%
Agents can escalate to a human	48%	77%	62%	75%
Can detect issues quickly	21%	49%	67%	88%
Consistent audit trail	24%	49%	57%	85%
Efficiency + outcome measurement	14%	50%	0%	83%
Confident measures reveal failure	54%	87%	89%	99%

13% Moving fast, weak foundations

Moving Fast organizations have extended less consequential authority than Flying Blind or Disciplined operators on most significant measures.

23% Deliberate adopters

Deliberate adopters are comparatively cautious about consequential authority.

31% Capable but flying blind

Flying Blind organizations grant substantial authority and have meaningful operational capabilities.

33% Disciplined operators

Disciplined operators pair substantial authority with the strongest support.

The differences among these operating models become clearest when organizations need to intervene, reconstruct what an agent has done, or determine whether it produced the intended outcome.

SIDEBAR

Who sits within each operating model?

The four operating models appear across organizations of different sizes, industries, and geographies, suggesting that firmographics do not fully account for these differences. A few differences stand out.

- 1 **Moving fast, weak foundations:** Skews somewhat more junior: **44%** manager/senior manager, versus **25%** among Disciplined operators.
- 2 **Capable but flying blind:** Has the highest C-suite concentration (**28%**) of the four operating models.
- 3 **Disciplined operators:** Skews more senior than Moving Fast, with stronger director-level representation and only **25%** manager/senior manager. Compliance respondents are also more concentrated in this group.
- 4 **Deliberate adopters:** Have the most balanced respondent profile, with no strong concentration by seniority or GRC function.

SECTION 4

When authority outpaces operating model

Greater agent authority creates a straightforward test for whichever operating model a business pursues:

Can the organization intervene when human judgment is required, reconstruct what an agent has done, and determine whether it produced the intended result? The research reveals significant gaps across all three.

Can someone intervene?

Human authority is nearly universal, according to survey respondents, with **98%** saying reviewers can stop or override an agent. But having the authority to intervene is different from having the conditions necessary to exercise meaningful judgment.

Among **Moving Fast** organizations, only **64%** say reviewers have enough time to evaluate agent actions properly, and **70%** say they have enough context. Among **Disciplined** operators, those figures rise to **99%** and **100%**, respectively. So, a human approval step can exist formally without necessarily functioning as meaningful oversight.

“One of the things I’ve become very careful about is human in the loop,” said an information security risk leader at a consumer goods organization.

“I think we have to use that judiciously... once you get fatigued as a consequence of too much authorization and approving clicks, people just click without thinking about it. The control efficacy is lost forever once you get to that point.”

Can the organization reconstruct what happened?

The operating models also differ sharply in their ability to understand and respond to unexpected agent behavior. Without that baseline, businesses can’t determine the what, where, and when of something potentially consequential, even if unintended, taking place.

Among **Moving Fast** organizations, **21%** say they can detect unexpected behavior quickly, **24%** maintain a consistent audit trail, and **32%** can remediate or reverse the action. Among **Disciplined** operators, the corresponding figures are **88%**, **85%**, and **72%**.

Visibility is not limited to understanding what approved agents have done. Twenty-nine percent of organizations definitely know that agents have been deployed without formal approval or review, and another **16%** suspect it. **Flying Blind** organizations are the most likely to report confirmed unauthorized deployment, at **37%**. Even among **Disciplined** operators, only **34%** say all deployed agents have been formally reviewed.

“Anybody and everybody can spin up their own agents, and the agents can talk to the internet immediately,” said an information security leader at a retail organization. “Their identity spreads like wildfire, and you can’t really get a handle on agentic identity.”

Organizations can’t reconstruct agent activity if they don’t know which agents are operating.

SECTION 4

Can the organization tell whether it's working?

The sharpest contrast appears in measurement. None of the **Flying Blind** organizations combine efficiency measures with business outcomes, compared with **83%** of **Disciplined** operators. Only **2%** of **Flying Blind** organizations primarily measure business outcomes alone; their measurement is concentrated instead on efficiency, satisfaction/adoption, and cost.

Yet **Flying Blind** organizations are highly confident in those measures. Eighty-nine percent believe their measurement would reveal a failing agentic AI initiative, only 10 points below Disciplined operators at **99%**. The organizations flying blind do not feel blind. **Moving Fast** organizations show a different pattern: just **54%** are confident their measures would reveal failure, the lowest of any operating model.


“How can we consistently quantify what value is coming out of it?...[we’re] placing a lot of bets on it, and if we can’t quantify the value and really showcase that, that’s a problem as well,” said an audit and risk executive at a health insurance organization.

External scrutiny makes these gaps more glaring. Three-quarters of GRC leaders say regulators or external auditors already ask about AI agent use in their function, and **32%** say they do so regularly. Among **Flying Blind** organizations — the group most likely to authorize agents to modify controls or policies — **86%** are already fielding these questions.


Senior leaders are also identifying clear limits to how far they are willing to go. Among C-suite respondents, the leading reasons to slow or pause deployment include security concerns (**45%**), regulatory or legal uncertainty (**44%**), a control or regulatory issue (**43%**), and inability to demonstrate ROI (**43%**). Just **5%** say nothing would cause them to slow deployment.

What intel reaches the board?

Nearly all C-suite leaders (**99%**) report at least occasional board updates on AI agents, with **73%** on a regular schedule. Notably, the amount of agentic information shared with the board varies widely. Business value/ROI is reported by **77%** of leaders, versus just **47%** for incidents and **45%** for unresolved risks. While ROI is the most commonly reported board metric, the underlying measurement discipline to track it is not universal: **69%** of C-suite respondents report formally tracking AI costs against value, while **30%** track costs less comprehensively or not at all.



~**69%** of C-suite respondents report formally tracking AI costs against value,



while ~**30%** track costs less comprehensively or not at all.

SECTION 5

The work is changing faster than the operating model

Nearly every GRC leader surveyed agrees the work itself should change.

Ninety-six percent say organizations should adapt workflows to leverage AI capabilities. But only **9%** have redesigned workflows from scratch around agents. Forty-three percent have modified existing workflows, while another **43%** have added agents without changing the underlying workflow at all.

At the same time, employees' work is already shifting. Fifty-eight percent say employees are spending more time on strategic work, and **52%** report a reduction in routine work. But adoption is also creating new work: **46%** say employees spend more time reviewing, validating, or correcting AI outputs, while **47%** report new AI-related responsibilities.

The interviews suggest leaders are becoming more selective about what they automate.

"My goal is never to replace my person," said an information security leader at a retail organization. "It's to get the minutiae out of their triage court. It's a very thoughtful process. I've done a lot of conversations, tabletops with people, really diagramming everything out to be clear about what makes sense to automate. And then there are those things you tell your leaders, 'Hey, you have to have people for this. You cannot automate this, and here's why.'"

Organizations that have redesigned workflows from scratch score substantially higher on our Operationalization Index: **7.22**, compared with **4.90** among organizations that have not redesigned from

scratch ($p < .001$). That relationship does not establish that redesign itself causes greater operationalization, but it shows how closely deliberate workflow design and more developed agentic operating models are associated.

Agents are commonly now undertaking document generation, reporting, evidence collection, and regulatory monitoring. Risk management reports particularly high production use in reporting and regulatory monitoring, while compliance leads in continuous controls and intake/triage. Internal audit trails other functions across many workflows but is actively piloting agent use.

"I think it's because as we think about how we're consolidating work, how we're deciding how teams are going to be realigned, AI is playing a huge role in that, and it's become an AI-first culture,"

said an information security leader at a retail organization.



The operating model is already changing, whether organizations are redesigning it deliberately or not.

SECTION 6

3 steps to a stronger agentic operating model

The research identifies three requirements for durable agentic operations. The goal is not maximum autonomy — it's an operating model where agent authority matches the organization's ability to govern it.

01

Design the work

Organizations must decide how workflows should operate when agents can do more than assist. That means starting with the workflow, not the agent. Design around the respective roles of humans and agents rather than simply inserting agents into existing processes. Determine which decisions require human judgment, which tasks are repeatable enough to delegate to an agent, and where humans and agents should collaborate. Also account for dependencies such as data, knowledge bases, permissions, systems, models, and other agents.

02

Design human intervention

Organizations need to define where human judgment is required and ensure reviewers have the time, context, and authority to exercise it meaningfully. Define what triggers escalation to a human, what information reviewers receive, and what happens when they disagree with an agent's action. In high-volume workflows, account for fatigue and rubber-stamping.


"There are certain things you would never want to insert an agent to take over, because it changes not just the experience, but the outcome. You have to think about where to focus more attention on people versus agents," said an information security leader at a retail organization.

03

Measure outcomes

Efficiency and cost improvements matter. But, these agentic advantages are not the same thing as an agent producing the intended business outcome. For instance, a regulatory-monitoring agent should pinpoint relevant changes and drive timely action. A control-testing agent should be measured on assurance and evidence quality.

A practical operating model should pair operational measures — speed, cost, adoption, throughput — with outcome measures specific to the workflow.




Conclusion

The four operating models show that these capabilities do not necessarily develop together.

Organizations can extend meaningful authority while remaining weak on measurement or auditability. They can keep humans formally in the loop without giving them the conditions needed for meaningful oversight. And they can build increasingly sophisticated agent programs while unapproved deployment persists elsewhere in the organization.

Regulators and auditors are already asking about agent use, while senior leaders' own concerns center on security, regulatory uncertainty, controls, and ROI. The question for organizations is increasingly: What should an agent be allowed to do, under what conditions, and how will the organization know whether it worked? The ones that get this question right will build strong governance frameworks, workflows, human intervention, and outcomes measurement around the enterprise agents they deploy.

[illegible]

Appendix

Research methodology:

The survey included 417 respondents sourced from a leading global online panel provider. They were selected from the panel based on their GRC function and role at their respective companies, as well as geographic quotas. We further screened selected respondents based on self-reported GRC and agentic AI knowledge and attentiveness to survey questions.

Role and function quotas:

Respondents selected the role that best described their primary responsibility from a list of four options, even if none was an exact fit or if they performed multiple roles. The sample was intentionally structured to reflect different perspectives, comprising C-suite / Executive (22%), VP / Head of Function (13%), Senior Director / Director (33%), and Manager / Senior Manager (32%), representing the enterprise-level decision-makers.

Additionally, respondents selected the function that best aligns with their work. We highlighted four distinct functions and consolidated the rest. The results illustrate the following breakdown: Internal Audit (17%), Risk Management (17%), Compliance/Governance (17%), Information Security/Cybersecurity (17%), and All Other (33%).

Geographic quotas:

The survey included respondents from two regions: North America (60%) and Europe (40%). Broken down, the target countries include the United States (35%), Canada (25%), the United Kingdom (20%), and Germany (20%).

Respondent screens:

Company size: All respondents must self-report that their companies have at least 250 employees. We excluded all potential respondents from smaller companies. In total, the survey includes 15% of respondents from companies with 250 to 999 employees, 49% from companies with 1,000 to 4,999 employees, 19% from companies with 5,000 to 9,999 employees, 16% from companies with 10,000 or more employees.

Information level: In our experience, “qualifying respondents” may still have too little information or knowledge about the space to

provide useful data from which to draw insights. We therefore also apply an “information” screen to respondents. Specifically, we ask whether respondents can identify the definition that best describes what an AI agent is. To qualify for this survey, a respondent must select “An AI system that can independently plan, make decisions, and take actions to achieve a goal, often across multiple systems or workflows, with limited human intervention.”

It is technically impossible and improper to list a margin of error for a survey of this type. The respondents for this sample were drawn from an online panel with an unknown relationship to the total universe, about which we also do not know the true demographics. As such, the exact representativeness of this, or any similarly produced sample, is unknown.

Customer Advisory Board Interviews:

To complement the quantitative research, Panterra Global conducted qualitative interviews with members of Optro’s Customer Advisory Board (CAB). Participants were executives and senior leaders with responsibility for areas including internal audit, risk management, compliance, and information security at large organizations.

The interviews explored how organizations are approaching agentic AI in practice, including where agents are being introduced into GRC workflows, the types of actions they are being permitted to take, the role of human judgment and oversight, and emerging challenges related to governance, measurement, accountability, and control.

The discussions were designed to provide real-world context for the quantitative findings and to capture how GRC leaders are thinking about the operational implications of agentic AI. Quotes included in this report have been anonymized to protect participant confidentiality and may be lightly edited for clarity and readability without changing their meaning.

About Optro

Optro (formerly AuditBoard) helps enterprises transform risk into opportunity through its AI-powered GRC Intelligence Platform. More than 50% of the Fortune 500 trust Optro to elevate how they approach audit, risk, and compliance in a new era of risk. Optro is top-rated by customers on G2 and was named a Leader in the 2025 Gartner® Magic Quadrant™ for Governance, Risk and Compliance (GRC) Tools, Assurance Leaders. To learn more, visit optro.ai.

