

Accountable Intelligence

What replaces span of control.

ENTRY № 02 · 21 May 2026 · 13 min read

The next decade will retire span of control as a serious unit of management. Not because hierarchy will end, or because reporting lines will dissolve, but because the unit no longer measures the thing that matters.

Span of control has been the quiet accounting system of management for the better part of a century. How many people report to you. How many seats sit beneath yours on the org chart. How much budget passes through your authority. It was never the whole story, but it was a reliable proxy for the size and complexity of a role, and senior leaders learned to use it intuitively. A general manager with two hundred staff was, by common understanding, doing a bigger job than one with twenty. Promotion paths, compensation bands and political weight inside organisations were all anchored, directly or indirectly, on this measure.

That proxy is breaking. The reason is not that AI is a more powerful tool than the tools we have had before, although it is. The reason is that AI changes what work is, who is doing it, and what kind of accountability a manager carries. When a portion of the work inside your team is done by software that drafts, analyses, decides and acts, the question “how many people report to you” begins to undercount what you are actually responsible for. It is the wrong unit.

I want to argue for the replacement. My proposal is that the new unit of management is accountable intelligence: the amount of work, done by people and AI together, that a single human can responsibly direct, verify and stand behind. The phrase is deliberate. It contains four ideas, work, mixed agency, verification and ownership, and it preserves what was valuable about span of control while updating what is no longer adequate.

A LEDGER PRINCIPLE

Accountable intelligence is the amount of work, done by people and AI together, that a single human can responsibly direct, verify and stand behind.

I structure the argument in four moves. First, why span of control is no longer fit for purpose. Second, what accountable intelligence is, and how to recognise it. Third, what concretely changes for managers when accountable intelligence becomes the operating measure. Fourth, why organisations cannot expect this shift to happen by drift.

The old unit, and why it is breaking

Span of control was invented for a workforce composed entirely of people. The early management theorists, V. A. Graicunas and Lyndall Urwick among them, were preoccupied with the limits of human supervision. How many subordinates could a manager realistically coordinate, given the geometric increase in relationships as the team grew. The answer, they suggested, was small, and the practical implication was that hierarchy must be deep rather than wide. Later writers softened the prescriptions, but the underlying picture, of management as the supervision of human acts, remained the silent assumption.

The model worked because work was, in fact, a chain of human acts. A manager could observe their team. They could check work by reading it, listening to it, watching a meeting, reviewing a report, sampling the output. Verification was part of presence. Risk was understood through experience and proximity. Judgement was carried by people in a hierarchy of people, and the hierarchy was the audit trail.


That picture is no longer accurate. Inside most organisations today, work is already produced by combinations of people and software in ways that span of control does not register. Some of this is benign and well-understood: copilots drafting documents, analytics tools generating dashboards, automated workflows routing transactions. Some of it is newer and less understood: general-purpose assistants synthesising information for executives, AI systems triaging customer enquiries, vendor software embedded inside enterprise platforms making decisions that nobody on the staff directly authorises. The accountable person sits above all of it, but the measure used to describe their job, the number of people who report to them, captures only one part of what they are responsible for.

The drift is already visible. Senior leaders in most organisations cannot answer simple questions about the AI in their own teams. How many AI-assisted decisions did your function make last month. What data is your team feeding into systems outside the perimeter. Which of your processes would fail if a single vendor's AI was unavailable for a week. These are not technical questions. They are the questions a manager used to be able to answer about their people. The shift is that managers now need to answer them about the work itself, regardless of who or what is producing it.

This is the gap that accountable intelligence is meant to close.

From Span of Control to Accountable Intelligence

What was implicit before must now become explicit



The new unit

The simplest way to define accountable intelligence is to anchor it on a single question. Instead of asking how many people a manager controls, ask how much work, done by people and AI together, that manager can responsibly direct, verify and stand behind.

Each word matters. *Work* rather than *people*, because the unit being measured is output, not headcount. *People and AI together*, because the modern manager almost always supervises both, and the share is moving steadily toward the second. *Responsibly direct*, because the manager is not merely allowing the work to happen but actively shaping it, setting limits, deciding what is in scope. *Verify*, because direction without verification is a fiction, and accountable intelligence cannot exist where the manager has no practical means of checking the output. *Stand behind*, because the test of accountability is the willingness, and the capacity, to take responsibility when something goes wrong.

This is not a metric. It is a lens. The point is not that organisations should now invent a scoring formula to replace headcount, although in time some will. The point is that the conversation about management should change. Promotion criteria, performance reviews, role design, governance and compensation should all begin to organise themselves around the new question.

Four practical questions make accountable intelligence visible. They are deliberately simple, because the goal is to make explicit what was implicit before, not to introduce a methodology.

The first is the question of scope. How much work, done by whom and what, sits inside the role. This is the natural successor to “how many direct reports do you have.” A manager today supervising fifteen analysts may also be supervising six automation workflows, four AI-assisted decision tools, two copilots used by every member of the team, and a vendor system that touches customer data. The honest answer to the question of scope is no longer a single number. It is a composition.


The second is the question of consequence. What happens if the work goes wrong. A team using AI to draft internal meeting notes is not in the same position as a team using AI to support a regulator-facing communication, or a customer-facing credit decision. The same tool can sit in radically different risk environments depending on the activity, the data, the decision and the consequence. Accountable intelligence treats consequence as part of the measure, not as an afterthought. A manager with high-consequence work, even at modest volume, carries more accountable intelligence than one with high-volume work of low consequence.

The third is the question of verification. Can the output be checked. This is the test that most quietly distinguishes real accountability from nominal accountability. A manager who can describe in technical detail what their AI system does, but who has no practical means of sampling its outputs and challenging them, does not in fact have accountable intelligence over that system. They have a label. Verification capacity is the difference between the two. It is observable through sampling rates, escalation paths, incident records and the manager’s own ability to recognise when something is wrong.

The fourth is the question of ownership. Can the accountable human stand behind the work. This is partly about understanding (does the manager grasp enough of the context to take responsibility), partly about authority (do they have the means to intervene), and partly about willingness (would they accept the consequence if the output caused harm). Ownership is the test that closes the loop. Without it, accountable intelligence collapses into responsibility-shopping, in which each layer of the organisation insists that the layer below, or the vendor, or the tool itself, owns the outcome.

The Four Questions of Accountable Intelligence

Together they make the implicit explicit



These four questions are not a checklist. They are the dimensions of a single judgement. Strength in one cannot substitute for absence in another. A manager who covers a wide scope without verification capacity is in a worse position than a manager with narrow scope and strong verification. A manager who can verify but does not understand the consequence of the work is in a worse position than one who understands the consequence even if their verification is sampled rather than complete. The point is the combination, not the parts.

A practical illustration helps. Consider three managers in different parts of the same organisation, each with twenty analysts reporting to them.

The first manager runs a regulatory reporting team. Each analyst produces structured outputs that follow defined templates. AI is used inside the workflow for drafting and consistency checking, but every output is reviewed and signed by a named analyst before submission. The manager has a clear line of sight into what the team produces, regular sampling, and the technical knowledge to challenge the outputs if needed. Their accountable intelligence is high, narrow and well-supervised.

The second manager runs a customer operations team. The same twenty analysts work alongside an AI system that triages incoming enquiries, generates draft responses, and routes a meaningful proportion of cases without human intervention. The manager can see aggregate metrics but does not, in practice, review individual interactions. The work has expanded. The accountable intelligence has also expanded, but it sits partly inside the AI system itself, and the manager's verification capacity has not grown with the scope. The same headcount conceals a larger risk profile.

The third manager runs a strategic analysis team. The twenty analysts use general-purpose AI tools to synthesise information, often from public sources, often combined with confidential internal data. The outputs feed into executive decision-making. The manager has no formal AI workflow inside the team, no defined templates, and no systematic verification process. Each analyst is, in effect, running their own micro-process. The accountable intelligence is theoretical. In practice, the manager is signing off on outputs they have no realistic means of verifying.

By the standard measure of span of control, these three jobs look identical. Twenty reports, similar structures, similar pay grades. By the measure of accountable intelligence, they are radically different. The second and third managers are carrying more work and more risk than the first, but neither the org chart nor the compensation system reflects it.

What changes for managers

The shift from span of control to accountable intelligence is not a paperwork change. It changes what managers actually do.

The first change is in the framing of work. A manager organised around span of control thinks in terms of what their people are doing. A manager organised around accountable intelligence thinks in terms of what is being produced, by what combination of people and tools, with what risk and what verification. The unit of attention moves from staff to work. The day-to-day vocabulary changes. Conversations in the team move from “who is working on this” to “what is producing this output, and what would we need to see to trust it.”

The second change is in delegation. Delegation in a human-only organisation is delegation to a named person. The manager retains accountability but transfers execution. Delegation in a human-AI organisation is more complex, because part of the execution is now delegated to systems that do not have judgement, do not learn from rebuke, and do not voluntarily escalate when things become uncertain. A manager who delegates to an AI system the way they delegate to a person is making a category error. AI systems require limits, structured oversight, defined escalation rules and explicit consequence handling. The manager who delegates well will write these into the work itself, not into the employee handbook.

The third change is in verification. In a human-only organisation, verification was often informal. The manager read the report, sat in the meeting, talked to the customer, glanced at the dashboard. In a human-AI organisation, informal verification breaks down quickly, because the work is faster, more uniform, harder to taste-test by eye, and easier to produce confidently while being subtly wrong. Verification has to be designed.

Sampling regimes, challenge sessions, paired reviews, structured second opinions and adversarial checking become part of the operating model. This is not a tax. It is the new shape of accountability.

The fourth change is in risk literacy. A manager in the old model could often rely on experience to recognise risk. The categories were familiar. The new categories of risk, hallucination, prompt manipulation, model drift, training-data leakage, are unfamiliar to most senior leaders. The manager does not need to become a technologist. They do need to develop enough literacy to ask the right questions and recognise warning signs. The accountable person who does not understand how their system can fail is, in the limit, not the accountable person.

The fifth change is in performance measurement. The old metrics tracked activity, output, headcount productivity. They were already imperfect. They become misleading in a human-AI organisation, because output can be increased by quietly outsourcing more of the work to systems whose risk is not visible. Performance measurement in the new model has to track quality, verified output, risk events and adoption, not just throughput. Managers who optimise for throughput at the expense of verification will eventually deliver consequence rather than performance.

What the Manager Tracks

From counting acts to seeing the work



None of these changes is exotic. They are extensions of practices that the best managers have always followed, made explicit and made systematic. The reason they need to be made explicit now is that the speed and scale of AI-enabled work have

outstripped the implicit forms of supervision that previously sufficed. The kind of judgement that used to live inside a manager's head, often without that manager being able to articulate it, now has to be written down, taught, audited and renewed.

Why this will not happen by drift

There is a comfortable view, popular in some circles, that the shift to AI-enabled work will resolve itself through normal management evolution. Individual teams will adopt tools, managers will adapt, and over time the new equilibrium will emerge without the need for deliberate intervention. The evidence so far does not support this view.

The default trajectory in most organisations today is not adaptive equilibrium. It is one of two failure modes.

The first is fragmentation. AI is adopted in pockets, with each team developing its own practices, its own risk tolerance, its own undocumented dependencies on tools that nobody is auditing. Visibility falls. Accountability becomes diffuse. When something goes wrong, the organisation discovers that nobody can answer the basic questions about who was supervising what.

The second failure mode is over-control. Reacting against fragmentation, the organisation imposes a heavy approval process. Every AI use must be reviewed. Every tool must be registered. Every workflow must be documented. The approval system becomes the work, and the actual adoption of AI moves underground, where it can no longer be observed at all. This is what happens when the response to a new form of work is to apply the supervision system of the old form.

Neither outcome produces accountable intelligence. Both produce its absence, dressed up as either freedom or control.

The way through is to make the shift deliberate. This does not require a vast new bureaucracy. It requires a small, senior, time-bound transition structure whose purpose is to do five things: map where AI-enabled work is already happening, redesign work and roles so that accountability is explicit, deploy priority use cases with proper governance, govern risk in proportion to the work rather than the tool, and embed the resulting capability into normal management systems. The structure is a means, not an end. It exists to dissolve once the new operating system is in place.

I am not writing about that structure here. It is the practical answer to a question that has not yet been widely asked. The question is the more important contribution. What unit are we now measuring management by, and what does it take to manage well in those units.

A closing claim

A senior leader in five years, I expect, will not describe their job the way they describe it today. The phrase “I run a team of X people” will start to sound the way “I run a typing pool” sounds now: accurate as far as it goes, but no longer the relevant description of what the role contains. The replacement phrase has not yet settled into common use, but the underlying concept can already be named. Accountable intelligence is what a leader stands behind. It is the work, including the work done by systems, that bears their signature. It is the new test of management.

My argument reduces to a single sentence. The next decade will not be defined by how much AI an organisation deploys, but by how much accountable intelligence its leaders are willing and able to carry. Those that recognise the shift early will design their roles, their teams and their governance around the new unit, and they will be running a different kind of organisation. Those that do not will discover, slowly and then suddenly, that their measure of management has stopped measuring the thing that matters.

WRITTEN BY

Veemal Gungadin

WITH ASSISTANCE FROM **Claude Opus 4.7** · Anthropic , **GPT-5.5** · OpenAI

Drafted and revised by the author. AI tools were used to sharpen structure and prose; conclusions are the author's own.

Published in *The Goodwill Ledger* at www.goodwill-led.com/essays/accountable-intelligence/