

From Job Titles to Capability

Why the Way You Describe Work Is
Holding Your Organisation Back

A workforce strategy perspective for HR leaders

Sadiq Merali

[Talent Strategy Partner](#)

Chris Illingworth

[Principal Solutions Consultant](#)



Executive Summary

Most organisations have two blind spots in how they think about talent, and the two are closely related. One is leaning too hard on inbound applications to fill roles that actually matter. The other is describing the workforce almost entirely in job titles, when titles were built to manage payroll and hierarchy, not to describe what people can actually do.

Together, these create a kind of optical illusion. Applicant volume looks like pipeline health. An org chart looks like workforce insight. Neither holds up once labour markets tighten, skills fragment, and AI starts rewriting what a role actually requires faster than anyone can update the job description for it.

None of this is really about hiring volume or data hygiene, even though it usually gets discussed that way. It's a strategy question, and increasingly a transformation one. What follows makes the case for building workforce strategy around observable capability instead of surface signals, treating that capability data as infrastructure rather than an HR side project, and rolling it out in a deliberate order rather than buying it as a single initiative.


the predictable outcome of making capability decisions with title-level data.



Part One: The Illusion of Progress

Applicant volumes have climbed for a few years now, and on paper that looks like progress. In practice it's done close to the opposite: buried hiring teams under noise while leaving them less room to operate strategically. Most applicant-led hiring is built for speed, volume, and low cost per hire, not for scarce specialist roles or decisions with real long-term stakes. So organisations end up hiring for whoever shows up, not for what the business actually needs.

The numbers back this up: global job openings fell 14% in late 2024 while applications rose 8%, and the World Economic Forum's Future of Jobs Report 2025 puts the share of core skills likely to be disrupted within five years at 44%.



**of core skills likely to be disrupted
within five years.**


(World Economic Forum, 2025)

Part Two: The Same Blind Spot, Twice

Job titles have the same problem in a different shape: a title tells you where someone sits on an org chart, not what they actually do or what skills that work takes today. That gap between title and capability isn't new; titles were never designed to capture it. What's new is how expensive the gap has become. AI is changing the substance of jobs faster than job descriptions get rewritten, so the distance between the label on a role and what it actually requires keeps growing. The World Economic Forum expects 39% of workers' core skills to be transformed or made obsolete by 2030, and McKinsey estimates close to 30% of hours worked across the US economy could be automated by that same point, driven largely by generative AI.



39% of workers' core skills to be transformed or made obsolete by 2030.



30% of hours worked across the US economy could be automated by 2030.

The title hasn't moved. The work underneath it has. Most organisations have no real way to track that, because whatever skills data they hold, if they hold any, reflects the role as it was defined when someone built the framework, not as it exists now. Applicant-led hiring and title-based workforce data are really the same failure showing up twice: decisions made on what's easy to see, rather than on what's true.



Part Three: The Real Cost of Getting It Wrong

Getting this wrong is not cheap, and the numbers hold up across every major study of the subject:

30%

of first-year salary, the US Department of Labour's baseline estimate for the cost of a single bad hire.

\$240k

Up to \$240,000 in total cost once lost productivity and replacement are factored in, per SHRM.

50-200%

of annual salary in pure replacement cost alone, depending on seniority.

30%

A roughly 30% morale hit to the surrounding team, by one widely cited estimate, when a bad hire lands in a critical role.

None of that is a fringe scenario. It's what happens, predictably, when capability decisions get made on partial information (a title, an application, whoever happened to apply that quarter) at exactly the moments when getting it wrong costs the most.



Part Four: The Talent Already Inside the Organisation

A good chunk of this risk is self-inflicted: most organisations sit on more internal capability than they use. Companies that do internal mobility well retain people at close to twice the rate of those that don't, and internal hires reach full competency 20% faster while costing 18-20% less to place. External hiring remains the default anyway, and external hires are 61% more likely to be terminated than internal ones, a gap that gets more expensive whenever the market tightens.

Part of why internal mobility stays underused is that nobody can see it, not that nobody wants it. Gartner has found that a third of employees think they'd have more impact in a different role at their own company, but most organisations can't act on that because they lack the data to know who's a fit for what. Gartner's other finding is blunter: only 8% of organisations have skills data on their current workforce they would call reliable.

That's the gap this paper is really about: not a shortage of talent, but a shortage of visibility into the talent already on the payroll.

Part Five: What a Skills Layer Unlocks

A skills layer doesn't replace the org chart. It sits underneath it, closer to an X-ray of what's actually there once you look past titles and reporting lines.

- Pay decisions get easier to defend, because you can point to what a role actually requires and what those skills go for in the market.
- Performance conversations get more concrete: progress ties to specific skills at specific levels, not a manager's impression.
- Hiring gets sharper, because a vacancy reflects skills the business needs today, not a description someone wrote three years ago.
- Internal mobility surfaces ahead of time, instead of getting discovered after someone has already quit.
- AI readiness stops being a guess: which parts of the workforce are ready for augmentation, which need reskilling, and where the biggest gaps sit.

Part Six: Why Tasks Solve the Credibility Problem

There's a fair objection to skills-based approaches: skills are easy to claim and hard to prove. Ask people to self-report and you get a mix of what they can do and what they'd like to believe they can do. Frameworks built purely on survey data age badly, and once the numbers start looking off, the leaders who were supposed to act on them stop trusting them. That's usually where these projects die: not at launch, when the data looks comprehensive, but twelve months later, when nobody quite believes it any more.

The fix isn't to give up on skills. It's to get there through tasks instead of surveys. A task is something you can point to and observe; it's the unit of work a skill produces. Demonstrate the task and you've effectively evidenced the skill: the task is the proof, the skill is what gets inferred from it. Take a skill like communication: instead of a survey where someone rates themselves "advanced," the task-based version looks like leading three cross-functional design reviews, documenting the decisions, and socialising them at the next all-hands — something a manager can point to and verify

happened, not a word chosen off a scale. That shift, from opinion to observation, is what makes the data credible to finance and operations leaders who need to trust it before they'll act on it.

It also changes how the framework gets maintained. Roles change because the tasks inside them change first, so a task-based framework updates as the work shifts instead of waiting for an annual review cycle. That's the difference between something that stays alive and something that expires the day it's finished.

This is also why task-based frameworks tend to survive leadership changes, budget reviews, and reorgs in a way that survey-based ones don't. A new CHRO or CFO can question a spreadsheet built on self-report. It's much harder to argue with a framework built on what people were actually observed producing, month after month, because the evidence trail doesn't depend on any one person's judgement staying in the room.

Part Seven:

Why This Belongs on the Transformation Agenda

Capability data doesn't behave like a recruiting project or a single system rollout. It behaves more like finance infrastructure: it changes how Finance models labour cost and risk, what IT has to integrate and govern across HR and planning systems, how business leaders plan capacity, and where the executive team decides to build, buy, or reskill.

Projects built entirely inside HR rarely survive that kind of scope. One function builds it, everyone else adopts it half-heartedly, and it falls out of date because nobody outside HR is on the hook for keeping it current. The programmes that stick tend to have sponsorship that sits above any single function, a steering group that spans HR, Finance, and IT, and a fixed spot in the planning calendar that already exists, rather than being asked to justify themselves as a standalone HR initiative every budget cycle.

Put simply, moving from titles to skills changes how the organisation operates, not just what data it holds. Treat it as a cleanup exercise and it will behave like one: it gets finished, and then it gets ignored.

This is also why the business case has to be framed in terms Finance and the executive team already track, rather than in HR-specific language. A capability programme pitched as “better talent data” competes for budget against every other HR initiative. The same programme pitched as “the input that makes workforce planning and reskilling spend defensible” competes in a different, better-funded conversation.



Part Eight: What This Looks Like in Practice

Picture a workforce leadership team in a board meeting, faced with a fairly ordinary question these days: how AI-ready is the workforce, and where are the biggest gaps?

Without a skills layer, what's on the table is headcount and titles: how many people, in which function, at what level. What's missing is which roles are most exposed to AI-driven task substitution, where a reskilling pound would do the most good, or whether the organisation can deliver its strategy with the people it has. The honest answer is "we don't know," and that answer is wearing thin, because workforce capability now sits close enough to capital allocation that boards are starting to treat it that way.

THE SAME MEETING, WITH THE DATA

With a skills layer, the same meeting goes differently. It becomes possible to show which functions are ready for AI augmentation and which aren't, flag the roles most exposed to substitution and quantify what reskilling would take, and point to where internal mobility could close a gap without going external. The board gets something to act on instead of something to defer, and Finance, IT, and the business finally work off the same numbers instead of three different versions of the workforce.



WHAT CHANGES FOR FINANCE AND IT

It also changes who else can use the same data. Finance stops treating workforce cost as a fixed number and starts modelling it against actual capability. IT gets a single source of truth to govern instead of three competing spreadsheets. And the business units stop lobbying for headcount in the abstract and start making the case in terms of the specific skills a plan requires, which is a very different, and much more fundable, conversation.

The real difference between those two meetings isn't the technology in the room. It's whether the data is current, tied to actual work, and plugged into the systems where decisions get made.



Part Nine: Why Traditional Approaches Have Failed

The usual route into a skills framework has been consultancy-led: six to twelve months of interviews and analysis, handed over as a spreadsheet, at a cost that's hard to ignore. Anyone who has been through this recognises the pattern:

- The framework is out of date roughly the moment it's delivered.
- A spreadsheet doesn't update itself when roles change. It's a photograph, not a system.
- Data that lives outside core HR systems can't feed the platforms where decisions actually happen.
- Finance struggles to justify the spend, because a skills framework in a spreadsheet doesn't look like infrastructure. It looks like a one-time HR project.

These aren't bugs in execution. They're built into the model. A consultancy delivers a document; it doesn't maintain a living asset. The organisations getting this right treat skills data the way they treat financial reporting: something that has to stay current, stay auditable, and plug into the systems where decisions get made.

The test worth applying to any skills initiative is simple: if the people who built it left tomorrow, would the data still update itself? If the honest answer is no, it was never really infrastructure, whatever it was called on the original business case.

Part Ten: A Maturity Model for Workforce Capability

Most organisations can place themselves on this spectrum without much argument. Knowing where you actually sit is more useful than any single new data project, because it tells you what the next investment should be for.

Stage	What It Looks Like
1. Ad hoc	Capability knowledge lives in managers' heads and old CVs. There's no shared taxonomy, so decisions get made on title and tenure.
2. Documented	A skills taxonomy exists somewhere, usually because a consultancy built one. It's static and lives outside the core HR systems.
3. Self-reported	People tag their own skills inside core systems. The data exists, but nobody fully trusts it, and it goes stale within months.
4. Verified	Skills get inferred from actual tasks and work product rather than self-report. The data is solid enough to inform pay, hiring, and mobility calls.
5. Predictive	Capability data feeds directly into workforce and financial planning. The organisation can model the workforce it will need, not just describe the one it has.

The jump that actually matters is from stage three to stage four: self-reported to verified. It's also the one most programmes skip, because it means changing how the data gets generated, not just where it gets stored.

Most organisations we talk to place themselves at stage two or three, comfortable that a taxonomy exists somewhere even if nobody fully trusts it. Very few make it to stage four without treating the move as a deliberate, resourced decision rather than something that happens on its own once the taxonomy is in place.

Part Eleven: Ten Reasons the Old Model No Longer Works



The best people usually aren't applying anywhere. High performers tend to be passive candidates who need to be engaged, not job ads waiting to be found.



More applications doesn't mean better ones. Volumes are up 8% even as roles have gotten scarcer, so the funnel is just noisier, not stronger.



Job ads describe a role that's already out of date. With 44% of core skills expected to shift by 2030, a posting can be stale before it's even live.



Skills predict performance better than a CV does, roughly five times better by some estimates than hiring on credentials alone.



The roles that matter most get treated like the roles that matter least, even though a small slice of them create most of the value.



Internal talent sits there unused, despite being both cheaper and faster to bring up to speed than an external hire.



Hiring runs on whatever requisition just opened, not on where the business is actually headed.



Every time someone leaves rather than moves internally, institutional knowledge leaves with them.



Growth plans stall because nobody had the talent insight early enough to act on it.



Speed keeps winning over value, which is exactly the wrong trade-off for the roles that matter most.

None of these are new observations in isolation. What's changed is that they now compound faster than they used to, because AI is shortening the shelf life of job descriptions and the skills behind them at the same time.

Part Twelve: From Pilot to Infrastructure

Transformation programmes in this space tend to succeed or fail on sequencing, not ambition. A few things show up consistently in the ones that stick:

1. **Start with one function, not the whole company.** Pick something business-critical where the stakes are obvious, ideally a function already being reshaped by AI, and prove the approach there before trying to scale it.
2. **Sort out governance before collecting more data.** Decide who owns the taxonomy, how often it refreshes, and who's actually allowed to act on it. Otherwise it's just another survey nobody trusts.
3. **Plug it into existing systems early.** Feed the data into wherever hiring, pay, and mobility decisions already happen, instead of building a separate dashboard that quietly stops getting checked.
4. **Talk about it in numbers Finance already cares about.** Time-to-fill for critical roles, internal fill rate, return on reskilling spend, not how many people completed the survey.
5. **Expect pushback, because it's coming.** Skills-based approaches take some control away from managers who currently gatekeep opportunity through who they happen to know. That's a change-management problem as much as a technology one.

Organisations that work through all five usually find the same thing on the other side: the programme stops being something HR has to defend every budget cycle, and starts being something Finance, IT, and the business units ask to expand into their own planning.



Part Thirteen:

A Different Way Forward

The organisations pulling ahead have stopped asking how to get more applicants and started asking whether they actually know where their best talent will come from before they need it. Getting to a real answer means doing three things: getting specific about which skills genuinely matter to future performance rather than which ones sound right on a job spec, building talent pools ahead of need so there's already a shortlist when the market shifts, and hiring for the gaps that are emerging instead of the ones that were true two years ago.

The same logic applies to the people already on the payroll. Org charts and job architecture are still useful for structure, for governance, for knowing who reports to whom, but they were never built to answer questions about capability, and using them that way is a mistake most organisations have already made at least once. The ones pulling ahead are building a skills layer underneath the org chart they already have: grounded in what people actually do rather than what they claim to do, run as a shared programme rather than an HR side project, and wired into the systems where hiring, pay, mobility, and planning decisions actually get made.



Closing Thought

Inbound hiring isn't going away, and neither are job titles. They're still useful for structure and for running payroll. The problem is treating either one as the main source of truth about the workforce. That's a bet on short-term signals and administrative convenience, and it's a bet more organisations are starting to lose.

The better questions aren't "how fast can we process applicants" or "what does the org chart say." They're simpler than that: do we know where our best people are actually going to come from, and do we know what the people we already have can really do?

CONTINUE THE CONVERSATION

If any of this sparked questions about your own organisation, there are two easy ways to keep going.

JOIN A ROUNDTABLE

We regularly bring small groups of HR leaders together to work through exactly this, both in person and over video call. It's an informal, practitioner-led conversation rather than a pitch. If you'd like an invite to the next one, let us know and we'll add you to the list.

GET A COMPLIMENTARY SKILLS ASSESSMENT

If you'd rather start with your own organisation's numbers, we offer a complimentary skills assessment: a short working session with our team to map out where you actually sit on the maturity curve set out in this paper, and where the biggest gaps are. Let us know if you'd like to do this and we'll get some time in the diary.

ABOUT THE AUTHORS



Sadiq Merali,
Talent Strategy Partner

Sadiq works with HR and talent leaders to close skills gaps and build workforce efficiency. He previously worked at PwC.



Chris Illingworth,
Principal Solutions Consultant

Chris works with clients to design technology solutions and drive system implementations. He previously worked at Workday.



Sources

- Statistics and research referenced throughout this paper are drawn from the following sources.
- World Economic Forum, The Future of Jobs Report 2025, January 2025 (skills disruption statistics).
- McKinsey and Company, Talent-to-Value research (5% of roles creating 95% of value); Generative AI and the Future of Work in America, 2023; A New Future of Work: The Race to Deploy AI and Raise Skills, 2024.
- Society for Human Resource Management (SHRM), The Real Cost of a Bad Hire; cost and cultural-impact research on poor hiring decisions.
- US Department of Labour, cost of a bad hire estimate.
- Gartner, HR Research: Talent Management Efforts Inhibit Optimal Employee and Organizational Performance, September 2024.
- LinkedIn, Future of Recruiting 2025: application volume trends and skills-based hiring predictiveness.
- LinkedIn Learning, internal vs. external hiring statistics (retention, cost, speed to competency).

THIS DOCUMENT IS FOR INFORMATIONAL PURPOSES ONLY. STATISTICS AND FIGURES CITED ARE DRAWN FROM THIRD-PARTY RESEARCH AND SHOULD BE INDEPENDENTLY VERIFIED BEFORE USE IN DECISION-MAKING. SEVERAL OF THE STUDIES REFERENCED HERE ARE UPDATED ANNUALLY, SO ANYONE USING THESE FIGURES IN A CUSTOMER-FACING OR BOARD SETTING SHOULD CHECK THE CURRENT PUBLISHED VALUES AGAINST THE PRIMARY SOURCE RATHER THAN CITING THIS PAPER DIRECTLY. © BEAMERY 2026.
