WHITE PAPER

AI, Workforce and the Operating Model

What boards and executive teams need to understand about AI-driven workforce transformation

Marco Ciobo
Principal, Digital Insights and Transformation | Board and C-Suite Advisor
June 2026

Executive Summary

Based on my advisory work with boards and executive teams across the ANZ region, I am consistently asked the same question: what does artificial intelligence actually mean for our workforce, our organisation and our operating model?

This white paper sets out my observations from the field. It is not an academic treatise. It is a practitioner's account of the frameworks that matter, the research you should know, and the strategic conclusions I believe boards and executive leadership teams need to reach, ideally before their next planning cycle.

The central argument is straightforward. AI and automation do not primarily eliminate jobs. They eliminate tasks, reshape workflows, and fundamentally alter the economics of how work is organised. Organisations that understand this at the task and workflow level will make better investment decisions, better workforce decisions, and better operating model decisions than those still thinking in terms of headcount and occupational categories.

> ***The strategic imperative has shifted from headcount optimisation to capability portfolio management. The organisations that navigate this well will be the ones that understand which uniquely human capabilities they need, where those capabilities currently sit, and how to develop and deploy them.***

Generative AI, which entered mainstream enterprise usage from 2022 onwards, has materially expanded the automation frontier into cognitive and knowledge-intensive work. This is not a refinement of earlier automation. It is a step change that requires boards and executive teams to revisit prior assumptions and take fresh stock of their workforce exposure and capability strategy.

1. Why the Debate Has Been Framed Incorrectly

The public and boardroom narrative around AI and workforce has been dominated, and in my view distorted, by a single headline. In 2013, Oxford University researchers Carl Benedikt Frey and Michael Osborne published a study estimating that approximately 47 percent of United States employment was at high risk of computerisation within 10 to 20 years. That number travelled everywhere and shaped a decade of enterprise AI strategy, workforce planning, and public policy.

The problem is not that the study was wrong in spirit. It was genuinely pioneering work that established a rigorous analytical framework for thinking about automation exposure. The problem is how the headline figure was applied. Boards and executive teams frequently treated it as a prediction of job losses rather than an exposure mapping exercise, and that led to a series of strategic errors that I still encounter regularly in my advisory work.

The most important corrective insight is this: automation does not replace occupations uniformly. It replaces tasks, and it does so unevenly. A customer service manager might have 30 percent of her daily tasks automated by AI tools while the remaining 70 percent, judgment calls, relationship management, escalation handling, context-sensitive decision making, remain firmly human. The occupation survives. The role evolves. The skill requirements shift.

> ***Jobs are not singular entities. They are bundles of tasks. AI and automation affect tasks unevenly. This is the foundational analytical insight that every board and executive team needs to internalise.***

Organisations that built their AI and automation strategies primarily around headcount reduction missed the more significant opportunity: deploying AI to augment human productivity, redesign workflows, and release human capability for higher-value work that machines cannot easily perform.

2. The Analytical Framework That Actually Works

The most useful analytical tool for understanding AI's impact on workforce comes from an unlikely source: the United States Department of Labor's Occupational Information Network, known as O*NET. Originally developed as a career guidance resource, O*NET became the foundational occupational taxonomy used by governments, researchers, and major consulting firms to decompose work at the task level.

O*NET structures each occupation across seven dimensions: the specific tasks performed, the skills required, the knowledge domains involved, the underlying human abilities drawn upon, the typical work activities, the work context, and the education and training requirements. This multi-dimensional structure is what makes it analytically powerful.

When you apply O*NET's task-level decomposition to AI impact assessment, you move from asking 'will this job be automated?' to asking much more operationally specific and useful questions:

- Which tasks within this role involve structured, rules-based processing that AI can handle?
- Which tasks require contextual judgment, ambiguity handling, or relationship intelligence that remain firmly human?
- Where does augmentation create the greatest productivity opportunity?
- How should roles be redesigned as the human-machine boundary shifts?
- What are the operating model and organisational design implications?

The consulting firms that used O*NET-based analysis most effectively during the 2015 to 2023 period, and BCG, Bain, Accenture, Deloitte and McKinsey all developed significant capability in this area, were able to translate abstract AI capability discussions into specific, measurable assessments of workforce exposure and operating model opportunity. That is still the right approach today.

Bain's approach is instructive here. As a global advisor to Bain, I have seen first-hand how the firm has positioned AI not as a replacement for consulting judgment but as a force multiplier for human expertise. Bain's strategic partnership with OpenAI, one of the earlier and more substantive partnerships between a major consulting firm and a leading AI developer, was built on a deliberate philosophy: keep human expertise in the loop. The firm's AI playbook is grounded in the view that generative AI accelerates the analytical and research-intensive dimensions of consulting work, but that the judgment, synthesis, client relationship and change leadership dimensions remain firmly human. That is precisely the augmentation dynamic I am describing in this paper, and it plays out identically inside enterprise organisations.

What has changed is the automation frontier. Earlier O*NET-based assessments focused on repetitive manual work, structured administrative processing, and rules-based transactions. Generative AI has moved that frontier substantially up the stack into cognitive work:

analysis, synthesis, writing, coding, decision support, and knowledge-intensive workflows that were previously assumed to be resistant to automation.

> ***Modern generative AI exposure analysis applied to O*NET occupational structures finds a substantially different pattern to earlier automation studies, with a far greater share of high-wage, high-education roles now facing meaningful exposure to AI augmentation or substitution.***

3. The Augmentation Reality

The dominant near-term pattern of AI deployment in enterprise is not replacement. It is augmentation. This distinction matters enormously for how boards and executive teams should be thinking about their AI investments and workforce strategies.

Across the organisations I work with, the AI implementations that are generating the most measurable value are operating as copilots, workflow accelerators, and decision-support systems. They compress the time required to perform specific cognitive tasks, improve the quality of outputs, and extend the productive bandwidth of skilled professionals. They do not, in most cases, eliminate the human professional.

BCG's research on AI deployment speed has highlighted a risk I see regularly in my advisory work: organisations are deploying AI capabilities without corresponding adjustments to role design, skill development, or operating model architecture. The AI investment lands. The workflow changes. But the organisational design, the role architecture, the supervisory structure, and the skill development agenda have not kept pace. Value leaks out through the gap.

The organisations that are capturing the most value from AI are those that have thought carefully about human-machine workflow design, not just AI capability deployment. They are asking: given what AI can now do, how should we redesign this workflow end to end? Where does human judgment add the most value? How do we train our people to work effectively alongside AI systems? What does the manager's role look like when a substantial share of analytical work is AI-assisted?

> ***The organisations capturing the most value from AI are not those who deployed it fastest. They are those who redesigned their workflows, roles and operating models around it most thoughtfully.***

4. The Human Premium

As AI capability extends across a wider range of cognitive tasks, the human capabilities that remain difficult to automate become more strategically valuable, not less. This is an insight that does not always land intuitively with boards and executive teams, but it follows directly from the logic of comparative advantage.

The capabilities that are genuinely difficult for AI systems to replicate include:

- Contextual judgment: the ability to read a complex, ambiguous situation and make a well-calibrated decision with incomplete information
- Ethical reasoning: the capacity to weigh competing values, identify unintended consequences, and navigate moral complexity
- Relationship intelligence: the subtlety of human interaction, trust-building, empathy and social perception
- Creative synthesis: the generation of genuinely novel ideas and approaches that draw on broad, contextual understanding
- Leadership: the ability to align people behind a direction, manage through uncertainty, and create the conditions for others to perform

These capabilities are not soft. They are the capabilities that determine outcomes in board rooms, client relationships, complex negotiations, and high-stakes operational decisions. The irony of the current AI moment is that these deeply human capabilities are becoming more valuable precisely because AI is absorbing more of the analytical and transactional work that previously surrounded them.

The strategic implication is that workforce planning needs to shift its centre of gravity from headcount optimisation toward capability portfolio management. The questions that matter are: which human capabilities do we need to compete? Where do those capabilities currently sit in our organisation? Are we developing them, retaining them, and deploying them effectively? What does our capability development agenda look like for the next three to five years?

5. Operating Model Implications

AI-driven workforce transformation does not stop at the individual role level. It propagates through the operating model in ways that boards and executive teams need to understand and govern.

Span of control assumptions change when AI systems handle a substantial share of analytical, monitoring, and reporting work. A manager who previously needed a team of six to synthesise weekly performance data, draft reports, and maintain dashboards may now need a team of three. The supervisory ratio shifts. The organisational layer structure changes.

Shared services economics change materially when transactional and analytical processing is AI-assisted. The case for large centralised processing teams weakens as AI platforms absorb volume. The residual human function in shared services shifts toward exception handling, quality assurance, and relationship management.

Governance requirements increase, not decrease, as AI systems assume a greater share of analytical and decision-support work. Boards need to understand what decisions are being influenced by AI systems, what the risk and bias characteristics of those systems are, and what human oversight mechanisms are in place. This is not a technology risk question. It is a governance question.

> ***AI does not reduce governance requirements. It changes them. Boards need clear sight of which decisions are AI-influenced, what the oversight mechanisms are, and where human accountability sits.***

Role architecture requires revisiting across most functions when an organisation makes substantive AI investments. The tasks that justified certain roles change. The skill profiles required in residual human roles change. The interfaces between human and AI-assisted work need to be designed deliberately rather than allowed to evolve ad hoc.

In regulated sectors, these operating model implications carry additional weight. Telecommunications, utilities, financial services, and government organisations face constraints that do not apply in the same way to unregulated industries. Regulatory obligations around data handling, decision accountability, explainability, and auditability shape what is permissible in AI-assisted workflows, not just what is technically possible. Operating model redesign in these environments must be developed with regulatory compliance built in from the outset, not retrofitted after the fact.

6. Data Architecture, Governance and Risk: The Prerequisites That Determine Everything

Much of the boardroom and executive conversation about AI and workforce focuses on capability, deployment, and organisational design. These are the right topics. But in my experience working across telco, utilities, banking, and government organisations, the conversation often skips past a more fundamental set of questions that ultimately determine whether any AI-driven workforce or operating model strategy can actually be delivered.

Those questions are about data architecture and its governance. Before an organisation can deploy AI to automate tasks, augment workflows, or redesign operating models, it needs to be able to answer a basic set of questions about its own data: Where does it live? Who owns it? Is it accurate, consistent, and current? Is it structured in a way that AI systems can use? And who is accountable when it is not?

In the organisations I work with, the answer to these questions is frequently less reassuring than boards and executive teams assume. Data is fragmented across legacy systems and platforms. Ownership is contested or unclear. Quality is inconsistent. Lineage is poorly documented. Governance frameworks exist on paper but are not operationally embedded. These are not edge cases. They are the norm across large, complex enterprises in regulated sectors, and they represent the primary constraint on the pace and depth of AI adoption.

> ***AI capability is advancing faster than data infrastructure in most of the organisations I advise. The gap between what AI systems can theoretically do and what enterprise data environments can actually support is the single most common source of underdelivered AI value.***

Data Architecture as Strategic Infrastructure

Workforce transformation powered by AI requires high-quality, well-governed data at every layer. Automation tools need reliable, consistently structured input data to function accurately. Augmentation platforms need rich, contextually complete data to generate outputs that human professionals can trust and act on. Decision-support systems need data lineage and provenance that is auditable when decisions are questioned, and in regulated sectors they will be questioned.

This means data architecture is not an IT consideration. It is a strategic infrastructure question with direct implications for AI investment returns, operating model viability, and regulatory compliance. Boards and executive teams that treat data architecture as a technology sub-agenda item are systematically underinvesting in the foundation that their AI strategies depend on.

The key architectural requirements I assess in client engagements include data integration across fragmented legacy environments, clear data ownership and stewardship models, data quality frameworks with measurable standards and accountability, metadata management and lineage tracking, and the governance structures that ensure all of the above are maintained over time rather than eroded by the pace of operational change.

Governance: The Accountability Layer

Data governance in the AI era requires a level of rigour that goes beyond what most organisations have historically applied. When a human analyst produces a report, the accountability chain is relatively clear. When an AI system produces the same analysis, the accountability chain becomes more complex and, in a regulatory context, considerably more consequential.

Effective data governance for AI-driven operations needs to address several questions that boards should be asking explicitly. Who is accountable for the accuracy and currency of the data that AI systems use? What are the quality standards, and how are they enforced? How is data lineage documented so that AI-generated outputs can be traced back to their inputs if a decision is challenged? What processes exist to identify, escalate, and remediate data quality failures before they propagate through AI-assisted workflows?

These are governance questions, not technology questions. The technology can be configured to surface quality issues and maintain lineage records. But the accountability structures, the ownership models, and the escalation pathways are organisational design decisions that need board-level endorsement and executive sponsorship to be effective.

> ***Data governance is not a data team responsibility. In the AI era it is a board-level accountability. The organisations that treat it as a technology sub-problem are accumulating risk that will surface, often at the worst possible moment.***

Risk Management: The Stakes Are Higher in Regulated Sectors

For the sectors I work in most intensively, these considerations connect directly and immediately to risk management frameworks. Telecommunications, utilities, financial services, and government organisations operate under regulatory regimes that impose specific obligations around data handling, system accountability, decision auditability, and the explainability of automated outputs. These obligations are not hypothetical. They are enforced, and the consequences of non-compliance are material.

In financial services, obligations around model risk management, explainability of credit decisions, and the auditability of AI-influenced outcomes are well established and tightening. In telecommunications and utilities, sector regulators are increasingly attentive to the use of AI in customer-facing processes, network management decisions, and the handling of sensitive operational data. In government, public accountability obligations and freedom of information requirements create a distinct set of constraints on what AI-assisted decision-making is permissible without specific governance controls in place.

The risk management implication is direct: AI deployment in regulated sectors cannot be treated as a pure capability and productivity question. It must be developed within a risk and compliance envelope that is defined in advance, maintained operationally, and governed at a level commensurate with the regulatory stakes involved. Organisations that deploy AI in these environments without that envelope in place are not just creating operational risk. They are creating regulatory risk, reputational risk, and in some cases legal risk.

From my advisory work in these sectors, the risk management questions I consistently raise at board and executive level include: What is the firm's AI risk taxonomy, and is it integrated into the enterprise risk framework? Are the data quality and governance controls adequate to support the accuracy and auditability standards that regulators will apply? What human oversight mechanisms are in place for AI-influenced decisions in regulated processes? Is there a clear escalation pathway when AI systems produce outputs that fall outside expected parameters? And how is regulatory change being monitored and incorporated into the AI governance framework as the regulatory environment continues to evolve?

None of these questions has a simple answer, and in most of the organisations I work with, the honest answer is that the framework is at an early stage of development relative to the pace of AI deployment. That gap is manageable, but it requires deliberate attention at the right level of the organisation. It will not close on its own.

7. What Boards and Executive Teams Should Be Doing

Based on my advisory work across boards and executive teams, including in heavily regulated sectors, I offer the following observations on where leadership attention should be focused.

Understand your workforce exposure at the task level, not the job title level

The relevant unit of analysis is the task and the workflow, not the occupation or the job description. Boards and executive teams that are working from job-level exposure assessments are operating on inputs that are too coarse to be useful. The task-level decomposition approach, grounded in frameworks like O*NET, provides the analytical foundation for credible workforce and operating model strategy in the AI era.

Distinguish between replacement risk and augmentation opportunity

Most AI deployment in enterprise today is augmentation, not replacement. Your workforce strategy, your AI investment thesis, and your operating model design should reflect this reality. Augmentation-led AI strategy has different organisational implications than replacement-led strategy, and the evidence from early adopters suggests that augmentation-led approaches are generating more sustainable value.

Update your assessments for generative AI

If your organisation conducted automation exposure assessments or workforce impact studies before 2022, those assessments require updating. The generative AI frontier is materially different from the automation frontier that prior studies were designed to measure. Roles and functions previously assessed as low-exposure may now be significantly exposed. The analysis needs to be refreshed.

Treat human capability development as a strategic priority

The capabilities that remain genuinely difficult to automate, contextual judgment, ethical reasoning, relationship intelligence, creative synthesis, and leadership, are also the capabilities that determine competitive outcomes in most of the markets I work across. Building, retaining and deploying these capabilities needs to be a board-level priority, not a human resources sub-agenda item.

Get honest about your data foundation before you commit to your AI strategy

Before committing significant capital to AI-driven workforce transformation, boards should demand an honest assessment of the data architecture and governance maturity that underpins it. The most common source of underdelivered AI value is not the AI technology. It is the data infrastructure it depends on. In regulated sectors, this is not a technical due diligence question. It is a risk management question, and it belongs on the board agenda with the same rigour as any other material risk.

Integrate AI risk explicitly into your enterprise risk framework

AI risk is not a technology risk sub-category. In regulated sectors particularly, it is an enterprise risk with regulatory, reputational, operational, and strategic dimensions. Boards should satisfy themselves that AI risk is explicitly identified, assessed, and governed within the enterprise risk framework, that human oversight mechanisms are adequate for the regulatory environment in which the organisation operates, and that the accountability structure for AI-influenced decisions is clear, documented, and auditable.

Establish clear AI governance at the board level

Boards that are not yet asking structured questions about how AI systems are being used in their organisations, what decisions they are influencing, what the risk and bias characteristics of those systems are, and where human oversight sits are running governance gaps that will become more significant as AI deployment accelerates. In regulated sectors, this is not optional. It is a regulatory and fiduciary responsibility, and the window for getting ahead of it is narrowing.

Concluding Observations

The AI-driven transformation of workforce and operating model is not a future event. It is underway now, across every sector I work in, at varying speeds and with varying degrees of strategic intentionality. The organisations that are navigating it best are not those with the most AI investment. They are the ones that have thought most carefully about the intersection of AI capability, human capability, workflow design, and operating model architecture.

For organisations in regulated sectors, that thinking must extend further. The data architecture that underpins AI deployment, the governance frameworks that maintain accountability and auditability, and the risk management structures that keep the organisation inside its regulatory envelope are not implementation details. They are the conditions under which AI-driven transformation is permissible and sustainable. Organisations that get the workforce and operating model strategy right but neglect the data and governance foundation will find, sooner or later, that the foundation gives way.

The analytical foundations are well established. The task-level decomposition approach, O*NET-based workforce analysis, the distinction between automation and augmentation, and the governance frameworks that regulated sectors require provide a rigorous basis for board and executive decision-making. What is often missing is the translation of those frameworks into operational strategy and governance practice.

That translation is where advisory work adds the most value, and it is where I focus my practice. If this paper has raised questions relevant to your organisation's AI, workforce, data or governance strategy, I welcome the conversation.

About the Author

Marco Ciobo is a senior transformation and strategy advisor with a BCG background and deep experience in technology strategy, digital transformation, workforce redesign, and board governance across telco, utilities, banking, and government sectors in the ANZ region. He is a Global Advisor to Bain & Company and advises boards and C-suite executives through his consulting practice Digital Insights and Transformation. He is the author of a forthcoming book on board governance in the digital era.

marcociobo.com