

THE DIVISION OF INTELLIGENCE

Shaping the World of Tomorrow in the Age of AI

A WHITE PAPER FOR SENIOR LEADERS AND AI PRACTITIONERS

Derived from the Basel Knowledge Creation Event, 18 March 2026

Hosted by Onesto Consulting in partnership with Enterprise Ireland

Marriott Hotel, Basel, Switzerland

CONTRIBUTORS AND PARTICIPANTS FROM:

*Novartis / Roche / Siemens / Bayer / Merck / Sandoz / Novo Nordisk / Abbott / EY /
Logitech / Agroscope / Boyden / Work Evolved / Xebia / InnoNext / Aspeya / Viv /
Enterprise Ireland / Embassy of Ireland in Switzerland*

Abstract

This white paper synthesises the perspectives that emerged from the Division of Intelligence Knowledge Creation event convened in Basel, Switzerland on 18 March 2026. Hosted by Onesto Consulting in partnership with Enterprise Ireland, the event brought together senior leaders from globally prominent organisations to examine the structural, human, and technical dimensions of AI adoption at enterprise scale. The event comprised a keynote fireside discussion, a presentation on the future workforce, a moderated panel discussion, and a structured peer-learning workshop addressing four organisational challenges that participants themselves identified as most pressing. The event was conducted under Chatham House rules; organisations are named as participants, but no contribution in this paper is attributed to any individual or organisation.

The paper advances five interconnected arguments — the Readiness Illusion, the Automation Trap, the Accountability Gap, Silent Skill Erosion, and the Outsourced Mind — each grounded in the lived experience of practitioners and substantiated by peer-reviewed and industry-level research. Together, these arguments form a practitioner-oriented framework for AI transformation that is simultaneously academically rigorous and operationally actionable. The paper is addressed to senior leaders — Chief Executive Officers, Chief Operating Officers, Chief Human Resources Officers, Chief Financial Officers, and Chief Technology Officers — and, equally, to the Non-Executive Directors and Board members whose oversight role in the AI era is no less consequential. It is also written for the practitioners, consultants, and governance professionals who support them.

The central thesis is that the dominant challenge of AI adoption is not technological but organisational. BCG's 10-20-70 framing — that AI value is roughly 10 per cent algorithms, 20 per cent data and technology, and 70 per cent people, processes, and operating models — captures a pattern that is empirically supported by McKinsey research showing that workflow redesign is the dominant predictor of AI value. Most enterprises are over-investing in the 30 per cent and chronically under-investing in the 70 per cent. The paper identifies where the real barriers lie, why they persist, and what leaders can do structurally to break through them — including, critically, the recognition that the operating models that brought today's leaders to their positions are not the operating models that will enable AI-augmented enterprises tomorrow. That change has already happened. The first task of leadership is to accept it.

Executive Summary

Eighty-eight per cent of organisations are now experimenting with artificial intelligence. Six per cent are extracting real, measurable value at enterprise scale. The gap between these two figures is not a technology gap. It is an organisational gap — a failure of operating models, accountability structures, skill development strategies, and leadership honesty.

This paper presents five insights that emerged from a Knowledge Creation event in Basel in March 2026, where thirty-five senior leaders from organisations including Novartis, Roche, Siemens,

Bayer, EY, and Sandoz spent an afternoon confronting the structural barriers to AI value. The insights are:

- The Readiness Illusion: organisations see the AI opportunity clearly but have not translated that recognition into honest readiness assessment. The gap is not a technology gap — it is an honesty gap.
- The Automation Trap: organisations are automating the processes they have rather than reimagining the processes they need. AI makes the wrong things faster, or accelerates one stage of a process while creating bottlenecks elsewhere.
- The Accountability Gap: AI is making or influencing decisions, but no individual or body is clearly responsible for outcomes. The accountability sink is widening as AI moves toward autonomy. The boundary between human-generated and AI-generated outputs is becoming porous, making accountability harder to assign by the day.
- Silent Skill Erosion: the expertise, judgment, and domain knowledge that organisations depend on are being quietly eroded by the very tools designed to enhance them. The damage is underway; the consequences are deferred.
- The Outsourced Mind: the risk of AI is not merely that it replaces human work, but that it replaces human thinking — and the organisation does not notice until the capacity for independent reasoning has already been lost.

For each insight, this paper presents what practitioners observed, what the research literature confirms, and what leaders can do. The paper does not stop at diagnosis. It closes with a synthesis framework, recommendations for the C-suite and the Board, and a set of fifteen diagnostic questions any leadership team can use to locate itself on the AI transformation journey. If you read this paper, you will know what to do — and you will know who in your organisation needs to do it.

1 Introduction: The AI Reckoning

This paper was developed under Chatham House rules. The thirty-five organisations whose senior leaders contributed to the discussions in Basel are named in the participant list. No view, observation, or quotation in this paper is attributed to any individual or organisation. The collective practitioner wisdom that emerged from the event is presented as such throughout.

The event that generated the analysis in this paper opened with a speech by the Irish Ambassador to Switzerland, a reminder that the conversation about artificial intelligence has consequences that extend beyond any single organisation or sector. Ireland and Switzerland — both small, open, innovation-led, knowledge-based economies that consistently punch above their weight in research, education, and global competitiveness — share a strategic interest in understanding how AI can be harnessed in ways that drive genuine value while protecting the rights, skills, and judgment of the people who must live and work within AI-augmented organisations. The knowledge-based nature of both economies amplifies their sensitivity to AI disruption: their competitive advantage rests on knowledge work, not manufacturing or natural resources, which makes the AI question existential rather than merely strategic. Stanford's Global AI Vibrancy rankings place both nations within the leading global cohort, and yet the practical challenges facing organisations in Basel, Zurich, Dublin, and Cork are strikingly similar to those facing organisations everywhere: how do you move from AI as a technology demonstration to AI as an organisational capability?

The event's title — The Division of Intelligence — named the core tension. Intelligence, in the age of generative AI, is no longer the exclusive property of human minds. It is being produced by systems that can generate text, analyse data, synthesise information, and make recommendations with a fluency and speed that no human can match. This division creates an extraordinary opportunity: to augment human capability, accelerate learning, and produce outcomes that neither human nor machine could achieve alone. But it also creates an extraordinary risk: that organisations, in their eagerness to capture the efficiency gains that AI offers, inadvertently overlook and eventually erode the human capabilities — the judgment, the expertise, the moral reasoning — that make those organisations worth running in the first place.

A central argument of this paper, made here explicitly and developed throughout, is that the operating models that brought today's leaders to their positions are not the operating models that will enable AI-augmented enterprises tomorrow. This is not a future possibility to plan for. It has already occurred. The leaders being asked to lead this shift are themselves products of the model that no longer fits the moment: their instincts, their governance habits, their definitions of value were all formed in a different context. This is not a failing. It is the human condition of transformation. But acknowledging it is the first step to addressing it, because the leaders who are waiting to decide whether to adapt are already behind.

The fireside dialogue framed around the question of AI transformation in highly regulated environments such as global pharma surfaced the central tension immediately. AI promises to augment medical decision-making, accelerate evidence generation, and reshape the talent

pipeline. But in environments where patient safety is non-negotiable and regulatory accountability is absolute, the question of where human judgment ends and algorithmic recommendation begins is not a governance nicety. It is a clinical and ethical imperative. The session posed six questions that would echo throughout the afternoon: where does AI genuinely augment expertise, and where does it fall short without human intuition? How should leaders think about the productivity paradox — speed versus rigour? Who owns the AI decision in a regulated world?

The workforce presentation that followed provided the structural frame. It articulated a shift from legacy operating models — linear, resource-constrained, time-bound — to AI-native models characterised by exponential outputs, predictive design, and computational intelligence. Critically, the presentation mapped the C-suite consequences: the CEO as chief orchestrator rather than chief decision-maker; the COO integrating AI into core processes rather than deploying it as a productivity tool; the CHRO leading the transition from headcount-based workforce planning to agentic workflows; the CFO shifting from cost optimisation to decision velocity as the primary metric; and the CIO/CTO building modular, scalable infrastructure for an organisation that will change continuously. This is not an IT transformation. It is a leadership transformation with IT as one of its instruments.

The panel discussion that followed sharpened the edge of these observations. Three panellists with experience spanning aviation, regulated research, and global consulting examined the AI reckoning from three angles: the theatre of innovation, the governance failure, and the human cost. The table workshop that concluded the afternoon asked participants to move from observation to action, confronting four challenges — the Accountability Gap, Efficiency Theatre versus Real Value, the Readiness Illusion, and Skill Atrophy — with the discipline of practitioners who must solve problems, not merely name them.

What follows is a critical synthesis of what was heard, confirmed by research, and distilled into a framework for leaders who see the AI era not as a technology investment but as an organisational transformation.

2 Insight 1: The Readiness Illusion

The readiness gap is not a technology gap. It is an honesty gap.

2.1 The Evidence from the Room

The consensus from Basel was uncomfortable but consistent, and subtler than it first appears. The issue is not that leaders are arrogantly overconfident about their preparedness. Participants were not claiming readiness; they could see clearly what was possible. The gap they could not name comfortably was between aspiration and honest self-knowledge. Governance frameworks exist — documented, approved, and filed. Ethical guidelines have been published. Compliance programmes are underway. What has not happened, in most cases, is translation: governance into behaviour, ethical principles into operational controls, accountability frameworks into named responsibility.

This is not a universal failing. Some organisations — a structured AI agent governance framework deliberately built by one of the larger pharma participants is a useful counterpoint — have moved past the readiness illusion. They have done so not by accident, but by design: they have constructed the structural conditions in which honest self-assessment is possible. The default trajectory of organisations, however, in the absence of that deliberate design, is toward performance rather than readiness. It is possible to break through. It requires explicit structural choices, not goodwill.

Perhaps the most counterintuitive finding from the peer workshop was that the workforce is often ahead of the systems. Curiosity and willingness to engage with AI exist throughout organisations. What is missing is the infrastructure, the operating model clarity, and the leadership courage to admit that the organisation is not ready and to treat that admission as a starting point rather than a career risk.

One observation from the room captured the dynamic precisely: nobody wants to be the person who says 'we are not ready.' What is meant by 'honesty' here matters, and it has three distinct dimensions. The first is understanding: do leaders actually know where their organisation is on the readiness curve? The second is capability: do they have the diagnostic tools and the operating data to assess readiness accurately? The third is willingness: do they know, but are not willing to admit it? All three dimensions matter. But the willingness dimension is the hardest to address and the most consequential, because no amount of better data or sharper tools resolves a structural reluctance to name what is true. In the absence of that admission, organisations proceed with AI deployments that produce visible activity — pilots, dashboards, use cases — without producing value.

2.2 Research Findings

The McKinsey State of AI 2025 survey of 1,993 executives across 105 countries confirms the quantitative dimension of this observation: 88 per cent of organisations now deploy AI in at least one business function, but only 6 per cent qualify as high performers who are extracting enterprise-level value. The gap between adoption and impact is not explained by technology access. Only 39 per cent of organisations attribute any EBIT impact to AI use, and among those, most report less than 5 per cent improvement.

Source: McKinsey & Company, The State of AI in 2025, November 2025.

The single strongest predictor of enterprise AI impact, McKinsey found, is whether an organisation fundamentally redesigned its workflows when deploying AI — not the sophistication of the model, not the scale of data access, not the technology budget. The structural implications of this finding are returned to in Section 7.3.4.

The World Economic Forum's Future of Jobs Report 2025 adds the workforce dimension: nearly 40 per cent of skills required on the job are projected to change, with 63 per cent of employers already citing the skills gap as the primary barrier to transformation. 170 million new roles will be created, 92 million displaced, and 22 per cent of all current roles affected by 2030. The transition will not be evenly distributed, and organisations that treat this as a spreadsheet exercise rather than a human challenge will lose the trust of the workforce they most need.

Source: World Economic Forum, Future of Jobs Report 2025, January 2025.

PwC's 28th Annual Global CEO Survey (2025) captures the leadership version of the illusion: 73 per cent of CEOs view AI as central to competitiveness, while only 37 per cent feel ready to integrate it effectively. IMD's 2026 AI Trends analysis identifies 2026 as the year in which this illusion must break, arguing that the most successful organisations will stop treating AI as a technology race and start treating it as a management revolution.

Sources: PwC, 28th Annual Global CEO Survey, 2025; IMD, 2026 AI Trends, December 2025.

2.3 The Practitioner Imperative

This is primarily a CEO and Board mandate, in close partnership with the CHRO and the COO. Without sponsorship at this level, the structural conditions for honesty cannot be built. Breaking through the Readiness Illusion requires four things. The evidence points to a sequence that cannot be shortcut.

Before even attempting to review your readiness for AI, ensure that you have done the hard work on your strategy. Although technology has always been an enabler for strategy in the same way an operating model is, leaders need to ensure both that AI implementations are aligned with strategy and, as the uses of AI further reveal themselves, that AI is shaping strategy in turn. This becomes a reinforcing loop in which AI is both shaping and reinforcing strategic intent.

First, conduct an honest diagnostic — not a maturity model exercise designed to confirm progress, but a genuine assessment of whether governance is lived or merely documented, whether accountability has a name or only a framework, and whether the operating model is designed for

AI-era decision-making or is a legacy structure with AI artificially added on. The diagnostic must be conducted with people who are empowered to report what they observe, not what they believe the organisation wants to hear. If the diagnostic produces only green lights, the process is broken. AI itself should be used as part of the diagnostic process. AI can surface patterns, anomalies, and gaps in governance data that humans might miss or be reluctant to name. This is not the same as the claim, advanced in some commentary, that AI will find its own paths through organisational chaos. It will not. AI only converges on good governance outcomes if it is deliberately designed and directed to do so. Used carefully, however, AI is a powerful instrument for surfacing what human diagnostics suppress.

Second, align the operating model before deploying technology. AI is only as good as the operating model it lands in. An AI tool deployed into a governance structure that lacks clarity, a decision architecture that lacks speed, or an incentive system that rewards the wrong behaviours will produce faster versions of existing dysfunction. Operating model readiness is not a downstream consequence of AI deployment — it is a precondition of it.

Third, create structurally enabled honesty. Asking for individual courage in a culture that punishes candour is unfair and ineffective. The organisations capturing AI value did not rely on heroic individuals. They built structural conditions for honesty: governance bodies constituted to challenge rather than confirm, incentive systems that reward learning from failure, and leadership behaviours that model vulnerability — the willingness to name gaps before they become crises. Culture cannot be changed by declaring it; it changes when the structures that produce it change.

Fourth, track the gap between stated and lived AI readiness. Most organisations measure AI adoption — tools deployed, pilots launched, training completed. Very few measure whether governance frameworks are actually being followed, whether named accountability owners understand and accept their responsibilities, or whether stated AI principles are shaping real decisions. Build a small set of indicators that test whether readiness is real or performed, and report them honestly to the board.

To make this actionable, the CEO and Board will need to grant explicit authority for the diagnostic, accept that the findings may be uncomfortable, and treat the resulting agenda as a leadership priority rather than a transformation programme to be delegated.

3 Insight 2: The Automation Trap

The most dangerous AI strategy is a successful one that automates a process that should not exist.

3.1 The Evidence from the Room

The room kept returning to one uncomfortable question: are we automating the right things? A widely cited airport example — the Emirates facial recognition boarding system, which abolished the ticket rather than accelerating the ticket check — landed directly: the innovation was not making the existing process faster. It was removing it. That is not an efficiency gain. It is a fundamentally different conception of what the organisation does. The distinction between incremental improvement and genuine redesign ran through every table discussion. One table formulated it explicitly: if you have the same leadership, the same people, and the same processes, will AI really make a difference? The room's answer was no.

Senior leaders acknowledged that their organisations were deploying AI to accelerate current processes rather than asking what the ideal process would look like. The bottleneck, as was observed, is not the process. It is the decision-making architecture around it. Automating a process that sits inside a broken decision structure produces faster versions of the wrong outcomes.

There is also a further observation, raised forcefully in the room, that practitioners need to confront: AI does not always reduce work. It often intensifies it. When AI accelerates a task that previously took hours, expectations rise. The bar for output is reset upward. The pressure to use the time saved on more work — not on better work, not on reflection, not on rest — follows immediately. That pressure creates an incentive to cut corners: to pass AI output off as fully validated work when it has not been, to skip the verification step that the AI was supposed to enable, to mistake speed for quality. This is the dynamic that connects the Automation Trap directly to the Accountability Gap and to Silent Skill Erosion. The productivity case for automation is more complicated, and less straightforwardly positive, than the case sheets suggest.

3.2 Research Findings

The RAND Corporation has found that AI projects fail at roughly double the rate of traditional IT initiatives, and Carnegie Mellon University's Human-Computer Interaction Institute reports failure rates approaching 90 per cent for enterprise AI projects. The consistent finding is that most failures trace to organisational and strategic causes — unclear vision, poor process design, lack of user adoption, cultural resistance — rather than technology limitations.

Sources: RAND Corporation, AI Project Failure Analysis; Carnegie Mellon HCI Institute.

Messeri and Crockett's landmark 2024 paper in Nature warns that AI tools can create 'illusions of understanding' — environments in which organisations believe they comprehend more about their operations and capabilities than they actually do. Organisations automating most confidently may

be those understanding least about what should change. The risk is not only doing the wrong things faster; it is believing transformation is underway while the wrong things are being preserved.

Source: Messeri, L. & Crockett, M.J. (2024), 'Artificial intelligence and illusions of understanding in scientific research,' *Nature*, 627, 49–58.

Mollick's 'Jagged Frontier' research reveals something counterintuitive: AI's capabilities are deeply uneven, excelling far beyond human ability in some areas while failing surprisingly at others, and the pattern rarely maps to what leaders expect. An organisation that automates an entire workflow because AI handles 95 per cent of it brilliantly may not notice the 5 per cent where it quietly breaks down — until the consequences surface. Mollick borrows the concept of 'reverse salients' from military history, where a front line advances unevenly and a single lagging section holds back the entire advance. In AI terms, one stubborn weakness can bottleneck an entire system — but when that weakness is finally solved, the whole front lurches forward suddenly and unpredictably. This means that operating models need to be designed for continuous adaptability, not optimised around today's snapshot of what AI can and cannot do.

Source: Mollick, E. (2025), 'The Shape of AI: Jaggedness, Bottlenecks and Salients,' *One Useful Thing*.

A recent Harvard Business Review article (February 2026) sharpens the practitioner concern raised in Basel: AI does not reduce work — it intensifies it. The article documents a pattern across knowledge-worker settings in which AI adoption accelerates individual tasks but raises expectations and intensifies workloads in the same motion. The productivity case for automation is therefore more complex than headline efficiency numbers suggest. Speed gains on the individual task can be absorbed by raised performance bars, by additional verification work created downstream, or by the displacement of effort into other, less measured activities.

Source: Harvard Business Review (February 2026), 'AI Doesn't Reduce Work, It Intensifies It.'

3.3 The Practitioner Imperative

This is primarily a COO mandate, in partnership with the CIO/CTO and the heads of the functions whose processes are being reimaged. Practitioner ownership without operational authority will fail. The position taken in this paper is that the answer is not to automate better — it is to reimagine first. The sequence matters, and getting it wrong is costly.

Before asking how to redesign a process with AI, ask the prior question: should this process exist at all? Many of the processes organisations are now considering automating exist because they were built around constraints that no longer apply, or to manage information flows that AI itself can render unnecessary. Process elimination is the most radical — and often the most honest — answer to the automation question. It is also the answer most likely to be missed, because eliminating a process means foregoing the visible activity of automating it.

First, reimagine before you automate. Leverage AI to think through how work should be done — with people and technology working together — before committing to automation. The right question is not 'how do we make this process faster?' but 'what should this process look like if we were designing it today?' The organisations that skip this step consistently discover that they have automated their bottlenecks rather than eliminated them.

Second, look at the whole process, not just the parts AI can accelerate. A critical and often underappreciated risk of automation is that AI may dramatically speed up one part of a process while creating new bottlenecks, quality problems, or compliance burdens elsewhere — consequences that only become visible when the entire process is examined holistically. An organisation that automates document drafting without examining the review and approval steps downstream may find that overall cycle time increases rather than decreases. Process redesign must consider the full end-to-end flow, including the human steps that sit before and after the AI-assisted ones, and the question of whether the process itself should be retained or trained out of existence. Speed in one stage without attention to the whole is not transformation; it is displacement of effort.

Third, once reimagined, use automation in service of a clear design rather than as a substitute for it. BCG's framework of Deploy, Reshape, and Invent is a useful reference: most organisations are stuck in Deploy because they have not done the foundational work of asking what their operating model should look like when AI is a core component, not an add-on. Moving into Reshape requires answering that question first — and committing to the answer even when it disrupts current structures.

Fourth, make process elimination as visible and celebrated as process automation. Every organisation tracks how many processes it has automated. Very few track how many it has eliminated entirely. The former measures speed; the latter measures wisdom. Leaders rewarded only for building and deploying will systematically over-automate and under-question. Structural incentives must make elimination — the recognition that a process should not exist at all — a respected and valued outcome.

Making this real requires the COO to hold cross-functional authority for process redesign, the CFO to track elimination as a metric alongside automation, and the Board to scrutinise the balance. Without that structural support, the gravitational pull is always back toward visible automation rather than invisible elimination.

4 Insight 3: The Accountability Gap

If the system is accountable, nobody is. And in the age of AI, 'nobody' is making more decisions than ever.

4.1 The Evidence from the Room

The panel surfaced the accountability question through a pointed exchange between experienced voices. It was argued that 'Copilot' was yesterday — today it is 'Autopilot', with agentic AI taking decisions autonomously and executing activities without human intervention. If competitors operate without human-in-the-loop oversight and move faster as a result, organisations that insist on human oversight risk being too slow to compete. It was countered that this trajectory is not inevitable and that organisations and regulators retain the ability to shape AI governance, not merely respond to its acceleration. This tension was not resolved in the room, nor should it be: it names the central governance question of the AI era.

The room also examined how organisations are beginning to treat AI agents as team members — setting them up within defined policies and boundaries, essentially as digital employees. The actual practice is more nuanced than blunt removal: where an agent operates outside its policy, it is first corrected or retrained, exactly as an underperforming employee would be coached and developed before any other action is taken. Only if that correction does not resolve the issue is the agent decommissioned. The language used in some organisations is 'fired.' This paper argues that this language is dangerous, and that the danger is sharper than it first appears. We are not simply borrowing the vocabulary of firing. We are replicating the full performance management cycle — coaching, corrective action, formal removal — around entities that are not people. The cognitive boundary between tool and person erodes not at the moment of decommissioning, but across the whole disciplinary process. AI agents should be decommissioned. The language of technology management must remain distinct from the language of people management, precisely because that boundary is becoming harder to see.

A live tension surfaced in the conversation also deserves explicit acknowledgement. Large organisations with deliberately flat or federated structures — the over 100 CEOs example raised in the room — face a specific accountability challenge: the governance demands of agentic AI may create pressure toward more explicit hierarchical accountability, even in organisations that have invested deliberately in flatter, holacratic, or decentralised governance models. This paper does not seek to resolve that tension. It names it as a live and unresolved organisational design question: can flat governance models accommodate the accountability demands of agentic AI, or does AI exert a structural pull back toward more hierarchical accountability? The answer will vary by organisation, and there is, at present, no settled view.

4.2 Research Findings

Dan Davies's *The Unaccountability Machine* (2024) provides the foundational concept of 'accountability sinks' — structures that absorb or obscure the consequences of decisions so that no one can be held directly accountable. Davies's fundamental law is precise: the extent to which you are able to change a decision is precisely the extent to which you can be accountable for it. AI is the ultimate accountability sink. The developer says the deployer is responsible. The deployer says the developer is responsible. The manager says the system made the recommendation. The system cannot be asked.

Source: Davies, D. (2024), *The Unaccountability Machine: Why Big Systems Make Terrible Decisions*, Profile Books / University of Chicago Press.

Conway's Law — Melvin Conway's 1968 observation that organisations design systems which mirror their own communication structures — provides the structural mechanism behind Davies's accountability sinks. If the organisational structure has unclear accountability lines, fragmented decision-making authority, and ambiguous ownership, the AI systems deployed within that structure will reproduce and amplify those same patterns. This is not a governance failure that can be corrected after deployment. It is a design inevitability. An organisation with over 100 CEOs, to use the Basel example, will produce AI systems with as many different accountability assumptions — unless the accountability architecture is explicitly designed before deployment begins. Conway's Law means that fixing AI accountability without fixing organisational accountability is structurally impossible. The system will always mirror the organisation that built it.

Source: Conway, M.E. (1968), 'How Do Committees Invent?', *Datamation*, 14(4), 28–31.

The EU AI Act (Regulation 2024/1689) creates a structural imbalance between provider obligations and deployer obligations. While Article 26 places monitoring and oversight duties on deployers, the heaviest obligations fall on providers. Moreover, the Act's high-risk system obligations do not become fully enforceable until August 2026. Regulation at this level will never keep pace with the speed of AI development. Compliance is the floor; ethical governance is the ceiling; and the space between them is where accountability lives or dies.

Source: EU AI Act, Regulation (EU) 2024/1689; Freshfields analysis of deployer vs provider obligations, 2025.

MIT Sloan Management Review research (Işık and Goswami, 2025) identified three structural obstacles preventing organisations from translating responsible AI principles into practice: the gap between stated ethics and actual implementation; the lack of governance mechanisms with real authority; and the tendency for responsible AI frameworks to serve as reputational window dressing. Research published in *Performance Management Review* (2025), analysing 115 real-world AI incidents in the public sector, found that deep-learning models — the dominant AI technique in deployment — are inherently opaque, producing decisions that are difficult even for technical experts to explain.

Sources: MIT Sloan Management Review (2025), 'The Three Obstacles Slowing Responsible AI'; *Performance Management Review* (2025), 'Accountability and AI: Redundancy, Overlaps and Blind-Spots.'

A further and increasingly urgent dimension of the accountability challenge concerns the blurring boundary between human-generated and AI-generated outputs. As generative AI becomes embedded in everyday workflows — drafting documents, analysing data, producing recommendations — the provenance of decisions is becoming harder to trace. At the moment of production, it may be genuinely unclear how much of a given output reflects human judgment and how much reflects the AI system's statistical inference. When a regulated report, a clinical recommendation, or a legal filing contains AI-generated content that has been lightly edited by a human professional, who bears responsibility for its accuracy and consequences? The answer, in most organisations today, is unclear — and that unclarity compounds with every tool adoption and every workflow integration. Governance structures that were designed for a world of clearly human-authored outputs are not equipped for a world where that boundary is porous and shifting by the moment.

It is worth being clear about how much autonomy organisations have, in fact, already granted to AI. The agentic autonomy that participants debated in Basel is only the visible edge of a much wider phenomenon. The moment generative AI was placed into the hands of employees without governance — to draft, analyse, summarise, and recommend without oversight on the use of AI itself — organisations effectively let AI loose into their decision-making fabric. The autonomous agent is visible and at least partially governed. The autonomous employee using AI without oversight is invisible and ungoverned. The accountability gap is therefore wider than the formal agentic AI debate suggests, and a great deal of the work of closing it lies upstream of agentic deployment, in the everyday governance of generative AI use by individuals.

Industry observers have also begun to highlight a striking practical consequence in the intellectual property domain: a significant and rising proportion of patents now being filed are AI-generated or AI-assisted. The legal and accountability consequences — questions of inventorship, ownership, liability, and disclosure — are not yet resolved by either patent offices or by the bodies that grant them. Specific figures vary by jurisdiction and source; the directional observation is robust. Current legal frameworks were designed for human inventors. They are visibly straining.

4.3 The Practitioner Imperative

This is primarily a Chief Risk Officer and Chief Compliance Officer mandate, in partnership with the General Counsel and the CIO/CTO. The Board has a direct oversight role and cannot delegate it.

First, build an AI inventory. Document every AI system in use across the organisation — not just what exists, but what each system is being used for, who owns it, which business process it operates within, and who is accountable for each specific use case. A single AI tool may be used for ten different purposes across five different teams. Each use case carries its own risk profile and requires its own named owner. If you cannot complete this inventory, that itself is the finding.

Second, govern by principles, not controls. In an environment where AI capabilities change quarterly, control-based governance is structurally behind from the moment it is written. Principles travel further. A principle such as 'every AI-assisted decision must have a named human owner' is

enforceable across any use case, any technology, and any regulatory regime — without needing to be rewritten every time the technology changes.

Third, test the human accountability chain before a crisis does. A common objection to accountability tracing is that AI decisions are probabilistic, not deterministic — and that this makes it impossible to trace how the AI reached its output. That objection is correct, but it misidentifies what should be traced. The accountability chain is human, not algorithmic. The question to ask is not ‘how did the AI reach this output?’ but ‘what human decisions allowed this AI system to operate in this context, with these parameters, without adequate oversight?’ Who authorised the deployment? Who defined the risk parameters? Who was responsible for oversight, monitoring, and intervention? That chain is traceable, and that is where accountability sits. If the answer to those questions is unclear, the accountability structure exists on paper only.

To make this actionable, the practitioner team needs explicit authority to require disclosure of AI use across the business, cross-functional partnership with the operational owners of each use case, and Board-level air cover when the inventory surfaces uses that are uncomfortable to acknowledge.

5 Insight 4: The Silent Skill Erosion

The damage is being done now. It will only be felt when it is too late to reverse it. This is the environmental crisis of organisational capability.

5.1 The Evidence from the Room

Skill erosion was the thread that participants found most difficult to name but most viscerally felt. It appeared as anxiety: about the quality of AI outputs that nobody was checking, about junior staff producing work faster but understanding it less, and about the growing distance between the people making decisions and the expertise needed to evaluate whether those decisions were sound.

One of the most honest moments in the workshop came when a senior practitioner observation captured a pattern many in the room recognised: AI deployment had made junior staff significantly more productive but had eliminated the tasks through which those staff would have developed the judgment needed for senior roles. The manual work — the painstaking analysis, the close reading of source documents, the construction of arguments from first principles — was precisely the work that AI had automated. And that manual work was precisely how expertise was built. The same tool that makes today better is making tomorrow worse. The room sat with this observation. There was no comfortable answer. There was only the recognition that doing nothing guarantees the outcome everyone most wants to avoid.

The succession question proved the hardest to resolve. Today's senior professionals can validate AI output because they spent decades doing the work manually. They developed the pattern recognition, the domain knowledge, the ability to sense when something is wrong before they can articulate why. If AI does that work now, where does the next generation build that judgment? You cannot validate what you have never learned to do yourself. This concern was raised most sharply in the context of software and code development — where AI-generated code is increasingly the default and junior developers may never write the basics from scratch — but the principle generalises across every domain where expertise is built by doing.

A related observation also belongs here, because it bridges directly into the Accountability Gap. Most organisations have not waited for formal agentic AI deployment to face an autonomy question. They created one the moment they allowed employees to pass generative AI output as their own work, without governance on the use of AI itself. The bar has shifted, expectations have risen, and the pressure to cut corners follows. When the verification step is also the very practice through which expertise would have been developed, every corner cut is also a unit of capability foregone.

5.2 Research Findings

Research published in *Communications of the ACM* (2025), drawing on Carnegie Mellon, Aarhus University, and UC Santa Barbara, found that knowledge workers reported AI made tasks seem cognitively easier, but were simultaneously ceding problem-solving expertise to the system. Critically, their confidence in using AI increased as their underlying capability declined. The combination — declining capability with rising confidence — is precisely the condition that makes erosion invisible to measurement.

Source: *Communications of the ACM* (2025), 'The AI Deskilling Paradox.'

Research in *Organisation Science* (Dell'Acqua et al., 2025), studying the 'Jagged Technological Frontier,' found that AI dramatically improves performance on tasks within its capability boundary but can reduce performance on tasks outside it, particularly when workers over-rely on AI in domains where it is unreliable. Crucially, AI flattens the expertise distribution rather than raising it — producing a workforce more uniformly competent at a lower ceiling. Ferdman's 2025 paper in *AI & Society* provides the structural framework: 'capacity-hostile environments' describe contexts in which AI mediation actively impedes the development and exercise of human capabilities. The erosion is not individual failure; it is a design consequence.

Sources: Dell'Acqua et al. (2025), 'Navigating the Jagged Technological Frontier,' *Organisation Science*; Ferdman, A. (2025), 'AI deskilling is a structural problem,' *AI & Society*.

The 'supervision deficit' identified by Coney (2026, practitioner analysis) names the generational compounding: each cohort of professionals has less unassisted analytical experience than the last, reducing the pool of qualified supervisors. Verification structures remain in place, but the humans executing them can no longer perform them effectively. The brittleness is invisible during normal operations and only manifests during disruptions — precisely when independent human judgment matters most.

Source: Coney, L. (2026), 'The Skill Erosion Paradox: Preserving Analytical Capability in AI-Augmented Teams' (practitioner analysis, not peer-reviewed).

5.3 The Practitioner Imperative

This is primarily a CHRO mandate, in partnership with the functional leaders of every business unit. Capability is built where work is done, not where training is delivered.

The evidence points to an uncomfortable but necessary conclusion: organisations must deliberately design environments that preserve and develop human expertise even as AI takes on more of the operational work. Left unmanaged, the natural trajectory is one-way. Four practical steps make this manageable.

First, protect deliberate practice. Define which tasks in your organisation are capability-building — the work through which judgment, pattern recognition, and domain expertise are actually formed — and protect a proportion of them from full automation, even where automation is technically possible. The short-term cost is marginal. The long-term cost of not doing it is a workforce that is confident and credentialed but unable to think independently when it matters.

Second, run AI-off resilience assessments. Once or twice a year, ask teams to demonstrate they can deliver their core function without AI support for a defined period. Treat it as a business continuity exercise, not a performance test. If teams cannot function without AI, the organisation has built a dependency, not a capability — and it is better to discover this in a controlled exercise than during a system outage or a regulatory intervention.

Third, track capability alongside productivity. Most organisations measure AI-assisted output obsessively. Very few measure whether the humans producing that output are developing or losing the underlying expertise. Add a capability dimension to performance conversations alongside the productivity metrics.

Finally, connect this to your AI inventory. For every AI system and use case documented, ask two questions: has this AI system enhanced or replaced the relevant human capability — and do you know which? And how is that capability being actively maintained? If the answer to either is unclear, the erosion has already begun. To act on this, the CHRO requires cross-functional authority to mandate deliberate practice protections, partnership with the Board on capability investment, and a budget structure that treats capability development as infrastructure rather than discretionary spend.

6 Insight 5: The Outsourced Mind

The greatest risk of AI is not that it will replace human work. It is that it will replace human thinking — and we will not notice until we have forgotten what it felt like to think for ourselves.

6.1 The Evidence from the Room

The cognitive dimension of AI risk was the thread that participants found hardest to articulate but felt most strongly. Several table discussions returned to the same observation: when AI generates the first draft of anything — a strategy paper, a market analysis, a regulatory assessment — it sets the frame. The human who reviews that draft is no longer thinking from first principles. They are thinking within the frame AI has established. Questions are shaped by the answers AI has already provided. The options considered are bounded by the options AI has surfaced. The thinking has been outsourced before anyone decided to outsource it.

One observation described this as ‘the death of the blank page.’ The blank page is uncomfortable. It forces you to decide what matters before any words are written. It requires judgment prior to any output. When AI fills the blank page, that judgment is skipped. The human role shifts from creator to editor. Editing is valuable. But it is not the same cognitive act as creation. An organisation whose people only ever edit AI’s thinking is an organisation that has outsourced its mind.

Multiple table discussions also noted the narrowing of inquiry. When AI is used for research, it returns what is statistically probable, not what is most interesting, contrarian, or unexplored. Participants noted that their teams were producing work that was more polished but less original since adopting AI tools. The thinking was smoother, more efficient, and noticeably less distinctive. The baseline of output quality had risen. The ceiling of original insight had lowered.

6.2 Research Findings

A 2025 study published in the CHI Conference proceedings, surveying 319 knowledge workers and collecting 936 first-hand examples of generative AI use, found that users with high AI confidence were significantly less likely to enact critical thinking. The researchers identified ‘mechanised convergence’: users with access to generative AI tools produce a less diverse set of outcomes for the same task compared to those without. Thinking converges toward AI’s statistical centre.

Source: Lee et al. (2025), ‘The Impact of Generative AI on Critical Thinking,’ CHI Conference on Human Factors in Computing Systems.

Gerlich’s 2025 study found a strongly negative correlation of $r = -0.68$ between frequent AI tool usage and critical thinking scores across 666 participants. Cognitive offloading was strongly correlated with AI usage ($r = +0.72$) and inversely related to critical thinking ($r = -0.75$). The more people trusted AI, the less they checked it.

Source: Gerlich, M. (2025), 'AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking,' Societies, 15(1).

Messeri and Crockett's 2024 paper in Nature describes the risk of 'monocultures of knowing' — environments where certain types of inquiry dominate because AI supports them well, while other modes of knowing are quietly marginalised. Organisations produce more but understand less. Doshi et al.'s 2025 research in Strategy Science demonstrates that AI-augmented strategic decisions may homogenise across firms. When organisations use the same AI tools to analyse the same markets, they converge on the same strategies. Competitive differentiation requires what AI cannot provide.

Sources: Messeri & Crockett (2024), Nature; Doshi et al. (2025), Strategy Science.

6.3 The Practitioner Imperative

This is a CEO and CHRO mandate jointly, with the leadership team modelling the protocol if it is to mean anything. Cognitive habits are set at the top.

Protecting independent thinking is not about rejecting AI — it is about being deliberate about when and how it enters the thinking process. The risk is not that AI is used, but that it is used too early, before the human mind has had the chance to frame the problem, form a view, or identify what actually matters. If AI sets the frame, the thinking will converge toward the statistical mean — and so will the strategy.

Four practical mechanisms make this operational. First, adopt an 'originate before you generate' protocol: for strategic, ethical, and client-facing decisions, the human point of view comes first — a written argument or clearly articulated position — before any AI tool is opened. Second, deploy AI explicitly as a challenger rather than an author: the question to ask is not 'write me an analysis' but 'here is my analysis — what have I missed and where is my reasoning weakest?' Third, implement a false positive protocol: introduce deliberate, known errors into AI outputs at unpredictable intervals and track whether people catch them. Declining catch rates are an early warning that critical evaluation has been replaced by passive acceptance. Fourth, develop AI abstinence fluency alongside AI fluency — in a world where everyone uses the same tools, knowing when not to use AI is the differentiator. Make the ability to work without AI a valued and measured competence, not a residual skill that quietly disappears through disuse. To make this real, the leadership team itself must use the protocol visibly. If it is not enforced at the top, it will not be practised below.

7 Synthesis: The 10-20-70 Framework and the Path Forward

7.1 The Investment Paradox

The five insights synthesise into a single diagnostic: organisations are systematically over-investing in the 10–20 per cent of AI transformation that concerns algorithms, data, and technology, while chronically under-investing in the 70 per cent that concerns people, processes, operating models, culture, and governance. This proportional framing is BCG's 10-20-70 — first published in BCG's 2022 article 'Five Rules for Fixing AI in Business' — a pattern drawn from BCG's extensive client experience rather than from primary research. It is a practitioner insight, and it should be treated as such. The strength of the pattern is that it is empirically supported by independent research — most notably the McKinsey State of AI 2025 finding that workflow redesign, not technology investment, is the dominant predictor of enterprise-level AI value. The tragedy is that investment in the 30 per cent is visible, measurable, and reportable. Investment in the 70 per cent is slow, structural, and produces results that appear in the next generation of leadership, not in the next quarter.

Source: BCG (2022), 'Five Rules for Fixing AI in Business'; McKinsey & Company (2025), The State of AI in 2025.

7.2 What the Voices in the Room Told Us

Several direct observations from participants crystallised the diagnosis: that there is no need to push the AI agenda — leaders will themselves decide when, where, and how to use AI when the conditions are right; that the human in the loop is essential for regulatory work and cannot be engineered away without consequences; that return on investment from AI is not in line with investment, pointing to the cost of the Readiness Illusion; that AI literacy for everyone and AI fluency for many is a prerequisite for value, not a downstream consequence; that the role of CEO, COO, CHRO, CFO, and CTO is to be an orchestra leader — or, in the more demanding metaphor developed below, a jazz band participant — for AI; that reskilling needs to be competency-based, not credential-based; and that the organisations where AI agents operate within certified, diligently checked compliance frameworks represent the emerging standard for responsible agentic deployment.

7.3 The Operating Model Imperative

The operating model is not something that enables AI. It is the thing that determines whether AI creates value or destroys it. Every other insight in this series — the Readiness Illusion, the Automation Trap, the Accountability Gap, Silent Skill Erosion, and the Outsourced Mind — is a downstream consequence of operating model choices.

The operating model was not a scheduled topic in Basel. It became the dominant theme because every other conversation led back to it. This was not coincidental. It is structural. The five insights presented in this paper are not five separate problems. They are five expressions of the same underlying condition: an operating model that has not been redesigned for the AI era. Understanding why that is so — and what it demands — is the purpose of this section.

The connection is explicit. The Readiness Illusion persists because organisations are not honest about where they are — and that dishonesty is an operating model problem, produced by incentive structures that punish candour and governance frameworks that reward the appearance of progress over its substance. The Automation Trap exists because organisations automate work within structures never designed for AI, treating technology as a layer to add rather than a reason to redesign. The Accountability Gap widens because decision-making architectures fail to assign clear ownership as AI moves from tool to agent — a governance failure at the heart of the operating model. Silent Skill Erosion accelerates when operating models optimise for productivity without protecting the deliberate practice that builds expertise. And the Outsourced Mind is the ultimate consequence of an operating model that has no mechanism for preserving independent human judgment as AI takes over more cognitive work. Fix the operating model and these problems become manageable. Leave it untouched and no amount of technology investment will close the gap.

An important preliminary note: before redesigning the operating model, organisations should ask a prior question — what is the impact on the business model itself? The shift from legacy models (linear, resource-constrained, time-bound) to AI-native models (exponential, driven by predictive design and computational intelligence) may require rethinking not just how the organisation operates, but what it does and for whom. The operating model must ultimately serve the business model, and in the AI era, both are in motion simultaneously.

7.3.1 WHAT THE OPERATING MODEL IS

The operating model is the means by which an organisation executes its vision and strategy. It is the system that connects purpose, people, process, technology, and governance into a coherent whole — the heartbeat of the organisation. This system identifies six interdependent elements: (1) skills and capabilities, which form the foundation and are the first question in any operating model design; (2) ways of working; (3) structure; (4) governance; (5) reward structure; and (6) ecosystem and customer engagement. These elements do not draw a line between people and technology. In a digitally integrated organisation, structure describes how people, data, and technology are connected as a unified architecture, not as separate concerns to be aligned after the fact.

Culture is not an element of the operating model — it cannot be designed directly. It is the emergent consequence of how these six elements interact as a system. Once established, culture reinforces the operating model that produced it and directly influences which operating model changes are achievable. This creates a feedback loop: the operating model shapes the culture, and the culture determines what changes are possible. The operating model can be designed; the culture is revealed by it. When the culture that emerges is not the one the organisation needs, the

answer is not a culture programme — it is a redesign of the elements that produced it. An AI strategy that sits outside the operating model is not a strategy. It is a project plan with ambition.

7.3.2 DECISION VELOCITY AND THE POWER PROBLEM

Two structural challenges dominate operating model transformation in the AI era, and both surfaced directly in Basel. The first is speed. The observation from the room was direct: you can have the leanest, most efficient processes in the world, and they are still stalled by a lack of speedy decisions by leadership. Traditional governance cycles — quarterly planning, annual budgeting, committee approvals — were designed for a world where decisions had time to mature. AI does not wait. The organisations capturing value are those that have redesigned their governance for decision velocity, not just decision quality. Research from the Journal of Decision Systems (2026), drawing on 27 executive interviews in a global multi-brand organisation, found that the primary barriers to embedding AI in strategic decision-making were not technical but organisational: strategic ambiguity, limited leadership awareness, and governance structures designed for deliberation rather than velocity. The infrastructure for decision-making, not computation, is the bottleneck.

This points to a structural observation that connects directly to the through-line of this paper. The leaders being asked to redesign governance for decision velocity are themselves products of legacy decision-making environments. Their mental models of what good governance looks like — the committees that should review what, the level at which trade-offs should be escalated, the cadence at which decisions should be revisited — were formed before AI. This is not a criticism. It is a structural observation. The implication is that leadership development, and self-awareness about this bias, is a prerequisite for genuine operating model change. It is not enough to redesign governance on paper. The leaders applying that governance must first recognise that the instincts they have trained over decades may now be working against the organisation they are trying to lead.

The second challenge is power. Leaders whose influence depends on headcount have a structural incentive to resist workforce redesign. Functions whose budgets depend on project volume have a structural incentive to resist process elimination. These are not irrational responses — they are rational responses to existing incentive structures, and they will not change through persuasion alone. The operating model question is, at its core, an incentives question. Organisations that want different behaviours must design different incentives. No strategy document, however compelling, changes this without structural support. The CFO role signals the transformation required: the shift from cost optimisation to decision velocity as the primary metric redefines what the operating model is optimised for.

7.3.3 GOVERNANCE: THE JAZZ BAND PRINCIPLE

The governance challenge for large, federated organisations was captured memorably in Basel by an observation that one participant's organisation effectively had over 100 CEOs. In such an organisation, the question is not about centralisation versus decentralisation. It is about coherence:

how do over 100 autonomous units make consistent decisions about AI deployment, risk, and capability development without either rigid central control or ungovernable fragmentation?

The jazz band metaphor captures the aspiration. Not an orchestra following a conductor — a model that is too slow and too centralised for the speed AI requires — but a jazz band where each instrument shapes the music based on shared principles that everyone has internalised. This model requires three things: deep individual expertise so each player can hold their own; shared principles rather than imposed rules that people understand and own rather than merely comply with; and the discipline to listen as much as to play — feedback loops that actually change direction, not dashboards that report on direction already taken. This is not loose governance. It is a harder form of governance than most organisations have today. It demands more of people, not less. And it is the only model that can operate at the speed AI requires while maintaining the judgment that regulated industries depend on.

Research confirms the value of this approach. Survey data from 127 organisations found that hybrid, federated governance approaches achieve 41 per cent higher user adoption rates than centralised models, and that comprehensive upskilling strategies produce 2.7 times higher implementation success. The jazz band is not just a metaphor. It is a superior governance architecture for the AI era.

7.3.4 THE 10-20-70 CONNECTION

The operating model imperative is the structural explanation for the 10-20-70 pattern. BCG's framing — 10 per cent of AI transformation value in algorithms, 20 per cent in data and technology, and 70 per cent in people, processes, and operating models — is consistent with what McKinsey's 2025 State of AI survey demonstrates empirically: the single strongest predictor of enterprise-level AI impact is whether an organisation fundamentally redesigned its workflows, not the sophistication of the model, not the scale of data access, not the technology budget. High performers are 3.6 times more likely to aim for transformational change and nearly three times as likely to have redesigned processes when deploying AI. They did not invest more in technology. They invested differently in how they operate.

The five insights in this paper map directly onto that 70 per cent. The Readiness Illusion is a failure of governance and leadership honesty — both operating model elements. The Automation Trap is a failure to redesign ways of working and structure. The Accountability Gap is a governance failure. Silent Skill Erosion is a failure to protect skills and capabilities, the foundation element of the operating model. The Outsourced Mind is the consequence of all five elements failing together. The 70 per cent is not a residual category. It is the work. The organisations in the 6 per cent who are extracting genuine AI value did not get there by deploying better technology onto existing structures. They got there by redesigning the structures themselves.

The most powerful intervention is at the root. Break the Readiness Illusion through honest diagnostic and structurally enabled candour, and the remaining four become tractable: the operating model can be redesigned, accountability can be named, skills can be protected, and cognitive independence can be preserved. Leave the root untouched and no amount of technology investment will close the gap.

8 Recommendations for Senior Leaders and Board Members

The following recommendations are addressed to the C-suite leaders — CEO, COO, CHRO, CFO, and CTO — and, equally, to the Non-Executive Directors and Board members who must hold this leadership team to account. Many of the following recommendations are also questions of governance oversight, not only of executive action.

FOR THE CEO: BUILD THE CONDITIONS FOR HONESTY — AND CHOOSE REDESIGN

The most important question for a CEO in the AI era is not 'are we deploying AI well?' It is 'are we genuinely willing to redesign how work is done and how the organisation is structured — not just our tools?' This is the prior commitment that determines whether any other recommendation is meaningful. Layered on that commitment is the most important behavioural task: creating an environment in which it is safe to say 'we are not ready.' This is not a cultural aspiration; it is a structural design requirement. Governance bodies should be constituted to challenge, not confirm. Incentive structures should reward learning over performance theatre. And the CEO's own communications should model vulnerability — the willingness to name gaps before they become crises — rather than projecting confidence about a transformation that is still underway.

FOR THE COO: REDESIGN THE OPERATING MODEL BEFORE DEPLOYING AI

The COO must resist the pressure to demonstrate AI progress through automation metrics and focus instead on the more demanding question: what does the organisation's operating model need to look like when AI is a core component, not an add-on? This requires examining each of the six operating model elements — skills and capabilities, ways of working, structure, governance, reward structure, and ecosystem — through the lens of AI integration. The test is whether the operating model enables decision velocity with governance guardrails, or whether it was designed for deliberation in an era that demanded it.

FOR THE CHRO: LEAD THE CAPABILITY AGENDA, NOT THE TRAINING AGENDA

Reskilling is necessary but insufficient. The CHRO's mandate in the AI era is broader: to design the conditions under which human expertise continues to develop, to protect the capability foundations that make AI accountability possible, and to plan for the succession crisis that will arrive in ten years if the silent skill erosion is not addressed today. Competency-based rather than credential-based approaches to AI readiness, deliberate practice rotations, and capability-tracking metrics alongside productivity metrics are the practical instruments of this mandate.

FOR THE CFO: SHIFT FROM COST OPTIMISATION TO DECISION VELOCITY

The CFO who continues to evaluate AI investments through traditional ROI frameworks will consistently undervalue the operating model, governance, and capability investments that generate long-term AI value, and will consistently over-fund the technology deployments that produce short-

term activity without long-term return. Decision velocity — how quickly the organisation can move from insight to action — is the new metric. The CFO who aligns investment decisions with this metric, rather than with cost reduction targets, is the CFO whose organisation will be capturing genuine AI value five years from now.

FOR THE CTO/CIO: BUILD MODULAR, ADAPTABLE INFRASTRUCTURE

Given Mollick's insight about reverse salients — that AI bottlenecks can break suddenly when a single weakness is resolved — operating models designed around today's AI limitations will be obsolete overnight. Infrastructure must be modular and adaptable. The technology function's contribution to the accountability agenda is creating the observability, auditability, and explainability that make governance possible. The AI inventory recommended in Insight 3 is a technology leadership responsibility as much as a governance one.

FOR THE BOARD: TREAT AI AS A FIRST-ORDER GOVERNANCE RESPONSIBILITY

This is a board-level conversation, not only an executive one. Boards need to understand their role in AI governance, accountability oversight, and the capability agenda — and they need to hold the executive team accountable for it. Three questions should sit on the board's agenda regularly. First: are you receiving AI governance reports that tell you what is really happening in the organisation, or what management wants you to hear? The Readiness Illusion is structurally easier to project upward than downward. Second: is AI on the risk agenda at the same level of seriousness as financial and regulatory risk — or is it filed under 'innovation' and reviewed less rigorously? Agentic AI, the blurring boundary between human and AI outputs, and the accountability sink described in Insight 3 are all risks that warrant the discipline of formal risk oversight. Third: does the board have the capability to evaluate what it is being told about AI — the literacy, the diversity of perspective, and the willingness to ask uncomfortable questions? If not, addressing that gap is itself the first board action. A detailed treatment of what good Board engagement looks like in practice will be the subject of a follow-up article. For now, the imperative is clear: engagement is required, and the questions above are where it begins.

9 Diagnostic: Fifteen Questions for a Leadership Team

The following questions integrate the diagnostics from the five insights into a single leadership conversation. A leadership team that can answer all fifteen with confidence and specificity — not with reference to governance documents, strategy presentations, or transformation roadmaps, but with reference to lived operational reality — is an organisation that has broken through the Readiness Illusion. Each question is followed by a brief note on what a leadership team gains by answering it honestly.

- 1** Where in your operating model is there friction, or opportunity for efficiency gain, that AI could address — and which of those are structural rather than technological?

So what: this question forces the team off the language of ‘use cases’ and into the language of operating model design, which is where AI value is actually generated.

- 2** Are your cross-functional decisions being made on the key critical factors — budget, incentives, accountability — or are they being deferred or fragmented across functional silos?

So what: cross-functional decision quality is the single best predictor of whether AI investments will produce coherent value or duplicative pilots.

- 3** Of your last ten AI initiatives, how many eliminated a process entirely, how many redesigned the underlying training and capability, and how many simply accelerated an existing process?

So what: this is the ratio that distinguishes wisdom from speed. Acceleration without elimination or capability redesign is the Automation Trap by another name.

- 4** Are your leaders incentivised by headcount and budget size, or by outcomes and capability development?

So what: incentives, not strategy documents, determine behaviour. The wrong incentives will quietly defeat the right strategy.

- 5** Could your organisation operate effectively for one week if your AI systems went down?

So what: this is the resilience test that distinguishes dependency from capability. If the answer is no, you have built a dependency, not a capability.

- 6** Who is named as accountable for each AI use case and the business process it operates within — not the tool owner, but the named human owner of the use case and the process it serves?

So what: accountability that names neither the use case nor the business process is accountability on paper only. Both are required.

- 7** For each AI system in use, has this system enhanced and/or replaced the relevant human capability — and do you know which?
So what: organisations that cannot answer this question accurately are eroding capability without knowing it. The first step in protecting capability is knowing where it is being replaced.
- 8** Is your AI governance organised around principles that people understand and own, or around controls they are required to follow?
So what: control-based governance lags every quarterly model update. Principles travel further and survive technology change.
- 9** When did you last conduct an AI-off resilience assessment — demonstrating your team can deliver core functions without AI support?
So what: if you have never run this exercise, you do not know whether you have a capability or a dependency.
- 10** Name the five most critical judgment-dependent roles in your organisation. How many of the people in those roles built their expertise before AI? Who replaces them?
So what: this is the succession question that the Silent Skill Erosion makes urgent. If the answer is unclear, the erosion is already advanced.
- 11** What is your organisation's false positive protocol — the mechanism by which you verify that people are critically evaluating AI outputs rather than passively accepting them?
So what: in the absence of such a protocol, declining catch rates of obvious errors will be invisible to leadership until a crisis exposes them.
- 12** In your last three strategic analyses conducted with AI assistance, what do your conclusions contain that a competitor using the same AI tools would not have produced?
So what: the answer to this question is the answer to whether AI is producing competitive advantage or homogenisation.
- 13** Does your organisation track the input-to-output ratio of AI-assisted work — how much original human input goes into AI-assisted outputs?
So what: this ratio is the operational measure of the Outsourced Mind. Without it, cognitive offloading is invisible.
- 14** What happens to your operations if your AI provider doubles their pricing, changes terms, or becomes unavailable?
So what: dependency on a single provider is a strategic risk that compounds quietly. This is a question Boards should be asking.

15 Is it safe in your organisation to say 'we are not ready'? If not, what would need to change to make it safe?

So what: this is the most important question in the diagnostic. If the answer is no, every other answer above is unreliable.

10 Conclusion: The Leadership Choice

The event in Basel produced a finding that the research confirms and that every participant felt: the biggest barrier to AI value is not the technology. It is the willingness to be honest about where an organisation actually is, and the structural conditions that make that honesty possible.

The organisations capturing meaningful value from AI are not the ones with the best models. They are the ones that redesigned their operating models before their competitors did, named accountability before it became a crisis, protected human capability before the erosion became irreversible, and insisted that their people continue to think — independently, critically, and with the judgment that no algorithm can replace.

AI is not arriving as a single 'big bang' moment, as the Ambassador observed in her opening address. It is arriving as a new layer of intelligence that sits inside the everyday workflows of business. The question for every leader in every organisation is not whether to engage with that layer. It is whether to engage with it on AI's terms or on their own.

Engaging on your own terms means doing the harder work: the honest diagnostic, the operating model redesign, the accountability structure, the capability investment, the cognitive governance. It means treating the 70 per cent as the strategy and the 30 per cent as the instrument. It means building the conditions for honesty, for learning, for human judgment to remain at the centre of decisions that carry human consequences. And it begins with the recognition that the operating model that brought today's leaders to their positions is not the one that will carry them forward. That change has already happened. Accepting it is the precondition for everything else.

The organisations that make this choice will lead. The organisations that defer it will continue to experience AI as a series of expensive pilots that never scale, a set of governance documents that describe accountability without creating it, and a workforce that is faster and less capable with every passing quarter.

The division of intelligence is not just between human and machine. It is between leaders who are honest about the challenge and those who are not. That is the only division that matters.

References

The following references are cited in this paper. Where original sources were accessed directly, they are listed as such. Where cited via secondary synthesis, the secondary source is noted. Practitioner analyses are flagged as such and are not presented with the authority of peer-reviewed research.

INDUSTRY AND INSTITUTIONAL RESEARCH

BCG (2022). Five Rules for Fixing AI in Business. Boston Consulting Group.

BCG (2025). How Agents Are Accelerating the Next Wave of AI Value Creation. Boston Consulting Group, December 2025.

Gartner (2025). AI Maturity Index and Pilot Abandonment Projections. Via IMD AI Maturity Index, 2025.

Harvard Business Review (February 2026). 'AI Doesn't Reduce Work, It Intensifies It.' Harvard Business Publishing.

IMD (2025). 2026 AI Trends. International Institute for Management Development, December 2025.

McKinsey & Company (2025). The State of AI in 2025: Organisations Are Investing — But Not Capturing Value. McKinsey Global Institute, November 2025.

McKinsey & Company (2024). Executive Digital Transformation Survey.

PwC (2025). 28th Annual Global CEO Survey: Navigating an Era of Transformation. PricewaterhouseCoopers.

RAND Corporation. AI Project Failure Analysis. RAND Corporation Research Reports.

World Economic Forum (2025). Future of Jobs Report 2025. World Economic Forum, January 2025.

World Economic Forum / Kaizen Institute (2025). Business Barriers to AI Implementation Survey.

ACADEMIC AND PEER-REVIEWED LITERATURE

Busuioc, M. et al. (2025). Algorithmic discrimination in public service provision. *Journal of Public Administration Research and Theory*.

Carnegie Mellon Human-Computer Interaction Institute. Enterprise AI Failure Rate Studies.

Chalmers, D. et al. (2026). The Acceleration of Artificial Intelligence: Rethinking Organization and Work in an Era of Rapid Technological Change. *Journal of Management Studies*.

Cognitive Research: Principles and Implications (2024). Does using AI assistance accelerate skill decay and hinder skill development without performers' awareness?

Communications of the ACM (2025). The AI Deskilling Paradox. Association for Computing Machinery.

Davies, D. (2024). *The Unaccountability Machine: Why Big Systems Make Terrible Decisions*. Profile Books / University of Chicago Press.

Dell'Acqua, F. et al. (2025). *Navigating the Jagged Technological Frontier*. Organisation Science.

Doshi, A.R. et al. (2025). *Artificial Intelligence and Strategic Decision-Making*. Strategy Science.

EU AI Act (2024). Regulation (EU) 2024/1689 on Artificial Intelligence. European Parliament and Council.

Ferdman, A. (2025). *AI deskilling is a structural problem*. AI & Society, Springer.

Freshfields (2025). *Provider vs Deployer Obligations under the EU AI Act*. Freshfields Bruckhaus Deringer.

Gerlich, M. (2025). *AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking*. Societies, 15(1).

Işık, O. & Goswami, S. (2025). *The Three Obstacles Slowing Responsible AI*. MIT Sloan Management Review.

Journal of Decision Systems (2026). *From pilots to decision systems: embedding generative AI into strategic decision-making through a socio-technical and governance lens*.

Lee, M. et al. (2025). *The Impact of Generative AI on Critical Thinking*. CHI Conference on Human Factors in Computing Systems.

Messeri, L. & Crockett, M.J. (2024). *Artificial intelligence and illusions of understanding in scientific research*. Nature, 627, 49–58.

MIT Sloan Management Review (2025). *Accountability and Responsible AI*. MIT Sloan Management Review.

Mollick, E. (2025). *The Shape of AI: Jaggedness, Bottlenecks and Salients*. One Useful Thing (Substack).

Mollick, E. et al. (2023). *Centaur and Cyborgs on the Jagged Frontier*. Working paper.

Natali, C. et al. (2025). *AI-Induced Deskilling in Medicine*. Artificial Intelligence Review, Springer.

Performance Management Review (2025). *Accountability and AI: Redundancy, Overlaps and Blind-Spots*.

Preprints.org (2025). *The Impact of AI-Based Decision Support Systems on Human Agency in Governance*.

S&P Global Market Intelligence (2025). *Voice of the Enterprise: AI & Machine Learning*.

Stadler, M., Bannert, M. & Sailer, M. (2024). *Cognitive ease at a cost*. Computers in Human Behavior.

Sustainability (2025). *Sustainable AI Transformation: A Critical Framework for Organizational Resilience*.

Xu, Y. et al. (2025). Generative AI and Organizational Structure in the Knowledge Economy. arXiv working paper.

PRACTITIONER ANALYSES

Coney, L. (2026). The Skill Erosion Paradox: Preserving Analytical Capability in AI-Augmented Teams. Practitioner analysis (not peer-reviewed).

Conway, M.E. (1968). 'How Do Committees Invent?' *Datamation*, 14(4), 28–31.



About This Paper

This white paper was developed by Onesto Consulting from the Knowledge Creation event 'The Division of Intelligence: Shaping the World of Tomorrow in the Age of AI,' convened in Basel, Switzerland on 18 March 2026. The event was hosted by Onesto Consulting in partnership with Enterprise Ireland, opened by H.E. Ambassador McGarry, Irish Ambassador to Switzerland, with grateful acknowledgement of the support of Enterprise Ireland's Zurich office and, in particular, Vincenz Wagner and Philip Walsh.

The event was conducted under Chatham House rules. The organisations whose senior leaders contributed are named in the participant list. No view or observation in this paper is attributed to any individual or organisation.

The event featured a fireside discussion between Ashwini Mathur (Onesto Consulting) and Dr Thomas Lang (Novartis); a keynote presentation on the future **workforce** in the age of AI by Nadine Desiere; a moderated panel facilitated by Patrick O'Donnell (Onesto Consulting) with panellists Robyn Agoston (Work Evolved), Fabian Wahl (Agroscope), and Patrick Naef (Boyden); and a peer-learning workshop in which 35 senior leaders from organisations including Novartis, Roche, Siemens, Bayer, Merck, Sandoz, Novo Nordisk, Abbott, EY, Logitech, Agroscope, Boyden, Work Evolved, Xebia, InnoNext (represented by Susanne Lauber Fürst), Aspeya, Viv, and Enterprise Ireland worked in facilitated table groups on the four key challenges.

The five insights published alongside this white paper — The Readiness Illusion, The Automation Trap, The Accountability Gap, The Silent Skill Erosion, and The Outsourced Mind —constitute the analytical foundation from which this synthesis is drawn.

Onesto Consulting Limited | Navigating Complexity, Building Capability

Insight Series: The Division of Intelligence | March 2026