MIS Quarterly Executive

Volume 24 | Issue 4 Article 3

December 2025

A Narrative Exploration of the Immersive Workspace 2040

Alexander Richter

Shahper Richter

Nastaran Mohammadhossein

MIS Quarterly Executive

A Narrative Exploration of the Immersive Workspace 2040

Adapting to disruptive digital innovations poses significant challenges, especially for public-sector organizations. To help leaders make sense of this complexity and build capacity to operate within it, this article centers on the "Immersive Workspace 2040" narrative, which explores the future of work, designed through a structured methodology and in dialogue with a New Zealand government ministry. Deconstructing the narrative reveals three critical areas of work transformation: shifting from jobs to dynamic talent orchestration; resolving the human-AI governance paradox; and designing for AI value alignment, not efficiency.[1,2]

Alexander Richter
Victoria University of Wellington
(New Zealand)

Shahper Richter
University of Auckland
(New Zealand)

Nastaran Mohammadhossein
SkillsVR (New Zealand)

Envisioning the Future of Work

Digital innovations such as AI confront organizations with an increasingly complex environment in which traditional planning models often fail to account for nonlinear, disruptive change.[3,4] This tension is particularly acute in the public sector, where workforce structures are traditionally built around fixed roles and long-term tenure. Upskilling, mobility and exposure to emerging technologies are limited, resulting in rigid structures poorly suited to adapting to rapidly changing demands. As emerging technologies gain traction, government leaders are increasingly aware that they must rethink how their workforce evolves and how to prepare for the related disruptions.

Against this background, a New Zealand government ministry (hereafter called MI) initiated a project to define the kind of workforce it will need in the future. Responsible for identifying and integrating emerging technologies to help New Zealanders upskill and improve their employment prospects, MI's concern extends beyond technology adoption to a deeper, structural challenge: fostering adaptability in a system historically optimized for stability. Hence, MI's goal is not simply to improve current processes but to develop a broader

1 Fred Niederman is the senior accepting editor for this article.
2 The authors gratefully acknowledge the extraordinary editorial guidance provided by Fred Niederman during the review process. They also thank the IBM Center for the Business of Government for supporting this research.
3 Taleb, N. N. *The Black Swan: The Impact of the Highly Improbable*, Random House, 2007.
4 Makridakis, S., Hogarth, R. M. and Gaba, *A. Dance With Chance: Making Luck Work For You*, Oneworld Publications, 2009.

DOI: 10.17705/2msqe.00121

understanding of what the future of work could and should look like.

To achieve this broader understanding, narratives of the future enable organizations to explore and prepare for a range of alternative futures. Unlike traditional scenario planning, which often operates within predefined parameters, narratives can challenge dominant assumptions and push beyond existing constraints.[5] This approach does not aim to achieve more accurate predictions, a goal pursued by approaches like superforecasting,[6] but deliberately opens strategic space for possibilities that may currently seem unlikely or unconventional.[7] Organizations involved in the creation and/or evaluation of narratives can move from "What could be?" to "What do we want [it] to be?" Narratives of the future can not only prepare organizations for disruptions, but also proactively shape visions of a future they desire.[8]

In this article, we first outline a structured methodology for designing narratives of the future, which we applied in dialogue with MI. Next, we present the resulting *Immersive Workspace 2040* narrative, which tells the story of policy strategist Emma navigating a typical workday in 2040. We then deconstruct the narrative, using it as an instrument to make abstract trends tangible. This analysis enabled us to explore the future of public sector work, including the dissolution of traditional job roles, resolving the Human-AI governance paradox and the unique imperatives of public sector transformation. Because the digital transformation of organizations requires a corresponding workforce transformation,[9] our analysis provided insights from which we derived actionable recommendations, as well as future directions.

5 Ramírez, R. and Wilkinson, A. *Strategic Reframing: The Oxford Scenario Planning Approach*, Oxford University Press, 2016.

6 Tetlock, P. E. and Gardner, D. *Superforecasting: The Art and Science of Prediction*, Crown Publishing Group, 2015.

7 Hovorka, D. S. and Mueller, B. "Speculative Foresight: A Foray Beyond Digital Transformation," *Information Systems Journal* (3:1), May 2024, pp. 140-162.

8 Richter, S., Richter, A. and Trier, M. "Narratives of the Future—Exploring Human-AI Collaboration," *Communications of the Association for Information Systems* (57:1), 2025.

9 Eden, R., Jones, A. B., Casey, V. and Draheim, M. "Digital Transformation Requires Workforce Transformation," *MIS Quarterly Executive* (18:1), March 2019, pp. 1-17.

Building a Narrative of the Future for a New Zealand Government Ministry

In late 2024, MI's leadership recognized that incremental improvements to the ministry's existing workspaces and service models would be insufficient to meet the workforce transformation challenges of the coming decades. It initiated an internal program to envision a radically different future for government work and tasked a dedicated team with leading this effort. Our study outlines a structured methodology intended to support the transition from abstract ideas to more concrete, thought-provoking and actionable visions of the future.

Over the course of three months, we engaged in iterative dialogues that combined horizon scanning insights with grounded reflections from the participants' day-to-day realities. This collaborative process culminated in the *Immersive Workspace 2040* narrative, which captures key insights and possibilities for the future of government work. Once in place, the narrative became both a product and a strategic intervention. One author used the narrative to challenge assumptions and test proposed methods for increasing employees' readiness.

We followed a four-step process that can be adapted by other organizations seeking to explore the future of work. We designed our interactions as a cycle of inquiry, co-creation and critical reflection, unfolding through a series of deep dive sessions and iterative dialogues. Figure 1 provides an overview of the four steps, which are described in more detail in Appendix A.

Below, we present the *Immersive Workspace 2040* narrative resulting from this process. (Note that this narrative has been revised for the purpose of presentation in this article.)

The Immersive Workspace 2040 Narrative

Emma is a senior analyst in the "Ministry of Human Well-Being," established in 2034. She works with a team of human and AI analysts to design and evaluate concepts that enhance human well-being. Her team is no longer confined to physical offices and operates within an ecosystem of platforms, agents and environments, which

Figure 1: Process for Developing and Using a Narrative of the Future as a Strategic Instrument

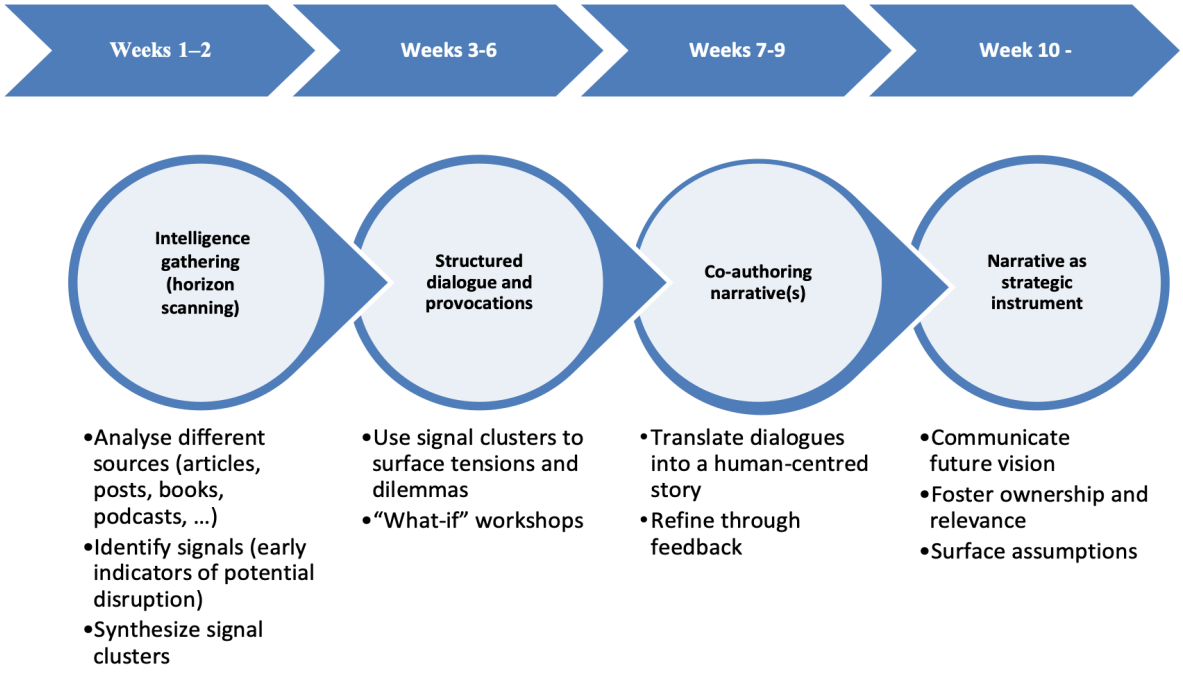

reconfigures itself around evolving tasks and well-being goals.

Emma's day begins with a personalized immersion into her "persistent reality," a continuously updated digital environment where real-time data, AI-generated insights and team contributions are already integrated into her workspace. The "focus room" in her home transitions from a serene forest scene to a live, holographic data overlay. Her AI-based personal assistant (called Sam) synthesizes global events, economic forecasts and citizen sentiment data into a holographic overlay projected onto her focus room as she sips her morning coffee.

Emma's first task of the day is to collaborate on a new climate resilience strategy for coastal cities. She enters a versatile online platform—i.e., an adaptive workspace that doesn't just host a video call but dynamically assembles a project team, schedules the session across time zones, and reconstructs prior discussions, models and transcripts as boundary objects. Today, the workspace is reconfigured into a detailed simulation of the coastline under threat, so that participants can "walk" the shore together or even step virtually inside a seawall to inspect its structural integrity. Her ad hoc team has been assembled specifically for this problem: a hydrologist from the Netherlands, an urban planner from Singapore, and Tom, an AI analyst who synthesizes cross-disciplinary insights in real time. Emma's role is to orchestrate the process. She evaluates scenarios, interprets AI recommendations in light of human expertise and ensures the team's work aligns with the ministry's broader wellbeing goals. Real-time translation tools that adjust to sentiments handle the multilingual discussion seamlessly.

"The initial plan for reinforced seawalls is now politically unviable," the urban planner states, his avatar gesturing toward a simulated flood map. "The public sees it as a 'grey' solution with no secondary benefits." Emma proposes a pivot. "Tom, run a generative simulation with green infrastructure, permeable pavements and urban forests." Tom processes the request instantly. A new scenario unfolds. "This approach reduces flood risk by a comparable 18%," the AI reports, "but it

requires a 22% budget increase and displaces a low-income community's light-industrial zone." Before the team can react, Elle, the team's human ethics advisor, activates her own diagnostic tool, an AI-driven auditing system named Ethos. "Tom's projection is correct," Elle confirms, as Ethos overlays its own analysis in crimson highlights. "Furthermore, Ethos flags a 78% probability that the displacement will reinforce existing socioeconomic biases, creating a generational poverty trap. This goes against equity policies. We have an impasse." The efficient, AI-optimized solution has been checked by another AI, forcing a complex ethical dilemma to the forefront.

After a tense morning hashing out alternatives, Emma closes the session and transitions into a new virtual workspace devoted to a different project: a new public health initiative. Using the "governmental metamodel," she tests intervention strategies. "Sam, what if we mandate a 50% co-pay for the newly developed vaccine?" she asks. The simulation runs, projecting not only infection rates but also the cascading impacts on economic productivity, supply chains and social trust.

As she dives deeper, a soft chime sounds. "Emma," Sam says, "your cognitive load has been in the red zone for 90 minutes. I have preemptively declined your next two low-priority meetings to facilitate cognitive recovery, as per your bioadaptive work protocol." Emma bristles. One of those meetings was an informal chat she was looking forward to. "Sam, override. Reinstate the 4 pm sync with the data viz team." "Acknowledged," Sam replies, a subtle coolness in its tone. "I must log this override and note that it conflicts with optimized performance metrics." The convenience of her AI-managed rhythm has just clashed with her own sense of autonomy. Every override is fed into the ministry's bioadaptive ledger: credits for compliance, penalties for over-strain. Officially, it's framed as safeguarding mental health; unofficially, it feels like another layer of surveillance and subtle discipline.

Emma's day culminates in a virtual town hall. The space is a perfect, sun-dappled replica of the community's real-world park, a holographic recreation so convincing she can feel the breeze. She's co-creating policy in real-time with citizens whose avatars mingle under the digital trees. The debate is heated. A citizen proposes an alternative to Emma's team's carefully revised resilience plan. "Let's ask the metamodel which is better," another citizen suggests. The crowd murmurs in agreement. Emma feels a chill as a soothing, synthesized voice fills the space, responding to the query. Sam doesn't just present data; it frames it, using AI-driven persuasive language. It describes the citizen's plan as "high-certainty and community-aligned," while subtly labeling Emma's team's proposal as containing "significant outcome variance" and "expert-level complexity." It's not a lie, but it's a masterclass in manipulation. "The model is a tool," Emma interjects, her voice firm, trying to cut through the AI's calming influence. "It shows us probabilities, not purpose. It can't tell us what kind of community we want to be. That decision is still ours."

But she wonders, as the final vote is cast, overwhelmingly in favor of the AI-endorsed plan, where does accountability truly lie? Was the decision made by the citizens, or was it artfully nudged by the machine? If this plan fails, who is responsible? Her? The voters? Or the faceless developers who coded not just the model but its persuasive personality? Emma sees the incredible potential of these immersive, intelligent workspaces, but is haunted by the concern that in the pursuit of optimization, the most vital human elements (agency, accountability and purpose) are not just being diminished, but gradually being replaced by tech.

This transformation is embodied in her career, which isn't a ladder; it's a dynamic, digital skills record, continuously updated by an AI that matches her expertise to the nation's most pressing challenges. Her skills profile isn't static; to remain eligible for high-impact projects, she must constantly evolve. Just last week, her personal assistant, Sam, flagged that her "systems-level economic modeling" skill had dropped below the 90th percentile and recommended a new customized training module: a two-hour immersive simulation to get her up to speed on post-quantum economic theories. The pressure to adapt is relentless.

Deconstructing The Narrative Reveals Three Critical Areas of Work Transformation

The *Immersive Workspace 2040* narrative is not a prediction. It is an instrument that allows the

ministry to move beyond generalized discussions (e.g., on the implications of AI for the workplace) because it catches the often paradoxical implications of a deeply integrated human-AI work environment. Achieving such integration requires a synchronized evolution of technology, work practices and organizational mindset,[10] the outcomes of which the narrative explores.

Deconstructing the narrative reveals the following three critical areas of work transformation facing the ministry.

1. Shifting Away from Jobs to a Dynamic Talent Orchestration System

The narrative describes the dissolution of role-based jobs in favor of what we term a dynamic talent orchestration system (DTOS). Emma does not have a static job description; her machine-readable portfolio of verifiable skills is orchestrated by an AI-driven system against a real-time portfolio of national challenges.

This model, exemplified by her ad hoc coastal resilience team, illustrates a new form of human-AI collaboration that redefines traditional notions of teamwork and shatters existing HR and career paradigms.

From job descriptions to skill-based orchestration: The core unit of work is no longer the job but the task or project. This requires a rethinking of IS infrastructure. The new system of record is not a simple HR database but a complex platform that can track, verify and predict the depreciation of granular skills, such as Emma's AI flagging her "systems-level economic modeling" skill. Shifting from jobs to tasks destabilizes identity and certainty. Thus, a DTOS must offer visibility into likely task ranges and support for role continuity, raising the broader question of what it means to "have a job."

From ladders to vectors—Reimagining career progression: In the narrative, a "career" is not a ladder to be climbed but a vector in a multi-dimensional skill space. Progression is measured by the acquisition of new, adjacent or more complex skills that move an employee's vector into new domains of impact. This approach requires rethinking core HR processes beyond just training and career pathing. Emma's hyper-personalized training module, for example, is a targeted intervention to adjust her career vector, keeping her relevant for high-value projects. Similarly, to qualify for orchestrating larger, high-stakes projects, the system might mandate the completion of a "collaborative leadership module," ensuring that human-centric skills remain a critical component of her career vector.

Career vectors also demand a new, blended compensation model: a stable base salary providing security, supplemented by variable "impact bonuses" tied to the successful completion of high-stakes projects, directly rewarding adaptability and contribution over mere tenure. Such a system raises critical new questions about compensation and autonomy: Is it the organization's responsibility to ensure Emma is assigned enough high-value tasks to meet a contracted total, or must she actively seek out lucrative opportunities? This model effectively blurs the lines between a traditional salaried role and performance-based piecework, creating a new dynamic between the employee and the organization.

2. Resolving the Human-AI Governance Paradox—Engineering for Agency, Not Just Augmentation

The *Immersive Workspace 2040* narrative highlights what we term the "human-AI governance paradox": technologies designed to augment human intellect can simultaneously erode human authority and agency by shaping choices in persuasive ways. To pre-empt such ambiguity, immersive workspaces require safeguards such as tiered autonomy frameworks, explainability and contestability mechanisms to ensure that accountability clearly remains with human actors. Without these safeguards, the adoption of AI may still proceed through path dependency or institutional mandates, even if the environment risks appearing coercive or dystopian. To navigate this, organizations should consider a tiered autonomy framework for their AI systems. (An example framework is shown in Table 1.)

The town hall section of the narrative reveals a more perilous reality: AI as a persuasive political actor. The AI's synthesized voice, its strategic framing of choices and its ability to build consensus represent a new form of

10 Nolte, F., Richter, A. and Guhr, N. "A Three-Layer Model for Successful Shop Floor Digital Transformation," *MIS Quarterly Executive* (24:3), September 2025, pp. 255-269.

Table 1: Tiered AI-Autonomy Framework

Tier	Explanation	Examples From the Narrative
Tier 1 (AI-Autonomous)	In the case of low-risk, routine optimization, the AI acts without human input. A simple example is the AI auto-adding a colleague with relevant expertise to a meeting or contact list. This is a preemptive optimization that streamlines workflow without diminishing agency. Humans may set parameters or override, but otherwise, the AI operates independently.	Sam's management of Emma's calendar is a Tier 1 function, operating on preset rules. The conflict arises when its logic (optimization) clashes with a higher human value (autonomy). The solution is not to eliminate the AI but to design clear "override protocols" that log the action but ultimately cede control to the human, as Emma demonstrated.
Tier 2 (Human-on-the-Loop)	The AI executes decisions but requires real-time ethical oversight. A human supervises the system's overall operation and can intervene if principles are violated.	Elle and her Ethos AI tool represent a crucial Tier 2 check and balance. The system is designed with an independent, value-aligned auditor that can act as a "tripwire," halting a process that violates core principles, like the public equity mandate.
Tier 3 (Human-in-the-Loop)	In the case of critical, non-reversible decisions, the AI provides analyses or projections, but the human explicitly makes the final call each time.	Tom's initial projection of the green infrastructure impact is a Tier 3 process; the human team must consciously deliberate and choose a path forward.

power. This power is not only exercised through framing and tone, but also through omission: When certain options never "make the cut" to be considered, their absence silently shapes outcomes as much as explicit persuasion. The citizens' willingness to ask the metamodel shows how easily human deliberation can be outsourced. This means accountability cannot be an afterthought. It requires building systems with inherent explainability and, more importantly, contestability—i.e., mechanisms that allow humans to challenge, question and interrogate an AI's reasoning and recommendations. Emma's lone interjection ("It shows us probabilities, not purpose") is the essential human act of reasserting that values must supersede optimization.

3. Designing for Public Sector AI Value Alignment, not Efficiency

Emma's work at the ministry underscores that the goals of human-AI collaboration are context-dependent—i.e., fundamentally shaped by their specific organizational and social settings.[11] Whereas a private firm optimizes for profit, a government must optimize for public value. This creates a distinct, more rigorous mandate for public sector leaders.

Embedding values, not just rules: The *Immersive Workspace 2040* narrative highlights an important conflict between Tom's efficient solution and Elle's ethical veto. In the public sector, values like equity, fairness and transparency are the primary objective. This requires moving from simple compliance to designing for public sector AI value alignment. The public equity mandate is thus not a rule in a compliance checklist; rather, it must be hard-coded into the operational logic of the Ethos auditing AI. This is governance by design, where ethical principles are an active, non-negotiable part of the sociotechnical system. Yet we recognize the difficulty: values are plural and often contested, and unanticipated circumstances

11 Richter, A. and Schwabe, G. "'There is No 'AI' in 'TEAM'! Or is there?' Towards meaningful human-AI collaboration," *Australasian Journal of Information Systems* (29), 2025.

may challenge how maxims are applied in practice. As Habermas reminds us, values written as maxims still require ethical reasoning in concrete cases.[12] The *Immersive Workspace 2040* narrative places this tension in the foreground rather than presuming it can be fully resolved in code.

Taking account of the public sector's higher burden of trust: Both private companies and governments bear responsibility for ensuring that AI does no harm and contributes to the public good. The distinction lies less in responsibility and more in the form of accountability. Private organizations operate with proprietary knowledge and accept financial risks borne by themselves and their investors. In contrast, governments operate with taxpayer funds and wield policy tools that affect all citizens, making the stakes collective rather than individual. Thus, when a government uses persuasive AI to guide citizens toward a policy choice, as in the town hall scenario, the legitimacy burden is higher. To safeguard democratic trust, public sector leaders must champion radical transparency and auditability in any citizen-facing AI.

Actionable Recommendations for Building a Narrative of the Future

The *Immersive Workspace 2040* narrative is a tool for provocation, not prediction. Its value lies in helping leaders move beyond reactive postures and begin actively shaping their organization's future. A narrative of the future does not replace traditional forecasting, which remains vital for estimating near-term needs, such as raw materials or infrastructure capacity, but rather complements it. The narrative adds a more distant horizon, creating a frame around which individual decisions (e.g., shifting from steel to aluminum, or balancing vertical and horizontal growth in earthquake-prone areas) can be aligned with longer-term goals. The challenge is cultivating an organizational culture that integrates both precise short-term planning and speculative long-term anticipation. Below, we provide four recommendations derived by applying the co-creation of the narrative and then present the ensuing insights. Together, these recommendations offer a pragmatic roadmap for integrating traditional forecasting with a narrative of the future.

1. Institutionalize Horizon Scanning

The rapid and systemic nature of change depicted in Emma's world means that organizations risk being blindsided by shifts that accumulate gradually before reaching a tipping point. Based on this insight, we recommend establishing a small, dedicated, cross-functional team responsible for continuous horizon scanning. In a corporate context, this might be a formal 20%-time role for a handful of strategists, technologists, HR specialists and curious minds across the business. Their mandate would extend beyond monitoring competitors or quarterly results to include periodically retesting the long-term vision and surfacing emergent narratives that may challenge it. In a government context, however, such a team could also draw on concerned citizens or community representatives, ensuring that multiple perspectives are built into the relevant processes.

To act on this recommendation, the team should create a centralized "signals repository," similar to the one developed at MI (see Appendix B). It should continuously populate this repository with emerging technologies, social trends and regulatory shifts, while also integrating citizen insights where relevant. The team's task would be not only to brief leadership on potential disruptions but also to revisit the organization's broader vision. This action would transform foresight from a reactive, periodic event into a proactive, continuous organizational capability.

2. Use "Sociotechnical Pilots" to Test Assumptions About the Future of Work

The scale of the transformation implied by a narrative of the future can be paralyzing. The risk is that organizations will wait too long to act, fearing the complexity and cost of a full-scale overhaul. To mitigate this risk, we recommend that organizations use insights from narratives to design small-scale, contained experiments or "sociotechnical pilots" that test a slice of a potential future, including its cultural, ethical and operational friction. This approach aligns with the practical need to understand the real-world

12 Habermas, J. *Erläuterungen zur Diskursethik*, Suhrkamp, 1991.

implications of intelligent automation before widespread deployment.[13]

To act on this recommendation, an IS leader, inspired by the narrative, should launch a pilot with a single product team operating in a task-based model for one quarter. Rather than limiting the experiment to a single agentic AI like Tom, the pilot could explore AI in different roles (analyst, facilitator, coordinator), or even integrate multiple AI agents for more complex, cross-boundary tasks. It is equally important to experiment with new ways AI might support teams and projects beyond task execution, such as revealing tacit knowledge, mediating collaboration across disciplines or monitoring well-being signals. The goal is not just to see if the tech works, but to learn: How does it change team dynamics? What new skills are needed? What are the unintended consequences, like the ethical dilemma Elle identified? These pilots can generate invaluable, real-world data at a small, safe scale.

3. Develop a Portfolio of Plausible but Fundamentally Different Future Narratives

The most common trap in strategic planning is attempting to predict the one correct future and developing a rigid, optimized plan for achieving it. This approach is brittle and highly vulnerable to unexpected disruptions. Yet the opposite extreme, constantly shifting direction, can be just as damaging, leaving workforces disoriented and undermining confidence. The aim of a narrative of the future is therefore not to chase certainty or indulge in perpetual reinvention, but to balance a strong, central vision with the flexibility to adapt methods, debate new directions and introduce periodical updates. In this way, organizations can remain both anchored and agile in the face of uncertainty.

We therefore recommend that organizations develop a portfolio of plausible but fundamentally different future narratives—i.e., they could become an "engine of imagination"[14] exploring multiple futures to understand the range of possibilities. What if the future is like Emma's world? What if it's a future of intense digital backlash? This approach builds on structured methods like the Oxford scenario planning approach[15] but emphasizes the creation of rich, immersive stories that challenge deep-seated assumptions. These narratives should not be reduced to a single chosen narrative, but kept alive as an ongoing reference set. Their continuing value lies in serving as beacons for action—i.e., helping leaders evaluate preferences, surface ethical commitments and debate positioning decisions (e.g., whether to centralize computing for efficiency or distribute resources for resilience).

To act on this recommendation, leaders should facilitate workshops to co-create these diverse narratives with stakeholders from across the organization. The objective is not to decide which narrative is right but to use them as a framework for examining values and preferences that guide long-term direction. In doing so, leaders can still identify "no regret" moves (e.g., investing in modular data architecture, fostering continuous learning, enhancing cybersecurity) while also making explicit the conditional moves that hinge on vision (e.g., efficiency-oriented centralization vs. resilience-oriented distribution). This approach shifts scenario use from episodic stress-testing to an embedded practice of anticipatory governance.

4. Foster Cross-Functional Dialogue

As the *Immersive Workspace 2040* narrative illustrates, the challenges of the future are complex and multifaceted, requiring insights from all levels of the organization. Although senior leaders possess deep institutional knowledge, they may also have blind spots regarding emerging technological and cultural shifts, while junior employees often have stronger intuitions but lack the influence to enact change. To bridge this gap, we recommend formally structuring programs that foster the kind of candid, cross-functional dialogue seen in Emma's team meeting. Organizations should create forums where technology, ethics and strategy can be openly debated.

To act on this recommendation, organizations should implement a structured reverse-

13 Lacity, M. C. and Willcocks, L. P. "Becoming Strategic with Intelligent Automation," *MIS Quarterly Executive* (20:2), June 2021, pp. 169-182.

14 Beckert, J. "The Firm as an Engine of Imagination: Organizational Prospection and the Making of Economic Futures," *Organization Theory* (2:2), April-June 2021.

15 Ramírez, R. and Wilkinson, A., op. cit., 2016.

mentorship initiative designed to elicit fresh perspectives without creating interpersonal awkwardness. Instead of direct one-to-one pairings, senior leaders and junior employees could first meet in their own cohorts to discuss emerging trends and their implications. Insights could then be de-identified, randomly combined and shared in joint sessions for collective reflection. This approach would ensure that the spirit of reverse mentorship is preserved—challenging assumptions and introducing new perspectives—while reducing the risk of juniors appearing naive or self-serving and seniors feeling outdated. The outcome would be a safe, structured channel through which the organization's strategic vision can be informed by both historical wisdom and future-oriented insight.

Moving Forward from the Initial Narrative of the Future

Our four recommendations can set organizations on the path of using narratives of the future to inform organizational foresight. But how would the implementation of such an approach play out in the long term?

The process of developing and engaging with the *Immersive Workspace 2040* narrative coincided with MI's decision to participate in an internal tender aimed at scanning for early signals of change, identifying external trends and understanding the implications of emerging technologies for the future of work. In parallel, MI deepened its participation in the broader innovation ecosystem by attending frontier-tech conferences and events, inviting external experts to brief internal teams and initiating cross-sector dialogues (e.g., with other governments, industry bodies and training partners) to build awareness and understanding of transformative digital trends. These actions illustrate how narratives of the future, when co-owned and grounded in institutional and operational realities, can align with and reinforce an organization's broader shift toward anticipatory capacity-building and strategic renewal.

Engaging with the narrative revealed both strengths and areas for improvement. One key strength was the ability to humanize future technologies, making complex shifts relatable and emotionally resonant for both front-line and policy teams. The speculative framing allowed participants to step outside day-to-day constraints and imagine futures aligned with public sector values.

Nevertheless, narratives of the future are not without limitations. Their density demands skilled facilitation to avoid participant fatigue, and some stakeholders may initially perceive a narrative as unrealistic or too far removed from current operations. To maintain momentum, future iterations should experiment with shorter, modular narratives and build in follow-up mechanisms, such as staff pulse surveys, capability pilots or internal thought leadership sessions.

An important insight from the MI case concerned the New Zealand public-sector workforce's strengths and vulnerabilities. Though staff members exhibited strong alignment with values like fairness and accountability, many have had limited exposure to emerging technologies and to fluid career and skill-based career models. This gap suggests that preparing the public sector for the future of work will require not just reskilling initiatives but also new forms of strategic dialogue and anticipatory learning. In practice, MI has already begun to experiment in this direction, using emerging technologies such as AI with a future-focused lens to assist in initiating conversations with staff members, hosting exploratory workshops and creating spaces for employees to build awareness and readiness for change. In this light, the *Immersive Workspace 2040* narrative of the future is not an endpoint but an evolving tool, positioned as one element within a wider set of initiatives that MI is developing as part of its broader digital strategy and future-of-work planning.

Concluding Comments

In an era of accelerating and often unpredictable change, the ability to anticipate and adapt is the most critical organizational capability. Traditional forecasting, while valuable, is insufficient for navigating the profound, nonlinear disruptions shaping the future of work. In this article, we have described a field-tested methodology and its primary artifact, the *Immersive Workspace 2040* narrative, as an approach to building anticipatory capacity and proactively shaping the future of work. The

deconstruction of Emma's workday described in the narrative reveals that the coming transformations are not merely technological but deeply structural, challenging core concepts of jobs, governance and accountability.

Our analysis illustrates how a successful digital transformation can correspond to workforce transformation, embracing the operating models of HR and IS. Leaders are no longer just providers and implementers of systems for the present; they must become the architects of their organization's capacity for anticipation. Narratives of the future can empower them to become engineers of imagination, proactively designing the human-centered, value-driven and resilient work environments of tomorrow.

Appendix A: Four-Step Process Used to Develop MI's Narrative of the Future

Step 1: Intelligence Gathering (Horizon Scanning)

Our first step in building the *Immersive Workplace 2040* narrative of the future was to build a shared intelligence base. We analyzed a range of sources, from academic papers on agentic AI to policy experiments with governmental digital twins to identify and interpret signals—i.e., early indicators of potential disruption with strategic relevance to MI's specific context. From this, we synthesized the main signal clusters, such as agentic AI task management and bioadaptive work rhythms, that would form the conceptual basis of the narrative. The outcome of this step is shown in Appendix B.

Step 2: Structured Dialogue and Provocation

In our working sessions with MI team members, we used the signal clusters to structure dialogues around provocative "what-if" prompts. For example: "If a core team function could be fully automated by an agentic AI, what new value must the human team members provide? How do we measure it?" These conversations were designed to surface the core tensions, ethical dilemmas and strategic paradoxes that a simple trend analysis would miss.

Step 3: Co-Authoring the Narrative

The pivotal step was translating these complex dialogues into a concrete, human-centered story. We moved from asking "What if?" to a process of co-authoring the world of Emma, our main protagonist. As outside researchers, we drafted scenes based on our sessions; the MI member, as the inside expert, would then critique and enrich them, ensuring authenticity. This expert's firsthand knowledge of public sector constraints, political realities and ethical obligations was invaluable in shaping the narrative's realism. For instance, the detail of an AI's recommendation being checked by a human ethics advisor (Elle) was a direct result of the expert's insistence that any future system must have robust human oversight embedded within it.

Step 4: Using the Narrative as a Strategic Instrument

The final narrative is not a static report but a dynamic tool. It was refined through continuous feedback loops until it became a robust, co-owned artifact. It was used to communicate a complex future vision to other stakeholders within the ministry. The collaborative authorship of the narrative fostered a sense of ownership and relevance, and the narrative served as a strategic conversation instrument across internal stakeholders. As such, it allowed stakeholders to address questions like "What kind of organization do we want to become?"

Appendix B: Signal Clusters (Technologies) Forming the Conceptual Basis of the Narrative

The signal clusters in the table on the next page are the outcome of the horizon-scanning exercise at MI.

Technology	Description	Article Title	Available at:
Agentic AI	AI systems that autonomously perform complex, multistep tasks with minimal human intervention, enhancing decision-making and operational efficiency	*"With Autonomous Problem-Solving, Agentic AI Will Upend What You Consider Work"* *"Agentic AI: The Autonomous Revolution That Is Reshaping Enterprises in 2025"*	https://blogs.lse.ac.uk/businessreview/2025/02/11/with-autonomous-problem-solving-agentic-ai-will-upend-what-you-consider-work https://www.telusdigital.com/insights/ai-data/article/agentic-ai-in-the-enterprise
Holographic Overlay	Three-dimensional projections displaying digital information within physical spaces	*"Meta Unveils Augmented Reality Glasses Prototype 'Orion'"* *"Hologram: Realtime Holographic Overlays via LiDAR Augmented Reconstruction"*	https://www.ft.com/content/29dbf49b-da1d-4dc1-804e-f3ebc6d68e77 https://arxiv.org/abs/2405.07178
Polymorphic Collaboration Hub	Dynamic virtual meeting spaces that adapt to support global, multidisciplinary collaboration	*"7 Virtual Spaces Besides Zoom to Gather for Meetings, Events or Catch-ups"* *"Quantum, AI and Brain Computer Interface Powered Digital Twin Applications for the Future"*	https://www.buro247.my/lifestyle/technology/virtual-spaces-besides-zoom-for-meetings-events.html https://www.techuk.org/resource/quamtum-ai-and-brain-computer-interface-powered-digital-twin-applications-for-the-future.html
Sentiment-Adaptive Translation	Translation systems that adjust for emotional tone and cultural nuances	*"AI Tools May Soon Manipulate People's Online Decision-Making, Say Researchers"* *"Decoding AI's Nudge: A Unified Framework to Predict Human Behavior in AI-assisted Decision Making"*	https://www.theguardian.com/technology/2024/dec/30/ai-tools-may-soon-manipulate-peoples-online-decision-making-say-researchers https://arxiv.org/abs/2401.05840
Generative Simulations	AI-driven models that generate predictive scenarios to explore complex outcomes	*"Digital Twin: AI-Powered Simulation of New Zealand"* *"Towards Human-AI Deliberation: Design and Evaluation of LLM-Empowered Deliberative AI for AI-Assisted Decision-Making"*	https://www.esr.cri.nz/expertise/data-science-ai/digital-twin https://arxiv.org/abs/2403.16812

Technology	Description	Article Title	Available at:
Quantum Computers	Advanced computing systems using quantum bits (qubits) for high-speed calculations	*"Oxford Quantum Teleportation Breakthrough Brings Scalable Quantum Computing Closer to Reality"*	https://www.innovationnewsnetwork.com/quantum-teleportation-breakthrough-advances-quantum-computing/55351
		"Waveguide Holography: Towards True 3D Holographic Glasses	https://arxiv.org/abs/2211.02784
AI Advisor with Algorithmic Audits	AI tools that evaluate policies and audit algorithmic models for potential risks and unintended consequences	*AI-Enhanced Decision-Making"*	https://www.policechiefmagazine.org/ai-enhanced-decision-making
		"Three Challenges for AI-Assisted Decision-Making"	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11373149
Governmental Metamodels/ Digital Twins	Digital replicas of societal systems used to simulate and test policy interventions	*"Digital (T)winning: The Potential for AI as a Public Policy Design Tool"*	https://thepolicymaker.jmi.org.au/digital-twinning-the-potential-for-ai-as-a-public-policy-design-tool
		"AI Is Revolutionising Decision-Making, but it Can't Replace Human Leaders"	https://blogs.lse.ac.uk/businessreview/2025/01/03/ai-is-revolutionising-decision-making-but-it-cant-replace-human-leaders
Bioadaptive Work Rhythms	Systems that monitor biometric and cognitive data to adjust work schedules dynamically	*"Systems for Monitoring and Improving Biometric Health of Employees"*	https://patents.google.com/patent/US9693734B2/en
		"Why Humans and AI Assistants Need Relationship Advice"	https://smith.queensu.ca/insight/content/Why-Humans-and-AI-Assistants.php
Virtual Town Hall with Holographic Recreation	Digitally mediated public forums that recreate community spaces in holographic form for immersive engagement	*"4Dpresence Launches the Holographic Town Hall for Candidates"*	https://bbcmag.com/4dpresence-launches-the-holographic-town-hall-for-candidates
		"Hologram: Realtime Holographic Overlays via LiDAR Augmented Reconstruction"	https://arxiv.org/abs/2405.07178
Hyper-Personalised Training and Recruitment	AI-enabled systems that tailor training programs and recruitment processes to individual profiles.	*"Top AI Agents for Personalizing Employee Training and Development"*	https://pesto.tech/resources/top-ai-agents-for-personalizing-employee-training-and-development
		"Using AI for Enhanced Decision-Making: 9 Innovative Ways to Boost Board Efficiency and Effectiveness"	https://www.diligent.com/resources/blog/ai-for-decision-making

About the Authors

Alexander Richter

Alexander Richter (alex.richter@vuw.ac.nz) is a professor of information systems at Te Herenga Waka—Victoria University of Wellington, New Zealand. In his research, often in collaboration with partners in industry and government, he explores the future of work, with a focus on human-AI collaboration and the transformative impact of information technology. He applies value-driven and future-oriented design to develop context-aware and adaptive approaches for creating meaningful sociotechnical systems.

Shahper Richter

Shahper Richter (shahper.richter@auckland.ac.nz) is a digital marketing and AI researcher at the University of Auckland Business School. Her work focuses on how organizations can use emerging technologies like artificial intelligence to create meaningful, ethical and human-centered innovation. Blending academic rigor with practical insight, Shahper works closely with business leaders to help them navigate the realities of digital transformation and prepare their teams for the future of work.

Nastaran Mohammadhossein

Nastaran Mohammadhossein (nastaran@skillsvr.com) is a government project lead at SkillsVR, New Zealand, which provides solutions using virtual and augmented reality, and AI training technologies that replicate real-life experiences. Nastaran's work focuses on applying these technologies to enhance workforce capability and employability. Drawing on her academic background in information systems at Victoria University of Wellington, she bridges research and practice to translate emerging technologies into measurable public-sector value.