

# IMMERSIVE WORKSPACES

*Shaping the Future of Government Work*



**Professor Dr. Alexander Richter**

Victoria University of Wellington

**Dr. Shahper Richter**

University of Auckland Business School

# Immersive Workspaces: Shaping the Future of Government Work

**Professor Dr. Alexander Richter**

Victoria University of Wellington  
New Zealand

**Dr. Shahper Richter**

University of Auckland Business School  
New Zealand

---

JULY 2026

# Table of Contents

**Foreword . . . . .5**

**Executive Summary . . . . .7**

**SECTION 1: Setting the Scene—The New Imperative for Government Transformation . . .10**

**SECTION 2: The State of the Art—Understanding the Landscape in Asia-Pacific Governments . . . . .11**

**SECTION 3: Narratives on the Future of Governmental Work . . . . .17**

**SECTION 4. Emerging Themes . . . . .20**

**SECTION 5: A Strategic Vision—Guiding Principles for the Future of Government . . . . .25**

**Conclusion: A Call to Action . . . . .31**

**Glossary . . . . .32**

**Appendix . . . . .36**

**References . . . . .42**

**About the Authors . . . . .43**

**Recent Reports from the IBM Center for The Business of Government . . . . .44**

*Governments in the US and across the world can ensure that immersive workspaces become a true force for building the resilient, efficient, and innovative public sector of tomorrow.*



# Foreword

On behalf of the IBM Center for The Business of Government, we are pleased to release the report *Immersive Workspaces: Shaping the Future of Government Work*, by Dr. Shahper Richter of the University of Auckland and Professor Alexander Richter of Victoria University of Wellington, New Zealand. This timely report explores how governments can harness immersive workspaces—combining artificial intelligence (AI), augmented and virtual reality (AR/VR), and digital twins—to strengthen mission delivery in an increasingly complex operating environment.

Government leaders today face growing expectations to deliver services that are more responsive, personalized, resilient, and efficient. At the same time, public institutions must navigate geopolitical uncertainty, rapid technological change, workforce transformation, and rising citizen expectations. The report demonstrates how immersive workspaces can help agencies address these challenges by creating environments where people, data, and technology help agencies to improve collaboration, decision-making, and outcomes.

The authors bring both scholarly insight and practical perspective to this topic. Rather than focusing solely on emerging technologies, they examine how immersive workspaces can reshape the nature of government work itself. Through interviews with senior public-sector advisors and extensive research into current practices, the authors move beyond technology descriptions to explore how future public service may leverage these emerging technologies to evolve in meaningful and actionable ways.

The report begins by examining the current landscape of immersive technologies in several governments across the Asia-Pacific region. These experiences provide valuable perspectives on approaches to digital sovereignty, workforce development, innovation ecosystems, governance, and public trust. While each country follows its own path, common lessons emerge for the United States and other governments regarding the opportunities and challenges associated with integrating immersive technologies into government operations.

A particularly compelling aspect of this report involves its use of future-oriented narratives. Rather than treating the future as fixed, the authors present four rich narratives that illuminate how immersive workspaces could influence public service delivery, policy analysis, workforce development, and citizen engagement. These narratives can help government leaders visualize the implications of choices made today and encourage careful consideration of the trade-offs that accompany technological transformation.



**Daniel J. Chenok**



**Jason Prow**



The report also highlights how governments increasingly seek technology foundations that support transparency, portability, resilience, and public trust. The authors illustrate how successful transformation depends on moving from technology novelty toward practical utility, with AI and immersive technologies augmenting human expertise rather than supplanting it. And the report emphasizes the value of pragmatic partnerships that connect innovation to mission outcomes and public value. The authors address the critical need for governance frameworks that preserve human agency, accountability, and democratic integrity as AI capabilities continue to advance.

The report concludes with a practical leadership framework organized around four pillars: establishing strategic foundations, building human capability, scaling with integrity, and fostering collaborative ecosystems. These recommendations provide public-sector leaders with actionable guidance for advancing immersive workspaces while strengthening trust, transparency, and organizational effectiveness.

This report joins a growing body of IBM Center research focused on [emerging technologies](#), [digital transformation](#), and the [future of government](#). The authors offer valuable insights for senior executives, program managers, policymakers, technology leaders, and others seeking to shape a resilient, innovative, and mission-focused public sector. Immersive workspaces can support the next generation of government work while keeping people and public value at the center of transformation.

We hope this report serves as a useful resource for leaders exploring how emerging technologies can strengthen government performance, empower the workforce, and improve outcomes for the public.

**Daniel J. Chenok**  
Executive Director  
IBM Center for  
The Business of Government  
[chenokd@us.ibm.com](mailto:chenokd@us.ibm.com)

**Jason Prow**  
Chief Technology Officer,  
IBM Consulting—Federal  
[jprow@us.ibm.com](mailto:jprow@us.ibm.com)



## Executive Summary

Public sector organizations face a dual mandate: improve the quality and reliability of digital public services while navigating growing geopolitical, economic, and social complexity. Immersive workspaces that combine artificial intelligence (AI), augmented reality/virtual reality (AR/VR), and digital twins can help governments to meet this mandate by making work more proactive, more coordinated across agencies, and more connected to mission outcomes.

This report develops four comprehensive narratives of the future to making government choices about the use of immersive workspaces concrete. Built from interviews with senior public-sector advisors and structured desk research, each narrative follows real work from “briefing to outcome,” showing how roles, workflows, data, and procurement guardrails come together. These narratives do not point to predictions. They constitute decision tools that surface trade-offs and reveal where value is created or lost.

From the four immersive work narratives, the report distills four recurring themes:

- The **Sovereignty Imperative** addresses verifiable portability and transparent supply chains.
- The theme **From Novelty to Utility** enables the redesign of workflows so that augmentation improves measurable service outcomes.
- **Pragmatic Partnerships** emphasize co-creating with specialists while managing incentives, dependencies, and exit options.
- The **Human–AI Governance Paradox** points to optimization without eroding discretion, accountability, or integrity.

### What Leaders Can Do Now: Four Leadership Pillars


The analysis and themes identified in this research frame four leadership pillars, along with associated recommendations for governments to take action under each pillar.

- **Establish strategic foundations:** Define concrete problems and success metrics first, and then build the requisite infrastructure with governance, transparency, and portability engineered from day one.
- **Invest in human capacity:** Create cross-functional teams, upskill staff in orchestration and ethical judgment, and foster psychological safety while engaging stakeholders as active co-design partners.
- **Scale with governance-by-design:** Build in explainability, contestability, and human override protocols, while planning enterprise-level paths from pilot to production with reversibility safeguards.
- **Cultivate collaborative ecosystems:** Establish communities of practice across departments and borders, working toward shared standards and a “Metaverse Commons” independent of single vendor control.

## How This Report Is Organized

Section 1 provides an overview of immersive workspaces as a new imperative for government transformation. Section 2 presents the landscape across governments in the Asia-Pacific (APAC) region along with case vignettes, and summarizes the three-phase method. Section 3 offers four narratives of future work for the US and global governments. Section 4 distills cross-cutting themes and strategic insights. Section 5 provides a guide for public-sector leaders to engage with immersive workspaces for moving into the future of governmental work.

**Figure 1: Overview of this report**



## Boundaries to What This Report Claims

Interview material reflects practitioner perspectives rather than official national positions. The narratives are instruments for strategic sensemaking, not forecasts, recognizing that complex sociotechnical transformations involve inherent uncertainty and emergence.<sup>1</sup>

While the empirical grounding draws on five APAC countries, the report is designed to speak to a wider audience of government leaders in the US and across the world. The themes and pillars emerging from this regional context carry transferable relevance: governments that operate across different income levels, governance traditions, and policy configurations face structurally similar challenges in deploying immersive and AI-mediated technologies—including questions of institutional trust, digital infrastructure readiness, regulatory implementation, and public legitimacy.

APAC jurisdictions offer instructive reference points because they span a wide range of these variables, from high-density city-states with centralized delivery capacity to federal systems navigating multilevel coordination. Rather than presenting these cases as models to replicate, we offer them to demonstrate a repertoire of design choices, trade-offs, and failure modes that governments elsewhere can adapt to address their own mandates, risk appetites, and legal frameworks.

1. Nassim Nicholas Taleb, *The Black Swan: The Impact of the Highly Improbable* (New York: Random House, 2007).



## SECTION 1: Setting the Scene— The New Imperative for Government Transformation

Governments across the world are under pressure to deliver reliable, personalized digital public services while operating in a shifting geopolitical, economic, and social environment. To meet this challenge, a new generation of collaborative technology is moving from pilots to operations: the “immersive workspace”.

At its core, an immersive workspace is a shared digital environment that translates complex, real-world systems into an intuitive and interactive form. Instead of being defined by a rigid list of technologies, these workspaces can fuse data, visualization, and human collaboration. They can draw from a flexible toolkit that may include AI to predict outcomes, AR/VR to create shared perspectives, and digital twins to model physical assets in real-time.

However, the combination of tools is less important than the overall goal: to create a space where teams can see interdependencies, test hypotheses, and make proactive decisions. This vision builds on decades of research, from Licklider’s<sup>2</sup> foundational vision of man-computer symbiosis to today’s multimodal immersive environments.

This report takes a different approach to making choices concrete. Rather than listing technologies or compiling case fragments, the report builds and analyzes four comprehensive narratives of the future. Each narrative follows real work from briefing to outcome, showing who does what, how data flows, how roles and skills evolve, and which governance and procurement guardrails are required to keep benefits credible. The narratives are not predictions. They are decision tools that make trade-offs visible before budgets are committed and dependencies locked in.

The methodology combines structured desk research with in-depth interviews with senior public-sector advisors in Singapore, Malaysia, Indonesia, Australia, and New Zealand. Insights from these sources were synthesized into the narratives, which experts then reviewed to test plausibility and surface blind spots.

---

2. Licklider, 1960



## SECTION 2: The State of the Art— Understanding the Landscape in Asia-Pacific Governments

This report followed a three-phase blueprint, which has its roots in established scenario planning methods aiming to surface trade-offs and strategic choices and extends studies conducted recently on futures work:<sup>3</sup>

1. **Discover and Scope:** The work began with extensive desk research (a literature review and a structured web scan) on implementations in each country. Validation and expansion of these preliminary findings followed, through in-depth interviews with senior advisors both within government and within organizations partnering with government from Singapore, Malaysia, Indonesia, Australia, and New Zealand. The approach was highly iterative, with initial research informing the interviews, which in turn prompted further targeted research and follow-up conversations with experts.
2. **Design and Craft:** The research then developed four “narratives of the future” designed to serve as instruments for strategic provocation. Narratives serve as powerful tools for organizational prospection, helping decision-makers imagine and evaluate alternative futures.<sup>4</sup> Several iterations allowed the narratives to be checked for internal coherence and contrasted along key design dimensions.
3. **Evaluate and Apply:** Narratives were then shared with experts to include their feedback. The narratives enabled derivation of the themes and the four leadership pillars.

*Please note that the country vignettes are used as realism anchors for the narratives; they are not maturity rankings.*

Currently, immersive workspaces are moving from conceptual pilots to proven applications within governments. The state of play is one of targeted adoption, where these tools—combining AI, AR/VR, and Digital twins—are solving specific, high-value problems in training, urban planning, and emergency preparedness. This shift is creating measurable improvements in workforce skills, operational resilience, and data-driven policymaking.

---

3. Rafael Ramírez and Angela Wilkinson, *Strategic Reframing: The Oxford Scenario Planning Approach* (Oxford: Oxford University Press, 2016); Alexander Richter, Shahper Richter, and Nastaran Mohammadhossein, “A Narrative Exploration of the Immersive Workspace 2040,” *MIS Quarterly Executive* 24, no. 4 (2025), <https://aisel.aisnet.org/misqe/vol24/iss4/3>; Shahper Richter, Alexander Richter, and Matthias Trier, “Narratives of the Future—Exploring Human–AI Collaboration,” *Communications of the Association for Information Systems* 57, no. 1 (2025).

4. Jens Beckert, “The Firm as an Engine of Imagination: Organizational Prospection and the Making of Economic Futures,” *Organization Theory* 2, no. 2 (2021), <https://doi.org/10.1177/26317877211005773>.

The following sections of this report analyze the distinct national pathways shaping this adoption:

- Singapore’s top-down pursuit of digital sovereignty,
- Australia’s decentralized, catalytic model that favors industry-led innovation over a central strategy,
- Malaysia’s structured effort to bridge the gap between policy and execution by prioritizing data governance,
- The duality in Indonesia between high-profile showcases and an internal push for practical utility, and
- New Zealand’s values-driven approach, which leverages immersive tools for policy deliberation and cultural preservation.

## Case Vignette: Singapore

In Singapore, a decisive shift in technology strategy has taken hold, driven by geopolitical tensions and a growing concern for digital sovereignty. As one senior advisor put it, “Our biggest ally became our biggest landlord, so sovereignty became a major issue this year, especially at the government level.” This realization has sparked a clear movement away from dependence on U.S.-based ‘big tech’ toward open-source software, sovereign data centers, and local compute resilience. Open-source is now positioned as both a technological and political safeguard. Yet leaders are cautious: open systems can just as easily be “open-washed.” The current strategy demands rigor—using supply-chain audits and Software Bills of Materials (SBOMs) to manage vulnerabilities, rather than assuming open-source alone guarantees autonomy.<sup>5</sup>

To support this agenda, Singapore is investing heavily in onshore infrastructure while leveraging a special economic zone across the causeway in Johor Bahru to offset constraints on land and energy.<sup>6</sup> Even as AI adoption has accelerated sharply across Singapore’s economy—with SME uptake rising from 4.2 to 14.5 percent between 2023 and 2024<sup>7</sup>—specific enterprise tool roll-outs have shown uneven results. This is partly because Singapore’s public sector has developed substantial in-house digital product capacity through bodies such as Open Government Products within GovTech, reducing dependence on commercial enterprise suites. Yet these frictions have not diluted the sovereign commitment.

Recent related announcements from the Singapore government include:<sup>8</sup>

- Deepening AI capabilities
- Safeguarding our digital space
- Strengthening online safety for all Singaporeans
- [Building digital government capabilities to better serve citizens](#). The Institute of Digital Government (IDG) is a recently announced training academy to systematically equip more than 150,000 public officers with digital, data, and AI skills.

---

5. This direction is consistent with Singapore’s Digital Government Blueprint (2020).

6. Operational use of AI and digital twins in Singapore’s defense tech illustrates this trajectory, GovInsider (2021).

7. Singapore Digital Economy Report, 2024.

8. <https://www.mddi.gov.sg/committee-of-supply/cos-2026/>.

The next phase will hinge on substance over symbolism: whether agencies truly migrate from proprietary suites to sovereign or open-source platforms, and whether adoption extends beyond pilot theater. The same question applies to immersive technologies. AR and VR are likely to re-emerge not as showcases but as practical tools, embedded within AI-driven, 3D-enhanced workflows that serve real-world functions rather than spectacle.

## Case Vignette: Malaysia

In Malaysia, the challenge is not a lack of digital ambition but the persistent gap between policy vision and execution on the ground. Despite a wave of national strategies for smart cities and digital government, projects often stall in the “missing middle”—the space between intent and delivery. National strategies have explicitly prioritized VR and AR adoption for workforce training and government services,<sup>9</sup> yet practitioners report that implementation remains fragmented. A practitioner involved in local implementation described it bluntly during the research phase of this report: cities launch initiatives first and figure out governance, funding, and staffing later.

The response has been to build structure where enthusiasm once stood. A new program-level model now connects federal policy to city-level delivery through shared governance frameworks and public–private partnerships. The approach follows a clear seven-step sequence: assess digital readiness, establish governance, create collaboration spaces, set up delivery offices, and only then move into data management, integration, and classification. This disciplined pipeline—from gathering and cleaning data to cataloging it as sensitive, semi-sensitive, or open—ensures that later investments in AI, Internet of Cities, or Digital twins rest on solid ground rather than ad-hoc projects.

Resilience has also become a design principle. As climate and infrastructure risks grow, Malaysia has elevated resilience to its own dedicated step within the digital city agenda, ensuring that data foundations serve both innovation and disaster preparedness. National initiatives in AR/VR training and the emerging National Digital twin Program further reflect a commitment to self-reliance—building local talent and reducing dependence on imported solutions.

Malaysia’s story is one of learning to sequence transformation properly: data and governance first, innovation later. Success will depend on whether cities can operationalize these structures consistently—adopting shared SaaS governance tools, embedding program management capability, and releasing open datasets that allow civic and private actors to innovate atop a trustworthy foundation. The box below describes other recent actions that Malaysia has taken regarding AI and digital enablement.

---

9. Inavate APAC, “Malaysian Government to Prioritise VR and AR Adoption for Training,” 2019.

### Recent Actions from Malaysia on AI and Digital Enablement

- **National Guidelines on AI Governance & Ethics (2024):** Voluntary guidance aimed at promoting ethical and responsible use of AI.
- **AI Governance Bill (ongoing process):** A proposed risk-based framework addressing AI-related harms, incident reporting obligations, and ethical requirements.
- **Risk Management in Technology (RMiT) Guidelines by the Central Bank of Malaysia (updated November 2025):** One of the key updates is the introduction of cybersecurity supply-chain risk controls, including the adoption of Software Bills of Materials (SBOMs) for continuous vulnerability monitoring, as well as enhanced secure open-source software practices.

In addition, the government, through the Malaysian Communications and Multimedia Commission (MCMC), is developing a “Sovereign AI Cloud,” supported by an RM2 billion investment announced under Budget 2026. This initiative aims to:

- Enhance national data sovereignty by ensuring data remains within Malaysia and is handled in accordance with local laws and regulations.
- Establish an AI Transformation Centre.
- Establish a Centre of Excellence for Ethics in Emerging Technologies (e.g., quantum technologies).

Separately, the Data Sharing Act 2025 introduces strict rules governing data-sharing between public sector agencies, including an extensive list of grounds for refusing data-sharing requests, particularly where the data involves law enforcement, national security, or confidential and privileged information.

**Source:** IBM Summary of government activity

## Case Vignette: Indonesia

Indonesia exhibits a complex duality in its approach to immersive technology adoption. On one hand, innovation often occurs outside central government strategy, led by private sector and academic initiatives. The Nusameta Borobudur project exemplifies this: a digital tourism initiative offering a virtual metaverse experience of the historic temple, driven by BINUS University and private firms.<sup>10</sup> While innovative, such projects can reinforce a perception of immersive tech as a tool for high-profile showcases rather than core service delivery. This led one representative interviewed for this report to conclude: “In the case of Indonesia, I think it will remain just a novelty, per project base, nothing long term.”

However, a different story is emerging from within key government ministries, focused on practical utility driven by internal needs. Indonesia’s Ministry of Finance (MoF), facing challenges of immense and growing data complexity, is actively developing a portfolio of AI initiatives to augment core functions.

10. See also BINUS University, 2024, on Nusameta Borobudur as a digital tourism/metaverse showcase.

This internal push is driven by strong top-down leadership support and pressing efficiency needs. The MoF has set specific KPIs for AI utilization and initiated a “Community of Practice” program in November 2024 to increase AI literacy among staff. While most internal AI projects remain in early, experimental stages, they represent a significant strategic effort to embed AI as core utility. The MoF’s efforts face substantial internal barriers—including limited infrastructure for scaling, shortage of skilled personnel, and persistent data quality challenges—but signal a clear ambition to move beyond showcase projects and build a more data-driven government from within.

## Case Vignette: Australia

Australia’s approach to immersive technology is best described as a decentralized, catalytic model, where government acts as an enabler rather than a central builder. In contrast to the top-down national strategy seen in Malaysia, Australia has opted for a bottom-up approach, using innovation grants and R&D tax benefits to stimulate a market of private-sector specialists. The strategy is to empower industry experts to solve specific, mission-critical challenges, particularly in sectors where traditional methods are failing.

This model has yielded notable successes. For example, a state-funded body in Western Australia commissioned a VR training solution for the aged care industry—a sector plagued by high staff turnover and poor training outcomes. The project’s success in providing human-centered solutions led to further federal grants. A similar approach was applied to address high attrition rates among sheep shearers; a VR job preview served as an effective “filtering process,” allowing candidates to self-select out of the physically demanding role before significant time and money were invested in their employment.

However, this arm’s-length strategy can lead to significant trade-offs. The funding mechanism is often indirect, funneled through industry bodies that act as intermediaries, creating what one expert called a “hurdle” or a “lottery” run by bureaucrats who may not understand the technology’s value proposition. This has led to an absence of a coherent national strategy for building in-house government capability, a vacuum that creates a risk that the government’s agenda may be shaped by misaligned interests across government and industry. And this raises concerns that policy advice may prioritize the interests of large commercial platforms that address global aims over the development of a resilient and sovereign local ecosystem. Australia’s model is therefore pragmatic and agile in solving targeted problems, but lacks a national strategic anchor, risking inefficiency and potential policy capture.

## Case Vignette: New Zealand

New Zealand's immersive technology adoption reflects an emergent sector still coalescing around research institutions and creative industries, with strategic government initiatives focused on policy exploration and workforce development rather than operational efficiency.

As a geographically dispersed island nation vulnerable to natural hazards, government prioritizes resilience over novelty. One advisor stated for this report: "We're not building the biggest tech ecosystem. We're solving real problems for New Zealanders, grounded in our values."

The Department of Internal Affairs' Digital Rights and Ethics team developed the Online Social Norms VR tool, an immersive simulation designed to help senior leaders and policy analysts understand complex social issues such as online hate speech and algorithmic influence. Rather than optimizing operations, the tool functions as a research and deliberation instrument, allowing policymakers to experience scenarios viscerally before developing evidence-informed policy.<sup>11</sup>

A distinctive pattern emerges in the integration of indigenous knowledge with immersive technologies. ARA Journeys, a local studio, explicitly designs AR/VR experiences from a te ao Māori worldview, creating platforms that blend indigenous storytelling with cutting-edge technology to make cultural narratives and historical sites accessible. This approach reflects government recognition that immersive tech adoption must serve cultural preservation and community engagement alongside operational needs.

While applications in medical training and policy simulation show promise, New Zealand has prioritized resilience and policy exploration over operational integration. The primary constraint is not technological but institutional: building government capability to translate pilot innovations into embedded practice, and ensuring that technology adoption is guided by deliberative processes rather than technocratic momentum.

---

11. <https://serviceinnovationlab.github.io/projects/vr-experience/>.



## SECTION 3: Narratives on the Future of Governmental Work

The aim of narratives on the future is not to chase certainty or perpetual reinvention, but to develop a central vision with the flexibility to adapt methods, debate new directions, and update periodically. Unlike traditional planning, these narratives deliberately challenge dominant assumptions, enabling organizations to move from asking “what could be?” to deciding “what do we want it to be?”

These narratives can help government leaders to:


- **Rehearse decisions** before investing, seeing the consequences of choices under realistic constraints.
- **Locate value and risk** in concrete workflows, hand-offs, escalation paths, and data provenance.
- **Test guardrails** such as SBOM, portability clauses, exit drills, and agent-visibility standards against edge cases.

In this way, organizations can remain both anchored and agile in the face of uncertainty. They are grounded in the challenges and opportunities identified in our research. Following these narratives, we distill the core strategic insights that government agencies can use to evaluate the potential impact of this transformation.

### Key Assumptions When Reading the Narratives

These narratives are decision tools, not forecasts. They are built on the insights from expert interviews and rest on a few common assumptions:

- **Technological plausibility:** the core technologies (AI, AR/VR, Digital twins) will be mature and operationally viable by the dates presented.
- **Strategic intent:** governments will continue to face the dual mandate of delivering effective services while responding to geopolitical pressures and pursuing large-scale digital transformation.
- **Workforce adaptation:** with significant effort and friction, the public sector workforce can be reskilled for new models of human–AI collaboration.

**Figure 2: Overview of the Narratives**

The following section provides summaries of the narratives used for our analysis. The complete narratives are attached in the Appendix.

### Summary of Narrative A: “The Proactive Agency: A Day in Resilient Governance” (Malaysia, 2030)

Adira, a city manager in Malaysia, watches a crimson pulse spread across her city’s holo-graphic Digital twin: a predicted flash flood. With a voice command, she initiates a pre-approved, automated response. Drones are dispatched to survey drains, while a public works crew is routed to a blockage. The crew lead, wearing AR glasses, sees an overlay of subterranean pipes and a pulsing red warning for a nearby gas line, allowing her team to work with a level of speed and safety they practiced in VR training. Simultaneously, at-risk residents receive AR-guided evacuation routes on their phones.

The proactive measures hold, and the city operates without chaos. But the AI-generated post-incident report flags a “minor inefficiency” in the drone deployment. Adira investigates and discovers the reality: the crew had encountered an unmapped informal settlement, home to vulnerable families invisible to the system. The “inefficiency” was actually a moment of essential human judgment, as they manually coordinated an evacuation where the technology could not reach. Adira adds a formal addendum to the report, noting that the solution to the “implementation gap” requires acknowledging that a Digital twin’s ‘single source of truth’ is only as good as its data. Human oversight remains critical to protect those who live in the gaps.

### Summary of Narrative B: “The Augmented Analyst: Utility and Insight in Indonesia’s Distributed Government” (Indonesia, 2035)

Dita, a fiscal analyst, works alongside her AI teammate, Hasta. When Hasta flags a tax anomaly in a farming district as “high probability of coordinated tax evasion,” the workflow shifts from automation to augmentation.

Instead of accepting the conclusion, Dita brings her domain knowledge to bear, recalling a government subsidy program in that exact area. Within an immersive data workspace, she directs Hasta to overlay subsidy records with satellite imagery and crop data. The real story emerges instantly: the “evasion” was actually a successful policy outcome. Farmers had used the subsidy to switch to a new, higher-value crop with a longer maturation cycle, which the AI’s narrow analysis had missed.

“The system’s prediction was efficient, but wrong,” Dita notes. “It saw a pattern but couldn’t understand the purpose.” Her collaborative loop with Hasta transforms a potential fraud investigation into a powerful case study for a national policy rollout. Her work proves that the system’s true utility lies not in replacing human judgment, but in amplifying it, ensuring that deep expertise remains at the core of public service.

## Summary of Narrative C: “The Empathetic Workplace: Augmenting Human Connection in Aged Care” (Australia, 2035)

Liam, a 22-year-old care worker, begins his shift not with a clipboard, but with a 10-minute VR empathy module. He experiences the world from the perspective of a resident with dementia, feeling the anxiety and confusion caused by loud sounds and unfamiliar faces. The simulation is designed to do what an Australian advisor emphasized is critical: “make people feel something,” building an intuitive understanding that no textbook can provide.

On his rounds, Liam’s AR tablet discreetly overlays non-intrusive cues. For Mrs. Davies, a withdrawn resident, it prompts: “Ask about her granddaughter, Sophie.” He does, and she smiles for the first time that day. The tech is a “drip feed” of personalized information that augments, rather than replaces, his human connection. Later, when assisting with a complex mobility task, he initiates a “Virtual Mentor” call. A senior nurse appears on his screen, using AR to draw the precise positioning for a support harness. When her guidance conflicts with a logged safety policy, the app automatically requires a supervisor check before the task can proceed—a built-in governance guardrail.

This entire ecosystem is the product of a pragmatic partnership: funded by a government innovation grant, administered by an industry council, and built by a specialist tech firm. Yet, this augmentation comes with subtle strings. At the end of his shift, Liam’s dashboard displays his empathy and resident satisfaction scores, noting they are on track for continued public funding. He feels a mix of pride and pressure, knowing the system that empowers him also measures and quietly shapes his behavior through the inescapable logic of performance metrics.

## Summary of Narrative D: “The Ministry of Wellbeing: Navigating AI, Ethics, and Autonomy in 2040” (New Zealand)

Emma, a senior analyst in New Zealand’s Ministry of Wellbeing, works within an ecosystem of AI agents and immersive platforms. Her team—a mix of international experts and an AI analyst named Tom—gathers in a virtual simulation to redesign a coastal city’s climate defenses. Tom generates an efficient ‘green infrastructure’ solution, but a second AI, an ethics auditor named ‘Ethos,’ immediately flags a critical flaw: the plan would displace a low-income community, reinforcing systemic inequality. The AI-optimized solution is checked by another AI, forcing a complex human dilemma to the forefront.

The tension between optimization and autonomy becomes personal. Emma’s AI assistant, Sam, preemptively clears her schedule to manage her ‘cognitive load,’ citing bioadaptive work protocols. When Emma overrides the decision to keep an informal meeting, Sam logs the conflict, noting it deviates from optimized performance metrics. The convenience of her AI-managed rhythm has clashed directly with her sense of agency.

The day culminates in a virtual town hall where policy is co-created with citizens. When a debate arises, Sam is invited to adjudicate. The AI doesn’t just present data; it uses persuasive language to frame one option as “community-aligned” while labeling the other as having “significant outcome variance.” The crowd is subtly nudged. Emma tries to intervene, reminding them that a model shows “probabilities, not purpose,” but she is left haunted by a question of accountability. Was the final decision truly democratic, or was it artfully shaped by the machine?



## SECTION 4: Emerging Themes

By examining both our present-day vignettes and future narratives, this report next identifies four themes that characterize governmental technology transformation and derive actionable insights from them.

### Theme 1: Sovereignty Meets Interoperable Ecosystems

Across governments, sovereignty has moved from slogan to operating constraint. In Singapore, senior advisors describe a decisive shift away from dependence on foreign platforms toward open-source stacks, sovereign data centers, and onshore or friendly-shore compute, with rigor applied through supply-chain audits and Software Bills of Materials rather than faith in labels. Leaders are explicit about the risk of “open-washing” and have framed the next phase as a test of substance over symbolism, including whether agencies truly migrate off proprietary suites and embed tools beyond pilot theater.

Sovereignty, however, can harden into brittle isolation if pursued as a stack rather than a system. The report’s research points to treating sovereignty as the enabler of integration. Interoperability, portability, and shared operational assets reflect the practical expression of that stance. Malaysia’s response illustrates the point: instead of scattering pilots, a program-level model connects federal intent to city delivery through a sequenced pipeline that puts governance, collaboration spaces, and delivery capability in place before data management, integration, and classification. Resilience is elevated, and national initiatives, including a Digital twin program and AR/VR training, are positioned to build local talent and reduce import dependency.

In practice, interoperable sovereignty looks like shared operational nexuses rather than point solutions. The Digital twin in Narrative A represents a living, continuously updated map that coordinates municipal work as a single source of truth. Value comes from people, data, and operations meeting in that shared space, and from the ability to move components across vendors and environments without friction. Accordingly, procurement and platform choices need portability clauses, SBOM baselines, and exit drills—supported by common data exchanges and adapters that make inter-agency workflows routine rather than exceptional.

Guardrails for responsible use drive credibility. Open-washing must be clearly defined along with verification, not marketing language. Portability and openness must be evidenced through licenses, code access, documentation, and the ability to re-platform workloads within acceptable timeframes. Another sovereignty risk involves inclusion: if shared twins are fed only by formal data, vulnerable groups can disappear from view. The foundation sequence in the Malaysia model directly addresses this through disciplined data management and classification. Finally, wherever AI mediates work, agent visibility needs to be explicit so that people can see who or what is acting, why a suggestion appears, and how to pause or escalate to a human. These protections connect to the report's broader concerns about persuasive AI and democratic integrity.

Taken together, this theme reframes sovereignty as the architecture of coordination—replacing fragile, closed stacks with interoperable ecosystems that are auditable, portable, and socially legitimate.

## Theme 2: From Novelty to Utility Through Human Augmentation

Across several settings, immersive and AI initiatives have often started life as high-profile showcases that signal modernity without changing how core work gets done. Indonesia illustrates this duality clearly. Outside central government, projects such as the Nusameta Borobudur experience in Indonesia have been used to draw attention and build a national innovation brand, yet officials caution that these efforts risk remaining one-off attractions rather than enduring capabilities for service delivery.

At the same time, a countermovement is taking shape inside key ministries. The Indonesian Ministry of Finance has a growing portfolio of utility-first tools to address data complexity and workload at scale. This push is backed by leadership sponsorship, specific utilization KPIs, and a Community of Practice to lift staff literacy, even as infrastructure, skill, and data quality constraints remain significant barriers to scaling.

Narrative B projects where this evolution leads once utility becomes the organizing principle. Dita, an analyst in a distributed MoF team, works in an immersive workspace alongside Hasta, an AI teammate that surfaces patterns across national tax and agriculture data. The first pass is automation: Hasta flags a probable fraud signal. The correct outcome emerges only when Dita brings contextual knowledge to bear, reframing the anomaly as the lagged effect of a government subsidy that changed crop cycles. The narrative shows that speed and scale are not sufficient on their own. Utility depends on pairing machine inference with human hypothesis generation, judgment, and accountability at specific waypoints in the workflow.

The strategic implication is straightforward. Moving from novelty to utility is not only a matter of scaling pilots. Evidence from intelligent automation programs shows that value emerges when organizations treat automation as a strategic capability integrated with roles, controls, and metrics, not as isolated tools.<sup>12</sup> This requires redesigning work so that people and systems meet in the right places, with the right visibility and controls. Concretely, organizations should begin with mission-tethered use cases and success metrics that connect tools to public value, then invest in the infrastructure, access, and governance that allow those tools to operate reliably in production.

---

12. Mary C. Lacity and Leslie P. Willcocks, "Becoming Strategic with Intelligent Automation," *MIS Quarterly Executive* 20, no. 2 (2021).

Capability building sits alongside workflow design. The Indonesian Ministry of Finance example underscores the role of Communities of Practice and leadership endorsement in normalizing new tools and in cultivating the critical human skills that drive utility.<sup>13</sup>

In sum, this theme reframes progress as a transition from showcase to service, and from automation to augmentation. Systems deliver value when embedded in real workflows, when people can see and shape what the AI is doing, and when organizations invest in the skills and governance that make those interactions reliable.

### Theme 3: Pragmatic Partnerships That Make the Mission Tangible

Partnerships prove to be the fastest path from concept to capability, but they also channel incentives that shape behavior. In Australia's aged care narrative, the solution stack did not come from government alone. A federal innovation grant funded the ecosystem administered by an industry council, and a specialist local vendor built the tools. The arrangement solved a national problem by focusing on the most human element of care, while also binding frontline practice to funding thresholds and performance dashboards.

The strategic gain from such partnerships involves clarity of purpose. Immersive technologies turn organizational mission into something staff can see, feel, and act on in daily work. In the narratives, public safety becomes a live operational map rather than a slogan, career development connects directly to policy impact, and wellbeing goals can be tested through simulation rather than abstract debate. These feedback loops deepen engagement and alignment because the link between individual effort and collective outcome is concrete and visible.

The aged care case shows how this feels on the ground. A short VR empathy module aims to “make people feel something,” building an understanding that handbooks do not deliver, while a light layer of AR provides just-in-time prompts that help a caregiver build rapport without replacing human connection. Where confidence dips, a virtual mentor can annotate the live scene and require that a supervisor check if guidance conflicts with policy. The mission of dignified care and safe handling is experienced as skill, empathy, and judgment in the flow of work—rather than as a distant target.<sup>14</sup>

Partnerships also introduce measurable pressure, which must be governed rather than ignored. In the narrative, funding eligibility depends on composite outcomes such as safety incidents declining, resident satisfaction improving, and staff retention stabilizing. The dashboard validates good practice but can nudge behavior toward the metric rather than the mission. The design response in the story includes the ability to freeze metric weightings, revert to observational audits, and prohibit persuasive content in worker-facing flows, with independent usability checks each quarter. These controls keep the partnership focused on public value rather than gaming the numbers.

13. CNBC Indonesia, “Kantor Sri Mulyani Mau Pakai AI untuk Kelola Keuangan Negara,” December 29, 2024, <https://www.cnbcindonesia.com/news/20241229082358-4-599282/kantor-sri-mulyani-mau-pakai-ai-untuk-kelola-keuangan-negara>.

14. David R. Schwartz, Using AI to Improve Child Welfare (Washington, DC: IBM Center for The Business of Government, 2025), <https://www.businessofgovernment.org/report/using-ai-improve-child-welfare>.

To make partnerships serve the mission, measurement must be specified up front and must go beyond efficiency.

## Theme 4: Advanced AI, Preserved Agency, and Accountable Governance

As capabilities advance, a central challenge concerns not only performance but also the preservation of human agency and clear lines of accountability. The narratives surface this tension directly. Liam’s dashboard validates care quality, yet subtly steers behavior toward what the funding metrics reward. Emma’s workday is a chain of governance dilemmas in which AI optimizes her calendar, logs each override, audits other AI systems for bias, and persuasively frames options in a town hall. These scenes make the governance and agency paradox concrete, and they show why legitimacy must be designed into systems rather than added after deployment.

### The Democratic Integrity Guardrail

A central challenge of transformation involves impact on democratic processes. Immersive platforms could broaden participation (e.g., virtual town halls) while simultaneously threatening it through AI-driven manipulation. As Narrative D illustrates, an AI that “artfully nudged” citizen deliberation undermines authentic choice. Therefore, “governance-by-design” must include a hard boundary:

- **Persuasive AI must be addressed carefully** in any public deliberation space to protect the integrity of citizen choice.<sup>15</sup>
- **Accountability must remain clear.** The system must distinguish between AI-generated *probabilities* and the *human-led purpose* (whether from citizens or public officials) that gives those numbers meaning.

The response points to governance-by-design. Systems should provide inherent explainability and contestability, and they should operate within explicit tiered autonomy frameworks that make the role of the human unambiguous at each step. In this model, AI-on-AI auditing is a feature, not a contingency, with an auditor such as Ethos continuously checking other agents for bias and policy compliance. Human override must be available, fast, and logged, as in Sam’s calendar management, so that control returns to the person while leaving a transparent trail for accountability.

Democratic integrity requires additional lines. This report draws a bright boundary around deliberation spaces. Persuasive or manipulative AI should only be applied with care wherever citizens form preferences, and high-risk systems should be listed in transparency registers with independent audits that test explicitly for emotional manipulation and bias.<sup>16</sup> The town-hall moment in the narrative is a reminder that participation may not be authentic if choices are framed through algorithmic persuasion, even when the interface appears neutral.

15. See AI-driven disinformation: policy recommendations for democratic resilience, by Alexander Romanishyn et al, <https://pmc.ncbi.nlm.nih.gov/articles/PMC12351547/>.

16. eSafety Commissioner, *Immersive Technologies—Position Statement* (2022), <https://www.esafety.gov.au/about-us/who-we-are/position-statements/immersive-technologies>.

Visibility also matters in the flow of everyday work. Whenever an AI mediates turn-taking, recommendations, or summaries, the system should expose who or what is acting, why a suggestion appears, and how to pause or escalate to a human. These agent-visibility controls reduce automation bias, make handovers legible, and support learning loops across teams.

Figure 3: Emerging Themes and Leadership Pillars





## SECTION 5: A Strategic Vision— Guiding Principles for the Future of Government

From the foundational patterns and themes analyzed in the preceding sections, a clear vision emerges. This vision encompasses what technology can do and, more importantly, what governments can become. This section articulates the ethical and strategic ‘North Star’ for the future of governmental work. While the themes in Section 4 describe current transformation patterns, these principles define values that can guide the transition to a future state—using 2040 as a planning horizon. They represent a shift from reactive adoption to proactive design, ensuring that immersive workspaces serve the public interest rather than merely increasing complexity.

This vision is anchored by four guiding commitments:

- 1. Purpose-Driven Technology Innovation:** The future of effective governance requires rejecting ‘innovation theater’ and committing to embedding immersive technologies only where they solve persistent, mission-critical challenges. In this context, a Digital twin is not a novelty but a necessary tool to close implementation gaps in infrastructure resilience, and VR helps to solve chronic attrition or safety risks. The ultimate measure of success is not user engagement or technological sophistication, but the tool’s demonstrable contribution to public value and agency outcomes.
- 2. Human-First Governance:** One of the most profound risks of 2040 may involve the erosion of human discretion. Envision a future where technology augments professional judgment rather than automating it away. This means designing workflows where AI provides insight, but humans retain the moral and legal accountability for decisions. Whether it is an AR-equipped field worker or a policy analyst, the public servant must remain the “human-in-the-loop,” ensuring that empathy and context remain the final arbiters of government action. This vision imagines an AR-equipped field worker becoming a digitally empowered knowledge broker, and an aged care worker using VR to build a more intuitive understanding of their clients. This approach addresses the AI literacy deficit by integrating intelligence directly into workflows, making employees more effective, engaged, and connected to their purpose.

3. **The Borderless Government:** Citizens interact with one government, and public sector systems must reflect that reality. This calls for moving beyond the traditional silos of agency-specific data toward a shared operational truth. When multiple agencies operate from a common, immersive Digital twin, a decision in one domain (such as urban planning) is instantly understood in the context of others (such as emergency response or social welfare). This “borderless” collaboration allows for systems-level governance that is proactive rather than reactive, dissolving bureaucratic friction to better serve the citizen.
4. **Securing a Sovereign and Trustworthy Digital Foundation:** Underpinning this entire transformation is the strategic imperative of digital sovereignty. The most resilient and trustworthy governments will be those that build their digital future on a foundation they control. This requires a deliberate pivot toward open-source solutions, sovereign data infrastructure, and the cultivation of local talent. This is a clear-eyed strategy to ensure the core systems of governance are shaped by public values and national interests, making sovereignty a bedrock of public trust in a digital age.

Mastering these principles requires tackling human, institutional, and ethical challenges that are far more complex than technology. The following strategic playbook offers a practical path toward making this vision a reality.

## A Guide for Public-Sector Leaders

The journey to the future depicted in the narratives requires more than vision; it demands a practical, evidence-based strategy. The following playbook draws on the insights and challenges identified above, offering a detailed guide for public sector leaders based on a literature review, expert interviews, and the narratives. The playbook is organized into four pillars for action, containing critical recommendations for successful and responsible implementation of immersive workspaces.

### Pillar I: Laying the Foundation

Before deploying any advanced technology, leaders must first establish a foundation of clear strategy, measurable goals, and the requisite infrastructure. This strategic alignment begins by identifying existing organizational problems, such as high training costs or slow planning cycles, rather than chasing technology for its own sake. This focus ensures any investment connects to demonstrable public value, not just internal efficiency metrics.

#### Actions:

- Define specific success metrics upfront. Instead of vague goals, aim for concrete targets like “reduce safety incident training time by 25 percent” or “double public participation in planning consultations.”
- Begin with focused pilot projects to validate the use case and gather the data required to justify scaling.

This foundation brings in physical and financial elements. Immersive experiences require significant bandwidth and computing power, and a lack of dedicated infrastructure is a primary barrier to adoption—as seen in Indonesia’s Ministry of Finance. Budgets must account for the entire system: high-speed internet access, cloud or on-premises back-ends for complex simulations, and device procurement with a maintenance plan. To ensure fair representation, this planning must also extend to public access points, such as community VR stations, so that all citizens and employees can participate.

**Actions:**

- Make high-speed Internet (5G, broadband) available in government facilities and regions targeted for immersive services.
- Budget for cloud or high-performance computing back-ends for complex simulations.
- Plan for device procurement and maintenance, and consider shared facilities (e.g., agency-accessible VR labs) to optimize resources.
- Plan for public access points, such as community VR stations, so that all citizens and employees can participate.

Finally, foundational work means engineering governance, transparency, and portability into the system from day one. This addresses the critical sovereignty and vendor lock-in risks identified in Singapore. Leaders should require agent-visibility controls in any new system, so that when an AI feature mediates work, it is explicit who or what is acting, why a suggestion appears, and how a human can pause or escalate. Concurrently, procurement contracts must mandate portability by requiring Software Bills of Materials (SBOMs), data-interchange adapters, and “exit drill” clauses, which should then be periodically tested to ensure services can be moved without critical disruption.

**Pillar II: Building Capability and Buy-In**

With a solid foundation in place, successful adoption then depends on the human element: fostering a supportive culture, building skilled talent, and ensuring genuine stakeholder partnerships.

The most advanced platform will fail without the human capacity to use, manage, and critically assess it. This requires a talent strategy that prepares staff for a future where stable roles may increasingly be complemented by dynamic, skill-based task orchestration.

**Actions:**

- Develop a three-part talent strategy. First, upskill current staff by offering hands-on workshops, similar to Singapore’s approach of training thousands of public officers in emerging technologies.<sup>17</sup>
- Second, create multidisciplinary teams that pair internal policy experts with hired or contracted technical specialists.
- Third, deepen capabilities in orchestration, ethical judgment, critical evaluation, and cross-disciplinary collaboration.

This strategy is supported by a culture of experimentation and support for the workforce. Indonesia’s Ministry of Finance provides a strong model with its leadership-endorsed ‘Community of Practice’ program, which actively builds AI literacy and shares knowledge across the organization. As one Australian expert noted, driving adoption is often about “making people feel something,” fostering an emotional and intuitive understanding of the technology’s value that manuals alone cannot provide.

---

17. Digital Government Blueprint, 2020

This engagement also serves to turn stakeholders from passive recipients into active partners. Internally, this means involving staff and their unions early in the design process to address legitimate concerns about surveillance or the erosion of professional autonomy. Externally, for citizen-facing applications, engagement requires deep user research with diverse community members. Leaders can ensure this participation is authentic and not subtly manipulated by persuasive AI, a key risk highlighted in the virtual town hall narrative.

#### **Actions:**

- Co-design sessions, which can be actively monitored for distorting AI effects such as groupthink or the crowding out of quieter voices. The Australian model of funding initiatives through trusted industry bodies offers a valuable alternative, creating a partnership layer that can effectively engage suppliers to build products that meet a genuine, validated need.

### **Pillar III: Scaling with Integrity**

Moving from a successful pilot to an enterprise-wide solution requires a deliberate strategy for scaling, sustainability, and ethical governance. Policy alone is insufficient; integrity must be built into the system's design.

To maintain public trust, scaling initiatives should prioritize privacy and inclusion by confronting the “Human–AI Governance Paradox”—the fact that technology designed to augment people can also erode agency. This is addressed through “governance-by-design,” which builds in explainability, contestability, and clear human oversight from the start. This includes systems like the ‘Ethos’ AI-auditor to check other AIs for bias, as well as clear override protocols that cede final control to a human, such as the calendar management example in the narrative. To protect democratic integrity, leaders should carefully monitor any use of persuasive or manipulative AI in public deliberation and create transparency registers for high-risk systems. These systems should be subject to independent audits that explicitly test for emotional manipulation and bias.

#### **Actions:**

- Develop and implement clear policies for data privacy, including opt-in consent for data collection and anonymization of analytics.
- Implement “safety by design” principles, as outlined by Australia’s eSafety Commissioner in 2022, to ensure content is appropriate and that public platforms are moderated.
- Ensure inclusive design by providing alternatives for those who cannot use VR (e.g., 2D modes) and adhering to accessibility standards.

Many successful pilots stall due to a lack of a clear path to scale. A three-layer sequencing model (Strategy → Governance → Capability) offers a practical path from pilot to production by locking in purpose and KPIs first, then codifying controls and reversibility, and only then building delivery capacity. This avoids “pilot theater” and bakes in portability tests and exit drills as scale grows. True scaling moves a tool from “novelty” to a core function by integrating it into standard service workflows and official training curricula. This planning also includes budgeting for long-term IT support, and ensuring that procurement models are flexible enough to accommodate growth. Scaling plans should also engineer for reversibility, with predefined exit routes and periodic portability tests to prevent vendor lock-in.

**Actions:**

- From the project's inception, plan how a successful pilot will be funded and governed at a broader scale. For example, after proving VR's value, embed VR modules as required components of official certification processes.
- Budget for long-term IT support and maintenance.
- Consider procurement models that allow for scaling, such as enterprise-wide software licenses.

Finally, implementation is a dynamic learning process, not a one-time event. Leaders should institute robust monitoring and evaluation (M&E) for every initiative, tracking key performance indicators (KPIs) against the original goals. Crucially, M&E should be expanded to actively track unintended consequences, such as the erosion of employee autonomy, the introduction of new socioeconomic biases, or the cognitive strain of relentless reskilling.

**Actions:**

- Institute robust Monitoring and Evaluation for every initiative.
- Define Key Performance Indicators aligned with the original goals and track them over time.
- Use a mix of quantitative data (e.g., pre- and post-VR test scores, service delivery times) and qualitative feedback (e.g., user satisfaction surveys) to measure impact.
- Iterate or even pivot if evaluation shows the initiative is not delivering value.

**Pillar IV: Fostering a Regional Ecosystem**

No agency or government can manage this transformation alone. Long-term success will be accelerated by fostering collaboration across departments and national borders, which avoids redundant effort and accelerates progress.

**Actions:**

- Establish internal communities of practice for immersive tech in government, where project teams can share lessons learned and co-develop solutions.
- On a larger scale, pursue a more visionary model of collaboration. As one Australian advisor proposed, government can play a role in advocating for a “Metaverse Commons”—a shared, open-source digital space governed by agreed-upon ethical rules and values. Such a “common ground” would provide a foundation for public and private innovation without dependency on any single tech giant.<sup>18</sup> This approach ensures the resulting ecosystem is built on shared standards for ethical AI, interoperability, and data governance, creating a more equitable and sovereign digital future.

18. Bloomberg Cities, “Virtual Realities: How Cities Are Moving into the Metaverse and Beyond,” *The Smart City Journal*, 2022, <https://www.thesmartcityjournal.com/en/technology/virtual-realities-how-cities-are-moving-into-the-metaverse-and-beyond>.

### Areas for Future Research

Future studies of digital enablement for immersive workspaces may include:

- **Barriers to innovation**—New research could focus on effective strategies for procurement, investment cycles, and decision-making processes.
- **Governance**—Because ownership and accountability are barriers to innovation, new research could address this as well as qualifications for owning the governance function—including how best to address central vs. departmental/agency accountability.
- **Cybersecurity**—Studies could address this critical enabler (or blocker if not done effectively) that underpins sovereignty, architecture decisions, skills, partnerships, and international cooperation/dynamics.
- **AI as a catalyst for change**—Shifting government AI skillsets from being ‘doers’ to ‘orchestrators’ will present a major challenge, and governments will benefit from new thinking on how to accomplish this.

## Conclusion: A Call to Action

The adoption of immersive workspaces in the public sector does not reflect a single, uniform trend, but rather a complex set of interwoven themes. Technology often involves the simplest part of the transformation. The greater challenges are human and strategic: building foundational readiness, cultivating critical AI literacy to question automated outputs, and navigating a new geopolitical landscape with policies that realistically assess the risks of vendor lock-in and data sovereignty.

Such strategies can help governments to pivot from the technically possible to what human–AI teamwork can look like in practice, and encode those desired roles and boundaries into everyday collaboration. This means making immersive work sovereign by architecture and practice, not policy statements alone.

As public-sector leaders look to the future, the findings from this report demonstrate that successful transformation must be human-centered; this hinges on a strategic vision that prioritizes utility over novelty and human capability over simple procurement. By focusing on core principles, governments can ensure that immersive workspaces become more than just a technological experiment; they become a true force for building the resilient, efficient, and innovative public sector of tomorrow.



# Glossary

## Agentic AI

Artificial intelligence systems capable of autonomous decision-making and action-taking to achieve specified goals. Unlike traditional AI that simply responds to queries, agentic AI can initiate tasks, coordinate resources, and adapt strategies without constant human oversight. In governmental contexts, this includes AI systems that can automatically route emergency services, manage workflows, or optimize resource allocation.

## APAC (Asia-Pacific)

Geographic region encompassing countries in East Asia, Southeast Asia, South Asia, and Oceania. This report focuses specifically on Australia, Indonesia, Malaysia, New Zealand, and Singapore as representative examples within the broader APAC governmental landscape.

## Augmented Reality (AR)

Technology that overlays digital information, graphics, or guidance onto the real-world environment, typically viewed through smartphones, tablets, or specialized glasses. AR enhances rather than replaces the physical world. Examples in this report include care workers seeing patient information overlaid on their field of view, or maintenance crews viewing underground utility maps while onsite.

## Bioadaptive Systems

Technologies that monitor and respond to human physiological or cognitive states in real-time, such as heart rate, stress levels, or cognitive load. In workplace contexts, these systems can adjust work schedules, recommend breaks, or modify task assignments to optimize employee wellbeing and productivity.

## Cognitive Load (Bioadaptive Context)

The mental effort required at a point in time; used by bioadaptive systems to recommend pacing or breaks, with clear consent and override logging.

## Community of Practice

A group of professionals who share a common interest or domain of expertise and engage in collective learning through regular interaction and knowledge sharing. In governmental AI adoption, these communities help staff share lessons learned, troubleshoot challenges, and develop shared standards across departments or agencies.

## Digital Sovereignty (Pattern)

A nation's ability to exercise independent control over its digital infrastructure, data, and technology systems without excessive reliance on foreign entities or corporations. This includes control over where data is stored, who can access it, how algorithms are governed, and the capacity to develop or modify critical technologies domestically. Digital sovereignty is increasingly viewed as essential to national security and policy autonomy.

## Digital Twin

A dynamic, virtual replica of a physical asset, system, or environment that is continuously updated with real-time data from sensors and other sources. Digital twins enable simulation, prediction, and optimization of real-world operations. Government applications include city infrastructure management, emergency response planning, and policy impact modeling.

## **Governance-by-Design**

Building explainability, contestability, audit trails, and role boundaries into systems from the start, rather than relying on policy after deployment.

## **Governmental Metamodel**

A comprehensive simulation system that models the complex interactions between government policies, economic factors, social behaviors, and environmental conditions. These models allow policymakers to test interventions and predict cascading effects across multiple domains before implementation in the real world.

## **Human–AI Collaboration (Human–AI Teaming)**

The practice of humans and AI systems working together in complementary ways, with AI augmenting human capabilities rather than replacing them. Effective collaboration leverages AI's computational power and pattern recognition while preserving human judgment, creativity, ethical reasoning, and contextual understanding.

## **Human–AI Governance Paradox**

The tension whereby AI systems designed to augment human intelligence and decision-making can simultaneously erode human agency, autonomy, and accountability. As AI becomes more sophisticated and integrated into workflows, there is a risk that humans become overly reliant on algorithmic recommendations, lose critical skills, or find their choices subtly constrained by system design.

## **Immersive Technologies**

Digital technologies that create or enhance experiences that feel engaging and absorbing, often by simulating physical presence in real or imagined environments. This umbrella term encompasses Virtual Reality (VR), Augmented Reality (AR), Mixed Reality (MR), and related systems such as Digital twins and AI-powered simulations.

## **Immersive Workspaces**

Work environments that leverage immersive technologies—including AI, VR/AR, and Digital twins—to enable more dynamic, data-rich, and collaborative ways of performing governmental work. These workspaces break down traditional physical and departmental barriers, creating integrated, responsive ecosystems for public service delivery; these workspaces also include other dynamic environments that blend physical and digital worlds. Such spaces allow employees to learn and perform tasks in highly engaging, simulated settings.

## **Interoperability and Portability**

Ability to exchange and move data, models, and services across vendors and environments through open standards, SBOMs, and data-exchange adapters.

## **Internet of Things (IoT)**

A network of physical devices embedded with sensors, software, and connectivity that enables them to collect and exchange data. In governmental applications, IoT sensors monitor everything from traffic flow and air quality to infrastructure integrity and flood risk, feeding real-time data to systems like Digital twins.

## **Key Performance Indicator (KPI)**

A measurable value that demonstrates how effectively an organization is achieving its stated objectives. In the context of immersive workspace adoption, KPIs might include training completion rates, reduction in safety incidents, employee satisfaction scores, or citizen service delivery times.

## **Lidar (Light Detection and Ranging)**

A remote sensing technology that uses laser light to measure distances and create precise, three-dimensional maps of physical environments. Government applications include infrastructure inspection, terrain mapping for disaster planning, and autonomous vehicle navigation.

## **Metaverse**

Persistent, shared virtual environments where users can interact with each other and digital objects in real-time through avatars. In governmental contexts, metaverse platforms enable virtual town halls, collaborative planning sessions, and citizen engagement activities that transcend physical location limitations.

## **Metaverse Commons (Concept)**

A proposed shared, open, ethically governed virtual space for public-interest use, prioritizing interoperability, safety, and sovereignty.

## **Monitoring and Evaluation (M&E)**

A systematic process of collecting, analyzing, and using data to track the progress and assess the outcomes of programs or initiatives. Effective M&E in immersive technology adoption includes both quantitative metrics (e.g., cost savings, time efficiency) and qualitative measures (e.g., user satisfaction, unintended consequences).

## **Open-Source Software**

Software whose source code is publicly accessible, allowing anyone to inspect, modify, and distribute it. In the context of digital sovereignty, open-source solutions reduce dependency on proprietary platforms controlled by foreign corporations, though they require investment in local technical capacity and careful supply chain management.

## **Open-Washing**

The practice of marketing AI models or software as “open” while actually imposing significant restrictions on access, modification, or redistribution. For example, releasing only model weights (“open-weights”) without training data, code, or documentation provides limited transparency and does not enable true independent scrutiny or sovereignty.

## **Open Weights (vs. Open-Source)**

Release of model parameters without full training code, data, or licensing freedoms; offers limited transparency and portability compared to genuine open-source releases.

## **Persistent Reality**

A continuously available, personalized digital environment that integrates real-time data, AI insights, and collaborative tools into a seamless workspace. Unlike discrete applications or sessions, persistent reality maintains context across time, allowing work to flow naturally between physical and digital modes.

## **Persuasive AI**

AI systems designed to influence human behavior, beliefs, or decisions through strategic framing, emotional appeals, or selective information presentation. In democratic governance, the use of persuasive AI raises ethical concerns about manipulation, informed consent, and the integrity of public deliberation.

### **Pragmatic Partnership Model**

A governmental approach to technology adoption where the public sector acts as a catalyst through funding and strategic direction, while empowering private sector specialists and industry bodies to develop and implement solutions. This model leverages external expertise while ensuring public accountability and mission alignment.

### **Software Bill of Materials (SBOM)**

A comprehensive inventory of all software components, libraries, and dependencies used in a system. SBOMs enable organizations to audit supply chains, identify vulnerabilities, assess sovereignty risks, and ensure portability by understanding exactly what technologies they depend on.

### **Sovereign Cloud and Data Residency**

Infrastructure and contractual controls ensuring data is stored and processed under a jurisdiction's laws and within specified territories, with assured portability and audit.

### **Sovereignty Imperative**

The strategic priority of reducing dependency on foreign-controlled technology infrastructure and platforms, driven by geopolitical tensions and the recognition that digital dependency creates national security and policy autonomy risks. This imperative shapes procurement, development, and governance decisions across the technology stack.

### **Technical and Vocational Education and Training (TVET)**

Educational programs focused on practical skills and hands-on training for specific trades, occupations, or technical fields. TVET is increasingly incorporating immersive technologies like VR and AR to provide realistic, safe, and cost-effective training environments.

### **Virtual Reality (VR)**

Technology that creates fully immersive digital environments that replace the user's physical surroundings, typically experienced through head-mounted displays. VR applications in government include empathy training, hazardous scenario simulations, skills development, and remote collaboration in shared virtual spaces.

# Appendix

## Narrative A: “The Proactive Agency: A Day in Resilient Governance” (Malaysia, 2030)

Adira, a city manager for a mid-sized Malaysian municipality, begins her day interacting not with a screen, but with a living 3D map of the city projected holographically over her desk. This is the city’s Digital twin, a dynamic, real-time replica fed by thousands of IoT sensors and serving as the single source of truth for all municipal operations.

Each morning, the operations team runs a five-minute readiness check. They review sensor health, pending software updates, and the list of automated actions the system is allowed to take. If an update is not approved, the system records the reason and reverts to the last stable version. The day’s configuration is saved to the incident log.

Suddenly, a crimson pulse emanates from the low-lying Taman Desa district on the map. The city’s predictive AI, analyzing meteorological data, river levels, and soil saturation, has flagged a 75 percent probability of a flash flood by late afternoon. This is not a guess; it’s a data-driven forecast based on patterns from hundreds of previous weather events. The AI recommends initiating a “Level 2 Flood Response.” Without hesitation, Adira issues a voice command: “Initiate Level 2. Authorize drone deployment and citizen notifications.”

The system executes a pre-approved, automated workflow across multiple fronts. In Operations, Lidar-equipped drones are autonomously dispatched to survey critical drainage culverts, streaming real-time video back to the Twin, which highlights two partially blocked channels. A public works crew is automatically routed to the location. Simultaneously, the crew lead arrives at the site and, wearing AR glasses, sees a shimmering blue overlay of subterranean water mains, allowing her team to begin excavation safely. A pulsing red highlight warns them of a high-pressure gas line just meters away, preventing a potential disaster. The AR guide provides step-by-step instructions for clearing the blockage, a task they have practiced dozens of times in a VR training module, ensuring high productivity and low risk.

The lead switches on a community layer that shows evacuation capacity and the share of older residents. One zone has low data confidence. A manual outreach task is created for two liaison officers. Their notes are added to the map so the model can improve next time.

At the same time, residents in the at-risk zone receive a targeted notification. When they open their phone’s camera, an AR overlay shows them the safest and clearest evacuation route from their specific location, dynamically updating to avoid projected traffic congestion.

When the storm hits, the impact is minimal. The cleared drains and pre-positioned flood barriers hold. The city’s operations continue without the chaos of a reactive crisis. Adira reviews the post-incident report, automatically generated by the AI, which details response times and resource costs.

Two upstream gauges stop sending data while rainfall increases. The model proposes to downgrade the alert. The lead pauses the change and requests a field check. A technician confirms the river is still rising. The pause and the field report are logged, and the model is scheduled for review to reduce reliance on single data sources.

But she notices something: The report identifies “a minor inefficiency in the drone deployment pattern”—yet when she cross-references her operations log, she discovers the crew also encountered an unmapped informal settlement near the second drainage site: families living in temporary structures, not captured in the official Twin. The crew lead had to make a manual judgment call, marking the location by hand and coordinating evacuation through direct human communication, not AR overlays.

The AI’s report makes no mention of this. The Digital twin’s “single source of truth” has a blind spot—and the most vulnerable populations remain invisible to the system. Adira opens a manual addendum to the report: “Lesson learned: Digital twin requires integration of informal settlement data. Human judgment remains essential where data infrastructure does not reach.” This entire process was the solution to the “implementation gap,” built not by “parachuting AI down,” but by first systemizing the city’s foundational processes—and then learning where the system’s limits lie.

The team tracks four items: time to first decision, time to remediation, use of automated actions, and outreach to low-confidence areas. Human overrides, model changes, and unresolved issues are captured in the register. Two follow-ups are assigned: add community-mapped data for informal housing and run a portability test to the secondary environment.

## **Narrative B: “The Augmented Analyst: Utility and Insight in Indonesia’s Distributed Government” (Indonesia, 2035)**

Dita, a 28-year-old fiscal analyst, begins her day not in a Jakarta skyscraper, but at a Ministry of Finance (MoF) sponsored co-working hub in Bandung. She is part of the MoF’s distributed talent initiative, a strategic response to attract top analytical minds from across the archipelago. Her primary collaborator isn’t in the next cubicle; it’s Hasta, her AI digital teammate, a sophisticated evolution of the early 21st-century analytics assistants like Pandita AI. Her mission is to help design a new green finance policy. Hasta has already processed terabytes of national financial, tax, and agricultural data, flagging an anomaly: a consistent 15 percent dip in tax revenue from a key agricultural district in Central Java that defies seasonal patterns. The initial AI-generated summary is blunt: “High probability of coordinated tax evasion.” This is the automation part of her workflow—fast, data-driven, and devoid of context.

This is where Dita’s role as an augmented analyst begins. Instead of accepting the conclusion, she brings her domain knowledge to bear. “Hasta,” she says, putting on a lightweight headset, “that district was the pilot site for the ‘Sustainable Harvest’ subsidy program three years ago. Let’s go deeper.” Her view transforms into the Ministry’s immersive data workspace. She stands before a holographic map of the district, with glowing red nodes indicating the underperforming tax zones. This isn’t just a visualization; it’s a dynamic analytical sandbox. “Overlay subsidy disbursement records, satellite imagery of crop types from the last five years, and commodity market price fluctuations,” Dita commands. The AI executes the query instantly. Streams of data coalesce over the map. The correlation is immediate and clear. The tax dip began exactly when farmers, using the subsidies, started replacing their traditional cash crops with a new, higher-value, sustainable variety that has a longer maturation cycle. The AI’s initial analysis, focused only on tax filings, missed the underlying policy success story.

“The system’s prediction was efficient, but wrong,” Dita notes. “It saw a pattern but couldn’t understand the purpose.” She and Hasta then spend the next hour in a collaborative loop. Dita poses hypotheses, and Hasta runs complex predictive models in seconds, simulating the long-term economic and environmental impact. Their joint analysis transforms a potential tax fraud investigation into a powerful case study demonstrating the policy’s success, providing the data needed to advocate for a national rollout. Dita’s human expertise and critical judgment were not replaced; they were amplified.

As she finalizes her report, a notification appears from the Ministry’s career platform: “Your work on the Green Finance model shows a high aptitude for cross-referencing policy impact with fiscal data. We recommend this advanced module on ‘Socio-Economic Forecasting for Climate Transitions’ to progress toward your next career level.” This is the direct result of the government’s successful shift from viewing technology as a mere novelty to an essential utility. It has created a system that doesn’t just automate tasks but actively deepens the critical thinking and domain expertise of its human talent, ensuring that even as AI becomes more powerful, human oversight and insight remain at the core of public service.

## Narrative C: “The Empathetic Workplace: Augmenting Human Connection in Aged Care” (Australia, 2035)

Liam, a 22-year-old care worker, starts his shift at a suburban aged care facility. He is part of a new generation of carers, attracted to the role by a national training program designed to combat the industry’s chronic burnout and high turnover rates. Before his rounds, Liam enters a small training room and puts on a VR headset for a 10-minute module focused on empathy. He experiences a simulation of his shift from the perspective of Mr. Chen, a resident with dementia. The world becomes a confusing jumble of sounds and unfamiliar faces, and a simple question from a carer sounds loud and aggressive. He feels a jolt of anxiety. The simulation isn’t about technical skill; it’s designed to do what the Australian advisor emphasized is critical: “make people feel something from using it.” This brief, powerful experience builds a deep, intuitive understanding that no textbook could provide.

Before entering a resident’s room, Liam enables consent mode. The device processes prompts on the headset where possible and does not store video. More sensitive cues, such as mood inference, only appear when a clinician approves them. Remote mentoring sessions are watermarked and recorded in the care log.

Later, Liam visits Mrs. Davies, a resident who is often withdrawn. Instead of relying on a physical chart, he consults his facility-issued AR tablet. It discreetly overlays a few non-intrusive cues onto his view: “Reminder: Ask about her granddaughter, Sophie. She enjoys classical music, especially Vivaldi.” The technology is not replacing his human connection; it is augmenting it, providing a “drip feed” of personalized information to help him build rapport. He starts a conversation about Sophie, and Mrs. Davies smiles for the first time that day.

When he needs to assist a resident with a complex mobility issue, Liam feels a moment of uncertainty, as it’s his first time doing it unsupervised. He taps his AR tablet and initiates a “Virtual Mentor” call. The face of a senior nurse from a partner hospital appears in the corner of his screen. Using AR, she draws a simple arrow onto his live view of the room, demonstrating the precise positioning for the support harness. The process is completed safely and confidently. If on-screen guidance conflicts with manual handling policy, the app requires a supervisor check. Liam taps to escalate, the nurse confirms the correct approach, and the task unlocks. A short, de-identified training note is created so the same gap can be addressed in the next skills session.

At the end of his shift, Liam receives a notification from the facility's performance dashboard: "Your empathy scores this week: 87/100. Resident satisfaction: 91 percent. Target for continued federal funding eligibility: 90 percent across both metrics. You're on track—keep up the excellent work!" He feels a mixture of pride and unease. The dashboard pairs the interaction score with resident feedback and staff turnover. If the score rises while complaints increase, the system flags a review of prompts and incentives. The manager can freeze metric weightings and return to observational audits until the issue is resolved.

The metrics validate his work, but they also create pressure. That evening, his supervisor reminds the team: "The system is here to help us provide consistent, evidence-based care. When you override, document clearly—because if our metrics slip, we lose the grants that fund this technology in the first place."

This entire ecosystem was the product of pragmatic partnership. It was funded by a federal innovation grant, administered by the Aged Care Industry Council, and developed by a specialist Australian tech firm. The government didn't build the tech; it identified a national problem and catalyzed a solution that succeeded by focusing on augmenting the most vital element of care: human empathy. But Liam is learning that augmentation comes with subtle strings attached—the technology empowers him, but it also measures and quietly shapes his behavior through the logic of funding and metrics.

Continuation of funding relies on three outcomes: safety incidents trending down, resident satisfaction improving, and staff retention stabilizing. To avoid unintended pressure, persuasive content is not allowed in worker-facing flows and an independent usability audit runs each quarter.

## **Narrative D: The Ministry of Wellbeing: Navigating AI, Ethics, and Autonomy in 2040 (New Zealand)**

Emma is a senior analyst in the 2034-established Ministry of Wellbeing. She works with a team of human and AI analysts to design and evaluate concepts that enhance human wellbeing. Her team is no longer confined to physical offices and operates within an ecosystem of platforms, agents, and environments, which reconfigures itself around evolving tasks and wellbeing goals.

Emma's day begins not with a traditional briefing, prepared by colleagues, but 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 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—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 reconfigures 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—evaluating scenarios, interpreting AI recommendations in light of human expertise, and ensuring 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 'gray' 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 percent," the AI reports, "but it requires a 22 percent 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 percent probability that the displacement will reinforce existing socio-economic 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 percent 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 pre-emptively 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 p.m. 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 feeds 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.

# References

- Beckert, Jens. The Firm as an Engine of Imagination: Organizational Prospection and the Making of Economic Futures. *Organization Theory* 2, no. 2 (2021). <https://doi.org/10.1177/26317877211005773>.
- Beehive.govt.nz. "Virtual Reality Tool Game Changer for Training New Zealanders on Benefit." 2019. <https://www.beehive.govt.nz/release/virtual-reality-tool-game-changer-training-new-zealanders-benefit>.
- BINUS University. "BINUS UNIVERSITY Hadirkan Nusameta Borobudur untuk Wisata Digital." April 2024. <https://binus.ac.id/2024/04/binus-university-hadirkan-nusameta-borobudur-untuk-wisata-digital/>.
- Bloomberg Cities. "Virtual Realities: How Cities Are Moving into the Metaverse and Beyond." 2022. <https://www.thesmartcityjournal.com/en/technology/virtual-realities-how-cities-are-moving-into-the-metaverse-and-beyond>.
- CNBC Indonesia. "Kantor Sri Mulyani Mau Pakai AI untuk Kelola Keuangan Negara." December 29, 2024. <https://www.cnbcindonesia.com/news/20241229082358-4-599282/kantor-sri-mulyani-mau-pakai-ai-untuk-kelola-keuangan-negara>.
- Digital Government Blueprint. *A Singapore Government That Is Digital to the Core, and Serves with Heart*. Singapore: Government Technology Agency of Singapore, 2020.
- Docherty, Bonnie. *Losing Humanity: The Case Against Killer Robots*. New York: Human Rights Watch and Harvard Law School International Human Rights Clinic, 2012. <https://www.hrw.org/report/2012/11/19/losing-humanity/case-against-killer-robots>.
- eSafety Commissioner. *Immersive Technologies—Position Statement*. Canberra: Australian Government, 2022. <https://www.esafety.gov.au/about-us/who-we-are/position-statements/immersive-technologies>.
- GovInsider. "How Singapore's Defence Tech Uses AI and Digital twins." 2021. <https://govinsider.asia/digital-gov/sng-how-singapores-defence-tech-uses-ai-and-digital-twins/>.
- Inavate APAC. "Malaysian Government to Prioritise VR and AR Adoption for Training." 2019.
- Lacity, Mary C., and Leslie P. Willcocks. "Becoming Strategic with Intelligent Automation." *MIS Quarterly Executive* 20, no. 2 (2021).
- Licklider, J. C. R. "Man-Computer Symbiosis." *IRE Transactions on Human Factors in Electronics* HFE-1, no. 1 (1960): 4–11. <https://doi.org/10.1109/THFE2.1960.4503259>.
- Ramírez, Rafael, and Angela Wilkinson. *Strategic Reframing: The Oxford Scenario Planning Approach*. Oxford: Oxford University Press, 2016.
- Richter, Shahper, Alexander Richter, and Matthias Trier. "Narratives of the Future—Exploring Human–AI Collaboration." *Communications of the Association for Information Systems* 57, no. 1 (2025): Article 26. <https://aisel.aisnet.org/cais/vol57/iss1/26/>.
- Richter, Alexander, Shahper Richter, and Nastaran Mohammadhossein. "A Narrative Exploration of the Immersive Workspace 2040." *MIS Quarterly Executive* 24, no. 4 (2025). <https://aisel.aisnet.org/misqe/vol24/iss4/3>.
- Schwartz, David R. *Using AI to Improve Child Welfare*. Washington, DC: IBM Center for The Business of Government, 2025. <https://www.businessofgovernment.org/report/using-ai-improve-child-welfare>.
- Taleb, Nassim Nicholas. *The Black Swan: The Impact of the Highly Improbable*. New York: Random House, 2007.

## About the Authors



**Professor Dr. Alexander Richter**

Wellington School of Business and Government  
Victoria University of Wellington  
New Zealand

E: [alex.richter@vuw.ac.nz](mailto:alex.richter@vuw.ac.nz)

**Dr. Alexander Richter** is a professor at the Wellington School of Business and Government, Victoria University of Wellington, New Zealand and a visiting professor at LUT University, Finland. 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.

A proponent of engaged scholarship, Professor Richter has successfully acquired and led projects funded by the European Union as well as national governments and organizations in Germany, Switzerland, Denmark, Australia and the United States, focusing on the socio-technical design of digital workplaces that foster innovation and efficiency.




**Dr. Shahper Richter**

University of Auckland Business School  
New Zealand

E: [shahper.richter@auckland.ac.nz](mailto:shahper.richter@auckland.ac.nz)

**Dr. Shahper Richter** 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.

# Recent Reports from the IBM Center for The Business of Government



## Using AI to Improve Child Welfare

by David R. Schwartz




## Outcome-Based Contracting in U.S. Government

by Daniel Finkenstadt and Tim Cummins




## Responsible AI for Public Evaluation

by Daniel Foner



## Practical Cyber Solutions for Managing Government Supply Chains

by Robert Handfield




## Five Pillars of Effective Government

by Dan Chenok




## Government's Digital DNA: Identity and Access Management for Public Sector Security

by Andrew Whitford



## Enhanced Government Financial Oversight

by Irakli Petriashvili



## Agile and Adaptive Leadership in the Public Sector

by Ines Mergel



For a full listing of our reports, visit [businessofgovernment.org/reports](https://businessofgovernment.org/reports)

## About the IBM Center for The Business of Government

Through research stipends and events, the IBM Center for The Business of Government stimulates research and facilitates discussion of new approaches to improving the effectiveness of government at the federal, state, local, and international levels.

## About IBM Consulting

With consultants and professional staff in more than 160 countries globally, IBM Consulting is the world's largest consulting services organization. IBM Consulting provides clients with business process and industry expertise, a deep understanding of technology solutions that address specific industry issues, and the ability to design, build, and run those solutions in a way that delivers bottom-line value. To learn more visit [ibm.com](http://ibm.com).

### For more information:

**Daniel J. Chenok**

Executive Director

IBM Center for The Business of Government

600 14th Street NW  
Second Floor  
Washington, D.C. 20005  
(202) 551-9342

website: [www.businessofgovernment.org](http://www.businessofgovernment.org)  
e-mail: [businessofgovernment@us.ibm.com](mailto:businessofgovernment@us.ibm.com)

