accenture

Rethinking Sovereign AI

Winning the AI era on Singapore's terms

A conversation for Singapore's public sector leaders



Executive Summary

Singapore is leading the race to adopt artificial intelligence (AI). The more difficult challenge now is determining the terms on which the nation governs, derives value and retains control over it. Sovereign AI is often framed as a procurement question centred on choosing the right cloud, model or technology vendor. However, this approach typically results in a collection of locally trusted components that fall short of genuine sovereignty, despite significant investment.

This paper argues for a different definition – **Sovereign AI is system-level control over how AI creates or erodes national value, exercised deliberately across four layers and matched by the organisational capability to operate it.** The central question is not whether every component should be owned, but where control creates value and where external capabilities should be strategically leveraged without compromising national interests. This paper translates that principle into a practical agenda for public sector leaders.

01 The sovereignty paradox

02 Why most sovereign AI programmes stall

03 Four plays to build sovereign AI that lasts

04 From control to confidence



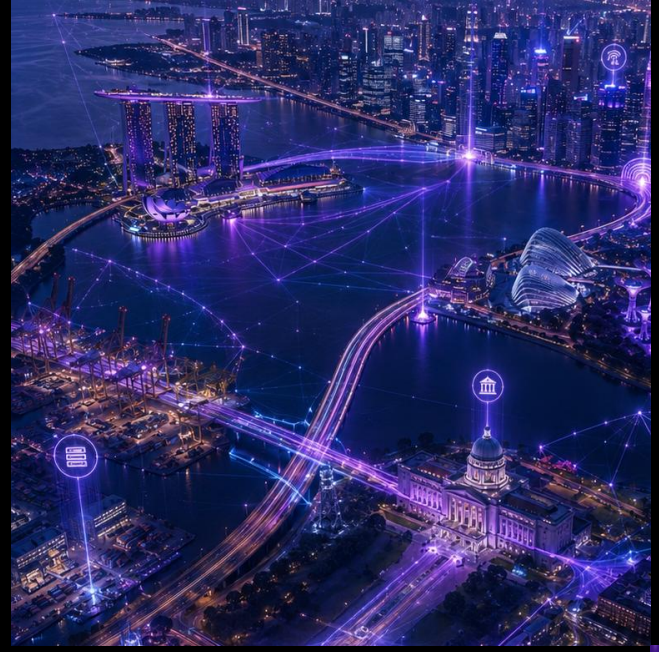
01

The Sovereignty Paradox

Singapore's public sector has responded more quickly than most. Through the National AI Strategy 2.0 – the world's first Model AI Governance Framework for Agentic AI and the National AI Impact Programme, the nation has established itself as a credible leader. The direction is clear. Whether Singapore reaches its destination now depends on decisions public sector leaders make today, beginning with their position on sovereign AI.

Here is where many jurisdictions go wrong. They treat sovereignty as a question of self-sufficiency and reach for the comfort of owning everything. That instinct is understandable. Building and maintaining the full stack internally is expensive, slow and ultimately fragile. Frontier models evolve on deployment cycles measured in weeks rather than years. Any nation that attempts to reproduce the entire ecosystem domestically risks paying today's costs for yesterday's capabilities.

On the opposite end of the spectrum, adopting whichever global AI solutions are fastest or cheapest, creates a different form of vulnerability. It gradually transfers control over the data, models and decision-making systems that increasingly shape public outcomes. Neither extreme serve public sector agencies well. The challenge is to define the boundary between them and manage it deliberately.



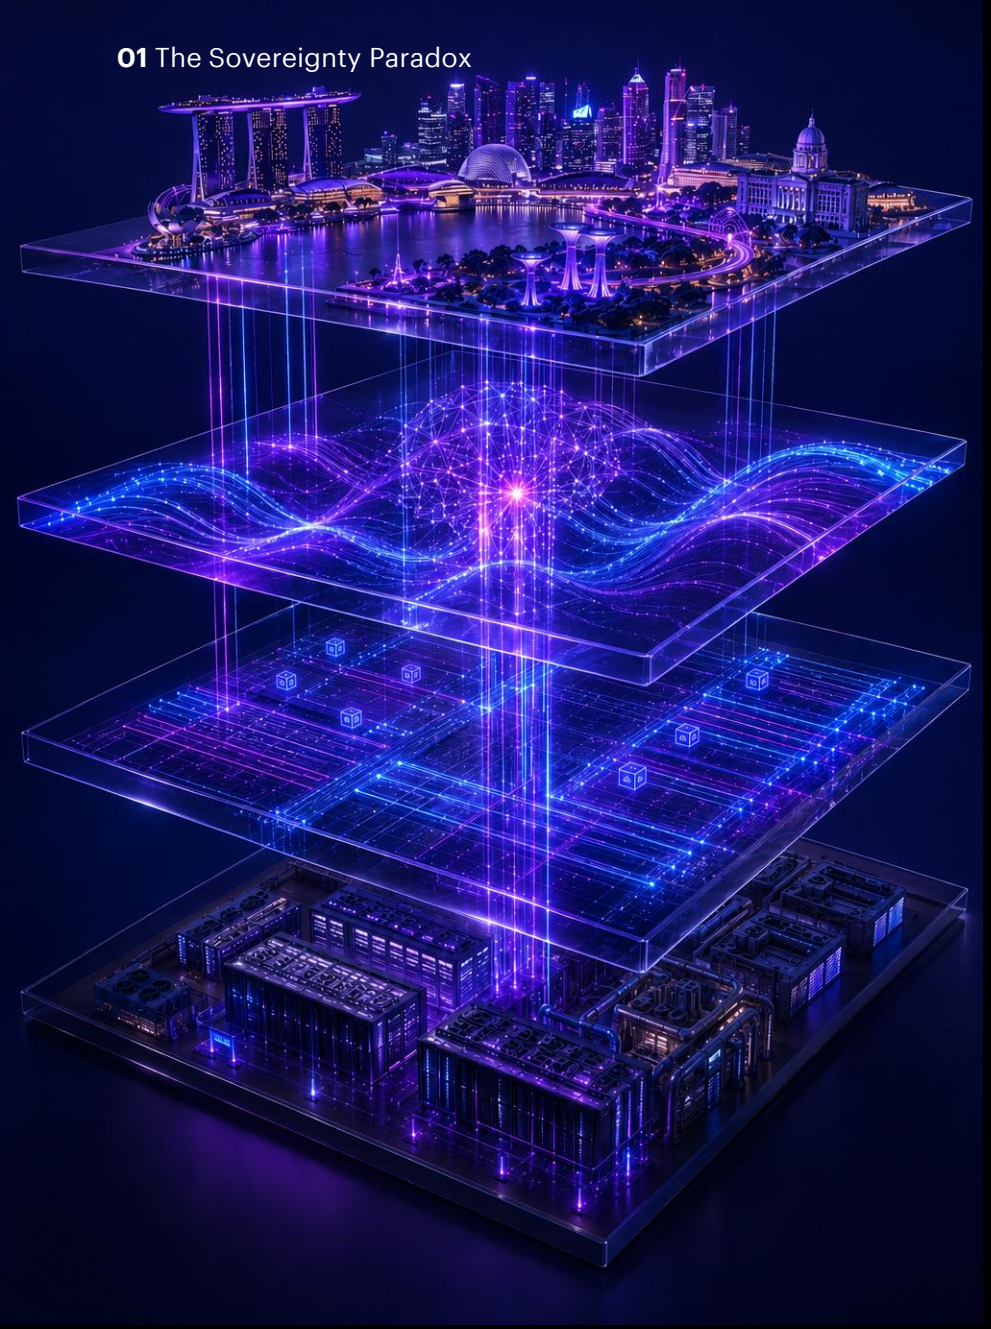
Sovereign AI is best understood not as ownership of any single component, but as system-level control over how AI creates or erodes national value.

61%

of nearly 2,000 business and government leaders surveyed by Accenture in 28 countries say geopolitical tensions make them more likely to seek sovereign technology solutions.

~1/3

On average, only about one-third of AI workloads require sovereign treatment. The remainder can run on global infrastructure with appropriate safeguards. This principle is termed selective sovereignty.



Technological stack:

01

Infrastructure

Where compute and data physically reside.

02

Data

What is collected, curated and made available for training and inference.

03

Models

Who owns, controls and can modify the intelligence itself.

04

Applications

How AI capability is built into the services, decisions and workflows that deliver public value, including the agentic platforms through which autonomous agents now plan, act and coordinate across those workflows.

That control depends on a system rather than a stack. At its core sit **four interconnected technological layers: Infrastructure, Data, Models, Applications.**

However, the four technological stack layers are not sufficient to deliver true sovereignty. Beyond technology, sovereignty depends on organisational enablers: a **future-ready AI workforce** capable of exercising informed oversight, **AI-native operating models** that evolve governance and decision-making alongside automation, and **sustainable AI adoption** practices that preserve human judgement, institutional knowledge and long-term capability.

Moreover, the technological layers must be deliberately **orchestrated** within each agency, as multiple vendors and technologies rarely come pre-integrated. At the same time, they must be **aligned across agency clusters,** where sovereignty compounds through shared governance and common operating models.

Together, these elements form a system, pictured in the diagram in the next page. The technological layers sit at the core, wrapped by the organisational enablers that make them function, orchestrated within each agency and aligned across agencies within a public sector domain (e.g., public safety, defence, healthcare).

True sovereignty is a system, not a stack, that compounds

Four layers of technological layers form the core: data, infrastructure, models and applications. They deliver sovereignty only when operated by a capable organisation, orchestrated within the agency and aligned across agencies within a domain.

TECHNOLOGICAL STACK

- Applications
- Models
- Infrastructure
- Data

Modular components only create value when integrated into a coherent system

1 3 4 5

ORGANISATION ENABLERS

- Future-ready AI Workforce
- AI-native Operating Model
- Sustainable AI Adoption

Organisational readiness turns a sovereign technology stack into a lasting capability

2 4 5

CROSS-AGENCY ALIGNMENT

Sovereign AI across the public sector requires agencies to move together

Agency A ⇄ **Agency B** ⇄ **Agency C**

Aligned Governance Approaches • Common Operating Models • Shared Workforce and Capability Baseline

1 3 4

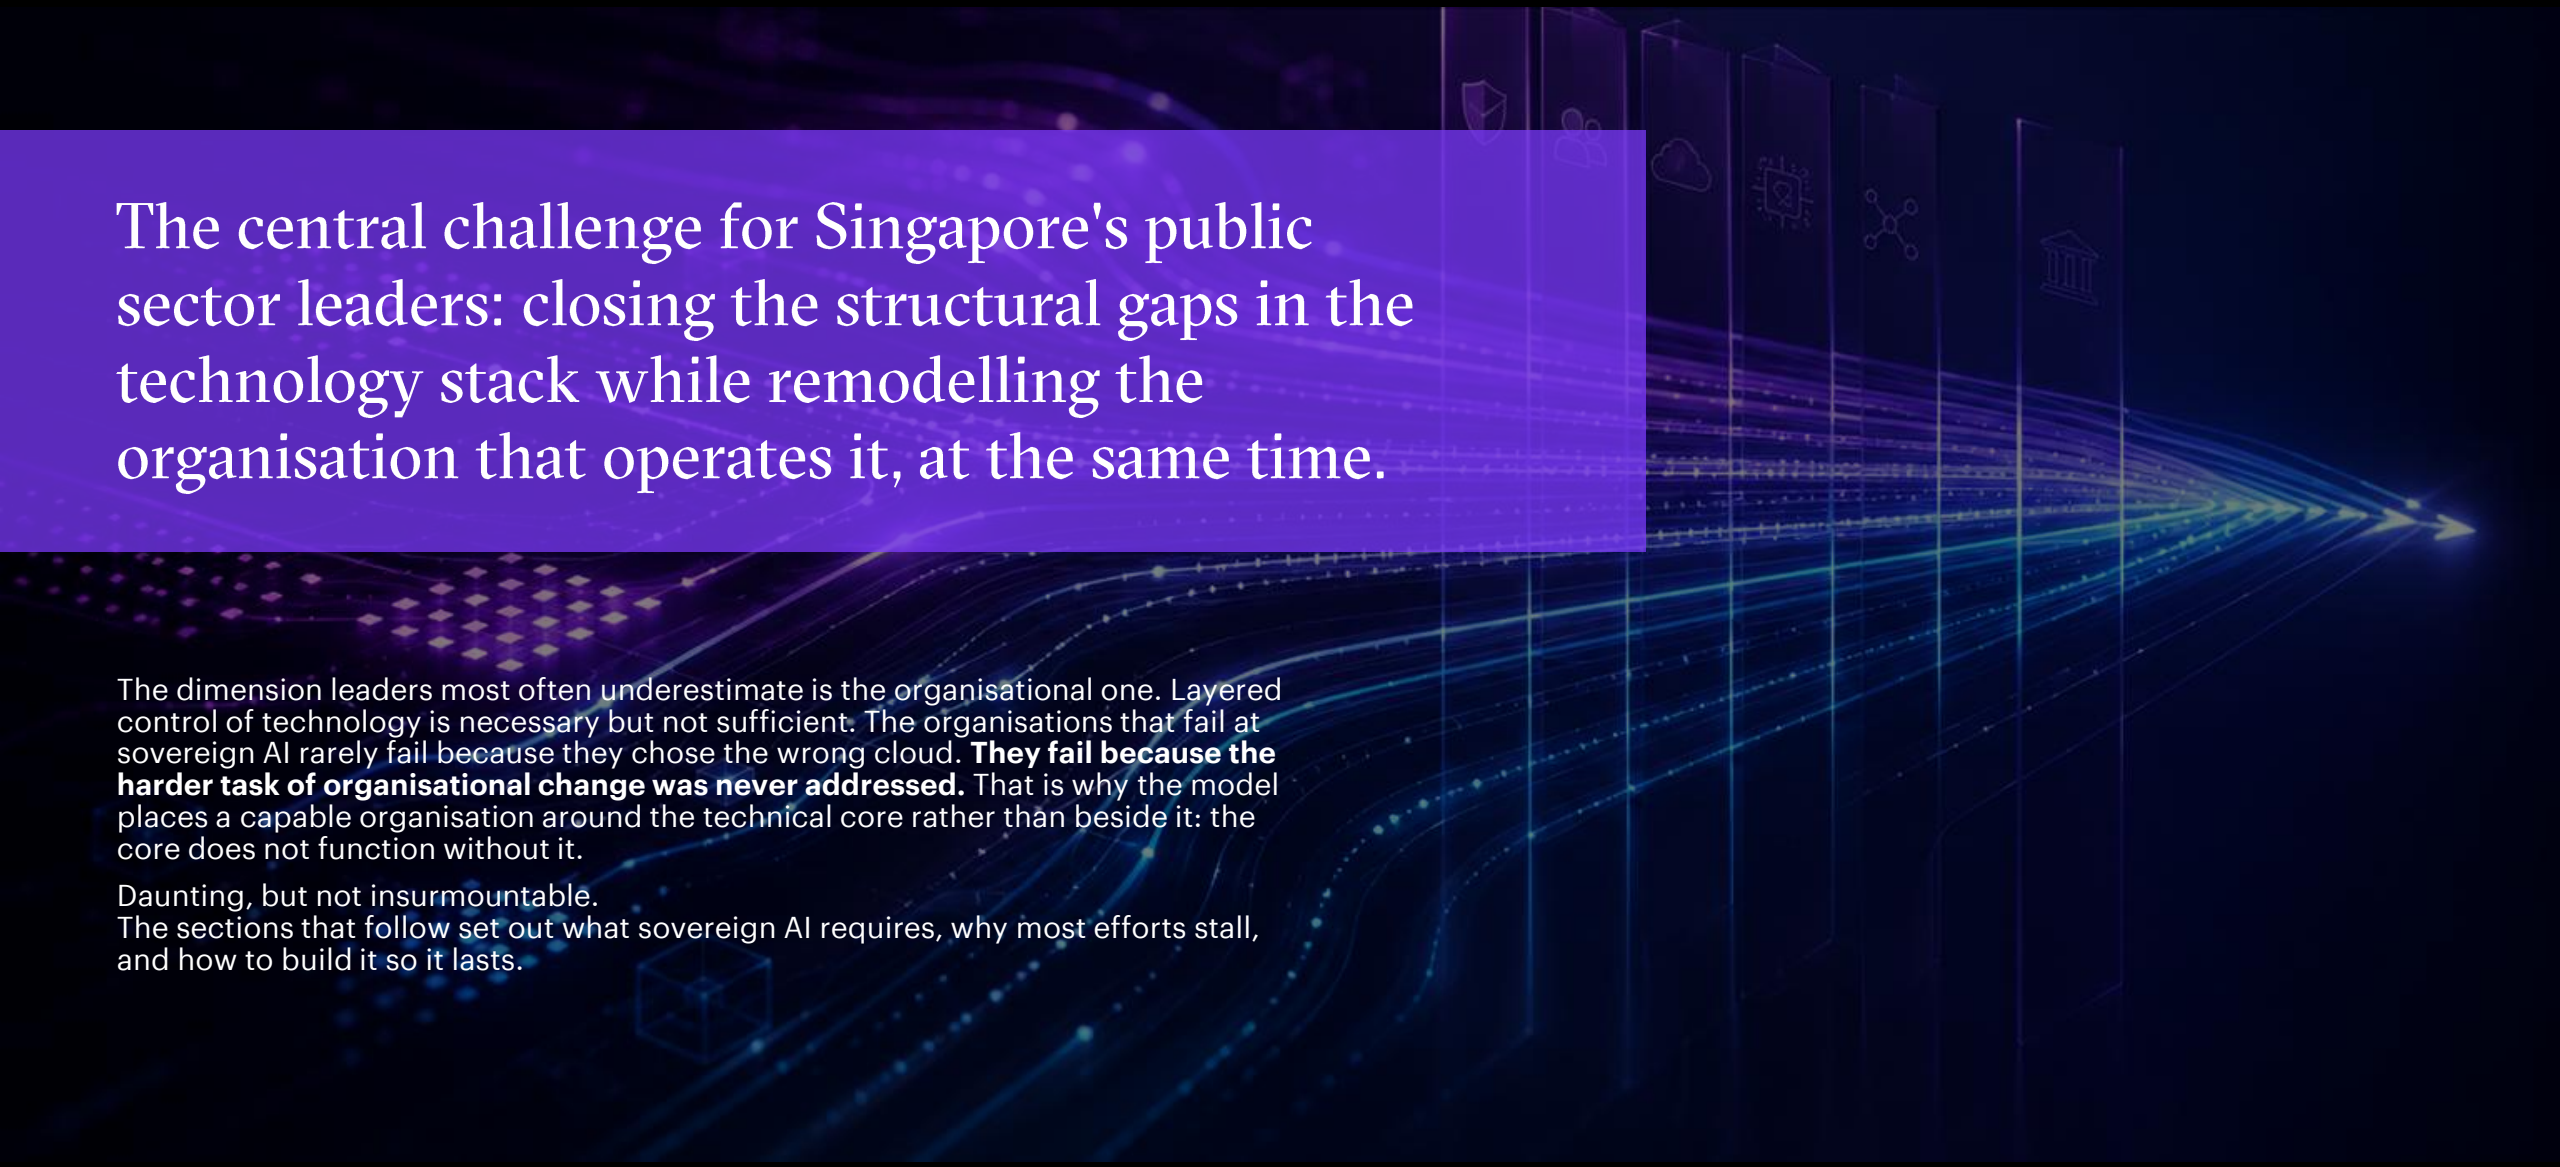
PUBLIC VALUE DELIVERED

1. **Improved Citizen Services**
   Faster and fine-tuned to context
2. **Empowered Workforce**
   Freed for higher-value work
3. **Trusted Decisions**
   Traceable and defensible
4. **National Resilience**
   Continuity under disruption
5. **Economic Capability**
   Skills & capabilities kept onshore

What this system does not require, however, is uniform application across every workload. On average, only about one-third of AI workloads require sovereign treatment. The remainder can run on global infrastructure with appropriate safeguards. This principle is termed **selective sovereignty**: applying control where the workload, sensitivity or strategic value demands it, and freeing budget and attention for everywhere it does not. This is the trade-off the rest of this paper helps leaders optimise.

The international evidence points the same way – nearly 2,000 business and government leaders surveyed by Accenture in 28 countries, 61% say geopolitical tensions make them more likely to seek sovereign technology solutions.[2] The European Union is pooling compute across member states.[3] India is combining national infrastructure with global model partnerships.[4] Saudi Arabia is co-building its national AI capability with frontier partners.[5] Each has converged on the same pragmatic conclusion: selective sovereignty over full self-sufficiency.





The central challenge for Singapore's public sector leaders: closing the structural gaps in the technology stack while remodelling the organisation that operates it, at the same time.

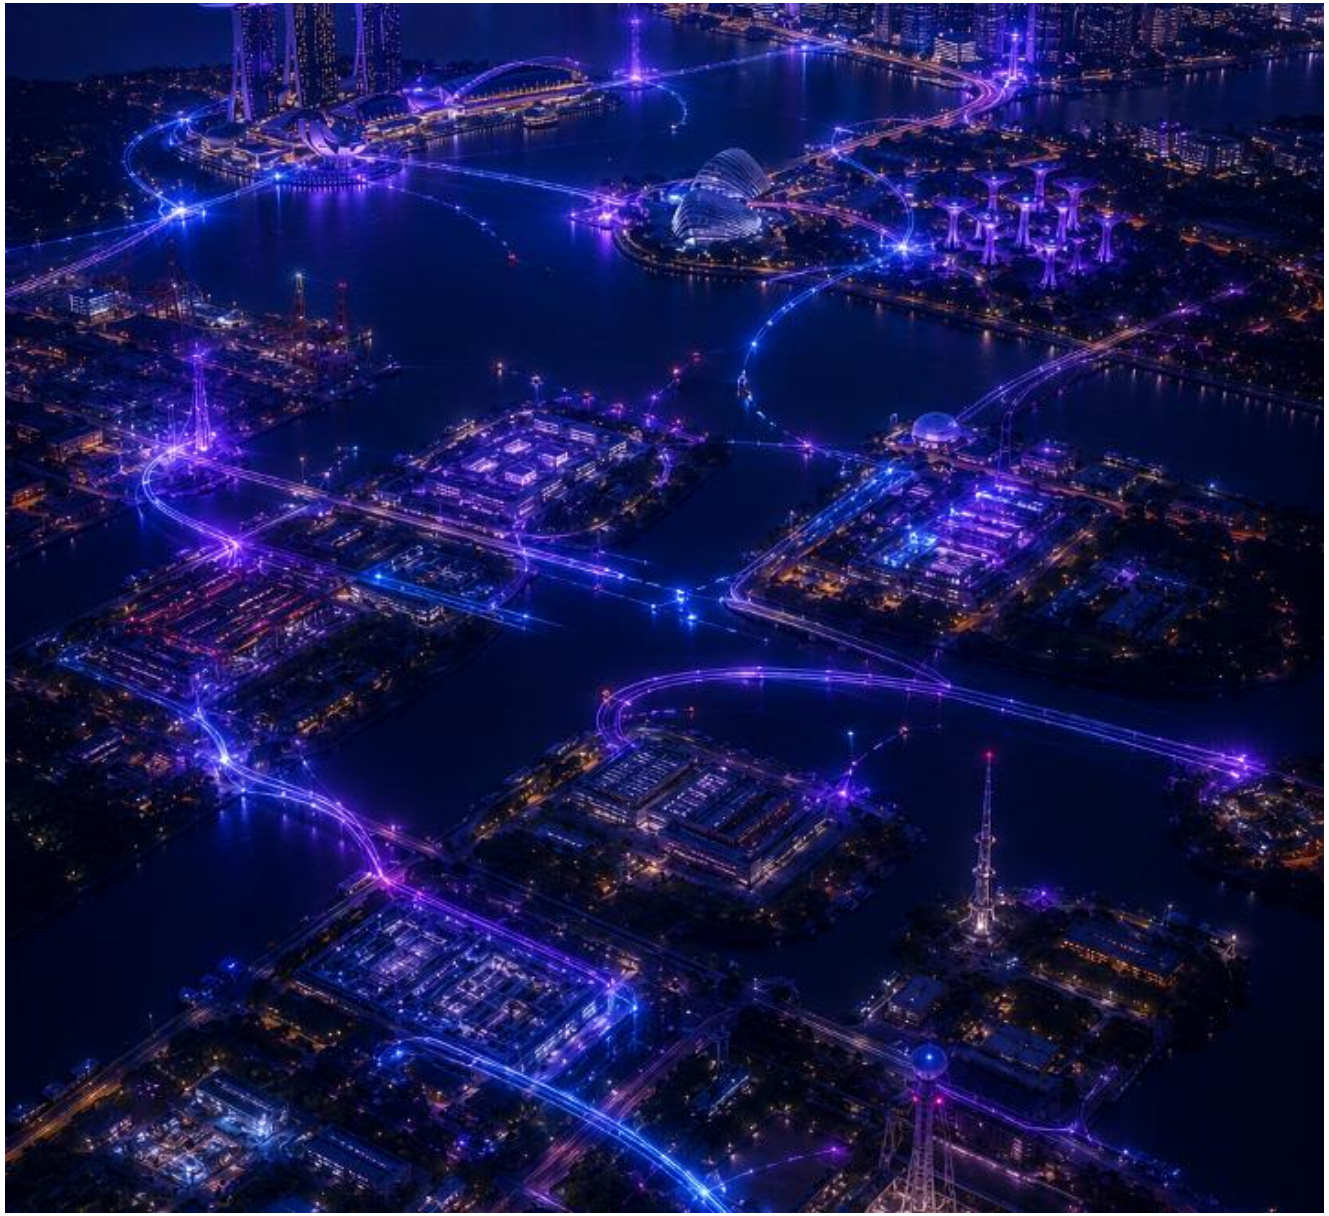
The dimension leaders most often underestimate is the organisational one. Layered control of technology is necessary but not sufficient. The organisations that fail at sovereign AI rarely fail because they chose the wrong cloud. **They fail because the harder task of organisational change was never addressed.** That is why the model places a capable organisation around the technical core rather than beside it: the core does not function without it.

Daunting, but not insurmountable.
The sections that follow set out what sovereign AI requires, why most efforts stall, and how to build it so it lasts.

02

Why Most Sovereign AI Programmes Stall

The failures that emerge in sovereign AI programmes rarely come down to ambition or investment. They fall into two groups: gaps in the architecture, and gaps in the organisation. Most technology partners will address the first, but few are honest about the second. The evidence, and Accenture's experience across comparable public sector programmes, is that the second decides the outcome.





The architecture gaps

Most sovereign AI programmes invest heavily in architecting the stack but still get it wrong. Three problems account for most of the failures:

A lack of systems thinking across the stack. Singapore's sovereign AI conversation today is largely confined to the bottom two technological layers: where data sits and compute runs. The two layers above receive far less attention but ultimately determines whether sovereignty compounds or breaks. Hosting a vendor's proprietary model on local compute is not the same as owning it; the model still answers to its vendor's training data and roadmap. Only 22% of organisations apply sovereignty oversight at the model layer, against 60% at data and 46% at infrastructure.[1] For high-stakes workloads, that gap is the difference between renting intelligence and owning it. The application layer is where sovereign intent is finally realised or lost: a locally hosted model still erodes sovereignty if the services built on it embed proprietary workflows, opaque decision logic or integration dependencies that cannot be audited or modified.

The pilot-to-production gap. AI that works in a controlled pilot breaks down when retrofitted onto decades of legacy infrastructure in operational environments. The very capabilities that demonstrate well in a lab, end-to-end lifecycle control, decision traceability, mission-critical reliability, are precisely the ones that fail under real conditions: brittle integration with ageing core systems, real-time service levels with no tolerance for downtime, and audit trails that must withstand regulatory scrutiny. Crossing this threshold is not a matter of model quality. It is a matter of architectural depth and operational readiness.

Fragmentation across domains. Even where an individual agency architects its own stack well, fragmentation emerges between agencies in the same public sector domain. Health agencies, Home Team agencies and financial regulators each face shared sovereignty challenges that no single agency can solve alone, yet governance standards, data-sharing protocols and model assurance practices are still being defined in silos.

The organisation gaps

Even with the architecture fully addressed, sovereign AI initiatives often fail because the organisation is not ready to absorb them. The constraint is no longer access to capable models but the capacity to operationalise them – a gap reflected in the rise of forward-deployed engineers (FDEs) embedded directly with delivery teams to move specific use cases into production, a model that works well at the workflow level but does not resolve the broader variation in data, models and governance that enterprise-wide AI demands. Three dimensions deserve direct attention from senior leadership.



Talent and workforce readiness. AI literacy is becoming a baseline competency across the public service, as officials increasingly make decisions with AI rather than about it. Capability must be progressively transferred from external partners into the public service, building lasting in-house expertise rather than long-term vendor dependence. Without this, sovereign AI risks remaining sovereign in name but outsourced in practice.

Operating model evolution. As AI becomes embedded in day-to-day decisions, operating models must evolve with it. Research across thousands of enterprise teams makes the point sharply: while AI has accelerated certain workflows, it has not automatically improved overall delivery. The bottleneck shifts downstream. Review queues grow, incidents accumulate, and teams spend more time on validation and oversight, because governance structures and delivery processes have not evolved at the same pace as AI-enabled activity.[6] For public agencies the lesson is blunt: faster AI does not produce better outcomes unless the operating model evolve.

The hidden costs of adoption. The most underestimated risk of AI is not technical failure but quiet organisational erosion. As AI handles a growing share of analysis, drafting and routine judgement, critical thinking can weaken in those who outsource it.[7] Junior staff lose the formative repetition through which professional judgement is built. Senior officers risk an illusion of understanding, accepting AI-generated outputs they would have scrutinised more rigorously had they produced the analysis themselves. Institutional knowledge can be silently displaced by outputs that are fluent but shallow. This is not an argument against adoption. It is an argument for adopting deliberately: preserving spaces for unaided judgement, ensuring outputs are challenged rather than accepted, and treating cognitive resilience as a national capability to be protected. This is the dimension most technology partners are not structured to address, and the one that most often determines whether a sovereign AI investment builds lasting national capability or an expensive dependency.

03

Four Plays to Build Sovereign AI that Lasts

The four plays below translate the architecture gaps and the organisational imperatives into a single, implementable national approach.

PLAY01

Architect for Selective Sovereignty

Selective sovereignty starts with knowing what to own, what to rent and when that decision should change. Four design choices give agencies the tools to do that.

PLAY02

Operationalise AI in Mission-Critical Environments

Many public sector AI in Singapore today remains early-stage. It demonstrates feasibility. It does not yet run the mission-critical workflows where outcomes are consequential, decisions must be explainable and the cost of failure is high.

PLAY03

Orchestrate the Ecosystem

A sovereign AI deployment for any agency draws on multiple providers at once. None of these arrives pre-integrated. Making them work together — technically, commercially and operationally — is where the hardest work of sovereign AI actually lives.

PLAY04

Align for Sovereignty at Scale

Where Play 3 focuses on orchestration within an agency, Play 4 focuses on alignment across agencies operating within the same public sector domain.

PLAY01

Architect for Selective Sovereignty

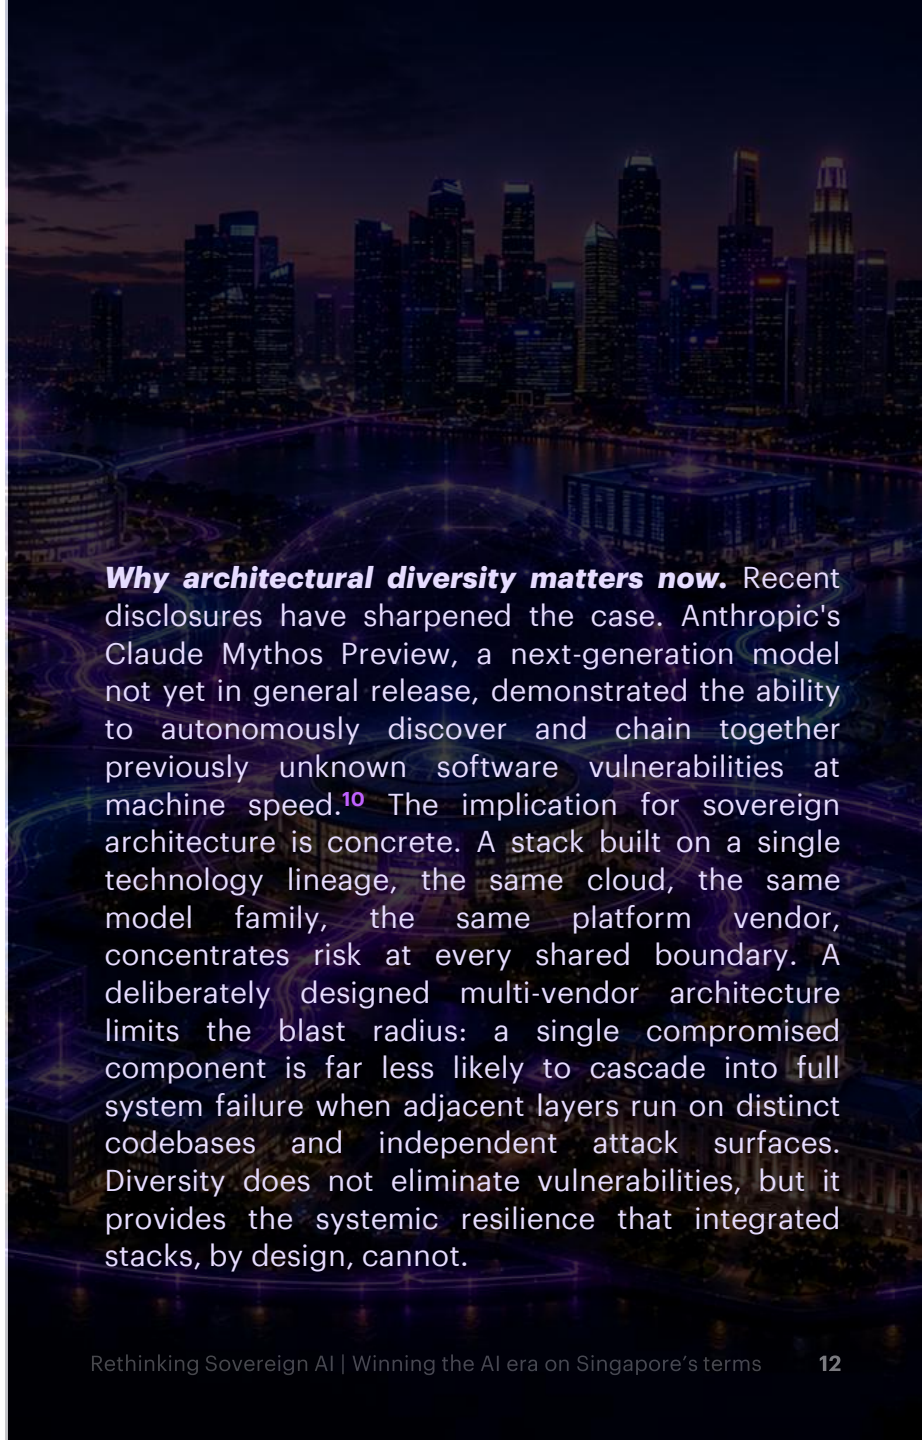
Selective sovereignty starts with knowing what to own, what to rent and when that decision should change. Four design choices give agencies the tools to do that:

Model-level control, including the ability to host and fine-tune open-weight models within national borders. Accenture's strategic partnership with Mistral AI, including a joint research lab based in Singapore, gives public sector clients direct access to production-grade open-weight models that can be hosted locally and fine-tuned on sensitive data.[8] Agencies control how the model runs, what it learns from and how it evolves, rather than depending on a proprietary vendor's training data and update cadence.

Multi-model architectures, combining local models such as SEA-LION (AI Singapore's regional large language model) and MERaLiON (A*STAR's multimodal speech model) for language and cultural context, global frontier models for non-sensitive workloads, and compressed on-premise models for air-gapped environments.

Multi-cloud flexibility, across sovereign options including Microsoft Azure Sovereign Cloud (already deployed for Singapore's Home Team Science and Technology Agency, HTX[9]), Google Cloud's sovereign offering and domestic providers where they add value.

Application-layer sovereignty, ensuring that services built on sovereign models do not quietly erode that sovereignty through proprietary workflows, opaque decision logic or integration dependencies that cannot be audited, modified or replaced.



Why architectural diversity matters now. Recent disclosures have sharpened the case. Anthropic's Claude Mythos Preview, a next-generation model not yet in general release, demonstrated the ability to autonomously discover and chain together previously unknown software vulnerabilities at machine speed.[10] The implication for sovereign architecture is concrete. A stack built on a single technology lineage, the same cloud, the same model family, the same platform vendor, concentrates risk at every shared boundary. A deliberately designed multi-vendor architecture limits the blast radius: a single compromised component is far less likely to cascade into full system failure when adjacent layers run on distinct codebases and independent attack surfaces. Diversity does not eliminate vulnerabilities, but it provides the systemic resilience that integrated stacks, by design, cannot.

PLAY02

Operationalise AI in mission-critical environments

Most public sector AI in Singapore today remains early-stage. It demonstrates feasibility. It does not yet run the mission-critical workflows where outcomes are consequential, decisions must be explainable and the cost of failure is high. Three capabilities define operational readiness in those environments.



End-to-end lifecycle control.

From data ingestion through deployment, monitoring and retraining, the full lifecycle must be instrumented, auditable and under agency control. This becomes more important still as agentic systems begin to act autonomously.

Decision intelligence as a first-class layer.

Deploying AI that merely generates outputs is no longer sufficient. Agencies need systems that produce traceable, explainable and contestable decisions, with full provenance across the data and logic involved. With AI agents, their actions must be independently assessed, and human authority preserved at the point of decision.

Accenture's ACCESS solution illustrates this pattern: an Event Monitoring Agent withholds escalation until independent signals converge, while an Incident Response Agent proposes a plan the commander must review and approve before any assets are dispatched.

Workforce readiness and human-AI operating models.

Mission-critical AI does not run itself. Agencies must design the operating model around it: who decides, when a human must intervene, how escalation works, and how officers are trained to interpret and challenge AI outputs.

As AI shifts from copilots to autonomous agents, CIO and CHRO accountability converge on co-designing a workforce of people and agents, managing realities such as skill atrophy and treating cognitive resilience as an explicit design objective.[11]

A Singapore proof point: Border intelligence at ICA

Singapore's land, sea and air checkpoints processed close to 245 million traveller movements in 2025, a 6.3% year-on-year increase, in an environment where decisions must be made in real time with zero tolerance for downtime. Accenture designed and implemented the operational data insights platform underpinning ICA's Next Generation Clearance transformation: a programme to make secure, automated and passport-less clearance the norm across all checkpoint types.

What was built

The platform performs dynamic, high-throughput entity resolution, consolidating fragmented traveller data into contextualised intelligence that feeds directly into live checkpoint workflows. It is not a parallel analytical layer, but a capability woven into the decisions officers make at the checkpoint itself. Advanced network analytics sit alongside entity resolution to surface insights that static, rules-only approaches cannot generate.

Mapped against the sovereign AI framework

Entity-level intelligence under agency control

The platform is built from ICA's own data, consolidating fragmented records, resolving identities and surfacing hidden networks from sources the agency owns. The intelligence is grounded in sovereign data, not imported inference.

AI embedded in sovereign decision-making

The hybrid decision framework combines static rules, purpose-built ML and generative AI and is built into ICA's existing assessment workflow. AI-assisted decisions are traceable to the agency's own logic, not to a black-box system it cannot see into.

Resilience that preserves decision authority

Isolated pipelines, intelligent fallback layers and end-to-end monitoring ensure operational continuity even under disruption. An architecture that cedes control when it matters most is not a sovereign architecture.

Capability transfer as a sovereignty condition

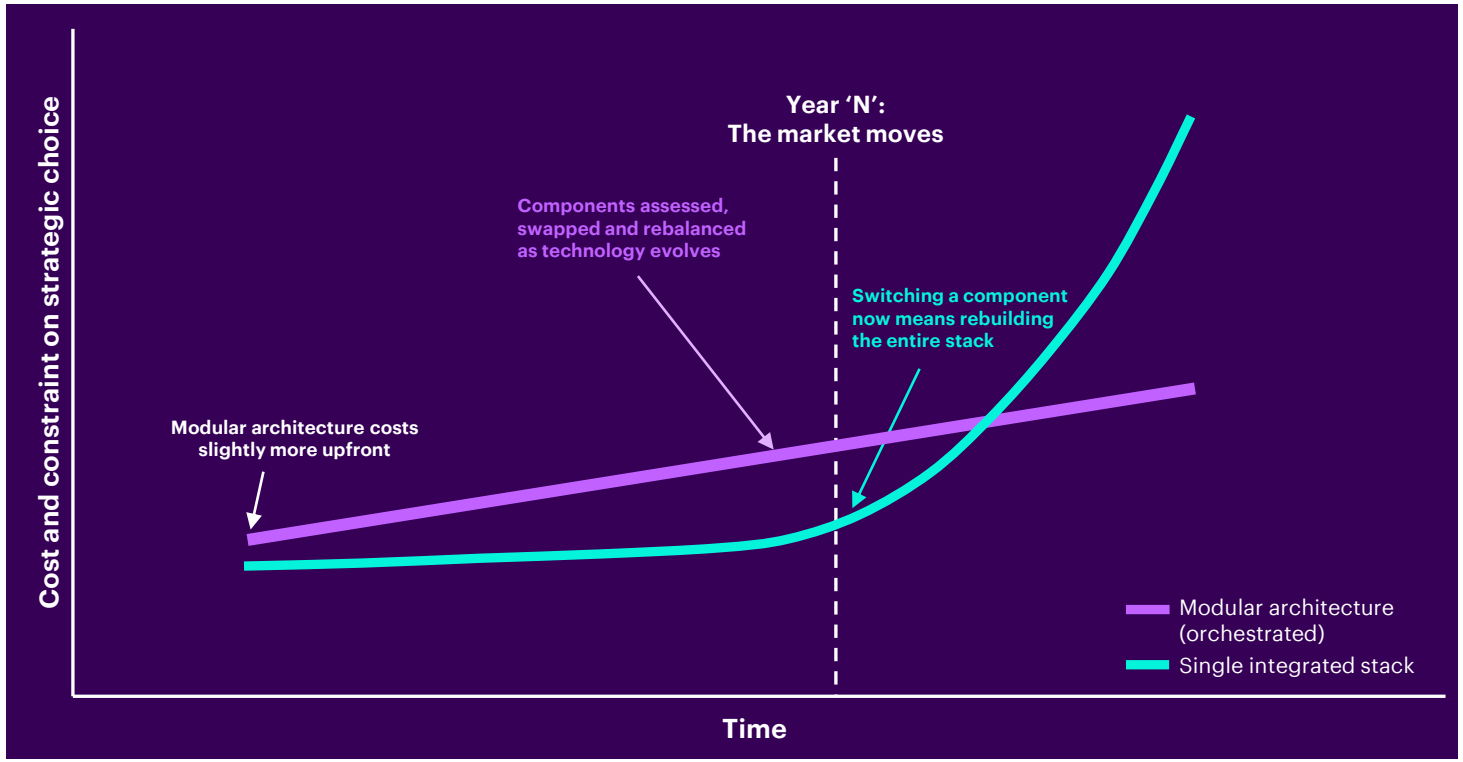
Sovereign AI that cannot be stewarded by the agency's own teams reverts to vendor dependence. The people-first programme, which included data literacy, job shadowing, AI upskilling, was a design requirement, not an afterthought. ICA's teams can now co-develop and own what was built.

PLAY03

Orchestrate the Ecosystem

Orchestrator or integrator: The cost of locking in

A single integrated stack feels efficient for the first few years. However, when the market moves and technology priorities evolve, the cost of switching any component becomes the cost of rebuilding the whole stack. A modular architecture keeps that strategic choice in the agency's hands.



A sovereign AI deployment for any agency draws on multiple providers at once: a sovereign cloud, frontier and open-weight models, a decision intelligence layer, local data platforms, cybersecurity tooling and existing legacy systems. None of these arrives pre-integrated. Making them work together, technically, commercially and operationally, is where the hardest work of sovereign AI actually lives.

Procuring an integrated stack from a single provider solves the orchestration problem by eliminating choice. For the first few years, this feels efficient. When the market moves and a better model has emerged, the vendor's pricing has shifted or the threat environment has changed and the cost of switching one component becomes the cost of rebuilding the entire stack.

An integrated stack simplifies deployment today. A modular architecture preserves flexibility to adapt as risks and priorities change tomorrow.

Accenture brings breadth of partnerships across the stack, sovereign cloud, open-weight models, decision intelligence, security and data, combined with the flexibility to combine, swap and rebalance components as the technology and threat landscapes evolve. Partnerships with Mistral on open-weight models, Microsoft and Google on sovereign cloud, Faculty on decision intelligence[12] and Palantir AIP on ontology-based governed deployment [13] mean agencies are never locked to a single vendor's roadmap. The practical value is concrete: Accenture's proximity to frontier AI development, is embedded directly in Cyber.AI, our AI-driven cybersecurity platform built to detect and remediate the new class of AI-enabled threats.

The measure of a sovereign AI architecture is not whether it works today, but whether it keeps strategic decisions in your hands tomorrow. A modular architecture, where components can be assessed, swapped and rebalanced as priorities evolve, is itself an expression of sovereignty. *This is what separates an orchestrator from an integrator, and it determines whether a sovereign AI stack ages into a national asset or a migration project.*



PLAY04

Align for Sovereignty at Scale

Where Play 3 focuses on orchestration within an agency, Play 4 focuses on alignment across agencies operating within the same public sector domain– for example, health agencies and healthcare clusters, the Home Team and public safety agencies, or defence organisations across the defence ecosystem. Even when each agency architects its own stack well, sovereignty at scale only compounds when adjacent agencies in the same public sector domain converge across these dimensions.

Aligned governance approaches. The IMDA framework provides the national foundation; operationalising it requires common assessment methodologies, audit practices and escalation protocols that agencies in the same domain genuinely adopt. MAS Project MindForge has shown what this looks like, bringing more than 30 financial institutions and technology partners together to construct a comprehensive AI Risk Management Operationalisation Handbook[14], a methodology directly transferable to the healthcare and homeland security clusters.

Common operating models. Agencies in the same cluster should not each reinvent how AI is procured, deployed, monitored and retired. A shared reference architecture for AI operations, covering model documentation, data lineage, incident response and human oversight, reduces risk and accelerates adoption at the same time.

Shared workforce and capability baseline. The organisational imperatives of Section 2 apply beyond the single agency. A health cluster that trains its AI workforce to consistent standards across MOH, Synapxe and the regional clusters will build more durable sovereign capability than one where each entity defines its own approach.

A national proof point: Saudi Arabia's Data and Artificial Intelligence Authority (SDAIA)

Established in 2019 and built in close partnership with Accenture, SDAIA illustrates what aligned national capability looks like at scale. Together, SDAIA and Accenture established the national cloud and the National Data Bank, now connected to more than 200 government systems, and launched Estishraf, a national insights and analytics platform that delivered an estimated value of more than US$10 billion across over 85 government entities. In parallel, SDAIA trained more than 45,000 professionals to build the sovereign workforce the strategy depends on.

The lesson for Singapore is not the scale of compute or budget. It is the operating model: a clear anchor for national coordination, shared infrastructure that delivers increasing returns across domains, and a deliberate talent agenda built for longevity.[15]

04

From control to confidence

Singapore's public sector stands at an inflection point on sovereign AI. The foundations are established, the policy signals are clear and the appetite for ambition is genuine. What remains is the question of execution.

If sovereign AI is pursued as a procurement question, which vendor, which cloud, which model, the result is a collection of locally trusted layers that fall short of genuine sovereignty. If it is pursued as a system question, how to maintain control, optionality and resilience across every layer of the stack, across every agency cluster and across time, the result is a national capability that compounds. The four plays of selective sovereignty, mission-critical operationalisation, ecosystem orchestration and cross-agency alignment are the practical substance of that system approach. They are mutually reinforcing, achievable within the current budget and policy envelope, and the difference between sovereign AI that creates national value and sovereign AI that creates national risk.



Next steps for public sector leaders:

1. **Define your sovereignty perimeter:** Identify which workloads genuinely warrant sovereign treatment and which can run safely on global infrastructure.

2. **Commission a four-layer architecture audit:** Benchmark current infrastructure, data residency, model dependencies and application-layer controls, and identify where sovereignty is real, where it is assumed and where it is missing.

3. **Convene your cluster:** Initiate a structured conversation with adjacent agencies on shared governance standards, common operating models and workforce strategies, and identity infrastructure for autonomous systems.

> The next decade will determine whether Singapore's digital advantage becomes sovereign intelligence: the models, decisions and governance that distinguish jurisdictions that command AI from those commanded by it.

Why Accenture, and why now

The partner question for Singapore's sovereign AI agenda is not simply who can deliver, but who can deliver in a way that does not recreate the dependencies the strategy is designed to escape. That requires credibility at both ends: globally connected to where AI capability is being built and frontier risks are being understood and locally embedded in how Singapore's public sector actually works. Few partners can show both ends as directly. Accenture worked with the Monetary Authority of Singapore to establish the Project MindForge initiative that define responsible AI readiness in financial services. Accenture also delivered the National Digital Identity programme behind Singpass, the identity infrastructure that secures citizens' interactions with government.

That local foundation is matched by global scale. Accenture invests at scale in AI, with 80,000 AI professionals and a US$3 billion investment since 2023. Singaporean professionals are leading frontier work in areas such as agent assurance and bringing those capabilities directly back into local delivery. Accenture Singapore's Public Safety Suite, spanning investigation, command and response, and rehabilitation, translates those global capabilities into solutions purpose-built for the operational reality of Singapore's Home Team agencies.

References

[1] Sovereign AI: Own Your AI Future, Accenture, November 3, 2025

[2] Ibid.

[3] AI Continent Action Plan & AI Gigafactories Call, European Commission, April 10, 2026

[4] Cabinet Approves IndiaAI Mission, PMIndia, March 7, 2024

[5] HUMAIN and NVIDIA Announce Strategic Partnership to Build AI Factories of the Future in Saudi Arabia, Nvidia, May 13, 2025

[6] DORA Report 2025, Dora Research, 2025

[7] AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking, Gerlich, M., 2025

[8] Partnership between Accenture and Mistral Enables Scalable AI that Delivers Strategic Autonomy for Customers, Accenture, February 26, 2026

[9] HTX and Microsoft Announce Strategic Agreement to Develop Singapore's First Sovereign Cloud, HTX, February 10, 2022

[10] Assessing Claude Mythos Preview's Cybersecurity Capabilities, Anthropic, April 7, 2026

[11] Agentic Labour and the CIO's New Mandate, Accenture, 2026

[12] Acquisition of Faculty to scale AI capabilities, Accenture, January 6, 2026

[13] Accenture and Palantir Expanded Global Strategic Partnership to Drive AI Reinvention, Accenture, December 16, 2025

[14] Project MindForge, Monetary Authority of Singapore, November 13, 2025

[15] Case study: Saudi Data & AI Authority (SDAIA), Accenture, n.d.



Authors



Eng Han Heng

Defence & Public Safety Lead, Accenture Singapore



Grace Yip

Health & Public Service Strategy & Consulting Lead, Accenture Singapore

Acknowledgements

The authors would like to thank the following people for their contributions to the report –

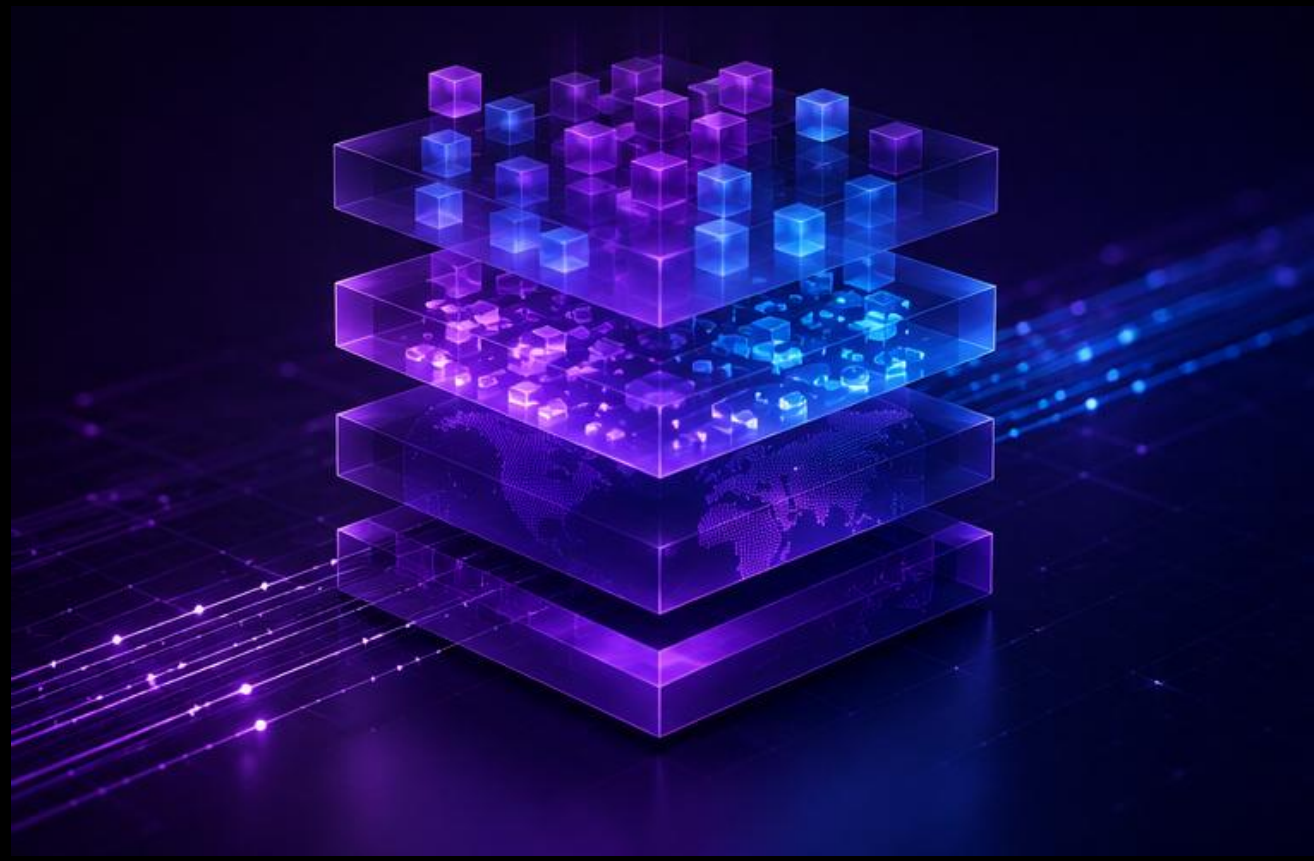
Research & Editorial: Jeromy Chia, Neo Hui Mei, Rebecca Tan, Minnie Chantpakpimon

Marketing & Communications: Aisha Hussain, Sharen Su

Special thanks to Louis Kwang and Eugene Siow for your insights and recommendations

About Accenture

Accenture is a leading solutions and global professional services company that helps the world's leading enterprises reinvent by building their digital core and unleashing the power of AI to create value at speed across the enterprise, bringing together the talent of our approximately 779,000 people, our proprietary assets and platforms, and deep ecosystem relationships. Our strategy is to be the reinvention partner of choice for our clients and to be the most AI-enabled, client-focused, great place to work in the world. Through our Reinvention Services we bring together our capabilities across strategy, consulting, technology, operations, Song and Industry X with our deep industry expertise to create and deliver solutions and services for our clients. Our purpose is to deliver on the promise of technology and human ingenuity, and we measure our success by the 360° value we create for all our stakeholders. Visit us at accenture.com.



Disclaimer: The material in this document reflects information available at the point in time at which this document was prepared as indicated by the date provided on the front page, however the global situation is rapidly evolving and the position may change. This content is provided for general information purposes only, does not take into account the reader's specific circumstances, and is not intended to be used in place of consultation with our professional advisors. Accenture disclaims, to the fullest extent permitted by applicable law, any and all liability for the accuracy and completeness of the information in this document and for any acts or omissions made based on such information. Accenture does not provide legal, regulatory, audit, or tax advice. Readers are responsible for obtaining such advice from their own legal counsel or other licensed professionals. This document refers to marks owned by third parties. All such third-party marks are the property of their respective owners. No sponsorship, endorsement or approval of this content by the owners of such marks is intended, expressed or implied.

Copyright © 2026 Accenture. All rights reserved. Accenture and its logo are registered trademarks of Accenture.