

In collaboration with Global Future
Council on Next Generation Computing



The Next-Generation Computing Framework

WHITE PAPER
OCTOBER 2026

non enim nigh euismod tincidunt ut laoreet dolore magna aliquam erat volutpat. Ut wisi enim ad

Contents

Foreword	3
Executive summary	4
Introduction	5
1 The next-generation computing acceleration framework	7
2 Sovereignty and interdependence	8
3 Cybersecurity and resilience	15
4 Sustainability and cross-council integration	19
5 Scorecard	23
Conclusion	31
Contributors	32
Endnotes	34

Disclaimer

This document is published by the World Economic Forum as a contribution to a project, insight area or interaction. The findings, interpretations and conclusions expressed herein are a result of a collaborative process facilitated and endorsed by the World Economic Forum but whose results do not necessarily represent the views of the World Economic Forum, nor the entirety of its Members, Partners or other stakeholders.

© 2026 World Economic Forum. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, including photocopying and recording, or by any information storage and retrieval system.

Foreword



Azeem Azhar
Chief Executive Officer,
Exponential View

Nations' ability to deliver public services and to make their economies more productive now depends on computing. Reliable and affordable access to information technology is now a condition of economic participation. However, the rapid evolution of computing, particularly artificial intelligence (AI), is turning it into a strategic infrastructure on which social and economic outcomes will depend.

Delivering on these capabilities is a determinant of future success. Next-generation computing requires more than access to it. It requires reforms to electricity supply, data access and skilled people. Bringing all this together requires synchrony across technology, training, market design, regulation, and public and private partnerships, as well as an understanding of how technologies are evolving.

Understandably, a large focus today is on AI. This must be a priority for any nation thinking about its economic development into the 21st century. Despite this, other technologies are maturing rapidly. Quantum computing, neuromorphic computing and other technologies are emerging



Jennifer Tour Chayes
Dean of the College of
Computing, Data Science,
and Society, University of
California at Berkeley

at different rates, and they may serve different parts of the market. Leaders must keep their options open to adopt these technologies as they mature.

Governments must ensure they retain agency over their choices in this domain, given the competitive and interdependent nature of technology ecosystems. They also need to ensure they consider security and resilience in their planning, alongside energy and sustainability considerations.

This report, written by the World Economic Forum in collaboration with its Global Future Council on Next Generation Computing, offers a practical diagnostic framework to support that work. Our scorecard sits across seven dimensions: compute and connectivity, infrastructure, data governance, talent, partnerships, cybersecurity and deployment. It is a diagnostic tool to help governments meet their ambitions and figure out what type of action is needed to achieve them.

Getting this right should result in tangible progress, more productive firms, more effective public services and more capable people.

Executive summary

As traditional computing reaches its limits, next-generation computing is expected to transform industries and societies over the next three to five years.

Computing is shifting from a product to a national infrastructure, much like telephone networks, electricity grids and road systems. It is becoming an essential utility, with entire segments of GDP (gross domestic product) relying on it. As it becomes more utility-like, broad and affordable access, reliability and consistency become increasingly important. Risks move from the buyer to the entire economy. Additionally, compute becomes a strategic asset subject to state power rather than merely market forces.

This transition is happening rapidly. In 2022, investment in computing infrastructure worldwide reached \$97 billion; in 2025, it reached \$448 billion. By 2027, however, it is expected to exceed \$1.5 trillion – a 15-fold increase in just five years.^{1,2}

Next-generation computing reflects the simple fact: computing is no longer something a nation buys; it is something nations depend on. While other pieces of infrastructure, such as sewer networks and road systems, last for decades or centuries with maintenance, computing infrastructure will be a continuously evolving, metamorphosing collection of capabilities.

Today's focus may be on graphics processing units (GPUs) and AI data centres, but other modalities are emerging rapidly. Consider compute focused on AI inference, the edge, space or new modalities like quantum computing or biocomputing. These are arriving on overlapping timelines, putting strain on energy grids, supply chains, talent pools, capital markets and planning policies. Most critically, they are arriving faster than traditional decision cycles. This white paper aims to help decision-makers evaluate their readiness and maturity for computing as infrastructure and next-generation computing. It treats compute, energy requirements, national autonomy, skills, data and security not as isolated, but as interlinked components comprising a national infrastructure.

The next three to five years constitute a critical investment and transformation window, with projections expected to increase almost tenfold. As public and private actors accelerate commitments to infrastructure, capacity and research, the imperative is clear: harnessing next-generation computing depends on building capabilities that are secure, sustainable, and inclusive.

The acceleration framework and scorecard support leaders to make informed infrastructure and policy decisions about next-generation computing, by understanding the timing, the impact, and the trade-offs between immediate needs (like AI/GPU deployment) and the emerging technologies (quantum, neuromorphic), while balancing sovereignty and access, cyber resilience and sustainability.

The paper's insights are drawn from the World Economic Forum's Global Future Council on Next Generation Computing, comprising diverse technology experts on next-generation computing, quantum computing and networks, space technologies and cybersecurity leaders and experts. We collaborated with and consulted experts and leaders from the Global Future Council on Cybersecurity, the Global Future Council on Energy Technology Frontiers, the Global Future Council on Generative Biology, and the Quantum Economy Network community.

Thanks to the council's work, the next-generation acceleration framework provides a useful reflection tool to inform decision-making by triggering conversations across ministries and business units and enabling greater collaborations across organizations. It is not designed as a ranking but as a framework to help understand status and priorities and to make informed infrastructure and policy decisions about next-generation computing.

Introduction

Next-generation computing scorecard to support reflection and decision-making.

“Next-generation computing is an evolving set of advanced and emerging technologies that expand the capabilities of traditional computing.


The next generation of computing will have a wider impact on the digital economy. Over the next three to five years, governments and organizations are likely to make major investments in next-generation computing systems that will underpin artificial intelligence (AI), advanced research and industrial transformation. These investments are predicted to exceed \$1 trillion globally by 2027.

While investment in AI infrastructure is dominated by the private sector, governments are also making substantial, multi-year commitments to other next-generation computing capabilities. For example, the European Union injected €8.2 billion into the European High-Performance Computing Joint Undertaking (EuroHPC JU) programme,³

the United Kingdom government is committing up to £2 billion through 2030 for the UK Compute Roadmap,⁴ the United States has planned \$2 billion for quantum incentives,⁵ Singapore has committed S\$300 million (Singaporean dollars) to their National Quantum Strategy⁶ and Canada’s parallel investments in sovereign AI compute and quantum technologies, which exceed C\$1.2 million (Canadian dollars).⁷

These investments will not only determine who leads but also who has access to more processing power – shaping the current digital and intelligent economy, energy consumption, critical infrastructure resilience, and autonomy and sovereignty in such a digitally interdependent world.

FIGURE 1 Computing architecture layers



Next-generation computing is not a single technology, an infrastructure category, nor a full replacement for traditional computing layers that are already market-ready. Rather, it refers to an evolving set of advanced and emerging technologies that expand the capabilities of traditional computing as it shifts from being one factor of production to a key infrastructure capability of the future economy.

It is a complementary and emerging base layer that, when combined with traditional ones, could open new transformational opportunities. It is a shifting constellation of compute paradigms, including advanced AI infrastructure, high-performance computing, quantum science, neuromorphic systems, edge and in-orbit computing, and biological or hybrid forms of computation. These systems can be combined to accelerate findings

“ Investment in computing infrastructure worldwide reached \$97 billion in 2022 and \$459 billion in 2025.

while differing in maturity, physical requirements, security risks and commercial pathways. Still, they share one strategic feature: they expand the ability of governments, firms and societies to model, optimize, automate, discover and decide at scales that were previously impossible.

The investment in computing infrastructure worldwide has been increasing rapidly and is expected to grow tenfold by 2027. The boost in infrastructure investments aligns with the need for more computing power to deliver new transformational value. Investment in computing infrastructure worldwide reached \$97 billion in 2022 and \$459 billion in 2025.^{8,9} Given the likely scale of investment in next-generation compute over the next three to five years, policy-makers and organizations must reflect on the transformational opportunities and risks for industries that are built on these new capabilities. These opportunities and risks can be sorted into three enabling pillars: sovereignty and interdependence, cybersecurity and resilience, and sustainability. These pillars are interlinked – decisions made in one will shape outcomes in the others. For example:

- Expanding access to next-generation computing can widen innovation, but it can also intensify demand for energy and infrastructure.
- Pursuing efficiency can lower costs, but it can also create new dependencies.
- Concentrating capability can speed deployment, but also increase exposure to geopolitical, operational and security risks.

The next-generation computing strategy discussion becomes largely qualitative through case-by-case assessments of specific technologies, geographies or risks without a common basis for comparing where a country stands relative to their peers, or its own ambitions.

This is a difficult balance for countries and organizations: having strong capabilities that translate into deployed capacity or having sustained economic, societal and strategic value.

This paper sets out a next-generation computing acceleration framework that helps countries and organizations to inform foundations and policy decisions on addressing the opportunities of next-generation computing. The framework comprises seven dimensions:



Compute and connectivity



Infrastructure



Data governance



Talent



Partnerships



Cybersecurity



Deployment

It examines the deeper strategic issues underlying those capabilities, grounded in three enabling pillars, to give decision-makers a practical way to think about the opportunities and trade-offs that will define the next generation of compute. Furthermore, the paper provides a scorecard approach to address this challenge by enabling a country or organization's capacity to access, adopt, deploy and derive value from next-generation computing capabilities.

The next-generation acceleration framework draws on consultations with selected Global Future Councils and C-suite communities, reflecting that next-generation computing decisions are no longer confined to government alone.

1

The next-generation computing acceleration framework

Seven dimensions, spanning supply enablers and demand outcomes, guide next-generation computing investment and policy decisions.

Given the likely scale of investment in next-generation compute over the next three to five years and beyond, the next-generation computing acceleration framework equips leaders and decision-makers with infrastructure and policy

decisions about next-generation computing, based on seven dimensions. The framework dimensions represent key areas for reflection, supported by a self-assessment scorecard that is technology- and time-bound-agnostic.

FIGURE 2 The next-generation computing acceleration framework



Sovereignty and interdependence

Next-generation computing complexity and interconnectivity are redefining the notion of sovereignty, from control to capability.

Ensuring agency over next-generation compute choices is key to preserving accessibility and the ability to decide where compute is hosted, who controls it,

and how organizations can exit or switch providers – while recognizing that hardware, data, talent and supply chains remain globally interdependent.

2.1 Why sovereignty looks different in next-generation computing

“ Technological sovereignty is grounded in accessibility, ownership and control of discrete assets – infrastructure, firms or intellectual property within national borders.

Sovereignty, as a formal legal and political concept, is aligned with the supreme authority of a state to govern within its territory and to make decisions independently of other states. Traditional notions of sovereignty have therefore been closely associated with territorial control and national autonomy.

The digital notion of sovereignty, per the Organisation for Economic Cooperation and Development (OECD), refers to the ability of governments to make autonomous decisions across data, software, standards, infrastructure and public services.¹⁰ Technological sovereignty is grounded in accessibility, ownership and control of discrete assets – infrastructure, firms or intellectual property within national borders. Next-generation computing is challenging this model on different fronts.

Agency and continuity

In this context, sovereignty is better understood as the preservation – the ability of a country, institution or organization to make and execute strategic choices about where systems are hosted, how they are used, who controls the data, how they are updated and under what conditions they can be exited or reconfigured. Agency shifts from nominal ownership to effective choice, including choice over hosting, use, licensing, standards, patch management, firmware and update life cycles.

For public-sector and critical-infrastructure systems, sovereignty also includes the ability to preserve accessibility and service continuity during disruptions through redundancy and guaranteed operational continuity.

Interconnected environments and supply chains

Compute today functions as an interconnected environment rather than a standalone sector or asset. Hardware, software, data, materials, energy, talent and capital evolve on overlapping and uneven innovation curves, creating deep structural interdependence across borders. As illustrated in emerging “third stack” strategies beyond the United States–China binary,¹¹ what travels across borders is not just technologies but entire stack configurations, model weights, licensing regimes, cloud dependencies, information data and governance architectures, which collectively shape who retains authorship and control.

There is also the concept of owning access to the supply chain. Even if a country owns the computing but has a risky supply chain spread across nations with single-point dependencies, then the risk to the supply chain and the sovereignty is high. Chips, fabrication capacity, rare materials, advanced packaging, cooling systems and hyperscale cloud services frequently sit outside national control. Export controls, licensing restrictions or geopolitical tensions can rapidly transform technical dependence into strategic vulnerability. Therefore, supply chain and value chain risk management are integral components of sovereignty, not peripheral considerations. Sovereignty risk is often misunderstood as a binary condition (dependent or independent), whereas in practice, it manifests as degrees of exposure across different layers of the stack. Second- and third-order dependencies are frequently underestimated – for example, reliance not only on foreign chip fabrication, but on firmware

“ Entangled innovation cycles across hardware, software, energy and data complicate linear sovereignty strategies.

updates, toolchains, maintenance ecosystems and embedded design standards that are harder to substitute once embedded.

Agency across the supply chain must therefore be tested, not presumed. A sovereign posture is only meaningful if countries and organizations can verify that they retain operational control when systems are disconnected, suppliers are unavailable, updates are restricted or geopolitical conditions change abruptly. This makes “red teaming” of critical supply chains an important governance practice: stress-testing whether autonomy actually exists in moments of crisis, rather than relying on peacetime assumptions of access. Such testing should explicitly validate assumptions embedded in supply-chain security and de-risking strategies, including logistics, weather, trade continuity, executive action and adversarial disruption.

The relevant risk horizon extends beyond classic supply-chain management. Critical dependencies may be disrupted by natural disasters, adversarial action, export restrictions, rapid executive decisions or sudden exclusion from previously available technologies and services. What was once optimized for efficiency has proved structurally fragile, and in next-generation computing, agility, attestability and the ability to operate through disruption increasingly matter as much as low cost. In this sense, sovereignty increasingly converges with assurance of supply and, in high-dependence environments, with survivability: the capacity to maintain access to essential compute capabilities when political, environmental or market conditions deteriorate.

Next-generation computing is a group of technologies, but a constellation of paradigms

Each technology it comprises – such as high-performance computing, quantum computing, neuromorphic computing, edge computing and in-orbit systems – has distinct infrastructure, regulatory and geopolitical characteristics. While considering which stacks are most useful in certain scenarios, nations, regions and states may need to assess compute against context and make strategic decisions about what sovereignty means in each context. For example, sovereignty in quantum computing may lie in cryogenic supply chains and intellectual property (IP) in error correction, while for edge systems, integration with domestic digital infrastructure will be key. Sovereignty is therefore differentiated across compute paradigms rather than being uniform. AI software dependencies can shift relatively quickly, whereas semiconductor fabrication ecosystems and quantum hardware supply chains embed path dependencies that can last for decades. Sovereignty strategies must therefore distinguish between fast-moving and slow-moving layers of the stack.

As the future compute landscape will see different types of systems coexist, each serving different problem classes and strategic functions, countries

will need differentiated strategies across paradigms rather than a single doctrine applied uniformly across all forms of computing. For dual-use areas such as quantum and space-related computing, the role of government is likely to be more pronounced, and sovereignty may be shaped more directly by security frameworks, export controls, alliance structures and public investment strategies than in more mature commercial technology domains.

Entangled innovation cycles across hardware, software, energy and data complicate linear sovereignty strategies

Physical constraints, such as power availability, water use, land, fabrication capacity and supply chains, limit national self-sufficiency; these constraints affect countries differently depending on size, income level and computing maturity. Energy, in particular, becomes a structural determinant of technology sovereignty. Next-generation computing systems require access to reliable power. Today’s large-scale AI data centres demand gigawatts (GW) of available power.¹² The United States’ net electricity generation was essentially flat between 2008 and 2024 and is now growing by approximately 9 terawatt-hours (TWh) per month. Expected load growth to 2030 has risen roughly sevenfold since 2022, from 24 GW to 166 GW, with data centres accounting for around 55% of that increase.¹³ Reliable electricity supply, flexible market rules and resilient grids will increasingly define who can scale advanced compute and under what conditions.

Energy, infrastructure and information

These three factors are especially consequential for supply chains for next-generation computing, and sovereignty depends on the interaction among them rather than on control of any one layer. Recent shifts in policy and market priorities suggest that the sovereignty–sustainability debate is increasingly being reframed through efficiency and energy allocation, as governments confront trade-offs between advanced compute expansion and other societal energy demands. The ability to attract and retain talent is also tied to these material conditions: researchers, engineers and firms increasingly cluster where compute infrastructure, information access and energy are available at scale. Talent mobility, therefore, becomes both a consequence and a determinant of sovereignty, with implications for firms, national ecosystems, universities and public research institutions that feed the broader compute economy.

Collaboration, risk management strategies and context-sensitive approaches to sovereignty in different types of computing may be meaningful. In this context, sovereignty in compute is less about eliminating dependence and more about shaping the terms, risks and evolution of dependence over time. Isolationist strategies, whether through full-stack nationalization or aggressive decoupling, are increasingly misaligned with the realities of compute

environments. For most countries, they are economically unviable; for others, they risk higher costs, slower innovation and reduced resilience. Instead, sovereignty increasingly resembles infrastructural authorship – the ability to configure, govern and adapt a stack, even when components are sourced globally, while retaining authority over deployment, data flows and system evolution.

This also means that market position alone can no longer be treated as a reliable proxy for agency. Large firms and markets may possess

scale, but face reduced room for manoeuvring when geopolitical shifts, supplier concentration or regulatory shocks undermine their ability to execute long-term plans. In some cases, smaller or more modular actors may be able to pivot more quickly than dominant players whose infrastructure, supplier relationships and demand commitments make them less adaptable. Sovereignty should therefore be evaluated not only at the level of states, but also through the interdependent agency of the firms, universities and infrastructure providers that countries rely upon to execute strategy.

2.2 From ownership to participation

Participation, rather than full ownership and control, is becoming the dominant logic for sovereign positioning in next-generation computing

As emerging “third stack” models demonstrate, states can retain agency not by replicating hyperscale stacks end-to-end, but by shaping how components are assembled, tuned and governed within their jurisdiction. It explores what meaningful participation looks like across different layers of the compute stack, how policy-makers can distinguish between strategic access and nominal control and where participation creates leverage, learning or agenda-setting power.

Meaningful participation

“Meaningful participation” is participation that preserves choice – choice over system design, standards, suppliers, licensing, upgrades, security controls and eventual exit. This is closely linked to architectural properties such as interoperability, modularity and composability, which make it easier to decouple from specific vendors, replace components and evolve systems over time without wholesale dependence on a single provider. Population-scale digital public infrastructure offers one illustration of this logic. Architectures inspired by modular, interoperable and composable design principles can create room for national adaptation and future reuse across multiple public-service functions, rather than embedding irreversible lock-in at the point of first deployment. Meaningful participation is therefore not only about being present in the ecosystem; it is about preserving the capacity to reconfigure the ecosystem on acceptable terms.

The tension between openness and strategic safeguards

While openness can accelerate innovation and resilience, poorly designed participation models can entrench asymmetric or exploitative dependencies. The analysis, therefore, focuses on

how governments can remain integrated into global ecosystems while preserving strategic choice and national interests. Licensing regimes, standards-setting forums and cloud governance frameworks increasingly become arenas where sovereignty is exercised indirectly, through configuration rather than exclusion. Open-source protocols, open standards and open-weight environments can strengthen sovereign participation by reducing dependence on proprietary bottlenecks and broadening the set of actors able to build, inspect, adapt and deploy key layers of the stack. Free and open-source software may face procurement barriers in some jurisdictions, yet it can materially improve visibility, resilience and substitutability in ways that closed systems cannot. Identity-layer and other foundational public digital infrastructure protocols illustrate how open, modular architectures can support sovereign scaling across multiple use cases, from identification and health systems to payments, insurance and public-service delivery. Open-weight AI models similarly create new forms of agency by allowing domestic fine-tuning, adaptation and experimentation even when frontier model development remains concentrated elsewhere. At the same time, openness does not remove material chokepoints: open models and open hardware still eventually depend on physical infrastructure, fabrication capacity, energy systems and manufacturing equipment that remain subject to geopolitical control.

Participation without leverage can create an illusion of sovereignty

Merely hosting infrastructure or attracting foreign R&D centres does not automatically translate into influence over standards, licensing or model governance. The same applies to strong research capabilities, public investment and infrastructure, which do not necessarily translate into strategic influence without commercialization pathways, industrial scale and effective coordination. Sovereign participation requires negotiated visibility into decision-making processes, access to retraining rights and influence over roadmap evolution.

“ Meaningful participation is about preserving the capacity to reconfigure the ecosystem on acceptable terms.

Participation also requires early and deliberate positioning in the parts of the stack where influence is created. Countries that remain only downstream users should not expect to exercise agency over upstream technologies, standards or system

direction. Agency is reinforced when investment, procurement and partnership strategies are intentionally distributed across upstream, midstream and downstream layers rather than concentrated only at the point of end use.

2.3 Key domains of interdependence

“Interdependence shapes compute sovereignty and highlights how segregated policy approaches create blind spots.”

Interdependence shapes compute sovereignty and highlights how segregated policy approaches create blind spots. Sovereignty should be considered an emergent property of interactions across domains, rather than something that can be achieved within any policy area. One practical implication is that policy attention should distinguish among upstream dependencies such as chips, tools and intellectual property; midstream dependencies such as platforms, clouds and orchestration layers; and downstream dependencies in applications, users and public-service delivery. Concentrating investment only in one of these layers weakens investment, especially when countries expect influence upstream while investing downstream.

Interdependence focuses on the following domains:

- Compute infrastructure and hardware supply chains, including chips, fabrication capacity, cloud infrastructure and energy systems
- Data assets and data governance, including data generation, curation, access and language coverage
- Talent, skills and mobility, shaped by education systems, migration regimes and research ecosystems
- Standards, platforms and interoperability, which influence market access and innovation pathways

- Capital, partnerships and ecosystems, including financing models and public-private collaboration

Decisions in one domain can amplify or undermine sovereignty outcomes in others

These domains interact dynamically: for example, data residency requirements influence cloud architecture; cloud architecture influences energy demand; energy constraints shape siting decisions; siting decisions affect geopolitical exposure. Sovereignty, therefore, emerges from system-level coherence, not isolated policy wins. This system-level view is increasingly necessary because long-standing assumptions about stability have weakened: global market position, open trade, predictable regulation and slow-moving political change can no longer be treated as dependable background conditions. The result is a more volatile environment in which economic power, geopolitical leverage and technical dependence are more tightly interlinked than before.

Fragmentation across ministries often undermines sovereignty objectives

Industrial policy, digital regulation, cybersecurity strategy and energy planning are frequently developed in parallel rather than in coordination. Without cross-domain governance, interventions in one domain (e.g. data localization) can unintentionally increase vulnerabilities in another (e.g. cloud concentration or grid strain).

2.4 Managing dependencies and chokepoints

Policy-makers can shift from implicit dependence to explicit governance of critical dependencies and chokepoints. This section explores how to identify single points of failure, assess trade-offs between diversification and concentration, and determine where redundancy or shared capacity is justified despite higher costs. These dependencies may be specific to the technology system.

Distinguishing between acceptable interdependence and unacceptable chokepoints is key

Not all dependencies are strategic risks; some generate resilience through diversification. The policy challenge is identifying which dependencies concentrate control asymmetrically and which distribute risk constructively. A central test is whether a dependency preserves or removes the ability to execute a strategic plan. When a country or firm cannot change providers, update a system, maintain operations or exit an arrangement on acceptable terms, dependence has become constraining rather than merely inconvenient. This requires closer attention to the difference between

resilience, robustness and survivability. For some critical missions, the relevant question is no longer only whether systems are robust in normal conditions, but whether access to compute can survive systemic shocks and exclusion scenarios.

Rather than aiming to eliminate chokepoints – which is often an unrealistic objective – governments may instead seek to map, monitor and negotiate them. This includes developing early warning mechanisms for export restrictions, concentration risks in model providers, hyperscaler dependence and strategic materials. Mapping should extend to points where political authority can override commercial expectations, including trade controls, procurement restrictions, executive actions and alliance-based access conditions. It should also identify where geographic concentration has become normalized through efficiency logic despite clear systemic exposure, as in cases where critical industrial capacity is clustered in a small number of locations that are vulnerable to geopolitical or environmental disruption.

Public-private coordination plays an important role, as many strategically significant assets are privately owned yet systemically important

Trust, transparency, contractual design and accountability are central to achieving resilience. Procedural sovereignty – the ability to set conditions for deployment, auditing, retraining and data handling – may provide leverage even

when physical infrastructure is sourced globally. Because many strategic dependencies are privately governed, sovereignty increasingly depends on whether public authorities and private actors can align around contingency planning, update rights, continuity obligations and access guarantees under adverse scenarios.

Contractual design is increasingly used as a sovereignty instrument

Access agreements, data-sharing clauses, retraining rights and exit provisions can determine whether dependence remains manageable or becomes structural. Sovereignty may therefore be embedded in procurement templates as much as in industrial strategy. Procurement is especially consequential because it can either lock countries into opaque and rigid systems or require interoperability, modularity, open interfaces, auditability, portability and meaningful exit rights from the outset.

Investments should be assigned to policies that enable agency, not merely to the acquisition of technology assets

This is particularly important in public-sector and dual-use settings, where the right of a service to continue operating should not depend on a single vendor, a single jurisdiction or a single untested chain of permissions.

2.5 Policy levers for sovereign participation

“ Given the structural interdependence of compute ecosystems, no country can secure the full stack independently.

Sovereign participation in next-generation computing requires coordinated action across policy, markets and ecosystem development, with the objective of preserving agency: the ability to decide where compute systems are hosted, how they are governed, how they evolve and when they can be adapted, switched or exited.

Policy levers for governments: Rather than advocating specific policy choices, these three sets of levers help governments shape their position within global compute ecosystems without defaulting to isolationist strategies. These levers can be combined into coherent policy portfolios that preserve flexibility and optionality over time.

Partnerships and alliances: Given the structural interdependence of compute ecosystems, no country can secure the full stack independently. Governments, therefore, must design bilateral and multilateral partnerships that diversify risk and provide access to complementary capabilities across infrastructure, talent and innovation systems. Increasingly, this includes participation in global or regional infrastructure initiatives and research collaborations that embed national priorities into shared architectures. Forming non-

traditional alliances can expand flexibility and reduce dependency.

Standards, interoperability and ecosystem governance: Standards-setting is not merely a technical exercise but a strategic arena in which influence over future architectures, upgrade pathways, and market access is negotiated. Participating in standards bodies can help shape future technical choices that enable reasonable access. With fast-moving technologies, like AI, standards processes may be too slow, so governments may want to insist on interoperability mandates to reduce lock-in and maintain the option of local industrial participation.

Procurement and market-shaping instruments: Public procurement is one of the most direct tools available to governments to embed sovereignty considerations into the design of compute systems. Procurement frameworks can require interoperability, modularity, auditability and portability, as well as meaningful exit options and control over updates, licensing, and data flows. In doing so, they shift sovereignty from an abstract objective to an operational requirement embedded in contracts and system architectures. Public-sector

“ Sovereignty is not only a function of structural position, but also of when and how countries intervene.

demand can also serve as an anchor for domestic ecosystems, shaping market incentives and lowering barriers to entry.

Investment and industrial policy levers: These are closely related. Strategic public investment, whether in infrastructure, research or ecosystem development, can determine where the stack creates agency. This includes investments in upstream capabilities such as semiconductors and advanced research, midstream platforms such as cloud and orchestration layers, and downstream applications and services. Importantly, investment strategies need to be intentional about their position in the stack: countries that invest only at the application layer should not expect influence over upstream standards or architectures. In emerging domains such as quantum or neuromorphic computing, early and sustained investment can shape long-term positioning more effectively than late-stage catch-up.

A further set of levers relates to openness and architectural choices: Open-source software, open standards and open-weight models can expand sovereign participation by increasing visibility into systems, enabling local adaptation and reducing dependence on proprietary bottlenecks. Modular, interoperable and composable architectures, often exemplified in digital public infrastructure, can preserve long-term flexibility by allowing components to be replaced or upgraded without systemic disruption. However, openness does not eliminate structural dependencies, particularly where compute ultimately relies on concentrated physical infrastructure, fabrication capacity or energy systems. Governments must therefore combine openness with an understanding of material constraints.

Talent and ecosystem development constitute another critical lever: Attracting, retaining and mobilizing talent is closely tied to accessing

compute infrastructure, energy and data. Policies that strengthen research ecosystems, enable talent mobility and facilitate collaboration between academia, industry and government can reinforce sovereign positioning. Conversely, a lack of access to infrastructure or restrictive institutional arrangements moves talent towards jurisdictions where compute resources are more readily available, weakening domestic capacity over time.

Governance, risk management and validation mechanisms play a central role in ensuring that sovereignty is not only designed but also operationalized: Governments can establish processes to map and monitor dependencies, identify chokepoints and develop early warning systems for disruptions. Increasingly, this includes the use of stress testing and “red teaming” approaches to verify whether the assumed agency holds up under adverse conditions. Contractual design, regulatory frameworks and public-private coordination mechanisms can further ensure that critical capabilities remain accessible, updatable and controllable when needed.

Across all these levers, timing is a strategic variable. If done early, participating in standard-setting, investing in emerging ecosystems and experimenting through pilot deployments can significantly shape long-term outcomes. Sovereignty is therefore not only a function of structural position, but also of when and how countries intervene. Taken together, these levers highlight that sovereign participation is less about maximizing control and more about preserving optionality. The most effective strategies will enable countries to adapt over time by reconfiguring partnerships, substituting components and evolving their roles within the global compute ecosystem as technologies, markets and geopolitical conditions change.

2.6 Illustrative examples

- **Countries pursuing modular “third stack” approaches that combine foreign hardware with domestic cloud governance and localized model tuning:** Companies like Amazon, Google, and Microsoft have set up data centres in India to meet the current and future requirements for data protection arising from India’s Data Protection and Data Privacy Act 2023.
- **Hybrid sovereignty models in which participation in global open-weight ecosystems enables contextual fine-tuning while maintaining procedural control over deployment:** Open-weight models can serve as sovereignty infrastructure, allowing countries with limited resources for frontier-model training to retain meaningful control over adaptation and deployment without disconnecting from

global innovation. This is sometimes described as AI sovereignty as managed interdependence rather than isolation. OpenEuroLLM consortium exemplifies this approach by developing open-source European AI models that can be adapted and deployed locally. This reflects the European Commission’s broader view of open-source AI as a potential lever for sovereignty because organizations can deploy and modify models without relying exclusively on proprietary model providers.

- **Ecosystem sovereignty means managing critical dependencies through trusted partners rather than seeking national self-sufficiency across the technology stack:** France builds their domestic capability at a strategic importance layer while accepting

that sovereignty comes from European and allied interdependence across the stack, rather than attempts to nationalize each layer. The 2026 ASML investment in Mistral, which develops open-weight models and is expanding European AI infrastructure, exemplifies the type of European industrial alliance intended to collectively strengthen control over strategic technologies.

- **Overconcentration on a single hyperscaler, fabrication partner or model provider might create systemic risks or catalyse the acceleration of domestic capability:** The extremely high concentration of semiconductor manufacturing in Taiwan, China, illustrates how

geopolitical disruption at a single critical node could constrain global access to computing technologies. Meanwhile, restrictions on China's access to GPU compute may accelerate the development of domestic alternatives enabled by its financial, technical and governance capacity.

These examples may also reflect situations where countries misjudged the stickiness of compute dependencies, for instance, discovering that switching cloud providers required rearchitecting entire application layers, or that migrating models required renegotiating licensing constraints. Such cases illustrate that sovereignty risks often crystallize not at adoption, but at attempted exit.

Cybersecurity and resilience

Next-generation computing introduces a qualitatively different risk landscape, while offering new opportunities to enhance cybersecurity.

The expansion of next-generation computing depends on multiple paradigms and their connectivity. Cybersecurity and resilience across infrastructure, technologies and supply chains

is a key requirement that supports governments and organizations in anticipating emerging risks, strengthening preparedness and mitigating evolving threats to next-generation computing.

3.1 Why cybersecurity and resilience matter for next-generation computing

“ Cybersecurity and resilience are systemic properties of the compute environment itself.

Next-generation computing introduces a qualitatively different risk landscape and offers opportunities to enhance security through computation. As compute systems become more complex, distributed, connected and interdependent, cybersecurity and resilience can no longer be treated as discrete technical issues or safeguards. They are systemic properties of the compute environment itself. Unlike earlier generations of digital infrastructure, next-generation systems (e.g. quantum, neuromorphic, hybrid stacks) combine physical sensitivity, software complexity and globalized supply chains in ways that expand both the attack surface and the scale of potential disruption. Advanced AI systems are already leading to a parabolic growth in the discovery of critical security vulnerabilities.¹⁴

Cybersecurity and resilience, therefore, extend across the full life cycle of compute systems, from early-stage research and design to deployment, operation and decommissioning, requiring integration into innovation ecosystems and value chains rather than isolated technical controls.^{15,16}

This section explains why cybersecurity and resilience in next-generation computing must be understood as systemic and cross-cutting. It will examine how growing interdependence and rapid technological change can amplify risks across systems, institutions and jurisdictions while also creating opportunities for enhanced security capabilities.

3.2 Key risk domains in next-generation compute systems

Cybersecurity and resilience risks are emerging, with vulnerabilities in one domain capable of propagating across the broader system. Risk domains should be understood as interdependent layers, in which vulnerabilities often arise from interactions among them rather than from isolated components. Increasing interdependence across global compute stacks can therefore intensify cascading failures, geopolitical exposure and unintended externalities.¹⁷

The analysis will cover risks across physical and digital infrastructure, globally distributed supply chains, data, models and algorithms, software

platforms and cloud dependencies, and human, organizational and governance factors. Particular attention will be paid to how these domains interact and how isolated approaches to cybersecurity can obscure systemic exposure. Limited visibility into upstream and cross-border dependencies can create blind spots in risk assessment, particularly where responsibilities are distributed across multiple actors and jurisdictions. The pace and diversity of technological change further challenge cybersecurity models designed for more stable, bounded systems, requiring greater emphasis on systemic risk and the conditions that shape resilience over time.

Infrastructure vulnerabilities include not only cyber intrusion but also hardware-level disruption, physical tampering, environmental interference and side-channel exposure. In emerging systems such as quantum platforms, hardware fragility and environmental sensitivity introduce new challenges (e.g. microwave, thermal or electromagnetic manipulation) that require dedicated testing and validation regimes. While side-channel attacks have been extensively studied in classical computing, equivalent threat models for emerging paradigms remain less mature, particularly in areas such as quantum communication and hardware-level attack surfaces.

Supply chain exposure extends beyond availability risks to embedded security risk: dependency on single-source or geopolitically concentrated components can create structural vulnerabilities.¹⁸ Compute infrastructure pipelines can themselves become cybersecurity challenges, as vulnerabilities in upstream materials, components, fabrication processes or integration layers propagate into downstream systems and create national security exposure. Supply chains, therefore, function not only as dependencies but also as transmission channels for risk across systems and jurisdictions. In cybersecurity contexts, supply chain fragility is an economic concern, a technology development hurdle

and a direct resilience issue. This includes both upstream dependencies (materials, fabrication, components) and downstream dependencies (integration, deployment, maintenance), as well as positioning relative to allied or partner ecosystems.

Data, models, control and algorithmic layers introduce additional risk domains, including model manipulation, data poisoning, privacy leakage and compromised integrity. As models become more central to decision-making systems, their robustness and verifiability become national resilience concerns. Increasing reliance on complex and opaque models may also introduce challenges for verification, auditing and trust assurance across different institutional and jurisdictional contexts.

The evolution of computing systems also creates a temporal risk domain. Resilience requires protection against known threats while anticipating changing threat models, maintaining long-term data confidentiality and developing migration pathways for transitions such as the adoption of post-quantum cryptography.¹⁹ It also requires forward-looking commitments from providers, including clear timelines for system migration and risk transition, to support long-term planning and reduce uncertainty. At the same time, adapting to new computing paradigms may create opportunities to strengthen security capabilities.

3.3 Resilience by design: policy levers

“ Resilience-by-design includes security, as well as building structures with forward-looking flexibility and redundancy to address future risks.

Policy levers shape incentives, architectures and behaviours across public and private actors. Resilience outcomes are shaped by early-stage incentives, including funding criteria, standards development and procurement requirements, which influence technology design choices before large-scale deployment.

Incentives can promote the development of information-secure computing methods, the use of standards, procurement, and domestic and international regulation to promote secure-by-design approaches; the importance of system architecture and design choices for long-term resilience; and the role of coordination mechanisms in incident response and crisis management.

Secure-by-design in this context includes integrating cybersecurity requirements at early technology readiness levels (TRLs), embedding resilience criteria into public R&D grants and procurement frameworks, and requiring providers to articulate forward-security timelines and migration strategies. The stage at which resilience considerations are introduced (e.g. early research phases versus post-deployment) can serve as an indicator of maturity in national approaches.

Resilience-by-design includes security, as well as building structures with forward-looking flexibility and redundancy to address future risks, such as national or collaborative secure computing testbeds, centres and ecosystems for emerging technologies. Such test environments allow governments and ecosystem actors to evaluate hardware-level vulnerabilities, side-channel exposure and systemic weaknesses before making large-scale deployment decisions. Testbeds can reduce the risk of premature strategic exclusion of technologies that may later become critical to resilience architectures. Maintaining such test environments also supports technical optionality, preserving expertise and capabilities in emerging technologies even where immediate deployment is not prioritized.

International collaboration is a core component of resilience, recognizing that no country can secure next-generation compute systems in isolation and that misaligned approaches can themselves create vulnerabilities. Resilience requires governance arrangements that manage interdependence rather than attempt absolute self-sufficiency, including structured cooperation with allied supply chains and shared testing frameworks. Comparative assessment across countries and alliances can further support understanding of relative dependencies, exposure and resilience positioning.

3.4 Assessment questions for the scorecard

The scorecard is intended as a diagnostic tool to support strategic reflection and internal dialogue, not as a benchmarking or ranking exercise.

Resilience-by-design maturity

These indicators assess the extent to which cybersecurity and resilience considerations are embedded early in next-generation compute strategies, rather than added after deployment.

- **Integration of resilience considerations:**
The degree to which cybersecurity and resilience are considered during the design and planning of next-generation compute infrastructure, systems and policies. This may include the stage (e.g. TRL level²⁰) at which resilience requirements are introduced and whether funding, procurement and regulatory processes systematically incorporate these criteria.
- **Policy-level alignment on resilience objectives:** Whether resilience objectives are coherently reflected across digital, industrial, research and security policy domains related to compute. This includes cross-agency coordination and the availability of technical expertise and tools within agencies not traditionally focused on cybersecurity. This also includes prioritizing the development of compute technologies with potentially greater security than mature compute technologies, as well as policies and funding programmes to develop secure emerging compute technologies.
- **Emphasis on preventive approaches:**
The extent to which secure-by-design and resilience-by-design principles are prioritized over primarily reactive or incident-driven approaches. This includes forward-security considerations, migration planning, and explicit articulation of long-term data protection horizons.
- **Standards:** Everlasting/long-term security/security by design and crypto agility

Awareness and governance of systemic risk

These indicators focus on understanding and governing systemic and cascading risks, rather than cataloguing individual threats.

- **Awareness of systemic interdependencies:**
Whether policy-makers have visibility into how cybersecurity risks can propagate across infrastructure, supply chains, platforms and organizational arrangements. This includes upstream and downstream supply chain dependencies, reliance on foreign or single-source components, talent and comparative positioning relative to allied or peer nations.
- **Recognition of evolving risk domains:**
The extent to which emerging risks related to

computation, algorithms, system architecture and physical sensitivity are acknowledged at the policy level. This includes recognition of hardware-level and side-channel risks in emerging computing paradigms that have not yet undergone extensive adversarial testing.

- **Institutional responsibility for resilience:**
The presence of clear governance arrangements or institutional ownership for overseeing cybersecurity and resilience of next-generation compute systems across public and private actors. This may include designated national coordination bodies, test infrastructure governance and structured collaboration between security agencies, research institutions and industry.

Comparison against other entities is also an important aspect. Domestic supply chains are interdependent with those of allied countries, involve multiple countries and create dependencies on other external entities. Compute infrastructure pipelines become cybersecurity challenges.

Ability to respond and adapt over time

These indicators assess readiness to adjust approaches as technologies, threat landscapes and geopolitical conditions evolve.

- **Preparedness for systemic disruption:**
Whether mechanisms exist to coordinate responses to incidents or broader disruptions affecting next-generation compute systems. This includes cross-sector crisis coordination, continuity planning for compute infrastructure, and stress-testing of interdependent systems.
- **Learning and adaptation mechanisms:**
The extent to which lessons from incidents, exercises or stress tests inform policy evolution and system design. This includes structured feedback loops between testbeds, research communities, regulators and deployment environments.
- **Anticipation of skills and capability needs:**
Awareness of how cybersecurity and resilience skill requirements are likely to evolve, and the ability to plan for capacity-building over time. This includes maintaining technical expertise in emerging domains, even when immediate deployment is not prioritized, to avoid a strategic loss of knowledge and capability.

In addition to defensive maturity, policy-makers may consider a complementary “offensive” or opportunity-oriented indicator: the extent to which next-generation computing is actively used to enhance cybersecurity capabilities (e.g. secure distributed architectures, advanced cryptographic toolkits, quantum-enabled secure communication, privacy-preserving computation).

3.5 Illustrative examples

- **Maintaining technical optionality through public testbeds:** Divergent national approaches to quantum key distribution (QKD) illustrate how strategic positioning can shape long-term capability. In contexts where QKD was deprioritized at the policy level, technical knowledge, deployment experience and supply chain familiarity have eroded over time. By contrast, jurisdictions that sustained public-sector testbeds and pilot deployments, even without full commercial endorsement, have preserved technical expertise and reduced the risk of future dependence on third-party infrastructure.
- **Embedding cybersecurity in broader infrastructure initiatives:** The development of next-generation 6G and AI infrastructure in the UK through large, multistakeholder consortia (academia, industry and innovation actors) demonstrates how cybersecurity considerations can be integrated into emerging network architectures. Although the primary objective was advanced communications infrastructure rather than cybersecurity per se, such collaborative platforms provide an opportunity to incorporate resilience testing and secure-by-design principles from the outset. Open, consortium-based infrastructures can complement more closed security processes by enabling broader technical scrutiny and innovation.
- **Strategic coherence across quantum domains:** National quantum strategies that initially treated quantum computing and quantum communications as distinct policy tracks are increasingly recognizing their interdependence, particularly in distributed quantum computing and secure communication architectures. Where communication infrastructure was not actively sustained, subsequent cybersecurity-oriented use cases may face capability gaps. This underscores the importance of avoiding rigid technology labelling and preserving cross-domain expertise to enable future architectural integration.
- **Whole-of-government capability building:** As next-generation computing becomes central to economic and industrial policy, agencies traditionally focused on economic development, innovation funding or infrastructure are being drawn into cybersecurity and resilience discussions. Where such agencies lack in-house technical expertise, resilience-by-design objectives may be unevenly implemented. This dynamic illustrates the need for cross-agency coordination mechanisms, shared technical resources, and accessible guidance tools to ensure that cybersecurity considerations are consistently embedded beyond traditional security portfolios.

3.6 Call to action

As next-generation computing scales, so does the cyber exposure it creates. To ensure organizations harness the opportunity, policy-makers should:

- Develop national and collaborative testbed infrastructures for next-generation computing that integrate multiple dimensions of security, including technology development security, supply chain security, talent resilience and cybersecurity; such environments can support the evaluation of emerging technologies under realistic conditions and reduce reliance on untested assumptions
 - Integrate cybersecurity and resilience criteria at the earliest stages of next-generation compute strategy development, including research funding, infrastructure planning and procurement processes
 - Establish or participate in collaborative test infrastructures to evaluate emerging technologies under realistic adversarial conditions
 - Avoid rigid technology labelling that forecloses future architectural options without empirical evaluation and instead preserve flexibility across emerging compute paradigms
 - Promote whole-of-government engagement, ensuring that agencies beyond traditional security portfolios are equipped with tools, expertise and guidance to address cybersecurity dimensions of compute policy
 - Shift from reactive, patch-based security models towards forward-looking architectures that incorporate long-term protection horizons, forward security principles and structured migration pathways
 - Strengthen capacity-building and shared understanding of emerging computing paradigms across government, enabling informed decision-making and more consistent integration of resilience considerations
 - Encourage alignment across policy domains and jurisdictions, including coordination with allied partners, to manage interdependencies in global compute ecosystems and reduce systemic vulnerabilities
- A coordinated, whole-of-government approach is required to ensure that cybersecurity and resilience considerations are consistently integrated across economic, industrial, research and security policy domains.

Sustainability and cross-council integration

Sustainability is a necessary condition for next-generation computing, while energy and infrastructure determine the deployment capabilities.



Sustainability should consider energy and infrastructure limits as design-stage constraints, not afterthoughts.

4.1 Why sustainability matters for next-generation computing

“Decisions about architectures, system workloads and deployment models have long-term implications for energy demand, resilience and the environment.

Sustainability is not merely an external constraint on next-generation computing strategies; it is a defining condition of their viability. As computing systems scale in power, complexity, and physical footprint, energy availability, environmental impact and infrastructure constraints increasingly shape which forms of computation can be deployed, where and by whom. This is already visible in the rapid expansion of AI data centres and the “race-for-scale” dynamic, which can push sustainability considerations down the priority list unless incentives or requirements change. This is also increasingly relevant for quantum computing: although many current systems still have modest energy footprints, the fault-tolerant systems expected to deliver large-scale industrial value may have far greater power and infrastructure requirements, making sustainability a design-stage concern rather than an issue to address only at commercialization.

Next-generation computing should therefore be understood as part of a coupled system linking compute infrastructure, energy systems, supply chains and public policy. Decisions about architectures, system workloads and deployment models have long-term implications for energy demand, resilience and environmental externalities. Treating sustainability as an afterthought risks locking in brittle and resource-intensive pathways. This also implies a need to distinguish between “useful” and “frivolous” compute demand, even if it is difficult to govern directly, and to improve transparency so that users, communities and regulators can see the consequences of different workload choices. In some paradigms, especially quantum computing, this coupled-system view is essential because the compute stack is intrinsically

hybrid: the quantum processor depends on significant classical resources for circuit generation, compilation, optimization, error correction, orchestration and repeated execution. In other words, the surrounding classical infrastructure will have a significant impact on the computer’s energy consumption, just as the novel processor does.

At the same time, advances in computing enable new approaches to sustainability challenges. High-performance and quantum-enhanced computing can accelerate discovery in energy materials, climate modelling, biological systems and infrastructure optimization. The policy challenge is to manage this dual role: mitigating the sustainability costs of compute while enabling its contribution to broader societal goals. Sustainability assessment should therefore consider system-level trade-offs rather than treating compute as an isolated footprint. In some cases, more capable simulation and modelling may reduce downstream reliance on costly physical testing, iterative laboratory work or inefficient industrial trial-and-error processes. This is particularly relevant in areas such as chemical and materials design, including batteries and fertilizers, where improved digital models could help reduce the time, materials and energy required across the broader development pipeline. In such cases, the relevant policy question is not only how much energy a computing system consumes, but whether it reduces total resource use across the full process it is intended to improve. Sustainability assessments should therefore include counterfactuals: in some domains (e.g. biological discovery), the “non-compute” alternative also has a substantial footprint (materials, consumables, logistics), so policy and procurement should compare net impacts rather than assess compute in isolation.

4.2 Energy, infrastructure and system constraints

Energy use is emerging as a binding constraint across multiple compute paradigms. AI data centres, high-performance computing (HPC) systems, and emerging fault-tolerant quantum computing architectures all place growing demands on power, cooling and physical infrastructure. In quantum systems, for example, energy consumption is distributed across cryogenics, control electronics, classical pre- and post-processing, and error correction, highlighting the importance of a full-stack perspective. Sustainability constraints are multidimensional: water use, materials intensity, siting constraints and governance and regulatory bottlenecks can each become limiting factors depending on the technology and location. For quantum computing specifically, energy demand may arise from three broad sources: 1) control electronics and photonics; 2) cryogenics and vacuum systems; and 3) classical computing needed to generate circuits, compile and optimize programs, manage error correction and process repeated probabilistic runs. This means that, in practice, “quantum” energy consumption comes primarily from classical components.

These dynamics reinforce the need for policy-makers to assess compute strategies not only in terms of performance and capacity, but also in terms of grid integration, infrastructure siting and long-term operational sustainability. Infrastructure choices made today can lock in energy demand profiles for decades. A key practical constraint is the “supply gap” challenge: even when actors call for cleaner power, the availability of scalable clean energy (and the speed of licensing/permitting) can lag behind compute deployment timelines, leading to a default reliance on expedient fossil-based supply unless policy intervenes. Capital formation and risk-sharing also matter: frontier clean-energy technologies are often capital-intensive, and the gap between off-take agreements and early-stage project finance can slow deployment unless incentives, public co-investment or blended finance mechanisms are in place.

For some next-generation compute systems, physical constraints may be as important as electricity supply. Large systems may require specialized buildings, highly stable temperature and vibration conditions, dust control and dedicated

cooling or cryogenic infrastructure that cannot simply be inserted into conventional data-centre environments. As a result, siting decisions should account not only for access to power but also for the location’s physical compatibility with the system’s operating requirements.

Uncertainty is itself a policy-relevant issue. In quantum computing, projected energy and power requirements for future fault-tolerant systems vary dramatically across technical approaches and vendors, with estimates spanning orders of magnitude. This reflects differing assumptions about qubit quality, connectivity, error-correction overhead, system engineering and the maturity of underlying components. Policy-makers should therefore avoid assuming a single energy profile for quantum computing as a category and instead require comparable, architecture-specific disclosure covering both the quantum hardware and all equipment necessary to operate it. Over time, this disclosure architecture should extend to application-based energy reporting – such as the quantity of energy required to solve a given chemistry problem.

From a policy perspective, power requirements also directly shape the feasible scale of deployment. Systems that require very high power levels and specialized infrastructure are likely to be limited to a small number of sites globally, thereby constraining adoption to a narrow set of actors. By contrast, architectures with lower power and infrastructure requirements can be deployed across a much broader range of locations, including existing industrial and public-sector environments. This creates a structural link between energy efficiency, deployability and market size, which should be explicitly considered in national compute strategies.

A further strategic consideration is that next-generation computing may be additive rather than substitutive. Quantum computing, for example, is unlikely to replace existing classical or HPC infrastructure wholesale in the foreseeable future; instead, it will often sit atop existing compute stacks and draw on them. This has important implications for energy planning: new paradigms may increase total compute demand, although they deliver valuable new capabilities.

“ Energy use is emerging as a binding constraint across multiple compute paradigms.

4.3 Sustainability as a cross-cutting governance issue

“ Addressing sustainability in next-generation computing requires coordination across policy domains that extend beyond computing itself.

Sustainability considerations cut across all dimensions of next-generation computing governance:

Sovereignty and interdependence: Access to energy, critical materials, critical data and infrastructure creates new forms of dependency that must be managed strategically. This interdependence extends to compute sovereignty intersecting with energy sovereignty: efforts to increase national independence in compute can materially reshape national energy demand and planning. It also intersects with technology choice: countries may back different compute architectures, not only on scientific grounds but on the basis of energy intensity, infrastructure compatibility and supply-chain resilience.

Cybersecurity and resilience: Energy-intensive and physically sensitive systems increase exposure to disruption, making sustainability and resilience mutually reinforcing issues. Energy-system volatility, local community impacts and physical infrastructure vulnerabilities (cooling, water, grid stress) are pathways through which sustainability risks translate into operational and security risks. In physically delicate systems, resilience also includes the ability to maintain temperature stability, control vibration and manage other environmental conditions necessary for reliable operation.

Innovation and competitiveness: Efficiency, life cycle management and system optimization increasingly influence the total cost of ownership and strategic advantage. Sustainability can be positioned explicitly as a competitiveness lever

(e.g. lower operating costs, faster approvals, stronger social license to operate), helping it travel across ministries and corporate KPIs. For some audiences, especially policy-makers and investors, competitiveness and market viability may be more effective entry points than environmental framing alone. Systems with lower power requirements are likely to be deployable in a larger number of sites and markets, while highly energy-intensive systems may remain confined to a small niche of well-resourced users. In this sense, energy performance is not only a sustainability metric; it can also become a market-shaping metric.

Rather than introducing a standalone sustainability pillar, this framework integrates sustainability considerations into each thematic area of the Global Future Council on Next Generation Computing, ensuring that trade-offs and interconnections are visible to decision-makers. Credible measurement is foundational: self-reporting without standardized definitions can be inconsistent, non-comparable and difficult to trust. Therefore, energy, water and emissions all require auditable, standardized accounting – ideally interoperable across jurisdictions. This is especially important in emerging domains where current systems may have modest footprints, but future commercial systems are still being designed. Standardized methods for measuring energy and power use can reduce uncertainty, improve trust and create pressure for vendors to compete on efficiency and performance. Public procurement, challenge programmes and government-backed calls can accelerate this by requiring energy and power disclosure.

4.4 Cross-council integration

Addressing sustainability in next-generation computing requires coordination across policy domains that extend beyond computing itself. This section draws on insights from other Global Future Councils, including those focused on energy technology systems and generative biology. This includes: 1) connecting compute roadmaps with clean energy scale-up and grid resilience planning; and 2) advancing compute-enabled applications in

generative biology (e.g. bio-based fuels, chemicals and materials) as potential contributors to longer-term energy transition, while recognizing that incentives may otherwise steer innovation towards more immediately lucrative applications (e.g. therapeutics). Cross-ministry “bridging institutions” (e.g. innovation authorities, research-and-innovation umbrella bodies) can align funding and policy across energy, compute and industrial strategy, helping to reduce isolation.

4.5 Illustrative examples

- **Integration of compute planning with energy infrastructure strategy to address growing compute demand:** China and the Middle East are expanding energy capacity alongside compute infrastructure, while in the United States, responses are emerging largely at the state level through data-centre planning and electricity-market interventions. The Nordics demonstrate how an existing renewable-rich energy system can create favourable conditions for compute infrastructure, attracting data-centre deployment in the region.
- **Use of advanced computing to accelerate sustainability-relevant research and public-sector capabilities:** Examples include energy materials discovery, grid optimization, climate and infrastructure modelling, and biological pathways for fuels, chemicals, materials and pharmaceuticals.
- **Governance approaches that balance performance, resilience and energy efficiency over time:** Examples include policy mixes that combine incentives for verified low-impact deployment alongside enforceable minimum reporting and audit requirements and streamlining licensing/permitting for clean energy facilities to close the timing gap between compute demand growth and clean supply scale-up.

5

Scorecard


This scorecard helps leaders assess readiness across five levels, from awareness to anticipatory.

This scorecard supports the next-generation acceleration framework by helping leaders assess their current readiness and to make informed infrastructure, investment and policy decisions.

The framework serves as a reflective guidance to help distinguish between early and long-term readiness across different stages of next-generation computing.

Level 5 Anticipatory Mastery of the current dimension paradigm	Level 4 Strategic Preparation for the next one	Level 3 Developing Real progress under way but strategic value not yet captured	Level 2 Nascent The country or organization knows about the problem	Level 1 Aware Initial awareness
--	--	---	---	---

FIGURE 3 | Next-generation computing readiness scorecard



Scorecard dimensions and core questions

Dimension	Type	Core question
 Dimension 1 Compute and connectivity	Supply	Can we access frontier compute when we need it?
 Dimension 2 Infrastructure (land, power, sustainability)	Supply	Do our physical foundations support current and future paradigms?
 Dimension 3 Data assets and governance	Supply	Can we legally and practically use our information assets?
 Dimension 4 Talent and skills	Supply	Do we have people who can deploy this – and a system to build the next cohort?
 Dimension 5 Sovereignty and partnerships	Supply	Can we access compute without asking permission from a single power?
 Dimension 6 Cybersecurity and resilience	Supply	Are we protected against current threats and preparing for the next ones?
 Dimension 7 Deployment and diffusion	Demand outcomes	Are all of the above actually being converted into economic and social value?



Dimension 1 Compute and connectivity

Can we access frontier compute when we need it?

TABLE 1 Compute and connectivity levels

Level	Descriptor	What this looks like
1	Absent/dependent	– Single foreign cloud provider; no alternatives – One or two subsea cable landing points – Severe urban-rural digital divide – Policy-maker has no visibility of national compute footprint
2	Emerging awareness	– Capacity available, on-premises capacity and multiple cloud providers accessible – Broadband strategy exists but patchy – No domestic compute infrastructure policy – Dependency recognized, unmitigated
3	Developing capacity	– Some domestic compute nodes – Reliable connectivity in major centres – Policy framework exists – Access to current frontier modalities via commercial contracts
4	Strategic deployment	– Diverse domestic and international capacity – Alternative connectivity with multiple peering points – National compute footprint actively mapped and managed – Access to emerging paradigms (cloud quantum, neuromorphic testbeds)
5	Anticipatory	– Significant part of the infrastructure accessible for the next paradigm or 10–100 times the current one – Global tier-one connectivity – Compute capacity used as a forward-looking strategic indicator

Sub-indicators (each scored 1–5, average = dimension score):

1.1 Access breadth – Diversity of compute modalities accessible nationally, such as cloud computing, HPC, quantum as a service (QCaaS) and neuromorphic computing. This assesses whether access is locked to one provider or genuinely diversified across multiple sources.

1.2 Connectivity capacity and resilience – Diversity and capacity of international connectivity: submarine cable landing stations, peering points, last-mile coverage. This also assesses the extent to which single points of failure are avoided.

1.3 Paradigm horizon – Whether the national compute strategy explicitly includes access planning for the next paradigm, not just the current one.



Dimension 2 Infrastructure (land, power, sustainability)

Do our physical foundations support current and future computing paradigms?

Different paradigms have radically different physical requirements: AI requires power density; quantum's cooling needs vary sharply by modality; biocomputing depends on wet-lab environments.

TABLE 2 **Infrastructure levels**

Level	Descriptor	What this looks like
1	Absent	– No compute-specific infrastructure policy – Grid unreliable – No sustainability mandate – Compute infrastructure entirely foreign-owned
2	Nascent	– Basic grid reliability in urban areas – Some domestic facilities (often foreign-owned) – No energy-compute co-planning – Sustainability metrics not yet tracked
3	Developing	– National strategy includes compute – Reliable grid, some renewable capacity – Energy-compute nexus in policy discourse – Carbon/water footprint beginning to be tracked
4	Integrated	– Energy and compute explicitly co-planned – High grid reliability, more than 50% renewable or low-carbon energy sources – Infrastructure roadmap accounts for different physical requirements of emerging paradigms – Metrics improving
5	Strategic	– Infrastructure planning anticipates the next paradigm's physical requirements or 10–100 times the current one – Predominantly renewable- or low-carbon-powered compute zones – Infrastructure is a competitive advantage – Internationally cited as a model

Sub-indicators:

2.1 Energy capacity, reliability and renewability – Grid capacity and reliability score and share of compute infrastructure powered by renewables or low-carbon energy sources.

2.2 Energy-compute co-planning – Whether energy and compute infrastructure are planned

jointly in national/regional/organizational strategy or sit in separate policy lanes.

2.3 Paradigm availability and scalability – Number of specialized national labs or compute facilities, and whether the physical infrastructure roadmap is scalable for requirements of the next paradigm (e.g. 10–100 times the current one, cryogenic for quantum, bio-lab for biocompute).



Dimension 3

Data assets and governance

Can we legally and practically use our information assets?

This includes, for example, training data for AI and annotated genomic databases.

TABLE 3 Data assets and governance levels

Level	Descriptor	What this looks like
1	Fragmented	– Data is departmentally isolated – No national strategy – No legal framework for compute applications – Governance gaps deter development
2	Aware	– Sector-specific data initiatives – Basic data protection law – Some open government data – Compute use of data not yet legally defined
3	Functional	– National data framework with clear legal basis for frontier compute – Cross-ministry sharing protocols – Open datasets in some priority domains – Governance enables development without excessive friction
4	Strategic	– High-quality, accessible national datasets across multiple domains – Genomic databases, industrial sensor networks, materials managed as strategic infrastructure – Governance balances openness and protection
5	Leading	– Information assets explicitly curated as national strategic resources across paradigms – Framework internationally cited as a model – Net exporter of data governance to other nations

Sub-indicators:

3.1 Asset quality and accessibility – Quality, completeness and openness of national datasets in priority domains. Are they usable or just held?

3.2 Legal framework clarity and security – Whether data law enables (not merely permits) use

of data for advanced compute applications without deterring development, and secure sharing.

3.3 Strategic curation – Whether information assets are actively managed as national strategic resources, not just departmental by-products.



Dimension 4

Talent and skills

Do we have people who can deploy this and a system to build the skills of the next cohort?

The focus here is not frontier researchers but practitioner depth. A level 5 country in 2026 may have AI practitioners; in 2031, it may have exceptional breadth but also advanced depth.

TABLE 4 Talent and skills levels

Level	Descriptor	What this looks like
1	Critical shortage	– Negligible domestic expertise – Severe brain drain – No national programme – Almost all expertise imported or unavailable
2	Nascent pipeline	– University programmes exist but are small – Limited industry-academia linkage – Most frontier expertise imported – No retention mechanisms
3	Growing practitioner base	– Sufficient domestic talent to operate (not just procure) current frontier systems – Active retraining – Industry-academia ecosystem developing – Early signs of retention
4	Deployed at scale	– Practitioner base sufficient to deploy across key economic sectors – not just technology – Strong talent flows across industry, academia and government – Strategy linked to deployment goals, not just research excellence
5	Anticipatory	– Systematic national mechanism for identifying skills needed at the next requirements and developing them before it arrives – Net talent attractor – Exports talent frameworks to other countries

Sub-indicators:

4.1 Practitioner depth – The number and sectoral distribution of practitioners capable of deploying (not just researching) current frontier compute tools.

4.2 Pipeline health – Strength and critical mass of education, retraining and retention mechanisms.

Is the country producing more practitioners than it loses?

4.3 Anticipatory mechanism – Whether there is a systematic national process for identifying which skills will be needed at the next paradigm and building them before it arrives.



Dimension 5

Sovereignty and partnerships

Can we access compute and build strategic partnerships without asking permission from a single power?

The dimension doesn't measure ownership, but the depth, diversity and quality of access arrangements. It applies equally to AI, quantum and future paradigms. A Level 5 country shapes the rules that others follow.

TABLE 5 Sovereignty and partnerships levels

Level	Descriptor	What this looks like
1	Dependent	– Compute infrastructure entirely foreign-owned – Fully dependent on a single foreign provider or ally – No alternative access routes – No strategic partnerships – No leverage in the compute supply chain
2	Aware	– Dependency recognized – Initial bilateral conversations begun – No formal access arrangements – Supply chain risk unmapped
3	Active participant	– Member of at least one multilateral compute arrangement – Some alternatives to single-source dependency – Supply chain risks mapped – Contributes to but does not shape international governance
4	Diversified and influential	– Diversified supply chain across multiple providers and allies – Active member of multiple access arrangements – No single-point-of-failure dependency – Contributes substantively to shaping international norms
5	Leading and shaping governance	– Shapes international frameworks for compute access and governance – Exercises strategic autonomy under major-power pressure – Exports its sovereignty architecture as a model – other countries adopt its approach

Sub-indicators:

5.1 Supply chain diversity, capacity and resilience – Number, capacity and geographic diversity of compute access routes. Resilience of supply chains and avoidance of a single point of failure.

5.2 Governance engagement – Level of active participation in multilateral compute governance and agency/company ability to execute strategy

plan/decisions along the full value chain. Not just membership but contributing to shaping norms. Support of at-scale research and infrastructure global/regional projects.

5.3 Strategic intent and ownership – The practical ability to make compute policy decisions and implement strategic decisions without requiring permission from a single major power. National manufacturing capability in strategic compute stack domains like semiconductors.



Dimension 6

Cybersecurity and resilience

Are we protected against current threats and preparing for the next ones?

A cryptographically relevant quantum computer could compromise widely deployed public-key cryptographic schemes, while its implications for symmetric cryptography differ and can be mitigated through appropriate parameters and

migration strategies. AI enables new attack vectors. Biocomputing introduces biosecurity risks. Level 5 is not secured today, but it is systematically anticipating tomorrow's threat landscape.

TABLE 6 Cybersecurity and resilience levels

Level	Descriptor	What this looks like
1	Unprotected	– No compute-specific cyber strategy – Infrastructure unaudited – Classical cryptography only, no post-quantum roadmap – No incident response – No awareness of compute-enabled emerging threats
2	Basic protection	– National cybersecurity law exists – Some incident response – Current-paradigm threats partially addressed – No forward-looking posture
3	Compute-aware	– National cyber strategy covers compute infrastructure – Supply chain audits under way – Post-quantum cryptography on the roadmap – Emerging threat landscape understood at policy level but not implemented at scale
4	Hardened and forward-looking	– Active monitoring and hardening – Supply chain verified – hardware and software – Post-quantum deployed in government and critical systems – Physical resilience tested – Emerging threat vectors actively addressed
5	Anticipatory	– Systematic process for identifying threats from the next paradigm before they materialize – Post-quantum fully deployed across all critical systems – Hardening against agentic AI threats – Physical, logical and supply chain resilience regularly stress-tested – Contributes to international security anticipation and standards

Sub-indicators:

6.1 Current threat detection and posture –

Detection and protection against threats from the current compute paradigm: AI-enabled attacks, supply chain compromise, physical infrastructure vulnerability. Weaponization of platforms and repurpose.

6.2 Cryptographic transition and quantum-

safe readiness – Concrete progress in identifying

quantum-vulnerable cryptographic assets and migrating government and critical systems towards post-quantum standards. This includes crypto-agility, implementation timelines, interoperability testing and protection against “harvest now, decrypt later” risks.

6.3 Next-paradigm anticipation – Whether there is a systematic process for identifying and mitigating threats from the next compute paradigm before they materialize and ensure the continuity of capacity.



Dimension 7

Deployment and diffusion

Are all the other dimensions actually being converted into economic and social value?

The demand-side counterpart to all other dimensions. The 30-year diffusion lag between technological feasibility and economy-wide

adoption is a pattern observed across steam, electricity, computing and AI. It is the universal signal that policy is (or is not) working.

TABLE 7 Deployment and diffusion levels

Level	Descriptor	What this looks like
1	Negligible	– Fewer than 5% of firms using frontier compute tools systematically – Government pilots only – No diffusion strategy – No measurement framework
2	Early clusters	– Visible adoption in technology and a selected vertical only – Government aware but no active strategy – Adoption not systematically measured
3	Sectoral deployment	– Deployment across multiple sectors – Government measuring and encouraging uptake – Productivity signal emerging – Firms moving from pilots to production – Adoption rate tracked as a policy metric
4	Economy-wide diffusion	– Over 40% of firms actively deploying – Sector playbooks published – Measurable productivity uplift at economy level – Public sector is a leading deployer – Adoption rate used as a leading economic indicator
5	Compute-native economy	– Frontier compute deployment is the default – Organizations systematically update practices as each new paradigm emerges – Deployment velocity is institutionalized – and is itself a signal that attracts further infrastructure investment

Sub-indicators:

7.1 Enterprise advanced compute adoption rate

– Share of firms actively deploying frontier compute tools in production and services – beyond pilots, beyond the technology sector. Assessments should define the problem, dataset, hardware configuration, error-mitigation methods, total classical and quantum resources and the best available classical comparator. Independent validation should be encouraged, particularly when results inform public investment, procurement or national strategy.

7.2 Cross-sector breadth – Number of distinct economic sectors with meaningful frontier compute deployment. Deployment beyond finance and tech is a positive sign.

7.3 Public sector leadership – Whether government itself is a leading deployer, not just a regulator, by measuring adoption (tracking of new applications) and using next-generation computing in its own operations.

Scoring and diagnostic logic

Score each of the 21 sub-indicators from one to five. Average the three within each dimension to get the dimension score. The total of the seven dimension scores produces the radar chart.

Interpretation:

Levels 1 and 2: Nascent/aware – the country knows the problem exists.

Level 3: Developing – real progress under way but strategic value not yet captured.

Levels 4 and 5: Strategic/anticipatory (4 = mastery of the current paradigm; 5 = preparation for the next one)

Binding constraint principle: A low score on any single dimension can act as a ceiling on all others. The lowest-scoring supply dimension is almost always the real bottleneck.

Conclusion

Next-generation computing is not a standalone technology race but a systemic opportunity. The acceleration framework and scorecard span over six supply enablers that can be translated into real economic and social values as demand outcomes. This helps decision-makers clearly reflect on where their true opportunities and chokepoints lie, and whether these capabilities have been reflected in actual outcomes.

Both the acceleration framework and scorecard serve as starting points and are not built for ranking or benchmarking, but to prompt honest internal dialogue and coordinated judgement

within governments and organizations across the technology, industrial, energy, education and security fields that all shape next-generation computing outcomes. These decisions require collaboration across energy planning, workforce and talent policy, procurement and cybersecurity frameworks.

Within this investment and transformation window, what determines success is whether a country or organization can identify their own strengths and binding constraints, and coordinate across domains to absorb and convert potential into actual strategic value.

Contributors

World Economic Forum

Filipe Beato

Manager, Technology and Innovation,
Centre for Frontier Technologies and Innovation

Arunima Sarkar

Head, Initiatives, C4IR Network and Partner
Engagement

Sophie Xiaoran Tang

Community Engagement Specialist,
Centre for Frontier Technologies and Innovation

Global Future Council on Next Generation Computing

The Forum would like to express its gratitude to the members of the Global Future Council on Next Generation Computing, who, through their work and discussions, have been an inspiration and guided many of the concepts presented in the paper.

Najwa Aaraj

Chief Executive Officer, Technology Innovation
Institute, Sole Proprietorship

Azeem Azhar

Chief Executive Officer, Exponential View

Giuseppe Borghi

Head, Φ -lab Division, European Space Agency (ESA)

Kirk Bresniker

Chief Architect, Hewlett Packard Labs; Fellow
and Vice-President, Hewlett Packard Enterprise

Nadia Carlsten

Member of the Board of Advisors,
Pacific Northwest Laboratory

Jennifer Chayes

Dean of the College of Computing,
Data Science and Society, UC Berkeley

Alessandro Curioni

Vice-President, Europe and Africa; Director,
IBM Research Europe

Jack Hidary

Chief Executive Officer, SandboxAQ

Mihai Adrian Ionescu

Dean of Engineering, Ecole Polytechnique Fédérale
de Lausanne (EPFL)

Elham Kashefi

Personal Chair in Quantum Computing,
University of Edinburgh

Dhiresha Kudithipudi

Director, MATRIX AI Consortium,
University of Texas at San Antonio (UTSA)

Satoshi Matsuoka

Director, RIKEN Center for Computational Science,
RIKEN

Michele Mosca

Co-Founder, President and Chief Executive Officer,
evolutionQ

Khalifa Al Qama

Director, Dubai Future Labs,
Dubai Future Foundation

Robert Speight

Director, CSIRO Advanced Engineering Biology
Future Science Platform, Commonwealth Scientific
and Industrial Research Organisation (CSIRO)

Kathryn Timmerman

Chief Executive Officer,
Chicago Quantum Exchange

Florian Wohlrab

Chief Executive Officer, OpenHW Foundation

Reena Dayal Yadav

Founder and Chief Executive Officer,
Quantum Ecosystems and Technology
Council of India

Acknowledgements

The World Economic Forum would like to thank the following individuals for their insightful reviews and feedback. The Forum also wishes to acknowledge the following contributions:

Olivier Ezratty

Co-Founder, Quantum Energy Initiative

Camille Georges

Deputy Director, EDHEC Quantum Institute (EQI)

María Fernanda González

Chief Executive Officer, PlanetAI Nature Space

Travis Humble

Director, Quantum Science Center, Oak Ridge National Laboratory, Institute of Electrical and Electronics Engineers (IEEE)

Michal Rosen-Zvi

Professor of Computational Medicine, Hebrew University of Jerusalem

Jennifer Schenker

Founder and Editor-in-Chief, The Innovator

Elvira Shishenina

Senior Director of Strategic Initiatives, Quantinuum

Rafael Sotelo

Dean of Engineering, Universidad de Montevideo

Elena Yndurain

Professor, IE University

Production

Laurence Denmark

Creative Director, Studio Miko

Charlotte Ivany

Designer, Studio Miko

Will Liley

Editor, Studio Miko

World Economic Forum Public Engagement

Maxwell Hall

Creative Editorial Lead

Floris Landi

Design Lead

Gayle Markovitz

Head, Written and Audio Content

Sybille Penhirin

Head, Video and Design

Endnotes

1. Azhar, A., Gildea, W., Petrovic, H., Warren, N., et al. (2026). *The State of the AI Economy*. Exponential View. <https://intelligence.exponentialview.co/>.
2. Thompson, E. J. (2026). *BofA Predicts Cloud and AI Infrastructure Spending to Reach \$1.5 Trillion by 2027*. <https://intellectia.ai/news/stock/bofa-predicts-cloud-and-ai-infrastructure-spending-to-reach-15-trillion-by-2027>.
3. European High-Performance Computing Joint Undertaking (EuroHPC JU). (n.d.). *Discover EuroHPC*. https://www.eurohpc-ju.europa.eu/about/discover-eurohpc-ju_en.
4. UK Department for Science, Innovation and Technology. (2026). UK Compute Roadmap. <https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap>.
5. National Institute of Standards and Technology (NIST). (2026). *Department of Commerce Announces Letters of Intent With 9 Companies for \$2 Billion to Accelerate U.S. Leadership in Quantum Computing*. <https://www.nist.gov/news-events/news/2026/05/department-commerce-announces-letters-intent-9-companies-2-billion>.
6. Prime Minister's Office Singapore. (2024). *DPM Heng Swee Keat at the Asia Tech x Singapore 2024 Opening Ceremony*. <https://www.pmo.gov.sg/newsroom/dpm-heng-swee-keat-at-the-asia-tech-x-singapore-2024-opening-ceremony/>.
7. Government of Canada. (2025). *Building a stronger Canadian economy*. <https://budget.canada.ca/2025/report-rapport/chap1-en.html>.
8. Azhar, A., Gildea, W., Petrovic, H., Warren, N., et al. (2026). *The State of the AI Economy*. Exponential View. <https://intelligence.exponentialview.co/>.
9. Thompson, E. J. (2026). *BofA Predicts Cloud and AI Infrastructure Spending to Reach \$1.5 Trillion by 2027*. <https://intellectia.ai/news/stock/bofa-predicts-cloud-and-ai-infrastructure-spending-to-reach-15-trillion-by-2027>.
10. Organisation for Economic Co-operation and Development (OECD). (2021). *Development Co-operation Report 2021*. https://www.oecd.org/en/publications/development-co-operation-report-2021_ce08832f-en/full-report/component-41.html.
11. Nguyen, D. (2025). *A Third Path for AI Beyond the US-China Binary*. Noema Magazine. <https://www.noemamag.com/a-third-path-for-ai-beyond-the-us-china-binary/>.
12. Sigalos, M. (2025). *Amazon opens \$11 billion AI data center in rural Indiana as rivals race to break ground*. CNBC. <https://www.cnbc.com/2025/10/29/amazon-opens-11-billion-ai-data-center-project-rainier-in-indiana.html>.
13. Azhar, A., Gildea, W., Petrovic, H., Warren, N., et al. (2026). *The State of the AI Economy*. Exponential View. <https://intelligence.exponentialview.co/>.
14. Sternstein, M. (2026). *Charts of the Week: Experience > Skills*. A16z. <https://www.a16z.news/p/chart-of-the-week-experience-skills>.
15. Ross, R., Pillitteri, V., Graubart, R., Bodeau, D., et al. (2026). *NIST SP 800-160 Vol.2: Developing Cyber Resilient Systems: A Systems Security Engineering Approach*. National Institute of Standards and Technology (NIST). <https://www.nist.gov/news-events/news/2026/05/department-commerce-announces-letters-intent-9-companies-2-billion>.
16. World Economic Forum. (2022). *Transitioning to a Quantum-Secure Economy*. <https://www.weforum.org/publications/transitioning-to-a-quantum-secure-economy/>.
17. Organisation for Economic Co-operation and Development (OECD). (n.d.). *Emerging technologies*. <https://www.oecd.org/en/topics/emerging-technologies.html>.
18. Boyens, J., Smith, A., Bartol, N., Winkler, K., et al. (2022). *NIST SP 800-161 Rev.1: Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations*. National Institute of Standards and Technology (NIST). <https://csrc.nist.gov/pubs/sp/800/161/r1/upd1/final>.
19. Ibid.
20. Manning, C. G. (2023). *Technology Readiness Levels*. <https://www.nasa.gov/directorates/somd/space-communications-navigation-program/technology-readiness-levels/>.



COMMITTED TO
IMPROVING THE STATE
OF THE WORLD

The World Economic Forum is the trusted global platform for public-private cooperation, bringing together leaders from business, government, civil society, international organizations, academia and the next generation to turn foresight and dialogue into meaningful action on the world's most pressing challenges.

World Economic Forum
91–93 route de la Capite
CH-1223 Cologny/Geneva
Switzerland

Tel.: +41 (0) 22 869 1212
Fax: +41 (0) 22 786 2744
contact@weforum.org
www.weforum.org