ALLIANCE FOR INNOVATION AND INFRASTRUCTURE

THE SOVEREIGNTY GAP IN COMPUTE

A seven-dimension framework for classifying what "sovereign" actually means at the infrastructure level

Why the market's missing vocabulary is a commercialization failure

APRIL 2026

Mothusi Pahl

ALLIANCE FOR INNOVATION AND INFRASTRUCTURE

The Sovereignty Gap in Compute

A seven-dimension framework for classifying what "sovereign" actually means at the infrastructure level. Why the market's missing vocabulary is a commercialization failure.

Executive Summary

The sovereign compute market is allocating hundreds of billions of dollars against a word that has no agreed-upon definition.[1] The foundation of that market is American infrastructure: power generation, grid capacity, transformer supply, and the supply chain dependencies that follow. So is most of the policy architecture shaping how the market operates: export controls, the AI exports framework, and the bilateral AI partnerships now treated as diplomatic instruments. When allied defense establishments and partner governments ask sovereignty questions U.S. providers cannot answer, that is an American innovation and infrastructure problem.

This piece introduces SAAFE-7[2], a framework for evaluating what "sovereign compute" actually means at the infrastructure level. The framework names seven dimensions that compound from the energy foundation through the engineering, industrial, legal, and human layers of the compute stack: power, data security, air-gapping, physical isolation, supply chain integrity, jurisdiction-alignment, and trust. Each dimension is a distinct dependency that a buyer either controls or does not. A sovereignty claim is testable only when the dimensions are named.

Applying the framework leads to the conclusion that most of what the market currently sells as "sovereign" is not. Providers using sovereignty as a marketing term routinely deliver on data security and physical isolation while their legal and industrial structure cannot support jurisdiction-alignment or supply chain integrity. Buyers purchasing under that ambiguity do not know which dimensions they have purchased and which they have only been told they have purchased. The gap between what is marketed and what is delivered is the sovereignty gap.

The piece is written for the capital allocators, infrastructure investors, defense procurement principals, and policy professionals operating in compute and AI infrastructure networks. The framework advanced in this paper is proposed as an open standard with auditable criteria at each dimension. The absence of an existing framework is leaving real money on the table for buyers, providers and investors alike, and is leaving American infrastructure investment, U.S. provider strategy and U.S. alliance policy operating against a market category that no one has defined.

Introduction

"Sovereign" appears in investor decks, government procurement documents and bilateral agreements between nations. In none of these contexts does it mean the same thing. McKinsey, Bain, the World Economic Forum, CSIS, Accenture, the Tony Blair Institute and NVIDIA have all published extensively on sovereign compute in the last two quarters. Their work is substantive. None of it is wrong. But none of it answers the question a buyer walks in with: Which dimensions of sovereignty does my threat model require and which am I actually purchasing?

This is not an analytical failure. It is a commercialization failure. And it is the kind of gap that leaves real money on the table for buyers, providers and investors alike.

The pattern emerged in Q4 2025 as I moved between meetings in London, Brussels and Washington. Every datacenter conversation started with power: supply, grid interconnection timelines, generation capacity, delivery infrastructure, and transformers, but every conversation ended somewhere else entirely. The power questions were real, but they were the tip of an iceberg. Beneath them sat a tangle of dependencies that my counterparts, many of them at compute companies and national security establishments, had started calling "sovereignty." The deeper I went, the clearer it became that the power question was leading to harder questions about control, jurisdiction and trust that nobody in the market had named.

To put it in Clayton Christensen's words: you need to understand the progress a customer is trying to make in a given circumstance.[3] When the market has no shared vocabulary for that progress, buyers overpay for dimensions they don't need and under-purchase on the dimensions they do need. Providers end up competing on the wrong features because they cannot clearly see the job to be done. This is precisely how hundreds of billions get allocated against a word nobody has defined.

* * *

The word sovereign has become the most expensive undefined term in infrastructure.

The compute market treats "sovereignty" as binary. You either have it or you don't. In practice, sovereignty is a spectrum with at least seven distinct dimensions. Glossing over these complexities and pushing a single marketing term has created a dangerous gap between what customers believe they are purchasing and what they are actually getting. A buyer who asks for "sovereign compute" might mean encrypted data at rest with customer-managed keys. They might mean an air-gapped facility with no internet connectivity. They might mean legal immunity from a foreign government's court orders. They might also mean all three simultaneously, plus physical isolation and independent power generation. These are fundamentally different requirements with different cost structures and the market currently has no shared vocabulary to distinguish them.

The consequences are not hypothetical. I have spoken with European industry leaders who believe they are purchasing shared, sovereign AI infrastructure when they are in fact purchasing U.S.-controlled AI infrastructure with shared access. The delta between what those senior leaders believe they are getting and what they are actually getting is what I call the **sovereignty gap**. Some of these gaps are marketed: the provider's branding implies coverage at a given dimension but the architecture structurally cannot deliver it. Others are unnamed: the gap exists but nobody, provider or buyer, has the vocabulary to properly identify it. This framework exists to convert unnamed gaps into named ones.

* * *

What the Market Has Published

A closer look at what the market has published makes the gap more specific.

In Q4 2025, CSIS published a paper arguing that sovereign clouds risk becoming “splinter clouds” with overcapacity.[4] McKinsey published “The Sovereign AI Agenda” with four stakeholder archetypes.[5] Accenture published a Sovereign Maturity Index, perhaps the closest to

actionable usefulness.[6] In January 2026, the WEF and Bain published a framework that segmented countries into five economy archetypes based on "posture" toward sovereign compute and framed sovereignty as "strategic interdependence."[7] Also in January, the Tony Blair Institute published seven strategic levers for governments.[8] NVIDIA, creator of the "sovereign AI" label, has published a series of country-level playbooks over the same period.[9]

All of this is useful. But none of it answers the question a knowledgeable buyer walks in with: "I need sovereign compute. Which dimensions of sovereignty does my threat model require, and which am I actually purchasing?" When I went looking for something that could answer that question, I could not find any classification framework for what sovereignty meant in operational terms from a customer perspective.

Are we really in a market that is allocating hundreds of billions of dollars and using a word that has no agreed-upon definition?

* * *

A Seven-Dimension Framework for Compute Sovereignty

The sovereign compute ecosystem is following a pattern I have watched play out many times in energy and infrastructure: brilliant technical founders and experienced boards over-index on technology and miss the commercialization problem. In this case, the problem is that the market has no shared vocabulary for what buyers are actually trying to purchase.

I am proposing the following framework as an open standard starting point for shared vocabulary with auditable criteria at each dimension. I am not claiming that this is complete. Seven dimensions may become nine or five once sharper minds pressure-test the boundaries. What I am claiming is that the market needs a framework of this kind and that the absence of one is leaving real money on the table for buyers, providers and investors alike.

The framework has seven dimensions. The first four are engineering problems. The fifth is an industrial problem. The sixth is a legal problem. The last is a human problem. Power is always the foundation. The remaining six shift in priority depending on the buyer's specific requirements, which is precisely why the market needs shared vocabulary to describe them.

* * *

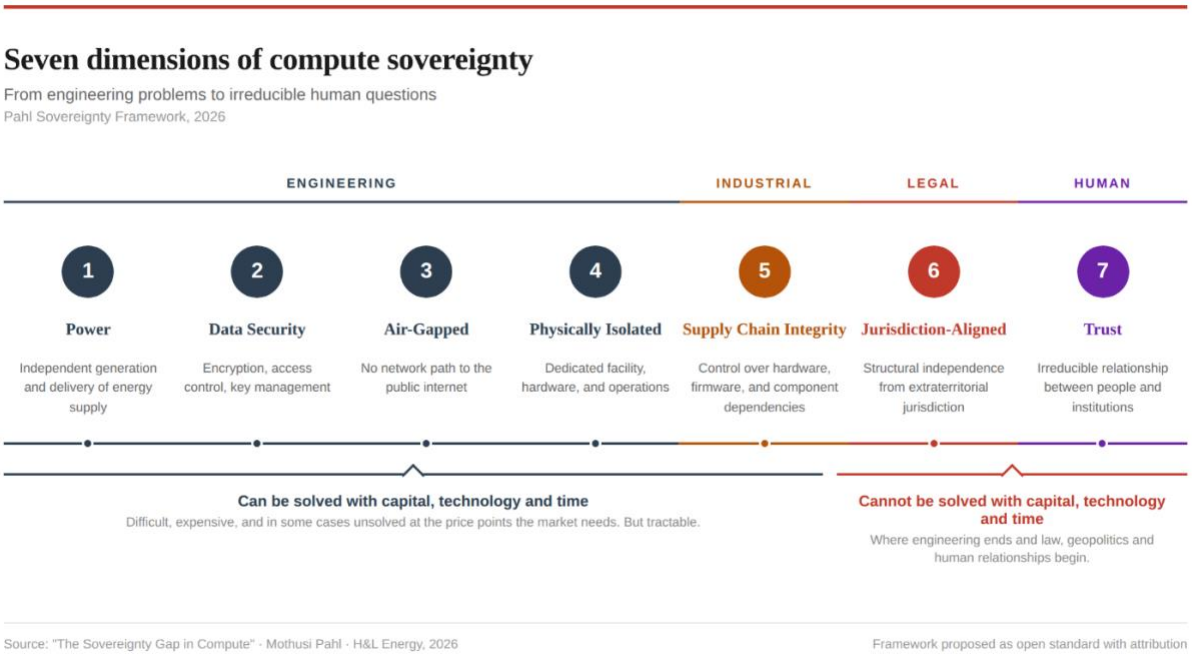

Dimension 1: Power

Independent generation and delivery of energy supply.

Something is not sovereign if it depends on others for the generation or delivery of power. If the grid connection can be interrupted by a third party's political, commercial or operational decisions, then the compute sovereignty has a single point of failure before you ever get to the data layer.

This is the dimension I know best and it is the one the market discusses least in the context of sovereignty. Power is treated as an input to the sovereign compute stack, not as a dependency within it. This framing is a mistake. The power supply chain is where sovereignty either starts or fails. For any customer whose threat model includes infrastructure disruption by a state-level adversary, it is the first place to look. A sovereign compute facility that draws power from a grid, a wellhead, or a solar cell factory it does not control is sovereign only until someone else decides it isn't. (See Nord Stream 2022.)[10]

For buyers whose grid interconnection timelines stretch four to seven years[11], there is a faster path: dedicated power supply with secured feedstock, deployed independent of traditional grid infrastructure and its associated permitting queues. This is not a novelty. The oil and gas industry, remote mining operations and military forward-deployed computing have been solving these exact constraints for decades. The operational playbook exists. What does not yet exist is a compute infrastructure market that applies it.

Dimension 2: Data Security

Encryption, access control, key management.

These are table stakes and the industry does them well. Most "sovereign cloud" marketing stops roughly here: data encrypted in transit and at rest, customer-managed encryption keys, access policies controlled by the customer. The problem is not that this dimension is poorly served. The problem is that many buyers believe purchasing this dimension means they have achieved sovereignty when what they have really achieved is security, a necessary but insufficient component of a much deeper stack.

Dimension 3: Air-Gapped

No network path to the public internet.

A full cloud stack deployed inside a customer's facility with no external connectivity. This is meaningfully different from "secure," but the market often confuses the two. A secure environment can be connected to the internet with strong access controls. An air-gapped environment has no network connectivity to the outside world. The operational implications, the cost structure and the threat model are fundamentally different. This is not just a stronger version of secure.

An important caveat: air gaps leak. Firmware-level implants, electromagnetic tracking, supply chain compromises, insider threats and facility breaches are a few examples. From well-documented incidents in national security contexts to ongoing red team findings, "no network connectivity" does not mean "no data exfiltration."[12] That assumption never survives contact with a competent and motivated adversary. This does not mean air gaps are useless. It does mean that buyers should understand exactly what air gaps do and do not guarantee, and providers should be explicit about where the boundary of their air gap security commitment lies. As Dimension 5 illustrates in more extreme terms, what enters the air gap matters as much as what crosses it.

Dimension 4: Physically Isolated

Dedicated facility, hardware, and operations.

Physical isolation means the building, power and cooling plant, hardware, logistics chain and operational staff are all dedicated to a single customer with nothing shared with any other tenants (sovereign or otherwise). This is distinct from air-gapping, which is a network property. An air-gap can sit inside a multi-tenant facility; a physically isolated facility can be connected to the internet. The threat model is also distinct: physical isolation addresses facility-level compromise, side-channel exposure from adjacent tenants, insider risk in shared operations teams and hardware tampering through shared logistics, none of which an air gap by itself resolves.

At roughly 5x the cost of a typical shared sovereign instance (unless you are sitting on a wellhead), only the largest nations and most sensitive workloads can justify this today.

Which creates a specific and underserved market segment that will define the next wave of sovereign infrastructure investment: the **middle-power sovereign buyer**.

This is a nation or institution that genuinely needs physical isolation because its threat model demands it and cannot afford what physical isolation costs today. The segment is large, it is growing, and nobody is building infrastructure specifically designed to serve it. Gulf states alone have committed over $2.5 trillion to AI infrastructure, with procurement requirements that go far beyond data residency.[13] ASEAN countries are establishing AI infrastructure independence as a matter of national strategy.[14] Allied defense establishments are explicitly asking how to train models without CLOUD Act exposure. And a significant number of European nations have regulatory frameworks that demand sovereignty their current providers cannot deliver.[15]

Geography	**2025**	**2026**	**2027**
China (Region)	37,539	47,379	58,544
North America	12,667	16,394	21,127
Europe	6,868	12,587	23,118
Mature Asia/Pacific	851	1,593	3,155
Japan Region	519	932	1,816
Emerging Asia/Pacific	430	755	1,326
Latin America	278	506	946
Middle East and North Africa	132	250	515
Sub-Saharan Africa	16	31	61
Total	**59,300**	**80,427**	**110,609**

Source: Gartner (February 2026). Sovereign Cloud IaaS Spending by Region. (Million USD)[1]

In commercialization terms, the middle-power sovereign buyer is the customer segment where money will flow in the next 24 months. These buyers are entering the market with requirements that the current vocabulary cannot describe and budgets that the current cost structure cannot serve. Any provider that can deliver physical isolation at a price point between shared sovereign cloud and bespoke national infrastructure will find a market that is both enormous and almost entirely uncontested. The framework I am proposing here exists in part to make this segment visible, because the current binary treatment of sovereignty hides it entirely.

* * *

These first four dimensions are engineering problems. They are difficult, expensive, and in some cases unsolved at the price points the market needs. But they are tractable. Given enough capital and enough time, they can all be built.

* * *

Dimension 5: Supply Chain Integrity

Control over hardware, firmware, and component dependencies.

This dimension addresses two distinct failure modes. The first is denial: someone controls your supply chain and restricts access. Export controls, entity list designations and the concentration of advanced semiconductor fabrication at TSMC (which produces over 50 percent of the world's advanced chips) create structural dependencies that no amount of engineering at the facility level can resolve.[16] The chip export license negotiations currently shaping relationships between the U.S. and Gulf states, including the G42 and HUMAIN arrangements, are evidence that this dimension is already defining market structure.[17]

The second failure mode is infiltration: someone compromises hardware within your supply chain before delivery. The 2024 pager supply chain attack illustrates the principle at its most extreme: hardware that functions as expected until it doesn't, weaponized before the buyer ever takes possession.[18] The same vector applies at the silicon level, where firmware-level implants can be introduced during fabrication, packaging or transit. An air gap means nothing if the hardware inside it has been compromised before arrival.

Virtually all large-scale AI training today depends on NVIDIA architectures, creating a single-vendor sovereignty risk that cuts across every other dimension.[19] A buyer who achieves jurisdiction-alignment, physical isolation and independent power generation still depends on the supply chain that delivers the silicon. This dependency compounds the jurisdictional problem of Dimension 6: even full legal compliance can be undermined at the chip level, and even the most carefully constructed trust relationships of Dimension 7 rest on confidence in the integrity of the underlying hardware.

* * *

The first five dimensions share a common property: they can be addressed with capital, technology and time. The next two cannot.

* * *

Dimension 6: Jurisdiction-Aligned

Structural independence from extraterritorial jurisdiction.

This is where engineering ends and law begins.

If a U.S.-headquartered company has any administrative pathway to data (possession, custody or control) and a U.S. court issues a valid order under the CLOUD Act, FISA Section 702, or a National Security Letter[20], that company is obligated to comply, regardless of where the data is

physically stored. This is not a criticism of any company or country. It is a statement about the structure of U.S. law, with the knowledge that equivalent structures exist in other jurisdictions. But the U.S. case is the most consequential for the sovereign compute market because most compute infrastructure depends on U.S. technology and the largest cloud and neocloud providers are U.S.-headquartered.

I have reviewed sovereign cloud marketing materials from every major provider and have yet to find one that clearly states: "Our air-gapped offering does not provide immunity from the legal jurisdiction of our home country." They discuss encryption, key management, facility security, operational controls. They do not discuss what happens when a court order arrives. The customer who believes they have achieved sovereignty because they hold the encryption keys is only looking skin-deep.

Structural control goes deeper than any single statute. Most compute infrastructure depends on U.S. technology at multiple levels: chips, firmware, orchestration software, development tools, model architectures. If the U.S. government decides to restrict access to any underlying technology through export controls, entity list designations or pressure on licensing partners, most compute infrastructure companies lose their technology supply chain. This is not a theoretical risk. The Huawei situation demonstrated precisely how this leverage operates when a technology relationship becomes a national security question.[21] Supply chain integrity, as described in Dimension 5, compounds the jurisdictional problem: the same silicon dependencies that create industrial risk also create legal exposure.

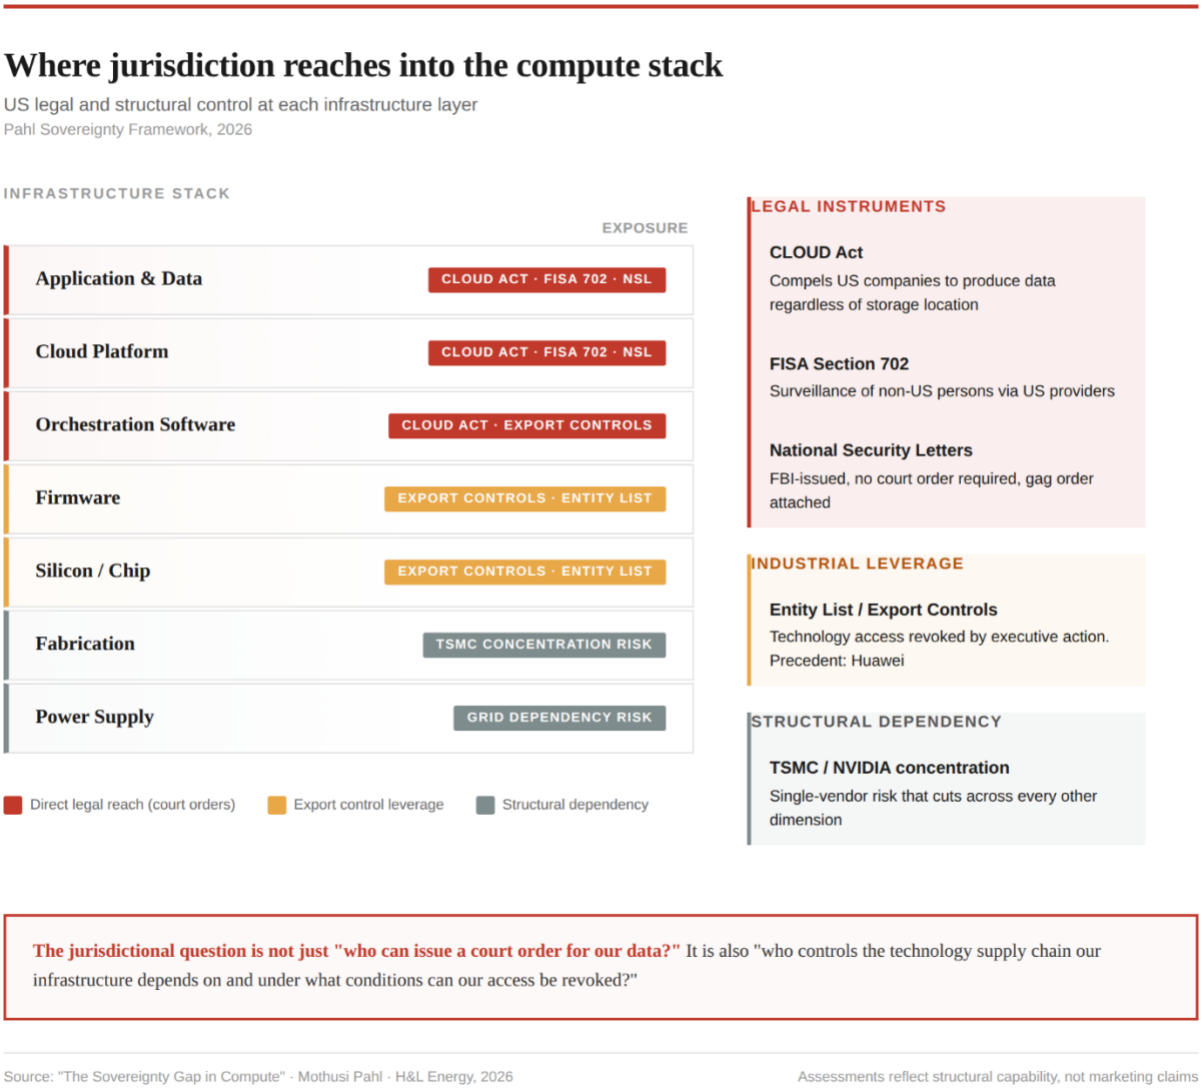

For non-U.S. buyers evaluating sovereignty, the jurisdictional question is therefore not just “who can issue a court order for our data?” It is also “who controls the technology supply chain our infrastructure depends on, and under what conditions can our access be revoked?”

The emerging structural alternative that several European and Gulf buyers are actively exploring is an architecture built around operating entities incorporated in non-Five Eyes jurisdictions with no U.S.-person involvement in data handling chains. This means non-U.S.-headquartered operating companies, non-U.S. technology dependencies at each stack layer where feasible and bilateral legal frameworks that provide enforceable jurisdictional boundaries.[22] I will not name specific procurement conversations here, but the pattern is consistent: the most sophisticated buyers are not waiting for the legal landscape to clarify. They are building around it. Any provider that understands this architecture and can deliver against it will find buyers who are ready to commit capital today.

Dimension 7: Trust

Irreducible relationship between people and institutions.

The final dimension is neither technical nor legal. It is personal. There is no place on earth truly free from all legal authority; even international waters are governed by UNCLOS, flag-state law and the law of the high seas.[23] True isolation means controlling everything, from potential energy through application and data output. Nobody does that today.

At the outer boundary of the sovereignty spectrum, what remains is a relationship between individual people and institutions. This trust is built over time and grounded in reliability. Trust cannot be legislated into existence or driven by executive order. It is what remains after every technical and legal control has been applied, and it is irreducible: even a buyer who achieves jurisdiction-alignment across every other dimension still depends on trust relationships with the entities that design, fabricate, package and deliver the infrastructure their operations run on. Any sovereignty framework that does not acknowledge this final human layer is incomplete.

* * *

The Buyers Nobody Is Building For

This framework reveals at least four categories of sovereign compute buyers that existing providers do not adequately serve. First, customers who need genuinely air-gapped environments with honest assessments of what air gaps can and cannot guarantee. Second, customers who need workloads on non-U.S. soil with non-U.S. operational control. Third, customers who need compute capacity but cannot get grid connections in the timelines they need. Fourth, customers who are simply capacity-constrained: the compute is not available at any dimension, at any price, in their geography.

These categories are not independent; they are entangled. A customer who needs jurisdiction-alignment almost certainly also needs physical isolation. A customer with grid-connection constraints may also need air-gapped operations. A capacity-constrained buyer who resolves the capacity problem may immediately discover they now have a jurisdiction problem.

The real question is one the market has not yet asked: what job is the buyer hiring sovereign compute to do? In commercialization terms, this is a *Jobs to Be Done* problem. A buyer who needs jurisdiction-alignment is not doing the same job as a buyer who needs air-gapped operations, even if both use the word “sovereign” to describe their requirement. These dimensions make jobs to be done visible. Without them, the mismatch I described above will keep growing.

The middle-power sovereign buyer sits at the center of this landscape. These are the Gulf states whose $2.5 trillion in combined AI commitments require sovereign infrastructure that goes far beyond data residency. These are the ASEAN countries establishing AI infrastructure independence as national strategy. These are the allied defense establishments explicitly asking how to train models outside CLOUD Act exposure. These are the European national AI programs whose regulatory frameworks increasingly demand what their current providers structurally cannot deliver. This segment is entering the market with capital and urgency, and no provider is competing for their business on the dimensions that matter most.

Neoclouds like Crusoe, Lambda and CoreWeave compete vigorously on performance, price and availability.[24] But none compete on the legal dimensions that engineering cannot resolve. The structural legal and jurisdictional questions of Dimensions 6 and 7 remain unaddressed in neocloud market positioning. They are U.S.-headquartered companies subject to U.S. law, which means the customers whose primary constraint is jurisdictional have no neocloud provider competing for their business on the dimensions that define their purchasing decision.

* * *

An Open Standard, Not an Invitation

This framework, while far from complete, is the first attempt to build it from the buyer's perspective. Without shared vocabulary, the sovereign compute market will continue to operate under a foundation built from a single, undefined, poorly understood word. The cost of that imprecision falls on everyone: buyers who purchase the wrong thing, providers who create the wrong product, and investors who do not fully understand the nuance and complexity of the industry.

This framework is proposed as an open standard with auditable criteria at each dimension, so that the conversation can move from marketing language to operational specificity. The engineering dimensions need validation from practitioners who have built and operated sovereign infrastructure in the environments where it matters. The legal dimension needs collaboration from attorneys and policy experts who work at the intersection of data sovereignty, extraterritorial jurisdiction and national security law.

What I have built here is a framework that makes an $80-billion-and-growing market[1] legible to the buyers who are trying to navigate it and the providers who are trying to serve it. The seven dimensions give everyone a shared vocabulary that the market currently lacks. For buyers, the framework is a diagnostic: which dimensions does your threat model require, and which are you actually purchasing? For providers, it is a positioning tool: which dimensions does your offering deliver and which does your architecture structurally prevent you from delivering? For investors,

it is a due diligence filter: does the infrastructure asset you are evaluating actually serve the sovereignty segment it claims to address?

But the framework is also a starting point for harder questions. What gets me power sovereignty faster? Do I need air-gapped, or is secure sufficient for my threat model? Which supply chain dependencies can I diversify in 12 months versus 36? Can I pursue jurisdiction-alignment in parallel with physical isolation, or are they sequential? Where might I locate to collapse timelines? These are the questions that convert a diagnostic framework into a deployment strategy, and they are the questions this market will spend the next several years answering.

If you work in sovereign infrastructure, defense procurement, national AI strategy or data jurisdiction law, I welcome your perspective on where this framework is wrong, where it is incomplete and what it takes to turn shared vocabulary into something buyers, builders and investors can actually use.

The conversation has started. The framework is on the table.

Mothusi Pahl serves on the Advisory Council of the Alliance for Innovation and Infrastructure and on the Board of Directors of the Great Plains Institute. He is principal at Hartwell & Loche, a strategic advisory practice focused on energy, regulated industries and commercial strategy.

Published with the Alliance for Innovation and Infrastructure (Aii). The views and opinions expressed are solely those of the author and do not necessarily reflect the views of Aii or its leadership.

* * *

Citations and Notes

1. Gartner, "Worldwide Sovereign Cloud IaaS Spending Will Total $80 Billion in 2026," press release, February 9, 2026, https://www.gartner.com/en/newsroom/press-releases/2026-02-09-gartner-says-worldwide-sovereign-cloud-iaas-spending-will-total-us-dollars-80-billion-in-2026. Gartner forecasts $80.4 billion in 2026 sovereign cloud IaaS spending, rising to over $110 billion by 2027. Regional sovereign cloud IaaS breakdown reproduced in the table in this article.

2. The SAAFE-7 framework (Sovereignty Architecture Assessment Framework for Exposure, seven dimensions), originally introduced as the Pahl Sovereignty Framework, and its seven-dimension classification system are proposed as an open standard for industry adoption. The framework, its terminology and its structural architecture were developed by Mothusi Pahl and are published here to establish priority and provenance. Commercial use of the framework is welcomed with attribution.

3. Clayton M. Christensen, Taddy Hall, Karen Dillon and David S. Duncan, *Competing Against Luck: The Story of Innovation and Customer Choice* (New York: HarperBusiness, 2016). Christensen's formulation defines a job as "the progress that a person is trying to make in a particular circumstance"; the phrasing in this article paraphrases that definition.

4. Bill Whyman, "Sovereign Cloud–Sovereign AI Conundrum: Policy Actions to Achieve Prosperity and Security," Center for Strategic and International Studies, December 4, 2025, https://www.csis.org/analysis/sovereign-cloud-sovereign-ai-conundrum-policy-actions-achieve-prosperity-and-security.

5. McKinsey & Company, “The Sovereign AI Agenda: Moving from Ambition to Reality,” December 18, 2025, https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/the-sovereign-ai-agenda-moving-from-ambition-to-reality.

6. Accenture, “Sovereign AI: From Managing Risk to Accelerating Growth,” November 3, 2025, https://www.accenture.com/us-en/insights/technology/sovereign-ai. Introduces the Sovereign Maturity Index, scoring digital and AI sovereignty on a 0–100 scale across 1,928 organizations in 28 countries.

7. World Economic Forum and Bain & Company, *Rethinking AI Sovereignty: Pathways to Competitiveness through Strategic Investments*, January 20, 2026, https://www.weforum.org/publications/rethinking-ai-sovereignty/. Frames AI sovereignty as “strategic interdependence” and identifies five economy archetypes for AI competitiveness.

8. Tony Blair Institute for Global Change, *Sovereignty in the Age of AI: Strategic Choices, Structural Dependencies and the Long Game Ahead*, January 2026, https://institute.global/insights/tech-and-digitalisation/sovereignty-in-the-age-of-ai-strategic-choices-structural-dependencies. Identifies seven strategic levers for governments to expand agency over AI.

9. NVIDIA, “Global Public Sector — National Transformation With Sovereign AI,” https://www.nvidia.com/en-us/industries/global-public-sector/. For a representative country-level program, see “NVIDIA to Help Elevate Japan’s Sovereign AI Efforts Through Generative AI Infrastructure Build-Out,” NVIDIA Blog, https://blogs.nvidia.com/blog/japan-sovereign-ai/. NVIDIA has published similar program announcements for Saudi Arabia, the UAE, South Korea, Germany, Italy, Thailand, Vietnam, India and the United Kingdom over the same period.

10. The September 2022 explosions damaged the Nord Stream 1 and Nord Stream 2 pipelines in the Baltic Sea. See “The Nord Stream Incident: Open Briefing,” United Nations Security Council Report, August 26, 2025, https://www.securitycouncilreport.org/whatsinblue/2025/08/the-nord-stream-incident-open-briefing-2.php. For a contemporaneous analytical treatment of the incident’s signaling implications for energy infrastructure, see Eugene Rumer, “Shock and Awe: Who Attacked the Nord Stream Pipelines?” Carnegie Endowment for International Peace, September 30, 2022, https://carnegieendowment.org/russia-eurasia/politika/2022/09/shock-and-awe-who-attacked-the-nord-stream-pipelines.

11. Joseph Rand et al., *Queued Up: 2025 Edition, Characteristics of Power Plants Seeking Transmission Interconnection As of the End of 2024*, Lawrence Berkeley National Laboratory, December 2025, https://emp.lbl.gov/queues. As of end-2024, approximately 2,290 GW of generation and storage capacity were actively seeking grid interconnection in the United States. The median duration from interconnection request to commercial operation

has more than doubled from under two years for projects built 2000–2007 to over four years (with substantial regional variation extending longer) for projects built 2018–2024.

12. See Jangseop Shin et al., "A Survey on Air-Gap Attacks: Fundamentals, Transport Means, Attack Scenarios and Challenges," *Sensors* 23, no. 6 (2023): 3215, https://www.mdpi.com/1424-8220/23/6/3215; and the body of work by Mordechai Guri (Ben-Gurion University Cyber Security Research Center) documenting electromagnetic, acoustic, optical, thermal and power-line covert channels for exfiltrating data from air-gapped systems. The 2010 Stuxnet worm is the canonical case of a hardware-level supply chain compromise bridging an air-gapped industrial control network.

13. Steven A. Cook, "For Saudi Arabia, Qatar, and the UAE, Investment in AI Is Partly About U.S. Protection," *Foreign Policy*, February 23, 2026, https://foreignpolicy.com/2026/02/23/gulf-states-investment-ai-american-protection-qatar-uae-saudi/. The combined Saudi, UAE and Qatari technology and AI infrastructure commitment figure of approximately $2.5 trillion is documented across multiple sources including Latitude Media's June 2025 analysis and tracking of the May 2025 announcements (Saudi Arabia $600B, UAE $200B over an existing $1.4T plan, Qatar $1.2T).

14. See "Advancing Southeast Asia's AI Future Through Sovereign AI Models," Fulcrum (ISEAS-Yusof Ishak Institute), February 25, 2026, https://fulcrum.sg/advancing-southeast-asias-ai-future-through-sovereign-ai-models/. Singapore (National AI Strategy 2.0), Indonesia (Sahabat-AI), Malaysia (ILMU LLM and AI-only data center policy) and Vietnam (December 2025 Law on Artificial Intelligence, effective March 2026) have each established sovereign AI infrastructure as a national priority.

15. European sovereignty regulations include France's SecNumCloud certification (ANSSI v3.2), which requires that cloud providers be structurally immune to non-EU laws and limits non-EU shareholder ownership. See "From Cloud Souverain to Cloud de Confiance: A Political Definition of Clouds," Aneo, https://www.aneo.eu/en/blog/cloud-souverain-cloud-de-confiance. Germany's BSI Cloud Computing Compliance Criteria Catalogue (C5:2020) provides equivalent national requirements. The European Cybersecurity Certification Scheme for Cloud Services (EUCS), under development by ENISA, would harmonize these requirements at the EU level. Compliance with the GDPR, NIS2 Directive and Digital Operational Resilience Act (DORA) creates additional structural constraints on US-headquartered providers serving EU regulated sectors.

16. TrendForce data shows TSMC held 70.2 percent of global semiconductor foundry revenue in Q2 2025 and approximately 90 percent of global capacity at advanced logic nodes (3nm and below). See "2Q25 Foundry Revenue Surges 14.6% to Record High, TSMC's Market Share Hits 70%," TrendForce, September 1, 2025,

https://www.trendforce.com/presscenter/news/20250901-12691.html. The article's "over 50 percent" framing is conservative across competing measurement methodologies.

17. US Department of Commerce, "Statement on UAE and Saudi Chip Exports," November 19, 2025, https://www.commerce.gov/news/press-releases/2025/11/statement-uae-and-saudi-chip-exports. Both G42 (UAE) and HUMAIN (Saudi Arabia) received Bureau of Industry and Security authorization to import the equivalent of 35,000 NVIDIA Blackwell GB300 chips each, conditioned on rigorous security and reporting requirements. The Microsoft-G42 partnership including a $1.5 billion equity investment and binding Intergovernmental Assurance Agreement is documented in Brad Smith, "Microsoft Is Expanding Its Investment in the UAE to $15.2 Billion," Microsoft On the Issues, November 3, 2025.

18. Eva Dou and Gerrit De Vynck, "Israel's Pager Attack on Hezbollah Reveals Power of Supply Chain Threats," *Washington Post*, September 19, 2024, https://www.washingtonpost.com/technology/2024/09/19/hezbollah-pager-attack-supply-chain/. On 17–18 September 2024, thousands of pagers and walkie-talkies distributed to Hezbollah operatives across Lebanon and Syria detonated simultaneously. Subsequent reporting established that the devices had been weaponized at the manufacturing stage through a shell-company supply chain originating with a Hungarian firm operating under license from a Taiwanese pager manufacturer.

19. Epoch AI estimates that NVIDIA chips account for over 60 percent of total AI compute capacity globally, with substantially higher concentration at the frontier of large-scale training runs. See "Data on AI Chip Owners," Epoch AI, https://epoch.ai/data/ai-chip-owners; and "Leading AI Companies Have Hundreds of Thousands of Cutting-edge AI Chips," Epoch AI, https://epoch.ai/data-insights/computing-capacity. AMD's AI chip revenue was approximately 4 percent of NVIDIA's in 2024.

20. The Clarifying Lawful Overseas Use of Data (CLOUD) Act, 18 U.S.C. § 2713, enacted as part of the Consolidated Appropriations Act, 2018, Pub. L. No. 115-141, Div. V, § 105 (March 23, 2018), authorizes US law enforcement to compel disclosure of data within the "possession, custody, or control" of a US-headquartered service provider regardless of where the data is physically stored. Section 702 of the Foreign Intelligence Surveillance Act, 50 U.S.C. § 1881a, authorizes warrantless surveillance of non-US persons reasonably believed to be located outside the United States, including communications passing through US-headquartered electronic communication service providers. National Security Letters under 18 U.S.C. § 2709 permit the FBI to compel production of subscriber and transactional records without a court order, with statutory nondisclosure provisions. Together these instruments establish that corporate domicile in the United States creates a jurisdictional reach that data residency outside US territory does not sever.

21. Huawei was added to the US Department of Commerce Bureau of Industry and Security Entity List in May 2019, with the restrictions extended via the Foreign Direct Product Rule in May and August 2020 to cover foreign-produced semiconductors made using US technology. See US Department of Commerce, "Commerce Addresses Huawei's Efforts to Undermine Entity List, Restricts Products Designed and Produced with U.S. Technologies," May 15, 2020, https://2017-2021.commerce.gov/news/press-releases/2020/05/commerce-addresses-huaweis-efforts-undermine-entity-list-restricts.html. For analytical treatment of the supply chain consequences, see Chad P. Bown, "The US Is Trying to Use Export Controls to Restrict Huawei's Access to Semiconductors," Peterson Institute for International Economics, October 2020, https://www.piie.com/research/piie-charts/us-trying-use-export-controls-restrict-huaweis-access-semiconductors.

22. The structural exposure created by US-headquartered corporate ownership has been the subject of recent legal analysis. A December 2025 legal opinion commissioned by the German Federal Interior Ministry and authored by the University of Cologne's law faculty concluded that physical data residency in Europe is legally insufficient to escape CLOUD Act reach when the cloud provider's parent is subject to US jurisdiction. For a typology of architectural responses — public cloud with enhanced controls, hyperscaler EU sovereign offerings, joint ventures with European-controlled operators and European-native providers — see "EU Cloud Sovereignty: Four Alternatives to Public Clouds," Unit8, June 9, 2025, https://unit8.com/resources/eu-cloud-sovereignty-four-alternatives-to-public-clouds/. Joint venture examples include Bleu (Microsoft/Orange/Capgemini, France), Delos Cloud (Microsoft/Arvato/SAP, Germany) and S3NS (Google/Thales, France). European-native providers operating with no US corporate chain include OVHcloud, Hetzner, STACKIT and Scaleway.

23. United Nations Convention on the Law of the Sea, opened for signature December 10, 1982, 1833 U.N.T.S. 397 (entered into force November 16, 1994). Articles 91–94 establish flag-state jurisdiction over vessels on the high seas; Articles 86–115 codify the regime governing the high seas more broadly, including the principle that ships are subject to the exclusive jurisdiction of the state whose flag they fly. The United States has signed but not ratified UNCLOS while observing most of its provisions as customary international law.

24. CoreWeave, Crusoe and Lambda are among the leading specialized GPU cloud providers (commonly termed "neoclouds") serving AI training and inference workloads, distinguished from hyperscalers (AWS, Azure, Google Cloud) by their dedicated focus on GPU-as-a-service. See Synergy Research Group, "Neoclouds Currently Growing by Over 200% per Year; Will Reach $180 Billion in Revenues by 2030," October 13, 2025, https://www.srgresearch.com/articles/neoclouds-currently-growing-by-over-200-per-year-will-reach-180-billion-in-revenues-by-2030. All three companies are US-headquartered.

* * *