



The AI Infrastructure Triad in Regional Governance: How Regions Balance Progress, Sustainability, and Equity

Citation

Carvão, Paulo and Tushar Kanade. "The AI Infrastructure Triad in Regional Governance: How Regions Balance Progress, Sustainability, and Equity." M-RCBG Associate Working Paper Series 2026.270, Harvard University, Cambridge, MA, May 2026.

Link

<https://dash.harvard.edu/handle/1/42739545>

Terms of use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Pass Through Licensed Materials, as set forth at

<https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFIOWI2M2EwYTg>

Accessibility

<https://accessibility.huit.harvard.edu/digital-accessibility-policy>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#)



HARVARD Kennedy School

MOSSAVAR-RAHMANI CENTER
for Business and Government

The AI Infrastructure Triad in Regional Governance: How Regions Balance Progress, Sustainability, and Equity

Paulo Carvão
Harvard Kennedy School

Tushar Kanade
Harvard Kennedy School

May 2026

M-RCBG Associate Working Paper Series | No. 270

The views expressed in the M-RCBG Associate Working Paper Series are those of the author(s) and do not necessarily reflect those of the Mossavar-Rahmani Center for Business & Government or of Harvard University. The papers in this series have not undergone formal review and approval; they are presented to elicit feedback and to encourage debate on important public policy challenges. This paper is copyrighted by the author(s). It cannot be reproduced or reused without permission. Pursuant to M-RCBG's Open Access Policy, this paper is available to the public at hks.harvard.edu/centers/mrcbg free of charge. Papers may be downloaded for personal use only.

The AI Infrastructure Triad in Regional Governance: How Regions Balance Progress, Sustainability, and Equity

Paulo Carvão¹, Tushar Kanade²

*This manuscript has been accepted for publication in **AI & Society** on May 7, 2026. The final published version will appear in the journal.*

Abstract

The rapid expansion of artificial intelligence infrastructure, including data centers and the energy, land, water, and labor systems that support them, presents regional policymakers with trade-offs that are poorly captured by the prevailing “innovation versus regulation” frame. This article develops the *AI Infrastructure Triad* as a conceptual framework for analyzing three competing priorities in regional AI infrastructure governance: Progress, Sustainability, and Equity. We argue that regions are unlikely to maximize all three simultaneously under current technological, institutional, and resource conditions. Drawing on prior work on the economic, physical, and moral limits of AI development, a previously coded dataset of 10,068 public comments submitted to the 2025 U.S. AI Action Plan and illustrative regional cases, the article interprets stakeholder and regional positions as different ways of prioritizing the triad’s frontiers. The evidence is used illustratively rather than as a full causal test. The paper’s contribution is to clarify the trade-offs that infrastructure decisions often obscure, distinguish deliberate triad governance from default allocation by market power or regulatory inertia, and propose a *Deliberate Triad Choice Framework* for policymakers considering AI infrastructure decisions of significant scale.

Keywords: AI infrastructure, data centers, innovation policy, regional economics, sustainability, governance, triad

AI-use Declaration: The authors used AI-based editorial tools for language refinement and structural review. The authors remain responsible for the argument, analysis, sources, and final text.

1. Introduction

In January 2026, two of the world’s most visible technology companies made commitments that appear to contradict the logic of profit maximization.

On January 13, Microsoft published what it called a “Community-First AI Infrastructure” plan (Microsoft 2026), pledging to pay full electricity costs for its data centers, replenish more water than its facilities consume, partner with trade unions for local job creation, pay full property taxes (explicitly forgoing the

¹ Senior Fellow, Mossavar-Rahmani Center for Business & Government - Harvard Kennedy School (paulo_carvao@harvard.edu)

² 2025-26 Polymath Fellow, Research Affiliate - Massachusetts Institute of Technology (tusharka@mit.edu)

abatements that municipalities routinely offer), and invest in AI education for host communities. A week later, on January 20, OpenAI unveiled “Stargate Community” (OpenAI 2026) detailing site-specific commitments for its planned data centers, including guarantees that operations would not raise local electricity costs, regional workforce training through OpenAI Academies, and, in Texas, a \$1 billion equity investment in SB Energy to build a 1.2-gigawatt Stargate campus with dedicated energy generation supplying most of its power.

By early March 2026, United States President Trump announced a “ratepayer protection pledge”, elevating the political economy of AI infrastructure to the highest levels of federal attention (The White House, 2026).

Why would firms that spent 2024 and 2025 lobbying to remove regulatory barriers now voluntarily adopt community-benefit obligations that no law required? The answer, we argue, lies in a set of constraints that have fundamentally altered the political economy of AI infrastructure, and that existing governance frameworks are poorly equipped to address.

The age of AI as a purely digital phenomenon is over. The assumption inherited from the Software as a Service (SaaS) era that digital infrastructure scales without significant material inputs has ended. Today, the training of state-of-the-art AI models requires hundreds of millions of dollars in compute, electricity sufficient to power a city, and significant water for cooling (IEA 2025). Global data center electricity consumption reached approximately 415 terawatt-hours in 2024, an amount comparable to powering tens of millions of homes for a year, and is projected to more than double by 2030 (IEA 2025). In the United States, data centers consume over 4 percent of all electricity (O’Donnell and Crownhart 2025); in Northern Virginia, facilities consume approximately 26 percent of the regional electric load, and an estimated 20 percent of planned global projects face delays due to grid capacity limitations (Goldman Sachs 2025).

These are hard physical constraints that may shape where AI infrastructure can be built, how fast it can be deployed, and who bears the costs. The dominant paradigm in U.S. AI policy, as codified in Executive Order 14179 “Removing Barriers to American Leadership in Artificial Intelligence” (January 2025), treats these constraints as obstacles to be eliminated. The prevailing assumption is that the path to AI leadership runs through the removal of regulatory friction: streamlined permitting, expedited environmental review, and federal preemption of state-level governance.

We propose a different approach and argue that regions balancing AI infrastructure development face three competing priorities unlikely to be simultaneously maximized.: Progress (scale and speed, maximizing the pace and magnitude of deployment), Sustainability (integrating infrastructure with environmental constraints and ensuring long-term resource viability), and Equity (distributing costs and benefits broadly across populations). We then address the policy question of how to build AI infrastructure while *balancing* this triad. Our approach does not assert that regions must choose one priority and abandon the others. It asserts that optimizing for all three simultaneously is difficult under current technological and institutional conditions, and that the distribution of effort across the three dimensions constitutes a consequential regional infrastructure policy choice.

Political economists have long recognized that governance systems face structural tradeoffs among competing objectives: Mundell’s (1963) “impossible trinity” in monetary policy, Rodrik’s (2011) political trilemma of globalization, and Ostrom’s (1990) work on commons governance each demonstrate that institutional choices involve irreducible tensions rather than optimal solutions. The AI infrastructure challenge is a problem of this kind.

This article integrates two bodies of prior work: an analysis of three fundamental limits to AI development (economic, physical, moral) detailed in Section 2.4 (Carvão 2025), and an analysis of 10,068 public comments submitted to the U.S. AI Action Plan that maps six stakeholder worldviews to organizational type, summarized in Section 3.1 (Carvão, Yashiro, and Jeloka 2025). The AI infrastructure triad synthesis translates these findings into a policy-relevant framework.

The paper makes four contributions. First, it develops the triad as a framework for regional AI infrastructure governance. The framework is analytically distinct from macro-level AI trilemma arguments, such as Ernst's (2022), because it focuses on state, metropolitan, and regional decisions about physical infrastructure rather than on economy-wide effects of AI adoption. Second, it clarifies how prior stakeholder-worldview research can be interpreted through the lens of infrastructure trade-offs, while acknowledging that this mapping is interpretive. Third, it uses selected cases to illustrate how regions can move toward different triad frontiers under different constraint profiles. These examples of constraint-driven policy can create a feedback loop in which physical constraints shape corporate behavior in ways that market incentives or traditional regulation may not produce independently, a topic for future development. Fourth, it offers a Deliberate Triad Choice Framework: a set of questions policymakers can use to make infrastructure trade-offs explicit before decisions of significant scale are made.

1.1 Methodological Approach

This article introduces a conceptual framework supported by illustrative empirical material. It does not claim to provide a comprehensive comparative case study, a causal test of regional outcomes, or a new coding of the full public comment dataset. Its primary contribution is interpretive, since it develops the AI Infrastructure Triad as a heuristic framework for analyzing regional governance of AI infrastructure.

The paper draws on three main sources. First, it builds on prior work that identifies three limits to AI development: an economic limit, a physical limit, and a moral limit.

Second, it uses a previously coded analysis of 10,068 public comments submitted in response to the 2025 U.S. AI Action Plan. That analysis identified six recurring AI policy worldviews and found that worldview orientations varied by organizational type. The prior study used a qualitative review and LLM-assisted classification. Comments were coded by institution type and six AI worldviews. GPT-4 Turbo and Gemini 1.5 Flash analyses were compared for validation, with Gemini used at scale. Misclassifications were manually corrected. Full methodological details appear in Carvão, Yashiro, and Jeloka 2025.

Third, we reinterpret those worldview patterns through the lens of infrastructure governance examples, using them to illustrate how different actors emphasize different triad dimensions rather than as independent validation of the triad. The cases were selected for theoretical variation, not statistical representativeness. Northern Virginia and Texas receive primary attention because they are U.S. regional examples connected to the federal policy context from which the public comment evidence is drawn. Taiwan, Singapore, and the Gulf States are used as comparative illustrations of different ways regions can prioritize long-term resilience, environmental constraint, or scale-first deployment. These cases are not intended to exhaust the global range of possible AI infrastructure strategies. Other cases could support or challenge the framework and should be examined in future research.

The classifications offered here should be read as interpretive and contestable. The triad acts as a conceptual device for making visible the trade-offs among scale, environmental/resource sustainability, and distributive fairness that are often obscured in AI infrastructure debates.

2. Theoretical Foundation and the AI Infrastructure Triad

This paper sits at the intersection of three literatures that have developed largely in isolation from one another: AI governance and policy, innovation ecosystems, and regional economics. We position our

contribution within each, then develop the theoretical framework by integrating the “three limits” of AI development with the triad synthesis.

The trilemma tradition discussed in the Introduction provides the conceptual logic for the AI Infrastructure Triad. Its main insight is that some governance problems have objectives that are individually desirable yet difficult to maximize simultaneously. The dominant “innovation versus regulation” framing treats policy choice as a single-axis problem of friction on technological deployment. The triad reframes the issue as a three-sided governance problem. Regions must consider the speed and scale of deployment, the long-term viability of energy, land, and labor systems, and the distribution of benefits and burdens across affected communities. The framework treats trade-offs not as implementation failures, but as constitutive features of regional AI infrastructure governance.

2.1 AI Governance and Policy

The AI governance literature has expanded rapidly, driven by the increase and fast deployment of large language models. The dominant analytical frame has been the tension between innovation and regulation, a binary that structures most U.S. policy debates. In both the EU’s risk-based AI Act (2024) and the shifting U.S. federal posture from Biden-era executive orders to the Trump administration’s emphasis on “removing barriers,” policy has oscillated between the two poles without resolving the underlying structural challenge.

This paper departs from the innovation-versus-regulation framing. We argue that AI policy should focus on the allocation of the physical resources upon which AI infrastructure depends. This shift from governance of algorithms to governance of infrastructure reflects the material turn in AI development described in the introduction. The AI governance literature has been slow to incorporate physical constraints as first-order variables, in part because the dominant Software as a Service (SaaS) era’s assumption that digital infrastructure scales without significant material inputs. As AI systems have become orders of magnitude more computationally intensive, this assumption no longer holds, and governance frameworks must adapt accordingly.

An empirical study identified six distinct “AI worldviews” among policy stakeholders: Accelerationist, Responsible AI Advocate, Open AI Innovator, Safety Advocate, Public Interest AI, and National Security Hawk (Carvão et al. 2025). A subsequent work demonstrated that organizational type (Big Tech, Small Tech, Civil Society) correlates with worldview orientation (Carvão, Yashiro, and Jeloka 2025). We argue in Section 3 that these worldviews map onto distinct preferences for triad priorities, revealing that the ostensible “innovation versus regulation” debate is also a contest over which triad frontier should be prioritized.

2.2 Innovation Ecosystems

The innovation ecosystems literature, rooted in AnnaLee Saxenian’s (1994) comparative study of Silicon Valley and Route 128, has emphasized the importance of network structures, knowledge spillovers, and institutional configurations to regional innovation trajectories. It has fundamentally reshaped how scholars and policymakers think about regional competitive advantage. Subsequent work by Feldman and Kogler (2010) formalized the geography of innovation, showing that knowledge spillovers are spatially bounded and that proximity facilitates the transfer of tacit knowledge essential for breakthrough innovation.

More recently, Collison and Cowen (2019) called for a “new science of progress” to understand the institutional, cultural, and scientific factors that drive innovation and how they might be strengthened to

accelerate technological and economic progress. This paper shares that normative orientation while seeking to identify how innovation ecosystems can be governed more deliberately.

However, the dominant frameworks in this tradition (Saxenian's networks, Porter's 1990 clusters) treat physical constraints as background conditions rather than active policy variables. This omission of physical constraints, harmless in the era of software-driven innovation, becomes untenable when AI data centers consume electricity at a scale material to national economies and require land, water, and specialized labor in quantities that compete with other regional priorities.

We introduce physical constraints as endogenous variables in innovation ecosystem dynamics and argue that the choices regions make about energy, land, and labor allocation for AI infrastructure are not peripheral to innovation outcomes. The regions that innovate most successfully in the AI era will not necessarily be those with the fewest constraints, but those that deploy constraints most deliberately as tools for shaping the kind of innovation ecosystem they wish to build.

2.3 Regional Economics and Infrastructure

The regional economics literature provides the third pillar of our theoretical foundation. Richard Florida (2002) established talent attraction as a central driver of regional economic dynamism, shifting analytic focus from firms to workforce. Enrico Moretti (2012) documented the growing divergence between "innovation cities" and the rest of the U.S. economy. Autor, Dorn, and Hanson (2013) showed how economic shocks distribute unevenly across regions, with lasting consequences for employment and social outcomes.

This body of work establishes that infrastructure decisions have distributional consequences that unfold across space and time. When a region decides to build data center capacity, it is not making a neutral technical choice. It is reshaping local labor markets, energy systems, land use patterns, and fiscal structures in ways that create winners and losers. The place-based policy literature has increasingly recognized this dynamic, arguing that policy interventions must account for the geographic specificity of economic outcomes (Autor, Dorn, and Hanson 2016; Kline and Moretti 2014).

Our work applies these insights to AI infrastructure, one of the largest place-based industrial investments of the twenty-first century. Yet the policy discourse has largely failed to engage with the distributional questions that regional economics identifies as essential. When data centers consume 26 percent of a region's electric load (Dominion Energy 2024), when grid connection wait times stretch to years (Camus Energy 2024), and when community opposition mounts (Brookings 2025), the distributional stakes become impossible to ignore. The triad framework provides a vocabulary for making these stakes explicit.

2.4 The Three Limits of AI Development

This paper builds on the identification of three fundamental limits on AI development, each of which constrains the trajectory of the technology in distinct ways (Carvão 2025).

The Economic Limit refers to the escalating capital expenditure required to develop and deploy frontier AI systems. Investment has reached a scale previously associated with sovereign states rather than private firms. Deutsche Bank projects cumulative global data center spending of \$4 trillion by 2030 (Deutsche Bank 2025). Because scaling laws (Sutton 2019; Kaplan et al. 2020) imply that performance gains become more expensive with each model generation, profitability is increasingly uncertain, and even well-

capitalized firms must invest selectively. This limit constrains which actors can participate in frontier AI development and shapes the geographic distribution of investment.

The Physical Limit constitutes a binding constraint that regional policy can directly address. Energy constraints are the most visible: as noted above, Northern Virginia’s data centers consume approximately 26 percent of Dominion Energy Virginia’s load, and the IEA projects global data center electricity demand will more than double by 2030 (IEA 2025). Land constraints manifest through zoning disputes, density requirements, and competition with other land uses. Labor constraints include both the shortage of specialized technical workers (up to 1.9 million manufacturing positions could remain unfilled by 2033) and the need for trades workers (electricians, HVAC technicians, construction workers) to physically build and maintain facilities (Deloitte and The Manufacturing Institute 2024).

The Moral Limit asserts that human agency must remain central to decisions mediated by AI systems. The risk of “moral outsourcing”, delegating decisions with ethical dimensions to automated systems without meaningful human oversight, requires enforceable accountability boundaries. Although not the primary focus of this paper, the moral limit undergirds the equity dimension of the triad and grounds the normative case for deliberate, democratically informed triad choices.

Together, these interconnected limits establish the possibility space for regional AI infrastructure policy: the economic limit shapes which actors can invest at scale, the physical limit shapes where and how, and the moral limit shapes whether the resulting systems serve broad or narrow interests.

Table 1: The Three Limits of AI Development

Limit Type	Core Constraint	Policy Implication	Triad Relevance
Economic	Capital expenditures approaching G20 GDP scale; marginal performance gains increasingly expensive	Profitability uncertainty forces selective investment	Determines which actors can participate; shapes geographic distribution
Physical (The New AI Triad)	Energy: grid capacity, renewables integration. Land: parcels, density, zoning. Labor: skilled workforce availability	Hard ceilings on capacity expansion; these are the primary policy levers	Constitutes the triad: trade-offs among scale, sustainability, equity
Moral	Human agency must remain central; risk of "moral outsourcing"	Enforceable accountability boundaries required	Normative foundation for equity dimension and deliberate governance

Adapted from Carvão (2025). The Triad relevance column is the authors’ synthesis.

2.5 The AI Infrastructure Triad

These limits are synthesized into a policy-relevant triad. The key concept is that regions pursuing AI infrastructure development face three competing priorities that are unlikely to be simultaneously optimized. They must instead make deliberate choices about how to weigh each dimension, recognizing that gains on one typically come at the expense of another. The distribution of effort across these dimensions is amongst the most consequential policy choices a region can make about its AI future.

We define the three priorities as follows:

Progress refers to the scale and speed of AI infrastructure deployment. It prioritizes data center capacity, compute availability, time-to-market, and the rapid expansion of energy and network systems needed to support AI development. Its primary beneficiaries are firms, investors, and national governments seeking technological leadership and regions seeking capital investment. Progress-oriented decisions often accept higher environmental pressure, concentrated infrastructure burdens, and reduced time for community review.

Sustainability refers to the capacity of ecological, energy, water, and land-use systems to support AI infrastructure over time. It prioritizes grid reliability, renewable energy integration, water stewardship, carbon intensity, land-use planning, and long-term resource viability. Its primary focus is not simply “social good” but the physical and environmental systems that make continued infrastructure operation possible. Sustainability-oriented decisions often accept slower deployment, higher compliance costs, or reduced near-term competitiveness.

Equity refers to the distribution of AI infrastructure’s benefits, burdens, and decision authority. It asks who receives employment, tax revenue, public investment, and access to compute, and who bears electricity costs, land-use disruption, environmental burdens, displacement risk, or procedural exclusion. Its primary emphasis is on workers, host communities, and the public institutions that bear distributional consequences. Equity-oriented decisions often accept longer deliberation, binding community commitments, higher developer obligations, or less flexibility for firms.

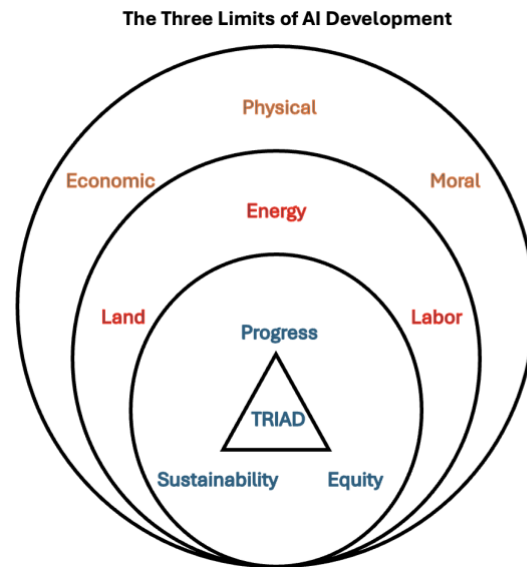


Figure 1 : The AI Infrastructure Triad and the Limits of AI Development

An example can illustrate the distinction between Sustainability and Equity. A state could require data centers to meet strict carbon- or water-efficiency standards while allowing the resulting compliance costs to be passed through to ratepayers or exclude host communities from siting decisions. Such a policy would improve Sustainability while worsening Equity. Conversely, a community benefit agreement could expand local hiring, tax revenue sharing, and public participation while still permitting construction in a water- or grid-constrained region. Such a policy would improve Equity while leaving Sustainability unresolved. The two vertices overlap but remain analytically distinct. The triad is not a binary of Progress versus socio-environmental constraint. It identifies three different questions: how much capacity and how fast; within what physical and ecological limits; and with what distribution of benefits, burdens, and voice.

Table 2: AI Infrastructure Triad: Priority Frontiers

Priority	Core Question	Primary Constituency	Typical Indicators	Trade-offs
Progress	How much AI infrastructure can be built, and how fast?	Firms, investors, national-security planners, growth-oriented regions.	MW deployed, compute capacity, permitting speed, capital investment.	Grid strain, environmental pressure, weaker community review.
Sustainability	Can physical and ecological systems support the buildout over time?	Environmental systems, grid operators, long-term planners, future users.	Carbon intensity, water use, PUE, renewable share, grid reliability.	Slower deployment, higher cost, possible loss of investment.
Equity	Who benefits, who pays, who decides and who is exposed to risk?	Workers, host communities, ratepayers, public institutions.	Local hiring, tax revenue use, ratepayer impact, community voice, access to compute.	Longer timelines, binding obligations, reduced firm flexibility.

National security is best viewed as a cross-cutting rationale across the three vertices of the triad. It often strengthens Progress claims because geopolitical competition rewards speed, scale, and domestic capacity. Yet national security also depends on Sustainability since infrastructure that strains grids or relies on fragile supply chains may undermine resilience. Similarly, from an Equity perspective, a national AI strategy that concentrates benefits narrowly or generates durable community opposition may weaken the social legitimacy needed for long-term technological leadership.

Each vertex has a primary constituency: Progress serves the interests of firms, investors, and governments seeking competitive position; Equity serves workers, host communities, and ratepayers who bear the costs of infrastructure sited near them; Sustainability serves the physical and ecological systems (grid stability, water supply, land capacity) on which all three groups ultimately depend. These constituencies can and do form cross-cutting coalitions, but they have distinct interests that map onto distinct policy demands, which is why the triad requires three vertices.

Two features of the framework merit emphasis. First, the triad is dynamic, not static. A region’s position within the triad space can shift over time as technologies mature, political coalitions evolve, and external conditions change. Second, the triad operates at multiple scales simultaneously. A national government may prioritize progress at the aggregate level while allowing or encouraging individual states or regions to pursue sustainability or equity-oriented approaches. The resulting variation is, as we argue in Section 4, a characteristic of federal systems rather than a deficiency.

The triad contrasts deliberate governance with default allocation. Without explicit choice, the balance is set by whichever actors hold the most capital or access, systematically favoring Progress almost always at a greater long-term cost.

3. Empirical Illustrations and Regional Cases

This section presents the empirical and illustrative material that informs the AI Infrastructure Triad. We draw on previous research and the analysis of public comments on the U.S. AI Action Plan (Carvão, Yashiro, and Jeloka 2025) to interpret the stakeholder worldviews and identify recurring affinities between these views and triad preferences. Then, we examine three cases, Taiwan, Singapore, and the Gulf States, as comparative illustrations of different triad balancing at the national and regional level. Finally, we analyze U.S. regional variation across Northern Virginia and Texas to show how constraint profiles can shape triad outcomes within a single federal system. Throughout, we point to an emerging cycle of constraint-driven policy innovation.

3.1 Stakeholder Worldviews and Triad Preferences

The White House AI Action Plan RFI from February 2025 generated significant public response, ranging from the largest technology firms to individual citizens. Carvão, Yashiro, and Jeloka (2025) analyzed these submissions and mapped each respondent onto a previously developed typology of six AI worldviews (Carvão et al. 2025). These worldviews (Accelerationist, National Security Hawk, Open AI Innovator, Responsible AI Advocate, Public Interest AI, and Safety Advocate) capture the distinct orientations that stakeholders bring to AI governance debates. We will discuss how organizational type and worldview orientation map to triad preferences.

This article does not re-code the original public comments dataset. Instead, it asks how the worldviews can be interpreted when the governance object is not only AI models or applications, but also the physical infrastructure needed to develop and deploy them. This interpretive step should be understood cautiously. The original authors of the typology did not intend to prove that organizational categories map cleanly onto a triad vertex.

The public comments are used to identify recurring affinities rather than to prove a one-to-one mapping between worldview and triad vertex. For example, progress-oriented submissions typically emphasized speed, federal preemption, permitting reform, and national competitiveness. Equity-oriented submissions emphasized community protection, civil rights, labor, accountability, and procedural voice. Sustainability-oriented submissions emphasized environmental review, grid reliability, energy use, water use, and long-term system resilience. These patterns help explain why infrastructure governance cannot be reduced to a single continuum from “more regulation” to “less regulation.”

Table 3: Stakeholder Worldviews and Infrastructure Implications

Worldview	Core Orientation	Typical Infrastructure Implication
Accelerationist	AI leadership depends on rapid deployment and minimal friction.	Prioritizes Progress.
National Security Hawk	AI capacity is central to geopolitical and military advantage.	Often prioritizes Progress but may support resilience-oriented Sustainability.
Open AI Innovator	Openness, competition and diffusion are central to innovation.	Can support Progress and Equity through broader access.
Responsible AI Advocate	AI should develop through standards, accountability and institutional trust.	Often bridges Sustainability and Equity.

Worldview	Core Orientation	Typical Infrastructure Implication
Public Interest AI	AI governance should prioritize public welfare, rights and democratic accountability.	Prioritizes Equity, often paired with Sustainability.
Safety Advocate	AI systems require stronger controls to prevent catastrophic or systemic harms.	May prioritize Sustainability understood as long-term societal resilience.

Worldviews by Carvão, Yashiro, and Jeloka (2025). Infrastructure implications by the authors.

Among Big Tech firms, the dominant worldview is Accelerationist (40.6 percent), followed by National Security Hawk (37.5 percent), with Responsible AI Advocate at 21.9 percent. These firms consistently advocate for the removal of regulatory barriers, federal preemption of state-level AI legislation, and rapid infrastructure deployment. Their policy positions, including calls for streamlined permitting, opposition to state-level safety mandates, and support for federal preemption, correspond directly to a Progress triad preference. OpenAI’s submission warns against a “patchwork of regulations that risks bogging down innovation”; Meta frames open-source AI as the foundation of continued U.S. leadership and condemns state-level regulatory complexity as a barrier to progress. The Responsible AI Advocate behavior, on the other hand, is typically associated with more mature enterprise technology providers, such as Microsoft or IBM. These companies balance a faith in the potential of the technology with reputational risk concerns, catering to a broader set of stakeholders, including other large enterprise clients. (Carvão, Yashiro, and Jeloka 2025)

Venture capital firms echo and amplify the Accelerationist orientation, with submissions from firms such as True Ventures and Andreessen Horowitz characterizing algorithmic impact assessments as punitive to startups and advocating for “light-touch” regulatory models (Carvão, Yashiro, and Jeloka 2025). The VC position maps onto the Progress frontier: portfolio companies depend on rapid market access, and the regulatory reforms VCs advocate (permitting acceleration, federal preemption, reduced environmental review) directly serve the Progress priority.

Civil society organizations present a starkly different pattern. The dominant worldview is Public Interest AI (43.7 percent), followed by Responsible AI Advocate (30.8 percent) (Carvão, Yashiro, and Jeloka 2025). These organizations emphasize accountability, transparency, trust, and risk mitigation. Their policy positions, support for enforceable standards, algorithmic impact assessments, community benefit requirements, and environmental review, correspond to a combined Sustainability and Equity triad preference.

Small tech firms are split between Accelerationist (25.7 percent) and Responsible AI Advocate (26.7 percent) positions. The Responsible AI Advocate tendency reflects enterprise software firms championing voluntary standards; the Accelerationist tendency reflects VC-backed startups focused on liability relief and federal preemption (Carvão, Yashiro, and Jeloka 2025).

This internal division challenges the monolithic “industry view” framing that dominates popular coverage of AI policy debates, suggesting that firm size and business model inform regulatory preference.

Table 4: Stakeholder Worldviews and Triad Preferences

Organization Type	Dominant Worldviews	Primary Triad Affinity	Secondary Affinity	Characteristic Policy Positions
Big Tech	Accelerationist; National Security Hawk	Progress	Sustainability when tied to grid reliability or reputational risk	Federal preemption, permitting reform, rapid deployment, infrastructure investment
Venture Capital	Accelerationist	Progress	Limited Equity through startup access and competition claims	Light-touch regulation, liability limits, opposition to impact assessments
Civil Society	Public Interest AI; Responsible AI	Equity	Sustainability	Accountability, transparency, community protection, rights, environmental review
Enterprise tech firms	Responsible AI Advocate	Sustainability	Equity	Standards, auditability, institutional trust, risk management
Small Tech	Mixed	Context-dependent	Context-dependent	Varies by business model, market position and regulatory exposure

Carvão, Yashiro, and Jeloka (2025). Triad affinity mapping is the authors' synthesis.

The prevailing “innovation versus regulation” frame suggests a single axis of disagreement, with stakeholders arrayed along a spectrum from more to less regulation. The triad mapping reveals that the actual disagreement is multi-dimensional: stakeholders differ not merely on how much governance AI requires, but on *which values* should govern infrastructure decisions, *who* should bear the costs of those decisions, and *what kind* of AI ecosystem should emerge.

3.2 Comparative Illustrations of Triad Balancing

The stakeholder analysis indicates that different actors favor different triad frontiers. The following case studies show that regions, too, make deliberate, though not always explicit, choices about how to balance the three priorities. The following comparative examples are intentionally brief. They are not presented as full case studies or as a representative global sample. Taiwan illustrates long-term institutional orchestration and patient capital. Singapore illustrates the use of physical constraints as sustainability-oriented governance. The Gulf States illustrate a scale-first strategy enabled by capital and energy abundance. These cases are used to clarify the framework before turning to U.S. regional variation.

Taiwan: Patient Capital and Institutional Orchestration

Taiwan’s semiconductor ecosystem, centered on the Taiwan Semiconductor Manufacturing Company (TSMC) and the Hsinchu Science Park, represents a successful example of deliberate triad balancing. Taiwan’s approach has prioritized sustainability, understood as long-term viability and resilience rather than exclusively environmental terms, through patient capital structures and institutional orchestration that operate on decades-long timescales.

The institutional foundation is the Industrial Technology Research Institute (ITRI), established in 1973 as a bridge between government-funded research and commercial application. ITRI incubated the

semiconductor industry, combining government seed capital with a state-orchestrated, quasi-patient capital approach: the government absorbed early-stage technological and financial risk, then reduced its direct role as commercial viability was demonstrated. The result was extraordinary manufacturing specialization. TSMC pioneered the foundry model, manufacturing chips exclusively for other companies rather than designing chips of its own, serving fabless firms such as Qualcomm and Nvidia that design chips but do not operate fabrication plants. Others, like Intel, pursued a vertically integrated approach. Taiwan's industrial policy and public-private financing mechanisms absorbed the long-term capital risk that such specialization entailed (Shih and Wang 2009), producing an ecosystem that achieved global dominance through process innovation and manufacturing excellence.

The Hsinchu Science Park exemplifies how deliberate spatial concentration can foster positive knowledge spillovers. Its planned density, concentrating manufacturing, and supporting industries within a compact geographic area, created conditions for knowledge transfer and rapid iteration consistent with the innovation cluster dynamics analyzed by Saxenian (1994) and Feldman and Kogler (2010). Generous R&D tax credits reinforced a virtuous cycle of process innovation (Shih and Wang 2009).

Taiwan's triad position is characterized as a sustainability-oriented balance. Taiwan did not pursue the fastest possible deployment (a pure "Progress" approach would have favored vertical integration and brute-force scaling), and the equity dimension remained constrained by the geographic concentration of benefits within the Hsinchu corridor. But the sustainability of the model, its resilience over five decades, its capacity to generate continuous process innovation, and its ability to maintain global technological leadership exemplify how a constraint-driven approach can produce durable competitive advantage.

Singapore: Regulatory Innovation

Singapore's approach represents an example of using constraints to drive policy innovation (Liu et al. 2025; IMDA 2024). In 2019, the Infocomm Media Development Authority (IMDA) imposed a moratorium on new data center construction, a decision that was criticized as a barrier to Singapore's competitiveness in the global digital economy. The moratorium remained in place until 2022, when it was replaced by a pilot Data Centre, Call for Application (DC-CFA), a competitive allocation mechanism for new capacity. The first round ultimately awarded about 80 megawatts of capacity to selected operators. In December 2025, Singapore launched DC-CFA2, expanding allocation to at least 200 megawatts with strengthened sustainability requirements (King & Wood Mallesons 2025).

The DC-CFA2 framework imposed binding sustainability requirements as operators must source at least 50 percent of energy from approved green pathways and meet a maximum Power Usage Effectiveness (PUE) of 1.25 for new facilities (setting the ratio of data center energy usage to 25% extra energy beyond computing equipment). Singapore simultaneously committed \$1 billion to AI investment over five years under its National AI Strategy 2.0, signaling that the moratorium was pro-sustainability rather than anti-technology (IMDA 2024).

Rather than weakening Singapore's position as a regional data center hub, the moratorium and subsequent DC-CFA framework restructured how new capacity is allocated. By constraining supply, the policy forced operators to compete on efficiency, sustainability, and technological sophistication rather than on rapid capacity expansion. The competitive allocation process then directed scarce capacity toward operators with the strongest environmental and operational proposals, embedding sustainability criteria directly into the growth of the sector.

Current industry reports consistently rank Singapore among the most attractive data center markets in Asia, with vacancy rates below two percent even during the moratorium period (IMDA 2025; Research and Markets 2026). Singapore's experience shows how the moratorium imposed a constraint on Progress that generated sustainability innovation without sacrificing Singapore's position as a major data center hub.

The Gulf States: Progress through Capital Abundance

The United Arab Emirates and Saudi Arabia (Middle East Institute 2025) represent the opposite of Singapore’s approach: a capital-intensive, scale-first strategy that leverages energy abundance and sovereign wealth to maximize the speed and magnitude of infrastructure deployment. Announced AI infrastructure and ecosystem investment ambitions across the UAE and Saudi Arabia add up to more than \$100 billion, including the Stargate UAE project (a 1-gigawatt data center cluster) and Saudi Arabia’s HUMAIN initiative (500 megawatts of AI factory capacity) (Foreign Policy 2025; Arab News 2025). The Gulf States’ approach is enabled by structural conditions that most other regions lack: abundant energy supply from hydrocarbon reserves, vast tracts of developable land, and sovereign wealth funds capable of absorbing the capital expenditure without the profitability constraints that limit private investment elsewhere.

Where Singapore uses scarcity as a governance tool, the Gulf strategy attempts to overwhelm constraints through scale, making triad trade-offs visible. The Gulf approach maximizes the Progress dimension but raises sustainability questions. Data centers rely on energy systems still dominated by natural gas and other hydrocarbons (Data Center Knowledge 2025). While the Gulf States position themselves as “neutral compute hubs” offering infrastructure to global AI firms without U.S. or European political constraints, the breadth of domestic benefit distribution remains uncertain. The workforce for data center construction and operation in the Gulf is heavily reliant on imported labor. The kafala sponsorship system, which ties a migrant worker’s legal residency to a specific employer and significantly limits their ability to change jobs or leave the country without employer consent, has been criticized for creating structural power imbalances (Human Rights Watch 2023; Wilson Center 2024).

For regions with energy abundance, limited domestic demand for that energy, and sovereign wealth capable of absorbing capital risk, a Progress approach may represent the optimal utilization of comparative advantage. The triad framework does not prescribe which frontier is “correct”; it provides a vocabulary for making the chosen trade-offs explicit and for evaluating whether the resulting distribution of costs and benefits is deliberate rather than incidental.

Table 5: Triad Balancing Strategies

Region	Triad Balance	Key Mechanisms	Trade-offs Accepted
Taiwan (TSMC/Hsinchu)	Sustainability + Patient Capital	ITRI as bridging institution (1973–present); foundry focus; 25% R&D tax credit; government early capital then step-back.	Slower initial scaling [vs Progress]; geographic concentration of benefits [vs Equity]; dependent on single-sector specialization [vs Sustainability/resilience].
Singapore	Sustainability + Regulatory Innovation	Data center moratorium (2019–2022); pilot DC-CFA (~80MW, 2022); DC-CFA2 (200MW+, 2025); 50% green energy mandate; 1.25 PUE requirement; \$1B AI investment over 5 years.	Slower deployment during moratorium period [vs Progress]; potential loss of operators to competing hubs [vs Progress]; higher operating costs [vs Progress].
UAE/Saudi Arabia	Progress (capital-intensive)	\$100B+ combined investment; Stargate UAE 1GW cluster; HUMAIN 500MW AI factories;	High carbon intensity [vs Sustainability]; imported labor dependency [vs

Region	Triad Balance	Key Mechanisms	Trade-offs Accepted
		energy abundance; sovereign wealth; geopolitical positioning.	Equity]; uncertain benefit distribution [vs Equity]; reputational risks [vs Equity].

3.3 Regional Variation

In the U.S., federalism allows regional variation in AI governance. We examine Northern Virginia as a Progress example and Texas as an emerging hybrid.

Northern Virginia: De Facto Progress

Northern Virginia’s “Data Center Alley” represents the paradigmatic Progress outcome within the United States, and a case study in how Progress priorities can be achieved through the *absence* of deliberate triad governance rather than through explicit choice. The region hosts one of the densest concentrations of data center infrastructure in the world, with facilities consuming approximately 26 percent of electricity demand within Dominion Energy Virginia’s service territory (Dominion Energy 2024). Data center leasing vacancy rates are below 1 percent, and grid stress has become a defining political issue for the region. Natural gas accounts for more than half of in-state electricity generation, meaning the region’s data centers operate on a grid with higher carbon intensity than in states with cleaner electricity mixes (O’Donnell and Crownhart 2025).

This trajectory illustrates what happens when infrastructure development outpaces governance capacity. The concentration of data centers was driven by market forces, proximity to federal government data needs, early infrastructure investment, and favorable business conditions, rather than by a deliberate policy framework that weighed scale against sustainability and equity. The result has been impressive in terms of infrastructure deployment but increasingly contentious in terms of community impact: residents face rising electricity costs, environmental concerns related to generator use during peak demand periods, and land use pressures as data center footprints expand into previously residential or agricultural areas.

The Northern Virginia case indicates that, in the absence of deliberate governance, the default outcome tends toward a Progress framing as the most well-capitalized actors shape outcomes through market power alone. Deliberate governance is necessary not because Progress is inherently wrong, but because the distributional consequences of unmanaged infrastructure concentration can be significant and unevenly borne.

Texas: An Emerging Hybrid

Texas occupies an intermediate position within the triad. The state’s energy abundance, deregulated electricity market, and pro-business policy environment favor Progress deployment. On the other hand, there is emerging legislative attention to data center accountability. Notably, Senate Bill 6 (SB 6), signed in June 2025, established cost-sharing and load curtailment requirements for customers exceeding 75 megawatts (Mayer Brown 2025; Data Center Frontier 2026), setting conditions for operators to incorporate sustainability elements into deployment plans.

This trajectory reflects a broader dynamic in which growing societal resistance to unconstrained data center expansion, manifested in community opposition, legislative scrutiny, and media attention to environmental costs, has become a constraint that shapes both government policy and corporate behavior.

Texas's case suggests that the triad frontiers are not static categories but dynamic positions that shift as technologies mature, costs change, and political coalitions evolve, as discussed earlier. A state that begins with a Progress orientation may gradually incorporate sustainability and equity elements, not necessarily through regulatory mandate, but through the market dynamics of renewable energy costs and the political dynamics of community engagement.

3.4 Constraint-Driven Policy Innovation: The Emerging Feedback Cycle

The January 2026 announcements from Microsoft and OpenAI are consistent with an emerging pattern of *constraint-driven policy innovation*. However, they should be read critically as alternative explanations are plausible. The commitments by the firms may reflect political branding, strategic risk management, investor signaling, efforts to preempt more stringent regulation, or attempts to secure local consent in communities increasingly skeptical of large data center projects. This action suggests that physical constraints, community resistance, and political scrutiny have become salient enough that leading firms now treat community benefit, energy cost allocation, and workforce development as part of infrastructure strategy.

As described in the Introduction, Microsoft's January 2026 "Community-First AI Infrastructure" commitment (full electricity cost payment, water replenishment, North America's Building Trades Unions–NABTU–workforce partnerships, full property tax payment, and community AI training) represents an example of a voluntary corporate framework. OpenAI's "Stargate Community" plans are designed to prevent local electricity cost increases, establish regional workforce training, and invest in dedicated energy generation (SB Energy 2026; OpenAI 2026).

This dynamic is not unique to AI infrastructure. Scholarship on innovation systems has shown that technological change often emerges from networked relationships among firms, public institutions, and research organizations, with government playing a central role in overcoming coordination and funding constraints (Block and Keller 2009). What may be distinctive about the AI infrastructure case is the speed and scale at which similar dynamics appear to be unfolding. Singapore moved from a moratorium (2019) through pilot allocation (2022) to a full competitive framework (DC-CFA2, December 2025) in six years; Microsoft's constraint-to-commitment cycle culminated in its January 2026 announcement. The near-simultaneous convergence of government-led and corporate-led governance models raises the possibility that constraint-driven innovation may operate faster than traditional regulatory cycles, a hypothesis that requires further empirical testing.

4. Policy Implications Suggested by the Framework

We advance three arguments. First, the prevailing "remove barriers" paradigm for AI infrastructure governance is incomplete; constraints should be reconceived as policy tools that shape innovation trajectories and distribute costs and benefits. Second, federal preemption of state and local AI infrastructure governance would reduce the institutional variation that drives policy learning and innovation; regional variation can be viewed as a benefit of federalism, not a deficiency. Third, we propose a Deliberate Triad Choice Framework that provides policymakers with a structured approach to making infrastructure decisions that are transparent, accountable, and responsive to the full range of stakeholder interests.

4.1 Reframing Constraints as Policy Tools

The dominant narrative in U.S. AI infrastructure policy treats physical constraints, permitting delays, environmental review requirements, grid capacity limitations, and community advocacy as barriers to be removed. The January 2025 executive order on “Removing Barriers to American Leadership in Artificial Intelligence” (The White House 2025a) codified this framing at the federal level. The shared assumption is that the path to AI leadership runs through the elimination of friction.

The examples presented in this paper challenge this assumption. Singapore’s data center moratorium did not destroy the country’s competitive position; it produced a higher-quality, more sustainable data center sector by forcing operators to compete on efficiency. Taiwan’s patient capital model did not sacrifice global leadership; it built a semiconductor ecosystem now indispensable to the global technology supply chain.

These cases suggest that constraints can function as innovation policy tools that shape the *kind* of AI ecosystem a region builds, beyond its *speed*. Singapore’s PUE requirement forces efficiency innovation; Taiwan’s Hsinchu Science Park concentrates activity to generate knowledge spillovers; Microsoft’s NABTU partnership ties investment to local workforce development to build political sustainability.

Table 6: Constraints as Policy Tools Examples

Constraint Type	Policy Mechanism	Innovation Outcome	Example
Energy limits	Efficiency mandates; green energy requirements; PUE standards.	Technological innovation in cooling, renewable integration, grid management.	Singapore: 50% green mandate, 1.25 PUE to drive efficiency competition.
Land limits	Density requirements; zoning as cluster policy; moratorium and competitive allocation.	Knowledge spillover concentration; ecosystem co-location; higher-quality operators.	Taiwan: Hsinchu density; Singapore: moratorium as quality filter.
Labor limits	Community benefit agreements; workforce development pipelines; union partnerships.	Local workforce upskilling; political support and reduced community opposition.	Microsoft-NABTU partnership; AI Action Plan DOL Workforce Hub.

We are not arguing that all constraints are beneficial, or that more constraints are always better than fewer. Some constraints reflect genuine inefficiencies in permitting and regulatory processes that should be reformed. The argument is that a policy that is *deliberately designed* to leverage existing constraints can produce better outcomes than an indiscriminate focus on barrier removal.

4.2 Against One-Size-Fits-All Federal Preemption

Several industry submissions to the AI Action Plan RFI advocated for federal preemption of state and local AI governance, a position advanced by the December 11, 2025, executive order “Ensuring a National Policy

Framework for Artificial Intelligence” (The White House 2025b). The preemption argument is that a “patchwork” of state-level regulations creates compliance costs, regulatory uncertainty, and competitive distortions that impede deployment. If the goal is to maximize AI infrastructure buildout, a uniform federal framework would eliminate the transaction costs of navigating fifty different regulatory environments.

The triad framework reveals why this argument, while internally coherent, is incomplete. Preemption presupposes that the correct triad balance can be determined nationally and imposed uniformly across regions with different constraint profiles. This presupposition fails on three grounds. First, different regions face different physical constraints: energy-abundant Texas does not face the grid limitations constraining Northern Virginia. Second, different communities hold different priorities, and federal preemption forecloses their ability to express these preferences through local governance. Third, policy experimentation drives learning: Singapore’s moratorium informed competitive allocation, and Virginia’s experience with data center concentration has prompted legislative proposals to address community impacts.

In the 2026 Virginia General Assembly session, multiple bills sought to shift infrastructure costs from residential ratepayers to data center operators and require state-level review of new high-load facilities. SB 253, the Fair and Affordable Electric Rates and Reliability Act, advanced through both chambers, authorizing the State Corporation Commission (SCC) to shift distribution and capacity costs onto high-load data center customers (Main 2026). SB 339 directed the SCC to investigate cost subsidization. But more ambitious proposals, requiring SCC pre-approval of high-load connections, such as HB 155 and SB 619, were rejected after industry opposition (Heckt 2026). Federal preemption would eliminate this experimental variation, substituting a single approach for the distributed policy innovation that federalism enables.

None of this argues against a federal role in AI infrastructure governance. Federal investment in grid modernization, workforce development (such as the AI Workforce Research Hub established by the AI Action Plan), and research infrastructure (such as the National AI Research Resource) all require federal coordination and funding. The argument is specifically against preemption, the displacement of state and local governance authority, as distinct from federal investment and coordination.

4.3 The Deliberate Triad Choice Framework

We propose that AI infrastructure decisions of significant scale should answer five questions.

Question 1: Which triad dimension is this decision advancing? Infrastructure decisions implicitly weigh the three priorities. A streamlined permitting process weighs Progress; an energy efficiency mandate weighs Sustainability; a community benefit agreement weighs Equity. Policymakers should articulate which dimensions a decision prioritizes and why.

Question 2: Who decided, and who was consulted? If infrastructure decisions are made through processes dominated by industry actors (whose preferences overwhelmingly favor Progress), the resulting balance will reflect those preferences regardless of stated objectives. Deliberate governance requires participation representative of the full range of triad preferences.

Question 3: What are the explicit trade-offs? Deliberate governance requires that trade-offs be stated explicitly rather than obscured by aspirational language. A decision to accelerate permitting should acknowledge the sustainability and equity dimensions that may be subordinated; a decision to impose stringent environmental requirements should acknowledge the potential cost in deployment speed.

Question 4: How will costs and benefits be distributed? This question addresses geographic distribution (will benefits and costs accrue to the same communities?), demographic distribution (will workforce

opportunities reach underrepresented populations?), and temporal distribution (will construction employment be matched by long-term operational employment?).

Question 5: What accountability mechanisms exist? What metrics will be tracked, what review processes will be established, and who is responsible for course correction if outcomes diverge from expectations? Accountability mechanisms should facilitate learning and adaptation rather than lock in initial choices.

The framework favors explicit deliberation instead of analytical complexity. The examples suggest that the most consequential triad failures occur not when policymakers make the “wrong” choice, but when they make no conscious choice at all, defaulting to the preferences of the most powerful actors by failing to ask who benefits, who bears costs, and what alternatives were considered. Northern Virginia’s data center concentration is not the result of a deliberate decision that Progress was the region’s optimal triad position; it is the result of an absence of deliberate governance. The framework’s purpose is to prevent such defaults.

Table 7: Deliberate Triad Choice Framework: Illustrative Application

Framework Question	Progress Decision (expedited permitting)	Sustainability Decision (green energy mandate)
1. Which balance?	Prioritizes speed of deployment; subordinates community review processes.	Prioritizes environmental integration; accepts slower deployment and higher cost.
2. Who decides?	Industry-led advocacy; community input via compressed comment periods.	Multi-stakeholder process including environmental and community groups.
3. What trade-offs?	Reduced environmental review may miss cumulative impacts; community concerns deferred.	Deployment may lag competing jurisdictions; higher operating costs may deter some operators.
4. Distribution?	Benefits (jobs, tax revenue) concentrated near facilities; costs (grid strain) may be regionalized.	Higher-quality jobs; lower environmental burden on host communities; costs shared through higher energy prices.
5. Accountability?	Deployment metrics (Megawatts online, time-to-permit); grid stress monitoring.	PUE tracking; renewable energy share; community benefit delivery; periodic review.

Illustrative examples, not exhaustive.

4.4 Governance Design Implications

The preceding analysis does not support a universal policy recipe. Instead, the triad framework provides design principles: AI infrastructure decisions should make their trade-offs explicit, identify who bears the costs, and create mechanisms for revisiting decisions as conditions change. The following recommendations should therefore be read as governance design options and examples rather than exhaustively validated best practices.

For State and Local Policymakers

State and local policymakers should consider community benefit agreements where data center projects impose significant local burdens or struggle to secure public support, specifying local hiring, workforce

training, infrastructure investment, and environmental mitigation. Distributional trade-offs should be made explicit when communities host infrastructure whose benefits may accrue elsewhere, with frank acknowledgment of costs in speed and flexibility.

Policymakers should also consider energy and water impact disclosure requirements. Without transparent information about expected load, water use, backup generation, and ratepayer impacts, communities cannot meaningfully evaluate triad trade-offs.

Property tax abatements are a Progress-weighted instrument, reducing developer costs while shifting tax burdens onto existing residents and reducing host-community fiscal capacity for mitigation. Where a region weights Equity more heavily, abatements should be evaluated against that standard, with explicit accounting of the Progress gains traded off.

For Federal Policymakers

Federal policy could set minimum standards and fund shared infrastructure while preserving state and local capacity for policy learning. Avoid blanket preemption that forecloses regional experimentation across different physical constraints and community preferences.

Federal policy should also fund the AI Workforce Research Hub (or equivalent facility) at a level sufficient for sustained, longitudinal analysis of workforce outcomes in data center communities, generating the independent evidence base that policymakers currently lack.

Federal policy should also expand and adequately fund the National AI Research Resource (NAIRR), initially launched as a pilot in January 2024 and later endorsed for expansion in the AI Action Plan (NSF 2024; The White House 2025c), to create a public-option compute infrastructure that addresses equity at the access level, ensuring that AI capabilities are available for research and public-benefit applications regardless of commercial ability to pay.

For Industry

Firms should treat community benefit, grid contribution, and workforce development as an infrastructure strategy, not only as philanthropy. Voluntary commitments can help build local legitimacy, but without transparent metrics, they remain difficult to evaluate.

Similarly, industry players should consider investing in renewable energy and grid infrastructure. As data center regional energy demand scales, the industry's relationship to the energy system can evolve from consumer to investor. Companies that invest in energy generation and grid modernization will be better positioned to secure the reliable, affordable energy supply that AI infrastructure requires over the long term.

4.5 Advancing the Equity Frontier

Advancing equity is a deliberate choice, not a retroactive correction of Progress-oriented policy. Current policy processes tend to avoid making that choice explicit. Equity-oriented instruments, such as community benefit agreements, mandatory disclosures, ratepayer protection requirements, and public computing

access, carry real costs. Where significant local burdens are imposed, the distributional question should be answered deliberately rather than by default.

The equity frontier remains the least explored. This does not mean that Equity is secondary to Progress or Sustainability. However, AI infrastructure policy has so far been driven by geopolitical considerations, capital deployment, energy availability, and environmental constraints, while the distribution of benefits, burdens, and decision authority has received less systematic attention.

Requirements for community consultation, local hiring, ratepayer protection, or public computing access can slow deployment, increase developer obligations, or reduce flexibility. Rendering these trade-offs explicit is why Equity belongs in the triad as a distinct frontier rather than as a general moral preference layered on top of the other two.

Consider Virginia's grid crisis: data centers consuming 26 percent of Dominion Energy Virginia's load have driven the need for major generation investments. Environmental groups have supported renewable buildout to meet that demand as a sustainability measure. But the cost of that buildout is being allocated to residential ratepayers rather than to the large industrial users that created the demand. A policy that succeeds on the Sustainability axis (more renewable generation, lower carbon intensity) can simultaneously fail on the Equity axis (regressive cost allocation that increases electricity bills for low-income households). SB 253 in the 2026 Virginia General Assembly session was precisely an attempt to correct this and make the grid work done in the name of sustainability serve equity rather than undermine it. The two goals pointed to different policy designs and different legislative coalitions.

In the Gulf States case, deploying solar and desalinated water systems for data centers addresses Sustainability, reducing hydrocarbon dependence. However, their development relies on migrant labor operating under inequitable conditions by most recognized standards. Improving the sustainability of the energy supply does not automatically improve the equity of the labor conditions, and improving labor conditions does not improve the carbon footprint. These are separate policy problems requiring separate interventions.

Distributed compute models, such as Crusoe Energy's modular data centers, which use proprietary technology to convert otherwise flared natural gas at oil and gas sites into electricity for co-located computing (Crusoe Energy 2025), can be used to alter the geographic distribution of both the economic benefits and environmental trade-offs associated with AI infrastructure.

The National AI Research Resource (NAIRR) public-option compute infrastructure can ensure that access to AI computing capacity is not exclusively determined by ability to pay (NSF 2024).

Oregon's evolving approach to data center taxation illustrates how tax policy can channel infrastructure revenues toward community investment. Data centers in Oregon have received over \$330 million annually in property tax exemptions through the state's Enterprise Zone and Strategic Investment Programs, prompting legislative debate in the 2026 session over whether to condition or curtail these incentives (OPB 2026).

The theoretical framework establishes why equity belongs in the triad; the case examples illustrate how current approaches prioritize scale and sustainability over equitable distribution; and the emerging models identified above suggest possible mechanisms for testing equity-oriented governance. Future research should prioritize longitudinal analysis of communities that have hosted data center investments, tracking employment, income, displacement, energy costs, and fiscal impacts over time, to build the evidence base for equity-oriented governance.

5. Limitations

This article has five limitations.

First, the triad is presented as a heuristic operational framework, not as a validated measurement instrument. The classifications offered in the paper are interpretive and should be tested against additional cases, indicators, and stakeholder data.

Second, the public comment analysis is used as an input from prior work rather than newly coded evidence. Although the underlying dataset is large, it reflects a specific moment in U.S. AI policy debate and may not capture how stakeholder positions evolve as infrastructure constraints become more visible.

Third, the regional cases are illustrative rather than representative. Taiwan, Singapore, and the Gulf States are used to clarify variation in triad strategies, but each would require a fuller treatment to support stronger comparative claims. The U.S. cases of Northern Virginia and Texas, although representative of the current environment, do not intend to capture the full range of state and local variation.

Fourth, the paper does not establish causality. Claims about constraint-driven policy innovation should be read as an emerging hypothesis. Corporate commitments by Microsoft, OpenAI, and others may reflect several overlapping motives, including political branding, strategic risk management, investor signaling, community resistance, and anticipation of future regulation.

Fifth, because no region has yet demonstrated an equity-first AI infrastructure strategy at scale, the equity vertex is illustrated through emerging mechanisms rather than mature cases. This is a limitation of practice that constrains what the paper can demonstrate. Future research should examine employment, ratepayer impacts, fiscal outcomes, displacement, procedural inclusion, and access to computing in communities hosting AI infrastructure.

6. Conclusion

A large AI infrastructure buildout is underway, and it is proceeding mostly without the analytical frameworks that policymakers need to govern it. This is forcing regional policymakers to confront choices that the familiar “innovation versus regulation” frame cannot adequately describe. Data centers, beyond being capital assets, represent claims on electricity, land, water, labor, public authority, and community consent.

This paper has introduced the AI Infrastructure Triad as a conceptual framework for making those claims visible. Progress, Sustainability, and Equity identify three distinct but interdependent priorities: the speed and scale of deployment, the capacity of physical and ecological systems to sustain that deployment, and the distribution of benefits, burdens, and decision authority. The framework does not prescribe a single correct balance. Its purpose is to make trade-offs explicit before they are locked in by market power, regulatory inertia, or crisis response.

The evidence presented here is illustrative. Stakeholder worldviews, U.S. regional variation, and comparative cases suggest that different actors and regions already prioritize different triad frontiers. The emerging pattern of constraint-driven policy innovation should be treated as a hypothesis for further research, not as a settled causal finding.

Regions will make triad choices whether they acknowledge them or not. Our goal is to make these governance choices deliberate, transparent, and democratically contestable, rather than allowing them to default to the preferences of the most powerful actors. The triad framework offers a vocabulary for that

deliberation and a starting point for future empirical work. The most consequential governance failures will come from inattention, not from poor choice.

References

- Arab News. (2025, December 3). AI infrastructure investments mark new chapter for region. <https://www.arabnews.com/node/2624905>
- Autor, D. H., Dorn, D., & Hanson, G. H. (2013). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review*, 103(6), 2121–2168.
- Autor, D. H., Dorn, D., & Hanson, G. H. (2016). The China Shock: Learning from Labor-Market Adjustment to Large Changes in Trade. *Annual Review of Economics*, 8, 205-240.
- Block, F., & Keller, M. R. (2009). Where do innovations come from? Transformations in the U.S. economy, 1970-2006. *Socio-Economic Review*, 7(3), 459-483.
- Brookings Institution. (2025). Local Implications of Data Centers for Rural Communities in the U.S. <https://www.brookings.edu/articles/local-implications-data-centers-rural-communities-us/>
- Camus Energy. (2024). Why Does It Take So Long to Connect a Data Center to the Grid? <https://www.camus.energy/blog/why-does-it-take-so-long-to-connect-a-data-center-to-the-grid>
- Carvão, P. (2025, December 10). AI Limits: The Three Limits to Artificial Intelligence. *Forbes*. <https://www.forbes.com/sites/paulocarvao/2025/12/10/ai-limits-the-three-limits-to-artificial-intelligence/>
- Carvão, P., Ancheva, S., Atir, Y., Zhou, B., & Jeloka, S. (2025, February). Governance at a Crossroads: Artificial Intelligence and the Future of Innovation in America. M-RCBG Associate Working Paper No. 251.
- Carvão, P., Yashiro, M., & Jeloka, S. (2025, June). The People Have Spoken: The Tech Industry, Civil Society, and the U.S. Artificial Intelligence Action Plan. M-RCBG Associate Working Paper No. 261.
- Collison, P., & Cowen, T. (2019, July 30). We Need a New Science of Progress. *The Atlantic*.
- Crusoe Energy Systems. (2025). About Crusoe. <https://www.crusoe.ai/about/company>
- Data Center Frontier. (2026, January 15). Texas Senate Bill 6: A Bellwether On How States May Approach Data Center Energy Use. <https://www.datacenterfrontier.com/energy/article/55298872/texas-senate-bill-6-a-bellwether-on-how-states-may-approach-data-center-energy-use>
- Data Center Knowledge. (2025, March). UAE Data Centers: Powering the Middle East’s AI Revolution. <https://www.datacenterknowledge.com/build-design/uae-data-centers-powering-the-middle-east-s-ai-and-cloud-revolution>
- Deloitte and The Manufacturing Institute. (2024, April). Taking Charge: Manufacturers Support Growth with Active Workforce Strategies.
- Deutsche Bank. (2025, November). AI and U.S. Economic Growth: Data Center Spending Projections. https://prod1.www.dbresearch.com/PROD/RI-PROD/PROD0000000000610225/AI_101%3A_Economy%3A_Five_ways_AI_is_driving_growth.pdf
- Dominion Energy. (2024). 2024 Annual Report. Dominion Energy, Inc.
- Ernst, E. (2022). The AI trilemma: Saving the planet without ruining our jobs. *Frontiers in Artificial Intelligence*, 5, 886561. <https://doi.org/10.3389/frai.2022.886561>
- Feldman, M. P., & Kogler, D. F. (2010). Stylized Facts in the Geography of Innovation. In B. H. Hall & N. Rosenberg (Eds.), *Handbook of the Economics of Innovation* (Vol. 1, pp. 381–410). North-Holland.
- Florida, R. (2002). *The Rise of the Creative Class*. Basic Books.

Foreign Policy. (2025, July 2). U.S.-China AI Competition Needs Data Centers in UAE, Saudi Arabia. <https://foreignpolicy.com/2025/07/02/data-centers-us-uae-partnership-saudi-arabia-ai/>

Goldman Sachs. (2025, November 10). GS SUSTAIN: AI/Data Center Power Demand: The 6 Ps Driving Growth and Constraints. <https://www.goldmansachs.com/insights/goldman-sachs-research/data-center-power-demand-the-6-ps-driving-growth-and-constraints>

Heckt, S. (2026, March 6). General Assembly nixes bills that required data centers to get SCC certificate. *Virginia Mercury*. <https://virginiamercury.com/2026/03/06/general-assembly-nixes-bills-that-required-data-centers-to-get-scc-certificate/>

Human Rights Watch. (2023, December 18). Gulf States Treat Migrant Workforce as Disposable. <https://www.hrw.org/news/2023/12/18/gulf-states-treat-migrant-workforce-disposable>

International Energy Agency. (2025, April). Energy and AI. IEA Special Report. <https://www.iea.org/reports/energy-and-ai/>

IMDA (Infocomm Media Development Authority). (2024). *Green Data Centre Roadmap*. Singapore. <https://www.imda.gov.sg/how-we-can-help/green-dc-roadmap>

IMDA. (2025, December 1). Launch of Second Data Centre - Call for Application. <https://www.imda.gov.sg/resources/press-releases-factsheets-and-speeches/factsheets/2025/launch-of-second-data-centre>

Kaplan, J., McCandlish, S., Henighan, T., et al. (2020). Scaling Laws for Neural Language Models. arXiv:2001.08361.

King & Wood Mallesons. (2025, December). Singapore Launches 200MW Data Centre Call for Application (DC-CFA2). <https://www.kwm.com/global/en/insights/latest-thinking/singapore-launches-200mw-data-centre-call-for-application-dc-cfa2.html>

Kline, P., & Moretti, E. (2014). People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs. *Annual Review of Economics*, 6, 629–662.

Liu, F.H.M., Lai, K.P.Y., Seah, B., & Chow, W.T.L. (2025). Decarbonising digital infrastructure and urban sustainability in the case of data centres. *npj Urban Sustainability* 5, 15.

Main, I. (2026, February 20). Halfway through the 2026 legislative session, there's still no consensus on data center bills. *Virginia Mercury*. <https://virginiamercury.com/2026/02/20/halfway-through-the-2026-legislative-session-theres-still-no-consensus-on-data-center-bills/>

Mayer Brown. (2025, November 24). Important Texas Regulatory Updates for Data Centers. <https://www.mayerbrown.com/en/insights/publications/2025/07/important-texas-regulatory-updates-for-data-centers>

Microsoft. (2026, January 13). Building Community-First AI Infrastructure. Microsoft On the Issues Blog. <https://blogs.microsoft.com/on-the-issues/2026/01/13/community-first-ai-infrastructure/>

Middle East Institute. (2025). From Crude to Compute: Building the GCC AI Stack. <https://www.mei.edu/publications/crude-to-compute-building-gcc-ai-stack>

Moretti, E. (2012). *The New Geography of Jobs*. Houghton Mifflin Harcourt.

Mundell, R. A. (1963). Capital Mobility and Stabilization Policy under Fixed and Flexible Exchange Rates. *Canadian Journal of Economics and Political Science*, 29(4), 475–485.

National Science Foundation. (2024). National Artificial Intelligence Research Resource (NAIRR) Pilot. <https://new.nsf.gov/focus-areas/artificial-intelligence/nairr>

O'Donnell, J., & Crownhart, C. (2025, May 20). We Did the Math on AI's Energy Footprint. MIT Technology Review.

OPB. (2026, March 2). Data Centers Are Cut, for Now, from a Bill Expanding Oregon Tax Breaks. <https://www.opb.org/article/2026/03/02/data-centers-cut-bill-expanding-oregon-tax-breaks/>

OpenAI. (2026, January 20). Stargate Community Plans. <https://openai.com/index/stargate-community/>

Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.

Porter, M. E. (1990). *The Competitive Advantage of Nations*. Free Press.

Research and Markets. (2026). Singapore Data Center Market Investment Analysis Report 2026-2031. Via Yahoo Finance. <https://finance.yahoo.com/news/singapore-data-center-market-investment-102000477.html>

Rodrik, D. (2011). *The Globalization Paradox: Democracy and the Future of the World Economy*. W.W. Norton.

Saxenian, A. (1994). *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*. Harvard University Press.

SB Energy. (2026, January 9). OpenAI and SoftBank Group partner with SB Energy to build and operate next-generation AI data centers to advance Stargate [Press release]. <https://sbenergy.com/openai-and-softbank-group-partner-with-sb-energy/>

Shih, W. C., & Wang, J.-C. (2009). *Upgrading the Economy: Industrial Policy and Taiwan's Semiconductor Industry*. Harvard Business School Case 609-089. (Revised December 2010.)

Sutton, R. (2019). The Bitter Lesson. <http://www.incompleteideas.net/IncIdeas/BitterLesson.html>

The White House (2025a, January 23). Executive Order 14179, Removing Barriers to American Leadership in Artificial Intelligence.

The White House (2025b, December 11). Ensuring a National Policy Framework for Artificial Intelligence.

The White House (2025c, July). Winning the Race: America's AI Action Plan.

The White House (2026, March 4). Ratepayer Protection Pledge. <https://www.whitehouse.gov/articles/2026/03/ratepayer-protection-pledge/>

Wilson Center. (2024). Changing the Tide for the Gulf's Migrant Workers. <https://www.wilsoncenter.org/article/changing-tide-gulfs-migrant-workers>