

Governing the Compute Frontier: Rival Regimes for Artificial Intelligence Infrastructure and the Case for a Third Path

Kolawole Anthony Fayemi¹

¹ Partner, Energy, Telecommunications and Climate Change Practice, Ihuoma Iheanacho & Co., Lagos, Nigeria;
Independent Non-Executive Director, Vas2Nets Technologies Limited

Correspondence: Dr. Kolawole Anthony Fayemi, Ihuoma Iheanacho & Co., Lagos, Nigeria.

Email: ksagacity@gmail.com; Photizogreen.com.ng

Abstract

The governance of artificial intelligence is usually described as a single global project, yet the three powers that lead in the field have built three different regimes for governing the infrastructure on which artificial intelligence depends. This article characterizes those regimes and asks what each treats as the governable object, what all three systematically neglect, and what a viable alternative for states outside the triad would require. The method is a qualitative, comparative, and doctrinal analysis of primary policy instruments and authoritative secondary analysis across the United States, the European Union, and China, read against the compute-governance literature. The analysis finds that the United States governs through market-led construction coupled to security-driven export control, the European Union through a threshold-triggered law of general-purpose models and a standards apparatus, and China through the state-coordinated construction of national computing capacity. Each regime is optimized for its own strategic position, and each is blind to the same two things: the energy and climate layer on which all compute runs, and the agency of the states that will host much of the physical infrastructure. The article argues that a governable compute frontier requires a third path defined by four criteria: governing the physical and energy layer rather than only the model, defaulting to interoperability and transparency, treating compute access as a development question, and distributing governance authority beyond the three powers. The implications concern global policy, the competition among the powers, and the design of workable governance for the majority of states.

Keywords: ai governance, compute, digital infrastructure, geopolitics, energy policy, global south

DOI: 10.7176/ISDE/16-1-04

Publication date: September 30th 2026

1. Introduction

Artificial intelligence has become a matter of statecraft, and the contest is fought increasingly over physical things: the chips that train models, the data centres that house them, the cables that connect them, and the electricity that powers all of it. For a decade the governance conversation concentrated on the conduct of models and the protection of users. That conversation remains necessary. It is no longer sufficient. The scale of the infrastructure now involved has turned compute, connectivity, and energy into instruments of national power. The International Energy Agency projects that global data-centre electricity demand will rise from roughly 415 terawatt-hours in 2024 to roughly 945 terawatt-hours by 2030, with artificial intelligence workloads the largest single driver of that growth (International Energy Agency, 2025). Private and public commitments to build compute now run to the hundreds of billions of dollars in both the United States and China (Congressional Research Service, 2025; RAND Corporation, 2025). Infrastructure at that scale does not sit outside politics. It has become the medium through which technological advantage is pursued and contested.

The dominant way of speaking about this contest treats the governance of artificial intelligence as a single, if contested, global endeavour, a matter of aligning national approaches toward common rules. That framing obscures more than it reveals. The three powers that lead in artificial intelligence, meaning the United States, the European Union, and China, have not built one governance project with three national accents. They have built three distinct regimes, each with a different theory of what should be governed, a different instrument for governing it, and a different external projection into the rest of the world. Naming those regimes precisely, and holding them side by side, is the first task of this article, because a comparison at the level of regime rather than of individual rule exposes a pattern that a rule-by-rule reading misses.

The choice of the regime as the unit of analysis is deliberate. A regime, in the sense used here, is more than a collection of rules. It is the pairing of a selected object of governance with an instrument for governing it and a projection of that pairing beyond the jurisdiction that built it. Reading at the level of the regime rather than the individual rule has two advantages. It reveals the internal logic that makes a set of measures coherent, so that an export control and a permitting reform can be seen as two moves in a single strategy rather than as unrelated

policies. And it makes omissions legible, because a regime is defined as much by the object it declines to govern as by the object it selects. A rule-by-rule reading can catalogue what each power has done. A regime reading can explain why, and can show what each has left undone. The pattern this article identifies, a shared neglect of the energy layer and of the agency of other states, becomes visible only at the level of the regime.

The article pursues three research questions. The primary question asks what distinct regimes for governing artificial intelligence infrastructure the United States, the European Union, and China have constructed, and what each treats as the governable object. The second question asks what all three regimes systematically undergovern, and with what global consequences. A third and supporting question asks what criteria a viable third path of artificial-intelligence-infrastructure governance would need to meet in order to be adoptable by states outside the three powers and to close the gaps the three regimes leave. The questions are linked. A regime is defined in part by what it chooses not to govern, and the shared blind spots of the three regimes define the space that a fourth model would have to occupy.

The thesis of the article is that the three powers have built three regimes for governing artificial intelligence infrastructure, each optimized for its own strategic position, and that all three neglect the same two dimensions: the energy and climate layer on which compute runs, and the agency of the states outside the triad that will host a growing share of the physical infrastructure. A governable compute frontier therefore requires a third path, meaning a fourth model available to the majority of states, defined by four criteria. It must govern the physical and energy layer, and not the model alone. It must default to interoperability and transparency rather than to enclosure. It must treat access to compute as a development question as much as a security question. And it must distribute governance authority beyond the three powers, through regional coordination and multilateral standards. The article develops each claim and then draws out the criteria for the third path.

The scope is bounded. The article addresses the governance of physical artificial intelligence infrastructure, meaning compute, data centres, connectivity, and the energy systems that power them. It does not address model safety or user-layer data protection except where these intersect with infrastructure. The frame is global and comparative, and the argument is conceptual and doctrinal rather than empirical. Its contribution is a typology of the three regimes, a diagnosis of their shared neglect, and a normative account of what a fourth model must satisfy.

2. Materials Studied and Literature Review

The relevant scholarship falls into five bodies of work that have developed largely apart. Reading them together, and against one another, is part of the contribution of this article.

2.1 *Compute as a governable object*

The first body concerns compute as a governable object. A central line of argument holds that computing power is a uniquely governable input to artificial intelligence because it is detectable, excludable, and quantifiable in ways that data and algorithms are not (Sastry et al., 2024). This insight has produced a research programme on compute as a policy instrument, including the design and functions of training-compute thresholds (Heim & Koessler, 2024) and a mapping of the open technical problems in operationalizing such governance, from verification to enforcement (Reuel et al., 2024). The Center for Security and Emerging Technology situates compute within a wider account of data, algorithms, and compute as the triad of artificial intelligence power (Buchanan, 2020). This body establishes the premise on which all three regimes draw, namely that compute is where governance can grip. It is written largely from the vantage of the leading powers and concentrates on advanced chips and hyperscale facilities.

2.2 *The geopolitics of the physical layer*

The second body concerns the geopolitics of the physical layer. Scholarship and policy analysis here treat infrastructure as strategic terrain. Work on semiconductor supply chains frames advanced chips as a point of acute geopolitical vulnerability, given the concentration of fabrication in East Asia (RAND Corporation, 2024; Atlantic Council, 2024b). A parallel literature treats submarine cables, which carry the large majority of intercontinental data through a small number of chokepoints, as contested infrastructure whose security has become an explicit object of alliance policy (Atlantic Council, 2024a; Carnegie Endowment for International Peace, 2024a; New America, 2025b). This body establishes that control over infrastructure, and not merely over models, is now a dimension of national power. It says less about how that control is translated into domestic governance regimes.

2.3 The three national approaches

The third body concerns the three national approaches directly. On the United States, official analysis documents a federal effort to accelerate the construction of data centres and associated energy, and a security posture built on export controls over advanced chips and equipment (Congressional Research Service, 2025; The White House, 2025a; The White House, 2025b; Bipartisan Policy Center, 2025; Center for Security and Emerging Technology, 2024). On the European Union, the Artificial Intelligence Act ties additional obligations to general-purpose models trained above a defined floating-point threshold, and a parallel effort extends the same measuring instinct to the energy performance of data centres (European Union, 2024; Center for Democracy and Technology, 2024; White & Case, 2026). On China, analysis documents a state-coordinated buildout of national computing capacity, with state carriers central to the network and a policy of siting compute where energy is available (RAND Corporation, 2025; U.S.-China Economic and Security Review Commission, 2026). The Stanford Institute for Human-Centered AI documents the resulting rise in state activity and the salience of compute access and sovereignty across jurisdictions (Stanford Institute for Human-Centered AI, 2026). This body describes each approach well. It rarely holds the three side by side as regimes with a common structure and common omissions.

2.4 Energy and climate

The fourth body concerns energy and climate. Data centres and networks are large and growing electricity loads, and advanced economies have begun to measure and regulate that load, as in the European Union rating effort (International Energy Agency, 2025; International Energy Agency, 2026; White & Case, 2026). Advocacy and policy analysis warn that demand may outrun energy systems and climate commitments absent intervention (Opportunity Green, 2026; Carbon Brief, 2026). Sector analysis documents heavy reliance on diesel at telecommunications sites in constrained-grid markets (GSMA, 2024). This body is largely technical or advanced-economy facing, and it is rarely connected to the governance regimes described in the third body.

2.5 The majority of states and the divide

The fifth body concerns the majority of states and the artificial intelligence divide. The International Monetary Fund identifies access to compute and data infrastructure as a determinant of which countries capture the benefits of artificial intelligence, and describes a concentration of advanced compute in a small group of economies (International Monetary Fund, 2025). Policy institutions examine whether developing regions can move from a divide toward a dividend (Center for Strategic and International Studies, 2025), and chart South-South cooperation on compute as one route (Carnegie Endowment for International Peace, 2026). Continental and regional efforts, including the African Union Continental Artificial Intelligence Strategy and the Smart Africa coordination on a continental artificial intelligence council, show that states outside the triad are building governance ambition of their own (African Union, 2024; Smart Africa, 2025). Regional statements have made artificial intelligence infrastructure an explicit item of collective diplomacy (BRICS, 2025). This body documents dependence and emerging agency with force. It speaks less to how the three dominant regimes shape the room in which that agency operates.

2.6 A gap across the literatures

A tension runs across these bodies that is rarely made explicit. The compute-governance and geopolitics literatures read infrastructure primarily through a security frame, in which the governable object is capability and the animating concern is advantage over a rival. The energy and divide literatures read the same infrastructure through a development and sustainability frame, in which the governable object is access and the animating concern is who is included and at what environmental cost. The two frames rarely meet, and the instruments they imply pull in different directions, one toward control and the other toward diffusion. The regimes examined below are built almost entirely within the first frame. The third path this article proposes is an attempt to hold both frames at once, to treat infrastructure as at once a security object and a development object, which is why it cannot be assembled from any single one of these literatures.

Across these bodies a gap is visible. The compute-governance literature supplies the premise but writes for the powers. The geopolitics literature treats infrastructure as terrain without analysing domestic regimes. The national-approaches literature describes each regime in isolation. The energy literature is disconnected from governance. The divide literature documents the position of the majority without theorizing the third path. No body brings the five together to ask how the three regimes differ, what they share in neglect, and what a fourth model for everyone else would require. This article occupies that gap.

3. Methodology

This article uses a qualitative research design that combines comparative policy analysis and legal and regulatory doctrine. This design suits a question that is institutional and conceptual rather than statistical, and that turns on how governance regimes are structured and how they interact. The study is not empirical in the human-subjects sense. It has no participants, sample, or experimental manipulation, and it reports no statistical tests. The subsections below describe the qualitative and doctrinal approach rather than an empirical protocol.

The design is a structured, focused comparison of three governance regimes. It reconstructs each regime from primary instruments and authoritative analysis, and it compares them along a fixed set of dimensions: the governable object each regime selects, the primary instrument it uses, its treatment of the energy layer, its external projection into third states, and its principal blind spot. A structured comparison permits a like-for-like reading of regimes that differ in legal form, and it makes the shared omissions visible in a way that a single-country study cannot. The dimensions are held constant across the three cases so that the comparison does not smuggle in the priorities of any one regime. Each regime is described in its own terms first, using its own instruments and its own stated objects, and only then read against the common dimensions, so that the typology emerges from the cases rather than being imposed on them. This two-step procedure, description in the regime's own terms followed by comparison on fixed dimensions, is the method by which the shared blind spots are identified without prejudging them.

The corpus is documentary. It comprises primary policy instruments and official actions, including United States executive actions and export-control measures, the European Union Artificial Intelligence Act and its energy-efficiency initiatives, and Chinese industrial-policy measures as documented in authoritative secondary analysis; reports from international institutions such as the International Energy Agency and the International Monetary Fund; analysis from established policy research organizations; and peer-reviewed and preprint scholarship on compute governance. Where an institutional document does not state individual authorship, it is cited by its publishing organization, consistent with accepted practice for grey literature.

The framework combines two lenses. The first is the compute-governance lens, which treats physical inputs to artificial intelligence as governable objects (Sastry et al., 2024). The second is a regime lens, which treats a governance regime as the pairing of a selected object with an instrument and an external projection, and which reads a regime as much by what it excludes as by what it governs. The article applies the first lens to identify what each power governs, and the second to identify the structure and the omissions of each regime.

The three cases are selected because they are the dominant regimes shaping the global compute frontier, and because together they set the terms within which other states act. The comparison is completed by reference to the position of states outside the triad, illustrated by continental coordination in Africa, in order to ground the account of the third path. The selection is purposive rather than representative, and the findings are used to build a typology and a normative argument rather than to establish statistical generality.

The analysis is interpretive and comparative, and it does not measure the magnitude of the effects it describes. It rests on secondary and grey literature, which carries the biases of its sources, and on regimes whose instruments are in motion, so specific provisions will evolve. The account of the third path is normative and does not claim that any state has yet built it. These are boundaries on the claims rather than defects in them, and the findings are framed accordingly.

4. Results and Findings

The findings are organized under five subsections. The first three characterize the regimes. The fourth builds a typology and names the shared blind spots. The fifth derives the criteria for a third path.

4.1 *The United States regime: market construction plus security control*

The United States governs artificial intelligence infrastructure through a pairing of market-led domestic construction and security-driven external control. Domestically, federal action has moved to accelerate the construction of large data centres and their energy supply, including measures to speed permitting for facilities that draw more than one hundred megawatts and prominent private commitments to build at scale (The White House, 2025a; The White House, 2025b; Congressional Research Service, 2025; Bipartisan Policy Center, 2025). The governable object here is capacity, and the instrument is the removal of friction from its construction. Externally, the United States has restricted the export of advanced chips and chipmaking equipment, using performance thresholds to define the controlled items and extending the reach to lithography systems, an approach widely read as the beginning of a control regime over the most capable compute (Center for Security and Emerging Technology, 2024; Atlantic Council, 2024b). The governable object here is the frontier of

capability held by a rival, and the instrument is denial. The regime is coherent on its own terms. It builds at home and it slows a competitor abroad. Its external projection into third states is conditional access: to the cloud services and chips of United States firms, on terms that carry the security logic with them.

The internal tension of the United States regime is worth naming, because it shapes what the regime exports. Construction and control pull in opposite directions. The drive to build at speed favours abundant, cheap energy and light permitting, while the drive to deny a rival favours tight control over the movement of chips and tools. The first logic is expansionary and largely indifferent to where compute ends up, provided it is built. The second is restrictive and intensely concerned with where the most capable compute ends up. For third states the two logics arrive together and can conflict. A government courting United States hyperscale investment receives the construction logic as opportunity and the control logic as constraint, and it discovers that access to the frontier of capability is conditioned on alignment with a security agenda it did not set. The regime is coherent for the United States. It is bifurcated for everyone else.

4.2 The European Union regime: threshold law plus standards

The European Union governs through law and standards rather than through construction or denial. Its Artificial Intelligence Act designates general-purpose models trained above ten to the power of twenty-five floating-point operations as carrying systemic risk, which makes compute a legal trigger for heightened obligations (European Union, 2024; Center for Democracy and Technology, 2024). The governable object here is the model, identified by the compute used to train it, and the instrument is binding law with extraterritorial reach through the market. The European Union has extended the same measuring instinct to physical infrastructure, advancing energy-efficiency requirements and a common rating scheme for data centres (White & Case, 2026). The regime is distinctive in treating a quantity of compute as a legal fact with consequences, and in exporting that logic through the size of its market rather than through construction or control. Its external projection into third states is a template: a set of standards that other jurisdictions can adopt or adapt, and a source of normative influence that competes with the other two regimes.

The European regime carries an internal tension of its own. By fixing the trigger for heightened obligation to a quantity of training compute, it governs the model through a proxy that is precise, auditable, and administrable, which is the source of its strength. Yet the same choice narrows the regime's field of vision. A threshold defined by training compute directs attention to the largest models and away from the infrastructure that trains and serves them, so that the physical layer enters European governance mainly through the side door of energy-efficiency ratings rather than as a governed object in its own right. The regime is powerful precisely where it looks, at the model, and thin precisely where it does not, at the machine room. Its export as a template carries that emphasis with it, which means that a jurisdiction adopting the European model inherits both its administrable clarity and its structural inattention to infrastructure.

4.3 The China regime: state construction of capacity

China governs through the state-coordinated construction of national computing capacity. Its industrial policy directs the building of computing hubs in provinces with available energy, positions state carriers at the centre of the network, and treats compute as national infrastructure to be planned rather than left to the market, with public analysis describing spending on the order of hundreds of billions of dollars over several years and a target of raising national capacity substantially (RAND Corporation, 2025; U.S.-China Economic and Security Review Commission, 2026). The governable object here is the national stock of compute, and the instrument is the state itself, acting through carriers, subsidies, and planning. Its external projection into third states is infrastructure on offer, financed and built, and an alternative to dependence on the incumbents.

The Chinese regime resolves the tension that divides the United States approach, because construction and control are unified in a single actor. The state that builds is the state that governs access, and it can plan capacity, direct siting, and condition participation through the same institutions. That unification is the regime's strength and the source of its appeal to third states, which are offered not conditional access or a template but finished infrastructure. It is also the regime's characteristic risk for a host. Infrastructure that is financed and built by a state actor arrives with dependencies of its own, in equipment, in operation, and in the standards embedded in the build, so that the alternative to dependence on the incumbents can become a different dependence rather than autonomy. The three regimes thus present a host with three distinct bargains, and none of the three is free.

The three regimes therefore select three different objects, capacity, the model, and the national stock, and deploy three different instruments, deregulation, law, and the state.

4.4 A typology and its blind spots

Table 1 sets the three regimes side by side. Reading across the rows, the differences are real and the powers compete on them. Reading down the final two columns, a pattern emerges that a single-country study would miss. None of the three regimes governs the energy and climate layer as a first-order matter. The United States accelerates energy supply for construction but does not condition capacity on its energy intensity. The European Union measures data-centre energy but treats it as an efficiency question adjacent to the main event of model regulation. China sites compute where energy is available but governs the buildout as an industrial rather than a climate matter. The demand figures make the omission consequential: data-centre electricity demand is set to more than double by 2030, and the growth strains grids and climate commitments alike (International Energy Agency, 2025; International Energy Agency, 2026; Opportunity Green, 2026). None of the three regimes, further, treats the agency of third states as a design concern. Each projects outward, through conditional access, template, or financed build, and each positions the rest of the world as a recipient of a regime made elsewhere. The shared blind spots are the energy layer and the agency of the majority.

Table 1. Three regimes for governing artificial intelligence infrastructure

Dimension	United States	European Union	China
Governable object	Domestic capacity; the rival's capability frontier	The model, identified by training compute	The national stock of compute
Primary instrument	Deregulated construction plus export control	Binding threshold law plus standards	State-coordinated construction
Energy layer	Supply accelerated, not conditioned	Measured as efficiency, adjacent to model rule	Sited near energy, governed as industry
External projection	Conditional access to firms and chips	Exportable template of standards	Financed and built infrastructure
Principal blind spot	Energy conditioning; third-state agency	Infrastructure and energy as first-order; third-state agency	Climate governance; third-state agency

Note. Compiled by the author from the cited instruments and analyses (Congressional Research Service, 2025; European Union, 2024; RAND Corporation, 2025; White & Case, 2026; International Energy Agency, 2025). The final column identifies the dimension each regime governs least.

Two clarifications sharpen the finding. The shared neglect of the energy layer is not a claim that the three powers ignore energy. Each attends to energy, but each attends to it as a means to another end: the United States as fuel for construction, the European Union as a metric of efficiency, and China as a siting input for industrial planning. In none of the three is the energy intensity of artificial intelligence infrastructure a governed object with its own binding standard and its own enforcement, which is what governing the layer would require. The shared neglect of third-state agency is likewise not a claim that the powers ignore other states. Each engages them constantly, as markets, as sites, and as adopters. The neglect is of their agency, meaning their capacity to set terms rather than receive them. A regime designed around conditional access, a template, or a financed build positions the third state as a taker in each case, and none of the three regimes treats the third state as a co-author of the terms. These two clarifications are what make the blind spots structural rather than incidental, and what make a fourth model necessary rather than merely desirable.

4.5 The criteria for a third path

The shared blind spots define the space a fourth model would occupy. Four criteria follow. First, a third path must govern the physical and energy layer, and not the model alone. In markets where the grid is constrained, the energy question precedes the artificial intelligence question, because what powers the infrastructure, at what cost, and with what emissions is settled before any model runs (GSMA, 2024; International Energy Agency, 2025). Energy-intensity reporting for data centres and networks, of the kind the European Union is adopting for facilities, is the concrete instrument, calibrated to constrained grids (White & Case, 2026). Second, a third path must default to interoperability and transparency rather than to enclosure, so that access to compute is not locked behind a single provider or a single standard. Third, it must treat access to compute as a development question, because the concentration of advanced compute is a determinant of who captures the benefits of artificial intelligence, and a security-only framing leaves the majority of states outside the calculation (International Monetary Fund, 2025; Center for Strategic and International Studies, 2025). Fourth, it must distribute governance authority beyond the three powers, through regional coordination and multilateral standards, so that host states set terms collectively rather than accept them individually (African Union, 2024; Smart Africa, 2025; Carnegie Endowment for International Peace, 2026; BRICS, 2025). These four criteria are the article's central constructive claim, and they are derived from the regimes rather than imposed on them.

Each criterion implies a concrete instrument and meets a foreseeable objection. Governing the energy layer implies mandatory energy-intensity reporting for data centres and large network operators, enforced by the sector regulators that already license them, with siting and grid-connection conditions attached to approval. The objection is that such conditions deter investment. The answer is that a constrained grid already imposes the cost through diesel, and that a clear reporting standard reduces uncertainty for serious investors while screening out the least efficient. Defaulting to interoperability implies procurement rules that favour open standards and the portability of data and workloads across providers. The objection is that the leading providers set the standards. The answer is that coordinated demand, aggregated across a region, is large enough to move terms that a single small market cannot. Treating compute as a development question implies public and pooled compute capacity, shared research access, and concessional financing for infrastructure that serves broad public purposes rather than a single firm. The objection is cost. The answer is that the cost of exclusion, measured as forgone participation in the artificial intelligence economy, is larger and is already being paid (International Monetary Fund, 2025). Distributing authority implies regional standard-setting bodies and coordinated negotiation with the three powers. The objection is that coordination is fragile and states defect. The answer is that the binding trade and standards institutions many states already belong to can raise the cost of defection, and that a shared standard is easier to hold than a set of separate bargains. The third path is therefore not an aspiration without instruments. It is a set of four instruments, each already visible in embryo, that the majority of states could adopt through bodies that already exist.

5. Discussion

5.1 Answering the research questions

The findings answer the article's questions. The first question asked what regimes the three powers have built and what each governs. The answer is that they have built three structurally different regimes, selecting capacity, the model, and the national stock as their objects, and deploying deregulation, law, and the state as their instruments. The second question asked what all three neglect. The answer is that they share two blind spots: the energy and climate layer on which all compute runs, and the agency of the states outside the triad. The third question asked what a third path requires. The answer is the four criteria derived above.

5.2 Situating the findings in the literature

These findings extend the existing literature rather than restating it. The compute-governance literature establishes that compute is governable but does not compare the regimes that have grown up around that premise (Sastry et al., 2024; Reuel et al., 2024). This article supplies the comparison and reads the regimes by their omissions. The national-approaches literature describes each regime in isolation (Congressional Research Service, 2025; European Union, 2024; RAND Corporation, 2025). This article holds them side by side and finds a common neglect that isolation conceals. The energy literature and the divide literature each document a problem that the governance literature does not address (International Energy Agency, 2025; International Monetary Fund, 2025). This article shows that the two problems are the same blind spot viewed from two angles, and that a fourth model must address both at once.

5.3 The stakes for host states

The geopolitical implication is that the competition among the three powers, precisely because each regime projects outward, widens the room for a capable host to set terms. A state or a region that governs its infrastructure competently, and that coordinates with its neighbours, can weigh conditional access from the United States against the European template against a financed build linked to China, and can attach its own conditions to any of them. The competition among the regimes is therefore more than a threat to the majority of states. For those with the capacity to use it, the competition is also a source of bargaining power. The risk is the mirror image. Absent coordination, host states compete against one another by relaxing energy and reporting requirements, which would erode the very standards that make governance meaningful and hand the advantage to whichever power is most willing to build without conditions.

5.4 Objections and replies

Three objections deserve a direct answer. The first holds that the three approaches are converging, so that a typology of distinct regimes freezes a moment that is already passing. The reply is that the objects and instruments are structurally different and are unlikely to converge, because each is anchored in a different strategic position: the United States as the frontier incumbent, the European Union as the regulatory power without a frontier industry of its own, and China as the state builder. Surface borrowing, such as each power's growing attention to energy, does not dissolve the underlying differences the typology captures. The second objection holds that the third path underestimates the coercive power of the three regimes, and that host states

have little real room to set terms. The reply is that the room is unevenly distributed and depends on capacity, which is precisely why this article treats capacity and coordination as the binding constraints rather than assuming them away. The third objection holds that energy is a transitional problem that efficiency gains will resolve, so that governing the energy layer is a solution in search of a problem. The reply is that efficiency improvements are real but that they redistribute the burden rather than removing it, because the questions of who bears the siting cost, who is on diesel, and who is required to report and reduce intensity are not answered by a falling energy cost per task (International Energy Agency, 2025; Opportunity Green, 2026). Meeting these objections strengthens the argument, because each points to a condition, namely capacity, coordination, and enforcement, that the third path must satisfy to work.

5.5 Implications for the three powers

The argument also carries implications for the three powers, and not for the states outside the triad alone. A regime that neglects the energy layer exports that neglect, and the aggregate result is a global build-out whose climate cost is governed by no one, which is a risk the powers bear along with everyone else (International Energy Agency, 2026; Carbon Brief, 2026). A regime that treats third states only as takers invites the coordinated response the third path describes, and a bloc of host states that sets common terms changes the environment in which the powers compete. The powers have an interest, then, in the emergence of a fourth model, because a governed compute frontier is more stable than an ungoverned one, and because standards that are adopted willingly are cheaper to sustain than terms that are imposed. The third path is usually framed as a demand of the majority against the powers. It can also be read as a shared interest in a frontier that does not outrun the capacity of any actor, including the powers, to govern it.

5.6 Limitations and future research

The analysis has limits. It is interpretive and comparative, and it does not measure the magnitude of the effects it describes. It rests on regimes whose instruments are in motion, and its account of the third path is normative rather than a description of any existing model. The four criteria depend on regulatory capacity and political will that many states lack, and the article does not resolve how that capacity is to be built and financed. These boundaries point to the questions that follow. Future work should test the typology against additional powers and against the middle powers now building compute, examine whether the third-path criteria can be operationalized through existing trade and standards bodies, and study the conditions under which regional coordination holds against the incentive to defect.

6. Conclusion and Recommendations

6.1 Summary of key findings

This article asked what regimes the United States, the European Union, and China have built to govern artificial intelligence infrastructure, what they neglect, and what an alternative for the majority of states would require. The central findings are three. The three powers have built three structurally different regimes, selecting capacity, the model, and the national stock as their governable objects. All three share two blind spots: the energy and climate layer, and the agency of states outside the triad. A viable third path is defined by four criteria: govern the physical and energy layer, default to interoperability and transparency, treat compute access as a development question, and distribute governance authority beyond the three powers.

6.2 Practical and Policy Recommendations

The practical contribution of this article is a set of instruments that named actors can adopt with powers they already hold. The recommendations move outward from the national regulator to the region and then to the multilateral order, and they are written to be implemented rather than merely endorsed.

At the national level, the fastest lever sits with sector regulators rather than with a future artificial intelligence statute. A communications or utilities regulator can issue an energy-intensity reporting rule that requires operators of data centres and large networks to measure and disclose energy use per unit of compute and per site on a fixed reporting cycle, and it can attach siting and grid-connection conditions to the approvals those facilities already need. The instruments are administrative and familiar: a licence condition, a reporting template, and an enforcement schedule. They do not wait on legislation, they screen out the least efficient operators, and they give serious investors the certainty of a known standard. In a constrained-grid market the same rule should recognize captive and mini-grid generation and reward a move away from diesel, so that the reporting obligation pulls investment toward cleaner and cheaper power rather than documenting the status quo. Nigeria's telecommunications and electricity regulators, which already license operators and govern grid connection, are a working example of where this authority sits.

At the regional level, the task is to convert separate national rules into a common standard, so that no single member competes for infrastructure by undercutting the others. A regional standards body or economic community can adopt a model energy-efficiency and reporting standard for data centres and networks, drawing on the European rating approach but calibrated to weaker grids (White & Case, 2026), and can annex energy and siting expectations to a binding trade instrument where one exists, which gives the standard the force that a strategy document lacks. Where a continental artificial intelligence body or a digital-trade protocol is already in motion, the practical step is to write the reporting standard into its implementation rather than to negotiate a new instrument from the beginning.

For industry and for the practitioners who advise it, the recommendation is to design for reporting from the outset. Operators that instrument their facilities for energy-intensity disclosure, publish the methodology, and help draft the standards they will be measured against convert a compliance burden into a competitive signal, because a credible efficiency record is now precisely what a host government and a climate-conscious client look for. For legal advisers, the reporting rule and the siting condition form the core of an emerging compliance practice at the intersection of telecommunications, energy, and technology regulation, and this article maps where that practice sits and what instruments it will turn on.

At the multilateral level, the standards and trade institutions that most states already belong to should carry energy-intensity and interoperability expectations for artificial intelligence infrastructure, so that the criteria acquire a forum rather than depending on the goodwill of the three powers (Carnegie Endowment for International Peace, 2026; BRICS, 2025). Because every instrument above presupposes the capacity to operate it, development finance institutions should fund the regulatory and technical capacity that reporting, open-standard procurement, and coordinated negotiation require. Capacity is itself a governable objective, and funding it is what moves the third path from a normative proposal toward an implementable programme. Coalitions of host states should coordinate procurement and compute access, so that they negotiate the terms of the three regimes from a collective position rather than singly.

These recommendations rest on a contribution to knowledge as well as to practice. The article offers a typology of the three dominant regimes and a method of reading a regime by what it declines to govern, and it joins the compute-governance frame to the ordinary law of infrastructure and energy regulation and to the theory of regional coordination. In doing so it treats the states outside the triad as agents with instruments rather than as terrain acted upon, which corrects an emphasis in both the compute-governance and the divide literatures.

The proposal is global by design. It derives, from the shared neglect of the three regimes, a model available to the majority of states rather than to the powers alone, and it reframes the constrained grid, often read as a deficit, as a variable that raises the bargaining position of a host that governs energy and siting well. The same logic applies to the Gulf states financing large builds, to South and Southeast Asian economies weighing cloud dependence against domestic capacity, and to Latin American states negotiating the siting of hyperscale facilities. The infrastructure of artificial intelligence is being built now, and the three powers are setting the terms on which the rest of the world will obtain it. The competition among their regimes will not wait for a fourth model to be designed. The task for the majority of states is to build that model deliberately, through the regulators, regional bodies, and coalitions named above, and the window is the current build-out. The compute frontier will be built either way. Whether it is governed, and by whom, is the choice in front of the world now.

6.3 Suggestions and Insights for Future Research

Three lines of inquiry follow from this analysis, and each also marks a limit of the present study. First, the typology should be tested against additional actors, including the middle powers now building compute, to establish whether the three-regime structure holds or whether a fourth regime type is emerging. The present article is confined to the three dominant powers, and that boundary is also its first invitation to further work. Second, future research should examine whether the third-path criteria can be operationalized through the specific trade and standards bodies that host states already belong to, and should document early attempts to write energy and siting expectations into binding instruments. This article proposes the criteria but does not track their adoption, because the instruments are only now in motion. Third, work should study the conditions under which regional coordination holds against the incentive to defect, since the force of the recommendations depends on regulatory and coordinating capacity that varies widely and that this article does not resolve how to finance and build. These are boundaries on the present claims and, at the same time, the agenda they set for the work that should follow.

References

- African Union. (2024). *Continental Artificial Intelligence Strategy: Harnessing AI for Africa's development*. https://au.int/sites/default/files/documents/44004-doc-EN-ContinentalAIStrategyJuly_2024.pdf
- Atlantic Council. (2024a). *Cyber defense across the ocean floor: The geopolitics of submarine cable security*. <https://www.atlanticcouncil.org/in-depth-research-reports/report/cyber-defense-across-the-ocean-floor-the-geopolitics-of-submarine-cable-security/>
- Atlantic Council. (2024b). *United States-China semiconductor standoff: A supply chain under stress*. <https://www.atlanticcouncil.org/in-depth-research-reports/issue-brief/united-states-china-semiconductor-standoff-a-supply-chain-under-stress/>
- Bipartisan Policy Center. (2025). *Strategic federal actions aim to strengthen AI and energy infrastructure*. <https://bipartisanpolicy.org/explainer/strategic-federal-actions-aim-to-strengthen-ai-and-energy-infrastructure/>
- BRICS. (2025). *BRICS leaders' statement on the global governance of artificial intelligence*. <http://www.brics.utoronto.ca/docs/250706-ai.html>
- Buchanan, B. (2020). *The AI triad and what it means for national security strategy*. Center for Security and Emerging Technology. <https://cset.georgetown.edu/publication/the-ai-triad-and-what-it-means-for-national-security-strategy/>
- Carbon Brief. (2026). *AI: Five charts that put data-centre energy use and emissions into context*. <https://www.carbonbrief.org/ai-five-charts-that-put-data-centre-energy-use-and-emissions-into-context>
- Carnegie Endowment for International Peace. (2024a). *Subsea communication cables in Southeast Asia: A comprehensive approach is needed*. <https://carnegieendowment.org/india/research/2024/12/southeast-asia-undersea-subsea-cables>
- Carnegie Endowment for International Peace. (2026). *South-South AI collaboration: Advancing practical pathways*. <https://carnegieendowment.org/research/2026/02/south-south-ai-collaboration-advancing-practical-pathways>
- Center for Democracy and Technology. (2024). *EU AI Act brief, part 5: General-purpose AI models*. <https://cdt.org/insights/eu-ai-act-brief-pt-5-general-purpose-ai-models/>
- Center for Security and Emerging Technology. (2024). *The National Security Memorandum on Artificial Intelligence: CSET experts react*. <https://cset.georgetown.edu/article/the-national-security-memorandum-on-artificial-intelligence-cset-experts-react/>
- Center for Strategic and International Studies. (2025). *From divide to delivery: How AI can serve the Global South*. <https://www.csis.org/analysis/divide-delivery-how-ai-can-serve-global-south>
- Congressional Research Service. (2025). *Data centers and cloud computing: Information technology infrastructure for artificial intelligence* (IF12899). <https://www.congress.gov/crs-product/IF12899>
- European Union. (2024). *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), Articles 51 and 55*. <https://artificialintelligenceact.eu/article/51/>
- GSMA. (2024). *Tower power Africa: Energy challenges and opportunities*. <https://www.gsma.com/solutions-and-impact/connectivity-for-good/mobile-for-development/uncategorized/tower-power-africa-energy-challenges-and-opportunities/>
- Heim, L., & Koessler, L. (2024). *Training compute thresholds: Features and functions in AI regulation* (arXiv:2405.10799). <https://arxiv.org/abs/2405.10799>
- International Energy Agency. (2025). *Energy and AI*. <https://www.iea.org/reports/energy-and-ai>
- International Energy Agency. (2026). *Electricity 2026: Analysis and forecast*. <https://www.iea.org/reports/electricity-2026>
- International Monetary Fund. (2025). *The global impact of AI: Mind the gap* (Working Paper No. WP/25/76). <https://www.imf.org/-/media/files/publications/wp/2025/english/wpica2025076-print-pdf.pdf>
- New America. (2025b). *Securing the backbone of AI*. <https://www.newamerica.org/insights/securing-the-backbone-of-ai/introduction/>
- Opportunity Green. (2026). *Data centres: How soaring demand threatens to overwhelm energy systems and climate goals*. <https://opportunitygreen.org/>

- RAND Corporation. (2024). *Supply chain interdependence and geopolitical vulnerability: The case of Taiwan and high-end semiconductors* (RRA2354-1). https://www.rand.org/pubs/research_reports/RRA2354-1.html
- RAND Corporation. (2025). *Full stack: China's evolving industrial policy for AI* (PEA4012-1). <https://www.rand.org/pubs/perspectives/PEA4012-1.html>
- Reuel, A., Bucknall, B., Casper, S., Fist, T., Soder, L., Aarne, O., et al. (2024). *Open problems in technical AI governance* (arXiv:2407.14981). <https://arxiv.org/abs/2407.14981>
- Sastry, G., Heim, L., Belfield, H., Anderljung, M., Brundage, M., Hazell, J., et al. (2024). *Computing power and the governance of artificial intelligence* (arXiv:2402.08797). <https://arxiv.org/abs/2402.08797>
- Smart Africa. (2025). *Smart Africa Steering Committee convenes in Kigali and endorses the establishment of the Africa Artificial Intelligence Council*. <https://smartafrica.org/smart-africa-steering-committee-convenes-in-kigali-and-endorses-the-establishment-of-the-africa-artificial-intelligence-council/>
- Stanford Institute for Human-Centered AI. (2026). *The 2026 AI Index report: Policy and governance*. <https://hai.stanford.edu/ai-index/2026-ai-index-report/policy-and-governance>
- The White House. (2025a). *Executive Order 14141: Advancing United States leadership in artificial intelligence infrastructure*. <https://www.federalregister.gov/documents/2025/01/17/2025-01395/advancing-united-states-leadership-in-artificial-intelligence-infrastructure>
- The White House. (2025b). *Accelerating federal permitting of data center infrastructure*. <https://www.whitehouse.gov/presidential-actions/2025/07/accelerating-federal-permitting-of-data-center-infrastructure/>
- U.S.-China Economic and Security Review Commission. (2026). *Two loops: How China's open AI strategy reinforces its industrial dominance*. <https://www.uscc.gov/>
- White & Case. (2026). *Data centres and energy consumption: Evolving EU regulatory framework and outlook for 2026*. <https://www.whitecase.com/insight-alert/data-centres-and-energy-consumption-evolving-eu-regulatory-landscape-and-outlook-2026>