

Compute as Critical Infrastructure: A Governance Framework for Protecting and Regulating the Physical Layer of Artificial Intelligence

Kolawole A. 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 physical layer of artificial intelligence, comprising data centres, submarine and terrestrial fibre, grid interconnections, and the telecommunications networks that bind them, has become materially and strategically critical, yet it sits inside a governance gap. Security law treats parts of it as critical infrastructure to be protected, artificial intelligence law treats a slice of it as compute to be regulated, and energy law treats another slice as load to be permitted, but no jurisdiction governs the physical layer as a coherent object. This article develops a governance framework that closes the gap. It identifies three governance framings now competing over the same infrastructure, namely national-security protection, compute regulation, and energy and environmental permitting, and argues that each is partial on its own. Drawing on the drone strikes against Amazon Web Services facilities in the United Arab Emirates and Bahrain in March 2026, the Baltic and Taiwan submarine-cable disruptions of 2024 and 2025, and the emerging regulation of data-centre electrical load, the article proposes a layered designation, protection, and regulation framework built on four criteria for a defensible designation of compute as critical infrastructure: essentiality, cascading interdependence, contestation, and governability. It uses Nigeria's Designation and Protection of Critical National Information Infrastructure Order 2024 as a worked instance of a developing-market state that has already designated telecommunications and data-centre assets as critical national infrastructure under the Office of the National Security Adviser, and reads that instrument against practice in the United States and the European Union. The analysis relies solely on publicly available information. The contribution is a portable governance architecture that treats the physical layer as a single regulated object while guarding against the over-securitisation of civilian infrastructure.

Keywords: critical infrastructure, ai governance, data centres, compute, telecommunications, national security

DOI: 10.7176/JETP/16-1-04

Publication date: September 28th 2026

1. Introduction

Artificial intelligence is commonly discussed as software, as models and their outputs, and the governance conversation has followed that framing toward the end-user layer of privacy, bias, transparency, and content. That attention is warranted and the field is well populated. It is also incomplete. Every model runs on a physical substrate of processors housed in data centres, fed by electricity from strained grids, and connected by fibre-optic cables that cross oceans and national borders. This physical layer has become a strategic asset, and in 2026 it became a target. On 1 March 2026, drones struck three Amazon Web Services facilities across the United Arab Emirates and Bahrain, causing structural damage, disrupting power delivery, and degrading more than one hundred services in the company's Middle East region (DefenseScoop, 2026; Marketplace, 2026). For the first time in modern conflict, commercial hyperscale data centres were hit as explicit objects of a military campaign. The event exposed a governance question that the ethics conversation had left unattended: how should states and the international community protect and regulate the physical infrastructure on which artificial intelligence depends.

The March strikes did not arrive without warning. Between November 2024 and early 2025, at least a dozen distinct submarine-cable damages were recorded in the Baltic Sea and around Taiwan, including the severing of the C-Lion1 and BCS East-West Interlink cables and repeated disruptions to Taiwan's external connectivity (Recorded Future, 2025; Al Jazeera, 2025). The North Atlantic Treaty Organization launched a dedicated Baltic surveillance operation in response. These incidents share a feature with the data-centre strikes: they concern the material carriers of digital services rather than the data or the algorithms, and they reveal that the layer beneath the model is contested terrain. The strategic literature has begun to register the shift, describing cloud regions,

data corridors, and cable landing stations in the same terms once reserved for shipping lanes and pipelines (World Economic Forum, 2026a; New America, 2025b).

The stakes are clearest through interdependence. Telecommunications is not one critical sector standing beside the others; it is the enabling layer on which the others increasingly run. Energy systems depend on it for grid telemetry and control, financial markets for clearing and settlement, hospitals for records and remote diagnostics, transport for signalling and logistics, and governments for the delivery of public services and emergency response. As artificial intelligence is embedded into these sectors, their dependence on the underlying compute and connectivity layer deepens, and a disruption at that layer cascades outward. The March strikes made the cascade visible when the degradation of a single cloud region interrupted banking and payment applications used by ordinary customers far from any theatre of conflict. The physical layer of artificial intelligence is therefore infrastructure in the fullest sense, meaning a substrate whose failure is felt across the economy and the state, and its governance is a matter of public interest rather than of private operational concern.

The governance response has been fragmented. National-security and critical-infrastructure regimes protect designated assets from physical and cyber threats. The United States organises its Cybersecurity and Infrastructure Security Agency around sixteen critical-infrastructure sectors, of which communications and information technology are two. The European Union has built the Critical Entities Resilience Directive and the second Network and Information Security Directive to raise the resilience of essential entities. Separately, artificial intelligence law has begun to regulate compute directly. The European Union's Artificial Intelligence Act ties obligations for general-purpose models to a training-compute threshold, making a quantity of computation a regulated legal fact (European Union, 2024; Center for Democracy and Technology, 2024). Separately again, energy and environmental law has started to govern the electrical load of data centres, as in the State of New York's 2026 moratorium on large data centres pending the development of binding standards (Hochul, 2026), and the European Union's data-centre energy-efficiency package (White & Case, 2026). Each of these regimes touches the physical layer of artificial intelligence. None governs it whole.

This article argues that the physical layer should be governed as a single critical-infrastructure object, and it builds a framework for doing so. The framework rests on a claim that the three existing framings, namely security protection, compute regulation, and energy permitting, are each partial, and that treating them as separate produces both gaps, where no regime reaches an asset, and conflicts, where regimes pull in opposite directions. The article develops four criteria for deciding when a class of compute infrastructure warrants designation as critical infrastructure, and it proposes a layered architecture of designation, protection, and regulation with institutional coordination at its centre. It grounds the framework in a worked instance from the Global South. Nigeria's Designation and Protection of Critical National Information Infrastructure Order 2024 classifies telecommunications towers, fibre-optic cables, and data centres as critical national information infrastructure and places their protection under the Office of the National Security Adviser (Federal Government of Nigeria, 2024). Nigeria therefore offers an example of a developing-market state that has moved ahead of many wealthier peers in the formal act of designation, while facing a protection-capacity and energy-governance gap that the framework must address.

Three research questions organise the analysis. First, what governance framings currently apply to the physical layer of artificial intelligence, and where do they leave gaps and generate conflicts? Second, on what principled basis should a class of compute infrastructure be designated as critical infrastructure, and what should designation entail beyond a label? Third, how can a designation, protection, and regulation framework be made adaptable across jurisdictions of unequal capacity, and what safeguards are needed to prevent the over-securitisation of civilian infrastructure? The purpose of the article is to supply a portable governance architecture for the physical layer of artificial intelligence, defensible in legal and institutional terms, illustrated across the United States, the European Union, the Gulf, and Nigeria, and attentive to the distinct position of constrained-grid states in the Global South.

2. Materials Studied and Literature Review

This section synthesises the materials on which the analysis draws, organised by theme. It covers the strategic reframing of compute infrastructure, the critical-infrastructure protection tradition, the compute-governance programme within artificial intelligence policy, the energy and environmental regulation of data centres, and the specific literature on infrastructure governance in the Global South. Across these themes the review identifies a consistent gap: the physical layer is studied in pieces, and no body of work treats it as a unified governance object.

2.1 Compute infrastructure as strategic terrain

A growing strategic literature treats the physical carriers of computation as instruments of national power. Buchanan (2020) framed data, algorithms, and computing power as the triad of artificial-intelligence capability, and subsequent work has sharpened the compute node of that triad into a question of infrastructure control. Analyses of the semiconductor supply chain describe chip fabrication as a concentrated strategic vulnerability (RAND Corporation, 2024), while studies of submarine cables document the physical connectivity layer as contested and poorly protected (Atlantic Council, 2024a; Carnegie Endowment for International Peace, 2024a). The World Economic Forum (2026a) has argued that artificial-intelligence infrastructure should be treated as critical infrastructure, and New America (2025b) has examined the security of the physical backbone on which artificial-intelligence services depend. This literature establishes that the physical layer is strategically significant. It stops short of specifying the governance architecture that significance requires.

2.2 The critical-infrastructure protection tradition

A separate and older tradition addresses the protection of critical infrastructure as such. In the United States, critical-infrastructure policy organises protection around defined sectors and lead agencies, with communications and information technology among them, and it has evolved from asset-hardening toward systemic resilience. The European Union has codified resilience obligations for essential entities through the Critical Entities Resilience Directive and the second Network and Information Security Directive, extending duties of risk management and incident reporting across member states. This tradition supplies mature concepts, including designation, essential-service definition, interdependence, and resilience, and it supplies institutional models for oversight. Its limitation for present purposes is that it predates the concentration of artificial-intelligence compute and does not address the specific dynamics of hyperscale data centres, cross-border cloud regions, or the electrical intensity of accelerated computing. The tradition also contributes a vocabulary of graduated obligation, in which designation triggers duties calibrated to the criticality of the asset, and a model of the lead-agency structure through which those duties are supervised. The United Kingdom's treatment of communications within its critical national infrastructure, organised around designated sectors and a national protective-security authority and complemented by a statutory telecommunications-security regime, offers a further reference point, as do the resilience and security standards developed by the International Organization for Standardization, the International Electrotechnical Commission, and the International Telecommunication Union. What the tradition has yet to absorb are three features that distinguish artificial-intelligence compute from the communications assets it grew up governing: the extreme concentration of capability in a small number of hyperscale operators, the electrical intensity that binds the facility to the grid in a way earlier assets were not bound, and the cross-border cloud architecture that dissolves the tidy correspondence between an asset and a single jurisdiction.

2.3 The compute-governance programme in artificial intelligence policy

Within artificial-intelligence policy, a distinct research programme argues that compute is a uniquely governable node in the artificial-intelligence supply chain because it is detectable, excludable, and quantifiable (Sastry et al., 2024). Work on training-compute thresholds examines how a measure of computation can serve as a regulatory trigger and where that approach reaches its limits (Heim & Koessler, 2024), and broader surveys map the open technical problems in operationalising compute governance (Reuel et al., 2024). The European Union's Artificial Intelligence Act operationalises this logic by attaching systemic-risk obligations to a training-compute threshold (European Union, 2024). This programme is concerned primarily with governing the capability of models through their compute inputs. It treats infrastructure as a lever for model governance rather than as an object of protection in its own right, which is the reciprocal concern this article develops.

2.4 Energy and environmental regulation of data centres

A fourth body of material addresses the electrical and environmental footprint of data centres. The International Energy Agency (2025; 2026) projects steep growth in data-centre electricity demand and identifies grid connection as an emerging constraint. Regulatory responses are appearing at sub-national and regional levels. The State of New York enacted a temporary moratorium on data centres capable of drawing fifty megawatts or more while it develops binding standards (Hochul, 2026), and the European Union is advancing a data-centre energy-efficiency package and rating scheme (White & Case, 2026). This literature treats the data centre as a regulated energy object defined by its load. It rarely connects that regulation to the security designation of the same facility, which is one of the coordination problems the present framework addresses.

2.5 Infrastructure governance in the Global South

A final theme concerns digital infrastructure in the Global South. The International Monetary Fund (2025) documents a compute divide between a Compute North and a Compute South, and analyses warn that data-localisation mandates without an accompanying energy and protection plan can deepen dependence rather than reduce it (New America, 2025a). Continental and national instruments are emerging, including the African

Union's continental artificial-intelligence strategy (African Union, 2024) and Nigeria's designation of critical national information infrastructure (Federal Government of Nigeria, 2024). This literature centres the agency of developing-market states and cautions against treating them as passive recipients of frameworks designed elsewhere. It has not yet examined designation of compute infrastructure as critical infrastructure as a governance instrument that constrained-grid states can wield, which this article takes up directly.

2.6 The gap

Read together, these five bodies of work describe the same physical layer from different vantages, each capturing part of it. The strategic literature establishes that the layer matters. The critical-infrastructure tradition supplies protective concepts but predates hyperscale compute. The compute-governance programme regulates model capability through infrastructure but not infrastructure as such. The energy literature regulates load in isolation from security. The Global South literature centres agency but has not connected designation to compute. The gap this article fills is the absence of a unified governance object and a framework that integrates protection, regulation, and energy conditioning across jurisdictions of unequal capacity.

3. Methodology

The study adopts a qualitative, interdisciplinary methodology that combines doctrinal analysis of legal instruments, comparative regulatory analysis across jurisdictions, and conceptual framework development. It does not involve human participants, experiments, or the generation of statistical data, and it makes no empirical claims of that kind. Its evidence base is documentary and entirely public.

The doctrinal component reads primary instruments, including Nigeria's Designation and Protection of Critical National Information Infrastructure Order 2024 and its parent Cybercrimes Act 2015, the European Union's Artificial Intelligence Act and its critical-entities and network-security directives, United States critical-infrastructure policy, and the State of New York's 2026 executive order, to establish what each instrument designates, protects, or regulates, and by which institution. The comparative component sets these instruments against one another to identify shared principles and points of divergence, following the logic of comparative regulatory analysis in which jurisdictions are read as varied responses to a common problem. Cases and jurisdictions were selected to span the relevant range of capacity and approach: the United States as a market-led and security-forward system, the European Union as a rules-led system, the Gulf as a site of recent physical contestation, and Nigeria as a constrained-grid developing-market state that has nonetheless legislated designation. The conceptual component develops criteria and a framework by abstracting from the comparative material and testing the abstraction against the worked instances.

The analysis relies solely on publicly available information, including official instruments and agency publications, peer-reviewed and institutional research, and reputable reporting of the 2026 data-centre strikes and the 2024 to 2025 cable disruptions. It does not draw on classified information, confidential reports, proprietary operational data, or any client matter, and it does not assess military tactics, operational vulnerabilities, or intelligence capabilities. Where the physical-security incidents are discussed, they are treated as illustrations of governance challenges rather than as objects of operational analysis. The principal limitation of the approach is that a documentary and comparative method cannot measure the effectiveness of the proposed framework in advance of its adoption; it can establish the framework's coherence, its fit with existing instruments, and its defensibility, and it can specify the conditions under which its effectiveness could later be evaluated. A second limitation is that the fast pace of events in 2026 means the incident record continues to develop, and the argument is therefore built on the governance implications of the incidents rather than on any contested attribution of responsibility.

4. Results and Findings

The analysis yields five findings. The first three establish the governance gap and the criteria for closing it. The fourth presents the comparative designation practice, including the Nigerian instance. The fifth sets out the proposed framework.

4.1 Three partial framings compete over one layer

The physical layer of artificial intelligence is currently addressed by three governance framings that operate largely in isolation. The first is critical-infrastructure protection, concerned with defending designated assets against physical and cyber threats and assuring the continuity of essential services. The second is compute regulation, concerned with governing the capability of artificial-intelligence systems by attaching obligations to the computation used to train or run them. The third is energy and environmental permitting, concerned with the electrical load, emissions, water use, and siting of data centres. A single hyperscale facility can fall within all

three at once. It may be a designated communications asset for security purposes, a site of regulated compute for artificial-intelligence purposes, and a fifty-megawatt load for energy purposes, governed by three different authorities that do not necessarily consult one another.

This separation produces two failure modes. The first is a gap, where an asset that matters escapes every regime, as with a submarine-cable landing station that is neither a designated critical entity nor a regulated compute site nor a permitted load, yet whose severance can disconnect a country. The second is a conflict, where regimes pull against each other, as when an energy authority accelerates data-centre permitting to attract investment while a security authority would prefer to slow and scrutinise the same build for resilience and ownership concerns. The interdependence among sectors sharpens the cost of this fragmentation. When a designated cloud facility is at once a large electrical load and a site of regulated compute, a decision taken under any one framing reaches into the others. An energy authority that grants a rapid connection to attract investment may seat a facility on a fragile part of the grid that a resilience authority would have avoided. A security authority that restricts foreign ownership may raise the cost of the capital that an energy transition requires. A compute regulator that concentrates obligations on the largest operators may deepen the very concentration that makes the layer fragile in the first place. None of these tensions can be resolved inside a single framing. The finding is that the physical layer is over-determined in principle and under-governed in practice, and that coordination, rather than the creation of a wholly new regime, is the central design problem.

4.2 The physical layer has become materially contested

The events of 2024 to 2026 establish that the physical layer has become contested terrain and can no longer be treated as a settled utility. The March 2026 strikes on Amazon Web Services facilities in the United Arab Emirates and Bahrain caused structural damage, interrupted power, required fire suppression that produced further water damage, and degraded more than one hundred services across a cloud region, including consumer banking and payment applications (DefenseScoop, 2026; Marketplace, 2026; Tech Policy Press, 2026). The significance for governance is that a commercial facility owned by a private hyperscaler produced cascading public consequences when it was struck, which is the defining signature of critical infrastructure. The submarine-cable disruptions in the Baltic Sea and around Taiwan across 2024 and 2025 make the same point at the connectivity layer, where a small number of vessels and anchors can sever the physical routes that carry a nation's external data, and where attribution and repair both remain slow and contested (Recorded Future, 2025; Al Jazeera, 2025). The finding is that materiality and contestation, two of the classic markers of critical infrastructure, are now present at the compute and connectivity layers of artificial intelligence, which strengthens the case for designation.

4.3 Four criteria for a defensible designation

If designation is warranted, it must rest on a principled basis rather than on the political salience of the moment, because designation carries costs, including security obligations, restrictions on ownership and operation, and the risk of militarising civilian assets. The comparative material supports four criteria. The first is essentiality, meaning that the asset or class of assets supports services whose interruption would cause serious harm to public safety, economic stability, or the continuity of government. The second is cascading interdependence, meaning that failure of the asset propagates into other sectors, as the AWS strikes propagated into banking and payments. The third is contestation, meaning that the asset faces a credible threat, whether physical, cyber, or from concentrated foreign control, which distinguishes genuinely critical assets from merely large ones. The fourth is governability, meaning that a designating authority can in fact impose and verify obligations on the asset, without which designation is symbolic. The finding is that these four criteria, applied together, discipline designation and provide a defensible test that can be applied consistently across jurisdictions.

4.4 Comparative designation practice, including Nigeria

Jurisdictions vary in how far they have moved from recognising the physical layer to designating it. The United States possesses a mature sectoral apparatus in which communications and information technology are designated critical-infrastructure sectors, and it has extended security scrutiny to data-centre supply chains and to the concentration of advanced compute, while its energy authorities separately manage the load. The European Union has raised resilience duties on essential entities through the Critical Entities Resilience Directive and the second Network and Information Security Directive, and it regulates compute capability through the Artificial Intelligence Act and data-centre energy through its efficiency package, though these operate on parallel tracks. Nigeria presents a distinctive case. Through the Designation and Protection of Critical National Information Infrastructure Order 2024, made under the Cybercrimes Act 2015, the Federal Government designated thirteen critical sectors and classified telecommunications towers, fibre-optic cables, and data centres as critical national information infrastructure, placing audit, compliance, and the development of a protection plan under the Office

of the National Security Adviser (Federal Government of Nigeria, 2024; Mondaq, 2024). Nigeria has therefore performed the formal act of designation at the physical layer that many wealthier jurisdictions have approached only in pieces. The finding is twofold. Designation of compute infrastructure as critical infrastructure is already practised, including by a developing-market state, which establishes its feasibility. And the value of designation depends on the capacity to protect and on coordination with the energy and telecommunications regulators, which is precisely where the Nigerian case reveals a gap, since designation under a security authority does not by itself resolve the electricity constraint that shapes how Nigerian data centres are built and powered.

4.5 A layered designation, protection, and regulation framework

The framework that follows from these findings treats the physical layer as one object governed through three coordinated functions rather than three separate regimes. The designation function applies the four criteria to identify classes of compute and connectivity infrastructure that warrant critical status, and it fixes the threshold, whether by electrical load, by service dependence, or by role in cross-border connectivity. The protection function attaches resilience, incident-reporting, and continuity obligations to designated assets, and it assigns a lead authority with the capacity to verify them. The regulation function conditions the operation of designated assets on energy, siting, and ownership requirements, so that a facility is not protected as critical while being permitted without regard to the grid it strains or the foreign control it concentrates. The distinguishing feature of the framework is institutional coordination: a standing mechanism through which the security authority, the telecommunications regulator, and the energy regulator share a single designation list and reconcile their obligations, so that the gaps and conflicts identified in the first finding are closed. The framework is designed to be adaptable across capacity levels, allowing a high-capacity jurisdiction to operate all three functions through mature agencies while a constrained-grid state sequences them and draws on regional mutual assistance for the protection function it cannot yet fully resource.

5. Discussion

The findings support a framework, and this section interprets them against the literature, addresses the strongest objections, and sets out the significance of the argument for theory and for policy.

5.1 From lever to object: reversing the compute-governance frame

The compute-governance programme treats compute as a lever for governing model capability (Sastry et al., 2024; Heim & Koessler, 2024). This article reverses the direction of that relationship and treats compute infrastructure as an object to be governed for its own resilience and security. The two directions are complementary rather than competing. A training-compute threshold governs what a model may do; a designation threshold governs how the facility that trains it must be protected and powered. Reading them together resolves an oddity in current practice, in which the European Union regulates the compute intensity of a model under its Artificial Intelligence Act while regulating the electricity intensity of the data centre under a separate energy instrument and the resilience of the same facility under a third directive, with limited connection among the three. The theoretical contribution is to name the physical layer as a distinct governance object and to show that governing it requires integrating the security, compute, and energy framings rather than choosing among them.

5.2 The over-securitisation objection

The most serious objection to designating compute infrastructure as critical infrastructure is that it militarises civilian assets, concentrates control in security agencies, chills investment, and entrenches incumbent operators who can bear the compliance burden. The objection has force, and the framework answers it in three ways. First, the four criteria discipline designation so that it reaches genuinely critical classes of asset rather than every large facility, which limits scope. Second, the coordination mechanism deliberately places the telecommunications and energy regulators alongside the security authority, so that designation does not default to a purely security logic and the civilian regulators retain their role. Third, the protection function should rest on transparent, auditable obligations rather than on discretionary security control, so that operators face predictable duties rather than opaque intervention. The Nigerian instance illustrates both the risk and the mitigation, since designation under the Office of the National Security Adviser must be paired with the ordinary work of the communications and electricity regulators if it is to strengthen resilience without narrowing the sector to a security preserve.

5.3 Sovereignty and the transnational asset

A second objection concerns fit between national designation and the transnational character of the assets. Submarine cables land in one state but traverse the high seas and connect many; cloud regions serve users across borders; ownership is frequently foreign. National designation cannot, on its own, protect an asset whose

vulnerability lies in international waters or in a supply chain controlled abroad. This is a real limit, and it defines the multilateral layer of the framework. Designation at the national level should be complemented by regional mutual-assistance arrangements for protection and repair, and by international recognition of submarine cables and cross-border cloud infrastructure as protected civilian infrastructure, so that the connectivity layer receives the kind of normative protection long extended to other shared systems. The 2024 to 2025 cable disruptions and the slow, contested repair that followed show why the national and the multilateral layers must operate together (Recorded Future, 2025).

The transnational character of the asset also points to the value of standards harmonisation as a complement to designation. Where states converge on common definitions of the essential service, common thresholds for designation, and common expectations for incident reporting and repair, an operator that runs infrastructure across many jurisdictions faces a coherent set of duties rather than a contradictory patchwork, and a state that lacks the capacity to develop its own technical standards can adopt an internationally recognised baseline. The International Telecommunication Union, the International Organization for Standardization, and the International Electrotechnical Commission already produce security and resilience standards that could anchor such a baseline, and the framework proposed here is intended to be read against them rather than in isolation. Harmonisation does not require a single global regulator, which is neither realistic nor desirable given the diversity of legal systems. It requires a shared reference that national designation regimes can each implement in their own form, which lowers the cost of governing the layer for the states least able to bear it and reduces the friction for the operators who run it across borders.

5.4 Designation without capacity

A third objection is that designation without the capacity to protect is symbolic and may even be counter-productive, creating an expectation of security that the state cannot meet. The governability criterion is the framework's answer, since it withholds the label from assets the authority cannot in fact govern. For constrained-capacity states, the implication is sequencing rather than abstention. A state can designate, attach reporting and continuity obligations that are within its administrative reach, and rely on regional mutual assistance and on obligations placed on well-resourced private operators for the protective functions it cannot yet perform directly. This reframes designation for the Global South as a staged instrument of agency rather than as an unaffordable imitation of high-capacity regimes, which aligns with the caution in the development literature against importing frameworks whole (New America, 2025a).

5.5 The energy dimension as a governance condition, not an afterthought

Because the physical layer is defined as much by the electricity it draws as by the data it carries, the framework integrates energy governance into designation rather than leaving it to a separate track. This matters most acutely in constrained-grid states. Where the grid is unreliable, data centres bridge the gap with fossil generation, which locks in emissions and cost for the life of the facility (International Energy Agency, 2025; GSMA, 2024). A designation regime that ignores energy would protect a facility as critical while permitting it to be powered in a way that undermines both climate goals and the reliability the designation is meant to secure. Conditioning designation on energy-intensity reporting and siting requirements, on the model of the New York moratorium and the European efficiency package adapted to local conditions (Hochul, 2026; White & Case, 2026), makes the energy dimension a term of criticality rather than an afterthought. This is where the article's infrastructure argument meets the energy and climate concerns that run through the wider research programme.

5.6 Significance for theory and policy

The theoretical significance of the argument is the identification of the physical layer of artificial intelligence as a unified governance object and the specification of the criteria and functions required to govern it. The practical significance is a framework that regulators can adopt incrementally and adapt to their capacity, and that international bodies can use to coordinate the protection of shared connectivity. For scholarship, the argument opens a line of inquiry at the intersection of critical-infrastructure law, artificial-intelligence governance, and energy regulation that has so far been pursued in separate literatures. For policy, it offers a way to respond to the events of 2026 that is more disciplined than ad hoc securitisation and more complete than the parallel-track regulation that now prevails.

The significance also extends to the distribution of governance capacity across the world. If the physical layer is governed only where mature agencies already exist, the states that host the fastest-growing share of new construction will remain the least able to protect and condition it, and the divide the development literature describes will widen from a divide in capacity into a divide in resilience. A framework that is explicitly adaptable across capacity levels, and that assigns the heavier protective duties to well-resourced operators where the state cannot yet perform them, is therefore not only a technical convenience. It is a condition for governing

the layer wherever it is built rather than only where it is easiest to govern. This is the sense in which the argument speaks to the wider research programme on infrastructure and equity, and it is the reason the Nigerian instance is treated here as a source of governance insight rather than as a regional exception to a framework designed elsewhere.

6. Conclusion and Recommendations

6.1 Summary of key findings

The physical layer of artificial intelligence has become materially and strategically critical, and the events of 2024 to 2026 have made it contested terrain. It is currently governed by three partial framings, namely security protection, compute regulation, and energy permitting, that operate in isolation and thereby leave gaps and generate conflicts. Designation of compute and connectivity infrastructure as critical infrastructure is feasible, and it is already practised, including by Nigeria, whose 2024 order designates telecommunications and data-centre assets under the Office of the National Security Adviser. Designation is defensible when it rests on four criteria, namely essentiality, cascading interdependence, contestation, and governability, and it is effective only when paired with the capacity to protect and with coordination among the security, telecommunications, and energy authorities. A layered framework of designation, protection, and regulation, coordinated through a standing institutional mechanism and adaptable across capacity levels, closes the gap that the three isolated framings leave open.

6.2 Practical and policy recommendations

The framework translates into a sequence of measures that ascend from the national regulator to the region and to the multilateral order, and that can be adopted incrementally.

At the national level, governments should designate classes of compute and connectivity infrastructure as critical infrastructure using the four criteria, and they should publish the thresholds they apply, whether an electrical-load threshold on the model of New York's fifty-megawatt trigger, a service-dependence threshold, or a connectivity threshold for cable landing stations. Designation should carry defined and auditable obligations rather than discretionary control, covering resilience, incident reporting, and service continuity, and a lead authority should be named with the capacity to verify them. Where a state has already designated, as Nigeria has, the immediate priority is to operationalise the protection plan and to connect it to the ordinary regulators, so that the communications regulator and the electricity regulator share the designation list and reconcile their requirements with the security authority. Critically, designation should be conditioned on energy terms, requiring designated facilities to measure and report their energy intensity and to meet siting conditions, so that criticality and sustainability are governed together rather than at cross purposes. For constrained-grid states, these obligations should be sequenced to match administrative capacity, beginning with reporting and continuity duties that are within reach and placing the heavier protective burdens on well-resourced operators.

At the regional level, states should establish mutual-assistance arrangements for the protection and repair of shared infrastructure, and they should converge on common designation thresholds and reporting standards so that operators face a coherent set of obligations across a region rather than a patchwork. In Africa, the continental artificial-intelligence strategy and the regional economic communities provide the institutional vehicles for this convergence (African Union, 2024), and a shared approach to cable landing stations and cross-border data routes would address vulnerabilities that no single state can protect alone. Regional cooperation also allows states to pool the scarce technical capacity that the protection function requires, which is the practical answer to the governability constraint for smaller economies.

At the multilateral level, the international community should recognise submarine cables and the cross-border elements of cloud infrastructure as protected civilian infrastructure, extending to the connectivity layer the normative protection long afforded to other shared systems, and it should coordinate the surveillance and repair capacity on which that protection depends. The events of 2026 give this recommendation urgency, because the strikes on cloud facilities and the disruptions to cables show that the physical layer of artificial intelligence is now a site of conflict, and because the layer is transnational by construction and cannot be secured by national designation alone. A multilateral compute-infrastructure resilience framework, aligned with the work of the International Telecommunication Union and standards bodies, would give the national and regional layers a common reference and would begin to build the shared expectation that this infrastructure is off-limits as a target.

Taken together, these measures make the case for treating the physical layer of artificial intelligence as a single governed object, protected and regulated and powered in a coordinated way, across jurisdictions of very different capacity. The argument matters beyond the security of the operators, because the services that depend on this

layer, including banking, health, emergency response, government, and the wider economy, are public, and their continuity is a public interest. The call is therefore to policymakers, regulators, standards bodies, and international organisations to move now, using instruments many of them already possess, to close the governance gap that the events of 2026 have exposed.

6.3 Suggestions and insights for future research

Three lines of inquiry follow from the argument, and each also marks a boundary of the present study. The first concerns measurement. Because a documentary and comparative method establishes coherence rather than effectiveness, future work should develop indicators for the resilience gains that designation and coordination are meant to produce, and should evaluate them where the framework is adopted. The second concerns the constrained-grid setting, which this article treats through the Nigerian instance but which merits dedicated comparative study across several developing-market states, to test how far the sequencing approach holds under different administrative and energy conditions and to refine the governability criterion. The third concerns the multilateral layer, where the normative protection of submarine cables and cross-border cloud infrastructure raises questions of international law that this article identifies but does not resolve, and which warrant sustained legal analysis. Pursuing these three lines would extend the framework from a coherent proposal toward an evaluated governance model, and would deepen the intersection of critical-infrastructure law, artificial-intelligence governance, and energy regulation that the physical layer of artificial intelligence now demands.

References

- African Union. (2024). *Continental Artificial Intelligence Strategy*. African Union.
https://au.int/sites/default/files/documents/44004-doc-EN-_Continental_AI_Strategy_July_2024.pdf
- Al Jazeera. (2025, March 10). As undersea cables break off Europe and Taiwan, proving sabotage is tough. *Al Jazeera*. <https://www.aljazeera.com/news/2025/3/10/as-undersea-cables-break-down-proving-sabotage-a-difficult-task>
- Atlantic Council. (2024a). *Cyber defense across the ocean floor: The geopolitics of submarine cable security*. Atlantic Council. <https://www.atlanticcouncil.org/in-depth-research-reports/report/cyber-defense-across-the-ocean-floor-the-geopolitics-of-submarine-cable-security/>
- 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/>
- Carnegie Endowment for International Peace. (2024a). *Subsea communication cables in Southeast Asia: A comprehensive approach is needed*. Carnegie Endowment for International Peace.
<https://carnegieendowment.org/india/research/2024/12/southeast-asia-undersea-subsea-cables>
- Center for Democracy and Technology. (2024). *EU AI Act brief, part 5: General-purpose AI models*. Center for Democracy and Technology. <https://cdt.org/insights/eu-ai-act-brief-pt-5-general-purpose-ai-models/>
- DefenseScoop. (2026, March 3). Commercial data centers emerge as targets in modern warfare after drones hit 3 AWS facilities. *DefenseScoop*. <https://defensescoop.com/2026/03/03/commercial-data-centers-drone-warfare-amazon-aws/>
- European Union. (2024). *Regulation (EU) 2024/1689 (Artificial Intelligence Act), Articles 51 and 55*. Official Journal of the European Union. <https://artificialintelligenceact.eu/article/51/>
- Federal Government of Nigeria. (2024). *Designation and Protection of Critical National Information Infrastructure Order, 2024* (made under the Cybercrimes (Prohibition, Prevention, etc.) Act, 2015). Federal Government of Nigeria.
- GSMA. (2024). *Tower power Africa: Energy challenges and opportunities*. GSMA.
<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. <https://arxiv.org/abs/2405.10799>
- Hochul, K. (2026). *Executive Order No. 62: Establishing a temporary moratorium on data centers in New York while the state develops higher standards for data center development*. State of New York.
<https://www.governor.ny.gov/executive-order/no-62-establishing-temporary-moratorium-data-centers-new-york-while-state-develops>

- International Energy Agency. (2025). *Energy and AI*. International Energy Agency. <https://www.iea.org/reports/energy-and-ai>
- International Energy Agency. (2026). *Electricity 2026: Analysis and forecast*. International Energy Agency. <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). International Monetary Fund. <https://www.imf.org/-/media/files/publications/wp/2025/english/wpica2025076-print-pdf.pdf>
- Marketplace. (2026, March 4). Drone strikes on Middle East data centers highlight their importance and vulnerabilities. *Marketplace*. <https://www.marketplace.org/story/2026/03/04/drone-strikes-highlight-data-center-vulnerabilities>
- Mondaq. (2024). The Federal Government has officially signed the Designation and Protection of Critical National Information Infrastructure Order 2024. *Mondaq*. <https://www.mondaq.com/nigeria/telecoms-mobile-cable-communications/1539672/>
- New America. (2025a). *Africa's digital sovereignty trap: The data center dilemma*. New America. <https://www.newamerica.org/insights/africas-digital-sovereignty-trap/>
- New America. (2025b). *Securing the backbone of AI*. New America. <https://www.newamerica.org/insights/securing-the-backbone-of-ai/introduction/>
- RAND Corporation. (2024). *Supply chain interdependence and geopolitical vulnerability: The case of Taiwan and high-end semiconductors* (RRA2354-1). RAND Corporation. https://www.rand.org/pubs/research_reports/RRA2354-1.html
- Recorded Future. (2025). *Submarine cable security at risk amid geopolitical tensions and limited repair capabilities*. Recorded Future. <https://www.recordedfuture.com/research/submarine-cables-face-increasing-threats>
- Reuel, A., Bucknall, B., Casper, S., Fist, T., Soder, L., Aarne, O., et al. (2024). *Open problems in technical AI governance*. arXiv. <https://arxiv.org/abs/2407.14981>
- Sastry, G., Heim, L., Belfield, H., Anderljung, M., et al. (2024). *Computing power and the governance of artificial intelligence*. arXiv. <https://arxiv.org/abs/2402.08797>
- Tech Policy Press. (2026). The legal and policy fallout from data center strikes in the Middle East war. *Tech Policy Press*. <https://www.techpolicy.press/the-legal-and-policy-fallout-from-data-center-strikes-in-the-middle-east-war/>
- White & Case. (2026). *Data centres and energy consumption: Evolving EU regulatory landscape and outlook for 2026*. White & Case. <https://www.whitecase.com/insight-alert/data-centres-and-energy-consumption-evolving-eu-regulatory-landscape-and-outlook-2026>
- World Economic Forum. (2026a). *It is time to start treating AI infrastructure as critical infrastructure*. World Economic Forum. <https://www.weforum.org/stories/2026/04/ai-infrastructure-critical-infrastructure/>