POLICY NOTE

EQUITABLE COMPUTE COOPERATION AND THE ROLE OF THE G20

Divij Joshi, ODI Global
Jai Vipra, IT for Change
Lucia Velasco, Oxford Martin School

WORKSTREAM

T20 U.S.
An informal engagement of global think tanks

2026

Equitable Compute Cooperation and the Role of the G20

Divij Joshi (ODI Global), Jai Vipra (IT for Change), Lucia Velasco (Oxford Martin School)

Abstract

Access to compute, which encompasses the specialised semiconductors, data centres, software and operational capacity that underpin AI systems, is concentrated at nearly every layer of the AI stack, and this concentration is producing system risks and inequalities, resulting in a new form of digital divide. Most countries are unable to build or govern compute capacity to support public priorities in AI development, and must rely on capacity that is allocated according to commercial priorities, with little leverage over pricing, availability, or allocation of compute infrastructure. National responses have not closed this gap, despite the emergence of 'sovereign' AI strategies. This brief argues that equitable compute cooperation is rooted in three coordination problems which require international cooperation to solve: First, current approaches prioritise autonomy and capacity, without addressing usability or control. Second, the lack of interoperability, governance frameworks and shared standards leads to overspend on, and underutilisation of, compute. Third, there is a lack of standardised information to evaluate and forecast compute supply and demand across countries.

The brief argues that the G20 is well placed to address these concerns through three lines of action: First, developing a common evidence base, working with international organisations such as the OECD, that measures both compute availability within a jurisdiction and the capacity subject to public allocation rights, beginning with voluntary self-reporting. Second, coordinating public financing and procurement so that public contributions across the supply chain include public access and benefit considerations, including through benefit-sharing principles, a toolkit of model contractual provisions, and the integration of public compute into multilateral development bank appraisal and financing of digital infrastructure. Third, establishing voluntary arrangements for pooled access and interoperability between public compute programmes, so that public-interest workloads can draw on capacity across participating jurisdictions and resilience is improved through portability.

Implementation Roadmap

Diagnosis

The concentration of compute infrastructure introduces inequalities and systemic risks within the global AI supply chain. Countries are currently unable to develop or govern compute infrastructure required for national priorities owing to a lack of standardised information, finance and governance of compute. This brief extends G20 commitments on systemic risks and equitable digital infrastructure and AI towards concrete action for cooperation on compute infrastructure.

Stakeholders

Lead/supervision: Sherpa Track working group or ministerial group responsible for digital and innovation, along with the G20 Finance Track for the financing elements of Recommendation II.

Execution: OECD, World Bank and MDBs and regional development banks as knowledge partners; national digital and innovation ministries; national and regional compute programmes.

Action Plan and Implementation Steps

- *Knowledge and Policy Products:* Index for forecasting and measuring public compute demand and supply. Interoperability standards for compute sharing infrastructure. Endorsement of benefit-sharing principles for compute, including a toolkit of model contractual and procurement provisions.
- *Funding:* MDB and regional development bank support for national and regional public compute projects, factoring in access and benefit sharing considerations in project appraisals.

Monitoring and Accountability

- *Process:* Annual reporting to a digital economy working group or ministerial, with multi-year roadmaps overseen by a partner institution.
- *Existing methods:* Build on OECD measurement dimensions and existing infrastructure funding and appraisal models.
- *Success indicators:* Endorsement of principles and toolkit on benefit sharing and model procurement; members reporting against the indicator set; share of national compute capacity subject to public allocation rights; amount of MDB operations including compute with access covenants; number of programmes participating in reciprocal access.
- *Failure indicators:* Lack of reporting members; lack of access arrangements between compute-rich and compute-poor members; lack of projects appraised on public access requirements.

Conclusion & Next Steps

This brief proposes concrete mechanisms for addressing inequalities and systemic risks arising from the current lack of coordination of compute infrastructure, through the development of measurement, financing, access sharing, and governance mechanisms. Immediate next steps include identifying and securing knowledge partners to develop the toolkit for benefit sharing, and encouraging member states to declare voluntary commitments for compute sharing mechanisms.

Introduction

The economic benefits of, and capacity to govern, AI are unequally distributed among countries (United Nations, 2026). This inequality leads to systemic dependencies and risks, particularly for developing economies, which undermine long-term developmental priorities linked with AI, which the G20 is well placed to address in its role as a forum for global economic cooperation and progress.

One area where these inequalities are particularly stark is in the underlying infrastructural capacities for production and use of AI systems — also known as 'compute'. 'Compute' refers to the specialised semiconductors, as well as the software, infrastructure and knowledge resources needed to run them (International AI Safety Report, 2026). The availability of suitable compute is crucial for the development and use of contemporary AI systems, including for training, fine-tuning, or running applications, particularly frontier models. However, this infrastructure depends on global supply chains with concentrated chokepoints, which are unevenly distributed. Moreover, operational control over compute – the determination of what gets built and who has computational power to develop or use frontier AI capabilities – remains highly unevenly concentrated. Many countries in the Global South, particularly, lack operational control over infrastructure and risk being increasingly dependent upon leasing expensive compute runtime from external providers, without adequate long-term reliability, security, or control over questions of allocation or pricing.

Chokepoints in AI supply chains include semiconductors, lithographic equipment, accelerator chips themselves, cloud services, and even the underlying energy resources (OECD, 2025a). These elements are unevenly distributed among countries and companies. For example, a country might host many data centers and still rely on foreign cloud service providers. Consequently, a number of countries are now dependent upon external provision of their capacity needs (for example, through renting cloud capacity from private providers) and therefore have little leverage over pricing, allocation, or availability (Lehdonvirta et al., 2025). Countries in both the Global North and South are affected by these limitations, although obtaining compute is much easier for those with substantial domestic capacity and bargaining power.

These dependencies cross borders and are not possible to resolve for any given country through national policy alone, and require global cooperation. However, global frameworks on AI cooperation tend to focus on AI governance, safety verification, technical standards, or technology transfers, and do not yet provide mechanisms for measuring and comparing compute capacity, coordinating equitable public access to compute, or managing risks associated with shared infrastructures and supply chains.

This brief examines how the G20 could help close this important gap, which is emerging as a new kind of global digital divide. It proposes three areas for cooperation: developing a shared evidence base on compute capacity and distribution; coordinating public financing and procurement conditions to expand equitable access to compute; and establishing arrangements for pooled access, interoperability, and resilience across compute infrastructure.

Why Cooperation on Compute Matters

'Compute' is the material precondition of AI development. Every stage of the model lifecycle, from training through fine-tuning to inference at scale, depends on access to specialised semiconductors such as in the form of GPUs, the data centres that host them, and the software, expertise and operational capacities that are built around them. However, while global compute has grown exponentially, most of this growth has been privately owned and operationalised, which is reflected in the price signals and availability of compute infrastructure for different sectors and priorities. As such, compute dedicated to public-interest purposes remains comparatively negligible, with consequences for the trajectory of AI development itself (Vipra and Davies, 2024).

This capacity is also extraordinarily concentrated, at nearly every layer of the compute stack required to deploy operational compute infrastructure for AI. The processing capacity for raw materials used in semiconductors and data centres is concentrated within a few providers and countries (OECD, 2025a). Chips move through an inelastic supply chain with single-supplier chokepoints in fabrication, lithography, and design (OECD, 2025c); the United States and China host a majority of the world's AI-capable compute infrastructure, and only 33 countries, mostly in the Global North, host inference-relevant public cloud AI compute, just 24 of them training-relevant (Hawkins, Lehdonvirta & Wu, 2025), and operational control over AI-relevant data centres is highly concentrated in a few firms (OECD, 2025a). Moreover, the electricity grid capacity and reliability required by AI data centres is also unevenly distributed among and within countries (IEA, 2025). Each of these chokepoints contributes to structural constraints in the capacities of countries to develop secure and reliable compute infrastructure which address public priorities. The consequence of this concentration is that most countries cannot build compute capacity, and more importantly, they cannot govern it, as capacity is concentrated and allocated according to commercial priorities. As such, countries dependent on renting access have limited mechanisms to direct its allocation towards public or developmental priorities.

Due to this concentration, AI supply chains are highly susceptible to systemic risks and disruptions. This structure of global AI compute compels most governments to rent capacity under terms set by private providers. As a result, countries lacking their own infrastructure face significant challenges in directing compute resources toward national priorities or public-interest objectives (Shearer, Davies & Lawrence, 2024; Open Future, 2025). Moreover, the costs and risks of this structure fall hardest on countries in the Global South, which already lag behind in compute availability, and where access is rented on revocable, provider-set terms, leaving these countries constrained in determining the trajectories of their own AI economies or directing capacity (that they have invested in) towards their own public priorities. Such constraints are also faced by countries in the Global North, even those that have made AI sovereignty a policy priority, as the concentrated compute market curtails state choices (Lehdonvirta et al., 2025; Open Future, 2025).

Coordination problem I: Current approaches prioritise autonomy and capacity, without addressing usability or control

Given the existing competitive dynamics for resources to develop compute infrastructure to address existing and potential AI needs, countries are stuck in a race to develop autonomous compute infrastructure, including investing in different aspects of the supply chain — from data centres to critical minerals. However, autarky in AI is almost impossible to achieve in practice for any one country or even region (Tanner et al., 2026). This is leading to countries investing in inefficient and unsustainable infrastructure, as well as exclusionary dynamics of countries bidding against each other against the same inelastic supply chain of compute, without gaining control or deciding allocation.

Existing responses to compute availability have not addressed this allocation gap. National approaches to securing AI compute differ depending on their goals. Some focus on attracting investment through data centers, while others seek domestic ownership of cloud providers or access to advanced chips. Public compute programs also use different combinations of direct provision through public facilities, distributed infrastructure through networks of smaller facilities, and subsidies for commercial cloud services (Davies and Vipra, 2025).

While 'sovereign' compute strategies now proliferate across the G20, with countries seeking to attract the requisite investments and capacities to develop long-term compute capacities in AI, they each seek build-outs against the same constraints in the supply chain, thereby worsening the procurement position of all, or build leverage through exclusivity over certain parts instead, creating a cycle of competition and a drive towards more exclusionary compute. It also promotes a race to the bottom, where competition to co-locate data centres results in public investment in compute without securing allocation rights over the same.

The rapid commercial compute buildout in recent times, moreover, hides public investments and public costs. AI infrastructure benefits from tax breaks, subsidised electricity, and public investment in electricity and other infrastructure. The negative externalities that commercial compute creates, such as the unsustainable extraction of natural resources, also represent public costs that are not currently adequately leveraged by public institutions. Public support for the commercial compute buildout need not be unconditional – public costs, which go unaccounted for, can be compensated for in the form of compute allocations for public benefit, which can be addressed through frameworks for cooperation (Adan et al., 2025).

Coordination Problem II: Lack of interoperability, governance frameworks and shared standards leads to overspend on, and underutilisation of compute

Where market allocation has not delivered public capacity, states have begun building public alternatives, from the US NAIRR and the UK's expanded AI Research Resource to IndiaAI's subsidised shared compute, Brazil's PBIA and the EuroHPC AI Factories (NAIRR Task Force, 2023; DSIT, 2025; IndiaAI, 2025; Government of Brazil, 2024; EuroHPC JU, 2026; Lemke & Schneider, 2025). However, these programmes remain national or regional in scope

— most restrict eligibility to users within the funding jurisdiction, and each operates its own access regimes, authentication, and allocation procedures, so that capacity cannot be shared across borders even where governments wish to, resulting in duplication of capacity without complementarity in its utilisation. Facilities in some jurisdictions run at capacity with significant unmet demand (Ghahramani et al., 2023) while capacity elsewhere may lie underused.

International initiatives address parts of the problem, but are limited in scope and only partially address concerns around compute. For example, Pax Silica focuses primarily on creating trusted supply chains and access among participating partners, while the World AI Cooperation Organisation emphasises development and capacity-building. These initiatives are currently siloed and inaccessible to a majority of the world, and do not provide mechanisms for comparing national capacity, coordinating procurement and financing, allocating compute for cross-border public-interest workloads, or governing shared infrastructure.

Coordination problem III: There is a lack of standardised information to evaluate and forecast compute supply and demand across countries

One barrier to addressing inefficient and inequitable allocation is the lack of available information about compute capacities and requirements. Governments today operate with a lack of information about compute capacity in their own countries and in other countries, and thus cannot measure the problem before they attempt to solve it (OECD, 2025). The scale and type of compute needed vary across models, workloads, and particularly between training, fine-tuning, and inference — from edge-inference to frontier training. While large-scale infrastructure is necessary for certain stages in the development of certain AI models, different scales and types of compute are useful for others depending on the model, the workload, and the stage, i.e., training, fine-tuning, or inference. However, the lack of standard indexes and measurement frameworks for compute is a barrier for governments in understanding or forecasting where investments in compute may be required.

Moreover, although frameworks and indices are being developed to measure compute, these usually attempt to measure general commercial availability and installed capacity, rather than publicly governed capacity over which public authorities hold allocation rights, through ownership or contractual conditions. In particular, public compute demands and mechanisms to determine allocation, matched to policy and developmental aspirations, are rarely mapped to existing capabilities. While it is evident that publicly determined compute capacity lags far behind commercial capacity, some measure of forecasting, proportion, and comparison can direct policies and public funding towards more appropriate allocation.

These coordination problems reveal a clear institutional gap: **no existing multilateral framework provides mechanisms to translate public-interest priorities for AI into transparent rules for eligibility, allocation and access to compute. Addressing this requires the G20 to coordinate the *measurement, financing* and *governance* of compute capacity.**

Recommendations

The G20 is well placed to address this coordination gap because its members contribute different inputs to the global compute supply chain, and benefit from cooperation across the compute stack. These include major chip and cloud economies, countries supplying materials needed for the infrastructure itself, and large or growing markets looking for greater compute access. All should have different and complementary interests in compute cooperation. Building on the foundations laid by previous G20 presidencies to enhance cooperation on AI infrastructure, including the Maceio Declaration in Brazil (2024), and the AI Task Force under the South African presidency (2025), the G20 is well positioned to establish frameworks and standards for compute cooperation among its members, and globally.

Each coordination gap identified above corresponds to a policy lever available to the G20, which are presented below:

Recommendation I. Develop a common evidence base for compute capacity

Measuring and forecasting compute capacity and needs is foundational to AI development. However, there is a lack of a standardized evidence base for comparing countries. The OECD's work is instructive in this regard; in 2023, its survey on a blueprint for building national compute capacity found that "no country today has sufficient data on, or a targeted plan for, national AI compute capacity" (OECD, 2023). Subsequent work by the OECD and World Bank on measuring domestic compute availability indicates that measurement is feasible but remains partial (World Bank, 2025; Lehdonvirta, 2025).

Building on the OECD's dimensions of measurement – across *capacity, effectiveness, and resilience* – the G20 should work with International Organisations like the OECD to support the development of frameworks and indexes that specifically account for compute capacity supply and demand, for purposes defined by public authorities and linked to national AI strategies. Frameworks should measure both the availability of compute within a jurisdiction, as well as the capacity over which public authorities hold allocation rights and how. Indicators could include both installed public and publicly accessible capacity by tier; utilisation rates and allocation outcomes; energy and grid pipeline constraints on planned facilities, the terms of access, utilisation and allocation, as well as forecasted demands, including projected demand from public-priority workloads. An initial survey could include voluntary self-reporting from select G20 member states, which may be gradually expanded to other members and non-members.

Recommendation II. Coordinate public financing and procurement in the public interest

As described above, the benefits of increased public investment in a concentrated market accrue primarily to a small number of large firms (Davies and Vipra, 2024; Open Future, 2025). Although several states currently contribute to the compute supply chain in different ways — as mineral suppliers, energy providers, infrastructure hosts, and financial investors, including fiscal incentives, grid capacity, or public infrastructural investments — they do so without corresponding claims on the compute capacity these enable. These investments across the supply chain should therefore be conditioned on reciprocal benefits, including reciprocity in compute access and technology transfer, which can be coordinated by the G20 and partner organisations through model contractual terms, supported by procurement standards that develop common bargaining power.

The G20 should endorse benefit-sharing principles for equitable compute coordination, and promote the development of mechanisms such as model contractual provisions, which condition public contributions in compute on public benefit. The G20 could work with knowledge partners like the World Bank or OECD to develop a toolkit with model contractual provisions on reciprocal obligations, public access commitments, and allocation set-asides for public-priority compute. Benefit sharing principles should endorse reciprocity between states and compute providers, with procurement, subsidy, and the use of natural resources carrying conditions on portability, auditability, and price stability, as well as model provisions on local allocation shares, transparent energy and water terms, and participation in allocation governance; and between capacity-rich and capacity-seeking members, through public access commitments and allocation set-asides for cross-border public-interest workloads.

The G20 Finance Track should work with MDBs, including the World Bank and regional development banks, to support the financing of national and regional compute projects, and integrate compute into existing infrastructure appraisal mechanisms. In countries with limited existing capacity, support should cover feasibility studies, technical assistance, energy and connectivity requirements, and the development of appropriate operating capacity. Grants and concessional lending are likely to be most appropriate during the early stages of a project. Guarantees, blended finance and commercial lending should be used where demand and potential revenue are more established (Velasco et al., 2026). Financing appraisals should consider transparent allocation, affordable access for public institutions, and capacity for agreed public-interest uses.

Recommendation III. Develop shared access and interoperability between public compute programs

Interoperability and the federation of compute infrastructure between public compute programs potentially reduces systemic risk by reducing high-risk dependencies and single points of failure, as reciprocal access agreements give participants an alternative when domestic capacity is unavailable, and shared allocation across jurisdictions with distinct infrastructure supply reduces the likelihood that all participants are constrained at once. Regional facilities and remote-access arrangements can expand access without having every country build equivalent infrastructure, leveraging economies of scale and providing insulation against variable demand and supply (Velasco et al., 2026). This could build on existing models which have been trialled at smaller scales, like the EuroHPC (European High Performance Computing Joint Undertaking) and LUMI projects in the EU, which establish frameworks for accreditation, ownership and allocation, and common and interoperable platforms for access.

G20 members should pilot voluntary arrangements that allow public-interest projects in participating countries to access public compute available in other participating countries, which could also be open to non-members. Participating programs could reserve an agreed share of their capacity for this purpose. Infrastructure would remain under national or regional ownership, while members would agree on common rules for which projects are eligible, how capacity is allocated, and how access is managed, along the lines of projects like the EuroHPC and India's MAITRII.

G20 members should also promote interoperability among participating public-compute programs so that workloads can be moved from one system or provider to another where technically feasible. This would allow countries to use alternative capacity when their infrastructure is unavailable. Members should therefore test workload portability and continuity during cloud, data-center, and electricity outages, including whether backup capacity depends on the same provider, location, or electricity system (Gelvanovska-Garcia, Mačiulė and Rossotto, 2024). Initial arrangements could focus on a limited number of public-interest workloads and expand as experience grows.

References

Adan, Sumaya Nur, Wiaterek, Joana, et al., Oxford Martin AI Governance Initiative, 'AI Benefit-Sharing Framework: Balancing Access and Safety.' 2025.

Gelvanovska-Garcia, Natalija, Vaiva Mačiulė, and Carlo Maria Rossotto. World Bank Group. 2024. 'Advancing Cloud and Data Infrastructure Markets: Strategic Directions for Low- and Middle-Income Countries'. https://doi.org/10.1596/978-1-4648-2065-6.

Lehdonvirta, Vili, Bóxī Wú, and Zoe Hawkins. 2025. 'Compute North vs. Compute South: The Uneven Possibilities of Compute-Based AI Governance Around the Globe'. *Proceedings of the 2024 AAAI/ACM Conference on AI, Ethics, and Society*. 828–38. https://dl.acm.org/doi/10.5555/3716662.3716734.

Lehdonvirta, Vili, Boxi Wu, Zoe Jay Hawkins, Celine Caira, and Lucia Russo. OECD. 2025. 'Measuring Domestic Public Cloud Compute Availability for Artificial Intelligence'. https://doi.org/10.1787/8602a322-en.

OECD. 2023a. 'A Blueprint for Building National Compute Capacity for Artificial Intelligence'. https://doi.org/10.1787/876367e3-en.

OECD. 2025a. 'Full Report: Competition in Artificial Intelligence Infrastructure'. <https://www.oecd.org/en/publications/competition-in-artificial-intelligence-infrastructure_623d1874-en/full-report.html.>

OECD. 2025b. 'Measuring Domestic Public Cloud Compute Availability for Artificial Intelligence.' https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/10/measuring-domestic-public-cloud-compute-availability-for-artificial-intelligence_39fa6b0e/8602a322-en.pdf.

OECD. 2025c. 'Mapping the Semiconductor Value Chain: Working towards Identifying Dependencies and Vulnerabilities'. https://doi.org/10.1787/4154cdbf-en.

Shearer, Eleanor, Davies, Matt and Lawrence, Mathew, Ada Lovelace Institute. 2025. 'The Role of Public Compute', https://www.adalovelaceinstitute.org/blog/the-role-of-public-compute/

Tanner, Brooke, Cameron F. Kerry, Andrew W. Wyckoff, Nicoleta Kyosovska, Andrea Renda, and Elham Tabassi. 2026. 'Is AI Sovereignty Possible? Balancing Autonomy and Interdependence'. Brookings Institution.

https://www.brookings.edu/articles/is-ai-sovereignty-possible-balancing-autonomy-and-interdependence/.

United Nations, 2026. 'Preliminary Report of the Independent International Scientific Panel on AI: Evidence-based assessment of opportunities, risks and impacts of artificial intelligence'

Velasco, Lucia, Sumaya Nur Adan, Muhammad Salar Khan, et al., 2026. 'Financing the AI Triad: Compute, Data and Algorithms'

Vipra, Jai and Davies, Matt, Ada Lovelace Institute. 2024. 'Mapping global approaches to public compute'.https://www.adalovelaceinstitute.org/policy-briefing/global-public-compute/

World Bank. 'Digital Progress and Trends Report 2025: AI Foundations'. https://www.worldbank.org/en/publication/dptr2025-ai-foundations.