Future Talent Insight Brief

# The Hidden Cost of Intelligence: AI Infrastructure and the Global South

*An integrated value-chain analysis and policy framework for AI infrastructure, critical minerals, and labour in the Global South*

Amay Pandey, '26, Kingston University London, UK
Derek Gunter '27, University of Tennessee Knoxville, USA
Júlia Rocha '28, Federal University of Alfenas, Brazil
Krystyna Herasymenko '28, University of Liverpool, UK
Nathaly Avalos '28, Universidad de La Libertad, Mexico

# Summary

Artificial intelligence runs on a physical infrastructure concentrated in the Global South, where flexible labour regulation, weak environmental enforcement, and proximity to mineral supply chains lower costs for those building it while concentrating risk for those living within it.

This brief asks a direct question: as demand for AI infrastructure drives the expansion of data centres and AI training facilities into the Global South, to what extent does this growth intensify existing pressures on labour and natural resources, and what governmental and corporate strategies can transform that structural dependency into sustainable value for these communities? To answer it, the brief traces the full value chain that produces AI — from artisanal mineral extraction, through hardware manufacture, to cloud revenue — and finds that costs are borne almost entirely at one end while value accumulates almost entirely at the other.

In response, advances integrated recommendations across the two sectors where that dependency is most acute: on labour, they advance digital worker protection and skills development; on natural resources, they redirect mineral value toward host communities through supply-chain transparency and benefit-sharing, deploying regulatory, corporate, and international instruments to convert structural dependency into shared value. These recommendations are grounded and validated by expert interviews with practitioners and researchers across industry, applied economics, and the Global South.

This brief report, covered in Summary, provides an executive summary of the principal findings and recommendations. Section 1, this introduction, establishes the framing that the brief is designed to address. Section 2 surveys the existing literature across the four clusters, and Section 3 details the methodology for the mineral demand estimates and the labour vulnerability analysis. Section 4 presents the findings, and Section 5 situates those findings in current governance debates through expert commentary. Section 6 sets out the recommendations. Section 7 contains the appendices and source documentation.

# 1. Introduction

The artificial intelligence systems that will reshape how our generation works, learns, and governs rest on a physical foundation — minerals, hardware, and human labour — deliberately concentrated in the Global South. That concentration follows the contours of an older geography, one in which the regions that supplied raw inputs to earlier industrial orders now supply the inputs to the computational one. AI did not create this pattern; it inherits it and, in inheriting it, deepens it. The cobalt at the mine gate and the content labeller's working day are bound into the same value chain as the cloud revenue at the far end, yet at no point along that chain does a binding mechanism return value toward the people who make it possible.

What gives this pattern its force is who is permitted to speak about it. The rules that will govern this value chain for the coming decade are being written now, and largely without the people who bear the cost of building it. Workers at the extraction face, the communities living alongside the mines, and the data labellers who render AI models usable are not external to the AI economy, yet they are absent from the forums where its governance is decided. They are incorporated into the system as inputs to it rather than as participants in it, with the conditions of their work determined by architectures built without their voice.

What follows from this is less a matter of optimism than of timing and motive. The will to act, however genuine, dissipates without a strategy precise enough to carry it. Governments and companies do not move on intention but on instruments, the concrete mechanisms that turn a stated commitment into a condition of doing business. The task, then, is to engineer the strategy that lets it act: career pathways that give digital workers a stake in the economy they sustain, benefit-sharing, and in-country beneficiation that return mineral value to extraction communities. Across both streams, inclusion becomes something the system is designed to produce, as a condition of market access and operating licence, rather than something it is asked to permit.

That is why the present moment matters; the frameworks that will govern this value chain are still being drafted, and it is what this brief asks its readers to take up. The chain can still become a bridge rather than a pipeline, and the actors with the power to make it one are the same ones now drafting the rules.

# 2. Literature Review

Artificial Intelligence is projected to contribute up to USD 15.7 trillion to the global economy by 2030 (PwC, 2023), integrating deeply into nearly every sector. Yet this promise carries costs that remain largely invisible in mainstream AI discourse and fall disproportionately on the Global South. Within the region itself, countries able only to run ready-made models on outdated hardware ("Compute South") and those with no public AI infrastructure at all ("Compute Desert") depend nearly completely on systems built and hosted elsewhere (Lehdonvirta, Wu, & Hawkins, 2024). That infrastructure is predominantly foreign-owned, so revenues and data move northward; in Africa, non-African telecommunications companies control most critical digital infrastructure and citizen data is stored on European servers, while local providers bear the access costs (Png, 2022). Absent binding obligations spanning both labour conditions and mineral supply chains, this dependency will deepen in step with AI expansion.

AI's environmental footprint is inseparable from mining (Regilme, 2024). Data centres are built in Global South countries for their land, lower operational costs, weaker regulation, and proximity to mineral supply chains, reducing costs for hyperscalers while concentrating environmental and social risk in host communities (Png, 2022; Lehdonvirta et al., 2024). Demand for the underpinning minerals is rising: the DRC supplies over 70% of global cobalt, and Indonesia nearly 50% of nickel (IEA, 2023), and the sourcing labour remains sub-human: across three mining sites, 93% of 931 people interviewed had experienced some form of exploitation, and the cobalt has been traced directly to batteries sold by major multinationals (ILO, 2024b).

AI products depend equally on content moderation and data labelling performed by low-paid Global South workers, labour as structurally essential as the minerals. These workers perform high-skill, high-hazard work for poverty wages with no recognition, career pathway, or upskilling support. A TIME investigation found that OpenAI outsourced text snippets, including descriptions of child sexual abuse and torture, to Kenyan workers paid USD 2 per hour, while paying the intermediary roughly nine times what workers received (Perrigo, 2023). Living-wage and mental-health gaps persist, with lasting psychological harm across the workforce, and intermediary platforms classify workers as independent contractors, stripping social protection and obscuring the industry's dependence on them.

International AI governance reproduces the same inequalities it claims to address. One concrete lever for Global South governments is the physical presence of data centres on their territory: in principle, states can set conditions of access (Lehdonvirta et al., 2024), but most lack the legal frameworks and technical capacity to do so. Proposed remedies, South-South cooperation, co-governance models, and new roles for non-Western actors remain ideas, none of which has a mechanism to compel corporations to honour commitments (Png, 2022).

The literature's central limitation is that even its strongest sources treat individual dimensions of exploitation without linking them into a single analytical chain. The Global South bears the costs of extraction, from hardware to model training, but these stages are not examined as structurally connected, and the proposed remedies remain equally fragmented (Png, 2022; Regilme, 2024). The chain this brief traces points instead toward integrated responses: supply-chain accountability that follows value to its source.

# 3. Methodology

The methodology integrates six primary datasets across labour, infrastructure, energy, environmental, and mineral domains. These quantitative sources are complemented by a series of expert interviews spanning industry practice in artificial intelligence, applied economics and academic research in information sciences with direct field experience in the Global South. Interviewees' responses were analysed against the report's findings and integrated as primary evidence.

## 3.1 Data Sources

The datasets construct an end-to-end chain from physical AI infrastructure to material and environmental extraction demand. Epoch AI's facility-level data anchors the demand side by translating confirmed GPU deployments into installed power capacity; USGS and IEA intensity factors convert that capacity into mineral demand estimates; the Mineral Ores Round the World dataset maps those figures against the active global producer base in a way aggregate statistics cannot support; and environmental benchmarks from EIA CBECS, EPA eGRID, and The Green Grid complete the chain. Sources were selected for public availability, regulatory standing, and independent verification, with proprietary sources used only where no comparable public alternative existed at the required granularity. The 45× revenue-to-mineral ratio is a metric derived for this analysis, combining these datasets to compare raw extraction value against AI cloud revenue at equivalent capacity.

Workforce vulnerability was assessed through a composite measure built from World Bank and ILO indicators, including social safety nets, legal protections, and economic stability, weighed against each country's share of the global AI labour supply chain across both tiers: mineral labour, drawing on ILO field evidence and UNESCO survey data, and digital labour, drawing on field-verified workforce assessments alongside market-sizing data. Two caveats apply. Mineral intensity factors are estimates built from USGS and IEA hardware norms rather than audited bills of materials, and should be read as directionally accurate at scale; and mine site counts are a proxy, recording only site existence and status, so active-producer counts indicate concentration rather than supply sufficiency.

### 3.4 Qualitative Analysis

The process began with the identification of a core research question, followed by a mapping exercise categorising potential interviewees by area of expertise across five domains: AI and technology, government and philosophy, Global South perspectives, human resources and labour, and natural resources. Candidates were identified through LinkedIn and existing professional networks and contacted through a formal invitation outlining the project's objectives, scope, and format. Interviews were recorded and transcribed with each participant's prior approval, with a structured question guide shared in advance. Following each interview, transcripts were reviewed, and key insights were documented in individual analysis notes capturing thematic findings and direct quotations. These notes formed the basis for the commentary in this section, organised thematically around the three central arguments of the brief.

## 4. Findings

This section presents quantitative signals across mineral supply and the workforce that trains AI systems. On the mineral side, the findings trace supply fragility across semiconducting and magnetic minerals and the widening gap between AI capital expenditure and mining investment. On the labour side, they examine the wages and protections of the data labelling and content moderation workforce.

### 4.1 Mining

The 45x Value gap in minerals is in commodity prices, not as strategic inputs. Raw mineral extraction generated approximately $315M per GW of AI data centre capacity at current market prices. That same GW produces over $14B in annual AI cloud revenue and a 45x multiplier that accrues almost entirely to hardware manufacturers and cloud platform operators, not to the countries supplying the ore. Copper is priced identically whether it goes into a residential water pipe or a $22B frontier AI facility. Commodity markets have no mechanism to reflect the strategic criticality of the end use, creating a structural undervaluation of the mining supply chain relative to its irreplaceability in AI infrastructure.

Three minerals are already under active supply disruption. Gallium, Germanium, and Rare Earths are rated as CRITICAL vulnerability, and critically, this is not a forward-looking warning. China imposed export controls on gallium in July 2023, germanium in August 2023, and graphite export licensing in October 2003, while simultaneously being the dominant refiner of all three. These controls are active and operating while AI infrastructure demand is accelerating at its fastest recorded pace. The global AI buildout is structurally dependent on supply chains that a single nation has already demonstrated willingness to restrict, with no near-term Western refining alternative at scale.

Producer site counts reveal single points of failure invisible in aggregate data. The mineral ore dataset identifies 304,632 mine sites across 166 countries and reveals that the entire global AI infrastructure build-out places material demand pressure on 9 active Gallium producer sites, 10 Germanium sites, and 6 Indium sites worldwide. These are not supply chains; they are single points of failure. A disruption at two or three facilities in any of these categories would meaningfully constrain the global supply of materials with no near-term substitutes in semiconductor manufacturing. This finding is invisible in standard commodity market analysis, which reports aggregate production volumes rather than facility-level concentration.

The investment gap is widening at a structurally dangerous rate. The big 4 AI companies (Amazon, Microsoft, Google, Meta) have committed a combined $315B in capital expenditure in 2025 alone, growing at 50% annually. The global mining sector invested approximately $50B in the same year, growing at 5%. The IEA estimated that meeting mineral demand through 2040 requires $500-600B in new mine capital expenditure, and that investment is not materialising at anywhere near the required pace. Permitting timelines of 7-20 years for new mines versus 12-14 months for new data centres mean the supply response cannot keep pace with demand even if investment decisions were made today.

## 4.1 Labour

In the regions most critical to sustaining the AI supply chain, the labour force is defined by informality and uncontracted work. This stratum operates entirely outside legal protection frameworks, faces serious harm without adequate safety equipment, and is more likely to fall below the poverty line, trapped in a cycle of precariousness in which the industry they sustain grows rapidly while their own wages and safety do not.

AI's minerals come from least protected workers

Analysing mining labour against national vulnerability indices.

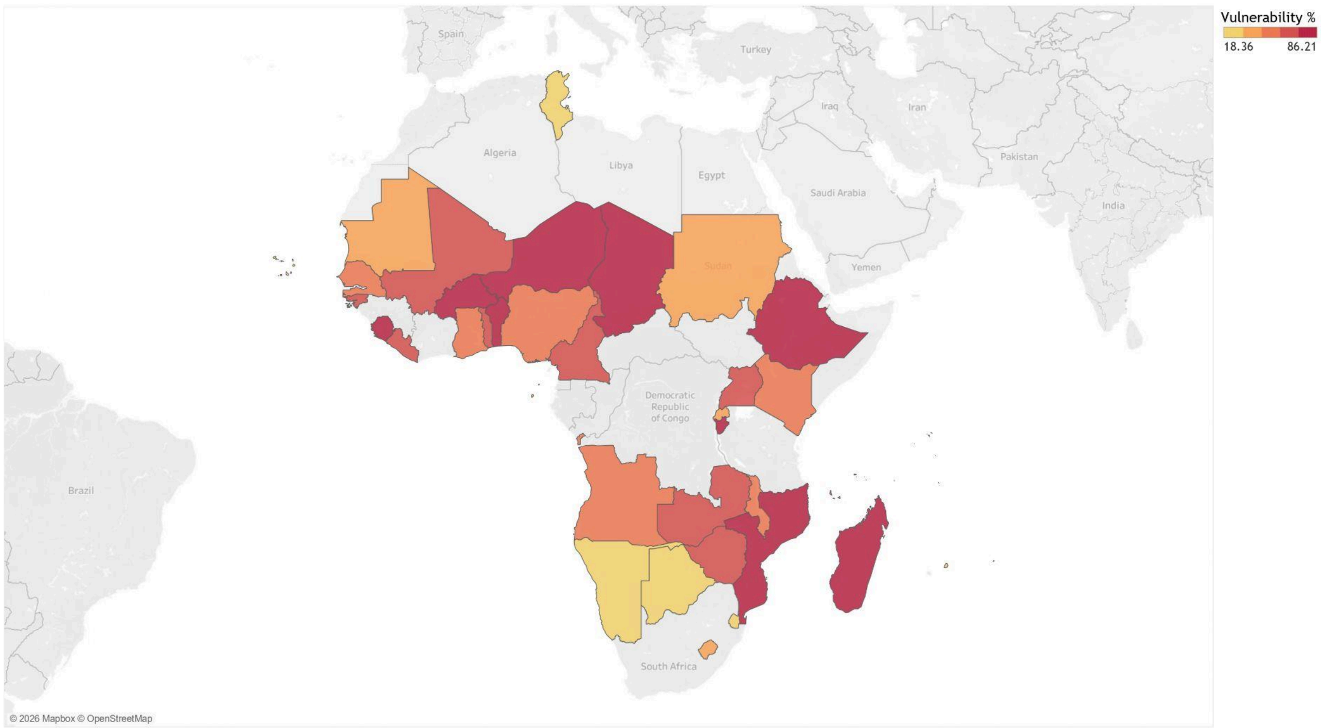


Source: ILO (2025)& World Bank Vulnerability Data.

*Key minerals are mined in countries where labour laws are weakest, especially in the global South. This map shows where the physical building blocks of AI are actually coming from. The colour indicates how much vulnerability those workers face, meaning a lack of legal rights or safety nets. To read this, look for countries with darker shades of red; these are the hotspots where the global tech industry is most dependent on workers who have the least protection.*

The Global North–South wage gap between data centre operators and mining labourers is among the most significant economic disparities in the modern economy, and it is widening: as AI demand and corporate revenue scale, extraction wages remain flat, held down by the absence of local ownership and infrastructure and by competition from established foreign companies that own the mines. The AI boom has concentrated economic activity at the tertiary level, leaving the primary worker who extracts the raw materials at the bottom of the value pyramid.

Transitioning out of extraction and into higher-paid digital roles requires bridging a wide skill gap. Against the competencies of data centres demand in power control and networking, mining work draws mainly on manual labour, leaving extraction workers outside the technical chain of production. Yet that distance is bridgeable: large-scale mining already requires logistical coordination and the operation and maintenance of heavy machinery, skills not far removed from those needed to run the hardware of an AI data centre and capable, with targeted technical and vocational investment, of being upskilled into their digital equivalents.

Even so, the inability to retain local value undermines the capital needed to fund training institutions. Bridging the gap by 2030 requires a deliberate connection between mining infrastructure and digital literacy programmes. Without it, the next generation of primary-sector workers in the AI revolution will go on supplying materials with no stake in the industry they sustain (World Bank, 2024).

# 5. Expert Commentary

To situate the brief's quantitative findings within practitioner and field-level judgement, this section draws on seven semi-structured expert interviews. The panel spans industry, applied economics, academia, and AI governance, and their commentary is organised around the brief's three central arguments.

## 5.1 The Value Gap: Structural, Not Incidental

The 45× ratio documented in Section 4 is a structural feature of how global value chains are designed. One framing situated it within infrastructure inequality, casting the Global North and Global South as operating with fundamentally different runway lengths, one requiring far greater resilience to achieve the same takeoff (A. Rorissa, 2025). Others placed it within a pattern of extractive positioning that predates AI, with African policymakers now naming the dynamic explicitly even as the instruments to address it remain underdeveloped; the more useful national AI strategies identify concrete mechanisms for learning and iteration rather than generic aspiration (K. Apeagyei, 2025).

The asymmetry was also read as the continuation of an older extractive dependency: governments reliant on corporate tax revenue stay reluctant to impose conditions that might drive investment away, handing corporations outsized political influence. "This relationship of dependency creates a real governance problem, because companies hold more influence than they should in developing countries" (C. A. dos Santos, 2025). The gap fits longer patterns of Global South integration into global commodity and manufacturing structures (L. Palafox, 2025; J. Polanco, 2025), with a linguistic dimension: the absence of AI systems trained on Spanish-language corpora risks leaving the region absent from the narrative infrastructure that future AI systems will encode (J. Polanco, 2025).

## 5.2 Labour and Workforce: The Gap Between Training and Opportunity

On labour, expert perspectives converged on a structural gap that upskilling alone cannot close. Kojo Apeagyei identified employment creation, not skills supply, as the critical missing variable in most Global South workforce development frameworks. His central concern was whether the jobs exist at the other end of training pipelines, or whether programmes are built on the assumption that supply will generate its own demand. The interventions that actually close that gap (internships, placements, mentorship systems) are, in his assessment, chronically underinvested relative to the volume of certificate provision.

Prof. Abebe Rorissa pointed to reskilling initiatives already underway in Ethiopia and Morocco as evidence that locally grounded responses are possible within constrained infrastructure environments. “AI won't replace you, but people who know AI will replace those who don't.” León Palafox located the equivalent gap in Mexico in the weak triangulation between government, industry, and universities, noting that technical graduates are being produced but that the durable skills and industry integration required to deploy them effectively are not. Jorge Polanco added organisational friction as an underappreciated constraint: the adoption of AI within Latin American companies has been slowed by generational clashes between traditional

IT leadership and newer development teams, leaving many organisations still building the data infrastructure that AI deployment requires.

### 5.3 Governance: Where Leverage Actually Sits

On governance, the experts converged most sharply with the brief's recommendations, beginning with a direct challenge to voluntarism: any instrument relying on corporate goodwill will not move the actors it targets. Genuine rights for workers outside a company's home jurisdiction require governments willing to impose extraterritorial duties on firms operating within their borders, mirroring the market-access logic of the brief's recommendations rather than the ESG commitments shown to have failed in Section 3: "Governments create rights. So insofar as it's genuinely about actual legal rights, those need to come from a legal framework. Governments can create duties on companies that operate within their border, that create rights for workers that operate outside their borders" (J. Morris, 2025). Enforcement, in turn, depends on credible measurement: economic output relative to a worker's starting position, employment statistics disaggregated by type of work, and direct engagement with unions, verified by a body independent of governments and corporations alike (K. Apeagyei, 2025).

Complementary mechanisms include regional bodies such as the African Union and UNECA exercising collective leverage, data trusts for community compensation, and retraining provisions embedded in operating licences to build "an ethical bridge rather than a one-way pipeline" (A. Rorissa, 2025). The precondition is disclosure: without mandatory reporting of data-centre energy and water use, evidence-based regulation is impossible (L. Palafox, 2025), since "without that transparency and accurate data, any planning around the impacts is compromised" (C. A. dos Santos, 2025). A market-oriented counterpoint held that frameworks attracting investment and preventing monopoly may distribute AI's benefits more broadly than redistributive regulation (D. Palacios de la Teja, 2025) — a tension the recommendations are built to navigate, conditioning market access rather than foreclosing it.

## 6. Recommendations

This section presents recommendations designed to redistribute value back to the communities that sustain the AI value chain. Organised into two streams, labour (R-L) and natural resources (R-NR), each converts a documented condition of dependency into an enforceable obligation, operating as a condition of market access or operating licence through instruments already available to governments, platforms, and multilateral institutions.

### 6.1 Protect and Invest in the Digital Workforce

#### R-L1 Require Supply Chain Accountability for Upskilling

AI platforms outsourcing data labelling or content moderation must be held jointly responsible for upskilling provision throughout the contracting chain. All such work — including micro-task, piecework, and freelance arrangements — must be governed by a written contract specifying pay rate, task scope, performance criteria, and termination conditions; absence of a compliant contract creates a rebuttable presumption of employee status in favour of the worker. Technology companies should additionally adopt a skills development levy of 1.5% of annual contract value, directed to nationally administered training funds with tripartite governance and geographically allocated to the regions where the labour is performed.

**Actors:** Technology platforms · BPO operators · National labour ministries

**Evidence base:** Fairwork (2025); Perrigo / TIME (2023); SOMO (2024)

#### R-L2 Establish a Credential-Linked Wage Premium Framework

Wages for digital workers should be tied to demonstrable skill. A structured four-tier credential framework should carry mandatory wage premiums of 10–20% above the verified living wage floor at each tier, with the living wage functioning as a non-negotiable floor. Progression through the credential framework should be supported by AI-augmented learning infrastructure: personalised learning paths that assess skill gaps and recommend specific micro-credentials with workers controlling their own data; AI-simulated practice environments for skill development before formal assessment; and a progression standard requiring both credential completion and demonstrated performance improvement.

**Actors:** Employers · Industry bodies · Labour market platforms

**Evidence base:** UNESCO (2026); ILO (2024); Perrigo / TIME (2023)

#### R-L3 Mandate Psychological Safety Standards for Content Moderation Work

Exposure to graphic and harmful content is an occupational hazard and must be governed with the same rigour as physical workplace hazards. Platforms should adopt ISO 45003-aligned psychosocial risk controls as a condition of market access in regulated jurisdictions: mandatory exposure limits of four hours per shift on the highest-severity content, 90-minute break intervals, and psychosocial risk assessment conducted regardless of contract type or worker classification. Algorithmic surveillance systems that monitor worker productivity must be subject to independent impact assessment, with workers informed of all data collected about them and given the right to contest automated decisions.

**Actors:** Platform operators · Regulators (EU AI Act; US DOL)

**Evidence base:** Fairwork (2025); Perrigo / TIME (2023); ILO AI Exposure Brief (2024)

### 6.2 R-NR: Close the Gap Between Mineral Value and Community Benefit

#### R-NR1 Require Full Mine-to-Data-Centre Supply Chain Disclosure

AI companies benefiting from mineral supply chains must disclose the full chain of custody from mine site to data centre, including extraction company identity, operating jurisdictions, royalty arrangements, and environmental assessments. Supply chain transparency should be adopted as a condition of market access in regulated jurisdictions, with a third-party audit every two years and mandatory public disclosure of findings. At the implementation level, all covered mineral supply chains must be disclosed to the level of country of origin and, for Tier 1 minerals, the mine of origin. Conflict mineral reporting should meet the OECD Due Diligence Guidance for Responsible Supply Chains as a minimum standard, with blockchain or equivalent immutable custody records for Tier 1 minerals required within three years of implementation.

**Actors:** AI companies · Hardware manufacturers · WEF / G20 governments

**Evidence base:** Stacciarini & Gonçalves (2025); USGS (2025);

R-NR2 Mandate In-Country Benefit-Sharing Agreements Before Extraction Commences

Communities close to active or planned extraction sites must be parties to a Community Benefit Sharing Agreement before operations begin, allocating a minimum of 3% of gross annual mineral revenue to a community-controlled fund with elected governance. Free, Prior and Informed Consent must be documented for all operations in or adjacent to indigenous territories, and that consent must be ongoing. Grievance mechanisms with a 30-day response commitment must be operational before first production, with non-retaliatory access guaranteed and independent monitoring of outcomes.

**Actors:** Host country governments · Mining companies · Development finance institutions

**Evidence base:** Banza Lubaba Nkulu et al. (2018); US DOL ILAB (2024);

R-NR3 Use Operating Licence Conditions to Mandate In-Country Beneficiation

Host country governments, supported by multilateral development institutions, should deploy the operating licence as the primary lever to rebalance this: mandatory in-country beneficiation requirements, minimum equity stakes in extraction joint ventures, binding technology transfer clauses requiring training of local engineers and geologists, and local procurement thresholds that escalate over time. Zimbabwe's 2023 prohibition on raw lithium exports and Zambia's 20–40% local procurement requirement demonstrate that these instruments can be deployed within existing legal and institutional frameworks. They do not require new international agreements — they require political will and multilateral support for governments exercising sovereign leverage over the infrastructure that makes AI possible.

**Actors:** Mineral-rich host governments · World Bank / AfDB · WEF as convener

**Evidence base:** Maheshwari / CIGI (2026); Regilme (2024); IEA Critical Minerals Dataset (2024)

# 7. Conclusion

This brief asked whether the expansion of AI infrastructure into the Global South intensifies the pressures already bearing on its labour and its land. The evidence answers plainly: it does, and it does so by design. Demand travels up the value chain directly, multiplying extraction volumes and labour requirements, while no binding mechanism anywhere along that chain returns value toward the people who generate it. The artisanal miner earns cents an hour, the data labeller two dollars, and the governance forums that legislate for these communities convene without them. One logic runs through all of it: those who supply the inputs are themselves treated as inputs. The brief reads the chain as a whole, from the mine gate to the hyperscaler, and shows it to be one continuous structure, answerable to one continuous response.

The argument turns here from diagnosis to obligation. What the chain lacks is a duty owed by those who profit from it to those who sustain it. The recommendations advanced here- worker protection, supply-chain transparency, and in-country beneficiation- assign that duty as enforceable conditions of market access and operating licence, built from instruments governments and platforms already hold. The brief stays bounded in what it claims: its mineral figures are directional, and the distribution of AI value reaches well beyond the chain mapped here. Its timing, though, is deliberate. The rules that will govern this value chain for the coming decade are being written now, and they are being written, still, largely without the people who bear the cost of building it. The chain can yet become a bridge rather than a pipeline, and the actors with the power to make it one are the same ones now drafting the rules.

## 8. Annotated Bibliography


ABI Research. (2026). Global data center capacity forecast Q1 2026. ABI Research. https://www.abiresearch.com

Proprietary report providing country-level AI data centre capacity projections and CAGR estimates through 2030. Used to contextualise regional investment trajectories and cross-validate JLL projections; figures are labelled as proprietary in the Methodology section.

Anwar, M. A. (2024). Value chains in AI. In Handbook on the global political economy of AI (Ch. 15). Taylor & Francis. https://www.taylorfrancis.com/chapters/oa-edit/10.4324/9781003391333-15/value-chains-ai-mohammad-amir-anwar

Edited handbook chapter analysing AI value chains with attention to data-labelling market growth and labour distribution, used as a reference for the global data labelling market valuation and projection figures in the literature review.

Apeagyei, K. (2025). Personal communication [AI Consultant and Researcher — MSc Media and Communications (Data and Society), London School of Economics; advisory experience with African governments on national AI strategies; AI Governance, Africa/UK].

Expert in AI governance and policy with firsthand advisory experience across African governments. Contributed comparative judgments on AI adoption criteria and governance frameworks in emerging markets.

Banza Lubaba Nkulu, C., et al. (2018). Sustainability of artisanal mining of cobalt in DR Congo. Nature Sustainability, 1(9), 495–504. https://doi.org/10.1038/s41893-018-0139-4

Peer-reviewed field study documenting health risks, environmental contamination, and labour exploitation at ASM cobalt sites in Lualaba Province, DRC, using soil/urine sampling and household surveys. Primary empirical evidence linking cobalt supply chains for AI hardware to documented community harm.

Energy Information Administration. (2024). Commercial Buildings Energy Consumption Survey (CBECS). U.S. DOE. https://www.eia.gov/consumption/commercial/

Nationally representative federal survey providing the PUE benchmark of 1.5 applied to convert IT load to total facility energy draw across all 38 frontier AI data centres analysed. Geographic limitation: US-based average may not reflect newer or non-US facilities.

Environmental Protection Agency. (2024). Emissions & generation resource integrated database (eGRID) 2024. U.S. EPA. https://www.epa.gov/egrid

The federal dataset providing the national average $CO_2$ emissions factor of 386 kg/MWh is used to translate AI data centre electricity consumption into carbon estimates. Limitation: The US-grid factor does not represent Middle Eastern, Asian, or Global South emissions intensity.

Epoch AI. (2026a, April 9). Frontier data centers [Dataset; CC BY 4.0]. https://epoch.ai/data/data-centers/

Satellite-verified, open-licence dataset covering 38 confirmed frontier AI data centres globally (26.9 GW total installed capacity), with facility-level power, GPU counts, capex, and location. Foundational infrastructure dataset underpinning all mineral demand, energy, and value chain analyses in this brief.

Epoch AI. (2026b, April 8). GPU clusters [Dataset; CC BY 4.0]. https://epoch.ai/data/gpu-clusters/

Global inventory of AI compute clusters complementing the Frontier Data Centers dataset, used to map compute concentration across Compute North, Compute South, and Compute Desert tiers. Cluster-level power capacity is sometimes estimated from hardware specs rather than being directly reported.

Fairwork. (2025). Fairwork AI supply chain assessment: Sama 2025. BMZ Digital. https://www.bmz-digital.global/wp-content/uploads/2025/12/Fairwork-Report-Sama_2025.pdf

Independent Oxford Internet Institute assessment of Sama's AI data labelling workforce using the five-dimensional Fairwork framework, documenting gaps in living wage compliance and mental health support. Most current field-verified evidence for labour conditions at a major AI data supply chain operator.

Farhad, S. (2025). Passengers in flight: AI governance capacity in the Global South. Digital Society, 4(2), 39. https://doi.org/10.1007/s44206-025-00195-6

Peer-reviewed Springer article arguing that Global South countries can develop AI governance capacity through institutional steering, distributional outcomes, and national learning. Provides the quantitative basis for the claim that 58% of AI governance initiatives originate from Europe and North America.

Fortune Business Insights. (2026). AI data centre market: Global market analysis, 2026–2033. https://www.fortunebusinessinsights.com/ai-data-center-market-110845

Proprietary report providing regional market sizing and growth projections for the AI data centre sector through 2033, used to triangulate against Epoch AI and JLL datasets. Methodology is not publicly auditable; all derived figures are labelled as proprietary.

International Energy Agency. (2023). Critical minerals market review 2023. IEA. https://www.iea.org/reports/critical-minerals-market-review-2023

Comprehensive IEA report on production, demand, pricing, and supply concentration for transition-critical minerals, including the tripling of lithium demand and a 70% rise in cobalt between 2017 and 2022. Used as a baseline cross-referenced against the updated IEA dataset for current projections.

International Energy Agency. (2024a). Critical minerals dataset [Dataset; CC BY 4.0]. https://www.iea.org/data-and-statistics/data-product/critical-minerals-dataset

Open-licence IEA dataset covering 37 minerals with demand projections under multiple scenarios, providing mineral intensity factors (tonnes/MW) applied to facility-level power data to derive mineral demand estimates. Intensity factors are technology-type averages and should be read as directional.

International Energy Agency. (2024b). Energy and AI. IEA. https://www.iea.org/reports/energy-and-ai

Flagship IEA report on AI's energy and water consumption implications, documenting that a typical AI data centre consumes electricity equivalent to 100,000 households and that global data centre investment has nearly doubled since 2022. Primary source for environmental consumption claims in RQ1.

International Labour Office. (2013). Measuring informality: A statistical manual on the informal sector and informal employment. International Labour Organisation.

Authoritative international reference for measuring informal sector employment and informal jobs, providing practical methodological guidance for national statistical offices on survey design, data collection, and dissemination. For this brief, the manual supplies the conceptual and statistical framework for disaggregating formal from informal employment in mining and digital labour supply chains. A limitation is that it predates the platform economy and does not address gig-based digital labour.

International Labour Organisation. (2024a, May 13). New ILO brief explains what AI exposure indicators reveal about jobs. ILO. https://www.ilo.org/resource/news/new-ilo-brief-explains-what-ai-exposure-indicators-reveal-about-jobs

ILO policy brief presenting AI exposure indicators distinguishing displacement-risk from augmentation-likely occupations, providing institutional framing for assessing data labelling and content moderation roles. Used alongside Fairwork and TIME sources to build the complete labour conditions picture.

International Labour Organisation. (2024b). Child labour in mining and global supply chains. ILO. https://www.ilo.org/publications/child-labour-mining-and-global-supply-chains

ILO field report based on interviews with 931 individuals at three ASM sites (DRC, Bolivia, Ghana, Philippines), finding 93% experienced labour exploitation, with direct supply chain linkage to multinational battery manufacturers. Primary empirical evidence for labour exploitation in AI mineral supply chains.

International Labour Organisation. (2025, January 2). World employment and social outlook: Trends 2025. https://www.ilo.org/publications/flagship-reports/world-employment-and-social-outlook-trends-2025

Authoritative annual labour market report identifying geopolitical frictions, climate costs, and sovereign debt risks as compounding pressures on labour market recovery, with decent work creation slowest in low-income countries. Provides the macroeconomic labour market backdrop against which digital worker and mining worker conditions should be understood. Limitation:

operates at country and regional aggregate level and does not capture occupation-specific or supply-chain-specific conditions.

JLL. (2026). 2026 global data centre outlook. JLL. https://www.jll.com/en-us/insights/market-outlook/data-center-outlook

Annual proprietary market report providing regional capacity forecasts including the 6,900 MW Middle East pipeline and 100 GW North American 2030 projection. Pipeline figures represent announced intentions rather than confirmed construction, distinguished from Epoch AI's satellite-verified data.

Lehdonvirta, V., Wu, B., & Hawkins, Z. (2024). Compute North vs. Compute South. Proceedings AAAI/ACM AIES, 7(1), 828–838. https://doi.org/10.1609/aies.v7i1.31683

Peer-reviewed AIES-24 paper providing the first systematic census of global public cloud GPU compute, classifying countries into Compute North, Compute South, and Compute Desert tiers. The primary basis for the three-tier model structuring RQ1's infrastructure distribution analysis.

Maheshwari, D. (2026). AI's Global South pivot: Equity, ethics and ecology (Policy Brief No. 225). CIGI. https://www.cigionline.org/static/documents/Policy_Brief_No._225_Maheshwari.pdf

CIGI policy brief examining equity, environmental, and governance implications of AI infrastructure expansion into the Global South, identifying actionable policy levers for host governments. Supports the prescriptive dimension of RQ2; limited long-term empirical evidence given the 2026 publication date.

Maurya, R. (n.d.). Mineral ores around the world [Dataset]. Kaggle. https://www.kaggle.com/datasets/ramjasmaurya/mineral-ores-around-the-world

Open dataset of 304,632 mine sites across 166 countries with commodity type, development status, and operator type, enabling supply-side analysis and employment structure disaggregation. Key limitation: no production volume data, so site counts function as concentration indicators rather than supply sufficiency measures.

Miceli, M., & Posada, J. (2022). The data-production dispositif. Big Data & Society. https://doi.org/10.1177/20539517251340600

Peer-reviewed article on the structures organising outsourced data production for AI systems, providing the analytical scaffolding for the labour-as-input argument in the literature review.

Morris, J. (2025). Personal communication [Director of Global Corporate Affairs and Advocacy — Former Prime Minister's Strategy Unit; Institute for Public Policy Research (IPPR); Strategic Communications, Global].

Specialist in strategic communications and global policy advocacy with senior-level government experience. Contributed expert judgments on policy-facing dimensions of AI workforce demand.

Our World in Data / Ritchie, H., & Roser, M. (n.d.). World energy consumption [Dataset]. Kaggle. https://www.kaggle.com/datasets/pralabhpoudel/world-energy-consumption

Dataset of 22,012 records across 306 countries (1900–2022) covering electricity demand, generation mix, and grid carbon intensity, used for regional energy trend context. Current AI-driven demand figures are drawn from IEA Energy and AI and Epoch AI, which carry more recent data.

Palacios de la Teja, D. (2025). Personal communication [Facilitator, Universidad de la Libertad — Public relations consultant and political analyst with over ten years of experience in public policy strategy; Policy and Governance, Mexico].

Practitioner with extensive experience advising on public policy and governance strategy in the Mexican context. Contributed judgments on regulatory and institutional criteria shaping AI skill demand.

Palafox, L. (2025). Personal communication [Head of AI and Innovation, GBM — Doctorate, University of Tokyo; former ML Director at Banorte and Grupo Salinas; Industry, Mexico].

Industry expert in applied artificial intelligence and machine learning within the Latin American financial sector. Contributed comparative judgments on technical skill demand criteria in industry contexts.

Perrigo, B. (2023, January 18). OpenAI used Kenyan workers on less than $2 per hour to make ChatGPT less toxic. TIME. https://time.com/6247678/openai-chatgpt-kenya-workers/

Orwell Prize-shortlisted investigation documenting conditions at Sama (OpenAI's content labelling outsourcer) using payslips and worker interviews, finding wages of $1.32–$2/hour and lasting psychological harm. The most important single source for the digital labour dimension of RQ1.

Png, M.-T. (2022). At the tensions between South and North. In Proceedings ACM FAccT 2022 (pp. 1434–1445). https://doi.org/10.1145/3531146.3533200

Peer-reviewed ACM FAccT paper mapping AI governance discourse from the Global South using critical and decolonial frameworks, quantifying that 58% of AI governance initiatives originate from Europe/North America versus 1.4% from Africa. Most important reference for the governance section's structural agency argument.

Polanco, J. (2025). Personal communication [Economist and Data Scientist — MSc Data Science, MBA; banking, fintech, retail, and industrial IoT across Latin America; Applied Economics, LATAM].

Cross-sector practitioner with applied experience in data science and economics across multiple industries in Latin America. Contributed judgments on the relative demand for analytical and technical skills across sectors.

Posada, J. (2022). Coloniality of datawork: Power inequality in outsourced data production [Doctoral dissertation]. https://liverpool.idm.oclc.org/login?url=https://www.proquest.com/dissertations-theses/coloniality-datawork-power-inequality-outsourced/docview/2762935719/se-2

Doctoral dissertation on Venezuelan data workers demonstrating that data labelling labour is interpretive rather than mechanical, with workers' expertise systematically excluded from the datasets they produce. Primary source for the interpretive-labour argument in the literature review.

PricewaterhouseCoopers New Zealand. (2023). Artificial intelligence study: Sizing the prize. PwC. https://www.pwc.co.nz/insights-and-publications/2023-publications/artificial-intelligence-study.html
Commercial PwC report projecting AI's GDP impact at up to USD 15.7 trillion by 2030, used solely as scene-setting context for the structural inequity argument rather than as primary evidence for any analytical claim.

Regilme, S. S. F. (2024). Artificial intelligence colonialism. SAIS Review of International Affairs, 44(2), 75–92. https://dx.doi.org/10.1353/sais.2024.a950958
Peer-reviewed Johns Hopkins article applying a structural colonialism framework to document AI-driven environmental damage, labour exploitation, and governance gaps in the Global South. Provides conceptual grounding for the structural dependency framing and supports the accountability gap argument.

Rorissa, A. (2025). Personal communication [Professor, University of Tennessee Knoxville School of Information Sciences — active in AI workforce development initiatives in Ethiopia; Academia, Global South].
Academic specialist in information sciences with active engagement in AI workforce development in the Global South. Contributed expert perspective on academic and international dimensions of AI skill demand.

Santos, C. A. dos. (2025). Personal communication [Professor Titular, Universidade Federal de Alfenas-MG (UNIFAL-MG) — researcher in Geosciences, Geotechnologies, and AI in Education; Academia, Brazil/Global South; interview conducted in Portuguese].
Brazilian academic researcher with expertise spanning geosciences, geotechnologies, and AI applications in education. Contributed expert judgments from a Global South perspective on AI infrastructure governance and educational equity.

SOMO. (2024, January 30). Big Tech sets unfair terms and conditions for AI data workers globally. https://www.somo.nl/big-tech-sets-unfair-terms-and-conditions-for-ai-data-workers-globally/
Corporate accountability report documenting systematic NDA use, task-based payment, and contractor misclassification across major AI data platforms via document analysis and worker surveys. Supports the labour section's argument that suppressive contracting is systemic rather than isolated.

Stacciarini, J. H. S., & Gonçalves, R. J. A. F. (2025). Data centres, critical minerals, energy, and geopolitics. Sociedade & Natureza, 37(1). https://doi.org/10.14393/SN-v37-2025-77215x

Brazilian peer-reviewed article mapping hardware-level mineral demand (gallium, germanium, tantalum, REEs) for data centre servers, chips, and cooling systems, including the finding that a single router requires up to 500 kg of raw materials. Primary reference for hardware-level mineral claims in the environmental section.

The Green Grid. (n.d.). Water usage effectiveness (WUE) benchmark. https://www.thegreengrid.org

Industry consortium benchmark of 1.9 L/kWh applied to translate AI data centre electricity consumption into water demand estimates. Acknowledged limitation: WUE varies substantially by cooling technology; figure used as a directional industry-standard estimate.

UNESCO. (2026). From rare earth to classroom: Building geoeconomic justice into digital education. https://www.unesco.org/en/articles/rare-earth-classroom-building-geoeconomic-justice-digital-education

UNESCO field report presenting primary survey data from DRC artisanal miners, finding average wages of $0.34/hour and that 87.9% entered mining due to the absence of alternatives. One of the most recent and direct wage figures available for the labour section.

U.S. Department of Labour. (2024). Congo, DRC: Minimal advancement. ILAB. https://www.dol.gov/sites/dolgov/files/ILAB/child_labor_reports/tda2024/Congo-Democratic-Republic-of-DRC.pdf

Annual federal country assessment classifying the DRC as "minimal advancement" on child and forced labour in mining under the Trade and Development Act. Institutionally validates ILO and Nature Sustainability findings on DRC cobalt labour conditions.

U.S. Geological Survey. (2025). Mineral commodity summaries 2025. USDI. https://pubs.usgs.gov/publication/mcs2025

Authoritative annual reference covering 90+ minerals across 180 countries, providing country-level supply concentration figures: DRC 70%+ cobalt, Indonesia ~50% nickel, Chile copper and lithium. Cross-referenced with the Mineral Ores dataset for site-level disaggregation.

World Bank. (2024). World development report 2024: The middle-income trap. Retrieved May 12, 2026, from https://www.worldbank.org/en/publication/wdr2024

The WDR 2024's "3i" strategy frames how passive resource extraction reinforces the middle-income trap for mineral-supplying countries unless paired with technology transfer and innovation. Limitation: this macro-level framework does not address the specific sector-governance mechanisms required for practical implementation.

World Bank. (n.d.). Commodity markets. Retrieved May 11, 2026, from https://www.worldbank.org/en/research/commodity-markets

Authoritative public source providing global commodity data and forecasts used to contextualise the price dynamics and volatility of AI-critical minerals. Limitation: relies on aggregated global benchmarks and does not capture local or informal artisanal-level pricing realities.

# 10. Appendix A

## Author Contributions

Júlia Rocha coordinated the project, contributed to the data analysis, developed the recommendations, and conducted the expert interviews. Derek Gunter contributed to the data analysis, designed the methodology, and developed the findings about natural resources. Amay Pandey contributed to the data analysis, developed the findings about labour, and produced the data visualisations. Nathaly Avalos led the expert interviews and the qualitative analysis. Krystyna Herasymenko conducted the literature review. All authors contributed to the writing of the final manuscript.