

AI at Half the Price: Harnessing Off-Grid Hydropower and Direct Cooling for Scalable Intelligence

Giorgi Mushkudiani

PhD Student (Informatics), Assistant, Georgian Technical University, Tbilisi, Georgia.

Email: giorgi.mushkudiani@eu.edu.ge

DOI: <https://doi.org/10.68218/rykpzz12>

Abstract

The rapid growth of artificial intelligence has significantly increased semiconductor investment, exposing energy and cooling infrastructure as the primary bottleneck for computational progress [1]. Using the Republic of Georgia as an illustrative case study, this paper proposes a framework for developing a decentralized network of AI data centers co-located with small and medium-sized hydroelectric power plants [11]. This model, which is equally applicable to other hydro-rich geographies such as Canada, Alaska, the Pacific Northwest, Kazakhstan, Nepal, Tibet, Scandinavia, the Alps and the Carpathians in Europe, targets the primary operational costs in AI: electricity consumption and thermal management. By eliminating grid-related transmission losses (10%-20%) and utilizing river water for direct cooling (further 25% to 50% electricity consumption reduction), our approach can reduce total electricity usage by up to 60% in optimal scenarios. This strategy offers a practical and economically sustainable method to support high-performance AI workloads globally.

Keywords: AI, Data Center Infrastructure, Off-Grid, Decentralized Computing, Free Cooling, Inference Cost Reduction.

Introduction

The rapid evolution of artificial intelligence has created the biggest investment boom in the history of the chip industry [3], [6]. As AI models grow in complexity, so do the energy and

cooling demands of the GPU based intensive servers required to train and run them. This has created a critical bottleneck: progress in AI is no longer limited by computational theory but by the physical constraints of power delivery and heat dissipation [1], [5]. Traditional, centralized data centers located in urban areas face several challenges: high electricity costs, grid instability, and inefficient cooling systems that consume up to 50% of operational energy [15]. These inefficiencies create unsustainable economics for high-performance AI workloads.

The Republic of Georgia, with its vast and largely untapped hydroelectric potential [7], [9], is uniquely positioned to solve this challenge. This paper outlines a comprehensive strategy to transform Georgia into a leading AI hub by deploying a network of decentralized off-grid data centers directly at the source of power, its rivers. By building smaller, modular data centers adjacent to hydroelectric plants, we can bypass the inefficient legacy power grid and provide AI services at a fraction of the current cost. Our primary contribution is a comparative model that distinguishes between physical energy savings and financial savings. It evaluates reductions in grid-delivery losses, cooling-energy demand, facility overhead, electricity purchase price, and the resulting energy-related cost of AI inference. Under the modeled assumptions, co-location may avoid approximately 10–20% of grid-delivery losses, support a facility PUE near 1.10, and reduce energy-related inference costs by approximately 20–50%. This will provide a decisive competitive advantage for Georgia (and similar rural regions) in the global technology landscape.

2. Background and Related Work

The energy consumption of large AI models, particularly transformers, is a well-documented problem [4], [5]. The operational cost of a data center is dominated by two factors: the cost of electricity to power the servers and the cost of the infrastructure required to cool them [15]. Conventional cooling methods, such as air conditioning, are notoriously inefficient and can account for 30–50% of a data center's total energy use [15].

While green computing initiatives frequently focus on renewable energy, sources like solar and wind are intermittent and ill-suited for the constant, high-power demands of AI workloads without expensive battery storage. Bitcoin mining has provided a proof-of-concept for co-locating intensive computation directly with cheap energy sources [8], but these efforts have typically lacked the strategic integration necessary for high-reliability AI services. Our work builds upon these precedents by proposing a formalized, decentralized off-grid architecture for AI data centers, that leverage the stability and cooling capacity of hydropower [2], [11], [14], [16].

3. The Object, Subject, and Methods of Research

The object of this research is the energy and cooling infrastructure of high-performance AI data centers. The subject is the strategic co-location of these data centers with small-to-medium-scale hydroelectric power plants (HPPs) in hydro-abundant regions [7], [9], [11].

To evaluate the viability of this strategy, this study employs a comparative theoretical and computational methodology, supported by preliminary testing at a small-scale pilot hydro-station. We utilize **mathematical efficiency modeling** to compare the energy losses and capital expenditures of a traditional, grid-dependent urban data center against a theoretical off-grid, hydro-cooled model. This comparative method was chosen because analyzing cumulative energy waste requires a strict mathematical breakdown of power transmission, transformer steps, and cooling thermodynamics. By isolating and calculating these specific variables, this method provides an objective, quantified assessment of how physical co-location reduces the total cost of AI inference.

The core architectural model rests on three pillars:

Decentralized, Co-located Architecture: In contrast to massive data centers integrated with large-scale power stations, this model evaluates numerous smaller, containerized units. This approach improves resilience and allows for a gradual scaling of AI compute capacity [2].

Vertical Energy Integration: Each theoretical AI node draws power directly from its paired hydroelectric plant. This eliminates the utility company intermediary, granting access to power at the base cost of generation. Furthermore, this direct connection provides a guaranteed revenue stream to the power producer, which incentivizes the development of new renewable projects [10], [17].

Direct Water Cooling: The modeled system uses a two-loop liquid-cooling architecture. The primary closed loop circulates coolant through the servers for direct-to-chip, on-die cooling of the GPUs. A secondary loop uses the same cold river water that passes through the hydroelectric plant to remove heat from the primary loop through a heat exchanger. This type of approach, also known as "free cooling" method [15] requires significantly less energy than traditional air cooling, reducing overhead and allowing for higher-density server racks [15].

4. System Architecture and Cost Reduction Analysis

To clearly distinguish the results of the analysis, the savings produced by the proposed architecture are divided into physical efficiency gains and financial cost reductions. The physical efficiency architecture aims to maximize useful electricity (E_u) through:

1. Grid-delivery energy-loss reduction
2. Cooling-energy reduction
3. Facility-overhead reduction (measured using Power Usage Effectiveness, or PUE)

Financial cost reduction architecture:

4. Electricity unit-cost reduction
5. AI operating-cost reduction (inference-cost reduction)
6. Capital-expenditure savings




Fig. 1 The Conventional Grid-connected vs Proposed Data Center Efficiency Scheme

4.1 The physical efficiency architecture

Unlike other renewable energy sources, hydroelectric power offers both stable electricity and an easily accessible resource of cold water. The proposed architecture targets several distinct sources of energy loss and financial cost in the traditional data-center model.

The useful electricity (E_u) reaching the servers in a conventional setup is a product of cascading efficiency losses: $E_u(\text{Electricity Useful}) = E_g(\text{Electricity Generated}) * T * L * C$

Where:

T = Efficiency per transformer stage, ranging from 95% (for older units) to 99% (for new units). Energy is lost at each voltage step-up and step-down. With an average 2 to 4 transformation stages (from the power plant to high-voltage lines, down to a substation, down to local distribution, and down to the data center), this loss is cumulative. Therefore, the total transformer efficiency is represented as T^n where n is the number of transformation stages

L = Efficiency of high-voltage transmission Lines, where energy is lost to resistance [17].

C = Efficiency of the Cooling system, representing the power used for cooling relative to the power used for computation [15].

Mathematical equation for electricity usage reduction:

$$E_u = E_g * T^4(95\% \text{ to } 99\%) * L(93.3\% \text{ to } 98\%) * C(70\% \text{ to } 40\%)$$

Best case scenario:

$$1MW * 0.99^2 * 0.98 * 0.70 = 0.672MW \text{ (33\% lost to inefficiencies)}$$

Worst case scenario:

$$1MW * 0.95^4 * 0.933 * 0.4 = 0.303MW \text{ (70\% lost to inefficiencies)}$$




Fig.2 Algorithm for electricity usage reduction

Our proposed architecture eliminates most of these inefficiencies:

$L \approx 1.0$ (negligible line losses due to co-location),

$T \approx 1.0$ (bypassing multiple transformers),

$C \geq 0.9$ (liquid cooling with river water).

By reducing total facility overhead to less than 10%, the proposed model maximizes the share of generated electricity delivered to computing equipment and achieves a hypothetical power usage effectiveness (PUE) of ≤ 1.10 . [16] These efficiency gains are complemented by lower electricity prices resulting from direct co-location with the hydropower plant and reduced dependence on conventional HVAC infrastructure.

4.2 Financial cost reduction architecture

Financial and Economic Modeling: The physical efficiency gains are complemented by significant economic benefits. For AI companies, the total cost of ownership is heavily influenced by the operational cost of running inference, which is a direct function of electricity prices [1], [5]. Bypassing the grid operator yields immediate financial benefits.

Electricity unit-cost reduction: The financial savings from purchasing electricity directly at the wholesale generation rate is calculated as:

$$S_{price} = P_{grid} - P_{hpp}$$

To calculate the percentage reduction in electricity cost, the formula is:

$$S_{price(\%)} = \frac{(P_{grid} - P_{hpp})}{P_{grid}} * 100$$

Where:

P_{grid} = The commercial grid electricity purchase price for urban data centers(36 tetri/kWh).

P_{hpp} = The wholesale market rate at which the grid operator purchases electricity from an HPP is roughly 8 tetri (3 cents)/kWh [11].

Calculation based on the Republic of Georgia case study:

(Bypassing the commercial grid rate of 36 tetri/kWh in favor of matching the wholesale HPP market rate of 8 tetri/kWh yields a unit-cost reduction of approximately 78%)

$$S_{price(\%)} = \frac{(36 - 8)}{36} * 100 = 77.7\%$$

Capital Expenditure Reductions: In traditional data centers, electrical and cooling infrastructure accounts for 55–65% of total upfront construction costs [3], [6]. Co-locating with HPPs severely reduces these upfront capital expenditures. By utilizing river water for direct liquid cooling, developers can entirely bypass the need to purchase and install expensive commercial HVAC infrastructure [9].




Fig. 3 Cost of Datacenter Infrastructure [6]

Furthermore, integrating AI data centers with new hydroelectric plants reduces the construction costs of the power plants themselves by 10% to 50% [7], [9]. Typically, 30% to

50% of the capital required to build a new HPP is dedicated to constructing high-voltage transmission lines and infrastructure to connect the plant to the national grid. By functioning entirely off-grid with an anchor tenant on-site, the power developer bypasses these grid connection costs and construction delays entirely.

Operating Cost Reduction: For AI companies, the total cost of ownership is dominated by the operational cost of running inference on trained models, which is a direct function of electricity prices [1], [5]. Bypassing the grid operator yields immediate financial benefits. In the Republic of Georgia, for example, the commercial purchase price of electricity for urban data centers is approximately 36 tetri (13.6 cents)/kWh, whereas the wholesale market rate at which the grid operator purchases electricity from an HPP is roughly 8 tetri (3 cents)/kWh[11].

Category	Affected Component	Evaluation Metric	Estimated Impact / Savings
PART I: Physical Efficiency Gains			
Grid-delivery energy loss	Transmission and transformer losses	Percentage of total generated electricity	10–20% of generated power preserved
Cooling-energy reduction	Cooling-system electricity usage	Percentage reduction vs. conventional air cooling	~90% reduction in cooling energy
Facility-overhead reduction	Non-IT facility electricity overhead	Power Usage Effectiveness (PUE)	$PUE \leq 1.10$ (IT energy share $\geq 90\%$)
PART II: Financial Cost Reductions			
Electricity unit-cost reduction	Commercial electricity price per kWh	Percentage reduction in unit cost	~78% reduction (based on Republic of Georgia case study)
Operating-cost reduction	Energy-related inference cost	Percentage of total operational expenditure	20–50% reduction in inference costs
Capital-expenditure (CAPEX) reduction	Cooling systems & grid-interconnection infrastructure	Percentage of upfront infrastructure CAPEX	10–50% reduction (depending on facility scope)

Table 1: Summary of Efficiency and Cost Reductions

By securing power at rates four-five times lower than commercial grid pricing and reducing cooling energy waste, this framework can reduce total inference costs by 50%. This massive reduction in operating expenses creates a highly attractive environment for compute-intensive sectors such as:

- **AI Cloud Service Providers:** Offering inference-as-a-service at globally competitive prices.
- **High-Performance Computing (HPC):** Supporting scientific and financial computations that are otherwise economically unfeasible.
- **Green AI and ESG-Driven Enterprises:** Providing a 100% renewable, zero-emission compute environment for global companies that must meet strict Environmental, Social, and Governance (ESG) and carbon-neutrality targets.
- **Decentralized AI and Web3 Networks:** Serving as a highly efficient, reliable backbone for Decentralized Physical Infrastructure Networks (DePIN) and blockchain protocols that rely on distributed compute nodes, including large-scale Bitcoin mining operations.
- **Autonomous Vehicle and Robotics Simulation:** Offering the cost-effective, large-scale GPU clusters required by automotive and technology companies to process massive sensor datasets and run millions of hours of driving simulations.
- **Bioinformatics and Genomic Sequencing:** Enabling medical research facilities, universities, and pharmaceutical companies to run complex protein-folding models (e.g., AlphaFold) and process massive biological datasets on tighter research budgets.
- **Large-Scale Rendering Farms:** Attracting the film, gaming, and spatial computing industries, which require thousands of GPUs for continuous, heavy 3D rendering and visual effects generation.
- **Disaster Recovery and Redundant Compute:** Acting as a secure and low-cost backup data center hub for European, Asian, and Middle Eastern enterprise networks.

Because these facilities operate independently of the national power grid, they are completely insulated from widespread transmission infrastructure failures and regional blackouts. Furthermore, integrating satellite internet systems, such as Starlink, provides vital network redundancy during severe regional outages.

If implemented, this framework could position Georgia (and similar hydro-rich regions globally) as highly competitive hubs for decentralized computing. Rather than merely supporting isolated data centers, this approach transforms latent hydroelectric potential into a highly resilient, foundational infrastructure for the broader digital economy.

4.3 Thermodynamic Feasibility and Ecological Impact

To fully contextualize the ecological benefits of this framework, it must be contrasted with the severe environmental externalities of traditional data center thermal management. Conventional hyper-scale facilities frequently rely on evaporative cooling towers, which permanently consume millions of gallons of fresh, potable water daily. This places an immense strain on municipal water supplies and exacerbates water scarcity in drought-prone regions. Alternatively, large-scale air-cooling systems depend on massive HVAC chiller plants and thousands of high-velocity server fans. This infrastructure generates pervasive, low-frequency acoustic noise (often exceeding 70–85 dB)[18], resulting in severe noise pollution that disrupts local wildlife migration and causes significant distress to neighboring human communities.

The proposed off-grid, hydro-cooled architecture eliminates both fresh-water consumption and acoustic pollution. However, verifying the environmental sustainability of this "free cooling" via river water requires addressing its own primary ecological concern: thermal pollution. If the discharge of heated water back into the ecosystem raises the local river temperature (ΔT) significantly, it can disrupt local aquatic life. In hydro-rich regions, mountainous and glacier-fed rivers typically exhibit water temperatures ranging from 6°C to 22°C. To evaluate the viability of our consumptive-free, silent liquid cooling model, we must mathematically model the 1:1 relationship between the hydroelectric flow required to generate power and the thermodynamic capacity of that same water to absorb computational heat.

Mini (100 kW to 1 MW) and small-scale (1 MW to 25 MW) hydroelectric power plants (HPPs) operate with an average hydraulic head (height difference between water intake and the generator turbine) ranging from 10 to 80 meters, with 50 meters serving as a standard

baseline for this model. To determine the water flow rate (Q) required to generate 1 MW of continuous electricity at a 50-meter head, we apply the standard hydroelectric power formula:

$$P = \eta \cdot \rho \cdot g \cdot H \cdot Q$$

Where:

P = Power output (1,000,000 W)

η = Turbine/generator efficiency (0.85)

ρ = Density of water (1,000 kg/m³)

g = Acceleration due to gravity (9.81 m/s²)

H = Net hydraulic head (50 m)

Solving for Q (Flow Rate):

$$Q = \frac{1,000,000}{0.85 \cdot 1,000 \cdot 9.81 \cdot 50} \approx 2.4 \text{ m}^3/\text{s}$$

Generating 1 MW of power under these conditions requires a continuous flow of approximately 2.4 cubic meters per second, or 2,400 liters per second.

In our proposed architecture, this 1 MW of generated electricity is consumed entirely by the co-located AI data center, which subsequently converts it into a maximum of 1 MW (1,000,000 Joules per second) of thermal energy. Even without the integration of a secondary Waste Heat Recovery System (WHRS) to recapture this energy, the maximum heat load transferred back to the river water via the heat exchanger is 1,000,000 J/s.

Using the specific heat capacity formula, we calculate the resulting temperature increase (ΔT) of the 2,400 liters per second of water used for cooling before it is returned to the river:

$$\Delta T = \frac{P_{heat}}{m \cdot c}$$

Where:

P_{heat} = Heat energy added (1,000,000 J/s)

m = Mass flow rate (2,400 kg/s, as 1 liter of water $\approx$ 1 kg)

c = Specific heat capacity of water (4,184 J/(kg·°C))

$$\Delta T = \frac{1,000,000}{2,400 \cdot 4,184} \approx 0.0996^{\circ}\text{C}$$

The mathematical model proves that absorbing the entirety of the 1 MW computational heat load only raises the temperature of the utilized river water by an imperceptible 0.1°C. For instance, if the upstream river water enters the facility at 15°C, it will return to the ecosystem downstream at a negligible 15.1°C.

Furthermore, this thermal efficiency scales favorably. For larger operations (e.g., a 25 MW facility), assuming the hydraulic head remains relatively constant, the additional energy generation relies primarily on a proportional increase in water flow rate. This higher volume of water further dilutes the waste heat being transferred, ensuring that the temperature delta remains at a fraction of a degree.

Because this temperature delta is an order of magnitude lower than the strictest global environmental regulatory thresholds (which typically mandate maximum allowable increases of 1.0°C to 3.0°C), the 1:1 power-to-cooling hydro-architecture is mathematically proven to cause no thermal pollution. By entirely avoiding municipal water evaporation, eliminating acoustic pollution, and maintaining baseline river temperatures, the proposed framework elevates from an economic cost-reduction strategy into a rigorously verified, zero-impact "Green AI" model, highly favorable for regulatory fast-tracking and stringent ESG compliance.

4.4 Limitation

While the proposed off-grid, hydro-cooled architecture offers profound economic and thermodynamic advantages, its implementation is constrained by several infrastructural, geographical, and logistical factors. Recognizing these limitations is essential for identifying the specific operational niches where this model is viable.

Network Connectivity and Bandwidth Constraints: The most significant limitation of remote co-location is the lack of traditional, high-capacity fiber-optic backbone infrastructure. While laying new fiber-optic cables to remote hydroelectric plants is less expensive (\$ 15,000 - \$30,000 per kilometer)[13] than building high voltage power lines (\$1M to \$3 Million+ per kilometer)[10] or constructing traditional data centers in highly populated urban hubs, it remains a notable capital expenditure. Consequently, this model relies on low-cost, decentralized data transmission alternatives, such as SpaceX Starlink (satellite internet), 5G cellular network towers, and Point-to-Point (PtP) wireless bridges utilizing long-distance Wi-Fi dishes with latency of 25 to 50 milliseconds.

Even if we combine all 3 of these alternatives, network cannot guarantee ultra-low latency (<5 ms) or massive bandwidth, thus our architecture is fundamentally unsuited for traditional data center workloads. Applications such as commercial web hosting, cloud gaming, video streaming, and remote-desktop provisioning are incompatible with this framework. Instead, the model is strictly optimized for asynchronous or low-bandwidth AI workloads. Specifically, the vast majority of Large Language Model (LLM) inference relies on text-based data inputs and outputs, which require minimal bandwidth. When coupled with edge-caching strategies, these alternative networks are more than capable of handling distributed AI inference and model finetuning tasks.

Geographical and Topographical Dependence: The framework is geographically constrained to regions with specific topographical and hydrological profiles. It requires abundant, fast-flowing, and cold water sources paired with mountainous or steep terrain to provide the necessary hydraulic head for power generation. While ideal for countries like Georgia, Nepal, Canada, north west USA and specific mountainous ranges (the Alps, the Carpathians), it cannot be universally deployed in flat, arid or tropical regions.

Hydrological Seasonality and Climate Vulnerability: Unlike grid-connected urban data centers that rely on a continuous, multi-source power mix, an off-grid facility is entirely dependent on the localized flow rate of its paired river. Hydroelectric power generation is

susceptible to seasonal environmental changes. While 10-15% fluctuation can be managed by underclocking or overclocking the servers. If during the planning phase the locations with stable base flows are not chosen, prolonged summer droughts or severe winter freezing can significantly reduce water flow, potentially leading to compute throttling, operational downtime, or the need for temporary reliance on backup power systems.

Maintenance Logistics and Physical Security: Decentralizing compute nodes into remote, off-grid terrains introduces complex operational logistics. Traditional data centers benefit from being close to IT talent pools and supply chains. In contrast, dispatching technicians to remote mountainous hydroelectric plants for hardware maintenance, GPU replacement, or server troubleshooting results in a higher Mean Time To Repair (MTTR). Furthermore, ensuring the physical security of expensive AI hardware in isolated, unattended locations requires additional investments in automated surveillance and access control systems.

Conclusion

The future of artificial intelligence and the future of renewable energy are closely linked [1], [4]. At the same time AI compute and blockchain mining represent a significant opportunity for the renewable energy sector, particularly in regions with abundant hydropower. By deploying off-grid AI data centers integrated directly with the hydroelectric sources and utilizing a novel direct water cooling system, this framework addresses the critical energy and thermal bottlenecks currently constraining the industry [15]. As our analysis demonstrates, bypassing the traditional power grid eliminates massive inefficiencies, reducing total electricity consumption by up to 60% and lowering overall inference costs by 20% to 50%.

This model provides a clear, actionable roadmap for building a sustainable, cost-effective AI hub. The architectural framework yields two primary benefits:

Elimination of Intermediaries and Infrastructure Inefficiencies: High-voltage transmission lines lose approximately 6.7% of their carrying capacity for every 1,000 kilometers

due to electrical resistance and ambient temperature [17]. When combined with a 2–5% energy loss at each of the multiple transformer stages, cumulative energy waste becomes substantial. Co-locating data centers directly with power plants eliminates the need for these intermediary lines and transformers, yielding an immediate physical energy savings of 10–20%. Furthermore, this vertical integration removes the utility middleman. In a market such as the Republic of Georgia, bypassing the commercial grid avoids the significant markup between the base generation cost of hydropower (approximately 8 tetri(3cents)/kWh) and commercial data center rates (around 36 tetri(13.6cents)/kWh), unlocking enormous operational savings and profit potential [12].

Symbiotic Economic Growth and Reduced Capital Expenditure: This co-location model creates a powerful economic incentive for new renewable energy development. Securing an AI data center as a guaranteed, high-demand anchor tenant significantly reduces financial risk for power developers. Because the power plant no longer needs to construct high-voltage transmission lines to connect to the national grid, initial capital expenditures can be reduced by 10% to 50% [11]. Simultaneously, the data center operator achieves significant electricity unit-cost reductions and facility-overhead reductions. By securing low-cost power and leveraging free-flowing river water for thermal management [15], the operator largely bypasses the capital expenditure and high maintenance costs associated with traditional HVAC systems [9]. To maximize these benefits, it is highly advisable for hyperscalers or data center developers to pursue co-ownership or deep integration with the power plant during the architectural planning phase to achieve seamless vertical integration.

While geographically constrained to hydro-rich topologies and operationally limited to low-bandwidth, asynchronous AI workloads (making it unsuitable for the training large-scale AI models), our model establishes a highly profitable framework for small and medium-sized AI-first datacenters. Beyond providing a 20–50% reduction in inference costs for hyperscalers, this strategic framework transforms AI infrastructure from a drain on public energy resources into a catalyst for renewable energy growth. It offers an economically superior and environmentally

sustainable path for regions with untapped hydroelectric potential to build a competitive advantage, securing their position in the global technological landscape of the 21st century.

REFERENCES

1. de Vries, "The growing energy footprint of artificial intelligence," *Joule*, vol. 7, no. 10, pp. 2191-2194, Oct. 2023, doi: 10.1016/j.joule.2023.09.004.
2. Z. Wang et al., "Power for AI data centers: Energy demand, grid impacts, challenges and perspectives," *Energies*, vol. 19, no. 3, Art. no. 722, Jan. 2026. [Online]. Available: <https://www.mdpi.com/1996-1073/19/3/722>.
3. Goldman Sachs Global Investment Research, "U.S. data center power demand projected to double by 2027: Powering the AI era," Goldman Sachs, May 2026. [Online]. Available: <https://www.goldmansachs.com/insights/articles/us-data-center-power-demand-projected-to-double-by-2027>.
4. Electric Power Research Institute, *Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption*. Palo Alto, CA, USA: EPRI, May 2024.
5. International Energy Agency, "AI is set to drive surging electricity demand from data centres while offering the potential to transform how the energy sector works," Apr. 2024. [Online]. Available: <https://www.iea.org/news/ai-is-set-to-drive-surging-electricity-demand-from-data-centres-while-offering-the-potential-to-transform-how-the-energy-sector-works>.
6. M. Zhang, "How much does it cost to build a data center?" *Dgtl Infra*, Nov. 2023. [Online]. Available: <https://dgtlinfra.com/how-much-does-it-cost-to-build-a-data-center/>
7. WorldData.info, "Energy consumption in Georgia," 2024. [Online]. Available: <https://www.worlddata.info/asia/georgia/energy-consumption.php>.
8. Cambridge Centre for Alternative Finance, "Cambridge Bitcoin Electricity Consumption Index (CBECI)," Judge Business School, University of Cambridge, 2023. [Online]. Available: <https://ccaf.io/cbeci/index>.
9. G. Mukhigulishvili, "Renewable energy sources," UNDP Georgia, pp. 1-120, 2022. [Online]. Available: <https://www.weg.ge/sites/default/files/undp-georgia-eu4climate-green-deal-assessment-2022-geo.pdf>.
10. Marcy, "EIA study examines the role of high-voltage power lines in integrating renewables," U.S. Energy Information Administration, 2018.
11. Galt & Taggart, "Electricity market in Georgia," Sector Research Rep., pp. 1-25, 2025. [Online]. Available: https://ramad.bog.ge/galt/EFjS4MMq-Energy-conference-presentation_GT_ENG_Site.pdf.
12. Georgian National Energy and Water Supply Regulatory Commission, "Tariffs for hydro power plants," Official Resolution on Marginal Tariffs of Electricity Generated by Hydro Power Plants, 2024. [Online]. Available: <https://gnerc.org/en/tariffs/tariff-el-energy/production/hidroeletrosadgurebis-tarifebi/21554>.
13. U.S. Department of Transportation, "Intelligent Transportation Systems (ITS) costs of fiber optic cable deployment," Federal Highway Administration, 2022.

14. P. Yin, Y. Guo, M. Zhang, et al., "Performance analysis of lake water cooling coupled with a waste heat recovery system in the data center," *Sustainability*, vol. 16, no. 15, Art. no. 6542, Jul. 2024, doi: 10.3390/su16156542.
15. S. Xu, H. Zhang, and Z. Wang, "Thermal management and energy consumption in air, liquid, and free cooling systems for data centers: A review," *Energies*, vol. 16, no. 3, Art. no. 1279, Jan. 2023. [Online]. Available: <https://www.mdpi.com/1996-1073/16/3/1279>.
16. H. E. Burford, L. Wiedemann, W. S. Joyce, and R. E. McCabe, "Deep water source cooling: An untapped resource," pp. 10-15, 2012.
17. World Bank, "Electric power transmission and distribution losses (% of output)," World Development Indicators, 2023. [Online]. Available: <https://data.worldbank.org/indicator/EG.ELC.LOSS.ZS>.
18. Occupational Safety and Health Administration, "Occupational noise exposure," Standard 1910.95, U.S. Department of Labor. [Online]. Available: <https://www.osha.gov/noise>.