Everest Advances in Science and Technology (EAST) Vol. 2, No. 1, 2026 (Special Issue)
ISSN: 3102-0410 (Print), 3102-0429 (Online)
DOI: https://doi.org/10.3126/east.v2i2.99477

Hydropower-Driven AI Data Center Infrastructure in Nepal: A Screening-Level Techno-Economic Assessment Under Run-of-River Seasonality

SAKSHYAM BANJADE[1]

Abstract. Nepal has large hydropower potential, but its run-of-river generation profile creates a reliability challenge for always-on digital infrastructure. This paper develops a hydropower-to-data-center (H2DC) screening framework for Nepal. The framework combines site screening, electrical topology, seismic allowance, cooling suitability, and an All-In Energy Cost (AIEC) equation. Its main contribution is methodological and case-applied: it translates national hydropower abundance into a practical pre-feasibility screen for AI data center siting under seasonal adequacy constraints. Official sector reporting cites 83,000 MW of theoretical hydropower potential and 42,000 MW of economically feasible potential, while installed hydropower reached 2,538.273 MW by FY 2023/24. For data centers, however, the decisive issue is not annual resource size but dry-season supply. Public evidence indicates minimum-to-maximum hydropower output ratios of about 0.33-0.34 for Nepal's run-of-river fleet. Three pathways are assessed: a flexible captive micro-grid (1-5 MW), a grid-backed campus (10-50 MW), and a larger hybrid campus with short-duration storage (50-200 MW). A one-way sensitivity screen shows that procurement tariff and financing cost dominate AIEC, while winter balancing determines whether a workload must be flexible or grid-backed. Under baseline assumptions, modeled AIEC ranges from USD 0.075 to USD 0.115/kWh. This range can support inference, sovereign cloud, and flexible batch computing, but it is not enough to support uninterrupted AI training without firm grid access and stronger market validation.

Keywords: AI infrastructure, data centers, hydropower, Nepal, run-of-river seasonality, seismic resilience, techno-economic analysis

[1] Independent Researcher, Kathmandu, Nepal

*Corresponding author: hi@sakshyambanjade.com.np, ORCID: 0009-0001-4283-9678

Manuscript received: 05 February 2026; Revised: 18 April 2026; Accepted: 09 May 2026.

1. INTRODUCTION

Data center electricity demand has become an energy-system issue. The International Energy Agency estimates that data centers consumed about 415 TWh in 2024, or roughly 1.5% of global electricity use [1]. Future growth will depend on power availability, grid access, and the energy intensity of AI workloads [1]. The recent United States Data Center Energy Usage Report also links rapid demand growth to AI servers and high-density compute clusters [2]. Because data-center energy estimates vary widely, screening studies need transparent assumptions and clear sensitivity tests [3, 4].

Nepal is a strong but complex candidate for hydro-linked data centers. National policy documents cite about 83,000 MW of theoretical hydropower potential and 42,000 MW of economically feasible potential [5]. Installed hydropower, however, stood at 2,538.273 MW by FY 2023/24, and most operating capacity is run-of-river. GIZ/AEPC 2024 reports minimum-to-maximum production ratios of about 0.33 for NEA-owned plants, 0.34 for independent power producers, and about 33.9% of installed capacity for run-of-river IPPs in the dry season [6]. For a continuously loaded AI campus, this seasonal shape matters more than the national headline potential.

This paper asks when Nepal is technically credible and economically defensible for hydropower-linked AI data centers. H2DC is used here to mean the joint screening of hydropower availability, grid-balancing options, data-center engineering requirements, and workload flexibility. AIEC means the delivered facility-electricity cost after infrastructure recovery, power procurement, spur/interconnection allocation, and dry-season balancing. The paper makes four contributions: (1) it defines an H2DC screening framework linking hydro proximity, cooling suitability, seismic exposure, fiber availability, and logistics; (2) it applies one AIEC model consistently across scenarios; (3) it reports a monthly hydro-availability proxy calibrated to public seasonal evidence; and (4) it connects the results to workload fit. The novelty is not a new hydrological forecast. It is an integrated decision screen for a hydropower-dominant developing power system where seasonality, grid rules, and compute flexibility jointly determine feasibility.

2. STUDY BOUNDARY AND H2DC FRAMEWORK

2.1. **Analytical Boundary.** The model begins at electricity procurement and interconnection. It then follows the cost path through on-site electrical and cooling systems to delivered facility electricity cost. It does not model chip performance, software economics, tax structuring, or cloud-service revenue. Site-specific flood, landslide, and geotechnical clearance are also outside the boundary and would be required in a later engineering study.

2.2. **Screening Criteria and Weights.** Candidate corridors are evaluated using the five variables in Table I: hydro proximity, cooling suitability, seismic exposure, fiber proximity, and road/logistics access. The weights reflect first-order cost and feasibility drivers. Hydro proximity receives the largest weight (0.30) because spur distance and interconnection work can sharply change delivered power cost. Cooling suitability follows (0.25) because ambient conditions affect facility design and operating cost. Seismic exposure (0.20) captures resilience requirements under NBC 105. Fiber (0.15) and logistics (0.10) complete the screen. These weights are transparent screening assumptions, not statistically estimated preferences.

2.3. **Illustrative Zone Scoring.** Table 2 uses ordinal scores: 1 = poor, 3 = acceptable, and 5 = best. The scores are based on public screening-stage indicators. Hydro proximity reflects adjacency to major hydropower corridors and likely spur length. Cooling reflects elevation and ambient-temperature advantage. Seismic exposure reflects relative hazard and expected resilience premium. Fiber reflects proximity to trunk or backhaul routes. Road/logistics reflects delivery and maintenance practicality. A ±10% one-at-a-time weight perturbation does not change the top-ranked zone. The scores therefore shortlist corridors for deeper study; they should not be read as investment grades.

TABLE 1. Site-Screening Criteria and Weights

Criterion	Weight	Engineering Rationale
Hydropower proximity	0.30	Controls spur cost and hydro-linked supply practicality.
Cooling suitability	0.25	Improves economizer hours; reduces mechanical cooling.
Seismic exposure	0.20	Drives anchorage, design category, resilience premium.
Fiber proximity	0.15	Influences latency, route redundancy, backhaul cost.
Road and logistics	0.10	Governs module delivery, O&M access, outage response.

TABLE 2. Illustrative Zone Scoring for Initial Screening

Zone	Hyd.	Cool.	Seis.	Fiber	Road	Score
Bhote Koshi corridor	5	4	2	4	3	3.80
Gandaki Valley	4	5	3	3	4	3.90
Karnali corridor	5	5	3	2	2	3.85
Pokhara Valley fringe	3	5	3	4	5	3.85
Eastern hill corridor	4	4	3	3	3	3.55

Note: Gandaki Valley is the most balanced first-screen candidate because it combines hydro adjacency, favorable cooling conditions, and road access. Close scores for Karnali and Pokhara Valley fringe show that the framework is for shortlisting, not final site selection.

2.4. **Cooling and Seismic Envelope.** Cooling is modeled using conservative air-side or indirect-evaporative economizer logic. Moderate elevation can improve economizer hours, but lower air density may increase fan work. The model captures this as a small PUE penalty of 0.01–0.03 for moderate-elevation sites. Server inlet conditions are assumed to remain within ASHRAE TC 9.9 A1/A2-class environmental envelopes where economization is feasible for much of the year [?].
Seismic design is treated as a mission-critical systems issue under NBC 105 [?]. Resilience measures must protect the building frame and acceleration-sensitive equipment, including GPU clusters, power electronics, and cable trays [?, ?]. The model applies a resilience premium of 6–10% for enhanced conventional detailing and 10–18% for cases that include base isolation.

3. Techno-Economic Model

3.1. **Scenario Definitions.** Three deployment pathways are defined in Table 3 to separate flexible operation from strict 24/7 operation.
Scenario A – Flexible micro-grid (1–5 MW): small captive or semi-captive campus serving inference, edge caching, or sovereign digital services. Winter curtailment or workload deferral is allowed. Diesel is restricted to emergency use in the base case.
Scenario B – Grid-backed campus (10–50 MW): hydro-linked facility with dedicated or allocated 132 kV spur and dry-season balancing from the national grid.
Scenario C – Hybrid campus (50–200 MW): larger facility using hydro-linked procurement, short-duration LFP storage, and grid balancing. BESS is modeled as a power-quality and ramping asset, not as seasonal storage.

3.2. **Model Assumptions.**

3.3. **Core Equations.** Annual IT electricity consumption is calculated as:

$$E_{\mathrm{IT}} = P_{\mathrm{IT}} \times LF \times 8760 \tag{1}$$

Annual facility electricity consumption at the supply boundary is calculated as:

TABLE 3. Principal Model Assumptions

ID	Parameter	Value / Range
A1	IT load factor	75%
A2	Annual hydro CF	45% screening average
A3	Wet-season procurement tariff	USD 0.040–0.055/kWh
A4	Spur electrical loss	3–5% (dedicated 132 kV)
A5	Facility CAPEX (excl. land)	USD 1,400–2,200/kW
A6	Seismic allowance	6–18% premium
A7	Fixed O&M	4% of CAPEX per year
A8	WACC	8–12%
A9	Economic life	20 years
A10	LFP BESS	USD 250–350/kWh; 90–92% RTE
A11	Emergency diesel	USD 0.18–0.28/kWh equiv.
A12	Grid carbon intensity	Low vs. fossil markets
A13	Dry-season balancing	USD 0.065–0.080/kWh
A14	Spur cost	USD 0.9–1.2 M/km

$$E_{\text{fac}} = \frac{E_{\text{IT}} \times PUE}{1 - \lambda_{\text{spur}}} \tag{2}$$

The consistent All-In Energy Cost (AIEC) equation used throughout the analysis is:

$$AIEC = LCOE_{\text{infra}} + C_{\text{PPA}} + C_{\text{spur}} + C_{\text{winter}} \tag{3}$$

where the infrastructure cost per delivered unit is:

$$LCOE_{\text{infra}} = \frac{I_0 \times CRF + O\&M_{\text{fix}}}{E_{\text{fac}}} \tag{4}$$

The capital recovery factor is calculated as:

$$CRF = \frac{r(1+r)^n}{(1+r)^n - 1} \tag{5}$$

The dry-season set is defined as:

$$D = \{\text{Dec}, \text{Jan}, \text{Feb}, \text{Mar}\} \tag{6}$$

The winter balancing term, C_{winter}, captures the additional cost of grid-supplemented or curtailed energy during these months. It is weighted by the normalized monthly hydro-availability factor, ϕ_m. In Scenario A, some curtailment is allowed, so the balancing cost is limited. In the always-on off-grid stress case, the balancing tariff becomes diesel-equivalent and the penalty rises sharply. This term connects hydrology to reliability: a low annual average cost can still be infeasible if the dry-season tail cannot support the workload.

3.4. **Sensitivity, Uncertainty, and Market Screen.** Plant-level SCADA data and contracted data-center revenue were not publicly available. For that reason, uncertainty is represented through bounded one-way sensitivity rather than probabilistic forecasting. The tested variables are procurement tariff, WACC, facility CAPEX, winter-balancing premium, PUE, and load factor.

The economic rule is simple: AIEC must fit within the power-cost allowance of the target service model after margin, network cost, operations, and risk premiums. This does not assume Nepal is cheaper than established Asian hubs. It tests whether low-carbon electricity and flexible workloads can offset Nepal's thinner grid redundancy, smaller domestic cloud market, and higher implementation risk.

3.5. **Reproducible Reference Case.** The Scenario B reference case uses $P_{\mathrm{IT}} = 25$ MW, $LF = 0.75$, $PUE = 1.25$, $\lambda_{\mathrm{spur}} = 0.04$, CAPEX = USD 1,800/kW, WACC = 10%, economic life = 20 years, and a procurement tariff of USD 0.047/kWh.
These inputs give $CRF = 0.1175$, an infrastructure component of approximately USD 0.033/kWh, and a total AIEC of approximately USD 0.090/kWh after procurement, spur allocation, and dry-season balancing. Table VII reports the full 3×3 parameter sweep.
The model is a cost and feasibility screen. A bankable investment case would also require service revenue, colocation utilization, taxes, land cost, foreign-exchange exposure, and contracted service-level penalties.

4. Seasonal Validation and Adequacy Screening

4.1. **Evidence Anchoring.** Table IV is anchored to three public sources. First, WECS 2024 confirms Nepal's hydropower-dominant system and installed-capacity trajectory [?]. Second, GIZ/AEPC 2024 reports minimum-to-maximum production ratios of approximately 0.33 for NEA plants and 0.34 for IPPs [?]. It also reports dry-season run-of-river IPP production at approximately 33.9% of installed capacity [?]. Third, the same study shows that Nepal's dry-season adequacy problem requires complementary resources, not only more nameplate capacity [?]. The monthly series is therefore calibrated to public seasonal evidence rather than fitted to proprietary plant telemetry.
The monthly series should be read as a documented proxy, not raw SCADA data. It can identify structural dry-season risk, but it cannot estimate outage probability, hourly ramping, spill, plant maintenance, or local congestion. The next validation step is clear: replace the proxy with monthly and hourly generation/dispatch data from the specific plants and substations serving the candidate corridor.

Table 4. Normalized Monthly Hydro Availability

Month	**Observed Proxy**	**Modeled Series**
January	0.34	0.35
February	0.32	0.33
March	0.30	0.31
April	0.33	0.34
May	0.46	0.45
June	0.62	0.61
July	0.82	0.81
August	0.93	0.92
September	0.95	0.94
October	0.78	0.77
November	0.56	0.55
December	0.39	0.40

Note: Values are normalized to wet-season peak = 1.0. December–March is the structural deficit window for always-on loads. Short-duration BESS can smooth ramps but cannot replace seasonal energy.

4.2. **Monthly Normalized Hydro Availability.** Figure 1 shows the main engineering result from Table 4. The normalized hydro profile falls to 0.30–0.40 in December–March while a 24/7 data-center load remains flat. This creates a structural deficit that short-duration batteries cannot bridge. Scenarios B and C are therefore stronger than a pure off-grid Scenario A for workloads that require uninterrupted operation. The figure also explains the workload conclusion: inference, batch processing, and sovereign cloud can be scheduled, throttled, or grid-backed, while continuous frontier-model training is exposed to the dry-season trough.

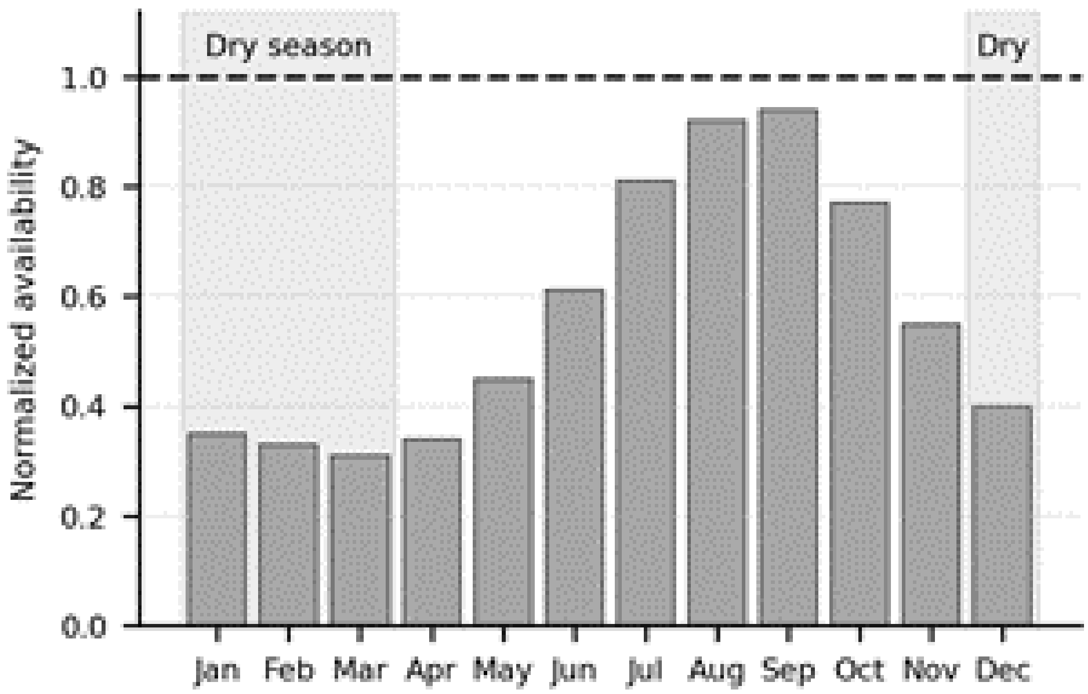

Figure 1. Seasonal adequacy screening using the modeled monthly hydro availability series. The shaded windows show the December–March dry-season deficit for always-on loads.

Table 5. Modeled AIEC Ranges by Scenario

Scenario	**Scale**	**AIEC (USD/kWh)**
A - Flexible micro-grid	1–5 MW	0.090–0.115
A - Always-on off-grid + diesel	1–5 MW	0.140–0.175
B - Grid-backed campus	10–50 MW	0.080–0.100
C - Hybrid + short-duration BESS	50–200 MW	0.075–0.095

Note: Scenario B is the most balanced near-term topology. Scenario A works only with flexible winter operation. The diesel stress case shows why hydropower proximity alone is not a reliability strategy.

Table 6. Base-Case AIEC Decomposition (USD/kWh)

Component	**Scenario A**	**Scenario B**	**Scenario C**
$LCOE_{\text{infra}}$	0.028	0.033	0.022
C_{PPA}	0.049	0.047	0.045
C_{spur}	0.016	0.004	0.005
C_{winter}	0.008	0.006	0.014
AIEC (Total)	**0.101**	**0.090**	**0.086**

Note: Procurement tariff and cost of capital dominate AIEC. Cooling efficiency and spur losses matter, but they do not drive the result within the assumed range.

Table 7. Scenario B AIEC by Procurement Tariff and WACC

WACC	**USD 0.040/kWh**	**USD 0.047/kWh**	**USD 0.055/kWh**
8%	0.080	0.087	0.095
10%	0.083	0.090	0.098
12%	0.087	0.094	0.102

Note: Procurement tariff has the largest effect, followed by WACC, facility CAPEX, and dry-season adequacy. PUE changes shift AIEC less within the assumed range but remain operationally important.

5. RESULTS

Nepal is best suited to AI inference, sovereign cloud, and flexible batch HPC. These workloads can tolerate grid-backed operation, location-aware scheduling, or time shifting. Recent AI-energy literature shows that training, fine-tuning, and online inference have different temporal and power-density profiles [3, 7]. Inference and sovereign cloud can run at high utilization with grid support. Flexible batch workloads can move away from dry-season or high-cost hours. Continuous frontier-model training is a weaker fit because it needs dense, uninterrupted power, strong route redundancy, and high utilization through the December–March trough. Latency-critical global hyperscale services are also weak because energy quality alone cannot replace carrier diversity, route redundancy, and a mature customer ecosystem.

The modeled AIEC range is therefore necessary but not sufficient for competitiveness. Regional data-center markets compete on power cost, uptime, tax treatment, fiber diversity, customer proximity, and permitting certainty. Nepal's advantage is low-carbon hydro-linked power. Its risks are seasonal adequacy, institutional readiness, smaller domestic demand, foreign-exchange exposure for imported equipment, and limited empirical cost benchmarks. A realistic entry path would begin with sovereign workloads, public-sector cloud, disaster-recovery capacity, inference services, and flexible compute before attempting hyperscale training.

6. DISCUSSION AND POLICY IMPLICATIONS

Nepal's main constraint is resource shape, not resource scarcity. Hydropower offers low-carbon electricity, but the seasonal profile is poorly aligned with nonstop AI training.

Grid access matters more than symbolic direct-hydro claims. The ERC five-year roadmap prioritizes open access in transmission, wholesale competition, and transparent tariffs [8]. The Open Access Directives 2082 (2025 AD) signal movement toward formal transmission access rules [9]. For a data-center operator, bankability depends on wheeling conditions, balancing treatment, curtailment priority, dispute resolution, outage compensation, and interconnection procedure clarity [10]. The key question is not only whether open access exists on paper. It is whether operators can obtain predictable capacity rights, transparent charges, and enforceable dry-season service quality. NEA's 2026 tariff-based competitive bidding documentation for transmission service providers shows active work on 132 kV and 400 kV projects under BOOT structures [11]. This supports the broader claim that transmission expansion is part of Nepal's current institutional trajectory.

Seasonal balancing should be treated as a design variable. Solar output tends to strengthen when run-of-river hydro weakens [6]. A hydro-plus-solar-plus-grid architecture could help preserve Nepal's carbon advantage. It would also introduce land, permitting, grid-code, and dispatch-coordination questions. This next-step design question should be tested with hourly hydro, solar, load, and transmission-congestion data rather than annual averages.

Policy readiness extends beyond electricity. Data sovereignty, cross-border data transfer, cybersecurity assurance, environmental permitting, cooling-water use, e-waste handling, and disaster resilience all affect whether customers would trust Nepal as a data-center location [15]. A practical policy package would combine open-access implementation, standardized interconnection studies, renewable-power certification, data-protection rules, and resilience requirements for mission-critical facilities.

7. LIMITATIONS

This paper remains a pre-feasibility screen. The monthly adequacy series is a documented proxy, not plant telemetry. It cannot replace hourly generation, dispatch, outage, spill, and substation-loading records. CAPEX, BESS, and spur-cost assumptions use public benchmarks rather than contractor quotations. Site-specific terrain, import duties, land acquisition, financing currency, tax treatment, and construction logistics could materially change actual cost. The connectivity assessment is directional and does not include a telecom engineering model. Climate-driven hydrological shifts over the 20-year facility life are not modeled. Revenue, colocation pricing, utilization ramp, foreign-exchange risk, sovereign-data requirements, cybersecurity compliance, environmental permitting, water use, e-waste management, and embodied emissions are also outside the model. These omissions do not invalidate the screen, but they mark the boundary between this paper and an investment-grade feasibility study.

Data availability statement: All numerical assumptions used in the screening model are reported in the manuscript tables. The hydro-availability series is a normalized proxy calibrated to public sector evidence. Plant-level SCADA, hourly dispatch, and commercial tariff contracts were not available to the author. Future validation should use utility-approved plant and substation data for the selected corridor.

8. CONCLUSION

Under the stated assumptions, AIEC spans roughly USD 0.075–0.115/kWh across the principal deployment pathways. This range justifies further screening, but it does not remove the need to solve financing, transmission, winter balancing, regulation, and customer-revenue risk.

Nepal's constraint is seasonal adequacy, not hydropower potential in the abstract. Dry-season output near one-third of wet-season peak is structurally difficult for uninterrupted AI training. Inference, sovereign cloud, regional cloud, and flexible batch workloads fit Nepal's resource profile more credibly.

The most defensible near-term topology is a 10–50 MW grid-backed campus in the Gandaki Valley or a similar corridor. Such a project would need transparent wheeling rules, firm interconnection rights, and a workload mix that can tolerate seasonal flexibility. The highest-value next steps are to obtain monthly and hourly plant-level generation data, benchmark regional data-center energy and service revenues, test a hydro-solar complementarity case, and extend the model from deterministic screening to probabilistic reliability analysis.

References

[1] International Energy Agency, "Energy and ai," International Energy Agency, Paris, France, Tech. Rep., 2025.

[2] A. Shehabi *et al.*, "2024 united states data center energy usage report," Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep., 2024.

[3] E. Masanet, A. Shehabi, N. Lei, S. Smith, and J. Koomey, "Recalibrating global data center energy-use estimates," *Science*, vol. 367, no. 6481, pp. 984–986, 2020.

[4] IEA 4E EDNA, "Data centre energy use: Critical review of models and results," International Energy Agency Technology Collaboration Programme, Paris, Tech. Rep., 2025.

[5] Water and Energy Commission Secretariat, "Energy sector synopsis report 2024 (fy 2079/80)," Water and Energy Commission Secretariat, Kathmandu, Tech. Rep., 2024.

[6] Alternative Energy Promotion Centre and GIZ, "Complementing hydropower with utility-scale solar electricity in nepal," Alternative Energy Promotion Centre and GIZ, Kathmandu, Tech. Rep., 2024.

[7] J. Dodge *et al.*, "Measuring the carbon intensity of ai in cloud instances," in *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, Seoul, Republic of Korea, 2022, pp. 1877–1894.

[8] Electricity Regulatory Commission, "Five-year roadmap for erc," Electricity Regulatory Commission, Kathmandu, Tech. Rep., 2024.

[9] ——, "Electricity regulatory commission (open access) directives, 2082 (2025 ad)," Electricity Regulatory Commission, Kathmandu, Tech. Rep., 2025.

[10] International Telecommunication Union, "Recommendation itu-t l.1304: Procurement criteria for sustainable data centres," International Telecommunication Union, Geneva, Tech. Rep., 2020.
[11] Nepal Electricity Authority, "Expression of interest for transmission service provider under tariff based competitive bidding," Nepal Electricity Authority, Kathmandu, Tech. Rep., Mar. 2026.