TELBORG

Data Center Briefing

February 07, 2026

Global

Key themes:
BYO and clean transition tariffs for large loads; Clean energy procurement structures tied to data centre growth; Post-quantum cryptography migration and policy agenda; Nordics positioned for renewable-powered AI data centres; GPU shortages driving interest in neocloud capacity; Operational efficiency research for LLM inference on H100s; Compute-load flexibility as a grid support concept; Latency/NUMA performance tooling for infrastructure operations; Data-intensive science workflows scaling storage and analytics

Global Data Centres Briefing (UTC 2026-02-07)

Audience: Institutional asset managers and infrastructure fund managers (data centres, power, grid)

Top news (3)

1. **New utility tariff structures proposed to accelerate clean power for data centres**

- An RMI insight brief recommends **Bring-Your-Own (BYO)** and **Clean Transition Tariffs** to speed up clean energy procurement for large loads (including data centres) while **insulating other ratepayers**. Examples cited include **PJM's Bring Your Own New Generation Program (>250 MW)**, **Evergy Kansas's Clean Energy Choice Rider**, **NV Energy's Clean Transition Tariff (used by Google)**, and **Georgia Power's Customer-Identified Resource option**: Bring-Your-Own and Clean Transition Tariffs for Data Centers.

2. **Google pushes for faster post-quantum cryptography (PQC) transition across critical infrastructure**

- Google says it has been preparing since **2016**, is **rolling out PQC across its infrastructure**, and expects to complete migration **within NIST's guidelines**. It also issued **five recommendations for policymakers**,

including making **AI systems PQC-ready** and prioritizing **cloud-first modernization**: Google urges policymakers to prepare for quantum-era cryptography transition.

3. **Nordics positioned as a renewable-powered, AI-ready data centre hub (with heat reuse highlighted)**

- A report frames the **Nordic region** as a model for sustainably scaling AI-ready capacity, citing **renewable energy**, **efficient power markets**, and **cool climate**, plus Denmark examples such as **atNorth using district heat reuse**: Nordic region powers AI data centers with renewable energy.

Key deals and projects

Nordics (regional theme rather than a single transaction)

- The Nordic region is positioned as supportive for AI data centres due to **renewable supply**, **power-market efficiency**, and **cool climate**, with **district heat reuse** cited in Denmark (atNorth example): Nordic region powers AI data centers with renewable energy.

Note: No M&A, financing rounds, lease transactions, or capex figures were included in today's story set.

Power and grid / interconnection highlights

US (tariff and utility program structures)

- RMI's brief focuses on procurement mechanisms designed for large loads:
  - **BYO structures** (example: **PJM Bring Your Own New Generation Program** noted as **>250 MW**) to connect new supply with large customers.
  - **Clean Transition Tariffs** (example: **NV Energy's Clean Transition Tariff used by Google**) aimed at enabling clean procurement while limiting impacts to other ratepayers.
  - Additional examples: **Evergy Kansas Clean Energy Choice Rider** and **Georgia Power Customer-Identified Resource option**.
  - Implication for investors: these programs can influence **time-to-power**, **contracting pathways**, and **who bears system costs** for incremental generation tied to new data centre demand: Bring-Your-Own and Clean Transition Tariffs for Data Centers.

Grid-support concept from compute operations (research)

- A research proposal suggests data centre GPU loads could be modulated for **distribution voltage regulation**, by adjusting **batch size in LLM inference** to change GPU power draw (reducing GPU power to alleviate **lower-voltage violations**; increasing power to mitigate **upper-voltage**

violations): GPU-to-Grid: Distribution Voltage Regulation via GPU Utilization Control.

Policy and regulation

Cybersecurity / national infrastructure readiness (global)

- Google's PQC update is framed as both an implementation program and a policy agenda:
  - Migration underway across Google infrastructure; target completion aligned with **NIST guidelines**.
  - Policy recommendations include making **AI systems PQC-ready** and prioritizing **cloud-first modernization**.
  - Investor relevance: PQC timelines can drive **security-related capex/opex**, procurement cycles, and compliance expectations for critical digital infrastructure: Google urges policymakers to prepare for quantum-era cryptography transition.

Compute capacity, supply constraints, and operational efficiency (market signals)

GPU supply constraints and "neocloud" positioning

- A vendor article argues **GPU shortages** are hindering university research and proposes **neocloud GPU rentals** as a mitigation, offering hourly access to enterprise-grade **NVIDIA GPUs** via pre-configured VMs and a stated **15% coupon code**; it also notes shortages can reshape research priorities and widen inequality across institutions: GPU shortages hinder university research, neoclouds as solution.

Efficiency and performance research (potential longer-term impacts on power density and cost-to-serve)

- **LLM inference energy-efficiency "sweet spots"**: an analytical model validated on **NVIDIA H100** using **TensorRT-LLM**, across **1B–9B** parameter models and **64–4096** token sequences; reported mean **MAPE 1.79%**, recommending alignment of sequence lengths to reduce production energy usage: Energy Efficiency Sweet Spots in Production LLM Inference.
- **Serving throughput improvement**: TIDE proposes integrating online draft adaptation into inference, reporting up to **1.15x throughput improvement** vs static speculative decoding and **1.67x reduction** in draft training time: TIDE: Temporal Incremental Draft Engine for Self-Improving LLMs.
- **Virtualization and latency dataset**: a published dataset varying CPU affinity, packet injection frequency, virtual network driver, load types, and VM counts to study end-to-end latency; positioned to support ML-based

admin decision-making: Data analysis of cloud virtualization experiments and network latency.

- **NUMA migration technique (user-space)**: proposes `page_leap()` for fine-grained NUMA page migration (small and huge pages), as an alternative to Linux automatic NUMA balancing and `move_pages()`: Efficient and Reliable Fine-Grained NUMA Migration in User-space.

Data-intensive science workload signal (storage and workflow scale)

- A synchrotron facility (CHESS) deployed a data acquisition and analytics framework on a secure server with real-time web tools, implemented across beamlines **ID3A/ID3B/ID4B**, managing **50–100 TB** and **10+ million files**, tested with **43 research groups** and **86 dashboards**: Large Data Acquisition and Analytics at Synchrotron Radiation Facilities.

2-line close

Tariff designs that pair new generation with large loads are being positioned as a practical lever to speed clean power procurement for data centres without shifting costs to other customers.
In parallel, operators are being pushed by both security transition requirements (PQC) and GPU-driven scaling constraints, with efficiency research pointing to incremental gains in power and performance over time.