

# Data Centre Briefing

February 17, 2026

Global

## Key themes:

NVIDIA GB300 NVL72 efficiency claims and cloud deployments; Rising focus on throughput-per-megawatt and cost-per-token metrics; Liquid cooling modularisation (cold plates, manifolds, CDUs) as AI enabler; Storage architecture narratives around linear scaling and validation (NCP); Vendor lock-in vs open integration positioning in private AI stacks; Sustainable design messaging: climate, renewables, and low-impact materials (Iceland)

## Top news (Global)

- 1. GPU platform efficiency claims are rising sharply.** NVIDIA says its new Blackwell Ultra system offers major throughput-per-megawatt and cost-per-token gains, and it is already being deployed by several large cloud providers: [NVIDIA GB300 NVL72 boosts agentic AI performance and efficiency](#).
- 2. Storage vendors are pushing “AI-scale” architectures and third-party validation narratives.** Dell argues its PowerScale discrete, scale-out design delivers predictable linear scaling and sustained bandwidth, citing NVIDIA Cloud Provider (NCP) validation and Dell internal testing: [Dell PowerScale: Discrete Architecture for AI-scale Storage and Efficiency](#).
- 3. Liquid cooling is being positioned as an essential enabling layer for AI data centres.** Boyd highlights modular liquid cooling building blocks (cold plates, manifolds, CDUs) and simulation tools aimed at scaling deployments from edge to large AI campuses: [Scalable liquid cooling essential for modern AI data centers](#).

---

## Key deals & projects

### Europe (Nordics) — Iceland

- **atNorth ICE03 (Akureyri, Iceland): design choices tied to climate and materials**

- Feature highlights the site's use of Iceland's cool climate and renewable energy to support **efficient direct liquid cooling**.
- Construction materials called out include **Glulam laminated wood** and **locally produced Icelandic rockwool**, positioned as lower environmental-impact choices while supporting scalable buildout: [atNorth's ICE03 in Akureyri showcases sustainable data centre design](#).

## Global — AI infrastructure platforms (vendor and cloud ecosystem signals)

- **NVIDIA GB300 NVL72 (Blackwell Ultra): deployment signals from major cloud providers**
    - NVIDIA claims **up to 50x higher throughput per megawatt** and **up to 35x lower cost per token** vs Hopper, targeting low-latency “agentic AI” and coding assistants.
    - Named cloud adopters “deploying in production” include **Microsoft, CoreWeave, and OCI**.
    - NVIDIA also points to the next-generation **Rubin/Vera Rubin NVL72** with a stated goal of another **~10x throughput-per-megawatt gain** vs Blackwell: [NVIDIA GB300 NVL72 boosts agentic AI performance and efficiency](#).
- 

## Power, cooling, and facility efficiency highlights

### Cooling (thermal as a primary scaling constraint)

- **Boyd: modular liquid cooling + simulation-driven design**
  - Positions **cold plates, manifolds, and CDUs** as modular components to scale installations.
  - Promotes **SmartCFD simulation** to support design and rollout.
  - Frames demand backdrop as data centre electricity consumption projected at **~448 TWh in 2025** and expected to **more than double by 2030**.
  - Notes global manufacturing capacity and regional service coverage across **North America, Europe, and Asia**: [Scalable liquid cooling essential for modern AI data centers](#).

### Storage density and power draw (IT layer efficiency messaging)

- **Dell: PowerScale discrete scale-out storage positioned for lower footprint/power**
    - Dell argues for **predictable linear scaling, sustained bandwidth, and stable write performance** for AI workloads.
    - Cites **NVIDIA Cloud Provider (NCP) validation** and a Dell internal analysis of **64 SU configurations (August 2025)**.
    - Blog claims **lower rack space and power draw** versus alternatives (also referencing competitor pivots such as “VAST's E-Box evolution”): [Dell PowerScale: Discrete Architecture for AI-scale Storage and Efficiency](#).
- 

## Market structure / platform positioning (vendor lock-in vs openness)

- **Dell critique of HPE Private Cloud AI**

- Dell characterises HPE’s “three-click” Private Cloud AI as **prescriptive, GreenLake-dependent**, and creating **vendor lock-in** plus **Day-2 scaling challenges**.
  - Dell positions **Dell AI Factory** as a more open, validated alternative intended to integrate with existing management tools: [The three-click illusion of HPE’s Private Cloud AI](#).
- 

## **Policy and regulation**

- No policy, permitting, grid interconnection, or regulatory changes were included in today’s stories.
- 

## **Two-line wrap-up**

Hardware/platform vendors are increasingly competing on measurable “throughput-per-megawatt” and operational scalability narratives.

Material choices, cooling architectures, and validated reference designs continue to be used to differentiate new AI-focused data centre builds and retrofit strategies.