

Data Centre Briefing

February 23, 2026

Global

Key themes:

Vertiv Q4 orders +252% as \$650bn spend cited; atNorth 300MW Sollefteå Sweden campus targeting H1 2028; FERC EL25-49 enables PJM non-firm “connect while building”; Gujarat ₹25,000 crore MoU for L&T Vyoma 250MW Dholera

Vertiv, Eaton, GE Vernova and Legrand aren’t just “benefiting” from data centre growth — they’re getting dragged into it at full speed. In one of the cleanest demand signals you’ll see, Vertiv’s Q4 orders jumped 252% and its backlog now covers more than 100% of its 2026 sales guidance, on a day when the industry is being told hyperscalers plan **\$650bn** of data-centre and power-equipment spending this year. That kind of pull-through changes who has leverage in the stack: not the operator, but the people who can actually ship switchgear, UPS, thermal, and grid gear.

The Big Stories

[Electrical equipment makers race to serve booming data centers](#) lays out just how fast the “pick-and-shovel” side is tightening. Vertiv says data centres are now 80–85% of revenue, orders are exploding, and the backlog is effectively a forward book for next year’s sales. The important part isn’t that demand is strong — it’s that OEMs are increasingly dictating delivery windows and (implicitly) design choices. Investors should read this as a constraint story: when the equipment queue lengthens, the whole commissioning calendar shifts right.

[Flexible onsite generation accelerates hyperscale data center grid interconnections](#) is a reminder that the interconnection bottleneck is turning into a product category. Enchanted Rock is pitching dispatchable onsite generation as a way for a **500MW** campus to run **three to five years sooner**, and claims **\$78m per GW** in grid-cost reductions. The policy hook matters: FERC Order **EL25-49** is directing PJM to create non-firm pathways so projects can “connect while building,” with other regions

sniffing around similar ideas. If this spreads, “temporary” onsite power stops being a stopgap and starts looking like a standard phase of hyperscale delivery.

[atNorth to build 300MW data center campus in Sollefteå](#) is another big Nordic swing: **300MW** on a **50-hectare** plot in Långsele (Sollefteå), targeted for operations in **H1 2028**, using atNorth’s modular architecture, renewables, and heat-reuse partnerships. The signal here is that the Nordics remain one of the few places where “AI-scale” announcements can still lean on a credible clean-power narrative. The watch item is execution timing: by 2028, grid access and delivery discipline will matter as much as land and green electrons.

[Gujarat signs MoU with L&T Vyoma for 250MW data centre](#) puts another large marker down in India: a ₹25,000 crore MoU for a **250MW** “green, AI-ready hyperscale” campus in **Dholera SIR**, expected by **2028**, with a feasibility study to follow. India has plenty of ambition; what’s notable is the continued move toward named, MW-sized commitments rather than generic “digital infrastructure” talk. For competitors, this is pressure to show comparable scale (and credible power plans) in India’s top industrial corridors.

[NVIDIA and partners bring AI-driven zero trust to OT](#) is NVIDIA pushing deeper into the operational layer of “AI factories,” not just the compute layer. NVIDIA is teaming with Akamai, Forescout, Palo Alto Networks, Xage Security and Siemens around BlueField DPUs, agentless segmentation, local inspection/enforcement at the edge, and centralized AI threat analysis — with demos slated for S4x26 in Miami. This matters because OT security is becoming a first-order availability problem as campuses add more on-prem power, cooling, and industrial control complexity; NVIDIA is positioning DPUs as part of the control plane, not an optional acceleration card.

Behind the Headlines

[Hot-water cooling gains traction for AI data centers](#) is less about a clever thermals headline and more about capex geometry. NVIDIA says its Vera Rubin processor can be cooled with **45°C water**, eliminating chillers and pushing designs toward warmer liquid loops. If that holds at scale, it changes the cost and footprint allocation between mechanical plant and IT — and it pulls the cooling ecosystem (IBM, Lenovo, LiquidStack, Accelsius, SuperMicro, Vertiv, Schneider) into a faster replacement cycle. The underappreciated angle: warmer-water designs don’t just save energy; they can also simplify where and how you can site capacity, because “chillerless” is a permitting and infrastructure story as much as an efficiency story.

[Legacy LonTalk protocol exposes building systems to cyber risk](#) is the kind of thing operators ignore until they can't. Claroty is pointing to LonTalk still embedded across BMS deployments, with Internet-accessible controllers exposing LonTalk-over-IP (CEA-852), and says about **75%** of organisations manage BMS devices with known exploited vulnerabilities. Data centres often treat BMS as “facilities plumbing,” but it's also an attack surface that can translate cyber risk into physical disruption. As sites get denser and more automated, the separation between “IT security” and “facility reliability” keeps shrinking — and legacy protocols are where those worlds collide.

[Sam Altman rebuts water claim amid AI environmental debate](#) captures the communications trap the whole sector is walking into. Altman is calling the viral “one ChatGPT query = a bottle of water” statistic “completely untrue,” while still conceding AI's energy use is a real concern; the piece points to Microsoft's **\$80bn** AI data centre commitment and an IEA view that data-centre electricity demand could double by **2026**. The deeper issue isn't whether a single stat is wrong — it's that public scrutiny is shifting from abstract carbon talk to local, countable resources like water and grid capacity. If the industry can't explain impacts in plain numbers (and with credible boundaries), campaigners and regulators will do it for them.