

Data Centre Briefing

April 01, 2026

Global

Key themes:

Nebius \$10B 310MW AI campus in Lappeenranta, Finland; Nvidia \$2B Marvell deal to extend NVLink Fusion rack-scale AI; Michigan approves 1,332MW BESS backing 1,383MW Oracle data centre; Amazon \$20.5M Eastern Oregon settlement over cooling-water pollution claims

Nebius just put a giant marker down in Europe: a \$10 billion, 310MW AI data centre campus planned for Lappeenranta, Finland, targeted to come online in 2027. What makes it more than just another “hyperscale in the Nordics” announcement is the context—Nebius is tying it directly to a multibillion-dollar agreement with Meta and a \$2 billion investment from Nvidia. This is the shape of the next phase of AI infrastructure: fewer, bigger bets, and the winners are the ones who can actually land power, sites, and hardware at scale.

The Big Stories

[Nebius unveils \\$10B, 310 MW AI data center in Finland](#) is the day’s headline for a reason. The company says the 310MW Lappeenranta campus will be online in 2027, coming after a multibillion-dollar agreement with Meta and a \$2 billion investment from Nvidia. Nebius also flagged additional AI capacity in France and said it expanded its Mäntsälä site to 75MW. The signal here is blunt: European AI compute is moving from “plan and permit” to “commit real capital,” and suppliers (power, grid, and kit) will start picking favourites based on who can execute.

On the hardware side, [Nvidia, Marvell partner with \\$2B to expand NVLink](#) is a tell about where rack-scale design is headed. Nvidia is putting \$2 billion into Marvell to integrate custom XPU, optical DSP, and silicon photonics into NVLink Fusion, explicitly pushing “heterogeneous, semi-custom AI infrastructure.” This matters because it’s not just about more GPUs—it’s about interconnect, optics, and bespoke accelerators becoming first-class constraints in data centre design, procurement, and timelines.

Grid-scale storage just got a very data-centre-shaped endorsement in the Midwest. [Michigan approves six BESS projects totalling 1,332MW capacity](#) for DTE Electric includes three company-owned BESS projects (332MW initially) that will support a 1,383MW Oracle (Green Chile Ventures) data centre, with Oracle covering development costs over 15 years while DTE owns and operates the assets. The key takeaway: utilities are increasingly willing to build “bespoke” flexibility to make mega-loads workable—but the commercial structure (who pays, who owns, who takes risk) is becoming as important as the megawatts.

Environmental externalities keep turning into line items. [Amazon agrees \\$20.5M settlement over data-center water pollution](#) resolves a class action with Eastern Oregon residents over alleged nitrate contamination linked to data-centre cooling-water discharges. The \$20.5 million will fund private well projects and public water-system treatment after attorney fees, while litigation continues against other alleged polluters. For operators and investors, this is a reminder that “water strategy” isn’t a sustainability slide anymore—it’s litigation exposure, permitting friction, and reputational risk in one package.

Edge and sovereignty are colliding in a very practical way. [Microsoft and Armada deliver sovereign AI to the edge](#) pairs Microsoft Sovereign Private Cloud and Azure Local with Armada’s Galleon modular datacenters, aimed at disconnected and regulated environments with resilient connectivity options (satellite, LTE/5G, RF, SD-WAN). The point isn’t just a partnership announcement; it’s that “sovereign AI” is moving out of policy decks and into shippable, modular infrastructure that can live outside the traditional metro data centre footprint.

Behind the Headlines

Ontario is starting to talk about AI data centres like a system-planning problem, not a real-estate story. [Telehouse Canada urges coordinated planning for Ontario AI data centres](#) points to a MaRS/Mantle Climate workshop and cites an IESO estimate that an additional 5,000MW of capacity may be required by 2035. The recommendations—waste heat recovery, energy storage, demand response, clean PPAs—are basically a checklist for turning data centres from “new peak load” into “manageable grid participant.” Read this as early groundwork for tougher interconnection requirements and more explicit expectations around flexibility and emissions attributes.

The most honest AI-infrastructure commentary today is that compute isn’t the gating factor—power delivery is. In [Power, not compute, is the AI data center bottleneck](#), Mission Critical Group’s Drew Gravitt argues time-to-power, grid

connections, and delivery drive what actually gets built, noting ~30% of planned capacity is being designed as behind-the-meter prime power. That's a big deal because it reframes the competitive set: the "winners" aren't only those with GPUs and land, but those who can engineer around grid delay—whether via on-site generation, storage, or more sophisticated load management.

Ethiopia is getting a clearer in-country cloud story, and it's explicitly about residency and predictable economics. [Wingu Africa launches Wingu Cloud Exchange in Ethiopia](#) introduces WCX as a private cloud platform offering in-country data residency and services including compute, Kubernetes, drive, and security, with Tier III-standard, carrier-neutral positioning and hybrid integration with Azure and AWS. The local-currency pricing detail is easy to skim past, but it's often the difference between "pilot" and "production" for enterprises and public sector buyers in markets where FX volatility can kill cloud adoption. Zooming out, this is the same sovereignty/latency narrative as the edge story—just expressed through national data residency and commercial packaging instead of modular hardware.