TELBORG

Data Centre Briefing

March 31, 2026

Global

Key themes:
Nxtra raises $1B to reach 1GW India; Mistral AI $830M debt for 13,800 Nvidia GPUs near Paris; Virginia push to end $1.9B annual data center sales-tax exemption; AWS waives March charges after drone damage in ME-CENTRAL-1 and ME-SOUTH-1

Airtel is effectively trying to buy time—$1 billion's worth. Its data centre arm Nxtra just raised fresh capital to jump from ~300MW to 1GW "in the next few years" across Pune, Chennai, Mumbai, and Kolkata, at a post-money valuation of about $3.1bn. In a market where power and land are hard but capital is now very available, this is a clear signal: India's next wave of capacity is going to be built at scale, and it's going to be AI-shaped.

The Big Stories

Bharti Airtel's Nxtra raises $1 Bn to scale data centre capacity in a strategic round led by Alpha Wave Global ($435mn) and Carlyle ($240mn), with Airtel itself putting in $290mn (plus Anchorage at $35mn). Nxtra says the money will fund expansion from ~300MW to 1GW and the build-out of AI-ready campuses in four major metros. What matters here isn't just the headline number—it's the implied race to lock in power and sites ahead of the next demand step-change, and the fact that large global capital is underwriting that sprint.

Mistral AI raises $830M debt to build European data centre, earmarked for 13,800 Nvidia GPUs and a facility near Paris (Bruyères-le-Châtel), expected online in Q2 2026. Mistral's broader plan is to secure 200MW of capacity in Europe by end-2027, financed here by a seven-bank consortium including Crédit Agricole, HSBC, BNP Paribas and MUFG. The key signal: Europe's

“sovereign compute” push is moving from policy talk to balance-sheet decisions—and debt (not just venture equity) is becoming a tool to secure scarce GPUs and near-term capacity.

In the US, the subsidy politics are getting louder. Virginia group urges ending $1.9B data center tax breaks, with the Piedmont Environmental Council mobilising residents ahead of the legislature’s April 23 reconvening to end or phase out the state’s annual sales tax exemption on data centre equipment. The PEC frames its argument around Dominion Energy’s cited 70GW of load requests and an estimated >$100bn grid infrastructure need (including $30bn for transmission), plus claimed health damages of $53–$99mn per year from on-site fossil generation. Even if the exemption doesn’t disappear overnight, this is a reminder that “Virginia as default location” comes with rising political risk—especially as grid build costs and local impacts become harder to ignore.

Cloud resiliency just got a very literal stress test. AWS waives March fees after Middle East datacenter drone attacks, covering usage-related charges for March 2026 for affected regions (ME-CENTRAL-1 and ME-SOUTH-1) after March 1 strikes damaged facilities and disrupted services. AWS says it did not delete billing/usage data and that usage data remains available to customers on request; impacted services included core primitives like EC2, S3, DynamoDB, Lambda, and RDS. For customers, the fee waiver is the headline—but the deeper takeaway is that geopolitical and physical security risk is now an operational variable in cloud region choice, not just a procurement footnote.

Europe’s AI colocation arms race is also spreading north. Bitdeer to develop 180 MW AI data center in Norway, via subsidiary Tydal Data Center AS with partner Data Center Installations AS, designed primarily for co-location of Nvidia Vera Rubin technology and built to Nvidia reference design. Completion is targeted for December 2026, and it’s positioned as Norway’s largest AI data centre and one of Europe’s largest by installed capacity. The competitive angle: AI facilities are being marketed less as generic megawatts and more as “GPU-native” environments—reference designs, specific platform alignment, and a clearer promise about what can be hosted.

Behind the Headlines

The “AI data centre on a brownfield” narrative is quickly becoming a template—and it’s already colliding with local reality. Proposal to turn Janesville GM brownfield into AI data center has Viridian Partners pitching an $8bn, 11-building, 800MW campus on a 250-acre former GM site, including funding an estimated $30mn cleanup and projecting ~600 permanent jobs and ~13,000 construction jobs. The story also points to January 2026 EPA guidance identifying 335 candidate brownfields for AI data centres, which effectively formalises a pipeline of “pre-justified” sites. The interesting tension is that brownfields solve one political problem (land use) while intensifying others (energy demand, emissions, and public health)—so these projects may move faster on permitting, but not necessarily faster on social licence.

Energy strategy is now the product, not a supporting slide in the pitch deck. In Energy strategy now defines data center location and growth, atNorth argues that land alone doesn’t decide where data centres get built—power does, and the wait for connections can exceed five years. The piece flags how AI racks have moved from 8.4kW in 2020 to 50–60kW today (with some >100kW), pushing operators toward phased power allocations, renewable procurement, and heat reuse. Read this as a warning to anyone underwriting growth plans off a single interconnect assumption: the gating factor is increasingly a multi-year energy roadmap, not a real estate checklist.

Supply chain is quietly retooling for the cooling-heavy future, with big capex now following the heat. Furukawa to invest JPY55bn in data centre cooling expansion across the Philippines, Thailand, and China, including JPY51bn for a new Thailand water-cooling factory and JPY4bn for air-cooling heat sink upgrades. The timelines matter: air-cooling mass production in China and the Philippines is slated from July 2026, the Philippines expansion targets January 2027, and the Thailand factory January 2028. If you’re tracking where the bottlenecks might shift next, this looks like a pragmatic bet that liquid cooling demand is not a niche—it’s an industrial-scale manufacturing wave that needs years of lead time.