

Data Center Briefing

March 02, 2026

Global

Key themes:

NeevCloud-Agnikul plan for 600 LEO AI satellites; Rackbank-backed ₹10 Cr pilot for orbital inference; Pakistan \$1B AI-by-2030 plan faces water-risk scrutiny; ITIF Tech Policy 202 trains staff on data centre-grid interactions

NeevCloud and Agnikul just sketched out one of the more audacious “data centre” ideas we’ve seen in a while: AI inference infrastructure in low-Earth orbit, starting with a pilot mission and scaling to a constellation of up to 600 satellites over three years. If they can even partially deliver—10–15 kW per satellite, solar powered, radiation-hardened ASICs, and sub-10 ms latency targets—the implication is that the edge is no longer just “near the user.” It’s literally above them.

The Big Stories

NeevCloud and Agnikul propose orbital AI data centres in LEO after signing an MoU in February to deploy AI inference infrastructure at ~300 km altitude. The plan centres on defence and remote applications, pitching sub-10 ms inference latency using solar-powered, radiation-hardened ASIC inference chips, with initial per-satellite power of ~10–15 kW and an operating life of 4–5 years. There’s also a concrete first step: a pilot mission costing up to ₹10 Cr, financed by Rackbank Datacenters.

Why it matters: investors should treat this less like “a weird data centre story” and more like an edge-compute architecture bet. If orbital inference becomes even modestly viable, it creates a new competitive axis—latency and availability in places where terrestrial backhaul is fragile, regulated, or simply

absent. The watch item isn't the headline-grabbing "600 satellites" number; it's whether the pilot can prove reliable power, thermal management, and chip resilience at meaningful utilization.

Behind the Headlines

Pakistan's \$1B AI plan comes with a water warning. The prime minister's goal—a \$1 billion investment in AI by 2030 to build a national AI ecosystem—sounds like the kind of announcement governments everywhere want to make right now. The critique is sharper than the usual sustainability hand-waving: the author argues that scaling AI and data centres risks worsening Pakistan's water crisis, and calls for environmental accounting, disclosures, and mandatory safeguards.

The deeper point is that "AI industrial policy" is quietly turning into "utility policy." It's not enough to announce capital and talent programs when the binding constraint in many markets is water (and, by extension, power and permitting). For data centre developers and financiers, this is a reminder that in water-stressed geographies, project timelines and social licence are going to hinge on measurement and disclosure as much as on land and fiber.

ITIF's Tech Policy 202 seminar series puts data centres and grids on the DC policy agenda. ITIF is running a five-session course for congressional and federal staff in Washington, DC from March 2–30, 2026, with topics ranging from space-based broadband and quantum ecosystems to AI in life sciences and children's online safety. Notably for this audience, one session explicitly targets data centre/grid interactions, and the program is positioned as a practical, credentialed on-ramp for staff (free, certificate for attending at least four sessions).

This looks like a small item, but it's actually a tell: "data centre power" is moving from a niche permitting fight into mainstream tech policy education. When staffers are being trained on data centre/grid interactions alongside quantum and space broadband, it's a sign that infrastructure constraints are becoming a legislative object, not just a local-planning skirmish. For operators, the implication is straightforward: expect more federal curiosity—and eventually more federal opinions—about how quickly large loads can connect, how they're disclosed, and who pays for the upgrades.

TELBO●RG