

# Data Centre Briefing

March 03, 2026

Global

## Key themes:

BlackRock GIP and EQT agree \$33.4bn AES acquisition; IEA flags voltage and protection failures amid data-center load surges; Malaysia halts non-AI data centres, Johor pipeline pegged at 4.6GW; AWS probes ME-CENTRAL-1 UAE outage after fire and objects strike

BlackRock's GIP and EQT agreeing to buy AES for a \$33.4bn enterprise value is a reminder that the "data centre power" trade has moved from project finance to platform finance. If you believe data-centre-driven load growth is durable, owning the utility/renewables developer starts to look less like an ESG bet and more like basic supply security. The interesting wrinkle: this lands on a day when the IEA is effectively warning grids are already struggling with the operational side of electrification, not just the capex side.

## The Big Stories

[GIP and EQT to acquire AES for \\$33.4 billion](#) is the headline deal. The consortium (with CalPERS and QIA as co-underwriters) is paying \$15.00/share for a \$10.7bn equity value, with \$22.7bn of net debt, targeting close in late 2026 or early 2027. The stated motivation is straightforward: AES needs more capital to fund US renewables and utility growth, including new generation meant to serve data centers. The signal here is that "power for AI" is now big enough to pull in mega-fund control capital—because the bottleneck isn't chips, it's electrons and interconnection queues.

Grid fragility also got an unusually blunt airing in [IEA: 2025 Blackouts Expose Voltage, Protection, and Transmission Risks](#). The IEA's Electricity 2026 report (and POWER's write-up) points to voltage instability, reactive-power deficits, and protection coordination failures—problems that get worse as renewable penetration rises and large loads (including data centers) show up in concentrated pockets. The examples are not small: Chile's Feb. 25 blackout affecting ~19m people, Brazil's

substation fire that triggered ~10,000 MW of load-shedding, and a Northern Virginia data-center event pegged at ~1,800 MW. For investors, this is the uncomfortable part of the AI buildout: grid risk is becoming an uptime risk, and regulators will respond with rules that can change project timelines and costs.

In Southeast Asia, policy is actively picking winners. [Malaysia AI approval policy accelerates Johor data centre consolidation](#) says the Prime Minister has halted approvals for non-AI data centres since 2024, pushing the market toward fewer, larger AI-focused facilities—exactly the kind of consolidation buyout funds like. BMI estimates a 4.6GW AI capacity pipeline and expects AI data centres to account for over 70% of electricity output growth through 2026, while “displaced” non-AI demand migrates to Indonesia, Thailand, and the Philippines. The immediate takeaway: Johor isn’t just competing on land and tax breaks anymore; it’s competing through planning permission.

Availability matters because outages still happen—even in marquee cloud regions. [AWS investigates UAE data centre outage after fire](#) after “objects” struck a data center around 4:30 AM PST on March 1, affecting ME-CENTRAL-1 (UAE), initially mec1-az2 and later mec1-az3, degrading EC2, RDS, and nearly 60 other services. Separately, a localized power issue in ME-SOUTH-1 (Bahrain) triggered elevated errors across 50+ services, with AWS routing traffic away from the impacted zone. Two incidents, two different failure modes—physical event and localized power—both landing as customers push more latency-sensitive and AI-adjacent workloads into-region.

Singapore is also experimenting with what “alternative power” for AI campuses could look like. [Bridge, Concord to pilot barge hydrogen for AI data centres](#) outlines an MOU to develop what’s described as the country’s first barge-based hydrogen power generation solution for AI-ready facilities, with NTU and public agencies involved. The work will cover power system architecture, hydrogen supply chain frameworks, and procurement options including renewable PPAs and storage, and is framed as something that could scale across Southeast Asia. This is notable less for near-term megawatts and more for intent: land-constrained markets are starting to treat floating/port-adjacent energy infrastructure as a serious resilience lever.

## Behind the Headlines

The carbon math of “AI boom = electrification” is getting messier. [U.S. coal rebound drives rise in power-sector emissions](#) reports US power-sector CO2 emissions rose as coal generation surged 13% in 2025, contributing to a 4% increase in power plant emissions. The story ties the rebound to seasonal factors, higher natural gas prices,

and policy moves that kept five coal plants open via DOE emergency orders—while rising data-center demand in states like Indiana and Texas is cited as part of what’s sustaining coal-fired generation. The investor read-through: even if hyperscalers sign PPAs, marginal grid supply can still drift dirtier when load grows faster than transmission and firm low-carbon capacity.

India, meanwhile, is doing the unglamorous work that actually unlocks capacity: wires. [India inaugurates ₹3,600 crore transmission projects in Rajasthan](#) covers POWERGRID projects worth over ₹3,600 crore to strengthen renewable evacuation, including a 20 GW REZ transmission system (Phase-III Part-B1) and an 8.1 GW solar evacuation strengthening scheme (Phase-II Part-E). It includes 765/400/220 kV substations and 765 kV double-circuit lines to Sikar-II. Data centres aren’t the headline, but this is the enabling layer for any serious digital infrastructure growth outside a handful of metros: generation targets are cheap to announce; transmission is the hard part.

The “AI in ops” story is also moving from pilots to production in carrier-grade environments. [Telstra and Dell deploy AI-enabled telco cloud resilience](#) describes an AI-enabled telco cloud using Dell’s Telecom Infrastructure Automation Suite, embedding agentic AI for autonomous network operations and self-healing. It runs a multi-vendor architecture supporting containerized and virtualized network functions, and includes Model Context Protocol (MCP) plus a conversational interface for SREs to manage real-time resource reallocation. Put simply: more operators are betting that reliability at scale won’t come from bigger teams—it will come from software that can diagnose and remediate faster than humans can page each other.