Compute diplomacy starts with physical dependencies

Cloud contracts depend on electricity, chips, network access and permissions.

RESEARCH BRIEFING · COMPUTE · 11 NOVEMBER 2026

AI strategy depends on infrastructure that a cloud contract cannot create on demand. Buyers need to examine power, chips, serving regions and permissions before treating an alternative provider as a credible fallback.

Key points

• Separate energy-demand projections from installed or available capacity.

• Check shared dependencies behind nominally different suppliers.

• Test substitution on the actual workload.

Demand is growing; delivery remains uncertain

The IEA's 2026 update estimates data-centre electricity use at 485 terawatt-hours in 2025 and projects about 950 TWh in 2030. It reports 17% growth in overall data-centre demand in 2025 and 50% growth for AI-focused centres. These are different aggregates. [1]

The 2030 figure is a projection, not booked demand or a measurement of GPU capacity. The same report identifies constraints across electricity connections, manufacturing and other parts of the infrastructure chain. It does not establish whether a particular supplier can deliver a contracted workload. [1]

Keep the maps dated

The consulting atlas retains the IEA's 2024 electricity shares: United States 45%, China 25% and Europe 15%. Their sum, 85%, describes historical data-centre electricity consumption. It includes non-AI workloads and supporting equipment. [2]

These shares should not be relabelled as current computing power. The semiconductor exhibit likewise retains its 2022 fabrication baseline and separate denominators. A supplier assessment needs product, foundry, packaging, memory and site evidence.

Ask what the alternative actually changes

For each critical workload, record the model service, hosting region, relevant hardware, network path and regulatory permissions. Mark which facts the supplier has confirmed and which remain assumptions.

Ask whether a second provider shares the same upstream constraints. Compare latency, quality, throughput, recovery time and migration effort under the same demand. A technically available alternative may still fail the organisation’s quality or legal requirements.

Buy a tested option

Run a bounded failover or migration exercise before assigning a continuity benefit to diversification. Report how much work transfers, at what quality and cost, and how long recovery takes.

Prepare a decision record covering capacity commitments, substitution limits and the owner of each unresolved dependency. Track the share of critical workloads with verified maps and tested alternatives. Contractual assurances and operational evidence should remain separate entries.

Sources & research notes

01 / IEA / Key Questions on Energy and AI (2026)

https://www.iea.org/reports/key-questions-on-energy-and-ai/executive-summary
485 TWh is the 2025 data-centre estimate; about 950 TWh in 2030 is a projection covering all data-centre workloads.

02 / IEA / Energy and AI (2025)

https://www.iea.org/reports/energy-and-ai/executive-summary
2024 historical electricity geography; not current GPU availability or AI-only demand.

Independent TACITUS analysis. Research checked 11 November 2026. Published findings retain their dates and limits. Proposed curricula, tests, schedules and worked calculations are labelled analytical examples, not client results.