

Executive Brief

The Vertically Integrated Energy-to-Compute Platform

Brad Robson, CEO & Founder | Dotted i Strategies | April 2026

THE GAP THAT MUST BE BRIDGED

AI compute at scale demands two things in equal measure: silicon and power. Industry has focused heavily on one and underinvested in the other. That is the strategic gap defining the next decade of digital infrastructure.

Hyperscalers are chasing GPU density while grid infrastructure plays catch-up with demand it was never designed to absorb. Power procurement, cooling, transmission, and compute continue to be treated as sequential decisions rather than a unified system. The result is capacity constraints, missed SLAs, and spiralling CapEx across markets that cannot afford the exposure.

THE CASE FOR VERTICAL INTEGRATION

The solution is not faster grid upgrades or larger procurement budgets. It is a structural rethink of how compute infrastructure is built and powered.

Vertically integrating energy generation with purpose-built compute infrastructure, by co-locating nuclear SMRs, gas peaking assets, or large-scale renewables directly with or adjacent to data centres, collapses the distance between electrons and inference. It reduces cost per token by eliminating utility margin, transmission charges, and grid congestion pricing from the cost stack. And it creates platform moats that capital alone cannot replicate once established.

This is not a theoretical model. It is the direction the most forward-thinking infrastructure operators are already moving toward, quietly and deliberately.

THREE STRUCTURAL SHIFTS ACCELERATING THIS

01. Grid constraints have become a strategic weapon

Firms that locked in power purchase agreements and secured grid access three to five years ago now hold a competitive position that capital alone cannot replicate on short timelines. In Alberta, utility CEOs are publicly acknowledging they have been drawing down sustaining capital just to maintain reliability under demand growth that outpaced every forecast. That dynamic is not unique to Alberta. It is playing out across North America, the UK, and Northern Europe simultaneously.

02. Co-location economics are inflecting

SMRs, advanced geothermal, and distributed gas peaking are making on-site or adjacent generation commercially viable at data centre scale for the first time. The economics are not yet universal, but the trajectory is clear. In Canada, the UK, and select US markets, the regulatory and resource conditions are converging faster than most operators have priced in. First movers will set terms that followers will spend years trying to match.

03. Sovereignty is forcing the issue

Defence, government, and regulated industries will not accept compute infrastructure that depends on shared or foreign-controlled grid assets. Sovereign AI compute demands sovereign energy. This is no longer a policy preference. It is becoming a procurement requirement, and the infrastructure that cannot meet it will be disqualified from the most strategically valuable and enduring contracts available.

STRATEGIC IMPERATIVE

The organisations building AI infrastructure over the next five years need to think less like technology companies and more like integrated utilities. The winners will not just own the model. They will own the megawatt.

For investors and operators, the due diligence question has shifted. It is no longer solely about the GPU roadmap. It is about the power strategy: where it comes from, who controls it, how it scales, and whether it is defensible.

If your organisation is planning AI infrastructure investment over the next three to five years, I would ask one question: what is your energy strategy, and can you defend it in front of the investors and board members who will ultimately fund it?