CAN ORBITAL DATA CENTRES OVERCOME TERRESTRIAL CONSTRAINTS?

Data centre construction is being constrained by access to power and land. Among the various attempts to overcome these constraints, orbital data centres are a particularly noteworthy example. In the context of SpaceX's recent IPO, this article explores the pressures that orbital data centres are a reaction to, and the technical and financial hurdles that they face. It also explores where orbital data centres are likely to sit in the data centre ecosystem.

Magnus Jeffery, Consultant

August 2026

Power and land are now the binding constraints on data centre build-out

The AI-driven scale-up of compute capacity is driving a record-breaking investment cycle, with Moody's projecting that hyperscalers plan to spend $700 billion on data centres in 2026 alone (Data Centre Magazine). However, this rollout is being stalled by access to the power and land required by these data centres. Reflecting the increase in AI use, global data centre electricity consumption is predicted to reach 565 TWh in 2026, a 26% increase year-on-year (Gartner); this is putting strain on electricity grids and limits on new grid connections. Furthermore, access to land in suitable locations is increasingly challenging. Permitting roadblocks are increasing in the face of local opposition to data centre construction, presenting a major risk to the aforementioned investment; Data Center Watch estimates that opposition contributed to delays or blocks affecting $152 billion of proposed data-centre investment during 2025.

The industry is already exploring ways around these constraints. The increasing importance of access to power, and controlling power costs, is leading to deployment of efficiency-improving technologies such as power-first site selection, liquid cooling, heat reuse and behind-the-meter generation. Other solutions such as Google and Microsoft's investments in nuclear fusion startups demonstrate that companies are starting to consider increasingly 'moonshot' ideas to overcome power constraints. High water consumption for cooling has led to the development of closed-loop 'zero water' liquid cooling. Land access issues have led to modular and floating data centres being developed. Another radical idea is to decouple data centres from these terrestrial constraints entirely; vocal proponents of orbital data centres, most notably Elon Musk, are betting that the technological and financial constraints of orbital data centres are easier to overcome than the power and permitting constraints of Earth data centres.

Orbital data centres trade terrestrial constraints for technical constraints

There are various theoretical advantages to orbital data centres. Space does not have the same land availability constraints that Earth does, has no local water constraint, and has abundant solar energy. Google estimates that, in a suitable sun-synchronous orbit, a solar panel could be up to eight times more productive than an equivalent terrestrial panel and generate power almost continuously, reducing battery requirements. Ambitions are high, with Blue Origin and SpaceX filing with the FCC for 51,600 and 1 million satellites each, and even startups such as Orbital Compute filing for 100,000. While these figures may be largely designed to stir up attention (and drive SpaceX's valuation to a current $1.63 trillion), it is worth considering the feasibility of the project.

Figure 1: Orbital data centres shift the constraint profile

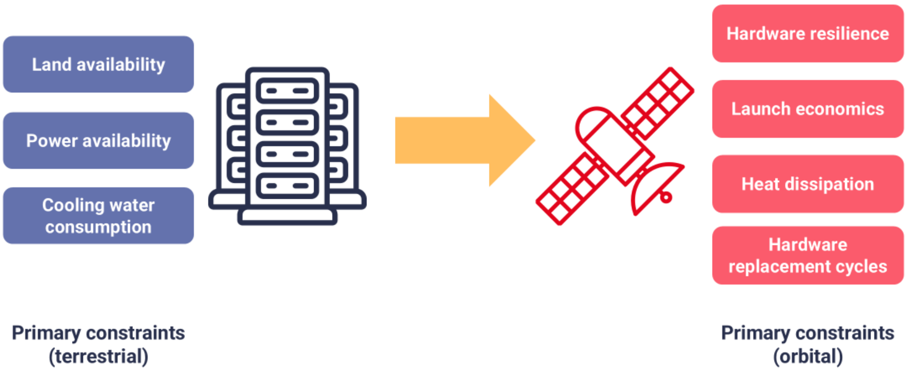

Source: STL Partners

The barriers remain substantial. First, launch economics require an order-of-magnitude improvement. A 2026 analysis estimated that, for a representative orbital system weighing around 40 kg per delivered kilowatt, the allowable combined spacecraft and launch cost would need to fall to $250–1,000 per kilogram. This is 3.4–13.5 times below the public Falcon 9 dedicated-launch benchmark, before communications and operating costs are included. Second, hardware must withstand radiation, thermal cycling and autonomous operation over five-year missions. Google's TPU radiation tests, where TPUs were irradiated with three times the estimated total mission dose, are promising but do not address scaling high-performance, high-bandwidth-memory systems. Third, the vacuum of space eliminates convection, requiring large, heavy radiators; these further complicate launch economics, with an estimated 2,500 m^2 of radiator area required for every MW of IT power (Turyshev). Fourth, communications constrain workloads involving frequent Earth-orbit data transfer, favouring space-native preprocessing, selected inference and orbital clusters. As space, and the RF spectrum, become more congested, the FCC and global governments are likely to increase regulations associated with space data centres, potentially eliminating one of the proposition's key value drivers (decreased political pressure). Finally, given the 2-year cycle of NVIDIA GPU upgrades and the inability to replace chips in space, orbital data centres may run second-rate chips for the majority of their estimated five-year life cycle.

Figure 2: Orbital data centres would act as an orbital edge layer

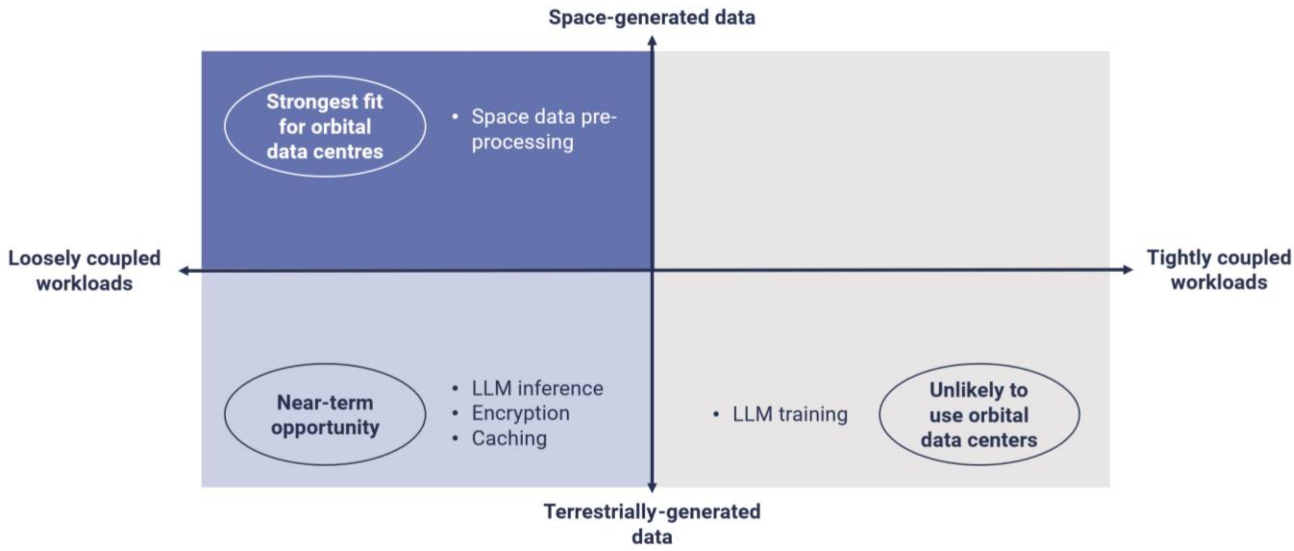

Orbital capacity is an edge layer, not a replacement for AI factories

In our Skywatch tracker, STL's tracker of telco-satellite partnerships, we identified that orbital data centres are unlikely to replace hyperscale training clusters. LLM training is a 'tightly coupled' workflow, meaning that computation steps are chained together in a time-dependent manner; this is unlikely to be run in an orbital data centre, because of the higher latency and chance of network drops. Instead, orbital data centres are better thought of as an edge layer, being used for workloads such as space-generated data preprocessing (e.g. Earth observation data), LLM inferencing, and caching of data for on-Earth processing. It is clear that, for the next several years, the centre of gravity will remain on Earth: the constraints of power access, planning permission, cooling and local acceptance will shape where AI infrastructure is built, rather than prevent it entirely. But the same pressures driving hyperscalers toward nuclear, edge sites and new geographies are also what make orbital data centres worth tracking. The success of orbital data centres partly depends on the extent to which political and power constraints worsen on Earth, incentivising data centre deployments in space.

Magnus Jeffery is a Consultant at STL Partners, specialising in AI topics.

Get in touch with the author to learn more
magnus.jeffery@stlpartners.com

Or visit STL Partners' Skywatch tracker
https://stlpartners.com/research/skywatch-satellite-tracker/