The working fluid problem

Why orbital data centers lose to subsea on the math

Subsea Cloud | July 2026

Orbital data centers and subsea data centers answer the same constraint: Earth's land, water, grid queues and permitting timelines are becoming the binding limits on compute. Both move the facility into an environment nobody is competing for. The comparison ends there. One of them sits in the densest practical heat sink on the planet; the other sits in a vacuum, which is an insulator. Almost every downstream problem in the orbital case, and almost every downstream advantage in the subsea case, follows from that single fact.

Cooling is the whole game

A data center is a machine for turning electricity into heat while doing arithmetic on the way through. Its viability is set by how cheaply it rejects that heat.

In vacuum there is no medium. The only rejection mechanism is radiation, governed by the Stefan-Boltzmann law: $q = \varepsilon\sigma T^4$, with $\sigma = 5.67 \times 10^{-8}$ W/m^2K^4. A radiator at 300 K with emissivity 0.9 rejects about 410 W per square meter, ideal conditions, before accounting for solar loading and view factors. Published estimates for a 40 MW orbital data center land at roughly 120,000 m^2 of effective radiator area, which implies around 330 W/m^2 in practice, and 400–800 tons of radiator hardware. That is the cooling plant alone: 15–20 football fields of panels, launched at $5,000/kg today (about $2–4B for the radiators) or $200–400M at aspirational Starship pricing.

The T^4 term looks like an escape route and is not. Radiating efficiently means running the radiators hot, but chip junction temperatures cap how hot the coolant loop can run. Push radiator temperature above the loop temperature and you need heat pumps, which consume additional power that itself must be radiated, and add mass that must be launched. The geometry does not compress.

Seawater removes the problem rather than solving it. Water is roughly 800 times denser than air and conducts heat about 23 times better; free convection film coefficients in water sit in the range of a few hundred to 1,000 W/m^2K. At a conservative h = 300–500 W/m^2K and a 15–20 K temperature difference against a 4–15°C ocean, a wetted surface rejects 5–10 kW/m^2. The same 40 MW load needs on the order of 4,000–9,000 m^2 of transfer surface: one to one-and-a-half orders of magnitude less area than the orbital radiator field, much of it the module hull itself, none of it launched, pumped or powered. This is why our modules carry no chillers, no CRACs and no working fluid loop, and why cooling contributes zero to our failure model's rotating-equipment risk: there is nothing rotating.

The thermal environment, and what it does to silicon

A body in low Earth orbit circles the planet every ~90 minutes, swinging from roughly +120°C in sun to −170°C in shadow: a 250–300°C cycle, thousands of times per year, accumulating fatigue in every joint, weld and solder ball. Add the radiation environment: high-energy particles flip bits, corrupt calculations and degrade transistors continuously, which is why orbital hardware needs shielding mass and error-tolerant architectures before it does any useful work.

The seabed inverts both terms. Water temperature at depth is stable to within a few degrees year-round, and ten meters of seawater is a better particle shield than anything you can afford to launch. Thermal stability shows up directly in failure rates: each 10°C reduction in junction temperature roughly halves the failure rate of power electronics (Arrhenius acceleration), and the only long-duration controlled experiment we have, Microsoft's Project Natick, retrieved 855 subsea servers after two years with one-eighth the failure rate of its land-based control population. Our reliability model uses that field correction on top of Telcordia SR-332 baselines; the orbital equivalent runs the correction in the other direction.

Availability math with no repair term

Steady-state availability is A = MTBF / (MTBF + MDT). The subsea case gets criticized on the MDT term, fairly: marine intervention takes 21–45 days against hours for a land technician. The model answer is architectural, and it is the same answer our RAM model quantifies at 99.995%: active parallel compute pairs inside each module, k-of-n redundancy across modules, dual diverse cable corridors, and planned interventions at Weibull-derived Year 5 and Year 10 thresholds. Expensive repairs, made rare and scheduled.

In orbit, MDT is not large. It is undefined. There is no repair: failed hardware is abandoned in place and the constellation routes around it, so availability collapses to pure reliability, R(t), decaying monotonically until the next launch window delivers replacement mass, along with its share of solar array and radiator tonnage. Over-provisioning is the only redundancy lever, and every spare is bought at launch prices. Subsea redundancy costs a module; orbital redundancy costs a rocket.

Bandwidth

Optical inter-satellite links are excellent; the atmosphere is the problem. The downlink degrades with cloud, weather and turbulence, and no current free-space optical path approaches fiber capacity. An orbital data center is therefore massive compute behind a narrow, weather-dependent pipe. A subsea module terminates on hard fiber, integrated with its power feed in a hybrid umbilical, on two geographically diverse routes, into the same cross-connect ecosystem every carrier already uses. The ocean floor is where the world's intercontinental bandwidth already lives; we deploy next to it.

The honest column for orbit

Orbital solar delivers roughly ten times the energy per panel of ground solar, with no night at the right altitudes and no interconnection queue. There are no permits in orbit and no neighbors. Launch costs are falling fast, and for workloads that are genuinely latency- and bandwidth-tolerant, sovereign, and worth almost any capex, a niche may exist. None of that touches the core problem: Stefan-Boltzmann is not on the cost curve. Launch prices fall; T^4 does not move.

Summary comparison (40 MW reference case)

	Orbital (LEO, published 40 MW concept)	Subsea (Subsea Cloud architecture)
Heat rejection	Radiation only, ~330–410 W/m²; ~120,000 m² of radiators, 400–800 t	Seawater convection, ~5–10 kW/m²; ~4,000–9,000 m² of wetted surface, mostly hull
Cooling plant	Launched, deployed, pumped; heat pumps to run radiators hot	None. Passive convection; no pumps, chillers or refrigerants
Thermal environment	250–300°C swing every ~90 min; radiation-driven degradation	Stable to within a few degrees; 1/8th land failure rate (Natick, 2-yr field data)
Delivery cost	~1,000 t launched; ≥$5B at $5,000/kg (excl. hardware)	~$1M per MW deployed; 12-week module build
Repair	None. Abandon in place, replace by launch	Retrieve-refresh-redeploy; ROV response ~12 h; planned Year 5/10 campaigns
Availability model	R(t) decay; over-provisioning is the only lever	99.995% via k-of-n redundancy and dual corridors (Monte Carlo, OREDA/Telcordia inputs)
Connectivity	Laser interlinks, weather-limited optical downlink	Hard fiber in dual diverse hybrid umbilicals

Conclusion

Both ideas relocate compute to escape terrestrial constraints. Only one relocates it into a medium that does the hardest job in the building for free. In orbit, cooling is the single largest mass line and the single largest unsolved scaling problem; underwater, cooling is the reason the architecture exists. Fifty years of subsea oil and gas operations supply the failure data, the tooling and the vessels; the orbital case is still waiting on a thermodynamic breakthrough that the Stefan-Boltzmann law does not offer. The seabed is not the conservative alternative to space. It is the version of the same idea where the physics cooperates.

Sources

Orbital figures (40 MW reference design, radiator area and mass, launch pricing, thermal swing, Starcloud/Starlink examples): Anastasi In Tech, "Space Data Centers: The Wildest Idea of the Decade," May 2026, anastasiintech.substack.com/p/space-data-centers-the-wildest-idea. Subsea availability, redundancy and lifecycle figures: Subsea Cloud internal models and data, inputs traced to OREDA, Telcordia SR-332, Quanterion NPRD/EPRD and ICPC. Natick field data: Microsoft Project Natick Phase II retrieval, 2020. Physical constants and convection ranges: standard engineering references; convection coefficients stated as ranges and flagged as assumptions, not measurements.