International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Space-Based versus Land-Based AI Data Centres: A Comparative Environmental Assessment Centred on Google's Orbital Compute Programme

Dr. Vijayalaxmi R. Shinde[1], Rajiv A. Shinde[2]

[1]Department of Environmental Science, M.C.E. Society's Abeda Inamdar Senior College, Azam Campus, Pune, Maharashtra, India
[2]Vishwakarma University, Pune,

Abstract

This paper compares space-based AI compute against terrestrial hyperscale data centres across six environmental parameters, using Google's published research programme on solar-powered TPU constellations in dawn-dusk low-Earth orbit as the reference case (Arcas et al., 2025). The comparison uses a dual functional unit — 1 MW-year of delivered IT power, and 1 MW of IT power over 1-, 5- and 10-year missions — under a cradle-to-grave boundary that includes launch and re-entry for the orbital case. Three findings emerge. First, the carbon result is **functional-unit dependent**: per kWh of energy delivered to the payload, orbital systems are 1.5–1.95× more carbon-intensive than terrestrial ones (52.1–66.3 versus 34.0 gCO_2e/kWh) (Ohs et al., 2025); but per MW of IT capacity against real-world grids, launch-phase emissions of 1,032–2,585 t CO_2e are repaid within the first year and amount to only 5.5–13.7% of five years of terrestrial operation at the 395.65 gCO_2e/kWh global data-centre average (Vries-Gao, 2025). Break-even occurs at grid intensities of 21.6–54.1 gCO_2e/kWh over five years. Second, water and land are where space wins categorically: no evaporative cooling, and no terrestrial land conversion of the tens-to-hundreds of hectares a hyperscale facility requires (Hlabisa, 2025). Third, the binding physical constraint in orbit is not compute silicon but radiative heat rejection — 33–67 cm^2 of radiator per watt of waste heat, or 3,333–6,667 m^2 per MW (Vicente, 2021) — which converts thermal inefficiency directly into launch mass and launch carbon.

Keywords: orbital data centres; data centre sustainability; water usage effectiveness; radiative cooling; life-cycle assessment; re-entry emissions

1. Introduction

The environmental footprint of AI infrastructure has become a first-order energy policy question. Data centres consumed approximately 415 TWh in 2024, with projections reaching ~945 TWh by 2030 — annual growth of 15%, roughly four times faster than heavy industry (Bermejo et al., 2026; Manassero et al., 2026). United States data-centre electricity demand is projected to rise from 176 TWh yr^{-1} in 2023 to 325–580 TWh yr^{-1} by 2028, corresponding to 37–66 GW of average load (Turyshev, 2026).

Against this backdrop, Google published a feasibility study for a space-based AI infrastructure system, publicly known as Project Suncatcher, proposing fleets of solar-powered satellites carrying tensor processing units, networked by free-space optical inter-satellite links and flown in tight formation in dawn-dusk sun-synchronous low-Earth orbit. The stated motivation is explicitly environmental as well as economic: the Sun emits 3.86×10^{26} W, more than 100 trillion times humanity's total electricity production, and a panel in the favoured orbit receives up to 8× more solar energy per year than one at mid-latitude on Earth. Google frames the long-run case as "the most scalable solution, with the additional benefit of minimizing the impact on terrestrial resources such as land and water" .

Precision about status matters. This is a research moonshot and feasibility study, not a deployed data centre. Its published results concern inter-satellite link budgets, orbital formation dynamics, TPU radiation tolerance, and launch-cost trajectories; thermal management, high-bandwidth ground communications, and on-orbit repair are explicitly deferred to future milestones. Any environmental claim about a "Google space data centre" is therefore a claim about a proposed architecture.

The literature contains isolated data points — a claim that space platforms improve PUE by 30.5% while cutting water footprint by 58.1% (Periola et al., 2021); a finding that in-orbit computing carries significantly higher carbon costs even under optimistic assumptions ; engineering assessments of radiator and array geometry (Kumar, 2026) — but no integrated comparison across parameters. That gap is what these paper addresses, and resolving the apparent conflict between those first two findings is one of its central tasks.

2. Reference Architectures and Analytical Framework

2.1 The two systems compared

The terrestrial reference is Google's own fleet, which offers unusual transparency. In 2024 the company consumed 30,825,600 MWh at its data centres at a fleet-average PUE of 1.09, with direct water consumption of 29,477,002 m³ . At that PUE, IT equipment electricity was approximately 28.28 TWh, implying a water usage effectiveness of 1.04 L per kWh of IT electricity. Google's PUE of 1.09 sits at the efficient end of the industry range of ~1.09–1.4, while legacy facilities still commonly exceed 2.0.

The orbital reference is the architecture in Google's feasibility paper: an 81-satellite planar cluster of 1 km radius at a mean altitude of 650 km, with V6e Trillium TPUs radiation-tested to the total ionising dose equivalent of a five-year mission and inter-satellite links targeted at ~10 Tbps aggregate per link. Independent engineering analysis of comparable orbital compute concepts converges on a mass intensity of ~30–50 kg of spacecraft per kW of delivered IT power once fixed spacecraft mass is included, with array area $\approx 4.4 \times 10^3$ m² and radiator area $\approx 1.9 \times 10^3$ m² for a 1 MW node.

2.2 Method

The assessment uses comparative life-cycle framing with a dual functional unit: (i) 1 MW-year of delivered IT power, and (ii) 1 MW of delivered IT power over the mission lifetime. Boundaries are cradle-to-grave, following the four-phase convention of manufacturing, transport, use, and end-of-life processing (Wu et al., 2024). For the orbital case the boundary must additionally include launch and re-entry, which the terrestrial boundary does not. Six parameters are examined: operational energy and carbon; water, direct and embedded; land occupation; embodied materials and circularity; atmospheric and re-entry emissions; and thermal rejection. Where published parameters permit, quantitative estimates are derived and clearly labelled as arithmetic; where they do not, the parameter is treated qualitatively and flagged as a data gap.

2.3 A note on metrics

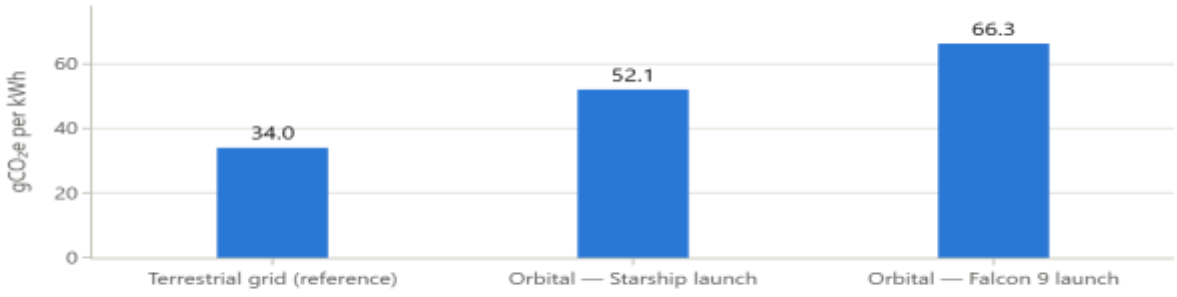
The comparison is complicated by the fact that the industry's own metrics are contested. PUE has been criticised because it considers only energy, can improve merely as IT load rises, and correlates poorly with carbon emissions (Mytton, 2020). WUE is criticised for its dependence on IT energy as denominator, seasonal variability, and failure to account for water *type* and geographic context — a limitation that matters acutely because the same withdrawal in a water-stressed basin is not environmentally equivalent to withdrawal elsewhere (Brand & To, 2025). Less than a third of data-centre operators track any water metric at all. Further, the largest component of a data centre's water footprint is often not on-site cooling but electricity generation, invisible without a source-based metric such as WUE_source (Mytton, 2021). A comparison reporting only direct WUE therefore understates the terrestrial burden and correspondingly overstates the space advantage.

3. Parameter-by-Parameter Comparison

3.1 Operational energy and carbon

The energy case for orbit rests on three physical facts: near-continuous illumination in dawn-dusk sun-synchronous orbit, absence of atmospheric attenuation and weather, and no battery cycling except for eclipse survival. Google estimates a mid-latitude terrestrial panel receives up to 8× less solar energy per year than one in the favoured orbital regime. Independent analysis identifies an optimal regime at a 70–80% sunlight duty cycle, in which high-power computation is limited to sunlit orbit and batteries are sized only for eclipse survival, yielding average delivered compute power of ~70–80 kW per satellite.

Figure 1. Carbon intensity per kWh delivered to the payload



The absence of a carbon-intensive grid in orbit does not, however, make orbital energy carbon-free. A lifecycle model of in-orbit computing found energy intensities of 34.0 gCO_2e/kWh on Earth versus 66.3 gCO_2e/kWh with Falcon 9 launch and 52.1 gCO_2e/kWh with Starship — orbital energy carrying roughly 1.5–1.95× the carbon intensity of terrestrial energy per kWh delivered, because launch and re-entry embodied emissions are amortised over delivered kilowatt-hours. Improved panel efficiency in vacuum is significant, "even outweighing the launch and re-entry cost," but it does not erase it. The rhetorical claim that "energy is carbon-free in space" is therefore not supported at the per-kWh level.

At the fleet level, terrestrial operational carbon is large and geographically variable. Global data centres were associated with an estimated 182 Mt CO_2 for 460 TWh in 2024, implying an average intensity of 395.65 gCO_2/kWh — below the IEA's 445 gCO_2/kWh global generation average, reflecting concentration in the United States (321 gCO_2/kWh) and Europe (174 gCO_2/kWh) . Applied to Google's United States

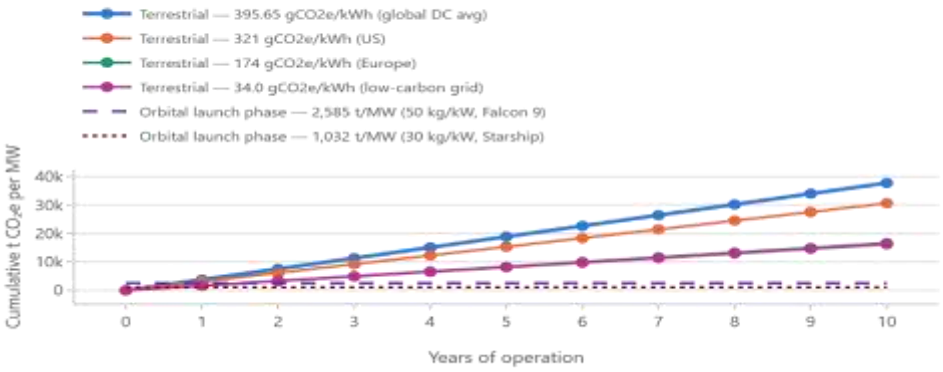
facilities, de Vries-Gao's site-level reconstruction estimated ~7.2 Mt CO_2 with ~57.5 billion L of embedded water, closely matching Google's reported North America scope-2 figure of 6.9 Mt CO_2.

Table 1. Terrestrial operational carbon for 1 MW of delivered IT power (PUE 1.09; author's arithmetic using the intensities reported in)

Grid carbon intensity	1 year	5 years	10 years
Global data-centre average, 395.65 gCO_2e/kWh	3,778 t	18,889 t	37,778 t
IEA global generation average, 445 gCO_2e/kWh	4,249 t	21,245 t	42,490 t
United States, 321 gCO_2e/kWh	3,065 t	15,325 t	30,650 t
Europe, 174 gCO_2e/kWh	1,661 t	8,307 t	16,614 t
Near-zero-carbon grid, 34.0 gCO_2e/kWh	1,623 t	8,117 t	16,234 t

Figure 2 plots those accumulation paths against the orbital launch pulse, so the crossover is visible rather than tabulated. The terrestrial lines rise with time; the orbital lines are flat, because the launch carbon is paid once.

Figure 2. Cumulative carbon: terrestrial operation vs one-off orbital launch emissions, per MW of IT power



3.2 Water consumption — the categorical asymmetry

Water is where the architectures diverge most sharply and where the space case is strongest.

Terrestrial data centres consume water both directly, for evaporative cooling, and indirectly through electricity generation. Direct consumption for a 15 MW facility is estimated at 80–130 million gallons annually, and a medium-sized data centre uses more water each year than three hospitals or two 18-hole golf courses (Gaston et al., 2023). The IEA reported about 140 billion L of direct water consumption against 360 TWh of data-centre electricity and a global average PUE of 1.43 in 2023, implying a WUE of 0.56. For Google specifically, 29,477,002 m³ against 28.28 TWh of IT electricity implies 1.04 L/kWh. Embedded water is larger still: the reconstruction of Google's United States facilities implies 2.97 L per kWh consumed.

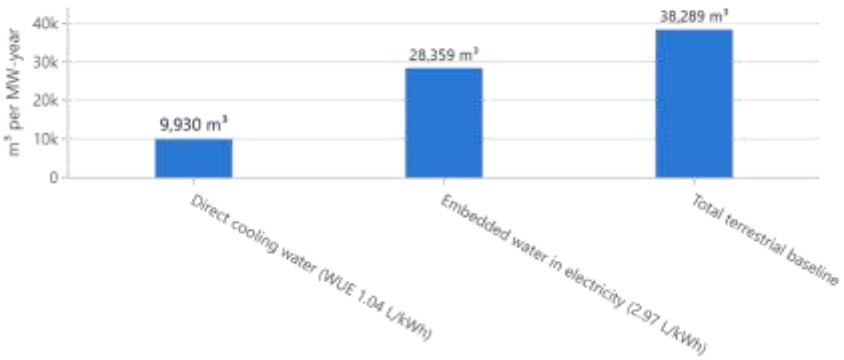
Table 2. Water consumption per 1 MW of IT load per year (Google fleet PUE 1.09; volumes are the author's arithmetic from the 1.04 L/kWh direct intensity and the 2.97 L/kWh embedded intensity)

Component	Volume per MW-year
Direct cooling water at WUE 1.04 L/kWh of IT electricity	9,930 m³
Embedded water in electricity at 2.97 L/kWh	28,359 m³

Total	**38,289 m³**

Figure 3 shows the split, which matters for how much of the water saving orbit can actually claim: roughly a quarter is direct cooling water that radiation genuinely eliminates, and three quarters is embedded water that a renewable-matched terrestrial facility could also avoid.

Figure 3. Terrestrial water consumption per MW of IT load per year



The orbital alternative consumes, in principle, zero water for cooling: heat is rejected by radiation rather than evaporation, and power comes from photovoltaics rather than thermoelectric plants. The benefit is categorical, not incremental. This underpins Periola et al.'s finding that a hybrid terrestrial–space architecture reduces water footprint by 58.1% and improves PUE by 30.5% on average , and their related result that a water-potential metric falls by 4.9–93% depending on configuration (Periola et al., 2022).

Two qualifications are needed. First, the *embedded* water figure is contingent on the terrestrial power source: because wind and solar generation have low water footprints, a grid-matched renewable terrestrial facility would already capture much of the embedded-water saving without leaving the ground. The non-substitutable saving is the ~9,930 m³ per MW-year of direct cooling water — significant, but an order of magnitude smaller than the embedded figure. Second, orbital systems displace water *stress* rather than water *dependency*: the upstream footprint of semiconductor fabrication, propellant production, and launch operations remains, and is unaccounted for in every study reviewed here.

3.3 Land occupation

Terrestrial AI infrastructure is spatially extensive. Data centres and fabrication facilities typically require tens to hundreds of hectares of serviced land with cooling systems, substations, and access infrastructure, and once established remain operational for decades, permanently converting agricultural, peri-urban, or semi-natural landscapes. In Northern Virginia, existing data centres occupy about 6,200 acres while proposed projects would add nearly 21,000 acres — more than tripling the regional footprint where roughly 80% of the state's facilities sit; the 152 proposed facilities average five times the size of the 191 existing ones (Gour et al., 2026). Land-use conflict has driven community opposition in Ireland and an early expansion moratorium in Singapore . Occupation also extends upstream: renewable capacity procured for the sector can claim a substantial share of regional territory, as analysis of Aragon's 47,719 km² illustrates (Torrubia et al., 2026).

Orbital compute removes this demand from the terrestrial surface entirely — a genuine advantage for water-stressed and land-constrained regions. But it substitutes a different spatial claim. The architecture depends on the dawn-dusk sun-synchronous low-Earth orbital band, which Google describes as capable

of hosting terawatts of compute capacity. That band is a finite, unowned common-pool resource, and within it the number of satellites in a cluster scales quadratically with cluster radius for a given minimum inter-satellite distance, so the orbital footprint is itself a governed quantity. Land displacement is real; land abolition is not.

3.4 Embodied materials and circularity

Both architectures share the same dominant embodied burden: the chips. For a Dell PowerEdge R740 rack server, 49.8% of lifecycle emissions come from hardware manufacturing, 80% of which is integrated-circuit manufacturing . Embodied carbon modelling for servers now typically follows the Boavizta framework combined with the Architectural Carbon modelling Tool, which combines process electricity, material embodied carbon, and die size to estimate chip-level emissions (Vanderbauwhede, 2025; Wadenstein & Vanderbauwhede, 2025).

The orbital case adds to this baseline and then subtracts from it. It adds the satellite bus, solar array, radiator mass (~4 kg/m² for a 1.5 mm aluminium panel), radiation shielding, and the launch vehicle's structural share. It also weakens the circularity loop: terrestrial servers are refurbished, resold, and materially recycled, whereas orbital hardware is designed to ablate completely on re-entry, converting engineered materials into atmospheric aerosol rather than recovered feedstock (Barker et al., 2024; Ferreira et al., 2024). Google acknowledges the operational analogue: failed TPUs are manually replaced on Earth, which is "relatively simple and low-cost," but "obviously impracticable in space," with redundant provisioning the simplest mitigation. Shortened service life and redundancy provisioning both raise embodied burden per unit of delivered compute — and embodied carbon is already the category where the literature cautions that accounting can overstate real reductions and distort operational incentives (Bashir et al., 2024).

3.5 Atmospheric and re-entry emissions

This parameter has no terrestrial counterpart of comparable scope, and is the least-regulated dimension of the comparison.

Rockets emit gases and particles directly into the middle and upper atmosphere where removal mechanisms are weak. Instantaneous radiative forcing from major rocket engine emissions — CO_2, H_2O, black carbon, and alumina — has been estimated at 16 ± 8 mW m^{-2}, with relative contributions of 70% from black carbon, 28% from alumina, and 2% from water vapour; alumina particles previously thought to cool the Earth by scattering solar flux were found to absorb outgoing longwave radiation, producing net positive forcing (Ross & Sheaffer, 2014). Rocket soot is estimated to be up to 500 times more effective at contributing to global warming than equivalent aviation soot, because it persists in the stratosphere for years (Olawade et al., 2025). Lifecycle assessment of reusable launch fleets concludes that climate impact is dominated by exhaust emissions rather than manufacturing or propellant production (Calabuig et al., 2024).

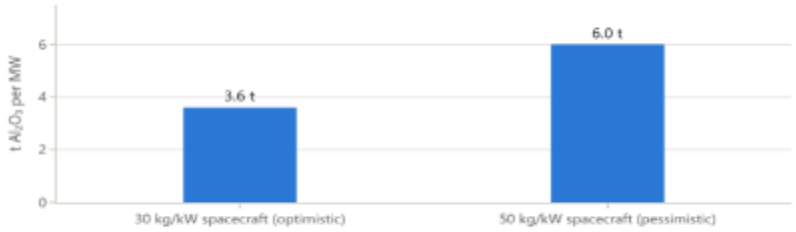
Re-entry adds a second, growing stream. A typical 250 kg satellite generates around 30 kg of aluminium oxide nanoparticles on demise, which may endure for decades; the entire satellite population re-entering in 2022 produced roughly 17 metric tons, while mega constellation scenarios point to over 360 metric tons per year. Aluminium oxides from satellite demise are known catalysts for chlorine activation that depletes stratospheric ozone. Under a scenario in which all mega constellations filed with the FCC as of 2022 operate, re-entry-ablated alumina reaches 2.52 Gg/yr, with an atmospheric lifetime of one to two years and a cooling radiative forcing of approximately -0.378 mW/m^{-2} (Jain et al., 2023). For satellite re-entry specifically, current near-Earth drag rates imply roughly 0.2 Gg/yr, with scaled growth scenarios reaching 0.8–2.5 Gg/yr and up to 5 Gg/yr including re-entering boosters (Revell et al., 2025). Projected burdens of

20–40 Gg stratospheric alumina imply 1.5 K temperature anomalies in the mesosphere and stratosphere at Southern Hemisphere high latitudes and a 10% reduction in polar vortex wind speed (Maloney et al., 2025).

Scaling this to an AI constellation is necessarily indicative, since no orbital compute system of relevant scale exists. Applying the ~30 kg Al_2O_3 per 250 kg of satellite mass ratio proportionally to the 30–50 t of spacecraft required per MW of delivered IT power yields roughly 3.6–6.0 t of alumina per MW of orbital compute at end of life. That is an order-of-magnitude indicator extrapolated across four orders of magnitude of structural complexity, not a modelled result.

Figure 4 visualises that indicative range against the mass intensity it scales with; the linearity is the point — re-entry alumina grows in direct proportion to the launch mass per megawatt.

Figure 4. Indicative end-of-life re-entry alumina per MW of orbital compute



The terrestrial comparison differs in kind rather than degree. Data-centre air-quality impacts are local and immediate: modelling suggests that by the late 2020s, emissions from United States data centres alone could be responsible for 1,300 premature deaths and 600,000 asthma cases annually, alongside noise and land-conversion pathways affecting nearby communities (Neto & Kazak, 2026). The trade is therefore not "impact versus no impact" but **local, regulated, measurable, and reversible** against **global, unregulated, persistent, and effectively irreversible on policy timescales**.

3.6 Thermal rejection — the binding physical constraint

This parameter most fundamentally differentiates the architectures and is the one most often omitted from popular comparisons.

On Earth, waste heat is rejected to the atmosphere by convection and, where water is available, evaporation — an effectively free, high-capacity sink requiring no apparatus beyond a cooling tower or dry cooler. Data centres dissipate more than 90% of their energy input as waste heat . In space, the only viable heat rejection pathway is radiative cooling governed by the Stefan–Boltzmann law $Q = \varepsilon\sigma T^4$, and "reducing heat generation is thus an indispensable avenue to reducing the operating temperature," because radiative rejection is fundamentally constrained (Abraham et al., 2026).

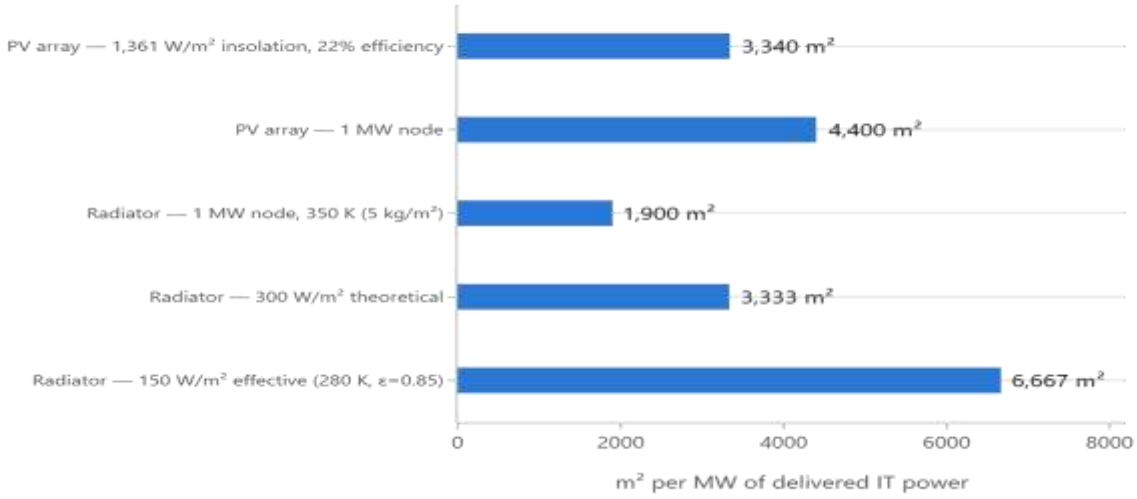
The engineering consequences are severe. A radiator facing outer space with infrared emissivity $\varepsilon = 0.85$ at 280 K dissipates approximately 300 W/m² in theory, halved in practice to 150 W/m² to account for equipment-to-radiator temperature gradients and illumination from the Sun or solar array . Heat rejection scales with the fourth power of radiator temperature and with radiating area, so design effort concentrates on elevating operating temperature and deploying larger surfaces, and the radiator must be positioned so it is not exposed to solar radiation, a major heat source (Selvadurai et al., 2022). Passive radiators suit heat rejection up to several hundred watts; deployable radiators are required beyond that, up to several kilowatts, bringing complexity, mass, and heater-power penalties (Tachikawa et al., 2022).

Table 3. Thermal and power area requirements per MW of delivered IT power

Parameter	Value
Radiator area at 150 W/m² effective (280 K, ε = 0.85)	6,667 m²
Radiator area at 300 W/m² theoretical	3,333 m²
Radiator area, 1 MW node at 350 K, 5 kg/m²	~1,900 m²
PV array, 1 MW node	~4,400 m²
PV array at 1.361 kW/m² insolation, 22% efficiency	~3,340 m²

Figure 5 makes the asymmetry Concrete: an orbital megawatt carries several thousand square metres of radiator and array area *that must be launched*, while the terrestrial equivalent's heat sink is the open sky.

Figure 5. Radiator and photovoltaic area required per MW of orbital IT power



Expressed per watt of waste heat, the requirement is 33–67 cm² of radiator surface. Independent assessment of a 100-kW orbital compute satellite with ~1 t spacecraft mass reached the same structural conclusion: solar array areas of ≈400–600 m² and radiator areas of ≈100–250 m² are practical, but "heat rejection — not compute silicon — emerging as the dominant constraint," with a hard feasibility boundary near ~1,000 kg total spacecraft mass. Google's paper likewise defers thermal management, describing only a system of heat pipes and radiators at nominal operating temperatures.

The asymmetry is therefore stark. A terrestrial 1 MW facility requires essentially no dedicated heat-rejection *area* — the atmosphere does the work. An orbital 1 MW facility requires 1,900–6,667 m² of radiator whose mass must be launched, whose surfaces must be shielded from solar illumination and from the infrared radiation of neighbouring satellites within the cluster , and whose materials must withstand strong ultraviolet radiation, cosmic rays, and atomic oxygen exposure over a mission life targeted at five years (Xu et al., 2025).

4. Statistical Analysis

4.1 Break-even carbon analysis

The central quantitative question is: under what conditions does orbital compute emit less carbon than the

terrestrial facility it displaces? Launch emissions are a one-off; terrestrial operational emissions accumulate annually. The break-even grid intensity — the intensity at which the two are equal over a mission — is the discriminating statistic. Above that value the orbital system is the lower-carbon option; below it, the terrestrial facility is.

Table 4. Break-even grid carbon intensity (gCO_2e/kWh) (author's arithmetic; launch intensities from, mass intensity from, PUE 1.09 from)

Spacecraft mass intensity	Launch vehicle	1-year mission	5-year mission	10-year mission
30 kg/kW	Falcon 9 (51.7 $kgCO_2e/kg$)	162.4	32.5	16.2
30 kg/kW	Starship (34.4 $kgCO_2e/kg$)	108.1	21.6	10.8
50 kg/kW	Falcon 9	270.7	54.1	27.1
50 kg/kW	Starship	180.1	36.0	18.0

Figure 6 presents the same result as a crossing problem. The rising line is five years of terrestrial operation; the two horizontal bands are the optimistic and pessimistic launch carbon. Wherever the line sits above the band, orbit wins on carbon — which, across the plausible grid range, is almost everywhere.

Figure 6. Five-year carbon crossover: terrestrial accumulation against the orbital launch band

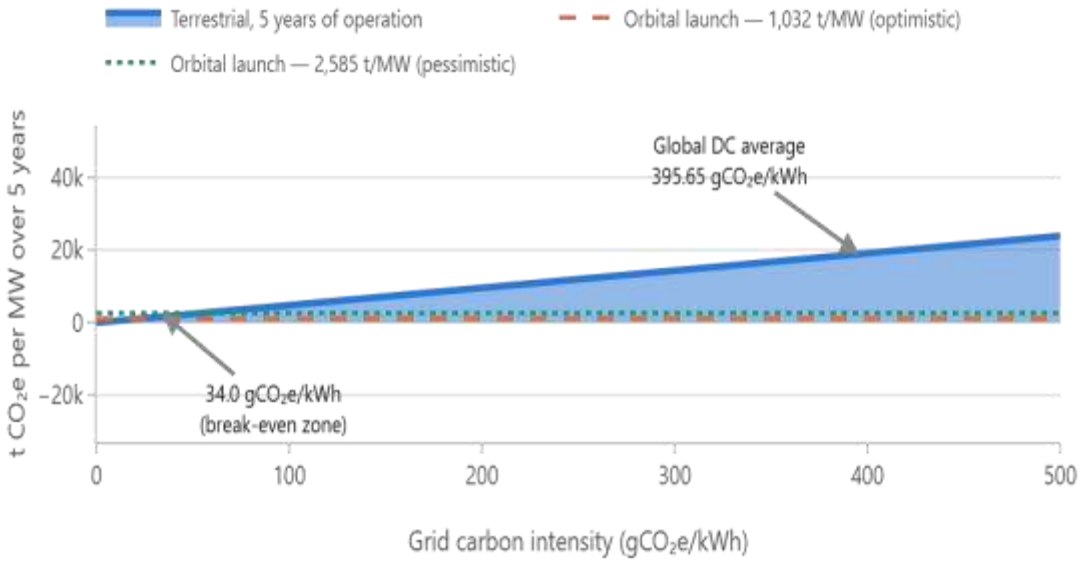
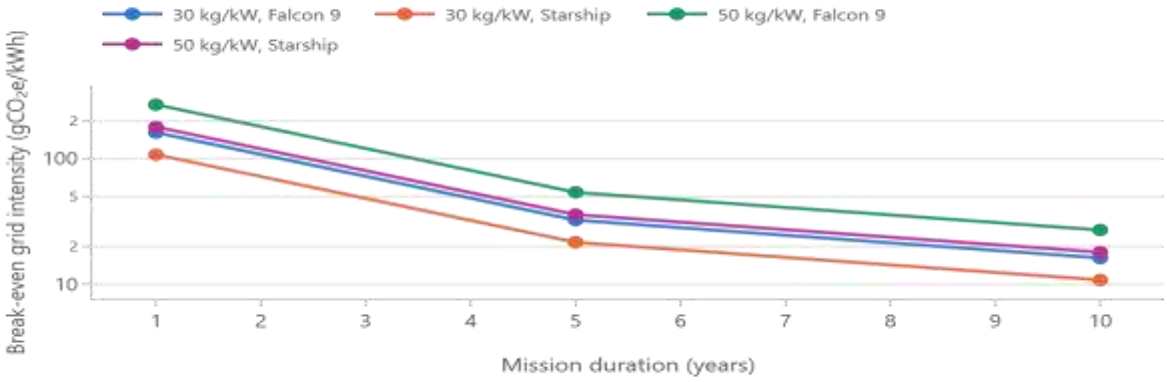


Figure 7 shows how the break-even threshold falls with mission duration — the longer the spacecraft operates, the more thinly the launch carbon is amortised, and the lower the grid intensity at which orbit remains the lower-carbon choice.

Figure 7. Break-even grid intensity by mission duration, mass intensity and launch vehicle



Launch-phase embodied carbon is 1,032 t CO_2e per MW at the optimistic corner (30 kg/kW, Starship) and 2,585 t CO_2e per MW at the pessimistic corner (50 kg/kW, Falcon 9). Against the 395.65 gCO_2e/kWh global data-centre average, that one-off equals 5.5–13.7% of five years of terrestrial operation: over that horizon the terrestrial facility emits 7.3× to 18.3× more than the launch. Against the United States average of 321 gCO_2e/kWh the ratio is 5.9× to 14.8×; against the European average of 174 gCO_2e/kWh, 3.2× to 8.0×.

The break-even values themselves establish where these reverses. Every one of the table's five-year values (21.6–54.1 gCO_2e/kWh) is far below the 395.65 gCO_2e/kWh average, so orbital compute is the lower-carbon option against any realistic operating grid — but it ceases to be so on grids already approaching the break-even band. That is exactly the regime in which Ohs et al.'s benchmark sits: their terrestrial comparison points of 34.0 gCO_2e/kWh yields 1,623 t over five years at 1 MW, against 1,551 t of launch carbon at 30 kg/kW with Falcon 9 — a difference of about 5%. The two literatures are therefore consistent, and the apparent conflict dissolves once the functional unit is made explicit.

4.2 Why the functional unit flips the sign

Three things must be said about that reconciliation, because it is the most consequential methodological point in this paper.

First, **per kWh delivered to the payload**, orbital energy is genuinely worse: 52.1–66.3 gCO_2e/kWh against 34.0 on Earth, or 1.5–1.95×. Any claim that space makes energy carbon-free fails at this level.

Second, **per MW of installed IT capacity over a multi-year mission against a real grid**, the launch carbon debt is small relative to avoided terrestrial emissions — 5.5–13.7% of five years at the global average. This framing favours orbit, because it credits the orbital system with full-time capacity utilisation.

Third, **per unit of compute work actually delivered**, the advantage narrows again, and Ohs et al.'s conclusion that in-orbit systems incur significantly higher carbon costs, driven by embodied emissions from launch and re-entry, becomes reachable. Duty cycling at 70–80% sunlight, redundant provisioning for unrepairable hardware, and a five-year replacement cadence all reduce delivered compute per kilogram launched, and each of these appears in the per-work framing but not in the per-capacity framing.

The honest summary is that the carbon verdict is **contingent, not structural** — and that it depends more on the displaced grid and the assumed utilisation than on any property of orbit itself.

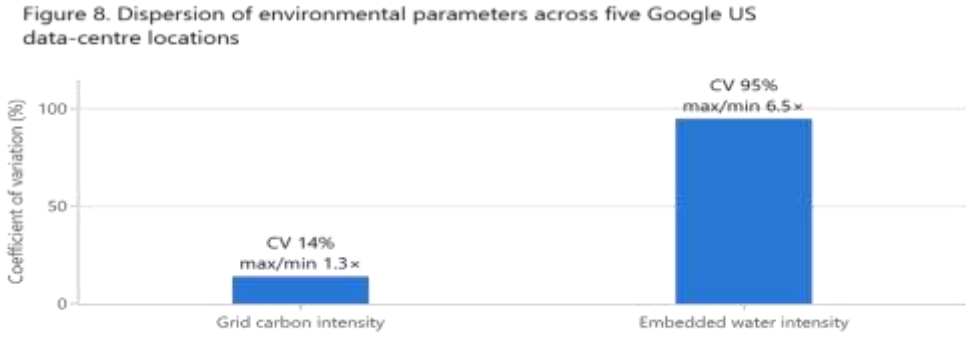
4.3 Dispersion across terrestrial sites

A second statistical observation strengthens the space case on water specifically. Reconstructing Google's United States data-centre locations shows markedly different dispersion across parameters.

Table 5. Dispersion of environmental parameters across Google United States data-centre locations (n = 5 sites; values from)

Parameter	Mean	SD	CV	Range	Max/min ratio
Grid carbon intensity (tCO_2/MWh)	0.398	0.057	14%	0.35–0.46	1.3×
Embedded water intensity (L/kWh)	4.240	4.015	95%	1.74–11.30	6.5×

Figure 8 contrasts the two coefficients of variation directly: siting decisions move carbon intensity very little but move embedded water intensity enormously.



Grid carbon intensity varies by only 1.3× across these sites, but embedded water intensity varies by 6.5×, with a coefficient of variation of 95%: 11.3 L/kWh at The Dalles, Oregon, against 1.74 L/kWh at Ashburn, Virginia. Two implications follow. Siting can substantially mitigate terrestrial water impact — but only if the embedded dimension is measured, which most operators do not do. More importantly for this comparison, the *terrestrial* water footprint is highly location-dependent and therefore physically and politically contingent, whereas the *orbital* water advantage is invariant. For a firm facing water-permitting risk in a specific basin, that invariance has real value a fleet average conceals.

4.4 Reconciling the contradictory literature

The apparent conflict between Periola et al. (58.1% water-footprint reduction, 30.5% PUE improvement) and Ohs et al. (significantly higher carbon costs in orbit) is not a contradiction at all. Periola's result is a water-and-PUE result with a hypothetical substitution of high-PUE terrestrial facilities by low-PUE space facilities; it makes no carbon claim and does not include launch emissions in its boundary. Ohs's result is a carbon-per-unit-of-compute result that explicitly includes launch and re-entry embodied emissions and finds these dominate. Both can be simultaneously correct. The error would be to generalise either to "environmental impact" as a whole — precisely the aggregation problem that multidimensional KPI frameworks were developed to address, distinguishing energy efficiency, cooling, greenness, thermal management, and other dimensions rather than collapsing them into a single score (Silva et al., 2023).

5. Discussion

5.1 What the comparison actually establishes

Read together, the six parameters resolve into a conclusion that is more conditional than either promotional or dismissive accounts suggest.

On **carbon**, the advantage is real but contingent. Launch-phase emissions of 1,032–2,585 t CO_2e per MW are one-off and are outweighed within a single year by terrestrial operation at the global average data-centre intensity of 395.65 gCO_2e/kWh. But the advantage evaporates on grids approaching the 21.6–54.1 gCO_2e/kWh five-year break-even band, and it narrows under per-work accounting because of duty cycling, redundancy, and replacement cadence. Space is not a general decarbonisation strategy; it is a low-carbon option for high-carbon grids and a roughly neutral one for clean grids.

On **water and land**, the advantage is structural rather than contingent. No amount of terrestrial engineering eliminates evaporative cooling demand in the way that radiative rejection does, and no terrestrial siting decision avoids the tens-to-hundreds of hectares of permanent land conversion a hyperscale facility implies. These are the parameters on which the project's environmental case genuinely rests.

On **the atmosphere**, the balance is unfavourable and ungoverned. Terrestrial impacts are local, measurable at the fence line, and subject to permitting; orbital re-entry and launch emissions are global, dispersed, unmonitored at the point of emission, and currently unallocated — and roughly 10% of stratospheric aerosol particles already contain elements unique to spent satellites and discarded rocket bodies.

5.2 The thermodynamic asymmetry that dominates

The single most important technical insight is that the two architectures are not symmetric in their constraints. On Earth the binding constraints are *resource and siting* constraints: grid interconnection, water permitting, land availability, community acceptance. In space the binding constraint is a *physical law* — heat rejection by radiation alone, at $Q = \varepsilon\sigma T^4$.

This has a counterintuitive consequence. The terrestrial industry's waste-heat problem — more than 90% of energy dissipated as heat, with reuse adoption at only 3–5% in the UK against a 10% Energy Reuse Factor mandate in Germany — is a problem of *underutilisation*. In orbit, the same waste heat becomes a *mass and area penalty* that propagates directly into launch mass, launch cost, and launch carbon. Every watt of inefficiency must be radiated away by 33–67 cm² of radiator surface at 150–300 W/m², and every kilogram of that radiator must be launched at 34.4–51.7 kgCO_2e per kilogram. In orbit the waste-heat problem and the carbon problem are the same problem. This is why heat rejection rather than compute silicon emerges as the dominant constraint in independent feasibility analysis, and why Google defers thermal management to future milestones.

5.3 Governance and the atmospheric commons

Alumina from satellite demise is a known catalyst for chlorine activation that depletes stratospheric ozone, and re-entry-ablated alumina has an atmospheric lifetime of one to two years. If orbital AI compute scales toward the terawatt capacity Google suggests the dawn-dusk band could host, the aggregate alumina and black carbon burden scales with launch and re-entry mass — that is, with the 30–50 t per MW mass intensity. No current framework allocates responsibility for that burden. The terrestrial alternative, whatever its failings, at least has an accounting protocol: the Greenhouse Gas Protocol, however imperfectly applied to cloud workloads, and emerging multidimensional KPI regimes that go beyond PUE to include CUE, ERE, and WUE.

5.4 Where the numbers are weakest

Three gaps materially limit confidence.

First, **no orbital compute system exists at relevant scale**, so all orbital figures derive from engineering models and analogous systems; Google explicitly treats thermal management, ground communications, and on-orbit repair as future work. The alumina scaling in Section 3.5 extrapolates a 250 kg-satellite demise ratio to a 30–50 t spacecraft and should be treated as indicative only.

Second, **embodied carbon accounting for semiconductors remains uncertain**. Methodologies have advanced through Boavizta and the Architectural Carbon modelling Tool, but over-emphasis on embodied carbon can overstate real reductions and distort incentives for operational decarbonisation. Since embodied carbon is where orbital systems are weakest, this uncertainty cuts asymmetrically against the space case.

Third, **water accounting is systematically incomplete**. WUE's denominator dependence, seasonal variability, and blindness to water type and geography make even the terrestrial baseline in Table 2 uncertain, and regulatory reliance on WUE risks underestimating true water impacts. Because the largest water footprint component is often electricity generation rather than on-site cooling, the terrestrial figure is highly sensitive to the generation mix assumed.

6. Conclusion

This paper compared space-based AI compute — represented by Google's solar-powered TPU constellation architecture in dawn-dusk low-Earth orbit — against terrestrial hyperscale data centres across six environmental parameters, using a dual functional unit and a cradle-to-grave boundary.

First, the carbon verdict is contingent on the displaced grid and the functional unit. Per kWh delivered to the payload, orbital energy is 1.5–1.95× more carbon-intensive than terrestrial energy (52.1–66.3 versus 34.0 gCO_2e/kWh). Per MW of IT capacity against real grids, launch-phase emissions of 1,032–2,585 t CO_2e amount to 5.5–13.7% of five years of terrestrial operation at the 395.65 gCO_2e/kWh global average, and break-even occurs only at 21.6–54.1 gCO_2e/kWh over five years. Per unit of compute work actually delivered, the advantage narrows toward Ohs et al.'s more pessimistic conclusion, because duty cycling, redundancy, and replacement cadence reduce delivered compute per kilogram launched. Space is a low-carbon option for high-carbon grids, not a general decarbonisation pathway.

Second, water and land are where the space case is structural, not contingent. Eliminating evaporative cooling avoids ~9,930 m^3 of direct water per MW-year, and eliminating thermoelectric generation avoids ~28,359 m^3 of embedded water against a terrestrial baseline of 38,289 m^3, while removing tens to hundreds of hectares per facility from terrestrial conversion. The terrestrial alternative cannot improve on these beyond its physical limits.

Third, the binding constraint in orbit is thermodynamic, not computational. Radiative rejection at $\varepsilon\sigma T^4$ requires 33–67 cm^2 of radiator per watt of waste heat, or 1,900–6,667 m^2 per MW, because space provides no convective sink. This converts waste heat directly into launch mass, launch cost, and launch carbon, making thermal efficiency the highest-leverage design variable for the environmental performance of orbital compute.

Fourth, the atmospheric dimension is a governance gap, not merely an emissions line item. Re-entry alumina and launch black carbon deposit directly into the stratosphere with one- to two-year residence times, catalytic ozone effects, and no allocation framework. Any scaling of orbital compute should carry lifecycle-based atmospheric accounting, since the terrestrial alternative — despite local air-quality impacts

that modelling suggests could cause 1,300 premature deaths and 600,000 asthma cases annually in the United States alone — is at least subject to permitting and measurement.

The practical implication is that orbital AI infrastructure should be evaluated as a **water-and-land intervention whose carbon position depends on the grid it displaces**, targeted at water-stressed, land-constrained, high-carbon-grid jurisdictions, and not marketed as a general sustainability solution. Framing it otherwise would repeat, at planetary scale, the error that PUE-only accounting already produced in the terrestrial industry: optimising the metric that is easy to measure while misallocating the resource that actually matters.

References:

1. Coyle, W. L., & Gee, K. L. (2024). An application of aircraft noise metrics to the Artemis-I launch. *NOISE-CON Proceedings*, *270*(5), 6304–6311. https://doi.org/10.3397/in_2024_3712
2. Gee, K. L., Mathews, L. T., McLaughlin, B., Anderson, M. C., Edgington-Mitchell, D., & Hart, G. W. (2023). Rocket noise: What does it mean for Australian spaceports? *The Journal of the Acoustical Society of America*, *154*. https://doi.org/10.1121/10.0023749
3. Gee, K. L., Pulsipher, N. L., Kellison, M. S., Hart, G. W., Mathews, L. T., & Anderson, M. C. (2025). Starship Super Heavy acoustics: Comparing launch noise from Flights 5 and 6. *JASA Express Letters*, *5*(2). https://doi.org/10.1121/10.0035925