Near-Continuous Solar Orbital Compute: A Falsifiable Project Thesis for 2029–2030

Research memorandum, 25 September 2026. Monetary values are real 2026 US dollars. Physical calculations are first-order engineering estimates. Numerical capital costs, accelerator efficiency, traffic, and financing terms are explicit scenario inputs, not observed vendor prices or signed commercial terms.

Abstract

The 2029–2030 opportunity in orbital AI is best stated as a *conditional* infrastructure thesis. Solar access in a carefully selected dawn–dusk low-Earth orbit (LEO) can reduce dependence on terrestrial electricity, but free irradiance does not imply cheap useful computation: photovoltaic arrays, radiators, shielding, deployment structures, communications, reliability and launch remain capital-intensive. The IEA's 2026 outlook projects approximately 950 TWh of data-centre electricity consumption in 2030, yet this global aggregate does not establish that every grid is constrained or that orbital compute wins a commodity-compute auction. The investable hypothesis tested here is narrower. A modular 250-kW IT satellite near 650 km, in a mission-verified high-beta sun-synchronous orbit, could create a first market by processing *data already born in orbit*, transferring compact results instead of large raw sensor streams and reducing time to insight. The long-run option is solar-powered orbital AI capacity for asynchronous terrestrial jobs, conditional on demonstrably low launch cost, high useful-FLOP efficiency, durable deployable hardware and adequate optical-network throughput. A ten-figure, reproducible model derives illumination, array/battery/radiator mass, normalized cost, break-even customer value, and fleet launch cadence. Its optimistic nominal case costs about USD 0.715 per $10^18 useful FLOP in orbit, versus USD 0.630 on a terrestrial grid at USD 80/MWh. A 10-MW fleet carrying USD 100 million of non-recurring engineering needs about USD 90 of *incremental realized customer value* per input TB at 100 TB/day/MW under a USD 200/kg launch scenario. This is a decision threshold, not evidence of a market. The conclusion is to build a commercial and technical proof path for orbital compute, while refusing to finance a gigawatt-scale constellation on solar geometry alone.

Keywords: orbital computing; sun-synchronous orbit; space solar power; thermal radiation; in-orbit data processing; levelized cost of compute; aerospace project finance.

1. Decision and unit of analysis

The location question must be asked at the level of a *workload and its data path*. For Earth-generated training corpora, a ground-to-orbit uplink and model-state exchange may nullify the benefit of orbital solar power. For satellite-generated imagery, hyperspectral cubes or RF measurements, computation before downlink can remove a constrained transport step. Both products may use the same spacecraft but have different revenue mechanisms. The minimum viable company proposition is therefore an orbital processing service for spacecraft operators, with exportable accelerator software and interfaces; the 250-kW spacecraft is a scale target, not the initial product.

The reference architecture is a 250-kW IT module at $h = 650$ km in a dawn–dusk sun-synchronous orbit (SSO), retrograde inclination near 98°, with a physically separate solar-collection wing and outward-facing radiator. A small constellation supports crosslinks to data sources and gateways. A single node should first run independent inference and preprocessing; synchronous multi-node training is deferred until throughput, clocking, formation safety and crosslink economics are proven. The SpaceX AI1 concept publishes 250-kW peak and 175-kW average compute-payload figures and a 75-kW/tonne

vehicle-efficiency figure. Google's Project Suncatcher examines an illustrative 81-node formation at 650 km, very short laser links, and a two-satellite learning mission. These are serious benchmarks and potential competitors; public concept specifications and planned missions are not proof of cost or operational availability.

Quantity	**Central scenario**	**Status and interpretation**
Orbit altitude; solar beta angle	650 km; 66 deg	Chosen geometry, not an ephemeris result
IT electric load; bus/IT ratio	250 kW; 1.15	Design assumption
Solar irradiance; PV efficiency; derating	1361 W/m²; 30%; 0.90	Engineering scenario
Radiator temperature; emissivity; absorbed flux	350 K; 0.90; 100 W/m²	Engineering scenario, one effective emitting face
PV, radiator areal mass	3; 3 kg/m²	Aggressive deployment assumptions
Battery usable depth; round-trip efficiency; pack density	0.80; 0.85; 200 Wh/kg	Sensitivity inputs
Useful accelerator efficiency; utilization	5×10^{11} FLOP/J; 0.75	Scenario, not a chip data-sheet peak
Ground/LEO service availability	0.995 / 0.970	Scenario, must be measured
Real discount rate; horizon	10%; 6 years	Financing scenario
Launch	USD 200, USD 500, USD 1000/kg	Scenarios; USD 200/kg is not a 2030 quote
Ground electricity	USD 80, USD 160, USD 240/MWh	Counterfactual scenarios

Each formula below uses *useful delivered* operations, not advertised accelerator peak. Define a workload-validity factor $\chi \in [0, 1]$, which penalizes error recovery, model incompatibility, communications stalls and work that customers cannot use. The numerical plots set $\chi = 1$ and are optimistic; operational validation must estimate it from real jobs.

2. Orbit choice and the non-negotiable eclipse calculation

For a circular orbit, $r = R_\oplus + h$, $n = \sqrt{\mu/r^3}$ and $P = 2\pi/n$. A first-order SSO sets the J_2 nodal precession equal to the mean solar angular rate:

$$\dot{\Omega} = -\tfrac{3}{2} J_2 n \left(\tfrac{R_\oplus}{r}\right)^2 \cos i \simeq \tfrac{2\pi}{365.2422\,\text{days}}.$$

At $h = 650$ km this gives $i \simeq 97.98^\circ$ in the simplified model. SSO preserves approximately local solar time; it does *not* prove permanent illumination. Let β be the Sun's angle to the orbit plane. In a cylindrical umbra approximation, full eclipse disappears only if

$$|\beta| \geq \beta_c = \sin^{-1}\left(\tfrac{R_\oplus}{r}\right),$$

and otherwise the fraction of one revolution in shadow is

$$f_e(\beta, h) = \tfrac{1}{\pi} \cos^{-1}\left[\tfrac{\sqrt{1 - (R_\oplus/r)^2}}{\cos\beta}\right], \qquad t_e = P f_e.$$

The 650-km circular calculation yields $P = 97.58$ min and $\beta_c = 65.15^\circ$. At an assumed $\beta = 66^\circ$ the ideal cylindrical model has no full-shadow eclipse; at 60° it gives 17.79 min per orbit. Penumbra, seasonal geometry, attitude, panel pointing, perturbations and station-keeping are outside this scalar calculation. A mission should use time-resolved Sun/Earth ephemerides for its exact orbit and launch epoch. The word “continuous” in a product specification requires that verified simulation and a storage margin, not an SSO label. NASA's eclipse/beta analysis supports treating beta and eclipse as design variables.

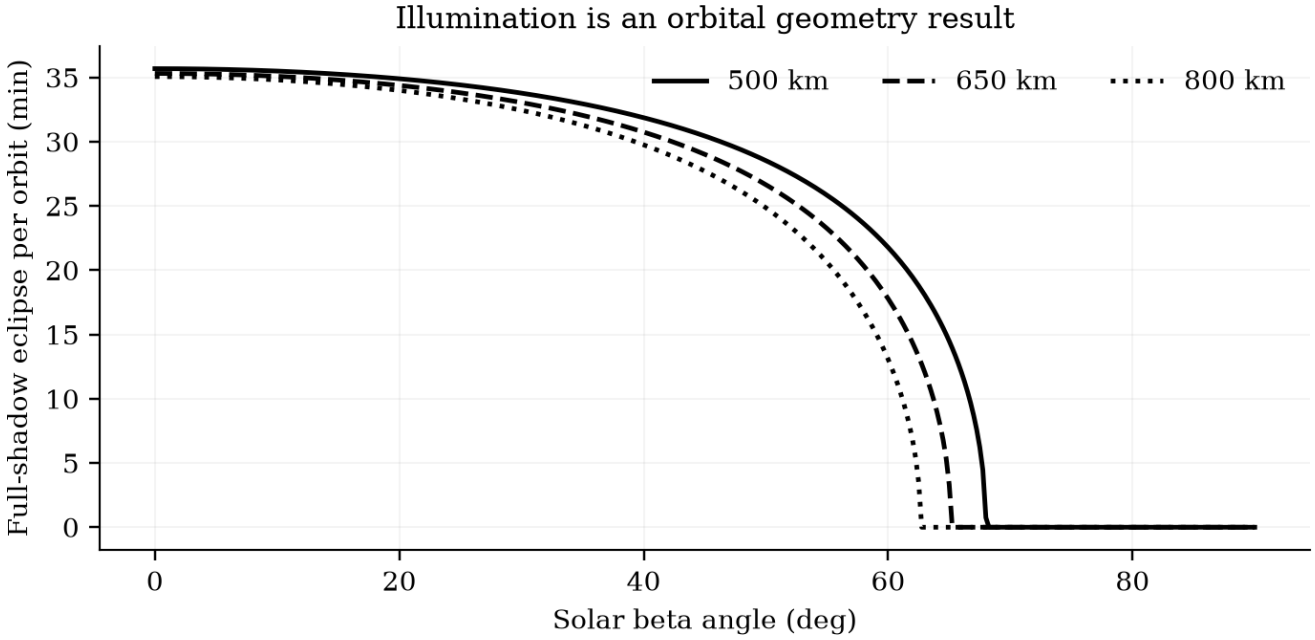


Figure 1. First-order cylindrical-shadow duration at three altitudes. The sharp zero-eclipse transition is an idealized geometry boundary; operational power design uses penumbra and time-varying beta.

3. Energy closure: the Sun is free, the collection system is not

Let P_{IT} be average IT power, $p = P_{bus}/P_{IT}$, G the incident irradiance, η_{PV} beginning-of-life conversion efficiency and d an aggregate end-of-life/pointing/electrical derating. If $f_e = 0$, the collection area is

$$A_{PV,\,0} = \frac{pP_{IT}}{G\eta_{PV}d}.$$

If a load must operate through eclipse, the daylit array must serve the contemporaneous load *and* replace energy drawn from storage. With battery round-trip efficiency η_{rt}, a per-orbit energy balance gives

$$P_{PV}(1-f_e)P = P_{bus}(1-f_e)P + \frac{P_{bus}f_eP}{\eta_{rt}},$$

$$A_{PV} = \frac{P_{bus}}{G\eta_{PV}d}\left(1 + \frac{f_e}{\eta_{rt}(1-f_e)}\right).$$

For allowable depth of discharge D, discharge efficiency η_d and installed specific energy e_b, the necessary nameplate capacity and battery mass are

$$E_{bat} = \frac{P_{bus}t_e}{D\eta_d}, \qquad m_{bat} = \frac{E_{bat}}{e_b}.$$

The code assumes $\eta_d = \sqrt{\eta_{rt}}$. At 250 kW IT, $P_{bus} = 287.5$ kW. The ideal $\beta = 66^\circ$ solution needs about 782 m² PV and zero *long-eclipse* storage in this model, but real designs still need ride-through energy and margins. At $\beta = 60^\circ$ it needs about 988 m² PV and 116 kWh of nameplate battery capacity (about 578 kg at 200 Wh/kg). Earth’s ground-level solar variation is thus an advantage for orbit only after conversion hardware, storage and end-of-life production are fully accounted for.

The PV mass assumption is particularly consequential. Three kilograms per square metre implies roughly 123 W/kg at this design point. NASA's spacecraft-power review describes flown solar-array missions clustering around approximately 30 W/kg, with a cited empirical upper end near 200 W/kg. The proposed array is within that reported envelope but substantially better than its cluster, and 250-kW deployability is not established by the existence of smaller arrays. At 30 W/kg, the array alone would approach 9.6 tonnes. Specific power, deployment reliability and aging are therefore first-class investment gates.

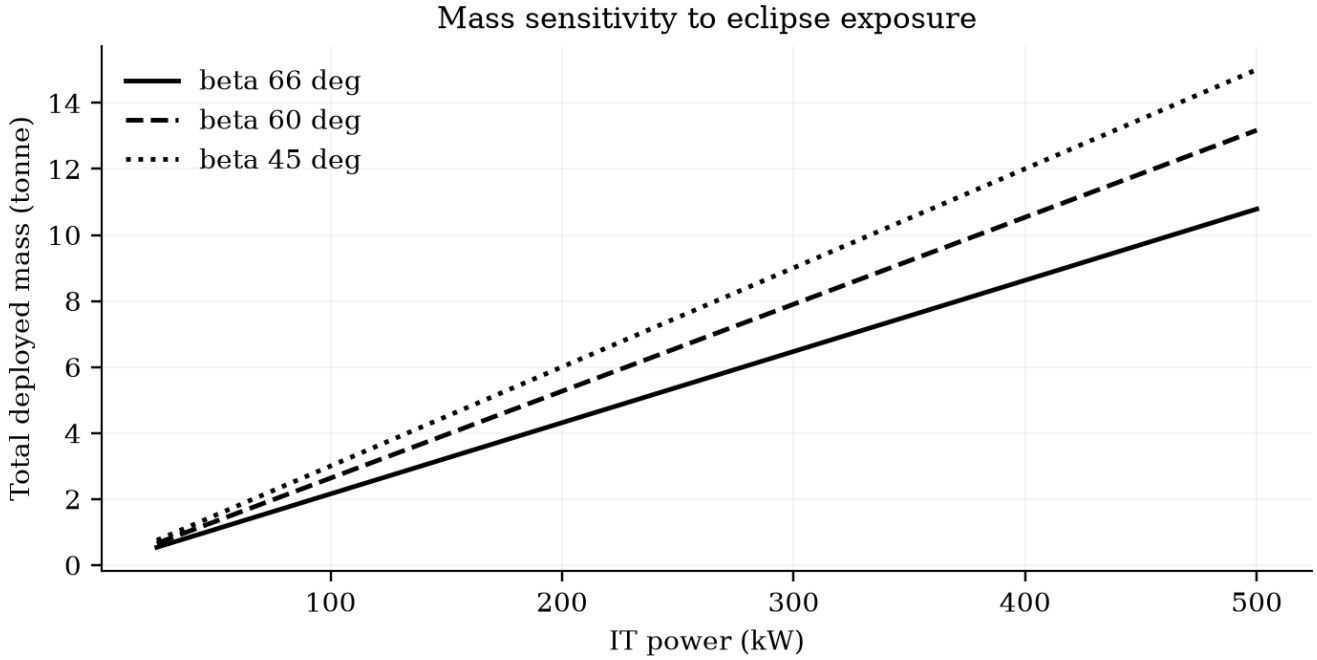


Figure 2. Deployed mass includes IT, bus, arrays, radiator and eclipse battery. It excludes detailed harness, propellant, shielding, packaging and reserve; the result is an optimistic mass floor.

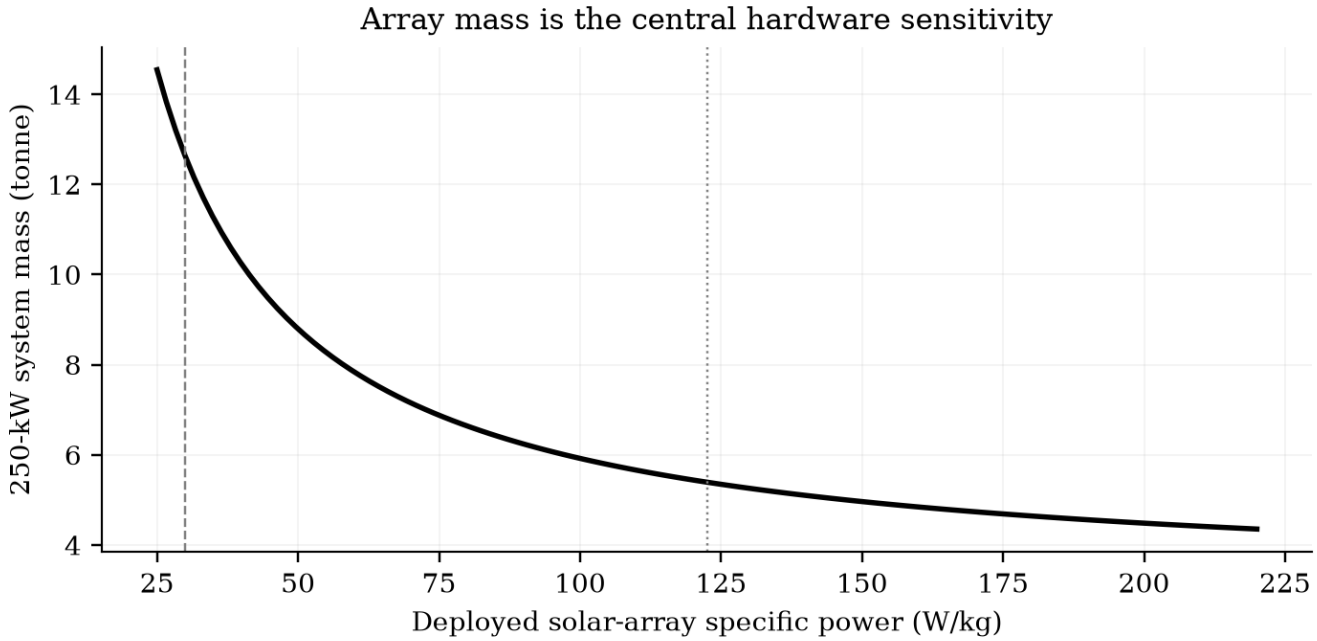


Figure 3. Solar-array specific power sensitivity for 250-kW IT at 650 km. The 30-W/kg marker represents the order of magnitude of NASA's flown-mission cluster, not a hard physical limit.

4. Thermal closure: radiative heat rejection is the second hard constraint

In vacuum, the platform cannot convect its compute heat to ambient air. For an effective one-sided radiator with view to cold space, conservative net rejectable flux is

$$q''_{rad} = \varepsilon\sigma T_r^4 - q''_{abs}, \qquad A_r = \frac{Q_{heat}}{q''_{rad}}.$$

Here $\sigma = 5.670374419 \times 10^{-8}\,\mathrm{W\,m^{-2}K^{-4}}$, $T_r = 350$ K, $\varepsilon = 0.9$ and a single lumped absorbed solar/albedo/terrestrial-IR flux $q''_{abs} = 100\,\mathrm{W/m^2}$. The result is $q''_{rad} = 665.82\,\mathrm{W/m^2}$. If all 287.5 kW bus input is ultimately dissipated as heat, then $A_r \simeq 432\,\mathrm{m^2}$, or approximately 1.73 m² per kW IT. This is a favorable orientation and heat-transport assumption; optical transmitters can export a small part of the energy, but the order of magnitude is unchanged. Higher radiator temperature lowers area as T_r^{-4} but forces higher semiconductor junction and coolant temperatures. Finite heat-pipe, pump, cold-plate and thermal-interface resistances therefore determine whether 350 K is achievable at the emitting surface while chips remain inside their limits. NASA's current thermal-control review specifically highlights mass, surface and deployable-radiator maturity as constraints.

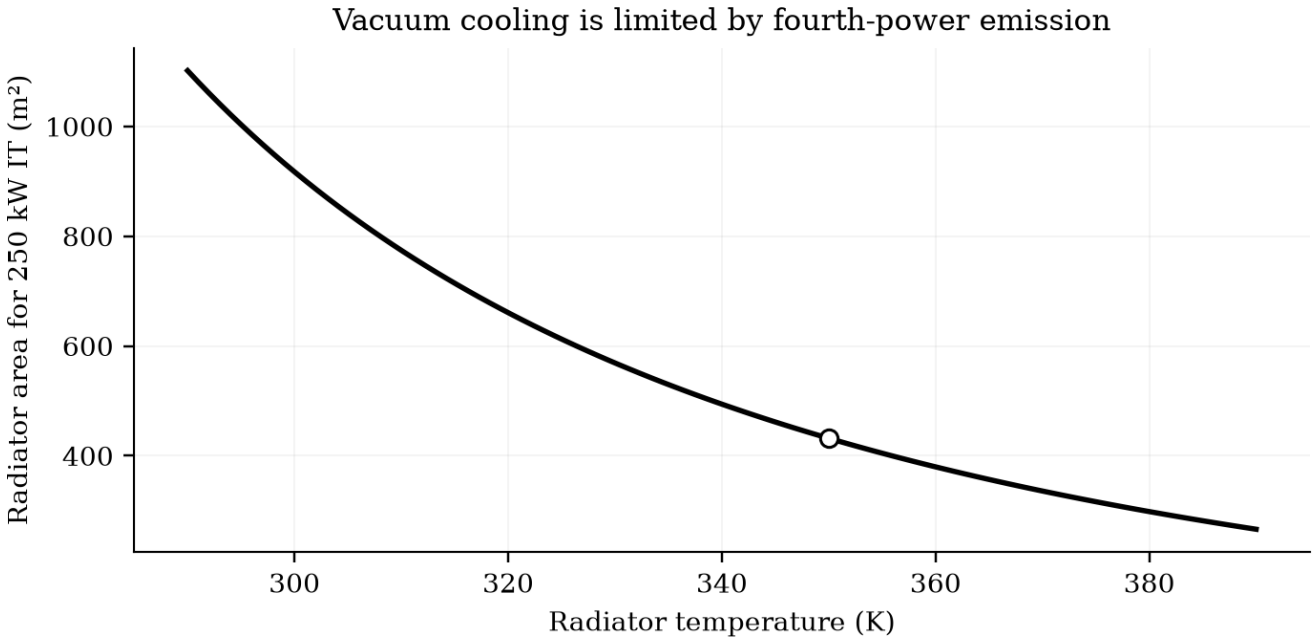


Figure 4. Required effective emitting area for a 287.5-kW heat load. Earth infrared, albedo and pointing are represented only by the 100-W/m² absorbed-flux scenario.

The reference deployed mass is $m_{PV} + m_r + m_{IT} + m_{bus} \simeq 5.39$ tonnes, or 21.57 kg/kW IT: approximately 2.35 tonnes PV, 1.30 tonnes radiator, 0.75 tonnes IT and 1.00 tonne other bus. This does not include a detailed launch-dispenser structure, shielding, propulsion, excess power redundancy, micrometeoroid margin or deorbit system. A hardware review must reconcile every subsystem, stowed volume and thermal path in a CAD-level mass/power/heat ledger. SpaceX's public 75-kW/tonne claim translates to 13.33 kg/kW under that particular definition; its payload-versus-bus power perimeter and operating point must be matched before a numerical one-to-one comparison.

5. Why 650-km orbital compute is energetically expensive to place

The circular speed at radius r is $v_c = \sqrt{\mu/r} \simeq 7.53$ km/s at 650 km. The ideal change in specific mechanical energy from rest at the equatorial surface to circular orbit is

$$\Delta e_{mech} = \left(-\frac{\mu}{2r}\right) - \left(-\frac{\mu}{R_\oplus}\right) = \frac{\mu}{R_\oplus} - \frac{\mu}{2r} \simeq 34.2\ \text{MJ/kg}.$$

This is a *mechanical lower-bound state change*, not rocket propellant energy, launch vehicle Δv or a billable launch price. Orbital insertion has atmospheric, gravity and trajectory losses, stage and propellant mass, and schedule risk. A 22-km aerostat only needs altitude potential energy of order $gh \approx 0.22$ MJ/kg but cannot access vacuum solar irradiance or remove the need for weather, buoyancy, station keeping and a physically huge airship. The decision variable here is delivered *economic service*, not launch energy alone.

With $m = 21.57$ kg/kW, one 250-kW unit is about 5.39 tonnes in the floor model. A 100-MW fleet would require roughly 2,157 tonnes of payload. If a launch system delivers 100 tonnes usable payload to the required orbit, 50% of each manifest is allocated to the service, and delivery success is 98%, the integer lower bound is

$$N_{flight} \geq \left\lceil \frac{P_{fleet}\, m_{spec}}{M_{payload}\, a\, s} \right\rceil.$$

That gives 45 launches for 100 MW and 441 for 1 GW before replacement satellites. Over a three-year build, those are approximately 15 and 147 launches annually. A startup cannot assume this cadence is contractually available. The 250-kW commercial test is a few-tonne satellite; the gigawatt narrative is a different industrial program.

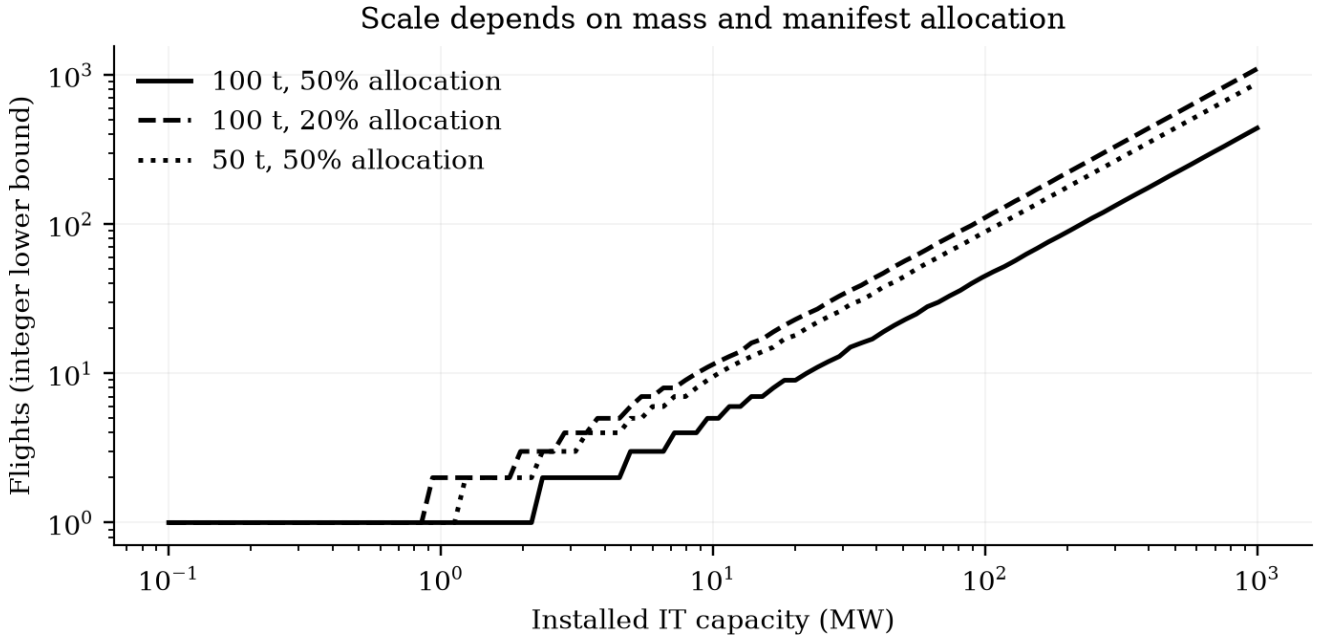


Figure 5. Integer flight lower bounds for different payload allocations. They exclude spare launches, commissioning attrition, replenishment and launch-slip buffers.

6. Optical network and end-to-end service rather than a speed-of-light slogan

For a diffraction-limited optical link at $\lambda = 1550$ nm, transmitter aperture D_t, receiver aperture D_r, distance R and efficiencies η_t, η_r, the far-field spot radius scales as $w \propto \lambda R / D_t$. A conservative Gaussian-style received-power approximation can be written

$$P_r = P_t \eta_t \eta_r \eta_{point} \eta_{atm} \left[1 - \exp\left(-\frac{D_r^2}{2w^2}\right)\right], \qquad w \approx \frac{2\lambda R}{\pi D_t}.$$

The exact constant depends on aperture illumination and beam definition; the purpose is to expose the dependencies. The design closes only if P_r exceeds the receiver sensitivity plus margin at the *coded* data

rate, with atmospheric attenuation η_{atm} for ground links, pointing loss η_{point} and availability of optical gateways. Short spacecraft-to-spacecraft links may carry raw sensor products to a shared accelerator node, while results reach Earth through a relay or ground station. Google's public Suncatcher analysis says tightly coupled distributed training would require tens of Tb/s between nodes; a commercial first mission should instead keep inference local and payload data paths simple. The Google design note reports an 800-Gb/s-each-way bench link, which is not a deployed space-network SLA.

For a numerical screening link at 1,550 nm, take 1 W optical transmit power, 10-cm transmit and receive apertures, 50% optical efficiency at each terminal and 50% pointing efficiency. The far-field spot radius is about 0.99 m at 100 km and 9.87 m at 1,000 km. The resulting average received powers are approximately −1.94 dBm and −21.93 dBm, respectively. These are *power estimates*, not throughput guarantees: sensitivity, background photons, detector noise, coding, acquisition time and safety margins still determine usable bits per second. The result makes distance and pointing an explicit service-design constraint.

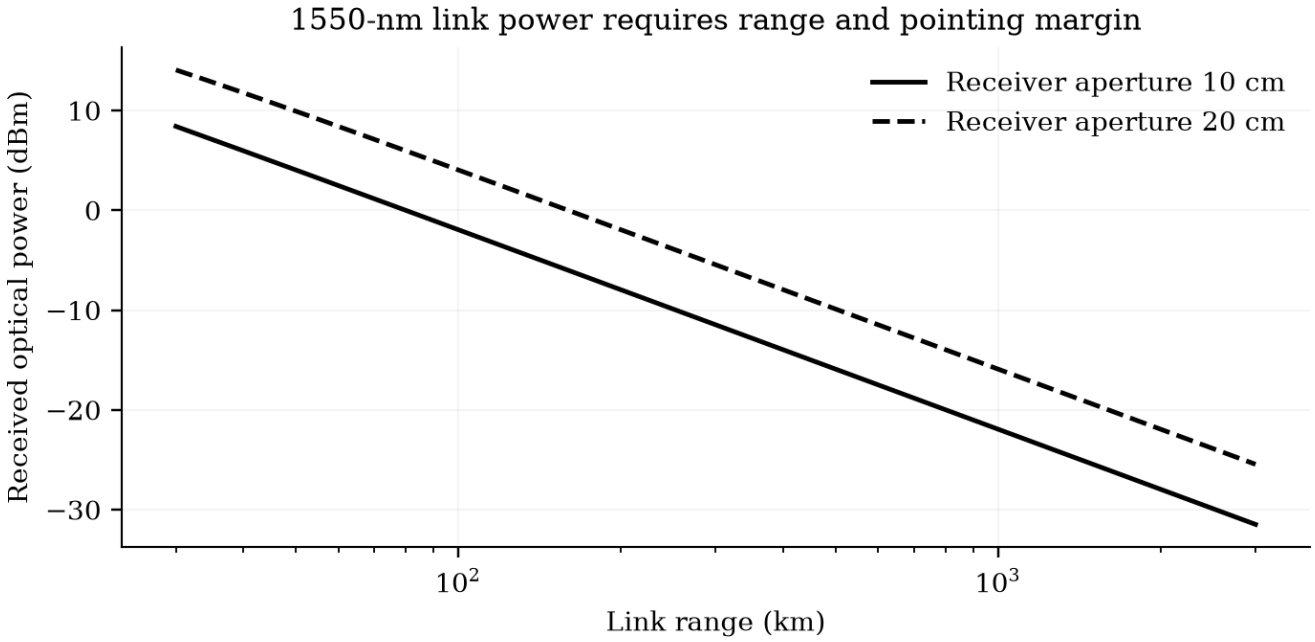


Figure 6. Received optical power for 1-W, 10-cm transmitter at 1,550 nm. The two curves vary receiver aperture; the model holds pointing and terminal efficiency constant and excludes atmospheric loss.

For a job of input size B_{in}, output size B_{out}, link rate C, queueing t_q, compute t_c and propagation R/c, one possible end-to-end bound is

$$t_{orb} = t_{crosslink} + \frac{B_{in}}{C_{in}} + t_q + t_c + \frac{B_{out}}{C_{down}} + t_{gateway},$$

$$t_{ground} = t_{raw\,downlink} + \frac{B_{in}}{C_{down}} + t_{ground\,queue} + t_{ground\,compute}.$$

At 650 km, the single-path vacuum light time to a nadir ground point is about 2.2 ms, but contacts, routing, clouded optical gateways and downlink serialization commonly dominate. Any asserted multi-minute latency advantage must be measured as a distribution by orbit and customer location; it cannot be inferred from R/c. Figure 7 is deliberately a *scenario sketch*, not benchmark data.

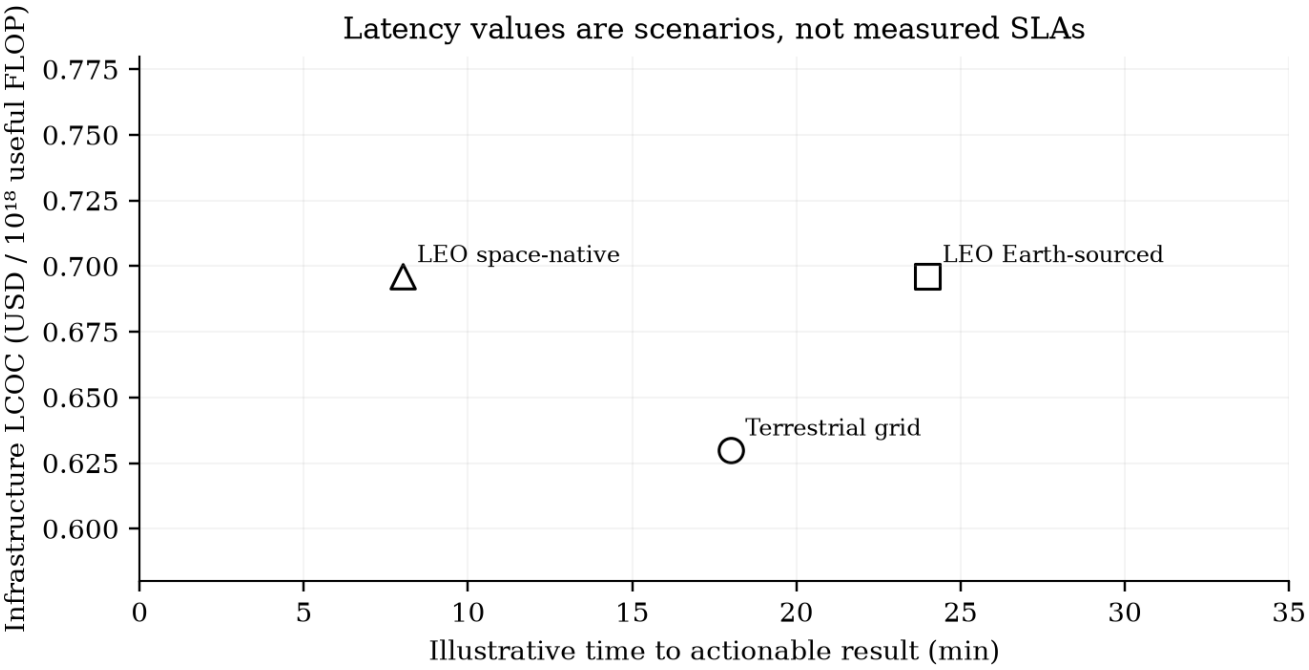


Figure 7. Illustrative job-level time-to-result assumptions are 18 min for the terrestrial path, 24 min for an Earth-sourced orbital path and 8 min for a space-native orbital path. These are unvalidated scenario inputs; infrastructure LCOC is shown without customer value.

7. Product-market-fit hypothesis: data born in space

ESA's operational Φsat-2 mission and NASA's Dynamic Targeting work already demonstrate why some sensing data should be analyzed before ground transfer. This is prior art, not evidence that a separate, multi-hundred-kW shared data centre has product-market fit. Small processors attached to each sensor may be cheaper and simpler. The proposed company must prove that one shared compute node creates extra value through a high-demand model, multi-sensor fusion, model refreshability, aggregation across satellites or a lower compute cost per validated result that outweighs crosslink integration.

Let B_r be raw source data per day, B_o the output downlinked after processing, v_t the customer's *actual marginal* value of transport capacity, V_{time} the monetary value of earlier insight and $V_{quality}$ the incremental value from a better product. For a given customer and service period,

$$V_{orb} = v_t(B_r - B_o) + V_{time} + V_{quality} - C_{crosslink} - C_{integration} - C_{rights}.$$

The project is commercially justified only when contracted incremental value covers the full orbital cost premium and a risk-adjusted capital return. Savings in bytes are not revenue by themselves: some operators own sunk ground capacity, cannot legally export raw data, or would run adequate edge inference on their own spacecraft. Customer interviews must collect actual volume, time sensitivity, ground-station marginal prices, error tolerance, data rights, API integration cost and procurement authority. A credible first customer is an *existing operator with a measured bottleneck* and a paid trial, not a hypothetical global AI buyer.

Workload selection matrix

Candidate	Orbital advantage	Primary failure mode	First evidence required
Earth-observation image filtering/fusion	Avoid raw downlink; act earlier	Onboard edge chips already sufficient	Paid multi-sensor or high-model-value workflow
SAR/hyperspectral preprocessing	Potentially high raw-data volume	Quality loss or sensor-specific integration	Reversible fidelity test and operator data rights
Spacecraft operations inference	Local reaction time	Hard safety certification	Bounded, non-critical pilot
Earth-sourced asynchronous inference	Continuous solar power	Uplink economics and ground competition	Binding price/volume above orbital LCOC
Distributed frontier training	Solar capacity at long horizon	Tb/s links, synchronization, refresh and cost	Multi-node test with delivered useful FLOP

The chosen launch wedge is shared, non-safety-critical processing of satellite-generated data. A venture-scale market is *not* established by this wedge alone. An investment committee must independently test obtainable customer spend and the expansion path to wider compute before financing manufacturing capacity.

8. Economic model and the break-even customer-value test

For horizon N, discount rate r, annual capital refresh C_t, annual operating cost O_t and delivered useful work F_t, define

$$\mathrm{LCOC} = \frac{C_0 + \sum_{t=1}^{N} (C_t + O_t)/(1+r)^t}{\sum_{t=1}^{N} F_t/(1+r)^t}, \qquad F_t = P_{IT}\,\eta_{useful}\,u\,a\,\chi\,T_t.$$

Here u is utilization, a service availability and T_t seconds in year t. For constant annual flow, the present-value annuity factor is $A(r,N) = \sum_{t=1}^{N} (1+r)^{-t} = 4.3553$ at $r = 10\%$, $N = 6$. The published code evaluates an illustrative 1-MW IT denominator with $\eta_{useful} = 5 \times 10^{11}$ FLOP/J, $u = 0.75$, $a = 0.97$ in orbit and $a = 0.995$ on Earth. It holds hardware efficiency flat for six years. That assumption is favorable to a hard-to-service satellite if terrestrial accelerators refresh quickly; a proper investment case must model upgrade cadence and residual value.

Terrestrial scenario capital is 26.0 million/MW IT, including an 18.0 million IT portion; fixed annual OPEX is 0.8 million/MW. Electricity uses $PUE = 1.15$, an idle power fraction of 0.20 when not utilized and $p_E = 80$/MWh. Orbital non-launch capital is an optimistic 26.607 million/MW IT and fixed annual OPEX is 1.1 million/MW. The revised deployed mass of 21.57 kg/kW is charged at the selected all-in USD/kg launch scenario. These are illustrative engineering cash flows, with no firm launch quotation, launch insurance, unusual deorbit liability, customer acquisition, data rights or revenue guarantee.

$$C_{orb,\,0} = 26.607\,\mathrm{MUSD/MW} + \frac{m_{spec}\,p_{launch}}{1000},$$

$$O_{grid} = 0.8\,\mathrm{MUSD/MW/yr} + \frac{p_E\,PUE\,[u + (1-u)f_{idle}]\,8760}{10^6}.$$

At a launch price of USD 200/kg, the reference yields orbital LCOC USD 0.715 per 10^{18} useful FLOP and grid LCOC USD 0.630 per 10^{18} useful FLOP. At USD 500/kg orbital LCOC is USD 0.844, and at USD 1000/kg it is USD 1.060. These cents-per-exaflop values are *model outputs*, not obtainable customer prices. At USD 80/MWh, the large fleet does not win commodity compute on LCOC. For a same-workload commodity comparison, the implied terrestrial electricity price at cost parity is approximately USD

204/MWh at USD 200/kg launch, USD 393/MWh at USD 500/kg and USD 708/MWh at USD 1000/kg, holding all other assumptions fixed. At a grid price of USD 80/MWh, launch would have to approach about USD 4/kg for nominal LCOC parity under this model. These extremes quantify why the initial business case must monetize space-native data value rather than rely on solar access alone. Ground sites with faster grid interconnection or cheap behind-the-meter supply can strengthen the counterfactual further. Scarce-grid premiums, carbon contracts, other terrestrial bottlenecks and launch learning could reverse the ranking, but must enter as measured site-specific cash flows.

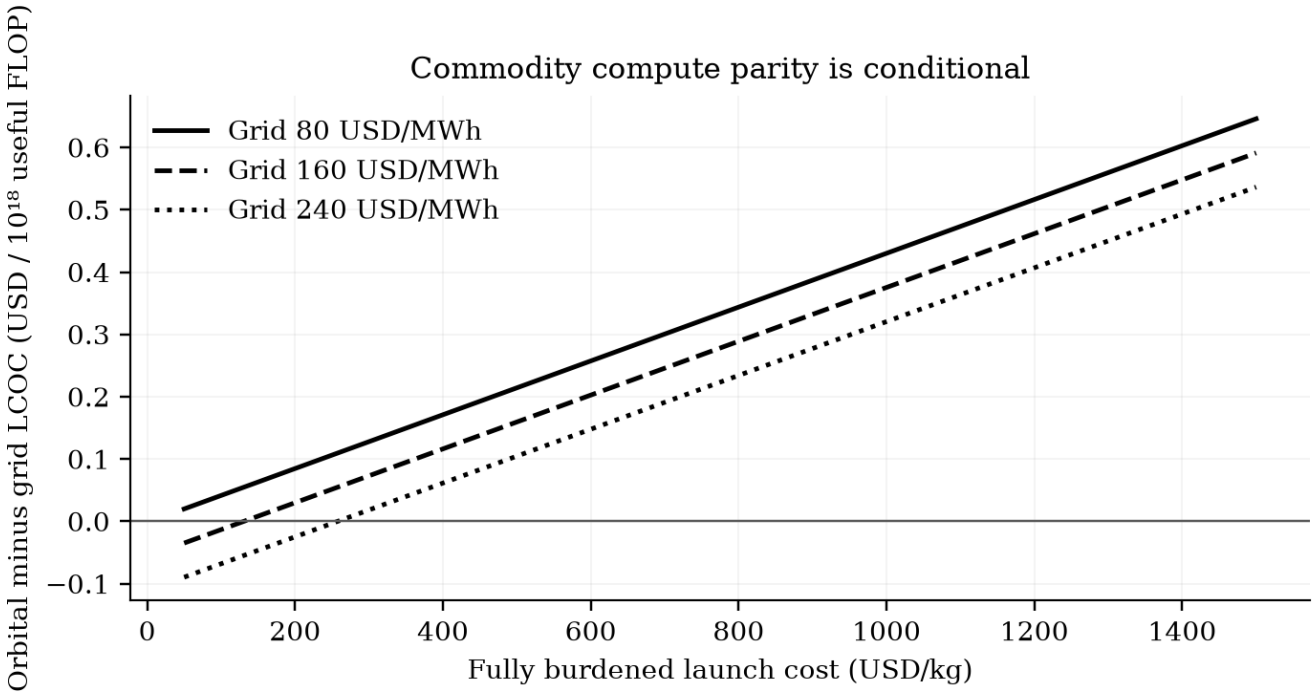


Figure 8. Orbital minus terrestrial infrastructure cost at three electricity-price scenarios. The USD 200/kg case is a sensitivity point, not a promised 2029–2030 procurement price.

For space-native service, normalize the orbital premium to annual 1-MW output: $F_{orb} = 11.471$ million units of 10^{18} useful FLOP per year, so the USD 200/kg infrastructure premium at 80/MWh is about USD 0.971 million per MW-year. Let first-generation non-recurring engineering be C_{NRE}, funded-fleet size M MW and incoming traffic D TB/day/MW. An intentionally simple break-even *incremental customer value per incoming TB* is

$$v_{min} = \frac{10^6\left[(\mathrm{LCOC}_{orb} - \mathrm{LCOC}_{grid})F_{orb}/10^6 + C_{NRE}/(MA)\right]}{365D}.$$

With $C_{NRE} = 100$ million, $M = 10$ MW and $D = 100$ TB/day/MW, the threshold is USD 89.5/TB at USD 200/kg, USD 130/TB at USD 500/kg and USD 198/TB at USD 1000/kg. At 1,000 TB/day/MW and USD 200/kg, it falls to USD 8.95/TB. These are *incremental value* thresholds; the first 250-kW pilot cannot amortize USD 100 million over an assumed 10-MW fleet unless financing explicitly prices the fleet expansion. A stand-alone pilot has far worse fully loaded economics. Conversely, lower development spend, customer-financed payloads or a platform partnership can improve the threshold. The formula excludes optical link capital and traffic-related costs, so it remains optimistic.

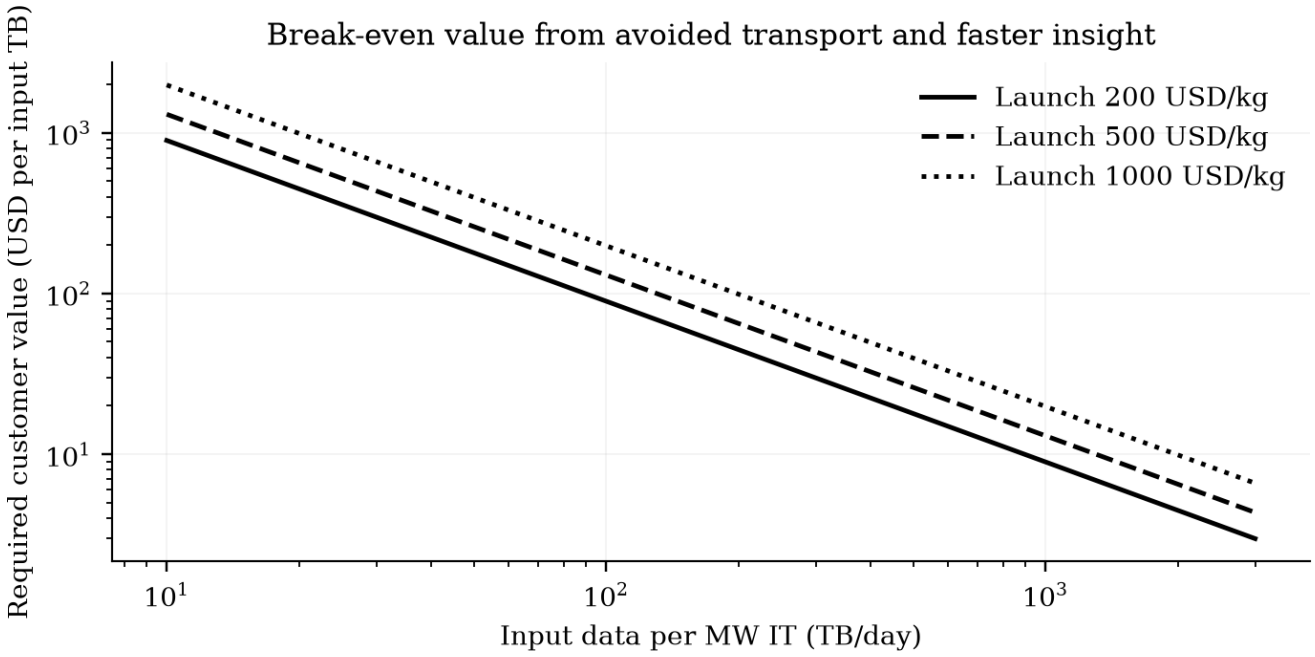


Figure 9. Required value includes the modeled orbital LCOC premium and USD 100 million non-recurring engineering amortized across 10 MW. The premise requires proved sellable data volume and paid value, not merely bytes.

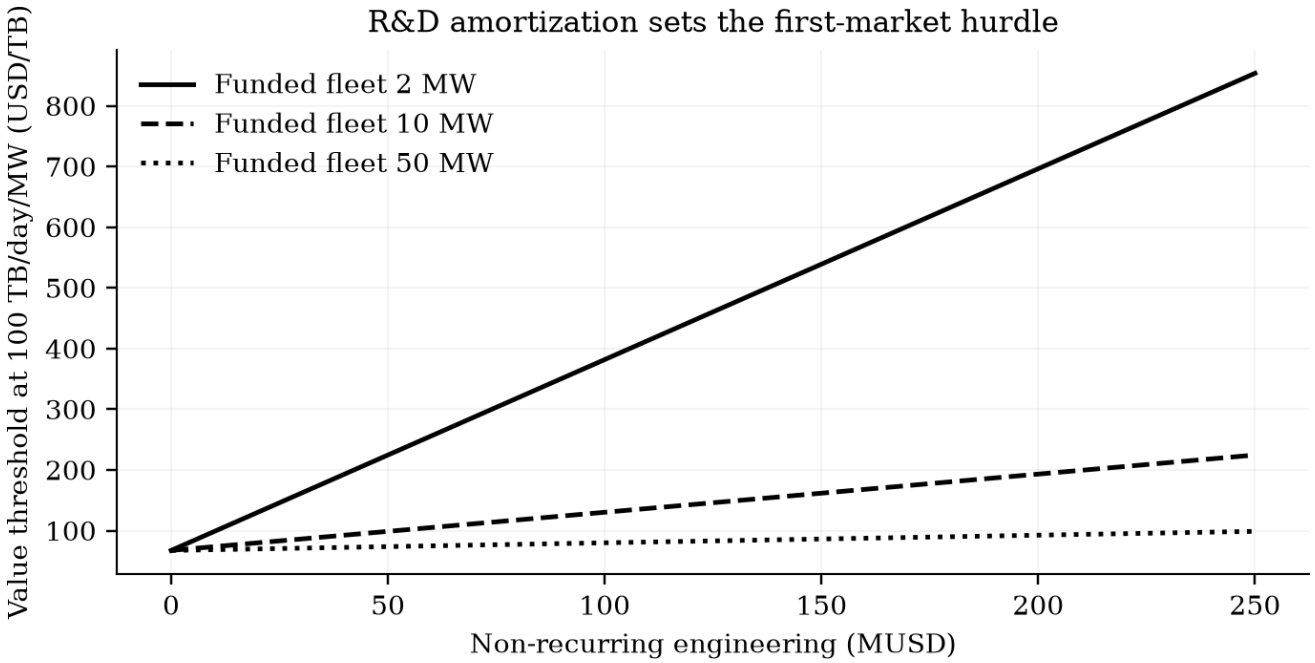


Figure 10. Development cost can dominate a small fleet. Raising money against a gigawatt total-addressable-market slide does not pay for a first customer contract.

The value estimate should be stress-tested with launch-cost distributions, loss of one satellite, short hardware life, radiation-induced faults, radiator coating aging, array degradation, queueing, crosslink unavailability, ground optical weather, and terrestrial accelerator refresh. A minimum financially coherent model includes tax, insurance, working capital, spare inventory, debt service reserve, replacement timing, residual value and a contracted revenue schedule. No project debt should be underwritten from a theoretical LCOC chart alone.

9. Benchmarks against SpaceX and Google

The strongest public orbital benchmark is vertically integrated: SpaceX can co-design satellite production, launch and a laser relay network. Its AI1 page asserts up to 250-kW peak compute, 175-kW average, continuous solar capture in SSO and large-scale factory plans. An independent startup cannot assume identical transfer pricing, launch priority or Starlink connectivity. Its potential advantage must be portability across launchers, payload hosts and sensor constellations; a narrow workflow and customer integration; or a specific power/thermal/computing component with defensible IP. Competing head-on to build a generic Starship-fed commodity orbital cloud would require capital and integration advantages the startup has not demonstrated.

Google Suncatcher identifies 650-km compact formations, near-constant solar exposure, laser-throughput scaling, TPU radiation tests and a less-than-USD 200/kg *mid-2030s projection* for launch. Its own account describes open challenges in thermal management, ground communications and reliability. The thesis here differs in sequencing: independent revenue from space-generated data precedes high-bandwidth distributed training. If Google or SpaceX can provide the same service at lower delivered cost, the proposed venture's independent infrastructure moat is weak; the business must then own a customer workflow, a specific technical subsystem or an integration position.

ESA Φsat-2 proves that orbital edge processing is already operational at small scale. It is the correct low-end benchmark. The startup must show a demand class for shared, higher-power, multi-tenant orbital computation that cannot be served more cheaply by a modest processor directly on each sensing spacecraft.

10. Project-finance architecture and investability gates

The first stage is a software-and-data-access business, not an infrastructure special-purpose vehicle (SPV). It requires an operator agreement granting defined rights to process sensor data, an integration plan, a paid trial with a baseline on the operator's existing edge/ground pipeline, and a jointly agreed metric such as dollars per actionable detection, time-to-insight distribution or recovered downlink capacity. A non-binding expression of interest is not an offtake contract.

Only after a repeatable buyer and a verified payload mass/power/thermal design should the company create a hardware SPV. A plausible capital stack is sponsor equity plus strategic customer prepayments and milestone financing during qualification, followed by limited project debt only when availability, launch, replacement and service revenues have credible contractual support. Debt has priority over equity, so one lost satellite or a delayed launch can threaten covenants even when the long-run technology remains promising. Capacity offtake should price *usable service* (data products or useful compute with an SLA) rather than theoretical kW. Customers should be able to reject invalid results, and the SPV should reserve for replacements and deorbit obligations.

A falsifiable 2026–2030 sequence is:

Decision date	Fundable proof	Kill or redesign condition
2026–2027	Paid satellite-data workflow on ground and/or existing edge hardware; measured raw/output bytes, value, demand owner and access rights	No measurable incremental value above customer integration cost
2027–2028	Full thermal loop, array deployment, radiation and FSO hardware tested under representative conditions; orbit-specific power simulation	Mass/thermal ledger fails payload and temperature margins
2028–2029	Small flight demonstrator with independently measured useful FLOP/J, availability, crosslink/gateway throughput and data-product accuracy	Service-quality or data-rights failure; cloud/edge baseline dominates
2029–2030	A 100–250-kW class pilot only after priced launch, two independent buyer commitments and a financeable replacement plan	Fleet financing depends on unquoted USD 200/kg launch or future gigawatt scale

These dates are *gates*, not a claim that a 250-kW mission has been booked. The mission definition should include formal failure modes: eclipse beyond power reserve, one array wing undeployed, radiator view-factor shortfall, failed pump/heat pipe, HBM upset, one optical terminal lost, no cloud-free gateway, launch slip and customer data embargo. Establish a measured requirement matrix and environmental tests before negotiating fixed-price commercial SLAs.

11. Strategic outlook and conclusion

The strongest defendable 2029–2030 thesis is **near-continuous-solar orbital compute for high-value, space-native data**, with an option to expand toward Earth-serving AI only after true delivered costs and network throughput are measured. The nominal 250-kW module closes algebraically at 650 km with about 782 m² of PV, 432 m² of radiator and a 5.39-tonne optimistic mass floor under a favorable 66-degree beta angle. A six-degree drop in beta creates nearly 18 minutes of eclipse, adding array and battery burden. These values turn the concept into a testable engineering object, not a claim of field-ready hardware.

The project's commercial test is harder than the physics: its proposed shared processing must outperform sensor-local edge chips and terrestrial pipelines for a buyer who pays. The nominal infrastructure model loses to an USD 80/MWh grid on generic compute, while a 10-MW space-native fleet could cover the modeled premium plus USD 100 million development cost if customers realize at least about USD 90/TB of incremental value at 100 TB/day/MW and USD 200/kg launch. No source establishes that this combination of launch price, traffic, willingness to pay and fleet utilization will exist in 2030. A company should proceed first with customer data and hardware proofs. It should reject a premature gigawatt financing plan and preserve the orbital solar thesis as an option unlocked by evidence.

Appendix A. Reproducibility and limits

Run `python3 figures_and_model.py` with NumPy and Matplotlib. The script generates Figures 1–10 as PNG files and writes `model_results.json` with the reference calculations. The code intentionally exposes all scenario coefficients. It is a deterministic screening model, not a high-fidelity orbit propagator, finite-element thermal model, communications simulator, mission reliability model or bankable financial forecast. Visual equality of two LCOC curves is not investment parity when the products have different data rights, development schedules and counterparty risks.

The model uses a cylindrical Earth shadow; a single-surface radiator flux; fixed solar conversion and array areal mass; one battery round-trip efficiency; fixed six-year accelerator performance; no taxation, salvage, insurance or launch failures in LCOC; constant USD/kg launch; fixed 97% orbital availability; and a hypothetical $100-million development program. Each of these must be replaced by measured or contractually supported inputs before a board commits to production capital.

References

• International Energy Agency, *Key Questions on Energy and AI*, executive summary, 2026.

• SpaceX, *AI1/Starmind*, product concept page, accessed 25 September 2026.

• Google Research, *Exploring a space-based, scalable AI infrastructure system design*, 2025.

• NASA, *Small Spacecraft Technology State of the Art: Power Subsystems*, 2026.

• NASA, *Small Spacecraft Technology State of the Art: Thermal Control*, 2026.

• ESA, *Φsat-2 begins science phase for AI Earth images*, 2025.

• NASA/JPL, *How NASA Is Testing AI to Make Earth-Observing Satellites Smarter*, 2025.

• NASA, *Sun-synchronous eclipse and beta-angle analysis*.

Appendix B. Complete Python / Matplotlib reproduction script

The following script is self-contained. Save it as `figures_and_model.py` and run it with Python, NumPy and Matplotlib. It creates all ten figures and the calculation ledger.

```
#!/usr/bin/env python3
"""Reproduce the illustrative orbital-compute project figures and checks.
All financial coefficients are explicit scenarios, not supplier quotations or forecasts.
Run: python3 figures_and_model.py
Requires numpy and matplotlib. Output: figures/*.png and model_results.json.
"""
from pathlib import Path
import json
import math
import numpy as np
import matplotlib
matplotlib.use("Agg")
import matplotlib.pyplot as plt
HERE = Path(__file__).resolve().parent
FIG = HERE / "figures"
FIG.mkdir(exist_ok=True)
plt.rcParams.update({
    "font.family": "DejaVu Serif", "font.size": 9, "axes.labelsize": 9,
    "axes.titlesize": 10, "axes.linewidth": 0.7, "xtick.direction": "out",
    "ytick.direction": "out", "savefig.dpi": 220,
})
RE = 6371e3
MU = 3.986004418e14
SIGMA = 5.670374419e-8
```

```
SOLAR = 1361.0
YEAR = 365 * 24 * 3600
EF = 1e18
def eclipse(alt_km, beta_deg):
    r = RE + np.asarray(alt_km, dtype=float) * 1000
    b = np.deg2rad(np.asarray(beta_deg, dtype=float))
    beta_c = np.arcsin(RE / r)
    ratio = np.sqrt(1 - (RE / r)**2) / np.cos(b)
    frac = np.where(np.abs(b) < beta_c, np.arccos(np.clip(ratio, -1, 1)) / np.pi, 0.)
    period = 2*np.pi*np.sqrt(r**3/MU)
    return frac, frac*period/60, np.rad2deg(beta_c), period/60
def design(p_it_kw=250, beta_deg=66, alt_km=650, radiator_K=350,
           pue=1.15, pv_eff=.30, derate=.90, battery_rt=.85,
           dod=.80, battery_Wh_kg=200., pv_kg_m2=3.0,
           rad_kg_m2=3., it_kg_kw=3., other_kg_kw=4.):
    f, mins, bc, period = eclipse(alt_km, beta_deg)
    f = float(f)
    power_W = p_it_kw * 1000 * pue
    pv_W = power_W*(1 + f/(battery_rt*(1-f)))
    a_pv = pv_W/(SOLAR*pv_eff*derate)
    e_bat_Wh = power_W * float(mins)/60/(dod*math.sqrt(battery_rt))
    m_bat = e_bat_Wh/battery_Wh_kg
    q_rad = .90*SIGMA*radiator_K**4 - 100
    a_rad = power_W/q_rad
    m = a_pv*pv_kg_m2 + a_rad*rad_kg_m2 + m_bat + p_it_kw*(it_kg_kw+other_kg_kw)
    return dict(pv_m2=a_pv, radiator_m2=a_rad, battery_kWh=e_bat_Wh/1000,
                battery_kg=m_bat, mass_kg=m, kg_per_kW=m/p_it_kw,
                eclipse_min=float(mins), eclipse_fraction=f,
                critical_beta_deg=float(bc), orbit_period_min=float(period),
                radiator_net_W_m2=q_rad)
def cost(launch_usd_kg=200, electricity_usd_MWh=80, nre_MUSD=100,
         installed_MW=10, stream_TB_day=100):
    """Six-year constant-performance comparison, USD per 1e18 useful FLOP."""
    n, rate, u, eta = 6, .10, .75, 5e11
    af = sum((1+rate)**(-t) for t in range(1,n+1))
    it_kg_kw = design()["kg_per_kW"]
    # 26.607 MUSD/MW non-launch is an optimistic installed-spacecraft
    # scenario from the component model; charge launch against revised mass.
    orb_cap = 26.607 + launch_usd_kg*it_kg_kw/1000
    orb_opex = 1.1
    ground_cap = 26.
    ground_opex = .8 + electricity_usd_MWh*1.15*(u+(1-u)*.2)*8760/1e6
    f_orb = 1e6*eta*u*.97*YEAR/EF
    f_ground = 1e6*eta*u*.995*YEAR/EF
    lc_orb = (orb_cap/af + orb_opex)*1e6/f_orb
    lc_ground = (ground_cap/af + ground_opex)*1e6/f_ground
    annual_premium_MUSD_MW = (lc_orb-lc_ground)*f_orb/1e6
    nre_annual_MUSD_MW = nre_MUSD/(af*installed_MW)
    threshold_TB = (annual_premium_MUSD_MW+nre_annual_MUSD_MW)*1e6/(365*stream_TB_day)
    return dict(lcoc_orbit_usd_EF=lc_orb, lcoc_ground_usd_EF=lc_ground,
                annual_orbital_premium_MUSD_MW=annual_premium_MUSD_MW,
                nre_annual_MUSD_MW=nre_annual_MUSD_MW,
                breakeven_value_usd_TB=threshold_TB,
                annual_factor=af, annual_EF_orbit_MW=f_orb)
def optical_received(R_km, D_t_m=.10, D_r_m=.10, P_t_W=1.,
                     eta_t=.5, eta_r=.5, eta_point=.5, eta_atm=1.):
    """Illustrative Gaussian far-field average power, not a coded-link SLA."""
    w=2*1550e-9*(R_km*1000)/(np.pi*D_t_m)
    capture=1-np.exp(-D_r_m**2/(2*w**2))
    return P_t_W*eta_t*eta_r*eta_point*eta_atm*capture
def style(ax):
    ax.spines[["right", "top"]].set_visible(False)
    ax.grid(alpha=.18, lw=.5)
def save(fig, name):
    fig.tight_layout(pad=1.2)
    fig.savefig(FIG / name, bbox_inches="tight")
    plt.close(fig)
def plots():
    # 1: actual dependence on beta; a dawn-dusk orbit requires ephemeris verification.
    betas=np.linspace(0,90,361)
    fig,ax=plt.subplots(figsize=(6.5,3.3))
```

```
for h,ls in [(500,"-"),(650,"--"),(800,":")]:
    ax.plot(betas,[eclipse(h,b)[1] for b in betas],ls,color="black",lw=1.6,label=f"{h} km")
ax.set(xlabel="Solar beta angle (deg)",ylabel="Full-shadow eclipse per orbit (min)",
       title="Illumination is an orbital geometry result")
ax.legend(frameon=False,ncol=3);style(ax);save(fig,"01_eclipse_geometry.png")
# 2: power hardware under eclipse, including charge-loss and battery mass.
fig,ax=plt.subplots(figsize=(6.5,3.4))
for b,ls in [(66,"-"),(60,"--"),(45,":")]:
    d=[design(p_it_kw=p,beta_deg=b) for p in np.linspace(25,500,80)]
    ax.plot(np.linspace(25,500,80),[x["mass_kg"]/1000 for x in d],ls,
            color="black",lw=1.6,label=f"beta {b} deg")
ax.set(xlabel="IT power (kW)",ylabel="Total deployed mass (tonne)",
       title="Mass sensitivity to eclipse exposure")
ax.legend(frameon=False);style(ax);save(fig,"02_mass_power.png")
# 3: radiator physics. One effective emitting side; flux includes 100 W/m2 absorbed.
fig,ax=plt.subplots(figsize=(6.5,3.3))
ts=np.linspace(290,390,100)
q=.90*SIGMA*ts**4-100
ax.plot(ts,287500/q,color="black",lw=1.8)
ax.scatter([350],[287500/(.90*SIGMA*350**4-100)],facecolors="white",
           edgecolors="black",zorder=3)
ax.set(xlabel="Radiator temperature (K)",ylabel="Radiator area for 250 kW IT (m²)",
       title="Vacuum cooling is limited by fourth-power emission")
style(ax);save(fig,"03_radiator.png")
# 4: price parity against terrestrial grid and launch sensitivity.
fig,ax=plt.subplots(figsize=(6.5,3.3))
launches=np.linspace(50,1500,100)
for e,ls in [(80,"-"),(160,"--"),(240,":")]:
    ax.plot(launches,[cost(x,e)["lcoc_orbit_usd_EF"]-
                      cost(x,e)["lcoc_ground_usd_EF"] for x in launches],
            ls,color="black",lw=1.6,label=f"Grid {e} USD/MWh")
ax.axhline(0,color="0.35",lw=.8)
ax.set(xlabel="Fully burdened launch cost (USD/kg)",
       ylabel="Orbital minus grid LCOC (USD / 10¹⁸ useful FLOP)",
       title="Commodity compute parity is conditional")
ax.legend(frameon=False);style(ax);save(fig,"04_cost_parity.png")
# 5: value per incoming TB needed for an orbital-native workload.
fig,ax=plt.subplots(figsize=(6.5,3.3))
traffic=np.geomspace(10,3000,100)
for l,ls in [(200,"-"),(500,"--"),(1000,":")]:
    ax.loglog(traffic,[cost(l,80,100,10,x)["breakeven_value_usd_TB"]
                       for x in traffic],ls,color="black",lw=1.6,
              label=f"Launch {l} USD/kg")
ax.set(xlabel="Input data per MW IT (TB/day)",
       ylabel="Required customer value (USD per input TB)",
       title="Break-even value from avoided transport and faster insight")
ax.legend(frameon=False);style(ax);save(fig,"05_native_value_threshold.png")
# 6: capacity scaling under launcher and allocation assumptions.
fig,ax=plt.subplots(figsize=(6.5,3.3))
mw=np.geomspace(.1,1000,100)
for payload,share,ls in [(100,.5,"-"),(100,.2,"--"),(50,.5,":")]:
    flights=np.ceil(mw*1000*design()["kg_per_kW"]/(payload*1000*share*.98))
    ax.loglog(mw,flights,ls,color="black",lw=1.6,
              label=f"{payload} t, {share:.0%} allocation")
ax.set(xlabel="Installed IT capacity (MW)",ylabel="Flights (integer lower bound)",
       title="Scale depends on mass and manifest allocation")
ax.legend(frameon=False);style(ax);save(fig,"06_launch_scale.png")
# 7: Pareto shows that user workload, not the location alone, determines value.
fig,ax=plt.subplots(figsize=(6.5,3.4))
pts=[("Terrestrial grid",18,.630,"o"),("LEO Earth-sourced",24,.696,"s"),
     ("LEO space-native",8,.696,"^")]
for name,x,y,marker in pts:
    ax.scatter([x],[y],facecolors="white",edgecolors="black",marker=marker,s=70)
    ax.annotate(name,(x,y),xytext=(5,5),textcoords="offset points",fontsize=8)
ax.set_xlim(0,35);ax.set_ylim(.58,.78)
ax.set(xlabel="Illustrative time to actionable result (min)",
       ylabel="Infrastructure LCOC (USD / 10¹⁸ useful FLOP)",
       title="Latency values are scenarios, not measured SLAs")
style(ax);save(fig,"07_workload_frontier.png")
# 8: decision-stage sensitivity: higher data rate lowers required value/TB,
# but mass and financing assumptions remain pivotal.
```

```
    fig,ax=plt.subplots(figsize=(6.5,3.4))
    nres=np.linspace(0,250,101)
    for capacity,ls in [(2,"-"),(10,"--"),(50,":")]:
        ax.plot(nres,[cost(500,80,x,capacity,100)["breakeven_value_usd_TB"]
                      for x in nres],ls,color="black",lw=1.6,
                label=f"Funded fleet {capacity} MW")
    ax.set(xlabel="Non-recurring engineering (MUSD)",
           ylabel="Value threshold at 100 TB/day/MW (USD/TB)",
           title="R&D amortization sets the first-market hurdle")
    ax.legend(frameon=False);style(ax);save(fig,"08_nre_scale.png")
    # 9: hardware feasibility sensitivity grounded in array specific power.
    fig,ax=plt.subplots(figsize=(6.5,3.3))
    sp=np.linspace(25,220,120)
    bus_W=250000*1.15
    fixed=design()["mass_kg"]-design()["pv_m2"]*3.
    ax.plot(sp,(fixed+bus_W/sp)/1000,color="black",lw=1.8)
    ax.axvline(30,color="0.5",lw=.8,ls="--")
    ax.axvline(bus_W/(design()["pv_m2"]*3),color="0.5",lw=.8,ls=":")
    ax.set(xlabel="Deployed solar-array specific power (W/kg)",
           ylabel="250-kW system mass (tonne)",
           title="Array mass is the central hardware sensitivity")
    style(ax);save(fig,"09_solar_specific_power.png")
    # 10: optical received power at 1550 nm; terminal losses held fixed.
    fig,ax=plt.subplots(figsize=(6.5,3.3))
    ranges=np.geomspace(30,3000,150)
    for d,ls in [(.10,"-"),(.20,"--")]:
        p=optical_received(ranges,D_r_m=d)
        ax.semilogx(ranges,10*np.log10(p/1e-3),ls,color="black",lw=1.6,
                    label=f"Receiver aperture {int(d*100)} cm")
    ax.set(xlabel="Link range (km)",ylabel="Received optical power (dBm)",
           title="1550-nm link power requires range and pointing margin")
    ax.legend(frameon=False);style(ax);save(fig,"10_fso_power.png")
if __name__=="__main__":
    assert eclipse(650,90)[1] == 0
    assert eclipse(650,0)[1] > 30
    assert design(beta_deg=60)["mass_kg"] > design(beta_deg=66)["mass_kg"]
    assert cost(1000)["lcoc_orbit_usd_EF"] > cost(200)["lcoc_orbit_usd_EF"]
    assert cost(200,240)["lcoc_ground_usd_EF"] > cost(200,80)["lcoc_ground_usd_EF"]
    assert cost(200,80,100,10,1000)["breakeven_value_usd_TB"] < cost(200,80,100,10,100)["breakeven_value_usd_TB"]
    assert optical_received(100)>optical_received(1000)
    plots()
    result={"design_250kW_beta66":design(),
            "design_250kW_beta60":design(beta_deg=60),
            "cost_200":cost(200),"cost_500":cost(500),"cost_1000":cost(1000),
            "break_even_1000TB_day":cost(200,80,100,10,1000)["breakeven_value_usd_TB"],
            "optical_100km_dBm":10*np.log10(optical_received(100)/1e-3),
            "optical_1000km_dBm":10*np.log10(optical_received(1000)/1e-3)}
    (HERE/"model_results.json").write_text(json.dumps(result,indent=2))
    print(json.dumps(result,indent=2))
```