RESEARCH / SR-2026-003 · FALSIFIABLE MODEL · V1.0 · 2026-09-27

Move the energy, not the data centre?

A falsifiable model of orbital energy for AI, and exactly what must change before space wins

By **Pavle Lazić** · ScalablyAI · Built with AI agents, reviewed by AI models (section 3)

Code and data: https://github.com/scalably-io/orbital-power-fabric · doi:10.5281/zenodo.22998545

1. Summary

Question. Should space first serve artificial intelligence as an energy layer, beaming power or reflecting sunlight to data centres on Earth, rather than as a compute layer with GPUs in orbit? The case for the energy layer is that GPUs depreciate in a few years while energy infrastructure could last decades, so one should launch only the long-lived part.

Method. We built an open model in which every input is a sourced claim with a verbatim quote, a numbered derivation with a test, or a labelled assumption. We fixed the metric and the verdict rules before computing any break-even. Independent reviewers with fresh context, who had not seen our reasoning, then tried to break the model; they were AI models, not human experts (section 3). Their reviews reversed several of our intermediate conclusions; the main reversals are in section 9 and all of them in `docs/NEGATIVE-FINDINGS.md`.

Answer, with today's hardware and prices.

1. **Microwave power beaming loses to GPUs in orbit on launched mass, from low orbit and from geostationary orbit, at every non-speculative tier.** The transmitter is the wall. The only measured transmitter specific power in our ledger, from a 2015 ground-test panel, is 28 W of radio-frequency power per kg (`R2-043`). Beaming from low orbit needs at least 147 to 194 W/kg even if the solar array weighed nothing. From geostationary orbit, with NASA's study-assumed conversion chain and 10 GW delivered, it needs at least 130 W/kg (4.7 times today); with today's arrays no transmitter is light enough.
2. **Orbital mirrors are lighter than GPUs in orbit only if two favourable conditions hold at once:** the areal density of a flown solar sail (200 g/m², NASA's ACS3, `R1-151`) and a purpose-built receiving field that turns every watt of mirror light into electricity at full-sun efficiency. Aimed at a real tracking solar farm, which converts about 5.9% of the light landing on its site (`R5-077`), sail-density mirrors are 105 kg per continuous kW over 15 years against 40 to 57, and only a bare membrane with no deployer or bus is lighter. Reflect Orbital's filed Earendil-1 system, at its own filed life of about 1 year, is 45.6 kg per kW even onto the ideal receiver with no clouds, inside the range of GPUs in orbit. Any mirror fleet must also hold 0.12 to 0.24 km² of mirror over each receiving site at once before a ground inverter at a site in darkness produces anything.
3. **At today's launch price, the launch alone of GPUs in orbit, of microwave beaming and of sail-class mirrors costs more than ground solar with storage sized for 97 to 99.4% of hours** ($120 to 192 per MWh of AI load); only speculative hardware passes on launch cost, and the last hours on the ground need backup we did not price. Of the options we compared, the ground is the best supported today: solar and storage built on site, with optics used for interconnect.
4. **Photonic matrix engines that stream weights cannot fix large-model inference, and our resident-weight photonic design fails with measured parts.** An engine that streams weights from high-bandwidth memory is capped at 1.9 to 6.4 times a B300 GPU in memory-bound decode even if its computation were free, depending on the memory's energy per bit. A resident-weight engine with direct photon-to-charge read-out beats a B300 by 10 times only in an idealised case that cannot be built as specified. With measured parts and precision required on real outputs, its whole system may use at most 2.65 times its optical-core energy at NVFP4 pairing, leaving the control, calibration, clocking and I/O electronics less than twice the core, and at FP8 pairing the core alone misses the gate. The architecture is not new; the negative result and its brackets are. Other photonic architectures are not ruled out by this analysis.

What would change the verdict. Table 9 lists 12 quantities, each with today's sourced value and the value at which a conclusion flips. Every measured or derived number in the prose is filled in by code from the ledger or the model, and design choices such as layer size are set in the build script: `./scripts/reproduce.sh` and `python3 paper/build_paper.py` rebuild it.

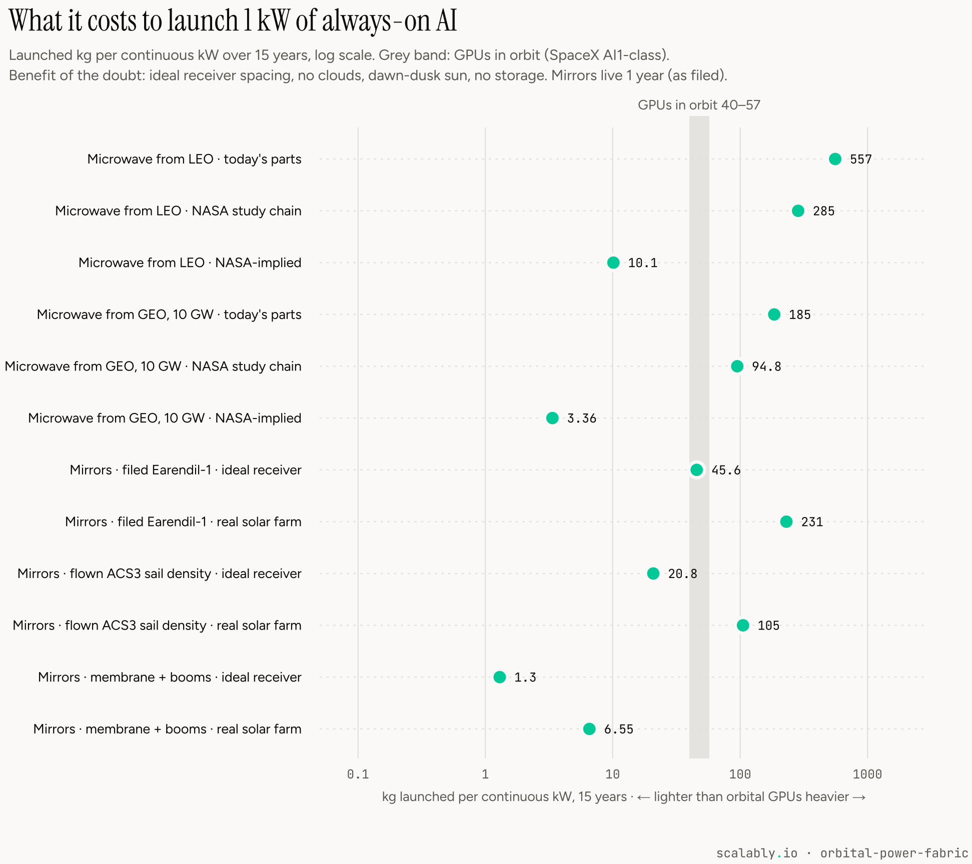

Figure 1. Launched kg per continuous kW of AI load over 15 years. Grey band: GPUs in orbit. Energy options receive ideal receiver spacing, no clouds, dawn-dusk sunlight and no storage. Mirrors are shown twice: onto an ideal receiver fully covered with PV, and onto a real tracking solar farm.

2. Question and preregistered metric

The unit of comparison is one kilowatt of AI accelerator load (IT power, GB300 NVL72 class) running continuously for L = 10, 15 or 20 years. Four branches supply it:

- **A, GPUs in orbit.** The accelerator flies and is relaunched at the shorter of its refresh cadence and the vehicle life.
- **B, microwave beaming.** Orbital solar arrays feed a microwave transmitter; a ground rectenna feeds GPUs on Earth. Low orbit (B-LEO, 625 km, 30° minimum elevation) and geostationary orbit (B-GEO).
- **C, orbital mirrors.** Mirrors reflect sunlight onto ground solar farms that feed GPUs on Earth.
- **T, the ground.** Solar and storage on Earth. The original handoff compared orbital energy only with orbital compute; any orbital option must also beat the ground.

Primary metric: lifecycle launched mass per continuous kW of AI load, counting every replacement as a whole spacecraft. Because the orbital-GPU baseline is a range, "lighter" means below its lower end, "heavier" means above its upper end, and anything between is inside the range. The verdict rules V1 to V4 and the metric were committed in `docs/PREREGISTRATION.md` before any break-even was computed; later changes are dated amendments.

Tiers, and where they depart from our own rule. The preregistration defines *current* as flown, qualified or buyable, *next* as best demonstrated and *speculative* as projected, and each B or C tier is compared with the same tier of A. The evidence does not fill every cell that way, so we state what each cell really is:

- A current is a vendor statement for a vehicle that has not flown (R1-005); A has no demonstrated next tier, so A next equals A current.
- B current combines flight and catalogue solar wings (median 88.2 W/kg, R2-093), a ground-tested transmitter (R2-043) and a conversion chain with two measured factors and five study assumptions. B next is a scenario: a catalogue wing and NASA's study-assumed chain, not demonstrated hardware.
- C current is a filing (Earendil-1, R1-112, R1-113). C next uses the density of a flown sail spacecraft, which is flown but is not a mirror. C speculative counts membrane and booms only.

Choices that favour B and C. No bus or structure mass for B; receivers spaced so that their visibility caps never overlap (an upper bound on coverage, D03); dawn-dusk sunlight; energy delivered to the receiver network counted as if it served one continuous load (D22); launched mass to geostationary orbit counted like mass to low orbit; no eclipse storage, rectenna exclusion zone or land for geostationary beaming; for mirrors, a receiving field fully covered with PV at its full-sun efficiency of 30% (we also show a real solar farm), zenith atmosphere and full capture. Where B or C lose with these choices, they lose without them.

3. Method: an auditable evidence ledger

The ledger `evidence/claims.csv` holds 1,214 claims (373 of them measurements) plus 70 sources we could not access, kept so that nobody repeats the dead end. Each row carries a source URL, a locator (page, table or section), a verbatim quote of at most forty words, a retrieval date and a status: MEASURED, VENDOR-STATED, FILING, STUDY-ASSUMED, DERIVED or PROJECTED. We speak of them differently: a vendor statement is not a measurement, and a study assumption is not a demonstration.

Every formula in the model has a numbered derivation (D01 to D22, appendix A) and a test that pins it; guard tests were mutation-checked by deleting the guarded code and confirming the test fails. When a quantity has no source, we either label the assumption (appendix B) or report the value at which the conclusion flips. Some model constants are code literals that carry their claim IDs in comments rather than being read from the ledger at run time.

How this was made. Scalably built this work with AI research agents (Anthropic's Claude) that searched primary sources and wrote ledger rows into separate files. Rows were merged only after an automated validator passed, with quotes spot-checked against the downloaded source text for each research batch, not every row (EXPERIMENTS E006...E029, E035). For a derived value the quote gives an input and the row's notes give the derivation. The model and this paper were then reviewed by AI models with fresh context (Anthropic's Claude Opus and OpenAI's Codex), not by human domain experts. We publish it so that people who build this hardware can check it.

To challenge a number, open an issue with the claim ID and a better source. The verdict, the tables and the prose regenerate from the ledger.

4. Results

The summary table compares the three energy branches with GPUs in orbit, tier by tier, over 15 years. Every B and C value uses the favourable choices listed in section 2.

Table 1. Launched kg per continuous kW of AI load over 15 years, with no clouds.

TIER	GPUS IN ORBIT (A)	MICROWAVE, LEO, IDEAL COVERAGE	MICROWAVE, GEO, 10 GW	MIRRORS, FULLY COVERED RECEIVER	MIRRORS, ONTO A REAL SOLAR FARM
current	40 to 57	557.1 (heavier)	185.1 (heavier)	45.6 (inside the range)	230.7 (heavier)
next	40 to 57	285.3 (heavier)	94.8 (heavier)	20.8 (lighter)	105.3 (heavier)
speculative	30 to 43	10.1 (lighter)	3.4 (lighter)	1.3 (lighter)	6.6 (lighter)

4.1 GPUs in orbit

SpaceX states 75 kW per ton for its Starmind AI1 vehicle (R1-005, vendor-stated). The page does not say whether the numerator is average or peak power, which gives 13.3 to 19.0 kg per continuous kW per launch; if the ton is a short ton, the lower end is 12.1. Starcloud's Starcloud-3 has been described as 2,000 kg in a filing (R1-073) and, in press accounts, as a two-tonne, 100 kW or a three-tonne, 200 kW spacecraft with the power basis not given (R1-076); each description gives 15 to 20 kg per kW, consistent with that range. Satellite filings give a 5-year life (R1-023), so 15 years takes 3 launches: 40 to 57 kg per kW. A filed target of 100 kW per ton (R1-020, projected) sets the speculative tier at 30 to 43.

4.2 Microwave beaming from low orbit

With perfect receiver coverage, beaming costs 557 kg per kW over 15 years with today's parts and 285 with a catalogue array and NASA's study-assumed chain, against 40 to 57 for GPUs in orbit. Only NASA-implied whole-system hardware (10.1) is lighter. With receivers only on land, the current tier rises to 1,914.

The transmitter decides it. At the measured 28 W/kg (JAXA's 2015 ground-test GaN panel, R2-043; no space-qualified value was found, R2-042), the transmitter alone costs 92 kg per kW per launch, 277 over 15 years, even with weightless solar arrays. The break-even transmitter is 147 to 194 W/kg (6.9 times today with the chain of today's parts). With today's arrays the arrays alone weigh 278 kg per kW over 15 years, so no transmitter, however light, rescues current-tier beaming.

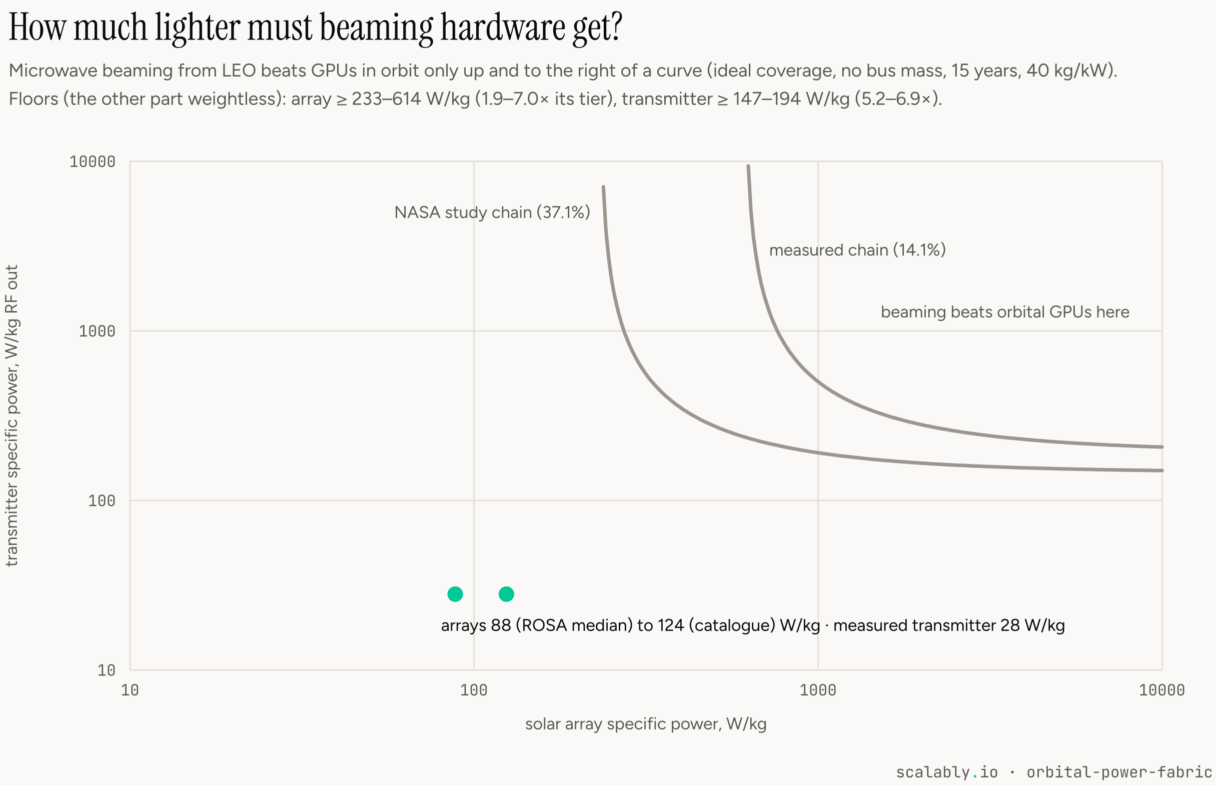

Figure 2. Microwave beaming from low orbit beats GPUs in orbit only above and to the right of a curve. The measured transmitter (28 W/kg) and today's arrays (88.2 to 124.1 W/kg) sit far below both curves.

The solar array is the second wall. NASA's own space-solar study implies 909 to 938 W/kg for the whole system (R2-018), about 11 times the 84.6 W/kg of the flown iROSA wing (R2-093). NASA concludes that its baseline is not cost-competitive even if access to space were free (R2-028).

A dawn-dusk orbit keeps the satellites in sunlight, but it also keeps them over the day-night line. An equatorial receiver lies within reach of a dawn-dusk orbital plane for only about 8.8% of the day, in two windows around dawn and dusk (D22), and any single satellite passes overhead for minutes. The receiver network as a whole can take energy continuously, but a continuous load at any one site would need storage or a grid linking sites, which we do not count.

4.3 Microwave beaming from geostationary orbit

Geostationary orbit removes the visibility problem: one satellite sees one rectenna almost all the time. It does not remove the transmitter. A panel of given specific power cannot radiate more than that power times its own mass (D20), so a one-kilometre JAXA-class aperture, which already needs a 13.1 km² rectenna, can radiate at most 0.98 GW of radio-frequency power. Delivering 10 GW needs 17.0 to 22.5 GW radiated, and the transmitter must grow with it.

Table 2. Microwave beaming from geostationary orbit, transmitter sized by the power it radiates (D20).

TIER	DELIVERED	RF TO RADIATE	KG PER KW PER LAUNCH	SERVICE LIFE NEEDED TO MATCH A OVER 15 Y (WHOLE LAUNCHES)	AT THE 15-YEAR DESIGN LIFE
current	1 GW	2.2 GW	185.1	none within the lifecycle	loses
current	10 GW	22.5 GW	185.1	none within the lifecycle	loses
next	1 GW	1.7 GW	94.8	none within the lifecycle	loses
next	10 GW	17.0 GW	94.8	none within the lifecycle	loses
speculative	1 GW	folded into system W/kg	3.4	1.9 y	wins
speculative	10 GW	folded into system W/kg	3.4	1.9 y	wins

GEO satellites are designed for a median 15 years (R12-022); no power satellite has flown. Counting whole spacecraft, no service life lets current- or next-tier GEO beaming match GPUs in orbit within 15 years: a single launch (94.8 to 185.1 kg per kW) already exceeds their 15-year total. Only the transmitter can change this (P1). The model is generous to GEO in several ways. It counts mass lifted to geostationary orbit like mass lifted to low orbit, although geostationary orbit takes far more launch energy. It delivers power through the eclipses around the equinoxes with no storage (D04). It keeps the one-kilometre receiving geometry, although a transmitter sized for 10 GW at the measured panel density would be about 4.2 to 4.8 km across. And it charges nothing for the rectenna site, although 10 GW on 13.1 km² averages about 1,237 to 1,636 W/m² of microwaves against a general-public exposure limit of 10 W/m² (R2-108), so the site would need an exclusion zone.

4.4 Orbital mirrors

A mirror in orbit reflects an image of the Sun, not a beam. At 625 km the image covers about 26.5 km² at zenith (D05), so a single mirror adds almost nothing to the irradiance on the ground. Utility inverters start producing at 2.6 W/m² with an ideal module and about 5 W/m² at half efficiency in low light (R5-069). A receiving site in darkness therefore needs 0.12 to 0.24 km² of mirror overhead at once, about 384 to 738 Earendil-1-size mirrors, before any power flows. The mass figures below assume a fleet at that scale. Like beaming, a dawn-dusk mirror serves each site near its own dawn and dusk (D22), which suits extending existing solar farms but not a continuous load at one site.

The only sourced life for an orbital reflector is Earendil-1's filing: about one year of operations and a passive drag deorbit in about five months (R1-115, R1-116). We apply that 1-year life to every tier, because lighter reflectors with more area per kilogram decay faster, and report the life each tier needs.

Table 3. Orbital mirrors: kg per continuous kW over 15 years with weather, and what each tier needs to match GPUs in orbit.

TIER	COVERAGE	CLOUD-FREE 0.66	NIGHT 0.62	GLOBAL 0.33	SAME-TIER A	CLOUD-FREE SHARE NEEDED	MIRROR LIFE NEEDED (WHOLE LAUNCHES)
current	ideal	69.1	73.6	138.2	40 to 57	cannot win	≥ 1.88 y
current	land-only	237.5	252.8	474.9	40 to 57	cannot win	≥ 7.50 y
next	ideal	31.5	33.6	63.1	40 to 57	≥ 0.52	≥ 0.79 y
next	land-only	108.4	115.4	216.7	40 to 57	cannot win	≥ 3.00 y
speculative	ideal	2.0	2.1	3.9	30 to 43	≥ 0.04	≥ 0.07 y
speculative	land-only	6.7	7.2	13.5	30 to 43	≥ 0.15	≥ 0.23 y

The filed Earendil-1 system (142 kg for 324 m², R1-112, R1-113) is 45.6 kg per kW onto the ideal receiver with no clouds, inside the range of GPUs in orbit, and 69.1, above the whole range, with the average cloud cover of solar regions; counting whole launches it would need a life of at least 1.88 years. At the density of the flown ACS3 sail spacecraft (200 g/m², R1-151) mirrors are lighter with ideal coverage and at least 0.52 of the site cloud-free, but not with receivers only on land. ACS3 is a sail, not a mirror: its optical quality as a sunlight relay is undemonstrated. All of these figures assume a receiving field fully covered with PV. A real tracking solar farm holds 59.3 MW of panels per km² at full sun (R5-077), so it turns only about 5.9% of the light landing on its site into electricity; aimed at such a farm, even with no clouds, the filed system is 231 kg per kW, the sail-density mirror 105 and the bare membrane 6.6. Mirrors that beat GPUs in orbit need a dedicated receiver built for them, whose area and cost this paper does not model. The weather factor is one minus the mean cloud amount, the expected share of a site in direct sun: 0.66 over all hours and 0.62 at night across 12 utility-PV regions (R12-016), and about 0.33 worldwide (R10-052). It ignores light passing through thin cloud and the loss when partial cover pushes part of a site below inverter start. Light pollution and debris risk from fleets of this size are outside the metric.

4.5 Cost against the ground

The terrestrial comparator is solar with batteries sized for flat output during 97% and 99.4% of hours in Las Vegas: $104 and $167 per MWh delivered (R5-027, R5-028, modelled with 2024 global prices). Per MWh of AI load that is $120 to 192. The last 2.4% of the year's energy costs about $2,700 per MWh more than the 97% case (R5-028), so full firmness is expensive on the ground too.

Table 4. Launch cost alone per MWh of AI load (hardware excluded), and the launch price at which each option matches the ground.

OPTION	KG PER KW PER LAUNCH	LIFE	LAUNCH $ PER IT-MWH AT $3,600/KG	AT $200/KG	BREAK-EVEN $/KG VS $120 TO 192
A current (13.3 kg/kW)	13.33	5 y	1,095.9	60.9	393 to 631
B-LEO current	185.69	5 y	15,262.6	847.9	28 to 45
B-LEO speculative	3.37	5 y	277.0	15.4	1,554 to 2,496
B-GEO speculative, 10 GW	3.36	15 y	92.1	5.1	4,677 to 7,510
C next (flown ACS3 density)	1.39	1 y	570.2	31.7	755 to 1,213
C speculative (membrane + booms)	0.09	1 y	35.5	2.0	12,130 to 19,478

At today's $3,600/kg (R1-057, a study assumption), launch alone for GPUs in orbit costs $1,096 per MWh of AI load, and the break-even launch price is $393 to 631/kg with hardware free. At $200/kg, projected for the mid-2030s (R1-053), it falls to $61. If launch were bought at $60/kg, a Starship cost scenario with rapid reuse in Google's Suncatcher study (R1-056, a cost to the launcher, not a customer price), it would be $18.3. At today's price two of the speculative rows pass: NASA-implied geostationary hardware and mirrors of bare membrane and booms.

4.6 Sensitivity

Table 5 draws the uncertain inputs across stated ranges and reports the share of draws in which each option is lighter than GPUs in orbit. It is not a probability of success: the ranges span tiers and regions, their shapes are assumptions, and no transmitter above 28 W/kg or array above 124.1 W/kg is sourced, so their upper ends are shown at several values.

Table 5. Share of draws lighter than GPUs in orbit over 15 years.

RANGES DRAWN	MICROWAVE, LEO	MICROWAVE, GEO 10 GW	MIRRORS
PV 88→124 W/kg (ROSA median→catalogue), transmitter 28→280 (≤10×), log-uniform	0.000	0.126	0.719
same, uniform	0.000	0.227	0.385
PV 88→124, transmitter 28→840 (≤30×), log-uniform	0.000	0.246	0.719
PV 88→500 (beyond any ROSA wing), transmitter 28→840, log-uniform	0.012	0.578	0.719
PV 88→500, transmitter 28→840, uniform	0.049	0.869	0.385

Low-orbit beaming is lighter in at most 0.049 of draws under the five priors shown. Geostationary beaming moves from 0.13 to 0.87 as the allowed transmitter improves: its verdict is a statement about transmitter specific power. Mirrors range from 0.39 to 0.72, before the site scale of section 4.4 is counted.

5. What the evidence supports today

Energy for AI, on the ground. Of the options we modelled, near-firm solar with storage on the ground is the cheapest per MWh (R5-027, R5-028). We did not compare it with grid power, gas, wind or nuclear, so we do not claim it is the best ground option. The binding constraint for new capacity is often time: the median utility project built in 2025 took 61 months from interconnection request to operation (R5-042). Off-grid generation on site avoids that queue; a modelled off-grid microgrid meeting 90% of a data centre's load from solar costs $109 per MWh (R5-029) and needs about 102 km² of land per GW of continuous load (R5-031). Coupling on-site PV directly to an 800 VDC rack bus is modelled to deliver 6.8% more energy than AC coupling and 12.4% more than remote PV over the grid (R6-122, excluding battery round-trip losses); that rack architecture is expected in production in 2027 (R6-016).

Space is not yet the energy layer. Beaming is transmitter-limited in both orbits, mirrors need sail-class hardware in fleets of hundreds per site, and at today's launch price only speculative hardware passes on launch cost. GPUs in orbit approach ground prices only when launch falls to a few hundred dollars per kilogram (Table 4).

Photonics' proven energy win today is interconnect. Photonic arithmetic exists (R3-042) but is not yet more efficient than GPUs at system level (section 6). Co-packaged optics move data at about 5 pJ/bit against 15 to 20 for pluggable transceivers (R9-030), at rack-to-cluster reach.

6. Photonic computing: where light wins and where it cannot

The memory wall. Generating one token with a large language model reads every weight once per batch step (D16). High-bandwidth memory costs 11.68 pJ/bit on a measured GH200, including the GPU's control logic (R3-078), and 18.0 pJ/bit through the GPU's caches (R9-005). A photonic engine that streams weights pays the same memory access, because light does not store weights. We bound its best case from above: a B300 draws at most 1,400 W (R3-067) while reading 8 TB/s of memory (R1-100, R1-094), so in memory-bound decode at batch 64 its whole board spends at most 2.73 pJ per multiply-accumulate with 8-bit weights. Even with a free compute core, the photonic engine is capped at 1.9 to 6.4 times the B300, depending on the memory energy assumed: 1.87 times with GH200 HBM access incl. control, measured (R3-078); 3.81 times with HBM3E incl. IO, study value (R3-082); 6.36 times with HBM3E, SK hynix via secondary citation (R3-081).

Table 6. Streamed-weight photonic engine (time-integrating read-out) against a B300 at its roofline, with the GH200 memory energy. Photonic support electronics are not included, so every ratio favours photonics.

WORKLOAD	PRECISION PAIR	PHOTONIC CORE + MEMORY, FJ/MAC	B300 BOARD, FJ/MAC	RATIO	CAP WITH A FREE CORE
decode, batch 64	FP4 (4-bit weights, 5-bit read-out)	826	1,367	1.65	1.87
prefill, 2,048-token reuse	FP4 (4-bit weights, 5-bit read-out)	77	187	2.41	8.18
decode, batch 64	FP8 (8-bit weights, 8-bit read-out)	1,594	2,734	1.72	1.87
prefill, 2,048-token reuse	FP8 (8-bit weights, 8-bit read-out)	120	560	4.69	12.27

Support overhead. The only measured full-stack photonic AI processor runs at 0.81 TOPS/W at system level, and its optical core draws 0.31% of the system power (R3-043, R3-045). That core excludes the laser and the converters. Counted the way we count cores below, with its laser at the 20% wall-plug efficiency we assume for H1' (R8-072), its whole system uses about 10 times its core (43 times if the laser's optical output is counted directly). It is a different architecture at a different precision, so it illustrates the size of support overhead in a real system; it does not bound H1'.

Sunlight cannot replace the laser. By the brightness theorem (D18) one spatial mode collects at most 64 nW of sunlight per THz of bandwidth at the telecom wavelength, whatever the collector. That supports far too few precise reads per second for a gigahertz engine; sunlight can feed only slow or massively multimode optics.

Precision. A b-bit analog read-out in this paper means fixed-point resolution with shot noise equal to one least-significant bit (D11, D21). That is closer to an INT format than to NVFP4 or FP8, which carry per-block exponents. We pair 5-bit read-out with NVFP4 and 8-bit read-out with FP8; whether either pair gives equal model accuracy is open problem P10, so "matched precision" below means these chosen pairings.

7. A testable hypothesis that fails: resident-weight photonic inference

H1'. Weights stay in place in non-volatile photonic cells (no hold power, no weight fetch). Each layer is square. Detectors integrate photocurrent directly onto a small capacitive node, with no amplifier (D12): the photon-to-charge conversion is the read-out. Detector-driven modulators pass activations between layers without converters, and the only data converters sit at the input of the first layer and the output of the last.

Table 7. The idealised case: H1' optical-core energy per multiply-accumulate, 8 layers, precision measured against full scale, a 1 fF node and 20 dB loss, paired with B300 NVFP4 (187 fJ/MAC) and FP8 (560 fJ/MAC) at chip level (R3-070).

LAYER	PAIRED WITH	READ-OUT	OPTICAL CORE, FJ/MAC	B300 ÷ CORE	SYSTEM ÷ CORE ALLOWED FOR 10×
256 × 256	NVFP4, 5-bit read-out	conventional receiver	42.3	4.42	0.442
256 × 256	NVFP4, 5-bit read-out	photon→charge node	3.63	51.5	5.15
1024 × 1024	NVFP4, 5-bit read-out	conventional receiver	10.6	17.7	1.77
1024 × 1024	NVFP4, 5-bit read-out	photon→charge node	0.907	206	20.6
256 × 256	FP8, 8-bit read-out	conventional receiver	319	1.76	0.176
256 × 256	FP8, 8-bit read-out	photon→charge node	21.7	25.8	2.58
1024 × 1024	FP8, 8-bit read-out	conventional receiver	79.7	7.02	0.702
1024 × 1024	FP8, 8-bit read-out	photon→charge node	5.42	103	10.3

The idealised case passes on paper but cannot be built as specified. With photon-to-charge read-out, a 0.91 fJ/MAC core at NVFP4 pairing lets the whole system use up to 20.6 times the core's energy before the 10-fold gain over B300 is lost, and 10.3 times at FP8. With a conventional receiver the allowance is at most 1.8 times the core at any size, which leaves almost nothing for the electronics around it. But a square single-wavelength layer shares its laser light across all outputs only as a lossless unitary (D13): meshes built from measured non-volatile elements lose tens to hundreds of dB (R8-126), and summing many waveguides incoherently onto one single-mode detector collects at most one of them (D13). A real layer would need one wavelength per input or accept those losses.

Measured parts do not pass. Table 8 changes one assumption at a time to a sourced value, cumulatively. At NVFP4 pairing three steps each take a large share: the measured-parts detector node (system allowance 20.6 to 12.2 times the core), the O-E-O device's measured input energy of 6.4 fJ per read (R7-046), which sets a floor on the light every output needs (to 5.1 times), and precision over real outputs (to 2.8 times). At FP8 pairing precision dominates. Measured against the largest possible sum of 1,024 terms, as in the idealised case, a typical output spans only a small part of full scale: about 0.17 of one least-significant bit at 5 bits, so it quantises to zero, and about 1.3 at 8 bits. Asking for the same bits over the real output distribution (D21, checked by a photon-counting simulation) makes the photons per multiply-accumulate independent of layer size, which at 8 bits multiplies the light needed many times over.

Table 8. From the idealised case to measured parts, 1,024 × 1,024, cumulative. Each row is one sourced change.

CUMULATIVE CHANGE (ONE SOURCED STEP EACH)	NVFP4 PAIRING: CORE FJ/MAC	SYSTEM ÷ CORE ALLOWED FOR 10×	FP8 PAIRING: CORE FJ/MAC	SYSTEM ÷ CORE ALLOWED FOR 10×
idealised H1' (A4): full-scale precision, 1 fF / QE 1 / 1,550 nm node, 20 dB	0.907	20.6	5.42	10.3
detector node from measured parts (R8-128)	1.53	12.2	6.65	8.42
O-E-O at its measured operating point (R7-046)	3.64	5.14	6.65	8.42
precision over the output distribution (D21)	6.77	2.76	97.3	0.576
buffer before the ADC at the network output (R3-062)	7.06	2.65	97.6	0.574
and at 26 dB optical loss instead of 20 dB	25.7	0.728	384	0.146
and at 30 dB optical loss instead of 20 dB	63.3	0.295	961	0.0582

Verdict. With today's measured components and our chosen precision pairings, resident-weight photonic inference does not reach a 10-fold energy gain over a B300. At NVFP4 pairing the whole system may use 2.65 times the core's energy, so control, calibration, clocking and I/O get less than twice the core; the only full system measured, a different design, uses about 10 times its core counted the same way. At FP8 pairing its core alone is 98 fJ/MAC, only 5.7 times below a B300 before any support electronics; with 7-bit or 6-bit read-out instead of 8, the system allowance is 2.2 or 7.9 times the core. More than 24.6 dB of optical loss leaves no allowance at all at NVFP4: the core alone would miss the gate. Not modelled, and each would lower the budget further: power to lock ring resonances (R8-103), static power at a stated clock rate, and splitting light to both arms of a balanced pair. The 10-fold gate is the gate we fixed in advance for claiming a large gain (docs/SYNTHESIS.md); the preregistered V4 test, full-system energy below a GB300-class GPU, remains open at both pairings because no full-system measurement exists.

What would overturn it. A measured resident-weight layer beating these budgets at the wall plug (P7, P8, P12), or a model evaluation showing that precision defined against full scale is enough for real networks (P10). The favourable definition is supported today only by simulations of small networks, at under one photon per multiply-accumulate (R12-030). Retention of non-volatile photonic weights is reported as test durations from 1,000 seconds (R8-112) to 4 days (R8-099); these are lower bounds, and some cells are extrapolated to years, and the best density we found is an upper bound of 3,183 weights per mm² for ring footprints against 4.76 million for SRAM (R8-104, R8-084). At that density the layer sizes we tested hold 0.5 to 8.4 million weights on 1.6 to 26 cm², so even the idealised case concerns small fixed models, and the larger case spans several lithographic reticles.

Prior art. The architecture is not new. MIT's US11546077B2 (priority 2021) describes static weights that attenuate photocurrents without consuming power and charge summation, with photodetectors driving modulators for linear summation while converters remain between layers; US10268232B2 (2016) and the University of California's US11144821B2 (2017) cover other combinations, the latter in spiking form (R11-001…019). Built hardware stops at pairs or triples of these elements. What no one has published is the measurement that decides the question: the full-system wall-plug support overhead of such an engine.

7.1 The open test chip, OPF-MZI4

A 4×4 Clements mesh of Mach-Zehnder interferometers with titanium-nitride heaters on the CORNERSTONE 220 nm silicon-on-insulator process, assembled from unmodified cells of the MIT-licensed cspdk design kit. It measures grating-coupler loss, heater power per π phase shift, thermal crosstalk, unitary programming and the heater power needed to hold a matrix. It realises unitary matrices only and does not test H1'. Because a power-only measurement cannot certify the complex matrix, the test plan uses an intensity-only fidelity. The layout passes our connectivity check, but the foundry's design-rule check has not been run, because the rules are released only at sign-up. We do not fabricate it ourselves; `chip/TAPEOUT-GUIDE.md` gives routes, prices with claim IDs, the test setup and the measurement order for any lab that wants to.

8. Open problems for teams that can build

Each row is a quantity that changes a conclusion of this paper if it moves past the derived target. Rerun `python3 scripts/verdict.py` with the new value to see where the verdict flips. Identifiers such as E021 refer to entries in the repository's experiment log, `EXPERIMENTS.md`.

Table 9. Open problems, generated by `scripts/open_problems.py`.

#	QUANTITY	TODAY (SOURCE)	TARGET (DERIVED)	GAP	EVIDENCE
P1	Space microwave transmitter specific power (W/kg RF out)	28 (R2-043, JAXA GaN, measured)	LEO: ≥ 147 to 194 even with a weightless array; GEO next-tier 10 GW, 15-y life: ≥ 130	4.7 to 6.9×	E021, E028, fig breakeven_bleo
P2	Space solar array specific power (W/kg, array level)	88 median ROSA-family wing (flight and catalogue); 124 OPAL catalogue wing; large flown class 71 to 98 (R2-093)	≥ 233 to 614 even with a weightless transmitter	1.9 to 7.0× (each vs its own tier's array)	E021
P3	GEO power-satellite service life (years)	15 y median DESIGN life of GEO comsats (R12-022); no power satellite flown	none: counting whole spacecraft, one current/next-tier launch (1 to 10 GW, transmitter sized by power, D20) already exceeds orbital GPUs' 15-y total, so no service life suffices; only the transmitter (P1) can change this	n/a	E021, E027, E028, E033
P4	Orbital-mirror scale per receiving site	Earendil-1: 18×18 m, 142 kg (R1-112/113)	≥ 0.12 to 0.24 km² of mirror over each spot at once (≈384 to 738 Earendil-1s) to clear the 2.6 to 5 W/m² inverter start (ideal / 50% low-light module)	n/a	E021, E028, R5-069
P5	Clear-sky fraction at mirror-fed sites	0.66 all-hours mean, range 0.44 to 0.84, 12 PV regions (R12-016)	filed Earendil-1 system at its filed 1-y life: cannot win even with ideal coverage (needs > 1.0 clear sky); flown-sail density (ACS3, 200 g/m²) onto a fully covered receiver: ≥ 0.52 ideal coverage, cannot win land-only	n/a	E023, E027, E028
P6	Launch price for orbital GPUs to match 97 to 99.4%-firm ground solar (launch cost only, per MWh of AI load)	\$3,600/kg (R1-057, STUDY-ASSUMED)	≤ \$393 to 631/kg (hardware cost excluded)	5.7 to 9.2× cheaper	E021 V3
P7	Photonic full-system support overhead (system ÷ optical core)	318× (R3-045, Lightmatter Nature 2025)	H1' at 1024², to beat B300 by 10×: system ÷ core ≤ 2.65× with measured parts and precision over real outputs (NVFP4 pairing; FP8 fails at any overhead); ≤ 10.3 to 20.6× only in the idealised case (full-scale precision, unsourced 1 fF node); 318× is an upper bound (its core excludes laser and ADCs)	≥ 15× lower (idealised); unattainable at FP8	E032 ladder (supersedes E028, E025, E018)
P8	Measured wall-plug energy of an integrating + receiverless photonic system	none published (R7)	a measured J/MAC including drivers, control and calibration	n/a	E022, R11
P9	Non-volatile photonic weight density (weights/mm²)	5.8 to 3,183 (R8-104 upper bound; N2 SRAM 4.76 M, R8-084)	≥ ~10^5/mm² to hold a 1 B-weight model on ~100 cm²; today H1' holds 0.5 to 8.4 M weights on 1.6 to 26 cm² (256² to 1024² × 8 layers)	≥ 30×	E028, R8-124
P10	Precision equivalence of analog fixed-point read-out	no ledger claim links b-bit analog outputs to NVFP4/FP8 model accuracy	a model run end-to-end on 5-bit fixed-point analog outputs (with measured noise) matching its NVFP4 or FP8 accuracy; note D11's σ = 1 LSB is ≈ b − 1.8 effective bits, and ENOB = b needs 12× the photons	n/a	D11 (E028), fresh-context AI reviews
P11	Orbital reflector service life (years)	~1 y operations, drag deorbit ~5 months (Earendil-1 filing, R1-115/116)	≥ 1.88 y (filed Earendil-1 density) / ≥ 0.79 y (flown-sail density) to match orbital GPUs, whole launches, all-hours cloud-free share, ideal coverage	1.9×	E028, E033

#	QUANTITY	TODAY (SOURCE)	TARGET (DERIVED)	GAP	EVIDENCE
P12	Optical loss of a 1024 × 1024 resident-weight layer, laser to detector (dB)	20 dB assumed (R3-037, PROJECTED); no measured layer at this size	≤ 24.6 dB for H1' (NVFP4 pairing, measured parts) for the whole system to stay within the 10× gate at all (system ÷ core ≥ 1); +6 dB over 20 dB leaves 0.73×	n/a	E032, E033

9. Negative results, including our own errors

Negative results are deliverables. The full list, with evidence classes, is `docs/NEGATIVE-FINDINGS.md`. The most consequential:

- **The original handoff omitted the ground as a comparator.** Every orbital option must beat solar and storage on Earth, not only orbital compute.
- **Space-to-ground power beaming has so far been shown as detection, not usable power.** Caltech's MAPLE beam was detected on the roof of its campus (R2-052), up to 20 dB above the ground station's noise floor (R2-057); its best in-orbit microwave transfer measured 3.37 to 3.66% DC to DC over about 30 cm inside one spacecraft (R2-053). Caltech's own estimate of today's sunlight-to-grid efficiency from geostationary orbit is 2.77% (R2-065), and NASA's baseline costs $0.61 to 1.59 per kWh (R2-023).
- **Our geostationary "win" was a modelling error.** An earlier version of our verdict kept the transmitter at one kilometre for any power and showed next-tier beaming winning at 10 GW. Sized by the power it must radiate, a single launch already outweighs GPUs in orbit over 15 years (reviewer finding, E028).
- **We gave Reflect's mirror five years it does not have.** We had borrowed the life of data-centre satellites. The mirror's own filing states about one year, and at that life it is not lighter than GPUs in orbit (E028).
- **We counted GPU memory energy twice.** B300's energy per operation is its total board power divided by peak throughput, which already includes memory; we had added memory energy on top. The corrected bound is in section 6 (E028).
- **We used fractional spacecraft for break-even lives.** Counting whole launches, the filed mirror needs 1.88 years, not less, and no life rescues geostationary beaming within the lifecycle (E033).
- **We let mirror light convert at full efficiency on every lit square metre.** A real solar farm covers only part of its land with panels; aimed at one, no mirror we could source is lighter than GPUs in orbit except a bare membrane (E034).
- **Our photonic hypothesis failed three times.** Time-integrating read-out needs a new weight every cycle, so it cannot combine with resident weights (E025). Our next design shared laser light across more inputs than outputs, which power conservation forbids in a single-wavelength network (E028). The design that survived both, H1', fails once built from measured parts (E031, E032; section 7).

10. Limitations, and the strongest case against our conclusions

Limitations. Cost comparisons count launch only; spacecraft hardware would widen every gap. The orbital-GPU baseline rests on vendor statements for vehicles that have not flown. The 5-year life of beaming satellites is borrowed from data-centre filings. Receiver coverage is an upper bound, the land-only case is an approximation, and continuous per-site supply is not modelled (D22). Mirror optics are assumed perfect. The photonic results are optical-core budgets at an assumed optical loss, precision equivalence is unshown, and resonance-locking and static power are not modelled. Laser power beaming is not modelled: the only in-orbit laser demonstration in our ledger delivered about 1.5 W at about 11% over 1.45 m inside one payload (R2-072), which is not yet a space-to-ground data point. No human domain expert has reviewed this work yet.

The strongest case against us.

- *Launch will become cheap.* That changes the cost comparison, not the mass ranking. Current-tier microwave beaming matches the ground on launch cost only below $28 to 45/kg (Table 4), below the customer price projected in Google's Suncatcher study but within reach of that study's cost scenarios for heavily reused Starships (R1-056). Beaming remains heavier than GPUs in orbit at every non-speculative tier.

- *Transmitters will improve.* That is the one lever that matters for beaming, and we give its target: 130 W/kg for next-tier geostationary beaming (P1). No such space transmitter is in the ledger.
- *Mirrors are not meant to power a continuous load.* Their business case is extending existing solar farms into dawn and dusk, which is also where a dawn-dusk orbit serves each site (D22), and onto such farms they convert only a few percent of their light (section 4.4). Our metric follows the handoff's question, continuous AI load, and so does not value shifting energy in time. A fair test compares a mirror-delivered MWh at dawn with a battery-shifted MWh, and we leave it open.
- *GPUs in orbit are a moving target too.* The filed goal of 100 kW per ton would lower A to 30 to 43 kg per kW, raising every bar in this paper.

11. Reproduce and contribute

```
git clone https://github.com/scalably-io/orbital-power-fabric && cd orbital-power-fabric
python3 -m pip install numpy scipy matplotlib pytest
./scripts/reproduce.sh          # validates the ledger, runs the tests, regenerates verdict, tables, f
python3 paper/build_paper.py    # regenerates this paper; fails on any typed number or unknown citatio
```

To change an assumption, edit it in `opf/scenarios.py` (each value carries its claim ID) and rerun. To correct a claim, open an issue with the claim ID, the better source, its locator and a verbatim quote. Code is Apache-2.0, text, data and figures are CC BY 4.0, and the chip design files are CERN-OHL-P-2.0.

Appendix A. Derivations

Each derivation in `docs/DERIVATIONS.md` gives a statement, the derivation, its assumptions, the code that implements it and the test that pins it.

ID	DERIVATION
D01	Slant range at elevation e
D02	Visibility cap and time fraction
D03	Receiver fabric coverage
D04	Eclipse fraction, circular orbit, beta angle β
D05	Image of the Sun after a flat mirror
D06	Oblique ground footprint
D07	Mirror tilt
D08	Airmass
D09	Diffraction and aperture-to-aperture collection
D10	Radiator floor
D11	Shot-noise photon floor for b-bit read-out
D12	Direct photon→charge read-out (no amplifier)
D13	Per-MAC energy of an N × N optical matrix engine
D14	Thermo-optic weight hold energy per MAC is independent of N
D15	Clements mesh
D16	Transformer decode operator counts
D17	Launched mass per continuous kW of AI load
D18	Brightness theorem: sunlight power per spatial mode
D19	Streamed-weight energy per MAC
D20	GEO transmitter mass: power-limited or aperture-limited, whichever is larger
D21	Photons for b bits over the output distribution (balanced incoherent read-out)
D22	Per-site service in a dawn-dusk orbit

Appendix B. Parameter registry

Values used by the model with their sources (`opf/params.py`). "Assumption" marks a value with no source, stated as such.

PARAMETER	VALUE	UNIT	SOURCE
`solar_constant`	1360.8	W/m2	R2-119
`sun_angle`	0.5331	deg	R12-032
`land_fraction`	0.291	-	R12-019
`clear_sky_fraction`	0.66	-	R12-016
`geo_design_life_years`	15	y	R12-022
`rack_kw`	142	kW	R1-095
`pue`	1.15	-	assumption: illustrative (a typical liquid-cooled data centre); not swept in v1.0
`pv_eff`	0.35	-	R2-001
`dc_dc_space`	0.9	-	R2-002
`dc_rf`	0.7	-	R2-003
`antenna`	0.9	-	R2-004
`atmosphere_rf`	0.98	-	R2-005
`beam_capture`	0.95	-	R2-006
`rectenna`	0.78	-	R2-007
`dc_dc_ground`	0.9	-	R2-008
`pv_w_per_kg_flown`	84.6	W/kg	R2-093
`pv_w_per_kg_today_caltech`	25	W/kg	R2-067
`reflector_altitude`	625000	m	assumption: handoff altitude, swept
`mirror_reflectivity`	0.9	-	assumption: no sourced value yet
`atm_clear_zenith`	0.64	-	R1-136
`ground_pv_eff`	0.3	-	assumption: full-sun value; low light is lower (R5)
`membrane_areal_density`	0.00425	kg/m2	R1-149
`inverter_start_irradiance`	2.6	W/m2	R5-069
`radiator_emissivity`	0.9	-	assumption: typical coating
`radiator_temp`	350	K	assumption: handoff value; AI1 needs >=322 K two-sided (R1-014)
`orbital_compute_kw_per_tonne`	75	kW/t	R1-005
`satellite_life_years`	5	y	R1-023
`server_life_years`	5.5	y	R1-035

Appendix C. Evidence rows cited in this paper

Each row is copied from the ledger (evidence/claims.csv, 1,284 rows in all), with its verbatim quote. Status says what kind of evidence it is: a vendor statement is not a measurement.

ROW	CLAIM	VALUE	STATUS	SOURCE	VERBATIM QUOTE
R1-005	SpaceX AI1 vehicle efficiency (compute power per vehicle mass)	75 kW/ton	VENDOR-STATED	www.spacex.com · Starmind page, AI1 OVERVIEW tab spec table	"VEHICLE EFFICIENCY 75 kW / ton"
R1-014	AI1 radiator heat flux and minimum radiating temperature at 175 kW avg	1094 W/m2 per m2 of listed area; T >= 322 K (two-sided) or 383 K (one-sided) W/m2; K	DERIVED	www.spacex.com · Starmind THERMAL + OVERVIEW tabs	"DEPLOYABLE LIQUID RADIATOR 160 m^2"
R1-020	SpaceX filing: compute per launched mass and annual build-out scenario	100 kW/t; 1 Mt/yr -> 100 GW/yr kW/t; GW/yr	PROJECTED	api-prod.fcc.gov · Application Narrative p.4 [PDF p.5]	"launching 1 million tonnes per year of satellites generatin 100 gigawatts of AI compute capacity annually"
R1-023	SpaceX ODC estimated operational lifetime per satellite	5 years	FILING	api-prod.fcc.gov · Schedule S (draft), S1.c Estimated Operational Lifetime [PDF p.1]	"SpaceX Orbital Data Center System Non-Geostationary (

ROW	CLAIM	VALUE	STATUS	SOURCE	VERBATIM QUOTE
R1-035	SpaceX estimated useful lives: satellites vs servers	satellites 3-5 yr; servers & networking 5-6 yr; data center infrastructure 20-25 yr years	FILING	www.sec.gov · 424B4, Notes to financial statements - Property, plant and equipment table	"Servers and networking equipment .... 5 - 6 - 5 years"
R1-053	Suncatcher launch-price break-even target and timing	<200 $/kg by ~2035 (mid-2030s) $/kg	PROJECTED	arxiv.org · arXiv 2511.19468v2, Sec 2.4 [PDF p.6]	"If the learning rate is sustained—which would require ~18 to <$200/kg by ~2035"
R1-056	Suncatcher: Starship cost to SpaceX with reuse; fuel floor	<=60 $/kg (10x reuse); <=15 $/kg (100x); fuel floor ~8 $/kg $/kg	PROJECTED	arxiv.org · arXiv 2511.19468v2, Sec 2.4 / 4.4 [PDF p.11]	"launch costs to LEO may drop to ≲ $60/kg (10× compone
R1-057	Suncatcher current launch price assumption (Falcon 9 reusable)	3,600 $/kg	STUDY-ASSUMED	arxiv.org · arXiv 2511.19468v2, Sec 4.4 [PDF p.12]	"we use current launch price $3,600/kg, based on Falcon 9
R1-073	Starcloud-3 mass (FCC DAS input)	2,000 kg	FILING	api-prod.fcc.gov · Starcloud 2026-09-16 response, attached NASA DAS 4.5-2 logs (Spacecraft = Starcloud-3)	"Final Area-To-Mass = 0.6510 (m^2/kg) Initial Mass = 2000
R1-076	Starcloud-3 power (press)	100 kW (Dec 2025) vs 200 kW (Mar 2026) kW	VENDOR-STATED	thenextweb.com · The Next Web 2026-03-30	"Starcloud-3: a 200-kilowatt, three-tonne spacecraft desi SpaceX built for launching Starlink satellites from Starship
R1-094	GB300 NVL72 GPU and CPU count	72 Blackwell Ultra GPUs; 36 Grace CPUs count	VENDOR-STATED	docs.nvidia.com · NVIDIA Enterprise RA, Components, NVLink section	"Since the GB300 NVL72 integrates 36 NVIDIA Grace CPU scale design"
R1-095	GB300 NVL72 maximum rack power	142 kW	VENDOR-STATED	docs.nvidia.com · NVIDIA Enterprise RA, Components, GB300 NVL72 rack highlights	"Full rack requiring up to 142 kW"
R1-100	GB300 NVL72 GPU HBM capacity and bandwidth; fast memory	20 TB HBM3E, up to 576 TB/s; 37 TB fast memory TB	VENDOR-STATED	www.nvidia.com · nvidia.com GB300 NVL72 Specifications table	"Fast Memory 37 TB GPU Memory | Bandwidth 20 TB | Up
R1-112	Earendil-1 spacecraft mass	142 kg	FILING	api-prod.fcc.gov · Reflect Orbital Application Narrative p.2, I.A Space Segment [PDF p.2]	"EARENDIL-1 is a 142 kg spacecraft designed to reflect su deployable, highly specular thin-film reflector."
R1-113	Earendil-1 reflector size and stowed/deployed envelope	18 m x 18 m (324 m2); stowed 840 x 857 x 400 mm m	FILING	api-prod.fcc.gov · Reflect Orbital ODAR, Sec II Satellite and Orbital Parameters, p.1 [PDF p.1]	"a deployable, highly specular 18 m x 18 m square thin-film 840.05 mm x 856.7 mm x 400 mm."
R1-115	Earendil-1 mission and license duration	~1 yr operations; 2-yr license term years	FILING	docs.fcc.gov · DA 26-706 fn 39 / Grant Condition 16 [PDF p.12]	"We grant Reflect Orbital authority to operate Earendil-1 f operations and less than one year of deorbit maneuvers"
R1-116	Earendil-1 reflector drag deorbit time after raising	~5 months	FILING	api-prod.fcc.gov · Reflect Orbital 2026-01-05 supplement, answer 1.a, p.2 [PDF p.2]	"Once the reflector has been deployed and raised to its op approximately five months."
R1-136	Atmospheric transmission for reflected sunlight (Hottel model)	~0.64 at zenith; ~0.13 at pass start/end; <=0.70 at elevated sites fraction	STUDY-ASSUMED	eprints.gla.ac.uk · Celik & McInnes 2022 ASR, Sec 3.3.1 / Fig 5 [PDF p.15]	"even at the zenith point, the transmission efficiency rema levels approximately 13% at the beginning and the end of a
R1-149	ACS3 quadrant mass (with hardware) and implied membrane areal density	85 g per ~20 m2 quadrant -> 4.25 g/m2 g; g/m2	DERIVED	ntrs.nasa.gov · Wilkie ISSS 2023 slides (NTRS 20230008378), slide 'Solar Sail Quadrant Design' [PDF p.17]	"Individual Quadrant Mass (w/ hardware): 85 g"
R1-151	ACS3 system masses	~16 kg total; SBS ~7.7 kg; bus ~8.3 kg kg	VENDOR-STATED	ntrs.nasa.gov · Wilkie ISSS 2023 slides, 'Spacecraft Elements' [PDF p.12]	"Total ACS3 Mass: ~16 kg"
R2-001	NASA OTPS SBSP chain factor: solar cell (PV) efficiency	0.35 fraction	STUDY-ASSUMED	www.nasa.gov · Table 6, PDF p.59 (printed p.42); also Table 4 PDF p.53 and Table 2 PDF p.28	"Solar Cell Efficiency 0.35 0.35 Rodenbeck et al."
R2-002	NASA OTPS chain factor: in-space DC-to-DC conversion	0.90 fraction	STUDY-ASSUMED	www.nasa.gov · Table 6, PDF p.59 (printed p.42)	"DC to DC Conversion Efficiency 0.9 0.9 Rodenbeck et al."
R2-003	NASA OTPS chain factor: DC-to-RF conversion	0.70 fraction	STUDY-ASSUMED	www.nasa.gov · Table 6, PDF p.59 (printed p.42)	"DC to RF Conversion Efficiency 0.7 0.7 Rodenbeck et al."
R2-004	NASA OTPS chain factor: antenna emission	0.90 fraction	STUDY-ASSUMED	www.nasa.gov · Table 6, PDF p.59 (printed p.42)	"Antenna Emission Efficiency 0.9 0.9 Rodenbeck et al."

ROW	CLAIM	VALUE	STATUS	SOURCE	VERBATIM QUOTE
R2-005	NASA OTPS chain factor: atmospheric travel	0.98 fraction	STUDY-ASSUMED	www.nasa.gov · Table 6, PDF p.59 (printed p.42)	"Atmospheric Travel Efficiency 0.98 0.98 Rodenbeck et al
R2-006	NASA OTPS chain factor: beam collection (capture)	0.95 fraction	STUDY-ASSUMED	www.nasa.gov · Table 6, PDF p.59 (printed p.42)	"Beam Collection Efficiency 0.95 0.95 Rodenbeck et al."
R2-007	NASA OTPS chain factor: rectenna array reception	0.78 fraction	STUDY-ASSUMED	www.nasa.gov · Table 7, PDF p.59 (printed p.42)	"Rectenna Array Reception Efficiency 0.78 0.78 Rodenbe
R2-008	NASA OTPS chain factor: ground DC-to-DC conversion	0.90 fraction	STUDY-ASSUMED	www.nasa.gov · Table 7, PDF p.59 (printed p.42)	"Convert on Ground DC to DC Conversion Efficiency 0.9 (
R2-018	NASA implied whole-system specific power at PV DC output (RD1 / RD2)	938 / 909 W/kg	DERIVED	www.nasa.gov · Tables 6 and 9, PDF p.58-61	"Incident solar energy (MW) 15,689.33 25,982.50 Calculat
R2-023	NASA baseline LCOE (RD1 / RD2)	0.61 / 1.59 USD/kWh	STUDY-ASSUMED	www.nasa.gov · Executive summary, PDF p.8 (printed vii); Table 22 PDF p.78	"The baseline lifecycle cost of electricity for RD1 is 0.61 $/
R2-028	NASA: SBSP not cost-competitive even with free launch (all else constant)	not competitive qualitative	STUDY-ASSUMED	www.nasa.gov · Section 4.4, PDF p.39 (printed p.22)	"SBSB still would not be cost-competitive with alternative assuming all else remains constant."
R2-042	JAXA/METI 2015 GaN phased array: amplifier efficiency, total DC-RF, rectenna RF-DC, panel mass	60.3 / 35.1 / 59; 16.1 kg per 60x60 cm panel percent; kg	MEASURED	ieeexplore.ieee.org · Section IV.C.1, p.239 (PDF p.11)	"Each panel included 76 GaN class-F amplifiers, whose av efficiency (DC–RF) was 35.1%."
R2-043	JAXA 2015 GaN transmitter: RF-out specific power and areal density	28 W/kg; 44.7 kg/m2 W/kg; kg/m2	DERIVED	ieeexplore.ieee.org · Section IV.C.1, p.239 (PDF p.11)	"The phased array was composed of four 2.5 cm × 60 cm
R2-052	Caltech MAPLE beam detected on Earth (C005)	detected on Caltech roof, 2023-05-22 (local) qualitative	MEASURED	www.caltech.edu · Press release body	"This transmitted energy was detected by a receiver on th Engineering on Caltech's campus in Pasadena on May 22."
R2-053	MAPLE in-space end-to-end DC-DC link efficiency (sidewall / broadside rectenna, ~30 cm)	3.37 / 3.66 percent	MEASURED	chic.caltech.edu · Section 3.1, Acta Astronaut. 224 p.233	"peak rectified powers of 231 mW and 251 mW correspon ηlink = 3.66% for the sidewall and broadside rectennas, re
R2-057	MAPLE beam-to-Earth received signal above noise floor	up to 20 dB above noise floor	MEASURED	chic.caltech.edu · Section 6, p.241	"the flexible aperture could focus power successfully to d noise floor"
R2-065	Caltech Joule 2025 end-to-end system efficiency, 24-h average, 10 GHz GEO (current / middle / asymptotic)	2.77 / 4.14 / 6.63 percent	PROJECTED	www.its.caltech.edu · Table 1, article p.5 (PDF p.6)	"hsystem 2.77% 4.14% 6.63%"
R2-067	Caltech Joule 2025 average specific power of power station (current / middle / asymptotic)	25 / 199 / 327 W/kg (avg delivered power per kg)	PROJECTED	www.its.caltech.edu · Table 2, article p.7 (PDF p.8)	"Average specific powera W/kg 25 199 327"
R2-072	NRL SWELL in-orbit laser power beaming: output, end-to-end efficiency, link distance (C011)	~1.5 W; ~11%; 1.45 m W; percent; m	MEASURED	www.nrl.navy.mil · News release body (retrieved via web.archive.org id_ snapshot; primary returned HTTP 403)	"SWELL has been providing nearly continuously about 1.5 end-to-end efficiency around 11%, more than ten times th distance is 1.45 meters."
R2-093	Redwire ROSA wing-level BOL specific power (flown iROSA and catalogue)	SUNSTONE 84.6; TANZANITE 97.8; SAPPHIRE 80.6; PHENACITE 71.0; OPAL 124.1; DIAMOND 91.7 W/kg	DERIVED	rdw.com · Specification table	"TANZANITE 24.8 180 17.6 4.15"
R2-108	FCC MPE, general population/uncontrolled, 1,500-100,000 MHz (covers 2.45 and 5.8 GHz)	1.0 mW/cm2 (10 W/m2), 30-min average mW/cm2	FILING	www.ecfr.gov · 47 CFR 1.1310, Table 1 to (e)(1), part (ii); eCFR as of 2026-09-22	"(ii) Limits for General Population/Uncontrolled Exposure .

ROW	CLAIM	VALUE	STATUS	SOURCE	VERBATIM QUOTE
R2-119	Measured total solar irradiance at 2008 solar minimum	1360.8 ± 0.5 W/m2	MEASURED	www.atmosp.physics.utoronto.ca · Abstract [1], GRL 38 L01706 p.1	"The most accurate value of total solar irradiance during th −2"
R3-037	Laser wall-plug floor for a 1000-input single-pass crossbar at 10 GS/s, 5-bit at the detector, 20 dB optical loss, 40% wall-plug efficiency; ADC/DAC neglected	10 fJ/MAC (laser only)	PROJECTED	media.springernature.com · Supplementary Note 1 "Optical Energy Overheads in Single-pass Architectures"	"Considering the plug energy efficiency of the optical pov total energy budget per MAC for such a system would be neglected)."
R3-042	Lightmatter photonic processor (4 photonic tensor cores 128x128 + 2 digital dies): full-system measured spec, running ResNet, BERT, Atari RL	65.5 TOPS at 78 W electrical + 1.6 W optical TOPS; W	MEASURED	media.springernature.com · Supplementary Table S3 (main article paywalled)	"This work 4×128×128 65.5 TOPS 7×11 (77) 500 MHz 349
R3-043	Lightmatter processor full-system efficiency (average)	0.81 TOPS/W (= 1.23 pJ/OP = 2.5 pJ/MAC)	MEASURED	media.springernature.com · Supplementary Table S4	"2024 This Work 65 avg, 262 peak 78 avg, 150 peak 0.81 a
R3-045	Lightmatter: ratio of full-system to photonic-core power	~318x (core = 0.31% of system power) ratio	DERIVED	media.springernature.com · Supplementary Tables S3/S4 and Section III	"78W electrical, 1.6W optical, 250mW / PTC"
R3-062	Optical Transformers component assumption: ADC energy per 7-bit sample at 10 GHz	3.17 pJ per sample	STUDY-ASSUMED	arxiv.org · Appendix D, list of assumed costs	"EADC = 3.17 pJ per 7-bit sample."
R3-067	NVIDIA Blackwell Ultra (B300/GB300) GPU maximum power	1,400 W (max TGP)	VENDOR-STATED	developer.nvidia.com · NVIDIA Technical Blog, Table 2 "NVIDIA GPU chip comparison"	"Max power (TGP) Up to 700W Up to 1,200W Up to 1,400
R3-070	B300 chip-level energy per FLOP at max TGP and peak dense throughput	93 (NVFP4); 280 (FP8) fJ/FLOP (= 187 / 560 fJ/MAC)	DERIVED	developer.nvidia.com · Derived from Table 2 (TGP and dense FP4/FP8)	"Max power (TGP) Up to 700W Up to 1,200W Up to 1,400
R3-078	Measured (microbenchmark-attributed) HBM access energy on NVIDIA GH200, total incl. control	11.68 pJ/bit	MEASURED	escholarship.org · Table 3 "Control, datapath, and total energy parameters", GH200 Total row	"Total 11.68 4.87 1.45 23.60 8.46 4.86 11.02 6.61 0.68 0.52"
R3-081	HBM3e energy per bit as reported by SK hynix (IEDM 2023), via secondary citation	3.44 pJ/bit	STUDY-ASSUMED	arxiv.org · RPU paper, Section on HBM-CO model validation	"We validate our HBM-CO model against HBM3e [43] repo
R3-082	HBM access energy range used in the LLMET LLM-serving model (HBM2E / HBM3E, IO included)	5.7-6.6 pJ/bit	STUDY-ASSUMED	arxiv.org · Section 4 insight text; Table 2 "Memory Access Energy"	"their access energy remains far below HBM energy (≈5.7-
R5-027	Ember: solar + battery sized for 1 GW flat output 97% of hours in Las Vegas (6 GW PV + 17 GWh), LCOE	104 $/MWh	STUDY-ASSUMED	ember-energy.org · Ember 'Solar electricity every hour of every day...' Chapter 3, section 3.1.1 and 'How close to 24/365' paragraph	"sunny cities can get as close as 97% of the way to full 24,
R5-028	Ember: 99.4% of 24/365 in Las Vegas needs 7 GW PV + 35 GWh per 1 GW; LCOE	167 $/MWh	STUDY-ASSUMED	ember-energy.org · Chapter 3, section 3.2 'How close to 24/365 solar generation is optimal?'	"This would need 7 GW of solar and 35 GWh of battery sto
R5-029	Paces/Scale Microgrids/Stripe: off-grid solar+storage microgrid meeting 90% of 24/7 data-center load from solar (gas backup), LCOE; vs off-grid gas turbines $86/MWh	109 $/MWh	STUDY-ASSUMED	www.offgridai.us · 'Fast, scalable, clean, and cheap enough' (Dec 2024), Summary	"A system supplying 90% of lifetime demand from solar is ($109/MWh versus $130/MWh)."
R5-031	Land per continuous kW of data-center load, 90%-solar off-grid design	102 km2 per GW continuous	DERIVED	www.offgridai.us · Derived from previous row	"Scenario 1: 90% Renewables / 10% Gas 500 1,600 100 2,5

ROW	CLAIM	VALUE	STATUS	SOURCE	VERBATIM QUOTE
R5-042	LBNL Queued Up 2026: median time from interconnection request to commercial operation, US generators built in 2025 (36 months in 2015; 22 in 2008)	61 months	MEASURED	eta-publications.lbl.gov · High-Level Findings slide, PDF p.4	"The median project built in 2025 took 61 months from the
R5-069	Irradiance at which a utility PV plant's inverter starts producing (plane-of-array), median CEC utility inverter, ILR 1.32	2.6 (IQR 1.9-3.4; range 0.5-7.6) at ideal module; ~5 at 50% low-light efficiency W/m2	DERIVED	raw.githubusercontent.com · Derived from previous row + ATB ILR	"Name,Vac,Pso,Paco,Pdco,Vdco,C0,C1,C2,C3,Pnt,Vdcmax
R5-077	Tracking utility PV power density in SI units (fixed-tilt 86.5)	59.3 MW_DC per km2	DERIVED	emp.lbl.gov · Derived from previous row	"Power density: 0.35 MWDC/acre (2.8 acres/MWDC) for fi tracking"
R6-016	NVIDIA 800 VDC architecture: end-to-end efficiency gain vs current 54 V in-rack systems	up to 5 percent	VENDOR-STATED	developer.nvidia.com · NVIDIA Technical Blog 2025-05-20, section 'Key benefits of 800 VDC', bullet 'Efficiency'	"Improves end-to-end efficiency by up to 5% compared to utilization."
R6-122	DERIVED on-site PV to the 12 V bus: AC-coupled (inverter, step-up + step-down transformer, UPS, PSU, 54->12 V) vs DC-coupled (MPPT DC-DC onto 800 VDC bus, 800 V->12.5 V)	0.8905 vs 0.9506 (x1.068); remote PV over the grid 0.8459 (x1.124) fraction; ratio	DERIVED	eta-publications.lbl.gov · Formula in notes; inputs R6-001, R6-002, R6-004, R6-007, R6-014, R6-023, R6-003	"MPPT & Charge Controller 97.4%5, 97.62, 3, 98.01"
R7-046	NTT O-E-O transistor, normally-on at 2.5 Gbit/s: total energy per bit with optical conversion gain 2.3	60.1 ± 5.7 fJ/bit	MEASURED	media.springernature.com · Supplementary Information, total energy/efficiency section	"Normally-on condition operating at a bit rate of 2.5 Gbit/ A/W and Vbias = 2 V"
R8-072	1310 nm SSC/SOA-integrated DFB laser: wall-plug (power) efficiency at >100 mW, 75 C	20 %	MEASURED	doi.org · Abstract (Optical Engineering 62(4) 046102, 2023; via Crossref)	"with record power efficiency of 20% at optical power of r
R8-084	Electronic weight density: 8-bit weights per mm2 in TSMC N2 SRAM	4.76 million 8-bit weights per mm2	DERIVED	marklapedus.substack.com · Derived from the 38.1 Mb/mm2 density	"resulting in a 38.1Mb/mm2 density."
R8-099	Non-volatile magneto-optic retention	4 days (measured) retention	MEASURED	www.nature.com · Results	"To confirm the long-term non-volatile stability of our dev shift of the CW and CCW modes after programming and s
R8-103	Memresonator phase-shift range and size: 20-um-diameter ring, 0.18 pi (0.08 nm) shift, estimated L_pi 0.35 mm	20 um diameter; 0.18 pi um; rad/pi	MEASURED	arxiv.org · Results, Device Characteristics	"The 20-μm diameter memresonator achieves about a 0.0 Supplementary Note S5) in the LRS, leading to an estimat
R8-104	Memresonator weight density upper bound (ring footprint only)	3,183 weights per mm2	DERIVED	arxiv.org · Results, Device Characteristics (20 um diameter)	"The 20-μm diameter memresonator achieves about a 0.0
R8-112	Ferroelectric HfAlO memory retention	1000 s measured; >10 years extrapolated s	MEASURED	pmc.ncbi.nlm.nih.gov · Results	"No noticeable change can be observed during the measu time is estimated to be more than 10 years"
R8-124	Area to hold all 72.71 B weights of Qwen2.5-72B (E005) one-weight-per-cell, by technology	N2 SRAM 0.015 m2; memresonator 22.8 m2; MEMS 327 m2; MO 571 m2; PCM tensor core 7,336 m2; BTO mesh 12,536 m2 m2	DERIVED	www.nature.com · Inputs: R8-084, R8-104, R8-123, R8-100, R8-117, R8-108	"The footprint of the fabricated circuit is 2 × 10 mm 2 and
R8-126	Loss of a Clements mesh (depth N) built from measured NV elements	BTO 1.48 dB/stage: 95 / 189 / 379 dB; MEMS 0.33 dB/stage: 21 / 42 / 84 dB (N=64/128/256) dB	DERIVED	www.nature.com · Input: R8-107 (1.48 dB per PUC), R8-121 (0.33 dB)	"The fabricated NV-PUC exhibits an insertion loss of 1.48 c

ROW	CLAIM	VALUE	STATUS	SOURCE	VERBATIM QUOTE
R8-128	Receiverless detector node from measured parts: electrons needed to swing C by 1 V vs the shot floor	2 fF/1 V: 12,483 e- (shot-limited to 6.8 bits); 8-bit at 1 V needs C >= 10.5 fF; 4-bit on 2 fF needs >= 443 e- (35.5 mV swing) electrons per read	DERIVED	www.nature.com · Inputs: R8-059 (2 fF), R8-026 (kTC), D11, D12	"Here we use a photonic-crystal platform to demonstrate integration at only 2 fF."
R9-005	Full streaming path energy for a GPU weight fetch, HBM + L2 + L1 (Table 3 HBM column is the HBM increment only)	18.00 pJ/bit	DERIVED	escholarship.org · Section 4.1, Eqs. 6-9 (e_HBM = e_HBM+L2+L1 - e_L2 - e_L1); Table 3 GH200 Total row	"To differentiate energy in the memory subsystem, we run
R9-030	Marvell authors' survey: pluggable DSP costs ~100-150 ns and ~15-20 pJ/bit; Broadcom TH6 CPO ~5 pJ/bit; Intel OCI 4 Tbps at 5 pJ/bit	15-20 (pluggable); ~5 (TH6 CPO) pJ/bit	STUDY-ASSUMED	arxiv.org · Section VII.C 'Photonic Interconnects in Datacenter'	"(~15–20 pJ/bit) from pluggable transceivers. Broadcom's ~5 pJ/bit"
R10-052	Global cloud fraction from the MODIS cloud mask (King et al., IEEE TGRS 51, 3826 (2013))	~67 percent	MEASURED	atmosphere-imager.gsfc.nasa.gov · Abstract	"Globally, the cloud fraction derived by the MODIS cloud r
R11-001	MIT patent (Shen, Harris, Skirlo, Soljacic, Englund et al.): optical nonlinearity unit in which a photodetector drives the optical cavity (detector-driven O-E-O activation) between ONN layers	photodetector-driven cavity = ReLU/sigmoid activation	FILING	patents.google.com · Detailed description, 'optical nonlinearity unit 924' (FIG. 9B), col. 9-10 (USPTO PDF p. 30)	"light incident on an optical cavity is detected by a photoc off-resonance. This opto-electronic element can impleme functions including ReLU and sigmoid"
R12-016	Clear-sky probability (1 - annual cloud amount) across 12 representative large utility-PV regions (US SW x2, Atacama, UAE, Saudi, Morocco, Egypt, Australia x2, India x2, Qinghai)	all-hours mean 0.66 (median 0.68, range 0.44-0.84); day mean 0.70 (range 0.45-0.88); night mean 0.62 (range 0.43-0.80) fraction	DERIVED	power.larc.nasa.gov · POWER parameter manager, CLOUD_AMT CLIMATOLOGY definition	"The average percent of cloud amount during the tempora
R12-019	Land fraction of Earth's total surface area, present day	29.1 %	MEASURED	pubs.usgs.gov · USGS Professional Paper 1386-A, Chapter A-1, Figure 9 caption (figure gallery 1)	"Three bar graphs showing percentages of land mass versu percent), present day (land, 29.1 percent), and no glacier ic
R12-022	Median expected (design) lifetime of GEO communications satellites in the UCS Satellite Database	15 (n=426; 72% list exactly 15 y; mean 14.35; P25=P75=15) years	DERIVED	www.ucs.org · UCS Satellite Database, text export 'UCS-Satellite-Database 5-1-2023', column 'Expected Lifetime (yrs.)'; filter Class of Orbit = GEO and Purpose = Communications	"The database contains 28 types of data for each satellite, (mass, power, launch date, expected lifetime)"
R12-030	Standard quantum limit (shot noise) per MAC for homodyne optical neural networks, MNIST benchmarks	0.5-1 aJ (5-10 photons) at N=100; 50-100 zJ (0.5-1 photon) at N=1000 J per MAC (optical)	STUDY-ASSUMED	arxiv.org · Hamerly et al. PRX 9, 021032 (2019); arXiv v2 p. 4, section 'Deep Learning at the Standard Quantum Limit', Fig. 2(c) discussion	"This is around 0.5–1 aJ (5–10 photons) for the small netw photon) for the large network (inner layer size N = 1000)."
R12-032	Apparent (angular) diameter of the Sun seen from Earth at 1 AU	1919 (max 1952, min 1887) arcsec	MEASURED	nssdc.gsfc.nasa.gov · NASA NSSDCA Sun Fact Sheet (D. R. Williams), table 'Observational Parameters', rows 'Apparent diameter from Earth'	"Apparent diameter from Earth At 1 A.U.(seconds of arc) 1 (seconds of arc) 1887."