Where the Space Economy's Rent Goes Next: Six Physical Walls Cheap Launch Cannot Move

Make orbit nearly free and scarcity jumps to the inputs that answer to thermodynamics, decay physics, and treaty-fixed spectrum, not to learning curves.

Frame

When a system scales, the money moves to the input that cannot scale with it. This board names that input, the date it starts to bite, and the line that would break the call.

Area

space and the space economy (launch vehicles and propulsion, satellites and constellations, in-orbit infrastructure and servicing, lunar and cislunar, space materials and manufacturing, ground segment and spectrum)

Horizon

2030 to 2040

Issued

2026-06-14

Method

Wide cast, adversarial gate, public resolution criteria.

Board summary

The cross-cutting read

We see the same trade six times, and the market misreads it the same way each time. The funded, visible bottleneck in space is always transport: launch cost, cadence, landers, GPUs. Cheap reusable launch is collapsing that layer, and the market is pricing the collapse as if it fixes the whole stack. It does not. It moves the binding constraint onto a short list of inputs that ride no learning curve because a physical or institutional invariant pins them. In vacuum, heat rejection is fixed by Stefan-Boltzmann T-to-the-fourth (P1). The 354-hour lunar night sets a linear energy-storage mass floor that only fission escapes (P2). Deep-space links need treaty-protected S/X/Ka bands and 70 m steel that takes a decade to pour (P3). LEO position error is driven by thermospheric density that a near-dead, non-commercial sensor base cannot forecast (P4). Every US radioisotope heat source is clad in one arc-remelted iridium alloy from a single reconstituted ORNL line, now bidding against the hydrogen buildout for the same 7-tonne-per-year metal (P5). And the stratosphere is a finite, 30-year-clearance alumina sink that a self-replacing megaconstellation fills monotonically (P6). In each case the scarce thing is one inelastic input: square-meters of deployable radiator per kilowatt, kilograms of night-survival storage per watt, large-aperture antenna-hours, in-situ density data, qualified iridium clad-vent-set throughput, qualified non-aluminum demisable structure. And in each case capital is still pricing the layer above it. The themes are getting loud (cooling is hard, Pu-238 is short, reentry chemistry sits in geoscience papers); the specific rent-bearing inputs are mispriced, absent from any equity model, sell-side note, or futures contract. We think that gap, between an audible theme and an unpriced input, is the whole edge.

At a glance

#	Claim	Binding constraint	Case	Call	Resolves
P1	The 2025-2026 orbital-data-center pitch (Starcloud raised $170M at $1.1B in March 2026, Lumen Orbit, China Three-Body...	Qualified deployable space-radiator AREA (m2 of micrometeoroid-survivable, fluid-loop-coupled panel) per...	82%	62%	2034-12-31
P2	The 2025-2026 lunar narrative is a transport story: Starship cadence, Artemis flight rates, who lands what mass when...	Night-survival energy/power mass per continuous watt across the 354-hour lunar night (kg per continuous-watt...	82%	70%	2038-12-31
P3	Cheap reusable launch and the Artemis crewed-lunar program drive a steep rise in deep-space mission count through the...	Large-aperture (34 m / 70 m class) antenna-hours on the ITU-protected deep-space S/X/Ka bands, jointly with...	78%	62%	2034-06-30
P4	Cheap launch and proliferated LEO are read as a traffic-management problem you solve with more radars or optical...	Real-time in-situ thermospheric neutral-density driver data (direct accelerometer-derived density plus...	72%	52%	2035-12-31
P5	The plutonium-238 production ramp at Oak Ridge absorbs every headline and every GAO citation on why NASA cannot fly...	Iridium DOP-26 alloy clad-vent-set fabrication capacity at ORNL, and the high-purity iridium feedstock...	72%	38%	2035-12-31
P6	Cheap launch and short orbital lifetimes turn every megaconstellation into a mass conveyor: all launched mass must...	Qualified non-aluminum demisable spacecraft structural material at constellation scale, plus the...	72%	52%	2038-12-31

Case is the strength of the structural thesis. Call is the probability on the exact dated clause.

P1

Orbital compute does not hit a launch-cost wall or a power wall first. It hits a radiator-area wall. We think the binding constraint on space-based AI compute through the 2030s is...

Domain: space

2034-12-31

Structural case	Our call, dated	Resolves
82%	62%	2034-12-31

Heat rejection in vacuum is radiative-only and obeys Stefan-Boltzmann: power rejected scales as emissivity times area times T-to-the-fourth. Radiating at high T means running the coolant hot, which is exactly what high-density AI silicon cannot tolerate (HBM stacks and flip-chip packages degrade above roughly 85-105 C), so practical reject temperatures sit near 290-320 K where flux is ~300-650 W/m2. That fixes a hard floor on required area per kilowatt that no algorithm, no chip node, and no launch-cost collapse can erase. It gets harder still. The radiator must be DEPLOYABLE (fold into a fairing), SURVIVABLE (micrometeoroid and debris flux punctures fluid channels over a multi-year life), and POINTABLE away from the sun while the solar array points at it, a geometric conflict on the same spacecraft. Panel-plus-pumped-loop areal mass has a materials floor around 2-7 kg/m2 and has barely moved in 40 years of active NASA/ESA/JAXA effort. ISS panels run roughly 3-4 kg/m2. The inelastic input is therefore square-meters of qualified, deployable, leak-tolerant radiator per kilowatt rejected, plus the pumped two-phase fluid loop and the talent that qualifies it. This is a thermodynamic and materials limit, orthogonal to compute and to launch economics, which is why cheaper launch or better GPUs do nothing to it.

The boom

The 2025-2026 orbital-data-center pitch (Starcloud raised $170M at $1.1B in March 2026, Lumen Orbit, China Three-Body, NVIDIA-in-orbit) rests on free sunlight and free cold space. Free sunlight is real. Free cooling is not. In vacuum the only way to dump heat is to radiate it, and a two-sided panel held near room temperature emits only ~600-650 W/m2. That is a physical constant, not an engineering target. Starcloud's own white paper concedes roughly 1,600 m2 of radiator PER MEGAWATT, and a 5 GW concept needs about 8 km2 of radiator, larger than Gibraltar. Every doubling of orbital compute requires a doubling of folded-then-deployed radiator that must survive micrometeoroids, atomic oxygen, and thermal cycling for years, at a panel-plus-working-fluid-loop mass floor of a few kg/m2 that has barely moved in 40 years. Chip power density keeps climbing, but radiator W/m2 is pinned by T-to-the-fourth and material emissivity. The two curves diverge forever. Orbital compute caps out one-to-three orders of magnitude below the gigawatt hype, and the scarce, rent-bearing asset is not the GPU, the solar array, or the launch slot. It is qualified large-area deployable radiator and its leak-proof two-phase transport loop.

Why it is not priced yet	Through mid-2026 the market has only just picked up cooling is hard as a theme (WEF June 2026, SatNews physics wall March 2026, IEEE Spectrum). That framing now reads as a qualitative caveat and is priced as one. What is not priced is the quantitative claim: that deployable radiator AREA-per-kW and its kg/m2 mass floor are a hard physical limit capping orbital compute one-to-three orders of magnitude below the marketed GW scale through the 2030s, so the rent accrues to deployable-radiator and two-phase-loop suppliers rather than GPU or launch players. Capital still flows to the compute-in-orbit story (Starcloud's $170M Series A at $1.1B valuation in March 2026) on the assumption that cooling is a solvable detail. It is a thermodynamic invariant, and the area-per-kW input is nowhere in the investment narrative.
Where the price sits today	Theme-saturated as a qualitative concern (WEF, IEEE Spectrum, SatNews all ran cooling-wall pieces in 2026), but capital keeps flowing to the compute/GPU narrative. Starcloud's $170M Series A in March 2026 was valued on compute density, not demonstrated radiator technology. No dedicated deployable-radiator supplier is publicly traded or taking named VC money as a scarcity play. The quantitative claim (m2/kW floor as the rent-bearing asset) is in no sell-side coverage or VC deck. Theme priced, input not.
The binding constraint	Qualified deployable space-radiator AREA (m2 of micrometeoroid-survivable, fluid-loop-coupled panel) per kilowatt of waste heat rejected, and the ~2-7 kg/m2 areal-mass floor of the panel plus its pumped two-phase transport loop. Not cooling as a theme. The specific inelastic input is deployable radiator area-per-kW and the loop hardware behind it.
What we are watching	Announced or as-built radiator area and radiator mass fraction of any orbital-compute spacecraft above ~100 kW IT load, expressed as m2 of radiator per kW rejected and as radiator-plus-loop mass as a share of dry mass. Track: (a) the largest single orbital-compute payload actually on orbit by power (currently sub-kW thermal on Starcloud-1, Nov 2025); (b) whether any flown system rejects >100 kW; (c) radiator areal mass in kg/m2 in flight hardware (currently ~3-7 kg/m2). Leading tell: every serious orbital-DC technical disclosure or teardown that names radiator area/mass (not power or launch) as the gating line item.
What would prove us wrong	Killed if, by 2034-12-31, any orbital data center actually rejects >100 kW of IT waste heat on orbit at a radiator areal mass at or below ~1 kg/m2 (a 3-5x materials breakthrough, e.g. droplet/liquid-sheet radiators reaching >400 W/kg in qualified flight hardware), OR a flown system demonstrates >1 MW IT load without radiator area/mass being the dominant mass and cost driver. Either outcome means the area wall was engineered around and the constraint migrated off the radiator.

How we tried to break it

We tried three kills, none landed. First, droplet/liquid-sheet radiators could hit 400-600 W/kg and erase the area wall. That is a real threat, and the kill criterion addresses it: no droplet radiator has ever flown in qualified form, and contamination and capture efficiency remain undemonstrated after 40 years of study. The 2034 kill bar sits correctly at >400 W/kg in flight hardware. Second, run chips hot (400 K coolant) for 3x more radiated flux via T^4. Real leverage but bounded. AI silicon junction limits sit at 85-105 C, and chip power density climbs faster than the T^4 gain can offset, so practical reject temperature stays near 310-330 K. Third, the $170M Starcloud raise supposedly means the market has already priced this in. The round is pitched on compute capacity; Starcloud-1 carries one H100 at sub-kW thermal load, while a 5 GW system needs ~8 km2 of radiator. The funding validates the compute story, not the radiator-scarcity thesis. The call survives all three.

Why we are making the call

We think the mechanism is right and the input is precisely named. The area wall is real, qualified deployable radiators are genuinely inelastic (two or three qualified suppliers globally, multi-year qualification timelines, no learning-curve analog to silicon), and capital is still pricing the compute side of the trade. The main risk to the dated call is the 8-year horizon, over which a droplet radiator or graphene-composite breakthrough could move the mass floor. The kill criterion sits correctly at a 3-5x improvement threshold. We promote it because the physics holds, the input is specific, the not-yet-priced check passes on the quantitative framing rather than the theme, and the attacks do not kill it.

If the call is right

If orbital compute caps one-to-three orders below the marketed gigawatt scale, the value lands on whoever holds qualified large-area deployable-radiator and two-phase pumped-loop IP, not on the GPU or launch line item. Today that is a thin bench: Advanced Cooling Technologies (private, US) and ARQUIMEA (private, Spain) on loop heat pipes and deployable structures, and Northrop Grumman, which holds the ISS-derived ~70 kW radiator-wing lineage Starcloud is adapting. The radiator outmasses the compute roughly ten to one at megawatt scale, so the thermal subsystem becomes the priced bottleneck of every orbital-DC bill of materials.

Who gains

Northrop Grumman (NOC): only flight-proven MW-class heat-rejection lineage (ISS radiator wings) plus deployable-structure heritage; sells the gating subsystem as orbital bus power scales.

Advanced Cooling Technologies and ARQUIMEA: private pure-plays on loop heat pipes and deployable radiators with no silicon-style learning curve; become acquisition or pre-IPO targets as the constraint is named.

Thales Alenia Space (Leonardo, BIT: LDO) and Airbus Defence and Space: incumbent satellite thermal-subsystem vendors that capture European bus-thermal rent as power per spacecraft climbs.

Who loses

Starcloud, Lumen Orbit, and China Three-Body Computing Constellation: TAM models assume thermal scales linearly with compute; the area wall caps deliverable IT load, so the GW pitch and its valuations are repriced down.

NVIDIA (NVDA) on the orbital-TAM slice only: H100 flew on Starcloud-1, Blackwell on Starcloud-2, but any orbital-GPU growth narrative is bounded by radiator area; terrestrial demand is untouched.

Launch providers counting on orbital-DC tonnage (SpaceX, smallsat-launch field): radiators are the mass, so an uneconomic MW tier shrinks the launch-cadence thesis.

What reprices

No clean public instrument prices the radiator scarcity directly. The expression is long Northrop Grumman (NOC) and Leonardo (LDO) on the thermal-IP side, and a fade of NVDA's orbital-TAM premium and of any future orbital-DC SPAC or IPO. The pure-plays (ACT, ARQUIMEA, Boyd) are private, so the trade is M&A and pre-IPO, not a ticker.

The next constraint it creates

Once area binds, the constraint moves to qualified leak-tolerant two-phase fluid-loop life under multi-year micrometeoroid and atomic-oxygen flux, and to the handful of engineers who can qualify it. If droplet or liquid-sheet radiators ever fly above ~400 W/kg, the constraint jumps again to contamination control and droplet-capture efficiency.

Earliest sign it has begun

The first orbital-DC technical disclosure or teardown that names radiator area or radiator mass fraction (not power or launch cost) as the gating line item, and the first flown system rejecting above 100 kW of IT waste heat with radiator-plus-loop mass dominating dry mass.

P2

Every kilowatt-class non-polar lunar surface asset through 2040 is gated by the 354-hour lunar night. We think the binding constraint is night-survival energy-storage mass per continuous...

Domain: space

2038-12-31

Structural case	Our call, dated	Resolves
82%	70%	2038-12-31

The lunar synodic day gives roughly 354 hours of continuous darkness over most of the surface. Energy storage to bridge that night scales linearly with both power level and duration: at 100-150 Wh/kg derated (the realistic flight-qualified specific energy for space-rated lithium-ion, after thermal management and depth-of-discharge derating), bridging one watt continuous needs roughly 2.4-3.5 kg of cells before thermal keep-alive mass. Regolith thermal storage and flywheels shift the constants but do not break the linear scaling. Fission is the only power source whose delivered mass is independent of night length, which is why it becomes mandatory above roughly a few hundred watts continuous. The inelastic input is therefore kg of qualified night-survival power capacity per continuous watt across 354 hours, and at kilowatt-and-up scale that collapses to the count of qualified, flight-ready fission surface power units, of which there will be a single-digit count through this decade.

The boom

The 2025-2026 lunar narrative is a transport story: Starship cadence, Artemis flight rates, who lands what mass when. The wall nobody is pricing is that landing the mass is the easy half. Keeping any non-polar surface asset alive through 14.77 Earth days of continuous darkness, every month, is the hard constraint. NASA's own analysis puts the battery mass to bridge a single lunar night at ~16 metric tonnes for 40 kW of continuous power, a mass that exceeds most landers' entire payload, for power a single data-center rack exceeds. Fission is the only source whose delivery mass does not scale with night length, which is why it becomes mandatory above a few hundred watts continuous. The US fission surface power program has already slipped from a 2026 target to 2030, the power spec jumped from 40 kW to over 100 kW with the same sub-6-tonne mass cap, and the January 2026 NASA/DOE MOU reset to a fresh downselect-to-two-designs process. The set of continuously-powered lunar surface assets is therefore capped at the reactor count, no matter how many landers fly.

Why it is not priced yet	The lunar investment and policy narrative through 2026 runs on transport (Starship/Artemis cadence, CLPS landers) and ISRU end-products (water ice, oxygen, regolith). Night-survival power shows up only as a known engineering caveat in NASA technical papers, never as the priced constraint that gates the whole surface economy. No sell-side equity coverage, futures, or capital-markets instrument captures the single-source fragility of the fission surface power program as the chokepoint deciding which lunar ventures are physically viable. The market assumes solar plus batteries, scale as needed. The 354-hour night plus the roughly 3 kg per continuous watt storage floor makes that false above toy power levels, and that mass-per-watt-to-reactor-count input is missing from the capital and policy story.
Where the price sits today	No priced channel. No sell-side coverage of night-survival mass-per-watt as an investable constraint. No futures or derivative on fission unit delivery risk. The NASA/DOE program shows up in trade press (Aviation Week, ANS, Lockheed Martin feature) as a program milestone, not as a binding economic chokepoint. Narrative-obscure and unpriced. The not-yet-priced test passes.
The binding constraint	Night-survival energy/power mass per continuous watt across the 354-hour lunar night (kg per continuous-watt for the storage route; equivalently, the count of qualified, flight-ready fission surface power units for the only route whose mass does not scale with night length). Not "lunar power" as a theme. The specific inelastic input is bridge-the-night mass-per-watt, and the single-sourced reactor program that escapes it.
What we are watching	(a) Count of fission surface power units that pass qualification and are manifested for lunar delivery (today: zero flown; one US program targeting Q1 2030 after slipping from 2026; design spec >=100 kWe, under 6 tonnes). (b) Best demonstrated night-survival energy-storage specific energy for a flight-qualified lunar system in Wh/kg including thermal keep-alive (today effectively battery-only, roughly 100-150 Wh/kg derated). (c) Count of non-polar lunar surface assets demonstrating continuous operation through a full lunar night at >=1 kW. Leading tell: any lander or base program that names night-survival power (not landing mass or launch cadence) as the schedule-gating line item, and any reactor-program slip past 2030.
What would prove us wrong	Killed if, by 2038-12-31, either (a) two or more independent lunar surface assets demonstrate continuous >=1 kW operation through a full lunar night without a fission reactor -- requiring a storage or beamed-power breakthrough crossing roughly 400 Wh/kg flight-qualified night-survival specific energy, or operational cislunar/orbital power beaming delivering kilowatt-class night power -- OR (b) three or more qualified fission surface power units are delivered to the lunar surface on schedule, dissolving the single-source bottleneck. Either condition means the night-survival mass-per-watt constraint stopped binding.

How we tried to break it

We ran three adversarial routes; all fail. First, polar peaks of eternal light make the argument moot. They do not. The permanently illuminated areas total a few square kilometers, are topographically constrained, and are contested by every lunar program at once; they cannot scale to a surface economy. Second, commercial nuclear enthusiasm (Oklo, Astral, X-energy) could spawn competing designs quickly. It cannot at the needed pace. Space qualification is a decade-long gauntlet, no commercial microreactor design is on a credible path to flight qualification by 2030, and the NASA/DOE MOU itself is resetting to a fresh downselect. Third, the dated call might fail because the kill events are too demanding. That is the design, not a flaw. The call resolves DEMOTE only if a genuine non-fission breakthrough or multi-reactor delivery dissolves the bottleneck, both low-probability on the 2038 horizon. The structural argument survives all three.

Why we are making the call

We think the physics is right and invariant, the constraint is absent from priced channels, and the adversarial routes fail. The single-source fragility of the fission program is real and documented by its own schedule history. The kill line is tight and appropriately demanding. The 12-year window leaves real uncertainty on a storage breakthrough (hence the dated call at 0.70 rather than 0.85), but the case for the mechanism is strong (conviction 0.82).

If the call is right

If kilowatt-class continuous power forces fission, the rent concentrates into the two teams that survive the US Fission Surface Power downselect (Phase 1 completed Feb 2026), because a downselect-to-two award is the entire addressable supply through the decade. The cleanest captures sit with the nuclear-fuel and reactor-core IP holders behind those teams, BWX Technologies (BWXT) and Cameco (CCJ) via its ~49% of Westinghouse, rather than with the lander primes. Every non-polar surface watt above a few hundred is gated by the reactor count, not by lander cadence.

Who gains

BWX Technologies (BWXT): nuclear fuel and reactor-core IP, FSP design subcontractor on the Lockheed Martin team and the dominant US naval-reactor fuel house; wins fuel and core rent on more than one downselect outcome.

Cameco (CCJ): owns ~49% of Westinghouse (eVinci heat-pipe microreactor) plus the uranium fuel cycle; captures rent if the Westinghouse/Aerojet team advances.

Lockheed Martin (LMT) and Intuitive Machines (LUNR): LMT as prime on one downselect team; LUNR is the high-beta lottery ticket on the IX (Intuitive Machines + X-energy) team.

Who loses

Astrobotic (LunaGrid solar-plus-wireless-power) and Volta Space Technologies (beamed power on Firefly Blue Ghost M2): bet the night-survival market is solved by solar arrays, batteries, and wireless transfer; stranded above a few hundred watts continuous if fission becomes mandatory.

Non-polar solar-plus-battery CLPS landers (Firefly (FLY), Astrobotic Peregrine-class, Intuitive Machines as a lander): structurally capped at the 354-hour night, roughly 14-day life, no continuous kilowatt-scale operation.

Oklo (OKLO): cited as a commercial microreactor name but is not on an FSP team; a narrative-correlated beta that does not capture the surface-power rent.

What reprices

The tradeable expression is concentration into BWXT and CCJ on a downselect-to-two award, with LMT and LUNR as the team-level proxies; FLY reprices down on the night-gated lander side. The night-survival mass-per-watt constraint itself has no futures or derivative; reactor-program slip risk shows up only as program milestones.

The next constraint it creates

Once fission is mandatory, the binding constraint moves to space-qualification throughput of the reactor units (a decade-long gauntlet, single-digit count this decade) and to HALEU/TRISO fuel supply behind BWXT and X-energy, then to the radioisotope and cladding supply chain that powers everything below reactor scale (see P5).

Earliest sign it has begun

Any lander or base program that names night-survival power (not landing mass or launch cadence) as the schedule-gating line item, the final FSP downselect-to-two award, and any reactor-program slip past the Q1 2030 target.

P3

By 2034, deep-space-band large-aperture antenna-hours are the binding supply constraint on missions past GEO, with demand already exceeding DSN capacity at peak and the deficit widening...

Domain: space

2034-06-30

Structural case	Our call, dated	Resolves
78%	62%	2034-06-30

Closing a link from the Moon, Mars, or beyond requires very large apertures (34 m and 70 m class) with cryo-cooled receivers, because received signal power falls with the square of distance. The dense commercial LEO ground-station fleets (AWS Ground Station, Leaf, RBC Signals) cannot serve this: their small dishes cannot close the link budget, and they run on commercial frequency allocations, not the ITU-protected deep-space bands. The protected deep-space allocations at S (roughly 2.2 GHz), X (8.4 GHz), and Ka (32 GHz) are fixed by international treaty and cannot be auctioned, expanded, or refarmed. The three 70 m dishes (Goldstone DSS-14, 1966; Canberra DSS-43, 1972; Madrid DSS-63, 1974) are single points of failure for the most distant links, cannot be upgraded to Ka-band, and each goes dark for years during overhaul. New large dishes take roughly a decade from funding to operation. Planned lunar relay constellations close the cislunar data relay tier but do not replace the ground aperture needed to close links to anything past the Moon, or to receive data from cislunar relays themselves. Supply is jointly inelastic in steel, spectrum, and trained operations staff.

The boom

Cheap reusable launch and the Artemis crewed-lunar program drive a steep rise in deep-space mission count through the late 2020s and early 2030s. Each crewed Artemis flight competes directly for the same S/X/Ka antennas as science flagships (Mars relays, outer-planet orbiters, JWST commanding) during the same windows, stacking hard contention spikes on top of a baseline that already exceeds capacity at peak by roughly 40%. The Goldstone 70 m dish (DSS-14) has been offline since September 2025 with no confirmed return date as of mid-2026, a live demonstration that supply has already failed in exactly the mode this call describes.

Why it is not priced yet

Space finance and the commercial space community price LEO downlink and ground-station-as-a-service. Deep-space tracking is treated as a NASA overhead problem, invisible to equity markets, with no traded proxy and no sell-side coverage. The market assumes we will build more antennas and lunar relay constellations fix it, missing that: (i) new large-aperture ground assets take a decade-plus to field and are not funded at the required scale; (ii) the protected spectrum bands are fixed by international treaty; (iii) the largest existing dishes are aging single points of failure already offline; and (iv) lunar relay constellations address cislunar data relay but not the ground-aperture constraint for anything past the Moon. The constraint is documented only in OIG audits and agency planning slides, in no instrument the market prices.

Where the price sits today	No equity proxy, no futures market, no sell-side coverage for deep-space aperture hours. The adjacent LEO ground-station-as-a-service market (the financializable layer) takes all the analyst attention. Deep-space tracking capacity is priced nowhere and traded nowhere. Not already priced.
The binding constraint	Large-aperture (34 m / 70 m class) antenna-hours on the ITU-protected deep-space S/X/Ka bands, jointly with those coordinated spectrum allocations and the trained operations staff needed to run them. Not LEO ground stations. Not relay satellites. Not launch slots.
What we are watching	(a) DSN demand-over-supply ratio: already reported at roughly 140% of capacity at peak by the 2023 NASA OIG audit, with the 2030s projection at 150% or above before new aperture comes online. (b) Count of missions reporting deferred, shortened, or descoped contact passes citing antenna contention rather than budget. (c) Number of operational 34/70 m-class deep-space-band apertures and their scheduled downtime: DSS-14 offline since September 2025 with no confirmed return. (d) ITU coordination filings for deep-space X/Ka as cislunar users multiply. Baseline state today: network already oversubscribed, the largest single dish offline for an indeterminate overhaul, demand curve steepening, and planned mitigations (LEGS cislunar antennas) acknowledged insufficient for the deep-space tier.
What would prove us wrong	The call is dead if by 2034: (a) a fielded constellation of new large-aperture or phased-array deep-space ground assets multiplies effective deep-space aperture-hours enough that demand-over-supply falls durably below 1.0 with no mission citing tracking time as a limiter; OR (b) optical (laser) deep-space communication is qualified and adopted as the primary link modality for the bulk of deep-space data volume, bypassing the RF aperture and protected-band constraint entirely; OR (c) the lunar and deep-space mission manifest collapses to the point that demand falls back to a level the existing network can absorb without rationing. Any one dissolves the needle.

How we tried to break it

The strongest kill is optical deep-space communication. NASA's DSOC experiment on the Psyche spacecraft proved laser links are technically feasible at deep-space distances. If optical is adopted at scale by 2030, the RF aperture constraint is partly bypassed. We push back: DSOC is a single experiment, optical ground terminals carry weather and cloud-cover limitations, and qualifying optical as the primary link for the bulk of deep-space data by 2034 needs a funded program-of-record adoption pace that does not exist in NASA's budget today. The kill criterion correctly requires adoption for the bulk of deep-space data, not mere demonstration. The second attack is Artemis demand collapse. If Artemis slips or is cancelled, the multiplicative demand growth weakens. But the base oversubscription at peak already exists from standing science missions regardless of Artemis, so the deficit survives partial manifest collapse, with the severity dampened. The call survives both vectors at the structural level. Our honest uncertainty is whether the constraint becomes visibly binding on the economy in the legible way the dated call implies, versus staying a government-internal rationing problem that mission planners absorb quietly.

Why we are making the call

We think the mechanism is right and well-documented. Inelasticity is real on three independent axes (steel, spectrum, staffing) with no funded path to elastic response on the 2024-2034 horizon. The market is genuinely blind to this layer. The constraint is already active today (DSS-14 offline, peak oversubscription confirmed by OIG). The kill criteria are specific and testable. We put the dated call below our conviction because the Artemis demand surge may slip and reduce severity, optical comms are moving faster than zero, and the exact binding-constraint-on-the-economy framing needs the consequence to be legible, not just buried in government audits.

If the call is right

If deep-space-band antenna-hours bind, rent flows to the one entity converting the DSN capacity shortfall into a commercial product. Intuitive Machines (LUNR) is the clean play: it holds the NASA Near Space Network relay IDIQ (max $4.82B, option period to Sept 30 2034) and is acquiring Goonhilly Earth Station plus COMSAT (close ~Q3 2026), folding in real deep-space dishes (GHY-6 32m, cryo GHY-3 30m, S/X/Ka). With DSS-14 folded into extended maintenance through roughly Oct 2028 and DSN demand already ~140% of peak capacity, missions either pay for commercial aperture or get rationed.

Who gains

Intuitive Machines (LUNR): captures both the cislunar relay tier (NSN IDIQ) and a commercial deep-space aperture (Goonhilly) in one entity as DSN saturates.

Goonhilly Earth Station (inside LUNR) and ESA Estrack (New Norcia DSA-4 35m, nominal ops Apr 2026): the few non-NASA apertures that can close a deep-space link budget on protected bands gain scheduling leverage.

KSAT (Kongsberg, OSE: KOG as parent) and SSC Space (state-owned, won ispace M3): build lunar and deep-space X/Ka networks that sell scarce contact hours; not directly tradeable except via KOG.

Who loses

NASA science flagships dependent on 70m-class downlink (Voyager 1/2 on DSS-43 for command uplink, JWST commanding, Mars relays, Europa Clipper, outer-planet orbiters): lose contact passes to Artemis crewed priority during contention windows; the mechanism is deferred and descoped passes, not a ticker.

Secondary and university deep-space missions without priority allocation: crowded off the schedule first.

NASA JPL DSN as an asset base: aging single-string 70m dishes (DSS-14, DSS-43, DSS-63) are the bottleneck, not a rent capturer, and stay chronically underfunded against a steepening demand curve.

What reprices

Intuitive Machines (LUNR) is the only clean public instrument that moves on cislunar and deep-space aperture-rent capture; Viasat (VSAT) touches near-space comms but is diluted. The constraint itself (deep-space aperture-hours) prices nowhere: no futures, no spot, no sell-side coverage.

The next constraint it creates

If commercial RF aperture relieves the steel constraint, the binding layer moves to the ITU-protected S/X/Ka spectrum coordination (treaty-fixed, cannot be auctioned or expanded) and to trained deep-space operations staff. If optical scales, it moves to weather-diverse optical ground-terminal networks and clear-sky site availability.

Earliest sign it has begun

Missions reporting deferred, shortened, or descoped contact passes citing antenna contention rather than budget; the confirmed DSS-14 return date (currently no return before ~Oct 2028); and the close of the Intuitive Machines / Goonhilly acquisition (~Q3 2026).

P4

By 2035, the binding constraint on operating proliferated LEO is not tracking sensors or autonomous-maneuver software but the accuracy of thermospheric neutral-density forecasts, and the...

Domain: space

2035-12-31

Structural case	Our call, dated	Resolves
72%	52%	2035-12-31

Drag physics in LEO is fixed: density forecast error propagates directly into along-track position uncertainty, which sets conjunction false-positive and false-negative rates. The number of required conjunction screens grows super-linearly with constellation size, while the in-situ thermospheric measurement base stays fixed and aging. Sensors that track where objects are do not measure the atmosphere that determines where they will be. Ground-tracking capex or maneuvering software cannot relax the constraint, only new in-situ measurement assets can, and those are slow, capital-gated, and currently non-commercial. One caveat weakens but does not kill the claim: a large-enough constellation operator (Starlink scale) can derive density from its own GNSS-accelerometry at near-real-time latency, potentially closing the gap internally. That makes the rent-migration story apply most cleanly to smaller third-party operators rather than universally.

The boom

Cheap launch and proliferated LEO are read as a traffic-management problem you solve with more radars or optical trackers plus onboard autonomous collision avoidance. The unpriced second-order fact: position uncertainty in LEO is dominated by atmospheric drag, and drag is set by thermospheric neutral density, which the operational community cannot forecast 24-72h out to better than roughly 20-40 percent (worse during geomagnetic storms). As object count and maneuver rate explode (Starlink running roughly 1,000 avoidance maneuvers per day, about 144,000 in six months Dec 2024 to May 2025), the system does not choke on tracking or compute. It chokes on density-forecast error producing a flood of false-positive conjunctions (wasted propellant, wasted operator decisions) and a residual of missed true ones. The physical input that fixes this is assimilated near-real-time thermospheric density driven by in-situ measurements: solar-EUV irradiance monitors and direct upper-atmosphere density sensors (the accelerometer lineage of CHAMP/GRACE/GRACE-FO/SWARM, now largely aged out, leaving the documented thermospheric gap). That driver data comes from a handful of single-string, government-budget instruments with no commercial supply chain. Value and regulatory leverage move to whoever can fly a commercial accelerometer and EUV density-sensing layer and sell assimilated forecasts, because the entire autonomous-collision-avoidance investment thesis sits physically downstream of this one inelastic measurement input.

Why it is not priced yet	The market frames LEO traffic as a tracking and coordination problem: build more radars and optical sensors, stand up space-traffic-coordination APIs, automate maneuvers. The space-weather community separately documents the thermospheric gap in academic venues (PMC surveys, EGU abstracts, AGU papers), but no constellation operator, SSA vendor, or SSA investor prices thermospheric density driver data as a binding constraint or a scarce investable asset. The only commercial entrant visible as of mid-2026 is Ensemble Space Lab, a startup at pitch-competition stage with a web API, not a flying sensor constellation. The shift from tracking sensors to atmosphere-measuring sensors is absent from any market framing we found in equity coverage or sell-side research. The not-yet-priced check passes.
Where the price sits today	No spot price, no equity coverage, no commodity market for thermospheric density data. HASDM (the operationally accurate U.S. government model) is restricted to government use and not purchasable. Ensemble Space Lab is pre-revenue pitch stage. The price channel is empty, which confirms the call is not yet priced.
The binding constraint	Real-time in-situ thermospheric neutral-density driver data (direct accelerometer-derived density plus solar-EUV irradiance) feeding assimilated short-horizon density forecasts. The specific scarce asset is the measurement base: GOES SUVI-class EUV monitors and the retired CHAMP/GRACE-FO/SWARM accelerometer density lineage with zero commercial successor in qualified operational supply as of 2026.
What we are watching	Operational thermospheric density forecast skill at 24-72h horizons (percent RMS error vs. assimilated truth) and the count of false-alert collision-avoidance maneuvers per active satellite per year attributable to density uncertainty. Baseline: density forecast error roughly 20-40 percent (much worse during storms); Starlink maneuver count climbing super-linearly. Secondary metrics: number of flying in-situ thermospheric density sensors (near zero post-SWARM), and any commercial density-data offtake contract signed by a major constellation operator.
What would prove us wrong	Kill the call if by 2035 any of the following: (a) operational 72h thermospheric density forecast error falls below roughly 10 percent through model or assimilation advances WITHOUT a new in-situ measurement layer, meaning the constraint dissolved by software or data-assimilation and density data never became the named gate; (b) autonomous maneuvering at full constellation scale proceeds with operators not citing density-forecast uncertainty as a material driver of maneuver burden or screening cost; or (c) a commercial in-situ thermospheric density-sensing constellation reaches qualified operational supply and prices the data as an ordinary commodity, meaning the constraint is relieved and the rent accrues only briefly. Also kill if the constraint resolves via a new government-funded replacement mission (a SWARM follow-on) that keeps supply public and non-commercial, negating the rent-migration thesis even if the physical gap closes.

How we tried to break it

Three adversarial angles carry partial force but do not kill the call. First, the self-measurement loophole: SpaceX can derive near-real-time density from Starlink GNSS-accelerometry (maximum 8h lag per published research), potentially closing the gap internally for the largest LEO operator without buying data. So rent may not migrate to a third-party sensor provider for Starlink, only for smaller operators. Second, the software-assimilation path: steady improvement in proxy-driven assimilation (solar indices, geomagnetic indices, ML) may narrow the forecast error gap without new hardware, though the kill criterion at 10 percent RMS error is a high bar that proxy-only assimilation has not approached. Third, a government replacement mission (a SWARM successor, a NASA small-satellite density program) could fill the sensor gap and keep the data public, relieving the commercial rent hypothesis while confirming the physical constraint. The physics argument survives all three; the rent-migration and the timing test carry genuine uncertainty.

Why we are making the call

We think the mechanism is sound and well-documented: density forecast error is the dominant LEO position-uncertainty driver, the in-situ measurement base is genuinely thin and aging, and no commercial supply exists today. The call is not yet priced by a clear margin. Our adversarial search turned up one pitch-stage startup (Ensemble Space Lab) and one academic research path (Starlink GNSS-accelerometry), neither of which qualifies as commercial supply. The weak points are the self-measurement loophole for mega-constellation operators and the chance that a government mission fills the gap without creating commercial rent. Those pull the dated call to 0.52 (timing and exact resolution mechanism are uncertain) while leaving our conviction at 0.72 (the structural constraint is real and likely to stay binding through at least the early 2030s). We promote it with the kill criteria intact.

If the call is right

If thermospheric density-forecast error is the binding LEO constraint, the value splits between the operator that self-supplies and the vendor that productizes assimilated forecasts. SpaceX/Starlink internalizes the rent by deriving fleet-wide density from its own GNSS-accelerometry at a cadence (~1,000 maneuvers/day) no agency matches, turning the constraint into a moat. The named commercial layer is Space Environment Technologies (SET), holder of a NASA SBIR for an operational density nowcast/forecast tool and the closest private HASDM successor. The upstream EUV and solar-wind feed (NOAA SOLAR-1 at L1, GOES SUVI) is government-funded and free at point of use, which caps how much rent any private layer can extract.

Who gains

SpaceX / Starlink: self-supplier; harvests near-real-time density from its own fleet ephemeris (~8h lag), avoids buying data, and out-screens smaller operators on the same physics.

Space Environment Technologies (SET): named commercial density-assimilation play (NASA SBIR) selling the processed nowcast/forecast product smaller operators cannot self-generate.

Whoever first productizes ML-density-from-EUV-imagery or density-from-debris-drag: captures the in-situ measurement gap without launching new accelerometer hardware.

Who loses

Position-only SSA and conjunction vendors (LeoLabs, Slingshot Aerospace, COMSPOC, Kayhan Space, SpaceNav, the five TraCSS Consolidated Pathfinder firms): screen geometry, not atmosphere; their false-positive rate is bounded by the density error they consume, so they are downstream price-takers of the density layer.

Smaller LEO operators (Eutelsat OneWeb, Planet, Spire, secondary cubesat fleets): pay the tax directly in propellant burned on false-alarm maneuvers and shortened life, lacking Starlink-scale data to self-correct.

Amazon Project Kuiper while it scales: large fleet without Starlink-grade accelerometry density inversion, so it eats the false-alert burden until it builds the same data layer.

What reprices

Nothing prices this cleanly. No spot market or equity proxy for thermospheric density data; HASDM is US Space Force-restricted and not purchasable. The closest public exposure is indirect, via a loser-side operator like Eutelsat (EPA: ETL) paying the LEO tax; SET, Kayhan, LeoLabs, and Slingshot are all private.

The next constraint it creates

Once density data binds, the constraint moves to qualified in-orbit accelerometer and EUV sensor supply (near zero post-SWARM, no commercial successor) and to data-assimilation talent. A government replacement mission (NASA Geospace Dynamics Constellation, a SWARM follow-on) would fill the physical gap while keeping supply public, relieving the commercial rent rather than transferring it.

Earliest sign it has begun

Operational 72h density-forecast error and the count of false-alert avoidance maneuvers per active satellite per year; the first density-data offtake contract signed by a major constellation operator; and any funded NASA GDC or SWARM-successor start that keeps the data public.

P5

By 2035, iridium DOP-26 clad-vent-set fabrication at ORNL, not plutonium-238 fuel rate, becomes the named gating constraint on US deep-space and lunar-surface nuclear power

Domain: space

2035-12-31

Structural case	Our call, dated	Resolves
72%	38%	2035-12-31

Iridium is a trace byproduct of platinum mining with a hard physical ceiling near 7 tonnes per year. No iridium-primary mine exists; expanding supply requires a new platinum project and a decade of lead time. South African PGM mines were cutting unprofitable output in 2024 and 2025, not expanding. DOP-26 clad-vent-set fabrication is a reconstituted single-site capability at ORNL with no qualified second source and no certified substitute material. PEM electrolyzer demand is a structural, policy-driven (net-zero mandate) draw on the same metal that does not retreat on a mission-planning horizon. This is a constraint-migration cascade one layer below the visible plutonium story, plus a cross-spine collision between space nuclear and the hydrogen buildout on a shared inelastic input. The collision is physically real and bureaucratically invisible today.

The boom

The plutonium-238 production ramp at Oak Ridge absorbs every headline and every GAO citation on why NASA cannot fly more outer-planet and lunar-surface missions. The metal that physically wraps every gram of that fuel never enters the discussion. Each Pu-238 pellet sits inside a cup of DOP-26 iridium alloy (iridium with 0.3 to 0.5 percent tungsten plus trace aluminum and thorium), formed into a clad vent set. That cup does the safety job that makes the source flyable: it contains the fuel through a launch-pad explosion or reentry impact and stops alpha-bred neutrons from escaping. There is no qualified substitute material and no second fabrication source. The clad-vent-set line at ORNL nearly lapsed and had to be reconstituted from scratch. Now add the upstream reality confirmed by live price data: iridium spot has risen 70% year-to-date in 2026 and 426% since 2020, driven by PEM hydrogen electrolysis demand already at 12% of world industrial consumption (up from 4% in 2021), with the hydrogen literature projecting 32 to 40 tonnes per year of PEM demand by 2030 against a hard physical ceiling of roughly 7 tonnes per year of global production, itself a byproduct of platinum mining that takes a decade to expand and was cutting output in 2024 and 2025. The space program draws only kilograms, but it draws a specific arc-remelted, drop-cast, single-crystal-controlled alloy from one reconstituted shop, and it now bids for the same metal against an energy-transition buyer willing to absorb any price. The binding constraint on deep-space and Artemis-era lunar power moves off the fuel and onto the cladding alloy and the sole fabrication line that forms it into clad vent sets. By 2035 a flagship outer-planet or lunar-surface mission slips or descopes its RPS unit count with iridium DOP-26 clad-vent-set throughput, or the iridium feedstock itself, named as a cause in a program-status or supply-chain document.

Why it is not priced yet	The iridium-for-PEM squeeze now sits in PGM sell-side notes, hydrogen market reports, and visible spot moves. That first-order claim is priced. The genuinely unpriced layer is the collision: that the same 7 t/yr stream is the sole cladding metal for every US space radioisotope power source, fabricated at a single reconstituted ORNL line, making the space-nuclear program an unhedged price-taker behind the hydrogen buildout. No space-sector analyst, no NASA RPS supply narrative, and no program-status document in public view joins these two demand streams or names iridium DOP-26 cladding as the gating node under the deep-space and lunar power story. The consensus space bottleneck is Pu-238 fuel rate; the cladding metal and its sole fabrication line are invisible in that discussion.
Where the price sits today	Iridium spot confirmed at roughly $278/gram as of April 2026, up 70% year-to-date in 2026 and 426% since January 2020. PEM electrolyzer demand growth is reflected in this price. The space-nuclear cladding collision is NOT reflected in any space program pricing or risk documentation we found, confirming the second-order claim remains unpriced.
The binding constraint	Iridium DOP-26 alloy clad-vent-set fabrication capacity at ORNL, and the high-purity iridium feedstock underpinning it, not plutonium-238 fuel rate and not the RTG or RPS system as a whole
What we are watching	(1) Iridium spot price and the PEM-electrolyzer share of world iridium consumption (2021: 4%; 2025: 12%; 2030 PEM demand literature: 32-40 t vs. 7 t/yr supply ceiling). Current live anchor: iridium at roughly $278/gram as of April 2026, up 70% year-to-date. (2) Any NASA, DOE, or GAO program-status or supply-chain document that names iridium DOP-26 or clad-vent-set fabrication throughput (as distinct from Pu-238 fuel production rate) as a schedule or unit-count limiter on a deep-space or lunar-surface mission. Baseline today: zero such citations exist in public program documentation.
What would prove us wrong	By 2035-12-31, either (a) no NASA, DOE, or GAO program document names iridium, DOP-26, or clad-vent-set fabrication (as opposed to Pu-238 fuel supply) as a limiter on any deep-space or lunar-surface RPS-powered mission, AND iridium real price stays within 30% of its 2025 level; or (b) a qualified second-source iridium CVS fabrication line stands up at a facility other than ORNL; or (c) PEM electrolyzer iridium catalyst loading falls below 0.1 mg per square centimeter at commercial scale, or a non-iridium anode catalyst reaches commercial deployment at gigawatt scale, eliminating the PEM demand pressure on world iridium supply. Any of these dissolves the squeeze that drives the call.

How we tried to break it

Three real attacks. First: the space program uses kilograms, not tonnes, so it could simply pay whatever price clears and never face a physical metal shortage. The claim therefore requires that the ORNL fabrication line itself (skills, tooling, throughput rate, not raw metal cost) becomes the named bottleneck, a more specific and harder-to-document claim than iridium being expensive. Second: PEM catalyst thrifting is advancing. Published loading reductions from 0.5 mg/cm2 toward under 0.1 mg/cm2, and ruthenium-based or mixed-oxide anodes at pilot scale, could collapse PEM iridium demand well before 2030 and dissolve the supply-squeeze driver. This is the most credible kill path. Third: even if the physical constraint is real, government program offices often absorb constraints quietly through budget workarounds or mission de-scopes attributed to other causes. The dated call resolves only if a public document explicitly names iridium DOP-26 or CVS fabrication as a mission limiter, which may never happen in that form even if the underlying constraint is genuine. The case survives attacks one and three in substance; attack two (catalyst thrifting) is the live uncertainty that most moves the dated call below our conviction.

Why we are making the call

We promote it. The mechanism is sound, the needle sits correctly at the fabrication line and feedstock rather than the system level, the upstream supply ceiling is confirmed by live price data, and the specific collision claim (space-nuclear cladding competing with the hydrogen buildout on the same 7 t/yr stream) is absent from all space-sector and program-status discourse. Our conviction is 0.72 because the case is well-grounded but catalyst thrifting is a live kill path. The dated call is 0.38 because the resolution criterion needs an explicit public naming in a government document over a nine-year horizon, and both silent absorption and thrifting-driven dissolution are plausible ways the physical reality never surfaces in that specific form.

If the call is right

If iridium DOP-26 clad-vent-set throughput, not Pu-238 fuel rate, becomes the named gate, the metal-side rent lands on the South African PGM producers and refiners that sit on a fixed ~7 t/yr byproduct stream they cannot scale, while the fabrication-side rent stays a non-tradeable strategic asset of DOE/ORNL, which holds the sole reconstituted clad-vent-set line and the Ir-0.3W-Th alloy process. Iridium spot at roughly $278/g (April 2026, +70% YTD, +426% since 2020) is being set by the PEM hydrogen buildout (12% of world iridium in 2025, up from 4% in 2021), so the US space-nuclear program bids for kilograms behind an energy-transition buyer willing to absorb any price.

Who gains

Valterra Platinum (JSE: VAL, ex-Anglo American Platinum), Sibanye-Stillwater (SBSW), and Impala Platinum (JSE: IMP): South African PGM producers for whom iridium is near-pure byproduct margin; the 2024-25 shaft cuts tightened supply directly.

Heraeus (private) and Johnson Matthey (LSE: JMAT): PGM refiners and recyclers that control high-purity iridium and the only meaningful secondary loop, pricing the spread.

DOE / ORNL: holds the sole qualified clad-vent-set fabrication capability and alloy IP; a strategic, non-tradeable winner whose throughput becomes the named line item.

Who loses

NASA RPS-powered missions, specifically Dragonfly (Titan, MMRTG, ~2028) and future outer-planet and lunar-surface RPS missions: unhedged price- and throughput-takers gated by the single ORNL line, exposed to slip or RPS-unit descope.

The US space-nuclear program broadly: an unhedged price-taker on a metal whose price is set by the hydrogen buildout, drawing kilograms it cannot substitute.

PEM electrolyzer makers (Plug Power (PLUG), Nel ASA (NEL), ITM Power (ITM), Siemens Energy, Cummins/Accelera, Ohmium): margin-squeezed by iridium cost, which is why they fund the catalyst-thrifting kill path.

What reprices

Iridium spot is the cleanest expression (~$278/g April 2026, no liquid futures, dealer-quoted by Heraeus and Johnson Matthey), so the proxy is physical metal or PGM equities: Valterra (VAL.JSE), Sibanye-Stillwater (SBSW), Implats (IMP.JSE), and Johnson Matthey (JMAT.L) for the refiner spread. The space-nuclear loser side is not cleanly shortable; NASA RPS risk surfaces only as mission slips.

The next constraint it creates

If the metal stays available at price, the constraint moves to the ORNL fabrication line itself: arc-remelt and drop-cast throughput, tooling, and the qualified machinists who form clad vent sets, with no second source. If catalyst thrifting (IrOx loadings below 0.1 mg/cm2, validated at 90-95% reduction in Plug-backed research) deploys at GW scale, the demand-spike leg softens and the squeeze relaxes.

Earliest sign it has begun

Any NASA, DOE, or GAO program document that names iridium, DOP-26, or clad-vent-set fabrication throughput (as distinct from Pu-238 fuel rate) as a mission limiter, plus the PEM-electrolyzer share of world iridium consumption crossing toward the projected 32-40 t/yr against a ~7 t/yr ceiling.

P6

By 2038, the binding constraint on proliferated-LEO throughput is the stratosphere's finite alumina-absorption capacity, forcing a regulated migration to non-aluminum demisable structures...

Domain: space

2038-12-31

Structural case	Our call, dated	Resolves
72%	52%	2038-12-31

Conservation of mass: in a cheap-launch regime with short orbital lifetimes, reentry mass flux scales directly with active-fleet size, and aluminum is roughly 30 percent of typical bus mass. Al2O3 from demise is chemically stable and settles from the stratosphere on a ~30-year timescale, so injection above the clearing rate accumulates monotonically. A finite slow-clearing sink plus a monotonically rising source is a textbook binding cap. Cheaper launch makes it worse, not better, by raising fleet size and replacement cadence. The only release valves are (a) non-aluminum demisable structures or (b) controlled reentry to ocean, both qualification-gated and slow. The 2026 Barker et al. paper (Earth's Future, AGU) now gives regulators the modeling basis; NOAA CSL issued a 2025 press release warning of stratospheric alteration within 15 years; an FAA public docket (FAA-2024-1395) already exists on ozone depletion from satellite demise. The physical chain is confirmed in peer-reviewed literature.

The boom

Cheap launch and short orbital lifetimes turn every megaconstellation into a mass conveyor: all launched mass must eventually reenter and burn. The dominant structural material is aluminum, and its combustion product, aluminum oxide, is chemically stable, settles from the stratosphere on a roughly 30-year timescale, and activates chlorine catalytically (ozone loss) while adding a persistent aerosol load. Reentry alumina already exceeded natural atmospheric aluminum by about 30 percent in 2022; full build-out of announced constellations implies roughly 360 metric tons per year of Al2O3 by the late 2030s. The stratosphere behaves as a finite sink with a slow clearance rate and a monotonically rising source, a textbook binding cap. Once the deposition signal is measurable and attributable (mid-2030s), it becomes a regulated quantity the way CFCs were. That converts an unpriced externality into a hard cap on annual alumina injection, and the rent moves to the inelastic input that keeps the conveyor running under the cap: spacecraft structures that fully demise without injecting alumina.

Why it is not priced yet	Reentry is narrated everywhere as a debris and ground-safety success story, with demisable design as the solution. Orbital sustainability talk fixates on Kessler dynamics and collision risk. The atmospheric-chemistry consequence of demise is now audible in geoscience (Barker et al. 2026 in Earth's Future, NOAA CSL 2025 press release, FAA-2024-1395 docket) but absent from constellation-operator design roadmaps, capital allocation, and SSA market framing. No sell-side coverage prices non-aluminum demisable structural material as a scarce input. The cross-domain chain, atmospheric planetary boundary to reentry-mass cap to structural-material rent, exists in zero market or regulatory framing as of mid-2026. The science is getting louder while the capital signal stays silent. That gap is the window.
Where the price sits today	No equity coverage of non-aluminum demisable structural materials as a category. No futures or derivatives market for reentry-alumina exposure. No licensing surcharge or deposition bond exists in any jurisdiction. FAA-2024-1395 is a public comment docket, not a rule. Science is audible; capital is silent. The unpriced gap is real.
The binding constraint	Qualified non-aluminum demisable spacecraft structural material at constellation scale, plus the stratospheric aluminum-deposition measurement and attribution capacity needed to set and enforce a reentry-mass cap. Candidates (magnesium alloys, novel polymer-composite hybrids) face 5-10 year space-qualification cycles for structural bus applications. There is no commodity market for this input; no constellation operator has committed to it; no sell-side coverage frames it as an investment category. The inelastic input is material qualification, not chips, launch, or bus electronics.
What we are watching	Measured stratospheric Al_2O_3 burden and annual reentry-injected alumina mass (tonnes/year); share of newly launched LEO mass certified to a low-alumina or non-aluminum demisable structural standard. Baseline: ~360 t/yr Al_2O_3 projected at full megaconstellation build; 2022 reentry already roughly +30 percent above natural atmospheric aluminum; effectively 0 percent of fleet on qualified non-aluminum demisable structures; no binding atmospheric-deposition licensing criterion exists anywhere. Track: first national or multilateral (FAA, EASA, ITU, UNOOSA/COPUOS) rulemaking that names reentry atmospheric deposition -- not ground casualty -- as a licensing criterion.
What would prove us wrong	Kill if by 2038: (a) no regulator has introduced reentry atmospheric-deposition or stratospheric-alumina limits as a licensing criterion AND no constellation operator commits to non-aluminum demisable structures citing this driver; OR (b) measured stratospheric alumina accumulation is shown empirically benign or below any actionable threshold through 2038 (the sink does not bind on relevant timescales); OR (c) mitigation migrates entirely to controlled de-orbit to ocean for the full small-LEO fleet WITHOUT repricing structural materials (constraint relieved via operations, not the named material needle); OR (d) reentry alumina is shown substitutable by an already-abundant, space-qualified material so no scarce qualified input emerges.

How we tried to break it

Three serious attacks. First: controlled de-orbit to ocean could relieve the constraint without repricing structural materials. This is the most credible kill path. We push back: for sub-200 kg LEO birds at constellation scale (thousands of vehicles), the fuel mass for a full deorbit burn competes directly with payload and bus mass budgets, making it uneconomical without a propulsion breakthrough. It is a live risk, not a closed one, and the kill criteria name it. Second: scientific attribution and regulatory timelines may slip past 2038. The Montreal Protocol ran 13 years from Rowland-Molina to binding rule, the closest historical analogy; the 2038 deadline is 12 years out and the FAA docket already exists, making the timeline plausible but not certain. Third: a cheap, already-space-qualified non-alumina structural material (magnesium alloy, for instance) could emerge quickly, making the input abundant rather than scarce. Possible, but magnesium alloys are not currently space-qualified for primary structure at constellation scale, and the 5-10 year qualification cycle makes a 2030s resolution realistic. The call survives all three with meaningful residual uncertainty, reflected in the dated call of 0.52.

Why we are making the call

We think peer-reviewed literature published in 2025-2026 confirms the mechanism. The unpriced gap is genuine: science is now audible while capital and regulatory pricing are silent. The named needle, qualified non-aluminum demisable structural material, is genuinely inelastic on a 5-10 year qualification cycle with no existing commodity market. The main competing resolution path (controlled de-orbit) is uneconomical at full constellation scale for small LEO birds and is correctly specified as a kill criterion. The 12-year horizon is long enough for the atmospheric signal to become attributable and for the first regulatory response to land, consistent with the Montreal Protocol precedent. The dated call sits at 0.52 rather than higher because the controlled-de-orbit mitigation path is real, and the exact resolution requires the constraint to bind via material qualification rather than operations, one path among several plausible outcomes.

If the call is right

If stratospheric alumina becomes a regulated quantity, the cost lands first on the largest reentry-mass injectors and the rent moves to whoever first qualifies non-aluminum demisable primary structure, a white space with no flight-qualified leader as of 2026. SpaceX Starlink is the dominant alumina source (>9,000 in orbit, ~5-year lifespans, Gen2 buses ~2 t each), so any per-satellite cap or deposition surcharge hits its replacement cadence hardest. The regulatory fight is live mainly at the FCC (the NEPA-exemption battle, operators versus 17 state attorneys general, after GAO told the FCC to reexamine its environmental review), not at the FAA, whose docket FAA-2024-1395 is actually the Starship-39A EIS.

Who gains

First firm to space-qualify non-aluminum demisable primary structure (magnesium alloys, polymer-composite hybrids): captures the qualification moat on a 5-10 year cycle; aerospace magnesium-alloy suppliers (Luxfer / Magnesium Elektron) are the watch-list but none is flight-qualified for primary structure today.

NOAA Chemical Sciences Laboratory and CIRES/University of Colorado: own the deposition-attribution science any regulated-quantity regime rests on, the regulatory gatekeeper position (not commercial rent).

Controlled-de-orbit and propulsion providers if mitigation shifts to ocean reentry: capture spend if operators avoid material re-qualification, though the thesis argues this is uneconomic for sub-200 kg birds at constellation scale.

Who loses

SpaceX Starlink: by far the largest alumina injector; a per-satellite alumina cap or deposition bond hits its replacement cadence hardest.

Amazon Project Kuiper, Eutelsat OneWeb, China Guowang and Qianfan/Spacesail (SSST): same aluminum-bus mass-conveyor model; Kuiper is already lobbying the FCC for a NEPA exemption, a tell that compliance cost is the live risk; foreign operators add an ITU and UNOOSA/COPUOS coordination gap.

Aluminum-bus incumbents (Airbus Defence and Space, Thales Alenia Space, Terran Orbital/Lockheed, York Space, and the in-house Starlink bus): forced into 5-10 year material re-qualification, since ESA's own Design-for-Demise guidance currently favors aluminum and would need re-basing.

What reprices

Nothing prices this cleanly. No equity proxy for non-aluminum demisable structure, no futures for reentry-alumina exposure, no licensing surcharge or deposition bond in any jurisdiction. Eutelsat (EPA: ETL) is the only listed pure-ish operator on the loser side; Luxfer (LXFR) is a thin, indirect magnesium proxy; the demisable-materials winner is not yet a public name.

The next constraint it creates

Once a cap binds, the constraint moves to qualified non-aluminum structural-material supply and the 5-10 year space-qualification throughput for primary bus structure, then to stratospheric deposition measurement and attribution capacity (NOAA CSL, ESA Cluster reentry calibration) needed to set and enforce the cap.

Earliest sign it has begun

The first national or multilateral rulemaking (FCC, FAA, EASA, ITU, UNOOSA/COPUOS) that names reentry atmospheric deposition or stratospheric alumina (not ground casualty) as a licensing criterion, the FCC outcome of the NEPA-exemption fight, and the first constellation operator committing to non-aluminum demisable structure citing this driver.

Seeds considered

These cleared the supply-side test but did not make the final board, usually because the trade was not clean or the move was already priced.

Seed	Physical case	Why not promoted
By 2035, radioisotope-power-source availability (Pu-238 heat-source supply) becomes the explicitly cited gating item for surviving the lunar night and distributed cislunar assets, with Pu-238 rationed against demand	Pu-238 single-line production ceiling (1.5 kg/yr, 4.8 kg per MMRTG, >$8M/kg) against distributed cislunar demand	Strong call, but it is the most central member of a five-way radioisotope cluster. P5 (iridium DOP-26 cladding) is the more specific and less duplicative pick because it sits one layer below the Pu-238 story everyone already cites, and the lunar-night demand driver here overlaps P2's night-survival mechanism. Dropped to avoid stacking the same input twice.
By 2037, AMPPEX-class Am-241 separation throughput at NNL/Sellafield becomes the binding constraint on Europe's sovereign cislunar and lunar-surface power program	Am-241 in-growth physics (14.35-yr Pu-241 half-life) plus single-site hot-cell separation ceiling (~8-10 kg/yr conceptual) against the headline 140 t plutonium stockpile	The mechanism is genuinely distinct (in-growth decay clock plus licensed hot-cell throughput) and the inversion is elegant, but it is a near-duplicate of the radioisotope-supply family, and the dated call depends on European lunar demand actually pressing the ceiling, which has a long slippage record. Lower defensible edge than the chosen six; held as the strongest radioisotope runner-up.
By 2035, the joint inelasticity of finite Np-237 feedstock and shared neutron-flux-hours at HFIR/ATR causes mission reschedules even while the 1.5 kg/yr production goal is nominally met	Non-replenishable Np-237 stockpile drawdown plus oversubscribed 1960s-vintage reactor flux shared with medical and defense isotopes	Sharpest double-inelasticity in the radioisotope cluster, but it sits upstream of P5's cladding needle and shares the same mission-class demand. The attribution problem (budget always takes the formal blame in 2035 documents) makes the dated call the weakest of the family. Near-duplicate, dropped.
By 2034, a US or multilateral regulator caps demisable-mass throughput or mandates design-for-non-demise, explicitly licensing reentry alumina footprint per replaced satellite	Same stratospheric alumina sink as P6, but resolving on a regulatory-cap event rather than a structural-material rent	Direct duplicate of P6's mechanism. P6 takes the slot because it names a cleaner rent-bearing needle (qualified non-aluminum demisable structure) and carries a dated call of 0.52 versus this call's 0.22, since the current US regulatory vector is running backward (FCC narrowing NEPA in 2025) and a multilateral treaty route likely overshoots 2034.
By 2035, radioisotope-power-source availability becomes the gating item for sustained lunar/cislunar presence, with Sr-90/Am-241 substitutes themselves separation-limited	Decay-heat kilograms from a single unexpandable federal line as the input you cannot launch around	Highest-level framing of the radioisotope thesis, which makes it the most theme-like and least specific of the cluster (qualified radioisotope heat-source kilograms rather than a single fabrication line or metal stream). Fission surface power (P2's escape valve) is its honest kill path. Subsumed by P5's sharper needle.

Each call is dated. The line that would prove it wrong is fixed when the board is issued.