

# Where the Space Economy's Rent Migrates: Six Physical Walls Cheap Launch Cannot Move

*When you make getting to orbit free, scarcity jumps to the inputs that obey thermodynamics, decay physics, and treaty-fixed spectrum, not learning curves.*

**Thesis spine:** Frontier → Capability → Dependency graph → Supply elasticity → Demand → Capital → Pricing → Policy → Outcomes. Rent accrues to the inelastic input. The edge is naming where the constraint moves before pricing catches up.

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

Method: generate wide and disruptive, then gate strict. Each call names the needle, not the theme.

Two probabilities per call: directional vision, and the strict dated clause scored at resolution.

Status: hardened candidates (survived the adversarial refute pass). Drafted 2026-06-14.

# The board: 6 hardened calls

---

## The cross-cutting read

Every one of these six calls follows the same migration logic and the consensus misses it the same way. 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 solves the whole stack. It does not. It relocates the binding constraint to a small set of inputs that do not ride any learning curve because they are pinned by a physical or institutional invariant: radiative heat rejection in vacuum 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 link budgets need treaty-protected S/X/Ka bands and 70 m steel that takes a decade to pour (P3); LEO position uncertainty is dominated by thermospheric density that a near-dead, non-commercial sensor base cannot forecast (P4); every US radioisotope heat source is clad in a single arc-remelted iridium alloy from one reconstituted ORNL line now bidding against the hydrogen buildout for the same 7-tonne-per-year metal stream (P5); and the stratosphere is a finite, 30-year-clearance alumina sink that a self-replacing megaconstellation fills monotonically (P6). In each case the needle is a specific inelastic input (square-meters of deployable radiator per kilowatt, kilograms of night-survival storage per watt, large-aperture antenna-hours, in-situ density-driver 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 now audible (cooling is hard, Pu-238 is short, reentry chemistry exists in geoscience papers); the quantitative, rent-bearing needles are not in a single equity model, sell-side note, or futures contract. That gap between an audible theme and an unpriced needle is the entire edge.

## At a glance

#	THE BOOM	BINDING CONSTRAINT (THE NEEDLE)	VISION P	CLAUSE P	RESOLVES
P1	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/m <sup>2</sup> . That is a physical constant, not an engineering target. Starcloud's own white paper concedes roughly 1,600 m <sup>2</sup> of radiator PER	Qualified deployable space-radiator AREA (m <sup>2</sup> of micrometeoroid-survivable, fluid-loop-coupled panel) per kilowatt of was	82%	62%	2034-12-31

MEGAWATT, and a 5 GW concept needs about 8 km<sup>2</sup> 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/m<sup>2</sup> that has barely moved in 40 years. Chip power density keeps climbing, but radiator W/m<sup>2</sup> 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.

P2	<p>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</p>	<p>Night-survival energy/power mass per continuous watt across the 354-hour lunar night (kg per continuous-watt for the sto</p>	82%	70%	2038-12-31
----	--	---	-----	-----	------------

#	THE BOOM	BINDING CONSTRAINT (THE NEEDLE)	VISION P	CLAUSE P	RESOLVES
---	----------	---------------------------------	----------	----------	----------

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 addressable set of continuously-powered lunar surface assets is therefore capped at the reactor count, regardless of how many landers fly.

#	THE BOOM	BINDING CONSTRAINT (THE NEEDLE)	VISION P	CLAUSE P	RESOLVES
P3	<p>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, creating 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 the call describes.</p>	<p>Large-aperture (34 m / 70 m class) antenna-hours on the ITU-protected deep-space S/X/Ka bands, jointly with those coordi</p>	78%	62%	2034-06-30
P4	<p>Cheap launch and proliferated LEO are read as a traffic-management problem solved by more radars or optical trackers and 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</p>	<p>Real-time in-situ thermospheric neutral-density driver data (direct accelerometer-derived density plus solar-EUV irradiation)</p>	72%	52%	2035-12-31

#	THE BOOM	BINDING CONSTRAINT (THE NEEDLE)	VISION P	CLAUSE P	RESOLVES
	<p>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"). This driver data is produced by a handful of single-string, government-budget-funded instruments with no commercial supply chain. Value and regulatory leverage migrate to whoever can fly a commercial accelerometer and EUV density-sensing layer and sell assimilated forecasts, because the entire autonomous-collision-avoidance investment thesis is physically downstream of this one inelastic measurement input.</p>				
P5	<p>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</p>	<p>Iridium DOP-26 alloy clad-vent-set fabrication capacity at ORNL, and the high-purity iridium feedstock underpinning it,</p>	72%	38%	2035-12-31

gram of that fuel is absent from 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 performs the safety function 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 layer on 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 is now bidding for the same metal against an energy-

#	THE BOOM	BINDING CONSTRAINT (THE NEEDLE)	VISION P	CLAUSE P	RESOLVES
---	----------	---------------------------------	----------	----------	----------

transition buyer willing to absorb any price. The binding constraint on deep-space and Artemis-era lunar power migrates 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.

P6	<p>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 contributing 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 Al<sub>2</sub>O<sub>3</sub> 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</p>	<p>Qualified non-aluminum demisable spacecraft structural material at constellation scale, plus the stratospheric aluminum-</p>	72%	52%	2038-12-31
----	---	---	-----	-----	------------

#	THE BOOM	BINDING CONSTRAINT (THE NEEDLE)	VISION P	CLAUSE P	RESOLVES
	(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 migrates to the inelastic input that keeps the conveyor running under the cap: spacecraft structures that fully demise without injecting alumina.				

Vision P = strength of the structural case. Clause P = calibrated odds the exact dated clause resolves true, scored with Brier. The gap is the honest timing and measurement tax, not timidity.

**P1 · Orbital compute does not hit a launch-cost wall or a power wall first. It hits a radiator-area wall. The binding constraint on space-based AI compute through the 2030s is deployable, micrometeoroid-survivable heat-rejection AREA per launched kilogram, governed by Stefan-Boltzmann T-to-the-fourth and a hard ~few-kg-per-square-meter panel-plus-fluid-loop mass floor that does not ride a learning curve. Multi-gigawatt orbital data centers are physically uncoverable this decade, and the rent migrates to whoever owns large-area deployable-radiator and two-phase loop technology, not to whoever owns the GPUs or the launch.**

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/m<sup>2</sup>. That is a physical constant, not an engineering target. Starcloud's own white paper concedes roughly 1,600 m<sup>2</sup> of radiator PER MEGAWATT, and a 5 GW concept needs about 8 km<sup>2</sup> 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/m<sup>2</sup> that has barely moved in 40 years. Chip power density keeps climbing, but radiator W/m<sup>2</sup> 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.

· Domain: space

DIRECTIONAL VISION  
**82%**

STRICT CLAUSE  
**62%**

RESOLVES  
**2034-12-31**

Heat-rejection in vacuum is radiative-only and obeys the Stefan-Boltzmann law: power rejected scales as emissivity times area times T-to-the-fourth. To radiate at high T you must run 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/m<sup>2</sup>. That fixes a hard floor on required area per kilowatt that no algorithm, no chip node, and no launch-cost collapse can erase. Worse, 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/m<sup>2</sup> and has barely moved in 40 years of active NASA/ESA/JAXA effort. ISS panels are roughly 3-4 kg/m<sup>2</sup>. 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 throwing cheaper launch or better GPUs at it does nothing.

**WHY IT IS PRE-CONSENSUS**

Through mid-2026 the consensus has only just adopted 'cooling is hard' as a THEME (WEF June 2026, SatNews 'physics wall' March 2026, IEEE Spectrum). That framing is now priced as a qualitative caveat. What

is NOT priced is the quantitative, structural claim: that deployable radiator AREA-per-kW and its kg/m<sup>2</sup> mass floor are a hard physical limit that caps orbital compute one-to-three orders of magnitude below the marketed GW scale through the 2030s, and that the rent therefore accrues to deployable-radiator and two-phase-loop suppliers rather than to GPU/launch players. Capital is still flowing to the compute-in-orbit story (Starcloud's \$170M Series A at \$1.1B valuation in March 2026) on the implicit assumption that cooling is a solvable detail. It is a thermodynamic invariant, and the area-per-kW needle is absent from the investment narrative.

#### HONEST PRICE CHANNEL

Theme-saturated as a qualitative concern (WEF, IEEE Spectrum, SatNews all ran cooling-wall pieces in 2026), but capital continues to flow to the compute/GPU narrative. Starcloud's \$170M Series A in March 2026 valued on compute density, not on demonstrated radiator technology. No dedicated deployable-radiator supplier is publicly traded or receiving named VC backing as a scarcity play. The quantitative claim (m<sup>2</sup>/kW floor as rent-bearing asset) is not in sell-side coverage or VC deck framing. Theme is priced; needle is not.

#### THE NEEDLE

Qualified deployable space-radiator AREA (m<sup>2</sup> of micrometeoroid-survivable, fluid-loop-coupled panel) per kilowatt of waste heat rejected, and the ~2-7 kg/m<sup>2</sup> 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.

#### LEADING METRIC

Announced or as-built radiator area and radiator mass fraction of any orbital-compute spacecraft above ~100 kW IT load, expressed as m<sup>2</sup> 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/m<sup>2</sup> in flight hardware (currently ~3-7 kg/m<sup>2</sup>). Leading tell: every serious orbital-DC technical disclosure or teardown that names radiator area/mass (not power or launch) as the gating line item.

#### KILL-CRITERION

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/m<sup>2</sup> (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.

#### REFUTE CHECK (SURVIVED)

Three attempted kills, none fatal. (1) Droplet/liquid-sheet radiators could achieve 400-600 W/kg and erase the area wall. Real threat, addressed in kill criterion: no droplet radiator has ever flown in qualified form; contamination and capture-efficiency remain undemonstrated after 40 years of study. The 2034 kill bar is set correctly at >400 W/kg in flight hardware. (2) Run chips hot (400 K coolant) to get 3x more radiated flux via T<sup>4</sup>. Real leverage but bounded: Al silicon junction limits sit at 85-105 C, and chip power density is climbing faster than the T<sup>4</sup> gain can offset. Net: practical reject temperature stays near 310-330 K. (3) The \$170M Starcloud raise means the market has priced this in. The round is pitched on compute capacity; Starcloud-1

carries one H100 at sub-kW thermal load. A 5 GW system needs ~8 km<sup>2</sup> of radiator. The funding validates the compute narrative, not the radiator-scarcity thesis. Call survives all three attacks.

**Why this call earned a place** The structural mechanism is correct and the inelastic needle is precisely named. The area wall is real, the supply of qualified deployable radiators is 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 primary risk to the clause is an 8-year horizon during which droplet radiator or graphene-composite breakthroughs could move the mass floor. The kill criterion is calibrated correctly at a 3-5x improvement threshold. Promoted because the physics is correct, the needle is specific, the pre-consensus check passes on the quantitative framing (not the theme), and the adversarial attacks do not kill it.

## P2 · Every kilowatt-class non-polar lunar surface asset through 2040 is gated by the 354-hour lunar night: the binding constraint is night-survival energy-storage mass per continuous watt, which forces fission surface power as the only viable route at useful power levels, and that route is single-sourced and schedule-fragile.

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 addressable set of continuously-powered lunar surface assets is therefore capped at the reactor count, regardless of how many landers fly.

Domain: space

DIRECTIONAL VISION

82%

STRICT CLAUSE

70%

RESOLVES

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 requires 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 needle 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.

### WHY IT IS PRE-CONSENSUS

The lunar investment and policy narrative through 2026 is dominated by transport (Starship/Artemis cadence, CLPS landers) and ISRU end-products (water ice, oxygen, regolith). Night-survival power appears only as a known engineering caveat in NASA technical papers, not as the priced gating constraint on the entire surface economy. No sell-side equity coverage, futures, or capital-markets instrument captures single-source fragility of the fission surface power program as the chokepoint that decides which lunar ventures are physically viable. The consensus implicitly assumes "solar plus batteries, scale as needed." The 354-hour night combined with the roughly 3 kg per continuous watt storage floor makes that false above toy power levels, and that specific mass-per-watt-to-reactor-count needle is absent from the capital and policy story.

### HONEST PRICE CHANNEL

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 appears in trade press (Aviation Week, ANS, Lockheed Martin feature) as a program milestone story, not as a binding economic chokepoint. Narrative-obscure AND unpriced -- the pre-consensus test passes.

#### THE NEEDLE

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.

#### LEADING METRIC

(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  $\geq 100$  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  $\geq 1$  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.

#### KILL-CRITERION

Killed if, by 2038-12-31, either (a) two or more independent lunar surface assets demonstrate continuous  $\geq 1$  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.

#### REFUTE CHECK (SURVIVED)

Three adversarial routes, all fail. (1) Polar peaks of eternal light make the argument moot: rebutted -- the permanently illuminated areas total a few square kilometers, are topographically constrained, and are geopolitically contested by every lunar program simultaneously; they cannot scale to a surface economy. (2) Commercial nuclear enthusiasm (Oklo, Astral, X-energy) could spawn competing designs quickly: rebutted - space qualification is a decade-long gauntlet; no commercial microreactor design is on a credible path to flight qualification by 2030; the NASA/DOE MOU itself is resetting to a fresh downselect process. (3) The clause might fail because the kill events are too demanding: this is actually a feature -- the clause 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 routes.

**Why this call earned a place** The physics is correct and invariant. The constraint is absent from priced channels. The adversarial routes fail. The single-source fragility of the fission program is real and documented by its own schedule history. The kill clause is tight and appropriately demanding. The 12-year window introduces non-negligible uncertainty on a storage breakthrough (hence clause\_p 0.70 rather than 0.85), but the structural case for the mechanism is strong (vision\_p 0.82).

## 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 structural deficit widening faster than new aperture can be fielded.

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, creating 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 the call describes. · Domain: space

DIRECTIONAL VISION

78%

STRICT CLAUSE

62%

RESOLVES

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 operate 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.

### WHY IT IS PRE-CONSENSUS

The space-finance and commercial space community prices 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 consensus assumption is "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, not in any instrument the market prices.

### HONEST PRICE CHANNEL

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) attracts all analyst attention. Deep-space tracking capacity is priced nowhere and traded nowhere. Not already priced.

## THE NEEDLE

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.

## LEADING METRIC

(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.

## KILL-CRITERION

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.

## REFUTE CHECK (SURVIVED)

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. Counter: DSOC is a single experiment; optical ground terminals have weather and cloud-cover limitations; qualifying optical as the primary link for the bulk of deep-space data by 2034 requires a funded program-of-record adoption pace that does not currently exist in NASA's budget. The kill criterion correctly requires adoption "for the bulk of deep-space data," not mere demonstration. Second counter: Artemis demand collapse. If Artemis slips or is cancelled, the multiplicative demand growth story weakens. Counter to that: the base oversubscription at peak already exists from standing science missions regardless of Artemis. The structural deficit survives partial manifest collapse, though the severity is dampened. The call survives both adversarial vectors at the structural level; the main honest uncertainty is whether the constraint becomes visibly "binding on the economy" in the legible, consequential way the clause implies, versus remaining a government-internal rationing problem that mission planners absorb quietly.

**Why this call earned a place** Physical mechanism is correct 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. Clause probability is set below vision probability 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 requires the consequence to be legible, not just documented in government audits.

## 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 inelastic input gating that accuracy is real-time in-situ upper-atmosphere driver data from a tiny, non-redundant, government-funded sensor base with no commercial supply.

The boom: Cheap launch and proliferated LEO are read as a traffic-management problem solved by more radars or optical trackers and 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"). This driver data is produced by a handful of single-string, government-budget-funded instruments with no commercial supply chain. Value and regulatory leverage migrate to whoever can fly a commercial accelerometer and EUV density-sensing layer and sell assimilated forecasts, because the entire autonomous-collision-avoidance investment thesis is physically downstream of this one inelastic measurement input.

Domain: space

DIRECTIONAL VISION

72%

STRICT CLAUSE

52%

RESOLVES

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, but the in-situ thermospheric measurement base is fixed and aging. Sensors that track where objects are do not measure the atmosphere that determines where they will be. The constraint cannot be relaxed by ground-tracking capex or maneuvering software, only by new in-situ measurement assets, which are slow, capital-gated, and currently non-commercial. One caveat that 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. This makes the rent-migration story apply most cleanly to smaller third-party operators rather than universally.

### WHY IT IS PRE-CONSENSUS

Consensus 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 as 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 rent migration from tracking sensors to atmosphere-measuring sensors is absent from any market framing found in equity coverage or sell-side research. Pre-consensus check passes.

#### HONEST PRICE CHANNEL

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, confirming the call is not yet priced.

#### THE NEEDLE

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.

#### LEADING METRIC

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.

#### KILL-CRITERION

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.

#### REFUTE CHECK (SURVIVED)

Three adversarial angles survive 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 needing to buy data. This means rent may not migrate to a third-party sensor provider for Starlink, only for smaller operators. Second, the software-assimilation path: continuous 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 structural physics argument survives all three; the rent-migration and timing clause carry genuine uncertainty.

**Why this call earned a place** The physical 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 pre-consensus by a clear margin. The adversarial search found one pitch-stage startup (Ensemble Space Lab) and one academic research path (Starlink GNSS-accelerometry), neither of which qualifies as commercial supply. The weaknesses are the self-measurement loophole for mega-constellation operators and the possibility that a government mission fills the gap without creating commercial rent. These lower clause\_p to 0.52 (the timing and exact resolution mechanism are uncertain) while leaving vision\_p at 0.72 (the structural constraint is real and likely to remain binding through at least the early 2030s). PROMOTE with the kill criteria intact.

# 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

P5

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 is absent from 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 performs the safety function 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 layer on 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 is now bidding 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 migrates 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. · Domain: space

DIRECTIONAL VISION

72%

STRICT CLAUSE

38%

RESOLVES

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 publicly visible plutonium story, plus a cross-spine collision between space nuclear and hydrogen buildout on a shared inelastic input. The collision is physically real and bureaucratically invisible today.

## WHY IT IS PRE-CONSENSUS

The iridium-for-PEM squeeze is now in PGM sell-side notes, hydrogen market reports, and moving visibly in spot prices. That first-order claim is priced. The genuinely pre-consensus 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.

## HONEST PRICE CHANNEL

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 found, confirming the second-order claim remains pre-consensus.

#### THE NEEDLE

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

#### LEADING METRIC

(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.

#### KILL-CRITERION

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.

#### REFUTE CHECK (SURVIVED)

Three genuine 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. That is 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/cm<sup>2</sup> toward under 0.1 mg/cm<sup>2</sup>, and ruthenium-based or mixed-oxide anodes at pilot scale, could collapse PEM iridium demand well before 2030, dissolving 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 silently through budget workarounds or mission de-scopes attributed to other causes. The clause 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 structural case survives attacks one and three in substance; attack two (catalyst thrifting) is the live uncertainty that most moves clause\_p below vision\_p.

**Why this call earned a place PROMOTE.** The physical mechanism is sound, the needle is correctly identified 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 hydrogen buildout on the same 7 t/yr stream) is absent from all space-sector and program-status

discourse. vision\_p is 0.72 because the structural case is well-grounded but catalyst thrifting is a live kill path. clause\_p is 0.38 because the resolution criterion requires an explicit public naming in a government document over a nine-year horizon, and silent absorption or thrifting-driven dissolution are both plausible ways the physical reality never surfaces in that specific form.

## 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 as the scarce, rent-capturing input. P6

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 contributing 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 Al<sub>2</sub>O<sub>3</sub> 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 migrates to the inelastic input that keeps the conveyor running under the cap: spacecraft structures that fully demise without injecting alumina. · Domain: space

DIRECTIONAL VISION

**72%**

STRICT CLAUSE

**52%**

RESOLVES

**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. Al<sub>2</sub>O<sub>3</sub> 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 of which are qualification-gated and slow. The 2026 Barker et al. paper (Earth's Future, AGU) now provides the modeling basis for regulators; 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.

### WHY IT IS PRE-CONSENSUS

Reentry is universally narrated as a debris and ground-safety success story; demisable design is the solution in industry framing. Orbital sustainability discourse focuses 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 is entirely 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 synthesis -- atmospheric planetary boundary translating to a reentry-mass cap translating to a structural-material rent -- exists in zero market or regulatory framing as of mid-2026. The science is getting louder while the capital signal remains silent; that gap is the pre-consensus window.

### HONEST PRICE CHANNEL

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 pre-consensus gap is real.

#### THE NEEDLE

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.

#### LEADING METRIC

Measured stratospheric Al<sub>2</sub>O<sub>3</sub> 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<sub>2</sub>O<sub>3</sub> 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.

#### KILL-CRITERION

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.

#### REFUTE CHECK (SURVIVED)

Three serious adversarial attacks. First: controlled de-orbit to ocean could relieve the constraint without repricing structural materials. This is the most credible kill path. Counter: for sub-200 kg LEO birds at constellation scale (thousands of vehicles), the fuel mass needed 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 is correctly named in the kill criteria. Second: the scientific attribution and regulatory timelines may slip past 2038. The Montreal Protocol ran 13 years from Rowland-Molina to binding rule, which is the closest historical analogy; the 2038 deadline is 12 years from now 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. This is 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 attacks with meaningful residual uncertainty, which is correctly reflected in clause\_p of 0.52.

**Why this call earned a place** The physical mechanism is confirmed by peer-reviewed literature published in 2025-2026. The pre-consensus gap is genuine: science is now audible but 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 materialize, consistent with the Montreal Protocol precedent. Clause\_p is set 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 -- a specific path among several plausible outcomes.

## Seeds considered and not promoted

Cleared the physical-constraint test but failed on investability or on the price channel. Logged because the discipline is to surface what was cut.

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) vs 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 needle-specific and less duplicative representative 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 a same-needle pileup.
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) vs the headline 140 t plutonium stockpile	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 its clause 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 is 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 its clause 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 wins the slot because it names a cleaner rent-bearing needle (qualified non-aluminum demisable structure) and carries clause_p 0.52 versus this call's 0.22, since the current US regulatory vector is running backward (FCC narrowing NEPA in 2025) and a

SEED	PHYSICAL CASE	WHY NOT PROMOTED
<p>By 2035, radioisotope-power-source availability becomes the gating item for sustained lunar/cislunar presence, with Sr-90/Am-241 substitutes themselves separation-limited</p>	<p>Decay-heat kilograms from a single unexpandable federal line as the input you cannot launch around</p>	<p>multilateral treaty route likely overshoots 2034.</p> <p>Highest-level framing of the radioisotope thesis, which makes it the most theme-like and least needle-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.</p>

Generated by the Pope System. Each call is a forward instrument: resolution date and kill-criterion fixed at creation, superseded never edited, clause probability scored with Brier at resolution.