SpaceX IPO, Orbital Data Centers, and Three Card Monte: The Cloud Is Not Above the World

Rob Tow

Nova Lux, New Mexico, USA, Sol III

17 June 2026

SpaceX's IPO names orbital AI compute as a central growth story and asks regulators for as many as a million data-center satellites. The phrase "data center in space" works like three-card monte; the money card is the familiar industrial noun, and the familiarity does the cheating. What is sold as one cooled, serviced, depreciated plant is a swarm of thousands of separate spacecraft, each carrying its own thermal, power, aging, and disposal problem. Follow the burdens the prospectus moves out of frame: radiators sized by the Stefan-Boltzmann law and exposed to debris; communications mass and weather-diverse ground stations; a dawn-dusk sun-synchronous shell that saves battery and buys crowding; compute that ages on one clock while the spacecraft amortizes on another, with no in-orbit servicing and no secondary market; a standing launch program below and a reentry crematorium above, both already measurable in the upper atmosphere; debris, near-Kessler traffic, Carrington-class common-cause risk, correlated insurance exposure, and overlapping sovereignty. Carried honestly, a five-year unit of useful compute costs three to five times its terrestrial equivalent. The cloud is not above the world; it is inside physics, chemistry, risk, and at the last inside the books. I did not invest in the IPO.

THE PITCH

I HAVE SPENT ENOUGH TIME around machines, and around the people who build them, and around the people who sell what those machines may someday become, to distrust any story in which the hard parts seem to dissolve all at once. Usually they have not dissolved. They have merely been moved out of frame. Something difficult to cool is being described as elegant. Something difficult to replace is being described as scalable. Something that lives under several ugly clocks at once is being spoken of as though it were already mature infrastructure. One learns, if one is lucky, to listen for that change in register. The nouns get cleaner. The object under discussion does not. Marketing departments lie. Once I worked for a company that shipped an empty box with a logo on the front to be displayed in a trade show booth, and took orders.

SpaceX's IPO materials say the company is "the ONLY company capable of building orbital AI compute at scale," and say it expects to begin deploying orbital AI compute satellites "as early as 2028." Reuters reported that this was not decorative color in the offering but part of the central growth story, with demonstration missions targeted for late 2027 and a regulatory request reaching as high as one million space-based data-center satellites. That is a large claim, and a loaded one. It asks the reader not merely to believe that computers can be put in orbit—of course they can—but to believe that orbit is ripening into a better industrial habitat for general compute than the planet that already contains our grids, our fiber, our maintenance crews, our warehouses, our lawyers, our spare parts, our roads, and our heat sinks.

My objection is not to ambition. Ambition is cheap, and can be useful. My objection is that the phrase "data center in space" misdescribes the thing so badly that honesty no longer seems essential to the product. It encourages the reader to picture a familiar industrial object in an unfamiliar location, as though altitude were the principal novelty and everything else would follow by some ordinary act of scaling. But what is actually being proposed is not a server farm moved upward. It is a launched, thermally constrained, damage-exposed, drag-exposed, depreciation-sensitive, disposal-shaped swarm undertaking whose operating logic runs from foundry to launch manifest to orbital shell to reentry corridor. SpaceX's own filing, whenever the prose stops preening long enough to speak legally, concedes the bones of the matter: no prior operating history for orbital AI compute, untested conditions in space, no ready repair or upgrade after deployment, and the risk of accelerated depreciation, decommissioning, and replacement. Plus liability.

I see perfectly well why the story has force. Building large facilities on Earth is increasingly entangled with land fights, power queues, water constraints, permits, and local resistance. The fantasy of stepping over those frictions has obvious appeal. Sunlight without county hearings. Expansion without neighbors. Scale without waiting for some utility commission to rediscover that electricity is finite. One need not be foolish to feel the attraction. One need only be tired of the earthly terms under which large systems now get built. That is why the pitch is effective. It starts from a real grievance. Earth has become slow, contested, and bureaucratically sticky. The prospectus offers not so much a solution as an escape hatch. Still, there's more than a bit of carnie going on.

This is precisely why the language deserves inspection. The problem is not one missing footnote or one engineering detail omitted for readability. It is a change in kind. Once the proposal leaves the slide deck and enters physics, ecology, finance, law, and strategy, the supposedly secondary burdens stop behaving like annotations and begin defining the enterprise itself. What is being sold remains legible only so long as those burdens are treated as background. The moment they are carried honestly in the account, the object changes. The investor story encourages the eye to glide past that change. The world will not be so polite. Neither will I.

The Con

There is another trick in the phrase "data center in space," and it is not minor. It works like three-card monte. For readers who have never seen the game: the operator shows the mark one winning card and invites him to keep his eye on it while the cards are moved. The point is not simply speed. The point is induced confidence. The mark is made to feel that the object remains visible and trackable, that careful watching should be enough, and that the game is therefore legible. Often there are shills nearby to thicken this confidence. But the apparent transparency is part of the fraud. What the mark thinks he is following is no longer what is actually in play.

That matter of what's in play is why the metaphor applies here. The money card is "data center." It is a familiar industrial noun, and the familiarity does most of the work. The reader hears it and forms a stable picture: one plant, one site, one coherent facility. He is then invited to believe that he can follow that same object through the rest of the pitch. The words move, the glamour increases, the location changes, but the reader feels he is still looking at the same class of thing. He is not. What SpaceX's materials actually describe is not one bounded orbital factory but a constellation of a very large number of separate satellites. A terrestrial data center is coherent because it is physically one place. It can be cooled as one place, maintained as one place, secured as one place, depreciated as one place, insured as one place. A compute constellation is something else: many separate spacecraft, each with its own orbit, power system, thermal system, exposed surfaces, aging schedule, failure modes, and end-of-life problem. The singular noun borrows the industrial dignity of a campus and transfers it, by verbal sleight of hand, to a fleet. The reader thinks he is following a plant through a change of address. In reality, the supposed plant has become a swarm. The singular is the cheat.

Earth, meanwhile, enjoys a home-field advantage so large that booster prose has to work hard not to notice it. Ground compute lives inside a vast support ecology: fiber, roads, cranes, technicians, substations, spare parts, warehouses, legal routine, and the plain fact that when something breaks a human being can go touch it with tools. That vulgar abundance is not incidental. It is part of what a terrestrial data center is. Orbit strips almost all of it away. The burden does not become cleaner because it rises. It becomes harder to service, harder to replace, and easier to romanticize. Ad astra.

Cooling, Communications, and Orbit Geometry

The cooling story shows how the cleansing by language works. "Cooling in space" sounds almost like a free gift from physics; getting rid of all that pesky water. But there is no convection or conduction in vacuum. Waste heat must be rejected by radiation, which means large area, structure, mass, deployment complexity, and vulnerability. Musk is reported as saying the satellites would use solar power and radiate heat into space. Quite so. But radiating heat into space is not the elimination of a burden. It is governed by physics—the Stefan-Boltzmann Law. As such it is the conversion of a burden into a very large piece of hardware. Public reporting on the first AI1 satellite concept points to roughly 120 kilowatts of average compute, 150 kilowatts peak, and on the order of 110 square meters of deployable radiator to get rid of the heat into space. That is already enough to kill the fantasy that cooling is free. And the burden is not even static over mission life. As coatings and surfaces degrade under ultraviolet light, atomic oxygen, and radiation, the area required to hold the same temperature rises. One is not merely launching enough radiator for day one. One is launching enough extra radiator to survive its own decay, which means more mass, more drag, and worse economics from the start. And once that hardware is spread across a large area, the thermal solution acquires a second meaning. A radiator is not only a cooler. It is also exposed surface. The larger the area through which the payload must dump its heat, the larger the area it offers to micrometeoroids and debris. NASA's thermal-control and MMOD experience has long treated radiator structures and cooling paths as vulnerable elements that must be designed for puncture, leak isolation, redundancy, and degraded operation. Even the supposed solution to the heat problem arrives already coupled to survivability and loss modes. The brochure says cooling. The engineer must add: through a large, fragile, puncturable object in orbit. Lacking service.

A clever engineer will object that the radiator need not be so large; run it hotter. Radiated power climbs as the fourth power of absolute temperature, so a panel held at seven hundred kelvin sheds roughly sixteen times the heat per square meter of one held at three hundred fifty, and on that arithmetic alone the area collapses to a sixteenth. The objection is correct, and it is exactly where the conservation law collects its toll. To hold the chips near three hundred fifty kelvin while dumping their heat from a surface at seven hundred, one must pump heat up the temperature gradient, and pumping heat is not free. The thermodynamic ceiling is unforgiving and real hardware falls well below it. An ideal refrigerator working across this span has a coefficient of performance of one, spending a joule of work for every joule it lifts; but no such machine flies. The best space-qualified closed-cycle coolers, the Stirling, pulse-tube, and turbo-Brayton machines with actual orbital heritage, reach only fifteen to twenty-five percent of the Carnot ideal, and a heat pump rejecting hundreds of kilowatts at seven hundred kelvin has no heritage at all. Take twenty percent and call it generous. The machine then spends about five joules of work for every joule it lifts, and every one of those joules turns to heat as well, so the radiator must reject the chip load plus the pump work, something like six times the original heat. The sixteenfold area saving collapses to less than threefold; the panel shrinks to roughly a third, not an eighth. That much, charitably, is real. But look at the bill. Five-sixths of what the radiator now sheds is waste from the pump, and that pump work must be generated, which means something like six times the electrical power and therefore something like six times the photovoltaic array, with its own mass, its own drag, its own debris-exposed area; the modestly smaller radiator has been bought with a vastly larger solar wing. And the heat pump is itself a machine, a compressor or turbo-Brayton loop carrying mass, drawing power, and adding a single-point failure to a system that cannot be serviced in orbit. A passive radiator fails gracefully, shedding efficiency as it ages. An active refrigeration loop fails like a refrigerator. Worse, the seven-hundred-kelvin surface that does the radiating must survive seven hundred kelvin; coatings, joints, and working fluids age faster hot than cold, so the degradation burden walks back in through the same door it was shown out of. Less radiator area is not less burden. It is the same burden, relocated into power, mass, and a fresh way to die.

Communications belongs in the same family of hidden burdens. The orbital-compute story quietly creates three different communications problems at once: how data reaches orbit, how results return to Earth, and how the constellation talks to itself. The intersatellite piece is the least absurd. Google’s Project Suncatcher openly envisions tightly clustered satellites tied together by free-space optical links. But even there the data path reaches backward into fleet geometry. If the links require tight formations, then communications becomes part of the station-keeping problem, the control problem, and ultimately the mass problem. The harder burden is the comm through atmosphere. Google, notably, still lists high-bandwidth ground communications as a major unresolved challenge. That is a useful confession. The vacuum ends at the air, and once one is trying to move serious amounts of data between orbit and Earth, the problem becomes weather, turbulence, clouds, background light, and the need for multiple ground stations in different conditions.

Weather is a problem that a mature terrestrial fiber network routes around by virtue of being underground. A ground optical link to orbit must route through it. Atlantic hurricane season, a broad Midwestern storm front, or a bank of coastal cloud does not merely degrade one neat laser path; at scale it forces traffic sideways across a much larger network of ground stations and terrestrial backhaul (fiber optics) to where the sky is clear. NASA's TBIRD demonstrations show that very high optical downlink rates are possible. They also show why they are not magically simple. Clear nighttime passes were favored because the atmosphere behaves better then. So the elegant story of continuous sunlight in orbit does not become an equally elegant story of continuous easy communication to the ground. The GPU may sit in perpetual light. The customer does not. Between orbit and customer lies atmosphere, weather, and a geographically redundant terrestrial network big enough to absorb the next hurricane season without dropping the act.

Communications, like cooling, also has mass. A record-setting NASA optical downlink payload still occupied a substantial fraction of a CubeSat. Commercial optical terminals are measured in double-digit kilograms per channel. Once one starts speaking seriously about multiple intersatellite links, redundant ground paths, pointing assemblies, buffering, route diversity for bad weather, and the control hardware needed to keep all of that aimed correctly, communications ceases to be an invisible abstraction and becomes another major hardware tree.

One might ask why not abandon the finicky laser and fall back on radio, which shrugs off the clouds that defeat an optical beam. High-gigahertz radio, the Ka, V, and E bands up past seventy gigahertz, does penetrate weather far better than light; but it buys that cloud tolerance at the price of bandwidth and spectrum. The frequencies wide enough to matter are congested and not the operator's to spend; they are parceled out band by band by the ITU and national regulators, and there is simply not enough licensed bandwidth to carry a data center's egress down from orbit on radio alone. What spectrum there is fights its own atmosphere, the sixty-gigahertz oxygen line and the water-vapor lines eating signal, with rain fade at these frequencies brutal. So the proposals reach for optics precisely because light carries terahertz of raw bandwidth in tight, unlicensed beams. The reach is rational; it also reimports the weather it was meant to escape, and drags the geographically redundant ground network back in behind it. Radio trades bandwidth for cloud tolerance; optics trades cloud tolerance for bandwidth; neither trade abolishes the ground. Both end in more dishes, more sites, more backhaul, more cost. Radiators are not free. Comms are not weightless. Wavelength allocations are not free, either.

The orbit choice, once stated honestly, also grows less free than the sales language implies. If one wants to avoid large battery mass, the attraction is not simply "LEO." It is a much narrower architectural choice: a dawn-dusk sun-synchronous orbit riding the terminator so that the spacecraft sees near-continuous sunlight and spends little or no time in Earth shadow. ESA describes this logic openly, and Google's Suncatcher concept makes the same move for the same reason. The gain is real. Battery mass that would otherwise be carried merely to survive eclipse can be reduced, and in a system already burdened by compute, radiators, communications, and control, that matters. But it is not a gift. It is a trade. The orbit that saves the battery does not save the link; the bits still leave only two ways, sideways in orbit through optical crosslinks or down into a ground dish and onward through terrestrial fiber, and either path is more infrastructure and more cost, not less. The architecture is being compressed into particular inclinations, altitudes, and local-time geometries in order to save mass. Once that happens, one no longer has the rhetoric of open orbital real estate. One has a much narrow favored family of shells.

The problem with favored shells is that they collect customers. Brian Weeden's old AMOS work on sun-synchronous orbit made this point years ago: because many Earth-observation and related systems want similar inclinations, altitudes, and local-time geometry, the orbit family develops recurring crowded crossing zones near the poles. The issue is not science fiction. It is geometry. Polar crossings occur every revolution, and if too many operators are drawn toward the same orbital family, conjunction burden rises as a matter of arithmetic. Weeden was already arguing in 2008 that the situation was serious enough to justify zoning proposals for SSO in order to reduce collision risk. That is the hidden price of the battery-saving elegance. Dawn-dusk SSO reduces one burden by compressing the design into a busier and more constrained traffic pattern.

I have not yet found one accidental collision that everyone cites as the canonical "sun-synchronous polar compression event," and I do not wish to pretend otherwise. But one does not need a famous corpse to understand the traffic pattern. The 2009 Iridium 33–Cosmos 2251 collision is caution enough. It was not a dawn-dusk SSO collision in the narrow sense, but it was an accidental high-inclination LEO collision at about 790 kilometers and 11.6 kilometers per second relative speed, and NASA called it the most severe accidental fragmentation on record at the time. More than 1,800 new trackable debris objects were created. That is what orbital traffic looks like when it stops being metaphor and becomes matter. The SSO literature says the geometry is already bad enough to worry about. The Iridium–Cosmos event shows what even one bad day in crowded LEO can do.

Once the power architecture, the communications architecture, and the orbital placement are all pulled toward the same favored shells, the supposed freedom of "just put the data center in space" begins to look like another three-card-monte move. The cards have changed. The battery burden may indeed shrink, but the conjunction burden grows, the communications geometry becomes less forgiving, and the service model becomes more dependent on terrestrial weather-diverse ground stations and sideways routing when conditions go bad. The battery you avoid in one place returns elsewhere as traffic, routing, and control burden. There's almost a conservation law here; you can squeeze one value and it shows up elsewhere.

Depreciation, Servicing, and Disposal

The depreciation problem sits even closer to the center. High-end compute ages on a different clock from spacecraft structure, and the two clocks do not naturally agree. Reuters reports that the first AI satellite is expected to use Nvidia chips with compute comparable to a single GB300 rack. That is an excellent anti-hype detail and it should not be lost. The first-generation public "orbital data center" is, in terrestrial terms, roughly only one modern high-end rack wrapped in a very large and delicate spacecraft—not a full data center of hundreds or thousands of racks. On Earth, that sort of hardware lives inside an ecology of refresh. A rack that is no longer frontier can still do useful work. It can slide from premium training to lesser training, from there to inference, then to internal use, resale, or salvage. Terrestrial compute ages badly, but it ages into a ladder of afterlives. Orbital compute does not. Once the payload upstairs is no longer economically first-rate, its choices are harsher: keep flying stale payload, replace it by launch, or destroy it. SpaceX's own filing warns of accelerated depreciation, decommissioning, and replacement for infrastructure that cannot be readily repaired or upgraded once deployed. That is not a footnote; it is a statement about industrial metabolism. The compute dies on one clock, and the plant wants to live on another. The rhetoric asks the reader to imagine that these clocks are friendly, but they are not. A terrestrial GPU can retire downward into secondary markets while an orbital GPU retires into the fire of disposal by reentry. The normal secondary market that softens terrestrial capital decline is absent. That missing afterlife belongs in the bookkeeping, harshly.

There is also the servicing question, and here the rhetoric should stop pretending. For this kind of hardware there is no in-orbit servicing at all. Not a little, not a primitive version waiting to mature, not some hidden industrial practice that orbital-enthusiasm pieces are merely taking for granted. There is no commercial regime in which somebody goes out to a compute satellite, pulls a bad board, swaps a failed optical terminal, changes a coolant pump, upgrades the accelerator, closes the panel, and sends it back to work. NASA's own ISAM program and the Northrop MEV work are proof that people are trying to build servicing capabilities in space; they are not proof that a routine orbital server-replacement business exists. Their very existence is evidence that it does not. Northrop's MEV extends life in GEO by docking and providing propulsion and attitude control. NASA's ISAM effort is still explicitly a survey of developing capabilities and an attempt to establish an economic ecosystem. That is not a service truck. It is a research program and a narrow life-extension niche. For orbital data-center hardware, the current operating model is not repair but simple disposal by burnup. When a node ages out, fails, drifts out of economic usefulness, or becomes suspect, the practical answer is only burn it and replace it.

Atmosphere, Reentry, and Solar Weather

This is also where the environmental account has to be told straight. The prospectus frame quietly invites the reader to compare orbit with the local ugliness of terrestrial data centers: water fights, substations, diesel backup, land use, local politics. But that comparison is malformed unless ascent and demise are kept inside the complete bundled inventory. Rocket exhaust is not external to the facility. It is part of the facility's ongoing operation. Nor is reentry some metaphysical vanishing act. It is a material process in the upper atmosphere. If the business model depends on routine launch, routine replenishment, and routine burn-up, then the atmosphere is being used not only as surroundings but as exhaust path and crematorium. That may be a different kind of ugliness from a cooling tower in the desert. It is ugliness all the same, and it belongs in the account. The thermosphere as waste dump has consequences, just like industrial sewage into Chesapeake Bay.

The best current modeling is no longer hand-waving. NASA's 2024 atmosphere assessment says plainly that launches and reentries emit gases and aerosols into every layer of the atmosphere from the surface to low orbit and that these emissions potentially affect climate, ozone, mesospheric cloudiness, and thermosphere-ionosphere composition. NASA also says launches can alter local atmospheric composition along the flight path depending on propellant. That is the right cautious language for the near-field question. One Starship launch won't rearrange the weather over the Atlantic in the ordinary synoptic sense. The better-supported claim is more systemic and more serious: if launches become frequent enough, the accumulated emissions alter the upper-atmosphere and stratospheric system in ways that feed back into circulation rather than into one evening's forecast. Globally.

A recent SpaceX fueling report describes a Starship V3 load of more than 5,000 metric tonnes of methane and oxygen propellant. Multiply that by the initial deployment arithmetic and the first buildout alone points to something like 155 to 188 million tonnes of propellant burned. Even the five-year replacement regime implies on the order of 31 to 38 million tonnes of propellant per year thereafter, with the three-year case harsher still. That is not a side effect on the way to the business. It is part of the business.

Again, the important thing is to resist magical thinking. A launch plume is not free because it disperses quickly near the pad, and it is not climatically irrelevant because the total mass is smaller than the global aviation fuel burn. Rockets inject unusual mixtures of gases and particles into unusual altitudes, along unusual trajectories, with a chemistry that is still being worked out. The local downstream question and the long-term climate question are not the same. Locally, the best cautious statement is NASA's: launches can change atmospheric composition along the trajectory and near the launch corridor depending on propellant. Long-term, the best current modeling says that enough launch traffic changes stratospheric temperature, ozone, and circulation. The right criticism is not "rockets make tomorrow's thunderstorm." It is that a sufficiently large launch regime ceases to be meteorological noise and becomes part of the background climate system. Part of geoengineering at scale, which we are already doing badly.

The reentry side is now being atmospherically modeled almost as seriously. NOAA's more recent aluminum-oxide work treats 10 gigagrams per year of $Al2O3$ from reentry—consistent with expected megaconstellation growth by 2040—as enough to alter middle- and upper-atmosphere temperatures and winds and to delay ozone recovery. Their companion analysis says those metallic aerosols can reach the ozone layer within roughly 1 to 3 years depending on particle size and reentry location. The observational warning light is already on: a 2023 PNAS paper found that about ten percent of sampled stratospheric sulfuric-acid particles larger than 120 nanometers already contained aluminum and other metals from spacecraft reentry. So this is no longer speculative in the weak sense. The atoms are already there. The question is whether we choose to turn a detectable pollutant stream into an industry spew stream down river.

The environment into which the proposal is launched is no less uncooperative. Low orbit is not a neutral shelf. Solar storms heat and expand the upper atmosphere. Drag rises. Orbits decay faster. What looked like reserve propellant becomes ordinary expenditure. What looked like a tidy asset-life assumption becomes partly a weather assumption. The Sun and its angry storms get a vote. This matters more here than it would in a simpler orbital business because the proposal is already trying to balance too many unfriendly clocks. The payload wants rapid refresh. The plant wants slow amortization. The disposal problem wants low surviving mass. The drag environment wants some altitudes less than others. There is no altitude at which all these claims become friendly at once. Lower orbit helps eventual cleanup, but worsens drag and storm sensitivity. Higher orbit eases drag, but makes eventual disposal slower, more deliberate, and more politically loaded. There is not a clean optimum.

And then there is the larger storm, the Carrington-scale problem that terrestrial data centers do not face in the same direct way. A severe geomagnetic event can disturb the ionosphere, alter the radiation belts, heat the thermosphere, raise drag, degrade communications, and damage or confuse spacecraft electronics. Of course a modern Carrington-class event would hurt the terrestrial world too—grids, telecom, GPS, logistics, all of it. But a ground data center is not itself flying through the disturbed medium. An orbital compute fleet is. It takes the hit in its own body: surface charging, single-event effects, communications outages, navigation degradation, changed drag, and the ugly certainty that many nominally independent nodes all begin misbehaving in the same weather. The terrestrial facility suffers by dependence. The orbital one suffers by immersion. A dawn-dusk SSO architecture chosen partly to reduce battery burden does not escape that fact. It becomes more dependent on a narrow power and orbit logic at exactly the moment the Sun reminds everyone who owns the medium. A lee shore in a storm is not a happy place.

The probability language here is ugly enough without melodrama. Older statistical work put the chance of a Carrington-class event in the next decade at roughly twelve percent. Later probabilistic work lowered that to something more like 0.46 to 1.88 percent per decade, while other forecasters land around four percent per decade. Engineering reviews often convert the uncertainty into return-time language of something like one in one hundred to one in two hundred years for a Carrington-scale storm. That is already bad enough for infrastructure built on ordinary business horizons. Worse still, Carrington is not the upper bound. Tree rings and ice cores record stronger solar-proton events. The Miyake events of 774/775 and 993/994 CE are established in cosmogenic isotopes, and more ancient spikes—most famously the event roughly 14,300 years ago—appear larger still. Not every isotope spike is a simple one-to-one analogue of a modern geomagnetic storm, and honesty requires saying so. But the paleoclimate warning is nonetheless plain: the instrumental era as recorded is not the full range of the Sun. Modern engineering cannot dismiss that, and a prospectus that treats orbital compute as though it lived in a placid celestial utility district is suppressing a real burden. Shit will happen, sometime.

Underwriters and Failure Arithmetic

That is where the underwriters enter, and they deserve more respect than they receive in most futurist treatments of space infrastructure. Insurance is where metaphor goes to die. FAA material lays out the categories plainly enough: launch failure, in-orbit loss, third-party liability. These are not exotic edge cases appended by timid lawyers. They are standard realities of space operations. Rockets fail, and when they fail the loss is usually not a repair event but a total-loss event: "sudden unplanned in-flight disassembly". But the underwriter's problem is larger than counting explosions. It is portfolio correlation. A large constellation is not simply many assets. It is many related assets exposed to related hazards in related regimes. Debris events, liability shocks, strategic demonstrations are shared, space-weather shocks are shared, common design flaws are shared, and regulatory and diplomatic reactions to falling debris or upper-atmosphere pollution are shared. The underwriter is not pricing one satellite. He is pricing a fleet as a correlated book of hazards, many of which are capable of turning ugly together.

If the insurance market grows less comfortable with collision clustering, falling fragments, atmospheric externalities, hostile interference, or Carrington-class common-cause loss, that discomfort will not land delicately on one unlucky satellite. It will spread across the book. That is why the failure arithmetic matters so much. When one is speaking not of dozens of launches but of tens of thousands, even a small failure probability ceases to be a comforting percentage and becomes a certainty of many very large claims. At 99.5 percent launch success, initial deployment still implies on the order of 155 to 188 launch failures. At 99 percent it implies 310 to 375. At the current global 2025 orbital-launch success rate of about 97.6 percent, the expected number of failures in the initial buildout rises toward 750 to 900. That is before replacements, before weather delays, before common-mode fleet issues, and before the first geopolitical event that makes underwriters suddenly remember they have exclusions. Those are not percentages. They are a very large number of explosions, insurance claims, local contamination events, and broken schedules. An underwriter is not paid to admire the future. An underwriter is paid to survive exactly this sort of arithmetic. The booster calls it scale. The underwriter calls it correlated exposure. And once that phrase is taken seriously, a great many heroic assumptions begin to look like exclusions waiting to happen.

In orbit the arithmetic is no kinder. The only real megaconstellation evidence we possess already shows that unexpected failures are not zero and that common design flaws can force large proactive retirement waves. Some failures propagate physically by creating debris. Others propagate operationally by turning maneuverable nodes into non-maneuvering hazards in a crowded shell. Still others propagate through homogeneity: a million similar units means that one flaw, once discovered, may suddenly belong to all of them at once. That is not melodrama. It is what scale does to engineering error. The larger the fleet, the less a failure remains local and the more it begins to resemble epidemiology.

SOVEREIGNTY AND JURISDICTION

There is also the sovereignty problem, and it is larger than the pitch wants to admit. Outer space is not a national territory; the Outer Space Treaty forbids appropriation by claim of sovereignty. But the same treaty, reinforced by the Registration Convention, says that the registry state retains jurisdiction and control over the object. So the orbital compute node is not legally nowhere. It is tied to a state, dependent on spectrum coordination, dependent on earth stations, and entangled in the legal geography of the customers it serves. A terrestrial data center lives inside one clear sovereignty problem. An orbital compute constellation, with global communications infrastructure, lives inside several at once. The prospectus whispers escape from Earth. The reality is overlapping jurisdiction with better lighting. You do not escape sovereignty by going to orbit. You exchange one tidy sovereignty problem for several untidy ones.

Targetability, Debris, and Near-Kessler

The strategic angle is harsher still. A compute constellation is not merely capacity. It is a distributed target set. U.S. Space Force doctrine is explicit that space capabilities can be attacked through terrestrial segments, link segments, and orbital segments, and Secure World Foundation's latest counterspace survey describes a world in which multiple states are developing or fielding direct-ascent, co-orbital, electronic, directed-energy, and cyber means of interference or destruction. One does not need to believe in daily orbital war to notice the implication. A destructive demonstration, a debris-generating strike, or even a credible pattern of interference could force operating doctrine, reserves, and replacement schedules to change across the fleet. What the futurist calls scale, the strategist calls targetability.

The nightmare here is Kessler syndrome, usually served in Hollywood portions. By that I mean a debris-collision cascade: collisions create fragments, fragments create more collisions, and the orbital shell becomes progressively dirtier, riskier, and more expensive to use. That is too dramatic in its full form, and therefore too easy for salesmen to bat away as a B movie. One does not need a permanent metal storm around Earth to wound this business. A near-Kessler regime is enough. A few major fragment-generating events, a more crowded shell, and a debris environment that becomes less forgiving, more twitchy, and more expensive to operate within—that is enough. The business does not need orbit to become impossible. It merely needs it to become ugly. A business can die long before an orbit does.

The anti-satellite record makes this painfully concrete. China's destruction of Fengyun-1C in 2007 created the classic high-altitude warning: thousands of tracked fragments and a debris population serious enough that NASA's Terra spacecraft had to execute an avoidance maneuver because of one of them. The 2009 accidental Iridium 33–Cosmos 2251 collision then demonstrated that high-inclination LEO traffic could create a similarly ugly debris legacy without anyone firing a missile. The U.S. destruction of USA-193 in 2008 showed the low-altitude contrast case: still ugly, but with much faster natural cleanup. India's 2019 Mission Shakti test, carried out against a target in sun-synchronous orbit, reminded everyone that even a lower-altitude event can throw fragments high enough to trouble the ISS and shift risk calculations immediately. Russia's destruction of Cosmos 1408 in 2021 added another 1,500-plus trackable fragments and likely hundreds of thousands of smaller ones. The lesson is not subtle. Altitude matters. Geometry matters. One destructive act can change the background risk for everyone else sharing the shell.

One need not imagine some grand orbital battle. A single launch by a rogue state releasing a cloud of dense pellets would be enough. At orbital velocity, small tungsten BBs become destructive projectiles. A crude dispersal could make an already crowded shell more dangerous to operate in, more expensive to insure, and harder to refresh or replace. The anti-satellite tests already performed by real states have done the explanatory work. The BB image merely strips the mechanism to its ugly essentials.

The arithmetic at scale makes this uglier still. If one takes the one-million-satellite filing seriously rather than poetically, the unit conversions are sobering. An AI1-class node at roughly 120 kilowatts average compute implies about 833 satellites for a 100-megawatt orbital "data center" and about 8,333 for a 1-gigawatt one. The one-million-satellite proposal therefore corresponds not to a clever little cluster but to roughly twelve hundred 100-megawatt facilities or about one hundred twenty 1-gigawatt facilities. That is the scale of a new off-Earth industry, not a colorful adjunct to Starlink.

The public 70 kilowatts-per-ton figure for AI1 should be treated as a floor, not a planning number. That number belongs to the world of roadshow efficiency, not to the world in which communications terminals, pointing assemblies, extra control authority, buffering, redundancy, packaging, and the usual aerospace margin all have to be carried honestly. Once one stops pretending those things are free, something like 2.5 to 3 tonnes per node is a more sober working mass. On that basis the full million-satellite proposal implies roughly 2.5 to 3 million tonnes on orbit. That is not a colorful extension of a communications constellation. It is the mass of an industry.

If Starship in fully reusable service can truly net something like 80 useful tonnes to orbit per launch after packaging and real-world inefficiencies, initial deployment alone runs to something like 31,000 to 37,500 launches. That is not cadence as metaphor. It is cadence as a permanent industrial regime. And deployment is the easy part. Then replacement begins. A five-year turnover means replacing on the order of half a million to six hundred thousand tonnes per year, which means roughly 6,250 to 7,500 launches annually just to stand still. A three-year cycle is harsher still, pushing toward 10,400 to 12,500 launches per year. Once the arithmetic is carried through, "data center in space" stops meaning compute and starts meaning a standing launch program that exists merely to keep the old hardware from becoming the current hardware.

The same mass then returns through the upper atmosphere if the fleet is treated as disposable. A five-year turnover of a 2.5 to 3 million tonne orbital industry implies roughly 500,000 to 600,000 tonnes reentering per year. A three-year cadence pushes that toward 830,000 to a million tonnes per year. Current atmospheric literature does not need to be stretched into apocalypse to make the point. NASA's broader atmosphere assessments and NOAA's work on aluminum-oxide aerosol already warn that launch and reentry products are becoming a scientifically significant upper-atmosphere burden at today's far smaller scales. At the scale implied by the filing, this ceases to be a marginal environmental question. The cloud in orbit turns out to require an industry below it and a crematorium above it. It is not merely compute. It is a standing atmospheric chemistry experiment.

BUNDLED COST: EARTH VERSUS ORBIT

The cost comparison, if done honestly, must be bundled and buyer-facing. It is not enough to compare terrestrial electricity cost to orbital sunlight, or one building permit to one launch. A serious customer—Walmart is a useful stand-in—buys a five-year plan for useful compute. The cleanest common unit here is one GB300-rack-equivalent of useful compute, because the public AI1 figures describe roughly 120 kilowatts of average compute and 150 kilowatts peak, which Musk himself equated to about one modern GB300 rack. The question is therefore not whether orbit can host compute in principle. The question is what a five-year unit of useful rack-class compute actually costs on Earth, and what the same unit costs in orbit, once the whole burden tree is carried honestly. It is worth separating two different forms of arithmetic here. Engineering total-cost-of-ownership thought experiments can be useful for showing that even optimistic assumptions still look bad. Buyer-planning economics are harsher. They are the numbers a serious customer would actually use to commit capital over five years.

On Earth, the bundled cost is ugly but legible. JLL's 2026 outlook puts average shell-and-core data-center construction at about $11.3 million per megawatt, with AI technical fit-out adding as much as $25 million per megawatt. At 0.12 megawatts per GB300-rack-equivalent, that implies roughly $4.36 million of terrestrial AI-ready capital burden before electricity. Five years of continuous power at a representative 2026 industrial electricity price around 8.6 cents per kilowatt-hour adds about $0.45 million more. Taxes belong in the account as well, though many terrestrial centers now enjoy sweetheart abatements generous enough to turn part of the burden into a public subsidy. Against that, Earth still grants something orbit does not: repair, rerouting, redeployment, salvage, resale. A terrestrial rack ages badly, but it does not ordinarily age into ashes. For planning purposes, a serious buyer should think in the neighborhood of roughly $5 million to $6 million over five years for one GB300-rack-equivalent on Earth.

Orbit carries a different bundle, and every time the prospectus says "data center" it quietly asks the reader not to total it. Start with mass. The public 70 kilowatts-per-ton figure for AI1 is a roadshow floor, not a procurement number. Once communications terminals, pointing assemblies, extra control authority, redundancy, packaging, and the ordinary sins of aerospace margin are priced honestly, a planning mass of 2.5 to 3.0 tonnes per node is the more sober assumption. At roughly $1,500 per kilogram—not Google's hoped-for mid-2030s threshold of under $200/kg, but a much more serious present-ish launch assumption—launch alone costs about $3.75 million to $4.50 million per rack-equivalent. That is before the spacecraft itself is paid for.

And the spacecraft is not free. One orbital rack-equivalent must still be manufactured, integrated, qualified, and wrapped in radiator structure, communications hardware, attitude-control burden, and packaging for launch. Stop being kind and the manufacturing and integration number is not lower than Earth. It is higher. For planning purposes, I would carry roughly $6 million to $9 million per node there. Then come insurance and retained-loss reserve: launch insurance, in-orbit insurance, exclusions, self-insured correlation risk, all priced against replacement value rather than fantasy. Carry another $1 million to $2 million over five years. Then communications: not just optical terminals on the spacecraft, but weather-diverse ground stations and the terrestrial fiber to route around cloud, storm fronts, and bad seeing. Carry another $1 million to $3 million there. Then terminal depreciation: no real secondary market, no graceful demotion path, no service truck. The asset does not age into resale. It ages into replacement or burn-up. Carry another $1 million to $2 million for that harshness alone.

That gets the five-year orbital planning number to something like $15 million to $25 million per GB300-rack-equivalent. Those are not vendor quotes or a clean market sheet. They are a buyer-side planning model built from sourced anchors and conservative burden allocations: launch mass, terrestrial construction cost, electricity, and the openly acknowledged communications problem, combined with explicit planning allowances for manufacturing, integration, insurance, correlation reserve, and terminal depreciation where no honest market quote yet exists. That is not a rhetorical multiplier. It is the consequence of refusing to give orbit free kindness on launch, manufacturing, communications, insurance, and terminal depreciation. Put next to the terrestrial planning number of roughly $5 million to $6 million, the comparison is no longer coy. The honest buyer's multiplier is roughly 3× to 5×. For the same useful compute, Earth gives you a costly asset. Orbit gives you a costlier consumable. IEEE Spectrum recently cited an ABI Research rough TCO exercise in which even a highly optimistic launch-cost assumption of $44 per kilogram still left space compute at least an order of magnitude more expensive than terrestrial compute on a per-year basis. That is not the same thing as a buyer's five-year planning number; if anything, it shows how little mercy the concept receives even when the assumptions are charitable.

That mass is what makes the sub-$200-per-kilogram story so slippery. Google's own Suncatcher note presents less than $200 per kilogram by the mid-2030s as a threshold at which space-based AI might begin to look roughly comparable to reported terrestrial energy costs on a per-kilowatt-year basis. That is not a present procurement number, and it is not full bundled parity. It belongs to the world of engineering threshold thought experiments, not to the world of buyer planning. To reach anything like that would require something very close to airline-like cadence and very mild refurbishment of the entire Starship stack, including the ship itself. I find that implausible. A reusable booster is plausible. The second-stage ship is where the hand-waving begins. It is not the booster's little brother, and not an airframe turned around by routine discipline. It is the orbital element: tanks, engines, structure, and thermal protection all carried deeper into the fight and brought home under harsher conditions. Mild refurbishment is therefore not a conservative assumption. It is part of the sales pitch.

I once toured Cape Canaveral while I was working in the Center Director's office at NASA Ames. I walked into a hangar where Discovery was being refurbished, hanging above me with tiles removed, and there was literally a man on scaffolding running a compressed-air wrench—brrrt!, brrrt!—like a car repair bay. It felt like walking into the fighter bay of Battlestar Galactica, except the scorched spaceplane above me was real. And that was a Shuttle, with all the power and seriousness of NASA behind it. It still took months to turn around. That memory is part of why talk of mild second-stage refurbishment leaves me cold. I have seen what real turnaround on a flown orbital vehicle looks like, and it does not look cheap, and it does not look mild. The ongoing test failures of Starship are the cold reality.

Blue Origin is a useful cautionary case. It studied a reusable New Glenn upper stage under Project Jarvis and then set it aside; New Glenn today advertises a reusable first stage and an expendable second stage. In April 2026 Blue Origin even flew a reused booster successfully while the mission still failed because the upper stage underperformed. That is not proof that full second-stage reuse is impossible, but it is evidence that the harder half of the problem remains the harder half. A more credible middle ground for large-launch economics is therefore something between Falcon Heavy and the Starship dream: reusable booster, expendable or heavily refurbished upper stage, and costs that fall substantially from today without ever approaching the fantasy threshold. If that is the actual regime, then a large fraction of the orbital-industrial mass still gets burned up, and the launch-cost floor stays far above the number used to make orbital compute look almost reasonable.

Earth gives you a costly asset that can be repaired, rerouted, depreciated, and resold. Orbit gives you a costlier asset that must be launched, insured, replaced, and burned, while living inside failure modes the ground largely suffers only by dependence: solar storms, debris, and enemy action. The economics are not adventurous. They are unreasonable. The three-card monte depends on the reader not doing this arithmetic.

INFRASTRUCTURE FRAGILITY AND THE NARROW CASES

I am not anti-space. I am strongly in favor of scientific spaceflight, communications constellations, lower ground-to-orbit cost, and the new commercial uses that cheaper lift will make possible. New capability will come. New businesses will come with it. My objection here is not to commercial space as such. It is to this specific fantasy of one million orbiting data centers, and to the prospectus language that tries to make that fantasy sound like ordinary infrastructure.

This is not really a space peculiarity. It is a recurring failure of managerial thought. Infrastructure vulnerability is almost always underpriced by people living on quarterly incentives. It sounds dreary and lacks glamour. People prefer throughput charts, installed capacity, and smooth growth curves. They do not like lead times, replacement paths, obscure single-source components, insurer appetite, or the question of how a system behaves after its first serious wound. Yet that is exactly where real infrastructure reveals its nature. Not in the rendering; in the wound. We have seen this pattern before. When the public conversation turns to Hormuz, it fixates first on crude. But the deeper injury often arrives through narrower channels and slower repair clocks: helium, sulphur, urea, specialized heat exchangers in Qatar with five-year lead times, the ugly calendar attached to replacing awkward but load-bearing industrial parts. Serious infrastructure is constrained not only by throughput, but by replacement time. The quarterly-bonus mind sees the big stream. The strategist inspects the missing part and the repair calendar. Orbital compute deserves the same inspection.

That, in turn, clarifies the credible use cases. If the useful customer is on Earth, then the useful result usually has to come back down. That fact does not kill every conceivable orbital compute use. It narrows the serious ones. The strongest cases are the narrow ones: computation close to data already gathered in space; filtering, compression, autonomy; some military or sovereign functions where ugly architectures are tolerated because other priorities dominate. I do not deny those cases. In fact, conceding them sharpens the knife. The strongest arguments for the concept are precisely the ones too small to bear the civilizational rhetoric now attached to it. The idea becomes most plausible exactly where it ceases to justify the prospectus. A niche can be genuine without being civilizational. An edge case can be useful without becoming the new cloud. The trouble is that once these modest possibilities are admitted, booster rhetoric immediately tries to inflate them into proof that the gigantic story is sound. That is how hope begins dressing as arithmetic.

PATTERN, PROSPECTUS, AND VERDICT

Nor is the rhetorical style of the prospectus unfamiliar. Musk has a long habit of speaking in destination tense, of describing immature systems in the grammar of near-arrival, of allowing aspiration and achieved capacity to live too close together for hygienic comfort. Tesla's self-driving saga is the obvious terrestrial example: years of rhetoric in the gravitational orbit of autonomy while regulators, courts, crash investigations, and even internal testimony kept insisting on unresolved burdens. Optimus belongs in the same family. A robot publicly framed as a step toward general-purpose labor still appears, at crucial moments, to depend on human teleoperation or carefully staged demonstrations. One sees why this style works. It keeps capital warm. It keeps followers animated. It keeps the future permanently within arm's reach. But it also teaches the audience a bad habit: to treat unresolved burdens as though they were already beneath notice. Here they are not beneath notice. They are the story. I do not believe this is innocent. Musk may be delusional about his own timelines. Or a skilled carnie. Gwynne Shotwell, the SpaceX CEO who has lived in launch economics and operations too long not to understand burden trees, is far harder to excuse.

So let me say it without embroidery. This does not read to me as an honest prospectus. It reads as a sales document in which the governing burdens of this complex arrangement have been repeatedly dimmed so that the abstraction can take the light. The phrase "data center in space" is part of that dimming. The talk of cooling, scale, and escape from earthly friction is part of that dimming. The migration of depreciation, drag, debris, targetability, disposal, insurer logic, atmospheric chemistry, and replacement time into the wings is part of that act of obscuration. None of this is an innocent omission. The omitted material is exactly what governs whether the project is industrially serious. That is why the dishonesty matters. The prospectus asks the reader to price the dream while quietly stepping over the machinery required to keep the dream from catching fire, failing, or poisoning its own medium.

The question is not whether computers can be placed in orbit. Of course they can. The question is whether orbit is a better industrial habitat for general terrestrial compute once the real account is carried honestly: launch, thermal rejection, thermal vulnerability, communications mass, weather-diverse ground links, sovereignty, compute obsolescence, missing secondary markets, the total absence of in-orbit servicing for this class of hardware, drag, debris, controlled or uncontrolled demise, hostile targetability, Carrington-class common-cause risk, liability, insurance, and the stubborn calendar of repair and replacement. I do not think it is. The undertaking is too heavy, too exposed, too coupled, and too dependent on too many favorable conditions remaining favorable at once. What is being sold as transcendence looks, on inspection, more like burden displacement with distracting illustrations.

But it gets still worse.

And that is where the three-card monte returns. The money card was "data center." The mark was invited to believe that if he watched closely enough he could follow one coherent industrial object through a simple change of address. But the hands moved, the glamour increased, and the object changed under his gaze. By the time the lift comes, the supposed plant has become a fleet, the supposed escape from Earth has become another legal and logistical entanglement, and the supposed clean cloud has acquired a launch regime below it, a debris regime around it, and a crematorium above it. The trick was not speed. The trick was induced confidence. The reader was encouraged to think he was still looking at one bounded industrial object when, in fact, he was being walked into a standing launch program, a standing replacement program, and a standing atmospheric burden. The cloud is not above the world. It is still inside physics, inside chemistry, inside risk, and at the last inside the books.

I did not invest in the IPO.

References

Primary reporting and official material

SpaceX S-1/A registration statement and risk disclosures on orbital AI compute, deployment timing, and lack of ready repair or upgrade once deployed. U.S. Securities and Exchange Commission, June 2026.
https://www.sec.gov/Archives/edgar/data/1181412/000162828026040364/space

SpaceX free-writing prospectus on orbital AI compute at scale. U.S. Securities and Exchange Commission, June 2026.
https://www.sec.gov/Archives/edgar/data/1181412/000162828026040610/space

"SpaceX aims to launch orbital AI computing tests by end of next year, sources say." Reuters, June 9, 2026.
https://www.reuters.com/business/media-telecom/spacex-aims-launch-orbital-ai-computing-tests-by-end-next-year-sources-say-2026-06-09/

"Insurance for Commercial Space Transportation." Federal Aviation Administration, 2002.
https://www.faa.gov/about/office_org/headquarters_offices/ast/media/q42002.p

"Space Warfighting: A Framework for Planners." U.S. Space Force, 2025.
https://www.spaceforce.mil/Portals/2/Documents/SAF_2025/Space_Warfighti
_A_Framework_for_Planners_BLK2_%28final_20250410%29.pdf

"Global Counterspace Capabilities: An Open Source Assessment." Secure World Foundation, 2026.
https://swfound.org/counterspace/

Scientific and technical papers / engineering sources

NASA thermal-control and MMOD vulnerability materials; radiator flow-path and cooling-loop puncture considerations.
https://tfaws.nasa.gov/wp-content/uploads/MMOD_Course_TFAWS_2014.pd

Google Research, "Exploring a space-based scalable AI infrastruc-

ture system design."
https://research.google/blog/exploring-a-space-based-scalable-ai-infrastructure-system-design/

Public reporting on early SpaceX AI satellite architecture and compute/radiator scale.
https://www.tomshardware.com/tech-industry/spacex-details-its-ai1-compute-satellite

IEEE Spectrum, "Orbital data centers will face a heat problem." Useful for radiator-area scaling, degradation over mission life, and the distinction between optimistic engineering TCO sketches and harder buyer economics.
https://spectrum.ieee.org/orbital-data-centers-heat

NASA TBIRD optical communications reports and related papers.
https://ntrs.nasa.gov/api/citations/20250003969/downloads/Riesing_TBIRD

NASA atmosphere assessment on launch and reentry emissions.
https://ntrs.nasa.gov/api/citations/20240013276/downloads/NASA-TM-20240013276-V6.pdf

NOAA Chemical Sciences Laboratory summary of modeling 10 Gg/yr Al_2O_3 reentry aerosol as a middle- and upper-atmosphere burden.
https://csl.noaa.gov/news/2025/427_0428.html

Jonathan McDowell, launch statistics and 2025 launch summary; used for current orbital-launch success rates and launch-count baselines.
https://planet4589.org/space/stats/launches.html
https://planet4589.org/space/papers/space25.pdf

SpaceX Starship official payload statement and roadshow payload-capacity material.
https://www.spacex.com/vehicles/starship
https://content.spacex.com/cms-assets/FINAL_Documents%20and%20Upda

ArXiv empirical work on Starlink constellation deployment, operational lifetime, and unexpected-failure behavior.
https://arxiv.org/html/2603.25835v1

ESA, “Types of orbits,” and related dawn-dusk sun-synchronous orbit material.
https://www.esa.int/Enabling_Support/Space_Transportation/Types_of_orbits
https://www.esa.int/ESA_Multimedia/Images/2020/03/Polar_and_Sun-synchronous_orbit

NOAA / SWPC and NOAA heritage material on geomagnetic storms, thermospheric heating, and Carrington-scale space-weather risk.
https://www.swpc.noaa.gov/phenomena/geomagnetic-storms
https://www.noaa.gov/heritage/stories/five-historically-huge-solar-events
https://www.nesdis.noaa.gov/news/safeguarding-satellites-how-noaa-monitors-space-weather-prevent-disruptions

Development of an Architecture of Sun-Synchronous Orbital Slots to Minimize Conjunctions. Brian Weeden, AMOS 2008. https://amostech.com/TechnicalPapers/2008/Posters/Weeden.pdf

2026 Global Data Center Market Outlook. JLL, 2026. https://www.jll.com/en-us/insights/market-outlook/data-center-outlook

Electric Power Monthly, Table 5.6.A, average price of electricity to ultimate customers by end-use sector, March 2026. U.S. Energy Information Administration.

Data Center Tax Abatements: Why States and Localities Must Disclose These Soaring Revenue Losses. Good Jobs First, April 2026. https://goodjobsfirst.org/data-center-tax-abatements-why-states-and-localities-must-disclose-these-soaring-revenue-losses/

Data center tax breaks becoming billion-dollar budget sinkholes. Good Jobs First, April 2026. https://goodjobsfirst.org/data-center-tax-breaks-becoming-billion-dollar-budget-sinkholes/

Ground Segment Design for Broadband Geostationary Satellite Systems Using Optical Feeder Links. Journal of Optical Communications and Networking, 2015. https://opg.optica.org/abstract.cfm?uri=jocn-7-4-325

Ground Station Network Optimization for Space-to-Ground Optical Communication Links. Journal of Optical Communications and Networking, 2015. https://elib.dlr.de/92544/1/2015_CloudAvail_final_p

The Collision of Iridium 33 and Cosmos 2251: The Shape of Things to Come. NASA, 2009. https://ntrs.nasa.gov/api/citations/2010000020

The Characteristics and Consequences of the Break-up of the Fengyun-1C Spacecraft. NASA/Hamilton Sundstrand, 2007. https://ntrs.nasa.gov/api/citations/20070007324/downloads/20070007324.p

Joint effort made satellite success possible. U.S. Strategic Command on USA-193 intercept, 2008. https://www.stratcom.mil/Media/News/N
Article-View/Article/983539/joint-effort-made-satellite-success-possible/

Indian Direct-Ascent Anti-Satellite Testing Fact Sheet. Secure World Foundation, 2019. https://www.swfound.org/publications-and-reports/indian-direct-ascent-anti-satellite-testing-fact-sheet

Russian Direct Ascent Anti-Satellite Missile Test Creates Significant, Long-Lasting Debris. U.S. Space Command, 2021. https://www.spacecom.mil/Newsroom/News/Article-Display/Article/2842957/russian-direct-ascent-anti-satellite-missile-test-creates-significant-long-last/

BACKGROUND AND CONTEXT

General background on the three-card monte confidence game. https://en.wikipedia.org/wiki/Three-card_monte

Additional FAA regulatory material on launch and reentry financial responsibility.
https://www.faa.gov/about/office_org/headquarters_offices/ast/licenses_perm

440.pdf

ISAM. NASA. Overview of in-space servicing, assembly, and manufacturing and the 2025 State of Play. https://www.nasa.gov/isam/

SpaceLogistics services using MEV for GEO satellite life extension. Northrop Grumman. https://www.northropgrumman.com/what-we-do/space/space-logistics-services

Impact of Spaceflight on Earth's Atmosphere. NASA Technical Memorandum, 2024. https://ntrs.nasa.gov/api/citations/20240013276/downloa
TM-20240013276-V6.pdf

Projected increase in space travel may damage ozone layer. NOAA CSL, 2022.

Investigating the Potential Atmospheric Accumulation and Ozone Impacts of Al2O3 from Satellite Reentry. NOAA/CIRES, 2025. https://repository.library.noaa.gov/view/noaa/71517/noaa_71517_DS1.pd

Metals from spacecraft reentry in stratospheric aerosol particles. PNAS, 2023. https://www.pnas.org/doi/10.1073/pnas.2313374120

On the probability of occurrence of extreme space weather events. Space Weather, 2012. https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/

Extreme geomagnetic storms: Probabilistic forecasts and their uncertainties. Space Weather, 2017. https://agupubs.onlinelibrary.wiley.com/doi/

Extreme space weather: impacts on engineered systems and infrastructure. Royal Academy of Engineering. https://raeng.org.uk/media/lz2fs5ql/space_weather_full_report_final.pdf

When trees talk: Researchers probe ancient wood for clues about massive solar storms. University of Arizona, 2024. https://news.arizona.edu/news/when-trees-talk-researchers-probe-ancient-wood-clues-about-massive-solar-storms

Why Tesla's AI trainers don't trust its self-driving tech—or its safety stats. Reuters, May 28, 2026. https://www.reuters.com/investigations/why-teslas-ai-trainers-dont-trust-its-self-driving-tech-or-its-safety-stats-2026-05-28/

Tesla Optimus robot takes a suspicious tumble in new demo. Electrek, Dec. 7, 2025. https://electrek.co/2025/12/07/tesla-optimus-robot-takes-suspicious-tumble-in-new-demo/

Space and satellite insurance solutions. Munich Re. https://www.munichre.com/en/solutions/for-industry-clients/space-and-satellite-insurance-solutions.html

Will We Really Put Data Centers in Space? Forethought, 2026. https://www.forethought.org/research/will-we-really-put-data-centers-in-space

JLL, 2026 global data center outlook; shell-and-core construction cost and AI fit-out ranges.
https://www.jll.com/en-us/insights/market-outlook/data-center-outlook

U.S. EIA monthly electricity prices; 2026 industrial electricity benchmark.

Good Jobs First on data-center tax abatements and revenue losses.
https://goodjobsfirst.org/data-center-tax-abatements-why-states-and-localities-must-disclose-these-soaring-revenue-losses/

NASA, The Recent Large Reduction in Space Launch Cost; Falcon 9 and Falcon Heavy historical launch-cost anchors.
https://ntrs.nasa.gov/api/citations/20200001093/downloads/20200001093.p

Blue Origin New Glenn official vehicle page; reusable first stage and expendable second stage.
https://www.blueorigin.com/new-glenn

Ars Technica reporting on Blue Origin's Project Jarvis reusable upper stage effort and its shelving after New Glenn's debut.
https://arstechnica.com/features/2025/01/after-the-success-of-new-glenn-blue-origin-to-focus-on-launching-frequently/

Reuters, April 2026, on Blue Origin reusing a New Glenn booster while the mission failed because the upper stage underperformed.
https://www.reuters.com/business/media-telecom/blue-origin-says-

it-has-landed-reused-new-glenn-rocket-booster-2026-04-19/

SpaceX Falcon 9 official vehicle page; reusable booster middle-ground reference.
https://www.spacex.com/vehicles/falcon-9/