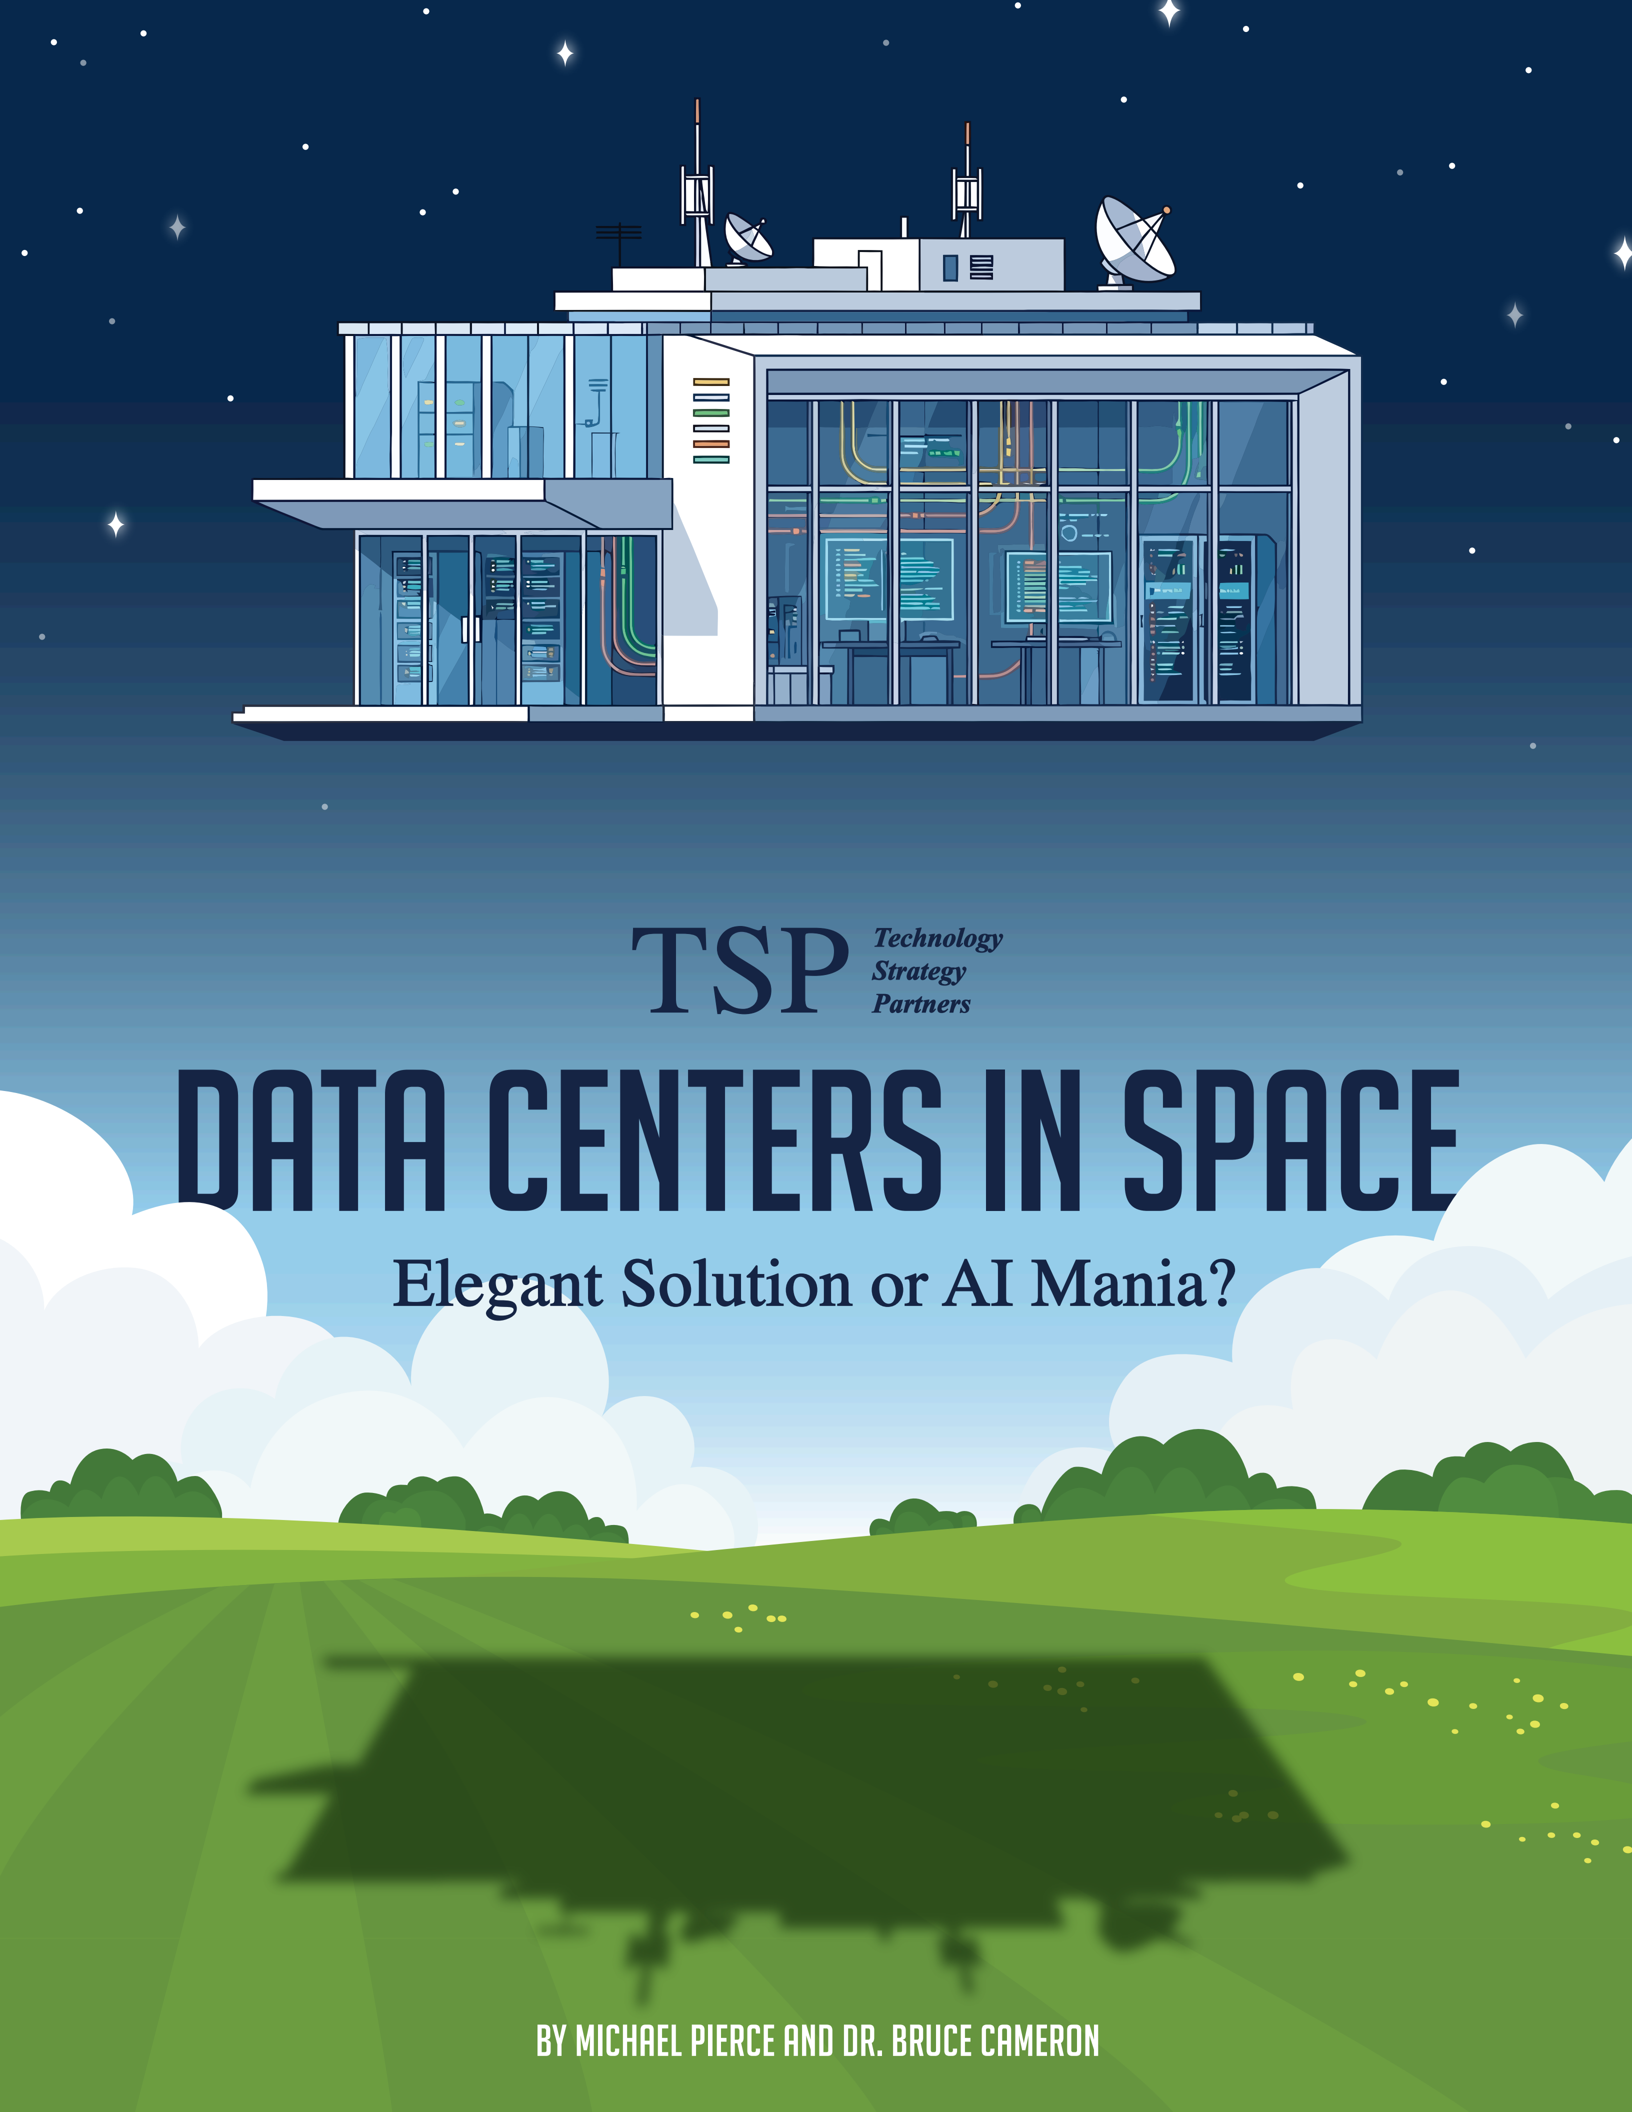

TSP Technology Strategy Partners

DATA CENTERS IN SPACE

Elegant Solution or AI Mania?

BY MICHAEL PIERCE AND DR. BRUCE CAMERON

EXECUTIVE SUMMARY

Data centers reached orbital velocity in 2025. Beyond the hundreds of billions committed to data centers on the ground, at least eight companies with significant backing are now proposing to put data centers in orbit. The pitch is seductive: unlimited solar power, free radiative cooling, no queues for grid power, no water consumption, and no NIMBY resistance. The question is whether the pitch can survive contact with physics and economics.

We built a detailed model to find out. The headline result: an orbital data center costs 5x more than a terrestrial data center today. This gulf is driven by the launch costs and brutal economics of chip lifetimes. We analyze two future scenarios, and conclude that they require radiator and solar panel technology that is 10+ years away, or a launch price of $135/kg, 96% lower than today's SpaceX Falcon 9. SpaceX has set ambitious targets for its new launch vehicle Starship, so data centers in space amount to a long odds bet on SpaceX.

Many past breakthroughs have surmounted the integration challenge, where the technical pieces existed separately, but needed to be architected and assembled to work together, such as the smartphone, Tesla's EVs, drones in the 2010s, and SpaceX's Starlink communications system. Data centers in space are not an integration challenge. This is an R&D challenge, necessitating technology breakthroughs across multiple technology domains, analogous to EUV lithography from 1986-2017, mRNA research from 2005-2020, and AR headset development 1993-2016. The history of space business should throw red flags: the 1990s satellite communications bust of Iridium and Globalstar was driven by a similar mismatch between technical ambition and market reality, and at a far lower capital risk than what orbital data centers are contemplating.

THE PITCH

When the Grid Says No

The global AI infrastructure buildout is placing extraordinary demands on terrestrial infrastructure. There has been 103 GW / $1T of global capacity built to date, with a further 100 GW / $3T estimated by 2030 [1]. One estimate suggests 2030 US AI data center power demand could exceed 50 GW, up from roughly 5 GW today [2]. Despite the DeepSeek R1 breakthrough in computing power, trends suggest larger models and more training expenses, not less.

Power availability from the grid and data center-to-grid shutdowns are already creating tensions. In 2024 alone, utilities in seven states passed more than $4.3 billion in grid upgrade costs onto ratepayers to accommodate data center loads, with billions more pending, Terrestrial Data Centers (TDC) are accused of outbidding home builders for land, spiking local electricity prices, and contributing to water scarcity [3,4].

Water consumption has become an additional bottleneck and a growing political liability. Google's data centers alone consumed 24 billion liters of water in 2024 [5]. In Mesa, Arizona, and Data Center Alley in Virginia, proposed expansions have faced pushback over municipal water infrastructure [6].

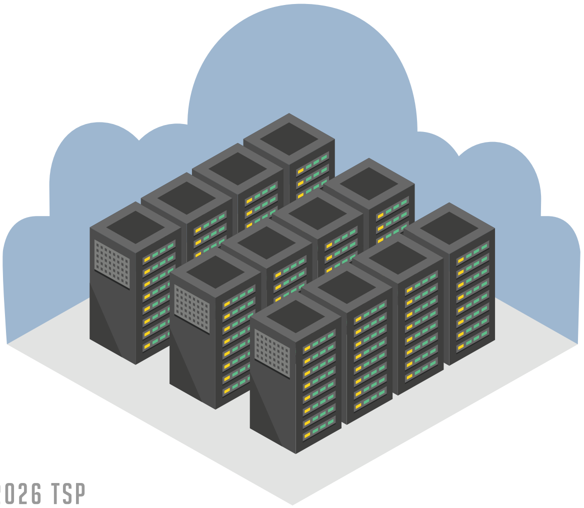

Against this backdrop, Orbital Data Centers (ODC) claim to solve three problems at once:

1. **Space has unlimited power.**

   Solar panels in space capture abundant power with zero marginal cost and have no grid interconnection queue.

2. **Space is cold.**

   More precisely, space offers an enormous heat sink via radiative cooling. Water resource dependence is eliminated, saving both environmental concerns and billions in HVAC costs.

3. **No NIMBYs in space.**

   There are plenty of orbital slots available and not nearly the same regulatory hurdles that are actively slowing data center rollout.

The question is whether the proposed solution - launching tens of thousands of metric tons of hardware into orbit - is remotely proportional to these problems. We will show that these proposals are not one small step for data centers – they are a giant leap.

The Landscape Today

At least eight companies of significant size are actively working on orbital data center plans.[*] The field includes players with the means to fund such an endeavor, and in SpaceX's and Blue Origin's cases, some are vertically integrated with launch capability, often cited as the primary cost driver. With investment across the commercial space sector reaching $45B in 2025, momentum is growing [7].

---

[*] SpaceX, Blue Origin, Starcloud, Google, Sophia Space, Axiom Space, Phantom Space, Aetherflux

The buzz is exciting. After decades of false starts, the commercial space industry has now produced two genuine success stories: communications constellations (Starlink foremost among them) and Earth observation (Planet Labs, Maxar). Space could be having its iPhone moment.

However, the space business remains modest despite these successes. The global commercial satellite communications sector is roughly a $25 billion industry - about half the size of Caterpillar, and approximately 4% of the terrestrial networking market [8,9]. Earth observation, including both optical and synthetic aperture radar, is an order of magnitude smaller at roughly $5 billion per year [10].

Orbital data centers would need to compete in a market measured in hundreds of billions.

OUR APPROACH AND ASSUMPTIONS

To rigorously evaluate the ODC opportunity, we built a detailed model that describes the economics and physics of data centers in space, which we compare against current data centers on the ground.

Three key principles guided the analysis:

1. **Stack the deck in ODC's favor.**

   We make many aggressive assumptions in favor of ODC, so that if the status quo wins out, the conclusion is hard to argue with. If the numbers are close, ODC would merit risk capital willing to take the bet. For example, we exclude solar heating on all surfaces, impose no compute performance penalty for elevated chip temperatures, assume perfect launch vehicle packing, and omit all non-recurring engineering costs – which disproportionately favor the ODC given the new technologies it requires. A full catalog of assumptions is provided in the Appendix.

2. **Transparency of assumptions.**

   Every metric is traced to sources with rationale, and assumptions are explicitly called out. The model is built entirely on publicly available data and published specifications, so that any reader can verify our inputs and challenge our conclusions.

3. **First principles analysis.**

   The model does the math on radiator sizing, compute power required, and sizing of solar panels based on the physics of the problem. For example, we avoided borrowing gross satellite hardware metrics ($/W, kg/W) from Starlink, whose telecommunications mission bears little resemblance to a heavy compute node, and instead decomposed each satellite subsystem into its constituent technologies. By anchoring to physics rather than any specific satellite design, the model sets the rules of the game; designers are free to optimize within them.

Apples to Apples Comparison

This analysis focuses specifically on AI compute data centers, purpose-built for training and inference workloads on GPU clusters, not general-purpose facilities for web hosting, storage, or enterprise IT. The distinction matters: AI data centers are power- and cooling-intensive, with GPU racks drawing 50-130+ kW per cabinet versus 8-12 kW for traditional enterprise workloads [11,12]. We do not distinguish between training and inference workloads, nor do we address latency considerations, though we note that ODC-based inference may hold a structural advantage for applications like autonomous maritime navigation or airborne surveillance processing where users are inherently far from terrestrial data centers.

Both the TDC and ODC cases were sized to have equivalent compute power, measured in watts. We model a 1 GW data center, on par with OpenAI's Stargate Abilene campus [13]. We decompose both the TDC and ODC into

COMPONENT	TERRESTRIAL SCOPE	ORBITAL SCOPE
Power	**OpEx** for electricity consumption from the grid	**CapEx** for solar panels
Compute	**CapEx** for GPUs and all associated IT equipment to run training and inference **OpEx** for failed GPU replacement and network/software costs	**CapEx** for GPUs, IT, spares
Cooling	**OpEx** for water consumption	**CapEx** for radiators
Deployment	**CapEx** for "powered shell:" land, building, HVAC, compute cooling, grid connection **OpEx** for all equipment maintenance, staffing, and property taxes	**CapEx** for integration hardware (propulsion, communication, structures) and launch **OpEx** for constellation operations and GPU refreshes every 5 years

four components : Power, Compute, Cooling, and Deployment (i.e. integration and operations).

Inputs are transparently sourced and cited, including NVIDIA datasheets, JLL construction cost guides, Turner & Townsend cost indices, NASA radiator research, and SpaceX pricing. Outputs are validated against nine major data center builds from 2017-2026, including xAI Colossus, Meta RSC, Oracle superclusters.

To understand how the economics might improve over time, we examine 3 *technology scenarios*:

SCENARIO 1: TECHNOLOGY TODAY

These are technologies that are commercially available today. We assume GaAs III-V solar panels, passive deployable radiators, and 2025 Falcon 9 capabilities (17.5T payload, $4000/kg, 167 Falcon 9 launches/year) [14].

SCENARIO 2: NEXT GENERATION

We classify this scenario as technologies that exist in prototype form but are not yet commercially available, but which may be ready in 5 years. We assume the adoption of lower-cost solar panel technologies adapted from terrestrial manufacturing (40% reduction in price), a hotter GPU (126C / 260F), and Starship comes into initial operational cadence and performance (100T payload, $900/kg, 52 Starship launches/year).

SCENARIO 3: FUTURE CONCEPT

The set of technologies where principles have been demonstrated but which will need 10 or more years of development. We assume simultaneous orders of magnitude improvements in solar panel specific mass (from 10 g/W to 0.33 g/W), increased radiator operating temperature (226C / 440F), and Starship achieves all its ambitions (200T payload, $100/kg, 365 Starship launches/year).

TODAY	TDC	ODC
CapEx	$48B	$238B
Annual OpEx	$6.4B	$32.9B
20-year Total Cost	$102B	$518B

RESULTS

Today: Not Even Close

Simply put, current technology suggests ODCs are not economically viable.

ODC costs are dominated today by launch cost at $131B - boosting satellites into space at 17,500 mph is expensive. At nearly 33,000 metric tons, this represents over 10 times the current annual global mass to orbit, or 5 Starlink constellations [15]. Solar panel weight alone is over 10 tons, while radiator mass is nearly 8 tons. Notably, the hardware material costs are $107B – in other words, it costs more to send materials to orbit than it does to make the materials themselves.

The cost of constellation-scale space-grade solar power represents the highest portion of material cost at $50B. For 1 billion watts of power to generate, the $50/W figure we use (on the low-end of current market prices) contributes substantially [16]. GPU hardware comes in as the next most significant capital expense at $24B. However, the largest factor of total cost comes from the recurring annualized cost of $33B; satellite and GPU lifetimes mean refreshing the entire constellation every 5 years.

Key drivers for the baseline 1 GW terrestrial data center are GPU capital cost at $24B, the associated non-GPU IT hardware at $12B, and the building's "powered shell" upfront cost at $12B. Replacing failed GPUs costs $4.7B/year, only 14% of the cost in space.

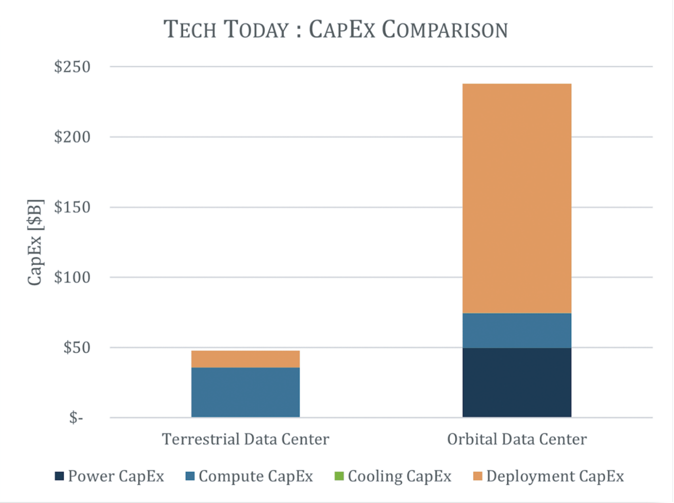

Next Gen: Closing the Gap, But Not Enough

The Next Generation scenario assumes that starting in 2031, launch costs fall by >75% with Starship hitting its stride through initial operations. Further, we assume techniques for qualifying terrestrial silicon solar panels come to the mass market (a 40% reduction in cost), and GPUs operate at 50C hotter temperatures than today with 10% lower cost and 17% lower density.

The net result is that ODC costs are cut by 72%, but they are still 24% more expensive than building on the ground.

NEXT GEN	TDC	ODC
CapEx	$56B	$90B
Annual OpEx	$7.2B	$6.6B
20-year Total Cost	$117B	$145B

The drop in launch prices is the largest lever, driving tremendous reductions in initial launch expense from $131B to $12B. Additionally, the estimated 40% reduction in solar panel cost would save another $20B. It's important to note that panel efficiency doesn't matter as much as cost or mass per watt - in a mega constellation, each satellite node can be sized to the power the packaging can deliver.

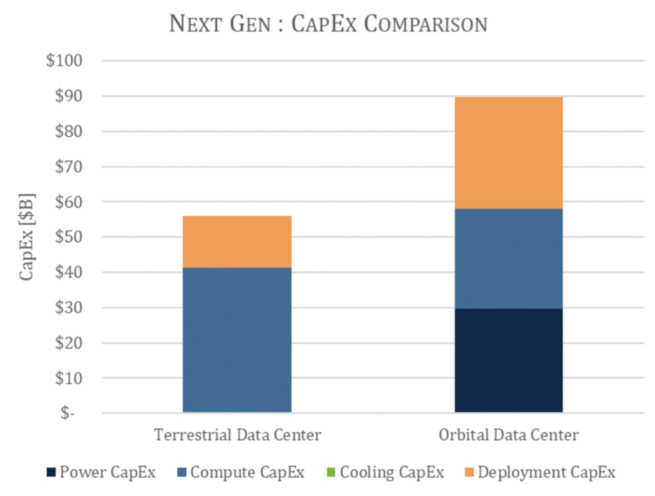

Note that TDC costs by comparison are expected to go up, not down. This is primarily due to power costs rising (from $0.06 to $0.08 per kWh), GPU costs rising (albeit, with greater computational power), as well as increases in maintenance and staffing costs beyond inflation rates.

This forecast suggests there is potential for convergence, but based on the favorable ODC assumptions made, our conclusion is that there aren't enough levers to make ODCs cost competitive in the next 5 years.

Future Concept: A Glimmer of Hope

In the third scenario, we assume Starship achieves its ambitious targets in mature operations and reuse (a 98% reduction in $/kg today), dirt-cheap space-ready solar panels are plentiful ($15/W), and GPUs push their operating temperature 150C hotter than today. All three of these are pushing the limits of physics and demand fundamental new technologies; at best these assumptions represent a 2036 world. For a fair comparison in 2036, we inflate TDC costs as well, due to rising electricity prices, land prices, chip costs, and operational costs.

Total ODC system mass falls by 87%, driven primarily by solar panel and radiator technology improvements. Combined with a 98% total reduction in launch price (from $4,000/kg to $100/kg), the effect compounds: the Future Concept 20-year Total Cost drops to $83B, roughly 84% below today's total cost of $518B.

With the extreme launch capacity and cadence of Starship, Future Concept ODCs also achieve full operation much faster than TDCs - in a matter of weeks, not years. This alone might be a winning argument: the ability to deploy new models on a dedicated data center may decide who wins the AI race. Lower launch costs mean constellation refreshes are more affordable, the ODC's

FUTURE CONCEPT	TDC	ODC
Total CapEx	$66B	$60B
Total Annual OpEx	$8.3B	$2.7B
20-year Total Cost	$136B	$83B

computational performance can match or even exceed its TDC counterpart at lower cost.

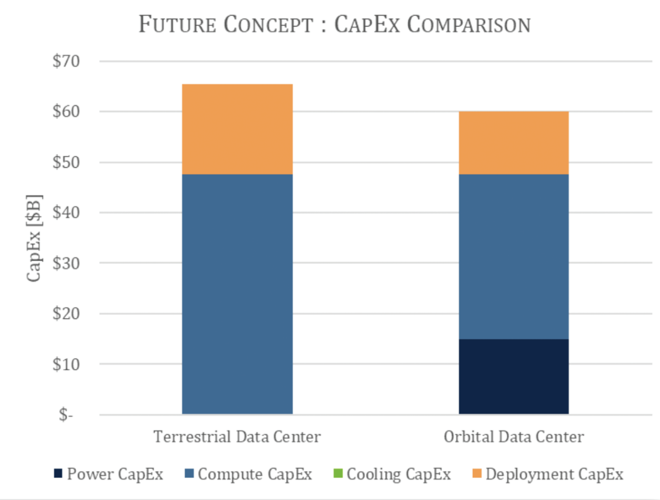

The $135/kg Question

The price of satellite launches ripples through initial constellation CapEx and refreshes. It is the single biggest lever for data centers in space: every kilogram of mass is amplified by expensive launch costs. The path to ODC viability seemingly runs through the rocket equation, not Moore’s Law. So at what launch price would an ODC become viable?

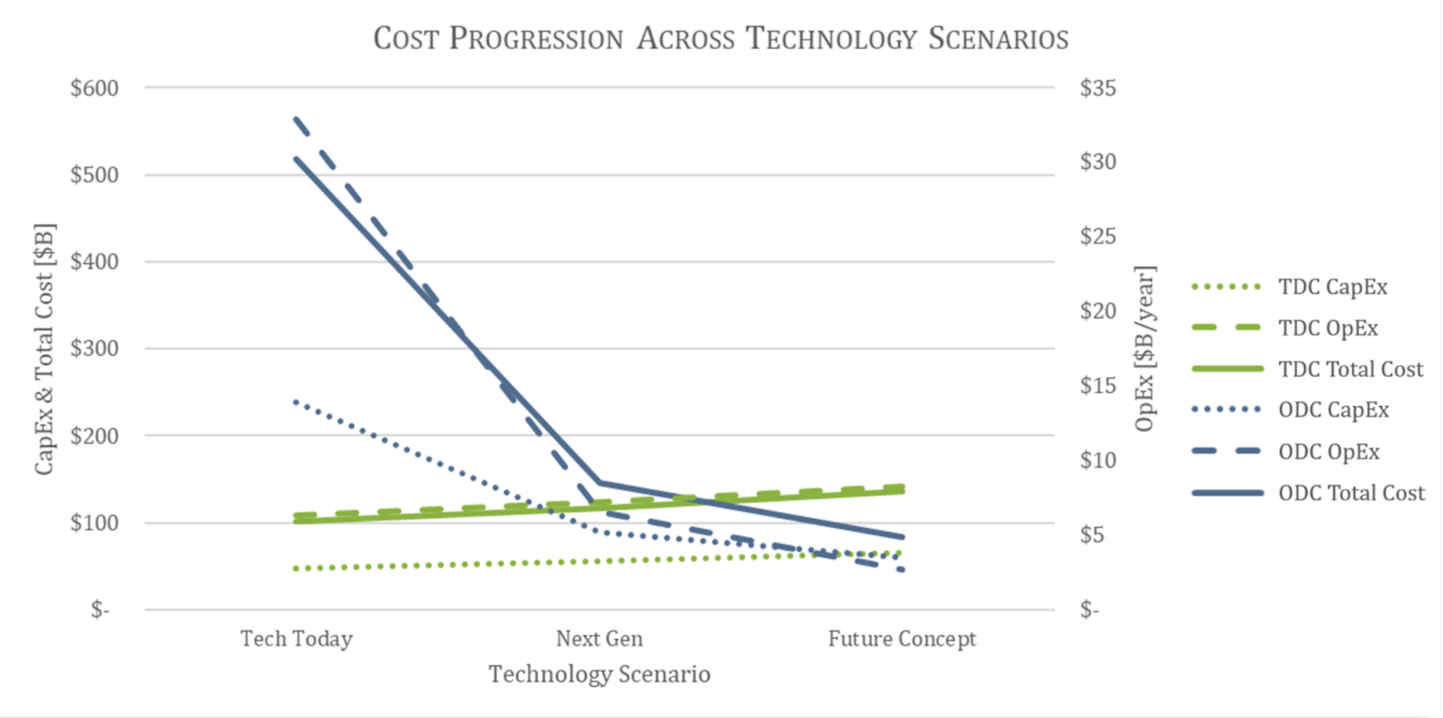

At the current state of GPU costs, infrastructure capabilities, and satellite technology, the answer is never. Launch could be free, and yet the remaining costs of manufacturing and operating satellites outweighs the cost of building a data center on earth and maintaining it.

For the next generation scenario (2031), we start to see ODCs become viable around the $135/kg threshold. For perspective, this is a 96% reduction over today's list price of space launch. This only holds if the cost of capital is 10%. If a 15% cost of capital is used to reflect the underlying technology and market risks, it makes the ODC unviable at any launch price.

In the future concept scenario (2036), ODCs leverage ultralightweight solar panels and very hot GPUs. If these ambitious technology goals are achieved, orbital data centers are even viable at today's SpaceX Falcon 9 list price of $4000/kg.

The Hidden Cost of Keeping Up

The minute a chip is launched into orbit, it is on its way to obsolescence. GPU performance has improved roughly 7× in efficiency (FLOPS per watt) over the last five years, with clear two-year generational cycles [17]. Jensen Huang of NVIDIA has been contradictory on this point, arguing on one hand that last year's Hopper chips are already obsolete, while also assuring customers that their GPU investments represent long-term assets.

In a terrestrial data center, the capital investment is 50% chip-related costs. The remainder is long-lived infrastructure—cooling systems, racks, power distribution—that changes far less frequently. When a new GPU generation arrives, TDCs can swap the boards and keep the building. We assume this GPU upgrade happens every 5 years, consistent with recent accounting treatment at data centers [18]. The key benefit of a TDC is that the infrastructure's 20-year useful life amortizes across multiple GPU generations [19].

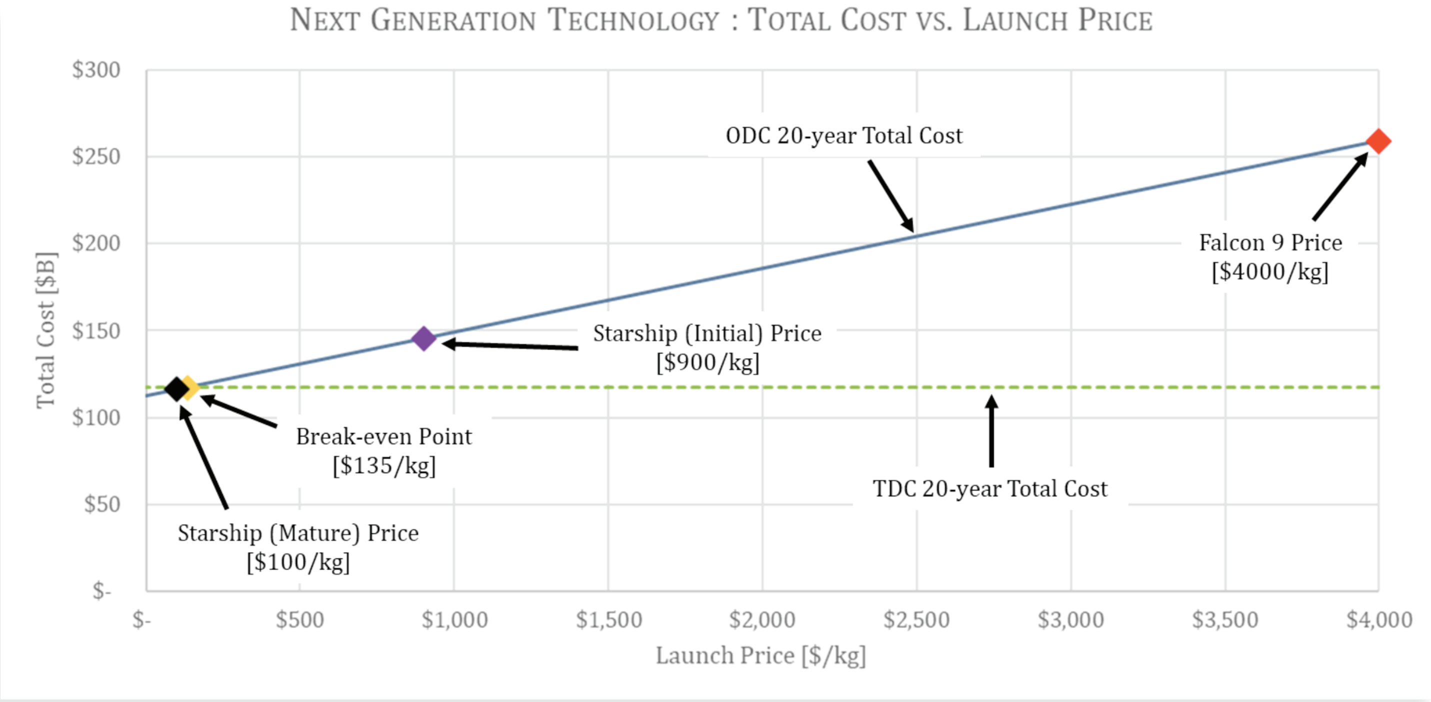

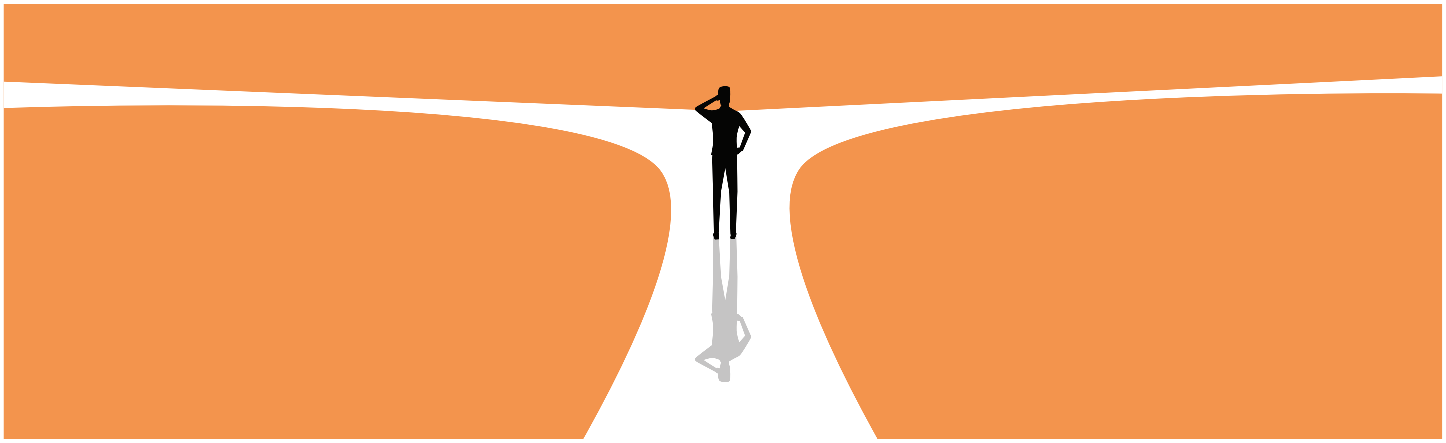

In orbit, the satellite lifetime and the chip lifetime are synonymous. This means a GPU refresh requires either replacing entire satellites or performing modular chip swaps in orbit (a technique that has yet to be demonstrated). Although chips could theoretically be designed for “self-eject” and plug-and-play replacement, we judge that ODCs are far more likely to pursue modularity at the satellite level, not the chip level. This means a GPU refresh requires replacing entire satellites. With a 5-year useful lifetime of a mega-constellation satellite, the entire constellation must be replaced 3 times to equate to the same useful life of a 20-year TDC [20].

Space is famously punishing : “tin whisker” growth causes short circuits, solar weather damages electronics, and micrometeorites and space debris remain an ever-present risk. We judge on-orbit self-assembly is at best a Future Concept (10-20 years away), and we do not model it here. Solar energy is free, but replacing the satellites containing the chips is not.

PICK YOUR PARLAY

For orbital data centers to reach cost parity with the incumbents in Virginia’s Data Center Alley, our analysis suggests two possible paths. Neither is easy, and both require multiple simultaneous breakthroughs.

Path 1: The Rocket Bet.

Starship achieves <$135/kg to orbit, next-generation solar panels drop roughly 40% in cost and mass, and GPU operating temperatures advance modestly beyond today's limits. For a startup without revolutionary IP, the point of departure is that Elon Musk’s projections come true on schedule. This path matches the technology progress in the “Next Generation” scenario plus Starship delivering on its most aspirational promises.

Path 2: The Physics Bet.

Launch costs stay roughly where they are, but a firm solves the “thermal problem” with order-of-magnitude radiator mass reductions and develops ultralight and cheap space-grade solar panels at scale. This is the "reinvent everything but the rocket" path, matching our “Future Concept” scenario without the Starship promises.

We examine the key assumptions in both paths below.

Path 1 Assumption : Starship Delivers on Audacious Promises

The biggest lever on the ODC business case is launch costs. At current Falcon 9 pricing, launch alone consumes 55% of CapEx. Starship must deliver an order-of-magnitude improvement over the vehicle that itself revolutionized the launch industry. Initial flight tests carried 35 metric tons, but the next generation aims for 100T, and Elon Musk's ultimate goal is 200T [21]. Musk is known for his use of stretch goals to motivate teams, and also for a pattern of overpromising and underdelivering (the 4680 battery cell, "Full Self-Driving is next year", Dragon propulsive landing).

In addition to launch price, ODCs require available launch capacity. A 1 GW data center in space would require 1,875 Falcon 9 launches or 42 Starship launches. At the current rate of 167 Falcon 9 launches, that would consume 100% of SpaceX capacity for over 11 years. If Starship launches 200T per day, the same deployment compresses to 3 weeks. The question then becomes whether that timeline aligns with ODC deployment plans. The ODC business case runs through a very small number of companies. Due to their vertical integration, SpaceX and Blue Origin may be the only entities capable of economically building ODCs in the next 15 years, meaning viability depends not just on rocket physics but on the strategic priorities of one or two firms. SpaceX has achieved the impossible before. The question is whether that track record justifies an all-in bet.

Path 2 Assumption : Radiators Shrink by an Order of Magnitude

Radiator mass grows to dominate ODC satellite mass as other subsystems improve, becoming a binding constraint on long-term viability. Radiative heat flux scales with temperature to the fourth power (T^4), making operating temperature the biggest nonlinear lever available. At current sustained GPU operating temperatures of 77C / 170F, a 1 GW data center needs 320 acres of radiator surface – 241 football fields. The fundamental physics limit for bulk silicon CMOS is around 175C / 346F before reliability mechanisms degrade unacceptably [22]. Wide-bandgap SiC/GaN technologies could push this to 227C / 440F or higher, shrinking total radiator area by roughly 77%. These are not engineering refinements; they are material science breakthroughs unlikely to arrive on the timescale of a venture capital fund.

The alternative to operating at higher temperatures is active cooling technologies, where heat pipes and fluid pumps allow the GPUs to run at current temperatures while radiators operate at higher temperatures. This is a proven technology on large spacecraft like the International Space Station, but it adds substantial mass, complexity, and cost not included in our analysis. Most satellite constellations today rely on passive radiators, and scaling active thermal management at the heat densities required by GPU compute remains an open engineering problem.

Unlike solar panels and semiconductors, where decades of investment across thousands of applications have driven costs down learning curves, space radiators serve a small customer base with no equivalent of Swanson's Law. The ODC industry would need to self-fund this development rather than inheriting gains from adjacent industries.

The semiconductor challenge is equally daunting. SpaceX's planned Terafab facility aims to manufacture advanced-node processors at a scale rivaling the world's leading foundries. Those foundries have spent decades and hundreds of billions of dollars reaching their current capabilities. Compressing that learning curve into years would be without precedent.

Path 2 Assumption : Solar Panel Weight and Cost

Solar panels have improved dramatically over recent decades, but the relevant gains are leveling off [23]. Today's space-grade standard gallium arsenide multi-junction cells are highly efficient but heavy and expensive. Our current-state analysis requires over 10,000 metric tons of solar arrays for a single 1 GW data center, roughly 10x the entire global space solar panel industry's annual production [24].

The next step down in cost and mass would be adapting terrestrial silicon cells for space. Starlink has done this with custom-built panels that sacrifice some efficiency for manufacturability, but these are proprietary and not commercially available [25]. An ODC operator would need to invest in or contract for a similar supply chain, but the technical path exists.

Reaching the future-concept envelope requires fundamentally new materials and manufacturing techniques that produce ultra-thin, lightweight panels at volume. Promising prototypes exist, but bringing novel materials from demonstration to the quality, volume, and price required for gigawatt-scale deployment within a 10-year window demands sustained investment with no guarantee of success.

The Reward Doesn't Match the Risk

For a substantial risk, equity holders would demand a substantial reward – either a durable cost advantage or the ability to lock competitors out through pricing, IP, or supply monopolies. Even if the technical bets pay off, the competitive moat is thin. Radiator or solar breakthroughs would yield IP, but neither is likely to prove a durable competitive advantage if orbital data centers become a real market. The strongest moat is probably a supply monopoly on launch services, which circles back to dependence on one of two providers.

The research on technology development is instructive: breakthrough cost reductions in solar cells, semiconductors, and other foundational technologies were largely incubated by government programs over decades, not by private risk capital over years. The gains private firms capture tend to be integration wins, assembling known technologies in new configurations (e.g. Tesla's EVs, drones in the 2010s, and Amazon Echo and Alexa) rather than fundamental advances. ODCs require not just integration but substantial R&D across multiple domains simultaneously.

Photo credit Steve Jurvetson, cropped, licensed under CC BY 2.0
https://creativecommons.org/licenses/by/2.0/deed.en

The history of commercial space reinforces the warning. Iridium and Globalstar in the late 1990s were technically successful but were wrong-footed by assumptions about cost trajectories and market adoption. Iridium filed for bankruptcy with $4 billion in debt. The parallel is uncomfortable: both propositions require massive upfront capital, depend on cost assumptions that have not materialized, and compete against a terrestrial alternative that is itself improving rapidly. The difference is that ODCs contemplate far more capital at risk than 1990s satcom ever raised.

We believe SpaceX may be the only entity that could plausibly attempt data centers in space. They have a stellar workforce, the ability to supply launch at internal cost, and potentially the capital. Starlink is a great act to follow — it has singlehandedly doubled the number of satellites in orbit, with annual revenue of $12B in 2025 — but Starlink competes with slow-moving telecom incumbents, not hyperscalers, and did not require the same magnitude of power and thermal technology development [26]. Even SpaceX would need to execute order-of-magnitude improvements across three simultaneous technology domains (solar, chips, and launch). That is not an integration challenge; it's a Manhattan Project without the government.

Niche Applications: Where ODCs Might Shine

There is a narrower case for in-orbit computation that has nothing to do with competing with terrestrial data centers. Processing data where it is generated—running compression algorithms on Earth observation imagery, performing edge inference for autonomous spacecraft, or filtering sensor data before downlinking—has genuine value for the existing space industry. The existing space business is not large enough to justify ODC-scale investment. The communication links between LEO constellations and ground stations are not near capacity saturation, and are being rapidly built out. The niche is real but small.

Proponents often highlight regulatory arbitrage as a key advantage: orbital data centers would not be bound by any single nation's environmental, zoning, or energy regulations. While this freedom from ground-based constraints is real, operators must still navigate the jurisdiction of their launching state and an evolving international regulatory landscape that may someday recognize the tragedy of the commons problem of rampant space debris in orbit [27].

CONCLUSIONS

This analysis leads to three conclusions.

1. Orbital data centers are not economically viable today, and the gap is not close.

At current technology, an ODC costs 5× more than a terrestrial equivalent. The driver is the compounding weight of mass-to-orbit economics: every kilogram of solar panel, radiator, and GPU must be launched at thousands of dollars per kilogram, then relaunched every five years to maintain compute parity with a building that lasts twenty years. The "zero OpEx" pitch of free solar and cold vacuum of space collapses when chip upgrades are accounted for. Even if launch were free, current-state ODCs never break even. These results, moreover, were derived under assumptions favorable to ODCs. A prudent investor would increase the discount rates to reflect the risk, allow for a 3% rate of satellite failures (as seen with early Starlink), assume chip radiation degradation, and account for other operational realities not modeled here [28].

2. Viability requires simultaneous breakthroughs, not incremental progress.

Sensitivity analysis identifies two possible paths to cost parity, and neither is straightforward. The Rocket Bet requires Starship to achieve $135/kg while next-generation solar and thermal technologies materialize on schedule. The Physics Bet requires revolutionary advances across GPU thermal limits, ultralight solar, and radiator technology, all within a decade. Both paths assume that terrestrial innovation stagnates by comparison. These are not integration challenges waiting for a clever architect; they are concurrent R&D programs – individually uncertain, collectively unlikely.

3. Terrestrial capabilities are unlikely to stagnate.

The case for orbital data centers implicitly assumes that terrestrial alternatives stand still while space technology leaps ahead. This is the wrong baseline. Improvements in solar panel efficiency, semiconductor manufacturing, and industrial automation benefit traditional data centers at least as much as orbital ones. ODC proponents are betting that terrestrial power bottlenecks persist and that alternative on-site energy never becomes viable. The evidence runs the other way. Heron Power is developing solid-state transformer technology that could reduce grid-to-chip energy losses by 4×. Crusoe is pioneering closed-loop zero-water cooling for OpenAI's Stargate campus. Boom Supersonic recently partnered with Baker Hughes to provide 1.21 GW of on-site jet-engine-derived power generation. Small modular reactors remain on the horizon. The data center industry commands a massive installed base of engineering talent, manufacturing capacity, and capital. Betting against this is betting against the U.S. industrial economy.

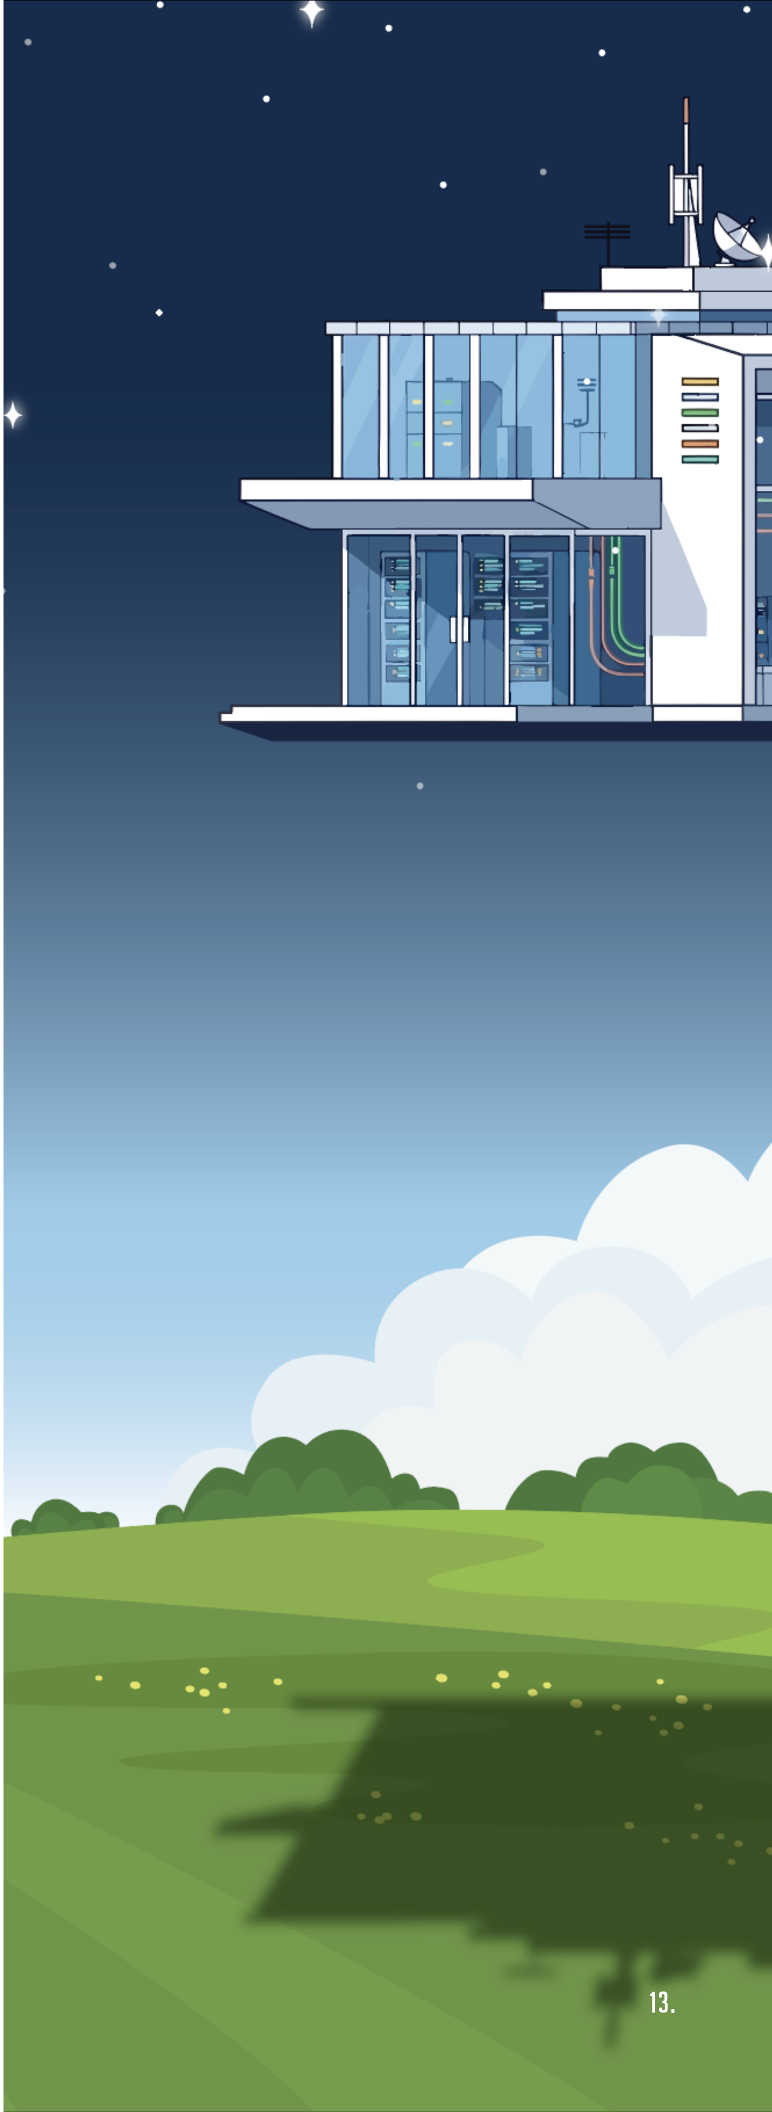

MICHAEL PIERCE IS A PRINCIPAL AT TSP IN DENVER, CO.
DR. BRUCE CAMERON IS A PARTNER AT TSP IN BOSTON, MA.

APPENDIX

Orbital Environment and Satellite Design

ASSUMPTION	EFFECT ON RESULTS
Sun-synchronous orbit with 100% solar illumination; zero eclipse time and zero battery mass.	Favors ODC. Lower orbits with eclipse periods would require batteries, adding mass and cost.
Mega-constellation of 10,000+ satellites, consistent with the scale most proponents are proposing. Provides aggregate performance and redundancy.	Favors ODC. Smaller constellations would perform worse on aggregate performance and fairing packaging.
Passive radiators require no active power. Solar heating on radiator surfaces is excluded.	Favors ODC. Actual thermal coupling depends on spacecraft geometry and orientation; real systems would reject less heat per unit area.
Solar cells do not degrade over the satellite lifetime.	Favors ODC. GEO-grade cells lose ~1–2% per year; LEO radiation belts will accelerate degradation.
Each satellite requires the same effective power as a terrestrial server node (~1,200 W per GPU). Satellite bus power (comms, GNC) is assumed to fit within this envelope rather than being added on top.	Approximately neutral. Understates ODC power slightly, but the margin is small relative to compute loads.
Higher radiator operating temperatures improve heat rejection per unit area and do not degrade GPU performance and reliability.	Favors ODC. Running hotter reduces radiator mass but real chips deliver fewer FLOPS per watt at elevated temperatures.

Launch and Deployment

ASSUMPTION	EFFECT ON RESULTS
ODC operators can purchase 100% of available launch capacity in a given year.	Favors ODC. Real manifest competition would stretch deployment timelines.
Launch cost is calculated on a per-kilogram basis with full utilization of payload volume and mass. Orbital maneuvering from injection orbit to operational slot is excluded.	Favors ODC. Real packing inefficiencies and on-orbit propulsion would add cost.
GPU failure policy assumes spare GPUs are brought up to orbit but turned “off” until needed.	Neutral. Assumes the satellite is sized appropriately for failures.

Compute and Supply Chain

ASSUMPTION	EFFECT ON RESULTS
All compute workloads are treated as equivalent. No distinction is made for latency-sensitive vs. latency-tolerant applications.	Favors ODC. In practice, ODC is better suited to latency-tolerant workloads (batch training, agentic services) than real-time inference.
Latency penalties and radiation-induced bit errors in LEO are excluded.	Favors ODC. Both effects would reduce ODC throughput.
ODC operator can solve coefficient-of-thermal-expansion (CTE) mismatch challenges for space-qualified GPU packaging.	Favors ODC. NVIDIA has encountered CTE issues in terrestrial liquid cooling; space thermal cycling is more severe.
Operators purchase as many GPUs as desired. Satellite production rate is not a bottleneck.	Approximately neutral. Enables the ODC to scale in ways that may not be realistic near-term.
GPU uptime and compute utilization at 100%.	Approximately neutral. Both architectures would experience downtime, but ODC failures are harder to remediate.

Terrestrial Data Centers

ASSUMPTION	EFFECT ON RESULTS
TDC electricity price escalation is conservative and does not account for behind-the-meter generation (on-site solar, gas turbines, SMRs).	Penalizes TDC. Behind-the-meter power could reduce TDC energy OpEx below grid rates.
TDC construction cost ($/W) does not decrease with scale. In practice, larger campuses tend to achieve lower unit costs.	Penalizes TDC. Real hyperscale economies of scale would widen the gap.
TDC OpEx buildup combines several proxy sources to approximate the right order of magnitude. Individual line items may not project technology cost curves accurately over time.	Direction uncertain. Errors could favor either architecture.

Financial and Scope

ASSUMPTION	EFFECT ON RESULTS
Non-recurring engineering, R&D costs, and the time required to achieve projected technology targets are excluded.	Favors ODC. The ODC requires far more novel engineering, notably space-qualified GPU packaging, autonomous operations, radiator manufacturing at scale.
Manufacturing facilities, embodied energy, and launch vehicle production costs are outside the system boundary.	Approximately neutral for the comparison, but understates absolute costs for both.
All costs are in real (inflation-adjusted) dollars.	Neutral.

REFERENCES

[1] Global data center sector to nearly double to 200GW amid AI infrastructure boom. JLL.
https://www.jll.com/en-us/newsroom/global-data-center-sector-to-nearly-double-to-200gw-amid-ai-infrastructure-boom
[2] Artificial intelligence power demand in US could top 50 GW by 2030. S&P Global.
https://www.spglobal.com/energy/en/news-research/latest-news/electric-power/081325-artificial-intelligence-power-demand-in-us-could-top-50-gw-by-2030-epri
[3] PJM Data Center Issue Brief. Union of Concerned Scientists.
https://www.ucs.org/sites/default/files/2025-09/PJM%20Data%20Center%20Issue%20Brief%20-%20Sep%202025.pdf
[4] Data center land deals and housing shortage. Wall Street Journal.
https://www.wsj.com/real-estate/data-center-land-deals-housing-shortage-81ea6e09
[5] Google emissions jump due to AI data center boom. Data Center Dynamics.
https://www.datacenterdynamics.com/en/news/google-emissions-jump-48-in-five-years-due-to-ai-data-center-boom/
[6] Data centers' water usage concerns in Arizona communities. YourValley.
https://www.yourvalley.net/chandler-independent/stories/data-centers-thirst-for-cooling-water-raises-alarms-in-arizonas-desert-communities,630751
[7] Why low Earth orbit is attracting billions in investment. CNBC.
https://www.cnbc.com/2026/03/22/why-low-earth-orbit-is-attracting-billions-in-investment.html
[8] Satellite communication market analysis. Global Market Insights.
https://www.gminsights.com/industry-analysis/satellite-communication-market
[9] Telecom equipment global market report. The Business Research Company.
https://www.thebusinessresearchcompany.com/report/telecom-equipment-global-market-report
[10] Commercial Earth observation market report. Credence Research.
https://www.credenceresearch.com/report/commercial-earth-observation-market
[11] AI data center design and deployment trends. Schneider Electric Blog.
https://blog.se.com/datacenter/2025/07/23/ai-data-center-design-and-deployment-are-moving-at-an-incredible-pace-3-ways-to-approach-a-changing-landscape/
[12] The path to power in data centers. Data Center Dynamics.
https://www.datacenterdynamics.com/en/marketwatch/the-path-to-power/
[13] Power is the moat: gigawatt-scale data centers. Global Data Center Hub.
https://www.globaldatacenterhub.com/p/power-is-the-moat-how-gigawatt-data
[14] Falcon 9 launch vehicle. Wikipedia.
https://en.wikipedia.org/wiki/Falcon_9
[15] Global space launch metrics and reports. BryceTech.
https://brycetech.com/reports
[16] Multi-junction gallium arsenide solar cells market report. Market Report Analytics.
https://www.marketreportanalytics.com/reports/multi-junction-gallium-arsenide-solar-cells-237612

[17] Machine learning hardware energy efficiency trends. Epoch AI.
https://epoch.ai/data-insights/ml-hardware-energy-efficiency
[18] AI GPU depreciation and lifecycle discussion. CNBC.
https://www.cnbc.com/2025/11/14/ai-gpu-depreciation-coreweave-nvidia-michael-burry.html
[19] Data center lifecycle and real estate trends. Principal Asset Management.
https://www.principalam.com/us/insights/real-estate/powering-digital-age-data-center-trends-and-opportunities
[20] Orbital data center analysis paper. arXiv.
https://arxiv.org/html/2603.25835v1
[21] SpaceX Starship launch vehicle. Wikipedia.
https://en.wikipedia.org/wiki/SpaceX_Starship
[22] High-temperature characterization of materials up to 450°C. ASM International.
https://meridian.allenpress.com/jmep/article/10/2/67/36791/High-Temperature-Characterization-up-to-450-C-of
[23] Solar cell efficiency records. NREL.
https://www.nrel.gov/pv/cell-efficiency.html
[24] Space-based solar power overview. Per Aspera.
https://www.peraspera.us/space-power/
[25] Satellite solar supply chain bottleneck. SpaceNews.
https://spacenews.com/modernizing-the-satellite-supply-chain-by-breaking-the-solar-power-bottleneck/
[26] Starlink revenue forecast and military contracts. SpaceNews.
https://spacenews.com/starlink-set-to-hit-11-8-billion-revenue-in-2025-boosted-by-military-contracts/
[27] Kessler syndrome (orbital debris risk). Wikipedia.
https://en.wikipedia.org/wiki/Kessler_syndrome
[28] Starlink satellite failure rates analysis. Universe Today.
https://www.universetoday.com/articles/about-3-of-starlinks-have-failed-so-far