Orbital Computing Can Reach $1 Trillion Addressable Market by 2030

Evaluating the Cost Curves Shaping AI's Future in Space

Author
Brendan Burke, Research Director

APRIL 2026

Key Points

- SpaceX Starship is on a trajectory to push launch costs below $100/kg by the end of the decade through full reuse and high flight cadence, collapsing the single largest cost input for orbital infrastructure. However, launch capacity will remain supply-constrained, with Starship slots contested across Starlink V3 deployment, lunar missions, and commercial payloads.
- Orbital data centers offer compelling structural advantages, but cost parity will arrive first for off-grid deployments, not for all AI compute. Grid-connected data centers, which benefit from existing electrical infrastructure and established construction supply chains, will remain the lower-cost option for the majority of workloads.
- If $3 trillion is dedicated to AI compute capex in 2030, we estimate roughly one-third – approximately $1 trillion – represents workloads where orbital deployment is economically justified: capacity facing multi-year grid queues, sovereign compute mandates in power-constrained geographies, and incremental demand beyond what terrestrial infrastructure can absorb. This is a massive market that will remain supply-constrained by launch throughput and satellite manufacturing capacity.

Recommendations for Vendors in Orbital Computing

1 **Define Your 'Space-Case' Addressable Market with Precision**. The orbital compute opportunity is not a wholesale replacement for terrestrial infrastructure – it is a structurally distinct segment defined by off-grid economics, power-constrained geographies, and grid interconnection bottlenecks. Vendors should model their total addressable market against the off-grid buildout pipeline specifically. Overstating the near-term opportunity risks misallocating R&D capital; understating it risks ceding a trillion-dollar market to faster movers.

2 **Communicate Orbital Strategy Now, Even if Deployment is Years Away**. Over the next decade, the compute infrastructure ecosystem will bifurcate between vendors with credible orbital roadmaps and those without. Vendors that align engineering investment and go-to-market positioning with emerging orbital opportunities will benefit from the fast pace of ecosystem formation.

3 **Pioneer Space-Hardened Compute and Optical Interconnect at Commercial Scale**. Silicon providers and networking vendors should invest aggressively in creating radiation-hardened GPUs, adaptive SoCs, and optical inter-satellite mesh networking systems. Many innovations required for space, including thermal efficiency and token-per-watt optimization, have direct applicability to terrestrial edge and off-grid deployments, creating dual-use R&D leverage.

What You Need to Know: Orbital Computing's Substitution Effect

Orbital computing pioneers anticipate that space data centers will become economically viable by 2028 or 2029. Elon Musk predicted at the World Economic Forum that orbital AI will be cheaper than terrestrial in 2–3 years. Starcloud CEO Philip Johnston raised $100 million in part due to expectations of shipping $1 trillion of orbital compute to orbit within the next 10 years, as laid out in Sequoia Capital's Training Data Podcast. These calculations take precedence from engineering breakthroughs at SpaceX to improve the economics of space launch and from assumptions about the worsening economics of terrestrial data centers.

While a hyperscale data center can be constructed in 12 to 18 months, acquiring grid-connected power generation often takes three to seven years, stranding billions of dollars in finished but unpowered silicon. To keep pace with AI workload demand that may require 300 GW of compute by 2030, operators are desperately pivoting away from reliance on the grid. In response to this crisis, two distinct and highly capital-intensive strategies have emerged to house the next generation of AI factories.

The first is the massive acceleration of terrestrial off-grid AI factories. Hyperscalers are increasingly moving behind the meter, pairing advanced liquid-cooled data centers with dedicated natural gas plants, fuel cells, or nuclear Power Purchase Agreements (PPAs). The second, far more radical strategy is the aggressive push toward orbital data centers. Backed by silicon giants alongside launch providers, the concept of space-based compute has rapidly transitioned from science fiction to a funded industrial race.

Orbital data centers promise limitless, unattenuated solar power and the infinite cooling sink of deep space, entirely bypassing Earth's regulatory and permitting gridlock. However, deploying a gigawatt-scale AI factory in Low Earth Orbit (LEO) incurs a 3–4x economic premium due to exorbitant launch costs and custom satellite manufacturing. As the industry stands at this infrastructure crossroads, the comparative unit economics of a gigawatt terrestrial off-grid facility and a gigawatt orbital data center are the most critical calculations for the future of cloud strategy.

Analysis

Deconstructing the 1 GW Off-Grid AI Factory

To achieve 1 GW of independent power, operators must deploy on-site generation, typically utilizing Combined Cycle Gas Turbines (CCGT) or large-scale fuel cell arrays, coupled with advanced switchgear and uninterruptible power supply (UPS) systems. Because of this complexity, we anticipate AI Factory costs significantly increasing in the coming years. The convergence of multiple bottlenecks across energy, labor, and memory defy the economies of scale common to technology.

For this reason, estimates of per-gigawatt costs vary widely, from as low as $10 billion to as high as $80 billion. We estimate that average per-gigawatt costs will grow from $40 billion in 2025 to $48 billion by 2030, according to Futurum Equities research. At a high level, terrestrial costs are likely to continue increasing while space costs decline. The intersection of these two cost curves is the central question for the orbital compute thesis – and we believe that crossover will arrive first, for off-grid deployments, not for all AI compute.

The majority of data centers under construction today are grid-connected facilities benefiting from existing electrical infrastructure, municipal permitting frameworks, and established supply chains. Grid-connected construction costs remain low enough that most data centers will remain on the ground in the long run. The orbital cost parity argument is most compelling for the subset of AI compute that would otherwise require purpose-built off-grid power infrastructure, with significant social costs, including environmental hazards and higher electricity prices.

The Orbital Premium: Economics of a 1 GW Space Data Center

In stark contrast, building a 1 GW data center in orbit replaces the constraints of Earth's grid with the unforgiving realities of astrophysics. Companies pursuing orbital data centers, such as Starcloud, aim to deploy thousands of interconnected satellite nodes in sun-synchronous orbits. These costs are likely to decline over time as launch cadence increases, satellite manufacturing scales, and the broader space supply chain matures, but the trajectory and timing of those declines carry substantial uncertainty.

Key cost line items include:

- **Launch costs:** Rapidly collapsing due to reusable rockets. SpaceX drove $/kg from over $30,000 in the Space Shuttle era down to approximately $2,000 with Falcon Heavy, and roughly $1,400/kg at commercial pricing. Starship is expected to push internal costs below $100/kg by 2030 through full rocket reuse, rapid restaging, and dramatically increased flight volumes.
- **Satellite hardware:**
  - **Radiators:** The single largest mass driver for orbital data centers. Because there is no air in space, all waste heat must be rejected via passive infrared radiation. At gigawatt scale, heat rejection requires deploying hundreds of thousands of square meters of radiator surface area – football-field-sized panels per satellite cluster. Starcloud claims 70% of its engineering effort is dedicated to solving the heat problem. They have developed custom heat sinks and large-scale deployable radiators for their Starcloud 3 architecture.
  - **Payload systems:** Current Starlink V2 Mini-class satellites cost approximately $22/W, down from $32/W for first-generation Starlink. Starcloud CEO Philip Johnston has stated a target of less than $5/W for Starcloud-3, excluding GPU costs.
  - **Solar arrays:** Space provides unattenuated solar irradiance with 99% capacity factors for solar panels, compared to 15–25% on Earth. At 70 W/kg specific power – Starcloud's near-term target – a 200 kW satellite requires roughly 2,860 kg of solar array mass, a significant but manageable fraction of the total spacecraft.
- **Maintenance:** Orbital data centers face unique servicing challenges. At low-Earth orbit altitudes, atmospheric drag requires periodic station-keeping propulsion. Radiation-induced component degradation necessitates either redundant systems or periodic replacement.
- **Monitoring:** Ground station networks for telemetry, tracking, and command (TT&C) are required to maintain continuous operational awareness.

We estimate that under current launch economics, space data center infrastructure costs for a system such as Starcloud 3 will fall around $72.1 billion per gigawatt, excluding the cost of compute itself. This compares to $16.2 billion for off-grid terrestrial data centers, given recent increases in gas turbine prices. Starcloud CEO Philip Johnston predicts that space infrastructure will drop to $5 billion per gigawatt, contingent on Starship achieving commercial launch costs of approximately $500/kg – a threshold Johnston expects to be reached around 2028–2029. We expect that even steeper launch cost declines and increased solar efficiency will be needed to hit that dollar-per-watt threshold.

Figure 1: Cost per GW Comparison: Orbital vs. Terrestrial Data Center Costs Under Current Economics (excluding Compute)

Source: Futurum Research, April 2026

While there is precedent for bringing down launch costs, additional cost savings will need to be realized in satellite hardware and manufacturing scale. Starcloud's Starcloud-3 spacecraft – a 200 kW, three-ton platform designed for Starship's deployment system – represents the first architecture purpose-built for cost parity with terrestrial facilities. If off-grid terrestrial construction costs surpass $20B/GW excluding the cost of compute, orbital infrastructure will first achieve cost competitiveness primarily against off-grid terrestrial deployments. This threshold can be reached in orbital data centers by doubling solar panel watts per kilogram, 50% decline in satellite hardware cost per watt, and $100/kg launch costs – all technically feasible milestones over a five-year horizon. It may take 15 years or more to compete with grid-connected data centers, where land, permitting, and power interconnection are already in place, and costs are closer to $10B/GW.

If $3 trillion is dedicated to AI compute capex in 2030, aligning with a bull case advanced by AI industry executives, we believe roughly one-third of that spend – approximately $1 trillion – could economically justify extraterrestrial deployment, representing capacity that would otherwise require new off-grid power generation, face multi-year grid interconnection delays, or encounter political challenges. This remains a massive addressable market, and one that will be supply-constrained by launch cadence and satellite manufacturing throughput well beyond 2030.

Emerging Orbital Computing Ecosystem

Economic transitions bring broad-based benefits across innovation ecosystems, even with dominant platforms taking early leads. While SpaceX is considered a major winner in space, the growth of Tesla's supplier ecosystem anticipates that complete vertical integration of an emerging supply chain will be too risky to attain. While space economy leaders have strong market positions, the unique optimization needs of AI workloads and cloud customers is likely to require additional entrants. The following ecosystem map identifies companies positioned across the orbital compute value chain.

Compute Clusters

- **AMD:** Versal AI Edge Gen 2 adaptive SoCs provide space-grade processing designed specifically for orbital deployments and extreme environments.
- **NVIDIA:** Space-1 Vera Rubin Module delivers up to 25x more AI compute for space-based inferencing than the H100, with a tightly integrated CPU–GPU architecture and high-bandwidth interconnect designed for processing massive data streams from space-based instruments in real time. Announced at GTC 2026, the module targets orbital data centers, geospatial intelligence, and autonomous space operations. Availability date has not yet been disclosed.

Semiconductor Hardening

- **Aitech:** Ruggedized electronic systems and command/data handling architecture for space-rated servers, providing radiation-hardened computing platforms for mission-critical orbital applications.
- **OrbitsEdge:** Adapts HPE Edgeline servers for high-radiation environments
- **Intel:** Recently fabricated the thinnest gallium nitride (GaN) chiplet at 19 micrometers as part of broader Terafab-era advanced packaging research. GaN's wide-bandgap properties – high electron mobility, radiation tolerance, and thermal resilience – make it a candidate material for future space-grade power electronics and RF systems.
- **Unibap:** Develops radiation-tolerant payload processing units optimized for autonomous on-orbit data processing and edge AI inference in LEO constellations.

Optical Interconnect

- **CACI International:** Produces optical terminals for defense and commercial crosslinks, with a heritage in secure government communications applied to inter-satellite laser links.
- **Kepler Communications:** Optical relay network serving as high-speed backhaul between satellites, providing the inter-node connectivity fabric that distributed orbital data centers require.
- **Mynaric:** Producer of laser communications terminals for air and space applications, scaling production of free-space optical systems for LEO mega-constellations.
- **Skyloom:** Developer of high-bandwidth laser networking infrastructure for real-time data relay between satellites and ground stations.
- **Starcloud:** Pure-play orbital computing connectivity leader, currently operating the most powerful commercial GPU (H100) ever deployed in orbit, with plans to scale to Blackwell-class and Vera Rubin-class compute on Starcloud-3 platforms.

Multi-physics Simulation

- **Cadence:** Provides 3D electromagnetic field solvers (Clarity) and computational fluid dynamics tools essential for modeling thermal management, signal integrity, and power distribution in space-grade electronic systems. Cadence's Celsius thermal solver enables full-system thermal analysis critical for validating radiator designs and chip junction temperatures in vacuum environments.
- **Synopsys (Ansys):** Following the completed Synopsys-Ansys merger, the combined entity offers the industry's deepest multiphysics simulation stack. Ansys HFSS electromagnetic simulation, structural mechanics, and thermal analysis software can model the performance of sensitive electronic systems in orbital environments.

Vertical Integrators

- **Amazon:** Integrator of Project Kuiper laser network with AWS terrestrial infrastructure, positioning orbital compute as a seamless extension of the AWS cloud.
- **Northrop Grumman:** Integrator across ground network, satellite manufacturing, and payload deployment, with decades of heritage in defense and intelligence satellite systems applicable to orbital data center infrastructure.
- **SpaceX:** The dominant heavy-lift and medium-lift launch provider, and increasingly a vertically integrated orbital compute player. Following the acquisition of xAI, SpaceX is both the infrastructure backbone and a potential direct competitor to every company in this ecosystem.
- **Rocket Lab:** Positioned as the leading small- and medium-lift alternative to SpaceX, with the Electron launch vehicle operational and the Neutron medium-lift vehicle in development. Rocket Lab's Space Systems division also manufactures satellite components and complete spacecraft, offering a vertically integrated alternative for customers seeking launch diversity.

What to Watch

- **Starship Commercial Cadence:** The entire orbital compute thesis depends on Starship achieving reliable, high-frequency commercial launches. Johnston expects commercial Starship access to open in 2028–2029. Delays push back the cost crossover point for off-grid deployments and constrain total addressable capacity.
- **Radiator and Solar Panel Mass-to-Power Ratios:** Improvements in specific power for solar panels (watts generated per kilogram) and radiators (watts rejected per kilogram) directly reduce launch mass and cost.
- **Grid-Connected vs. Off-Grid Cost Divergence:** Our thesis holds that orbital compute first achieves parity with off-grid terrestrial deployments. If grid interconnection timelines shorten dramatically – through policy reform, modular nuclear deployment, or accelerated transmission buildout – the addressable market for orbital compute narrows. Conversely, if grid bottlenecks worsen, the $1 trillion addressable estimate may prove conservative.
- **SpaceX Vertical Integration Risk:** SpaceX's xAI acquisition and million-satellite filing signal intent to compete directly in orbital compute, not merely provide launch services. This creates platform risk for independent players such as Starcloud, analogous to AWS competing with customers. Watch for SpaceX pricing signals to third-party orbital compute providers and for any changes in launch-priority allocation.
- **Semiconductor Radiation Hardening at Scale:** NVIDIA's Space-1 Vera Rubin Module and AMD's Versal AI Edge Gen 2 XQR represent the first generation of purpose-built space AI silicon. Their cost, availability, and performance in actual orbital conditions will determine whether orbital data centers can run the same workloads as terrestrial facilities, or remain limited to specialized inference and edge processing tasks.

Other Insights from Futurum

Can Intel Foundry's Advanced Packaging Bring the Terafab Vision to the Stars?

Will Starcloud's Orbital Data Centers Solve NVIDIA's Terrestrial Energy Crisis?

AI Grid Constraints Will Push Over 33% of Data Centers Off-Grid by 2030 Subscribers Only

About Us

About the Authors

Futurum Research Director Brendan Burke covers semiconductors, supply chain, and emerging technology for Futurum. He is an expert in AI computing with experience advising Fortune 100 companies on generative AI strategy. He is an accomplished analyst credited with the most comprehensive market landscapes in the high-performance semiconductor market. At Futurum, he helps clients understand sophisticated technologies and how to make the best use of them for success in their organizations. He can be reached at bburke@futurumgroup.com.

Futurum

About The Futurum Group

Every day, The Futurum Group's analysts, researchers, and advisors help business leaders worldwide anticipate tectonic shifts in their industries and leverage disruptive innovation. Unlike traditional analysts, The Futurum Group works not only in analysis and research but also takes that insight and knowledge even further, engaging all the way through the go-to-market process.

Futurum Research provides in-depth research and insights on global technology markets using advisory services, custom research reports, strategic consulting engagements, digital events, go-to-market planning, and message testing. It also creates, distributes, and amplifies rich media content that all stakeholders read, watch, and listen to.

See more details on The Futurum Group at futurumgroup.com.

Copyright & Use License

Copyright Notice

License Notice

Limitation of Liability Notice