TECHNOLOGY, MEDIA, AND TELECOMMUNICATIONS

Space-Based Data Centers: More Than Hype, but Not a Revolution

By Val Elbert, Thibault Werlé, Marc Nasr, and Clark O'Niell

ARTICLE AUGUST 27, 2026 12 MIN READ

Data centers are facing mounting constraints on land and power. Yet the industry's future may be looking up—literally. Space-based data centers promise virtually unlimited solar energy, room to grow, and no need to lay down stakes on land. Compelling as the vision is, however, it also divides opinion. Orbital data centers are often seen as either the inevitable next frontier of computing or little more than a science-fiction fantasy.

The reality is more nuanced. Our analysis indicates that space-based data centers could become technically feasible at scale within the next five to ten years. But it also suggests that they are unlikely to compete with terrestrial infrastructure on cost. Today, orbital data centers would carry a 2.5x to 3x cost premium. Even under realistic improvement trajectories, that premium narrows to approximately 1.5x over the next decade, but it does not disappear.

Yet cost is only part of the equation. For specific AI workloads, satellites offer advantages that justify the price of admission. Case in point: processing sovereign data that must never cross national borders. The key is to think of orbital data centers not as a replacement for terrestrial infrastructure but as a complementary layer.

It won't be a trivial layer, either. We estimate that orbit-advantaged use cases are substantial enough for satellites to capture 10% to 15% of the global AI data center market by 2040.

For business leaders, the strategic implication is clear: don't overcorrect your terrestrial roadmap, but don't ignore orbit, either. For investors and policymakers, this complementary layer opens new opportunities and raises new strategic considerations. But one thing is certain: AI infrastructure will no longer be confined to Earth.

Technical Barriers Are Falling

The idea of space-based data centers is not new, but proof of concept is. The first commercial data center satellite launched in late 2025, and multiple players have filed plans to deploy more than one million satellites in total. Meanwhile, China has announced plans to build a "space cloud" by 2030. Clearly, space-based data centers can be built. But can they be deployed at the scale, cost, and reliability required for widespread adoption?

To help answer that question, we assessed six key technical constraints spanning the space-based data center value chain. Using NASA's technology readiness level (TRL) framework, we evaluated each of these categories, looking at both current readiness and the likelihood of reaching full operational readiness (TRL 9) within the next five to ten years.

The results reveal that some constraints will ease faster than others, with launch efficiency approaching operational readiness while cooling and in-orbit maintenance remain the primary bottlenecks. (See Exhibit 1.)

EXHIBIT 1

The Six Technical Constraints Can Be Overcome, but the Solutions Are Maturing at Different Rates

Constraint	Leading solutions	Current TRL	Likelihood of reaching TRL 9 in 5 to 10 years	Rationale
Launch efficiency	Increase in reusable spacecraft volume and capacity Expansion of launchpad infrastructure	TRL 8 TRL 8	High	SpaceX's Starship has completed 13 test flights and is expected to begin commercial operations in 2027
Cooling	Higher-efficiency lightweight radiators Two-phase fluid loops	TRL 6 TRL 5	Low	Current radiators are in the range of 20 kW; increasing to 100 kW+ will require significant improvement
Battery life	Space batteries with higher energy density Enhanced spacecraft propulsion and control systems	TRL 7 TRL 8	High	Satellite deployments in dawn-to-dusk sun-synchronous orbit have been proven to be feasible
Radiation tolerance	Improved radiation-tolerant design Targeted shielding	TRL 6 TRL 5	Low	While promising, solutions remain in early stages and require further development
Connectivity	Adoption of optical communication links Higher-frequency communications systems	TRL 6 TRL 7	Medium	Optical intersatellite links have been deployed, but large-scale ground-to-space connectivity has yet to be demonstrated
In-orbit maintenance	Standardized robotic in-orbit servicing Adoption of longer-lifetime components	TRL 4 TRL 5	Low	Technologies are still in the early stages of development and testing and require modular satellite architectures

Source: BCG analysis.
Note: TRL = technology readiness level. Number of Starship test flights is current as of mid-August 2026.

- **Launch Efficiency.** Launch is one of the largest contributors to the price of space-based data centers. At current launch costs of approximately $1,500 per kilogram, putting 1 GW of compute into orbit would require roughly $30 billion, making large-scale deployment economically impractical. But we expect launch costs to fall dramatically, to perhaps $100/kg, driven by fully reusable, high-capacity launch systems such as SpaceX's Starship. Expanded launch infrastructure should further increase flight cadence and improve economies of scale.
  Current readiness (TRL): 8
  Expected time to TRL 9: 0 to 2 years

- **Cooling.** In the vacuum of space, heat must be radiated away, requiring much larger cooling systems than those of terrestrial data centers. A 100 kW satellite requires a radiator of roughly 400 m^2—far larger than the ones used on today's commercial satellites—driving significant increases in mass and system complexity. But next-generation lightweight radiators could reduce radiator area by as much as 30%, while high-capacity two-phase cooling loops could reduce it by an additional 40%. SpaceX's recently unveiled 110 m^2 radiators suggest that cooling technology may be maturing faster than previously expected, although it is likely to remain a principal bottleneck to large-scale deployment.
  Current readiness (TRL): 6
  Expected time to TRL 9: 5 To 10 years

- **Battery Life.** Satellites in low-Earth orbit spend approximately one-third of their time in Earth's shadow, relying on batteries during eclipse periods. Under these conditions, supporting AI workloads would require battery capacity roughly 5x to 10x greater than today's

space-grade cells provide. Two developments are particularly promising. Next-generation lithium-ion batteries could increase storage capacity by approximately 5x. And advances in propulsion and control systems are making dawn-to-dusk sun-synchronous orbits increasingly practical, enabling near-continuous solar exposure.
Current readiness (TRL): 7
Expected time to TRL 9: 2 to 5 years

- **Radiation Tolerance.** Radiation levels in low-Earth orbit are approximately 20x to 30x higher than on Earth's surface, meaning unprotected electronics can fail within 30 to 50 days. Space-qualified chips can improve resilience by up to 2x, while targeted shielding using multilayer graded-Z materials can extend operational life by up to 1.5x. The focus now is on optimizing these solutions for the low-mass designs required for commercially viable space-based data centers.
Current readiness (TRL): 7
Expected time to TRL 9: 2 to 5 years

- **Connectivity.** Space-based data centers depend on high-capacity communications links. Yet today's RF systems deliver only 1 to 10 Gbps, far below the capacity of terrestrial fiber networks, and available spectrum becomes increasingly congested as satellite constellations expand. Optical links are one solution, having demonstrated speeds exceeding 200 Gbps per link. And higher-frequency Q/V-band systems could increase available bandwidth by 3x to 5x. The challenge: honing these technologies for real-world operating conditions. Optical links remain sensitive to atmospheric conditions, while higher-frequency RF systems require additional power to penetrate the atmosphere.
Current readiness (TRL): 7
Expected time to TRL 9: 2 to 5 years

- **In-Orbit Maintenance.** Today's satellites typically operate for five to seven years before being deorbited. For space-based data centers, this model—replace rather than repair—increases both operating costs and orbital debris. Robotic servicing satellites could extend operational life by 2x to 3x, while longer-lifetime components could reduce replacement frequency by a similar amount. The issue is that demand for these technologies does not yet exist at scale, limiting investment. Large-scale deployment of space-based data centers will create the demand, but given this dependency, we expect advances to evolve more slowly in in-orbit maintenance than in the other categories.
Current readiness (TRL): 4 to 5
Expected time to TRL 9: 10+ years

Beyond Technology Readiness

Overcoming the six technical constraints will establish technical feasibility, not commercial viability.

Deployment will require tens of thousands of satellites, demanding mass production of satellite buses, solar arrays, radiators, and other space-qualified subsystems. No operator today has fully built this capability. Scaling production from prototype to fleet is itself a technological challenge, requiring each subsystem to be redesigned for high-volume manufacturing without sacrificing performance.

But the greatest benefits will go to players that combine mass manufacturing with vertical integration. Today's value chain spans satellite manufacturing, launch services, ground infrastructure, and AI compute operations, with each stage typically capturing its own margin. Operators that integrate these capabilities can retain those margins and reinvest them into cost reductions and capability development. Achieving that transformation while coordinating advances across launch, manufacturing, and operations represents the sector's central execution challenge. Those that succeed will hold a lasting structural advantage.

The Persistent Cost Premium

Commercial competitiveness also depends on the economics. Even under favorable assumptions, space-based data centers will be more expensive than their terrestrial counterparts.

To compare the economics of space-based and terrestrial AI data centers, we modeled their 20-year total cost of ownership (TCO). Notably, the two architectures have fundamentally different profiles for capital and operating expenditures. Whereas terrestrial data centers balance upfront capital investment with decades of opex—including electricity, cooling, water, and maintenance—space-based data centers are dominated by capex. GPUs account for roughly half of TCO, with launch costs contributing about one-fifth. Opex remains comparatively low because electricity and on-site operations are largely eliminated.

Our analysis assumes that the technical and manufacturing challenges discussed above have been overcome, allowing for a direct comparison of the underlying economics. Even under those assumptions, space-based data centers carry a 2.5x to 3x cost premium over terrestrial infrastructure today, with a 20-year TCO of approximately $660 million to $750 million per MW, compared with $230 million to $300 million per MW for terrestrial data centers.

Future improvements will narrow the gap, but are unlikely to eliminate it. To understand what would be required for space-based data centers to approach commercial viability, we conducted a sensitivity analysis across three key variables: launch cost per kilogram, satellite weight per 100 kW of compute capacity, and terrestrial electricity prices. But we also incorporated a fourth factor that deserves attention: satellite failure rates. Unlike terrestrial data centers, where failed hardware can typically be replaced, failures in orbit may require retiring an entire satellite.

As a result, failure rates materially affect the economics of space-based data centers. Under a realistic five- to ten-year improvement trajectory—with a 67% reduction in launch costs and a 50% reduction in satellite weight—the cost premium is approximately 1.8x if failure rates remain at 30%. It narrows to 1.6x if failure rates decline to 20%, and to 1.5x if they fall to 10%. (See Exhibit 2.) Under a more aggressive scenario consistent with full Starship economics and larger-capacity satellites, the cost premium falls to 1.2x at a 10% failure rate and narrows to 1.1x if failure rates decline to 5%, approaching parity with terrestrial data centers.

EXHIBIT 2

The Cost Premium Will Shrink but Not Disappear, Even Under the Most Aggressive Improvement Scenario

TCO sensitivity (measured by space-based/terrestrial multiplier expected in 5 to 10 years)

Launch cost ($/kg)	Satellite weight (tons/100 kW): 2 tons/100 kW (baseline)	1 ton/100 kW (50% improvement)	0.5 ton/100 kW (75% improvement)
$1,500/kg (baseline)	2.3X	1.6X	1.3X
$500/kg (67% improvement)	1.9X	(1) 1.5X	1.2X
$100/kg (93% improvement)	1.8X	1.4X	(2) 1.2X

(1) Optimistic scenario (2) Aggressive scenario

Source: BCG total cost of ownership (TCO) sensitivity analysis.
Note: The optimistic scenario assumes a realistic five- to ten-year improvement trajectory, with a 67% reduction in launch costs and a 50% reduction in satellite weight; the aggressive scenario assumes even further reductions, as shown here. Exhibit assumes that the space-based failure rate falls to 10%, terrestrial power tariff remains at current baseline of $0.12/kWh, and space-based compute costs hold steady at $220 to $260 million/MW (1.2x terrestrial).

Ultimately, narrowing the cost premium will depend not only on continued technological progress, but also on business model choices—particularly the ability of vertically integrated operators to capture efficiencies across the entire value chain.

Where Orbit Creates Value

If space-based data centers succeed, it won't be because they are cheaper. Instead, adoption will depend on fit. Some AI workloads benefit enough from operating in orbit to justify the premium, while others remain fundamentally better suited to terrestrial infrastructure. In our most likely scenario, orbit-advantaged workloads could enable space-based data centers to capture 10% to 15% of the global AI data center market by 2040, representing $240 billion to $320 billion in annual revenue. (See Exhibit 3.) Three types of use cases stand out:

- *Latency-tolerant inferencing* encompasses AI applications in which near-real-time response is not required, allowing compute to shift to orbit without affecting user experience. Use cases include batch document, image, and video generation, enterprise back-office AI (such as contract review), scientific inference, and bulk translation and tagging.
- *Sovereign AI workloads* are applications in which nationally sensitive or regulated data must be processed without transiting foreign networks or jurisdictions. Use cases include government AI services using citizen data (such as health and tax records), nationally regulated datasets, and resilient backup compute for critical systems.
- *Space-generated data processing* co-locates compute with the data source in orbit, with the results then transmitted to Earth. Use cases include Earth observation, climate and environmental monitoring, disaster response, and space debris and satellite tracking.

EXHIBIT 3

Workload Suitability Will Determine Adoption and Market Size of Space-Based Data Centers

		Conservative scenario	Optimistic scenario	Aggressive scenario
		SBDCs remain limited to niche use cases (such as sovereign AI workloads and space-generated data)	SBDCs become a complementary layer, capturing nonsensitive inferencing and niche use cases	SBDCs become a major market segment for nonsensitive inferencing
2030	Overall SBDC capacity	Negligible; focused on niche use cases	1 GW–2 GW	2 GW–3 GW
	Share of total AI data center market	Negligible; focused on niche use cases	1%–2%	2%–3%
	Total revenue	Negligible; focused on niche use cases	$8 billion–$16 billion	$16 billion–$24 billion
2040	Overall SBDC capacity	10 GW–20 GW	30 GW–40 GW	50 GW–100 GW
	Share of total AI data center market	3%–5%	10%–15%	20%–30%
	Total revenue	$80 billion–$160 billion	$240 billion–$320 billion	$400 billion–$800 billion

Source: BCG analysis.

While sovereign AI workloads and space-generated data represent relatively specialized opportunities, latency-tolerant inference is a much larger addressable market, accounting for approximately 40% to 45% of the global AI data center market by 2030.

Applications requiring real-time responsiveness—including interactive AI assistants and autonomous systems—are better suited to terrestrial infrastructure because low-Earth orbit introduces unavoidable communication delays. Similarly, training large foundation models is likely to remain Earth-based, as it requires tightly coupled compute clusters and power densities that orbital systems are unlikely to match.

Space-based and terrestrial data centers aren't going to do battle. They're going to coexist, with workloads matched to the environment that fits them best.

What Comes Next

The question is no longer whether space-based data centers will emerge, but how they will fit into the future of AI infrastructure. The decisions organizations make today will shape their position in orbital computing for the decade that follows. Three groups—business leaders, financial investors, and policymakers—should prepare now.

Business Leaders. Organizations will likely need to continue investing in terrestrial AI infrastructure. The displacement narrative is overstated, and space-based data centers are unlikely to achieve cost parity this decade. Near-term buildouts remain the right strategy.

At the same time, leaders should begin incorporating orbital computing into long-term planning, focusing on the workloads where it offers a clear advantage. Organizations with sovereign data obligations or significant space sector exposure should explore partnerships with emerging orbital providers while the market remains prescale. Early engagement preserves optionality at relatively low cost.

Financial Investors. This group has two distinct opportunities. One focuses on the enabling technologies that make space-based data centers possible. The other focuses on integrated operators seeking to build and operate orbital data centers themselves.

The enabling ecosystem offers the lower-risk opportunity. Technologies such as launch services, optical communications, thermal management, and radiation-hardened semiconductors have applications well beyond orbital computing. They offer paths to value creation even if space-based data center adoption proves slower than expected.

Integrated operators offer the higher-risk, higher-reward opportunity. Players that successfully combine mass manufacturing with vertical integration stand to capture a greater share of the value created. But realizing that opportunity will require coordinating manufacturing, launch, and operations at unprecedented scale.

Policymakers. Governments can establish regulatory and licensing frameworks before orbital data centers reach scale. They can also address issues such as data sovereignty, international liability, and orbital debris. Early action will reduce uncertainty for industry participants and enable more orderly market development.

Simultaneously, policymakers can identify the sovereign data, defense, and space-generated data use cases where orbital data centers provide a strategic advantage. For these applications, governments should be prepared to accept a cost premium in exchange for structural data security.

Even if everything goes right—cheaper launches, lighter satellites, more reliable hardware—space-based data centers are unlikely to outperform terrestrial infrastructure on cost anytime soon. But for the right workloads, cost is the wrong question. The winners won't be those that choose between Earth and orbit. They'll be the ones that know which workloads belong in each.

The authors thank Rüdiger Schicht, Troy Thomas, Raed Saab, and David Nagy for their contributions to this article.

Authors

Val Elbert

Managing Director & Senior Partner; Global Leader, Technology, Media, & Telecommunications Practice
New Jersey

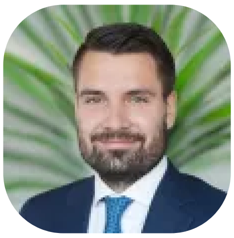

Thibault Werlé

Managing Director & Partner
Toronto

MN

Marc Nasr

Managing Director & Partner
Dubai

Clark O’Niell

Managing Director & Partner
San Francisco - Bay Area

ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group bridges the gap between ambition and outcomes for the world's leading companies and organizations. We are built for this era of unprecedented change — bringing strategic clarity rooted in over 60 years of deep domain knowledge, combined with applied AI shaped by our practitioners. BCG works shoulder-to-shoulder with CEOs across industries and geographies to deliver transformative impact at scale: stronger returns, transferred capabilities, and change that sticks. For more information, visit bcg.com.

