WHITE PAPER

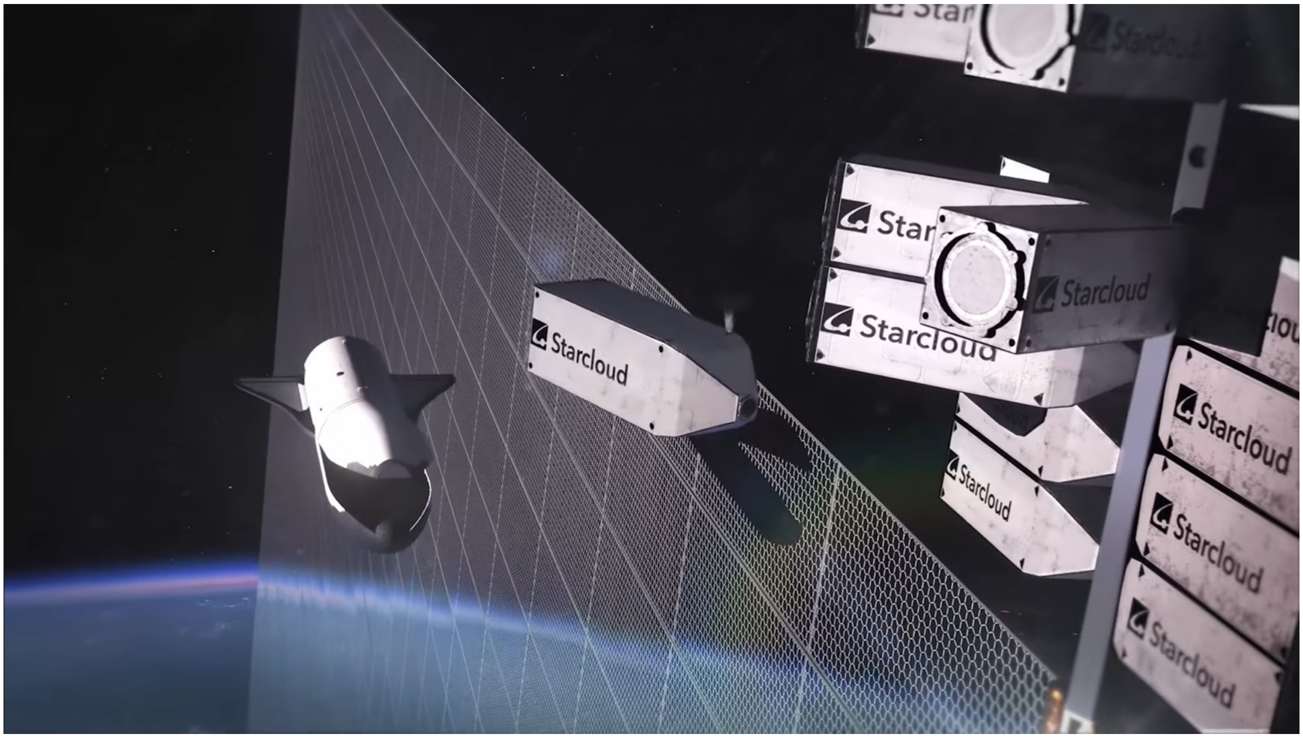

Pigs in Space: The Muppet Sketch That Turned Out to Be a Business Plan

The strategic case for data centres in space, and what boards and technology leaders need to understand about the orbital computing revolution

Marco Ciobo
Principal, Digital Insights and Transformation | Board and C-Suite Advisor
June 2026

Executive Summary

Miss Piggy did it in 1977.[1] Data centres are only now catching up.

This white paper examines the structural, economic, and environmental case for orbital data centres. It draws on analysis from the World Economic Forum, McKinsey & Company, NPR, Starcloud, and documentary research covering the full landscape of space-based computing. The conclusion is clear: moving compute infrastructure to low Earth orbit is not science fiction. It is engineering in progress — and it may be the most consequential infrastructure shift since the advent of cloud computing.

The central argument is straightforward. AI and the race toward artificial general intelligence (AGI) are placing energy demands on terrestrial infrastructure that Earth cannot satisfy sustainably. Data centres already consume approximately 1.5 percent of global electricity. That figure is projected to double by 2030, requiring 50 to 100 gigawatts of new energy in North America alone — the equivalent of 50 to 100 new nuclear power plants. Space offers something Earth cannot: continuous, abundant solar energy at greater than 95 percent capacity factor, passive cooling to deep space temperatures, and unlimited modular scalability unconstrained by land, permitting, or grid capacity.

> ***The strategic imperative is not whether to take orbital computing seriously. It is how quickly decision-makers, investors, and technologists understand that space is becoming the infrastructure layer for artificial intelligence.***

Key findings: Solar capacity factors in orbit exceed 95 percent versus a median of 24 percent for US terrestrial solar farms. Effective energy cost in orbit is estimated at approximately $0.005 per kilowatt-hour — up to 15 times lower than today's wholesale electricity prices. Infrastructure costs could fall below $5 million per megawatt, versus $12–15 million for US terrestrial facilities. A 5 GW cluster — required for next-generation AI models — is achievable in orbit but not on Earth's current grid infrastructure. The orbital data centre market, valued at approximately $500 million in 2025, is projected to reach $15–20 billion by 2030 and potentially $75 billion by 2035.

[1] Pigs in Space was a recurring sketch on The Muppet Show (1977–1981), featuring Miss Piggy, Link Hogthrob, and Dr. Julius Strangepork aboard the spaceship Swinetrek. It was comedy. It was also, it turns out, infrastructure foresight.

1. The Terrestrial Energy Crisis

The AI revolution is, at its core, an energy problem. Training and running large AI models demands compute clusters of unprecedented scale. Models such as GPT-6 and Llama 5 require multi-gigawatt clusters that stress electricity grids already operating at capacity. The International Energy Agency projects that global data centre power consumption will roughly double to nearly 1,000 terawatt-hours by 2030.

The constraints are not merely financial — they are physical. Utilities in the Western world cannot adapt at the required pace. Permitting timelines for new gigawatt-scale energy projects span a decade or more. Starcloud CEO Philip Johnston has stated plainly that without new infrastructure solutions, companies will soon be 'leaving chips in warehouses because they don't have power for turning them on.'

The sustainability challenges compound the energy problem. A single 40 megawatt terrestrial cluster can consume over one million tonnes of water annually for cooling. Hyperscale facilities require substantial land footprints, and community opposition, environmental review, and grid connection requirements add years to development timelines. The existing Earth-based model does not scale or sustain the future of AI.

> ***The organisations and nations that solve the energy problem for AI will define the next competitive era. The solution is not to consume more of Earth's resources. It is to leave them behind.***

The fundamental challenge is a scale mismatch. AI demand grows exponentially; terrestrial energy and land infrastructure grows linearly at best. A 5 GW compute cluster — necessary for the AI models of 2028 and beyond — is simply not achievable on Earth within any realistic planning horizon. It is, however, achievable in orbit.

Orbital vs Terrestrial: The Core Metrics at a Glance

Metric	Orbital (LEO)	Terrestrial (US)
Solar capacity factor	>95%	~24% (median)
Effective energy cost	~$0.005 / kWh	$0.07–0.12 / kWh
Water consumption (40 MW)	Zero	>1M tonnes / year
Land requirement	None	Significant + permitting
Infrastructure cost	<$5M / MW (projected)	$12–15M / MW
Cooling method	Passive radiative	Active chillers + towers
Scale-up timeline	Months (modular launch)	Decade+ (grid, permits)
Operational carbon	Near zero	Grid-dependent

2. Why Space? The Structural Advantages

Space is not merely an escape from terrestrial constraints. It offers a genuinely superior environment for high-efficiency, sustainable computing. The advantages are structural, not incidental, and they compound in ways that make orbital infrastructure increasingly attractive as AI compute demands accelerate.

Abundant, uninterrupted solar power

In a dawn-dusk sun-synchronous orbit, a satellite follows Earth's terminator line and remains in near-continuous sunlight — no night, no cloud cover, no weather. Orbital solar arrays achieve a capacity factor of greater than 95 percent, compared to a median of just 24 percent for terrestrial US solar farms. Peak power generation in space is roughly 40 percent higher due to the absence of atmospheric losses. The net result: the same solar array in space generates more than five times the energy as on Earth.

Starcloud estimates this translates to an effective energy cost of approximately $0.005 per kilowatt-hour — up to 15 times lower than today's wholesale electricity prices. Elon Musk summarised the logic directly: 'You're power constrained on Earth. Space has the advantage that it's always sunny.'

Free radiative cooling

Cooling is one of the most energy-intensive and water-hungry aspects of terrestrial data centre operations. In space, this problem is resolved by physics. The deep vacuum of space serves as a gigantic cold heatsink at an effective temperature of approximately –270 degrees Celsius. Orbital data centres use simple, passive radiative cooling — circulating coolant through chips, transferring heat to radiator panels, and emitting it as infrared radiation into deep space.

This eliminates energy-intensive chillers, cooling towers, and, critically, fresh water consumption entirely. The environmental benefit is material: no water drawn from stressed local aquifers, no heat island effect on surrounding communities.

Unlimited scalability

There is no land to acquire, no community to consult, no grid to upgrade. Orbital data centres can be scaled linearly and almost indefinitely through modular design. Compute containers dock to a central spine providing network, power, and cooling connectivity. Starcloud's modelling suggests a 5 GW cluster could be deployed in two to three months using reusable heavy-lift vehicles. The equivalent terrestrial infrastructure would take a decade or more.

Economic competitiveness

Terrestrial data centres carry infrastructure costs of $12–15 million per megawatt in the US. Orbital systems, which require no batteries, backup power, chillers, or cooling towers, could reduce that figure to below $5 million per megawatt at scale. Launch costs are the dominant current disadvantage — approximately $1,000 per kilogram to low Earth orbit today. SpaceX's Starship targets $10–20 per kilogram at full operational cadence. Google's analysis puts the break-even threshold at around $200 per kilogram. Break-even analysis from multiple sources suggests orbital data centres could reach cost parity with terrestrial alternatives within five to seven years of deployment at scale.

> ***The economics of orbital computing are not yet superior to terrestrial alternatives. They are improving faster than most analysts predicted two years ago, and the trajectory is clear.***

3. How Orbital Data Centres Work

An orbital data centre is a modular, maintainable system built around the constraints and opportunities of low Earth orbit. Understanding the architecture is important for decision-makers evaluating the credibility and timeline of this technology.

Core architecture

The fundamental building block is the compute container: a standardised module housing server racks, networking equipment, liquid-cooling infrastructure, and power distribution systems.

Containers dock to a central spine providing network, power, and cooling to the entire cluster. The modular design allows incremental expansion — new containers are launched and docked as capacity requirements grow, without disrupting existing operations.

Power comes from expansive thin-film solar arrays sized for the computational load. Waste heat is dissipated via large passive radiators pointed toward deep space. The combination of near-continuous solar generation and passive cooling creates an extraordinarily stable power environment — arguably more stable than on Earth, where weather and grid variability introduce constant operational uncertainty.

Data transport and connectivity

Connectivity between orbital clusters and Earth relies on two complementary approaches. For high-throughput AI training workloads, data uplinks and downlinks use laser-based optical communications with mega-constellations such as Starlink and Kuiper. Laser communication systems offer more than 1,000 times the bandwidth of radio-frequency alternatives. For transporting petabytes or exabytes of initial training data, physical 'data shuttles' — small docking modules carrying storage hardware — provide a practical near-term solution. Inter-satellite laser links also allow distributed compute clusters to communicate with near-zero latency within a closely flying constellation.

Workloads and suitability

Almost all AI inference workloads are viable for orbital data centres. The exception is applications requiring latency below approximately 50 milliseconds — a threshold relevant for high-frequency trading and certain real-time gaming scenarios but not for the vast majority of AI applications. Chat, search, voice agents, video generation, back-office automation, and satellite data processing are all well-suited to orbital compute. AI model training — large batch jobs rather than real-time interaction — is particularly well-matched, as latency is essentially irrelevant to training performance.

Radiation management and reliability

Operating in low Earth orbit requires radiation shielding for sensitive electronics. Traditional space hardware uses radiation-hardened chips that are extremely expensive and at least 100 times less capable than state-of-the-art commercial options. The leading approach, pioneered by Starcloud, is to use commercial chips and adapt them for space through extensive testing under simulated five-year radiation doses using particle accelerators. Real-world telemetry from satellites in orbit now supplements laboratory data. Autonomous software-based fault tolerance — routing around failures rather than repairing hardware — mirrors the approach already standard in large terrestrial data centre operations.

4. The Market Landscape and Key Players

The orbital data centre market is rapidly transitioning from concept to early deployment. Valued at approximately $500 million in 2025, analysts project the sector to reach $15–20 billion by 2030 at a compound annual growth rate of 85–90 percent, and potentially exceed $75 billion by 2035 if commercial viability is demonstrated. Several organisations are already past the conceptual stage.

Starcloud

Starcloud, co-founded by Philip Johnston (a McKinsey alumnus) and Ezra Feilden, is the most advanced orbital data centre company in the Western world. Its first spacecraft — Starcloud One — launched in November 2025 carrying an NVIDIA H100 GPU, the most powerful processor ever deployed in space. The company demonstrated the ability to run a version of Google's Gemini AI model directly from orbit. A second spacecraft with 100 times the power generation of the first is planned for late 2026, followed by Starcloud 3 — a 200 kilowatt, three-tonne system designed for Starship deployment — around 2028. Starcloud's long-term vision is a constellation of up to 88,000 satellites providing approximately 20 GW of compute capacity. The company has raised $450 million in Series A funding from investors including NVIDIA, Sequoia Capital, and Lux Capital.

Google Project Suncatcher

Google's Project Suncatcher envisions an 81-satellite cluster built in partnership with satellite-imagery company Planet. The project focuses on flying satellites in tight clusters to minimise inter-satellite latency and validate distributed machine learning in orbit. Two prototype satellites are scheduled to launch in early 2027, with a long-term vision of a 1,000-satellite constellation providing 10 GW of AI computing capacity by 2032. Will Marshall, Planet's CEO, has stated: 'Orbital data centres are an idea whose time has come.' Google has committed $2–3 billion to the initial phase.

SpaceX and xAI

In March 2026, Elon Musk announced that SpaceX — recently merged with his AI company xAI — would put data centres into orbit. Musk unveiled the first-generation 'AI Sat Mini' spacecraft featuring solar arrays spanning approximately 180 metres, and proposed launching upward of one million satellites into polar orbit. SpaceX's vertical integration — combining Starship's payload capacity, Starlink's communications infrastructure, and xAI's AI workloads — positions it as uniquely formidable. SpaceX subsequently filed confidentially with the SEC for an IPO, citing orbital AI infrastructure as a core growth driver.

European ASCEND and China's Three-Body Constellation

The European Space Agency's ASCEND programme is studying CO2 reduction via orbital computing, funded with €300 million through 2027, with an explicit focus on European data sovereignty. China is taking a distinct approach, initially focusing on processing satellite data in orbit rather than transmitting raw data to Earth. Twelve satellites have already been launched, with plans for 2,800 more. China's programme is government-funded, government-mandated, and embedded in international policy frameworks — a contrast to the privately-led approaches dominant in the United States.

5. Challenges and How They Are Being Addressed

A balanced assessment of orbital data centres requires honest engagement with the obstacles. The case is compelling, but the engineering challenges are real and the timelines are uncertain.

- Launch costs. At current rates of approximately $1,000 per kilogram, large-scale orbital data centres are not yet cost-competitive. The business case depends on Starship and other next-generation vehicles delivering on their launch cost reduction promises. Starcloud's Johnston has made the dependency explicit to investors: 'If you don't think Starship is going to work, don't invest in us.'
- Scale of infrastructure required. The International Space Station's solar arrays — half the size of a football field — produce approximately 100 kilowatts of power. A 100 megawatt data centre in space would require a facility 500 to 1,000 times larger. MIT Professor Olivier de Weck notes that this is technically feasible but not achievable on short timelines.
- Radiation and hardware reliability. Commercial chips require extensive testing and shielding for LEO operation. Starcloud now has real-world telemetry from its first satellite providing empirical data on chip behaviour in orbit. Autonomous software-based fault management mirrors approaches already standard in terrestrial hyperscale operations.
- Maintenance and upgrade cycles. Orbital systems must operate autonomously for their design life — typically five years for first-generation systems. Modular designs and robotic servicing missions (being demonstrated by companies such as Northrop Grumman) offer longer-term solutions.

– Orbital debris. Mitigation strategies include responsive manoeuvrability, state-of-the-art tracking, and end-of-life demise designs. The dawn-dusk SSO chosen by Starcloud is a relatively uncrowded orbital regime, reducing collision risk.
– Regulatory frameworks. International space law — anchored by the 1967 Outer Space Treaty — was not designed for commercial data centre operations. Questions of data jurisdiction, export controls, and spectrum allocation require updated frameworks. The World Economic Forum's January 2026 Annual Meeting at Davos identified this as a topic requiring urgent multilateral attention.

> ***The organisations and nations that solve the regulatory and technical challenges of orbital computing first will hold a structural advantage that compounds over time. The obstacles are real. They are also being worked on systematically by well-resourced teams.***

6. The Sustainability Imperative

The sustainability case for orbital data centres deserves particular emphasis. The defining environmental challenge of terrestrial AI infrastructure is not merely carbon emissions — it is the combination of carbon, water, land, and waste heat. Space-based systems address all four simultaneously.

Orbital data centres powered by continuous solar energy carry near-zero operational carbon emissions. Lifecycle analysis suggests the carbon embedded in launch vehicles is offset within two to three years of operation. Water consumption is eliminated entirely. Land consumption is zero. Waste heat is exported beyond Earth's atmosphere rather than contributing to urban heat island effects.

A single 40 MW terrestrial cluster consumes over one million tonnes of water annually. An equivalent orbital facility consumes none.

Philip Johnston frames the opportunity in direct terms: 'By moving the most energy- and water-intensive parts of AI computation beyond Earth's atmosphere, we turn an existential challenge — energy demand — into a profound opportunity for decarbonisation, resilience, and a truly sustainable digital future.'

As environmental regulations on data centre operators tighten across the EU, UK, and elsewhere, the sustainability advantages of orbital computing will become a material business differentiator rather than a marketing claim. Organisations that understand this now will be better positioned when the regulatory pressure becomes unavoidable.

7. Geopolitical and Strategic Implications

Space-based AI computing carries strategic weight that extends well beyond commercial economics. Nations are beginning to recognise control of orbital computing infrastructure as a competitive advantage comparable to historical contests over territory, resources, and maritime routes.

The United States and China are engaged in a space computing race with both economic and military dimensions. China's government-mandated approach — with 12 satellites already in orbit and 2,800 more planned — provides a coherent, well-funded programme. The US privately-led approach offers greater innovation velocity but less coordination. Experts note that space-based AI infrastructure could provide militaries with computing capabilities beyond the reach of physical and cyber attacks on Earth.

Data sovereignty concerns add a further dimension. Orbital data centres operating in international space create complex legal questions around jurisdiction and data handling. For multinational AI operations aggregating data across multiple regulatory jurisdictions, neutral orbital territory may offer practical advantages. The European ASCEND initiative is explicitly motivated in part by the desire to avoid European AI processing being dependent on US or Chinese orbital platforms.

> ***Whoever controls orbital computing infrastructure in 2035 will hold a structural advantage in AI capability that is extremely difficult for latecomers to overcome. The window to shape that outcome is open now, and it will not remain open indefinitely.***

A less-discussed implication is the potential democratisation of AI compute for nations without terrestrial data centre infrastructure. Satellite connectivity combined with orbital compute could provide world-class AI capability to regions currently excluded from the AI revolution by infrastructure deficits.

8. The Technology Roadmap

The development of orbital data centres will proceed through distinct phases, each building on demonstrated capability. Understanding the timeline helps decision-makers position their organisations appropriately.

Phase 1: Proof of Concept (2025–2027)

Starcloud's first two satellites establish the feasibility of data-centre-grade compute in orbit. Google's Project Suncatcher prototypes launch in early 2027. China's edge computing constellation expands. Key milestones include demonstrated model training in orbit, validated cooling and power systems, and initial real-world telemetry on radiation effects.

Phase 2: Early Commercial Deployment (2028–2030)

Starcloud 3 and equivalent systems from Google and SpaceX begin providing commercial GPU capacity. Launch costs decline as Starship reaches operational cadence. First commercial orbital cloud services are offered to hyperscalers and AI neocloud providers. Hybrid Earth-orbit architectures become standard for large AI training runs.

Phase 3: Mainstream Adoption (2031–2035)

Launch costs approach $10–50 per kilogram, establishing cost parity with terrestrial alternatives. Hundreds of operational satellites provide 20–100 GW of orbital compute. Major cloud providers offer orbital computing as a standard service option. A significant portion of new AI training shifts to space-based infrastructure. Philip Johnston's prediction — that within a decade most new compute capacity will be deployed in space, representing approximately $1 trillion per year in capital expenditure — begins to look conservative rather than optimistic.

9. Space as the Infrastructure Layer for AI — A Board Governance Perspective

Based on my advisory work with boards and C-suite executives, orbital computing represents exactly the kind of structural technology shift that boards are consistently underprepared for — not because they lack intelligence, but because the framing arrives too late and too technical. By the time a technology appears on a board agenda as a risk, the window to influence it has usually passed.

The argument I make to boards is straightforward: space is not a novelty item for the AI conversation. It is the plausible infrastructure layer that resolves the single most binding constraint on AI at scale — energy. Boards that understand this now will ask better questions of their management teams, make better capital allocation decisions, and avoid being caught flat-footed when orbital compute becomes commercially mainstream within the decade.

> ***The organisations that will lead in AI are not necessarily those with the most compute today. They are those whose leadership understands, right now, where compute is going — and why.***

Five questions every board should be putting to the executive team

In my advisory work, I have found that the quality of a board's questions on technology strategy is a reliable leading indicator of organisational readiness. On orbital computing specifically, the following five questions are ones no executive team should be unprepared to answer by 2027.

– What is our energy plan for AI at scale, and what are the contingencies if terrestrial grid capacity does not keep pace with our compute demand? The energy constraint is not a future risk. It is a present reality in many markets. A board that does not understand its organisation's exposure to this constraint is not adequately governing the AI strategy.
– Which of our competitors or hyperscaler partners are investing in orbital computing, and what is the strategic implication if they achieve a structural cost or sustainability advantage through space-based infrastructure? Competitive intelligence on this topic should be a standing agenda item, not a one-off briefing.
– What is our position on data sovereignty as orbital computing matures? Data processed in orbit sits in an ambiguous regulatory jurisdiction. For organisations in regulated sectors — financial services, healthcare, government — this question will not remain academic for long. The answer should be developed now, not in response to a regulator's inquiry.
– Are we engaged in the policy and standards processes that will shape how orbital data centre regulation develops? The organisations that shape frameworks have a structurally better outcome than those that adapt to them. The window to participate meaningfully in international standards development is open. It will not remain open indefinitely.
– What is our three-year scenario plan if Starship achieves its cost targets and orbital computing reaches price parity with terrestrial alternatives by 2028–2029? This is Musk's stated timeline. Most analysts consider it optimistic. But scenario planning for an optimistic outcome — even a 20 percent probability outcome — is sound governance when the strategic implications are material.

Beyond these five questions, I observe that the boards most likely to navigate this well are those that already have directors with genuine technology infrastructure literacy — not digital marketing or social media experience, but a working understanding of how compute, energy, and data architecture interact at enterprise scale. If that literacy is absent from the board, building it is a governance priority.

Concluding Observations

When the writers of Pigs in Space sent Miss Piggy into orbit in 1977, they were going for laughs. As it turns out, they were going for accuracy. The infrastructure agenda that has followed is anything but a joke.

The most urgent challenge of the 2020s — how to power the AI revolution sustainably — has an answer that is, quite literally, out of this world. Orbital data centres are a technically grounded, economically credible, and strategically important response to the compounding failures of terrestrial infrastructure to keep pace with AI demand.

The energy arithmetic is compelling. Space-based solar is more than five times as efficient per array, costs up to 15 times less per kilowatt-hour, and requires no water, no land, and no decade-long permitting process. The organisations investing in this space — from Starcloud and Google to SpaceX, the European Space Agency, and the Chinese government — represent the full spectrum of the global technology and space industries. Their collective judgement is that the shift to orbital computing is not a question of if, but when and how fast.

Philip Johnston summarised the long-term view directly: 'If you look far enough out, the potential scale of compute demand is very large, with the upper bound essentially the total energy output of the sun. Within the next decade, I expect most new compute capacity to be deployed in space.'

Earth's resources are not sufficient for AI-driven data centres at the scale that AGI requires. Space offers not just an escape from those constraints, but a genuinely superior environment for computing at scale. The pigs are in orbit. The data centres are following. And the organisations that understand this now will be the ones best positioned when the rest of the world catches up.

Sources

This white paper draws on the following primary sources:

– Philip Johnston and Ezra Feilden, 'How Data Centres in Space Sustainably Enable the AI Age,' World Economic Forum, January 2026.
– Philip Johnston, interview with McKinsey & Company, 'The Case for Data Centers in Space,' McKinsey Technology Practice, 2025.
– Geoff Brumfiel, 'Will data centers in space work? Elon Musk says yes,' NPR, March 2026.
– 'Space-Based Data Centres: The Future of AI Computing in 2025,' independent market analysis, 2025.
– Documentary research: 'Data Centres in Space,' YouTube documentary transcript, 2025–2026.
– Starcloud Blog, 'How Data Centres in Space Sustainably Enable the AI Revolution,' 2025.
– International Energy Agency, Global Energy and AI report, 2025.

About the Author

Marco Ciobo is a senior transformation and strategy advisor with deep experience in technology strategy, digital transformation, and board governance across telco, utilities, banking, and government sectors in the ANZ region and globally. He is a Global Advisor to Bain & Company and advises boards and C-suite executives through his consulting practice Digital Insights and Transformation. He is the author of a forthcoming book on board governance in the digital era.

marcociobo.com