NOVASPACE

Space Data Centers

Exploring the Next Frontier in Computing

May 2026

Table of Contents

NOVASPACE

1 | Foreword

As humanity enters an era defined by artificial intelligence (AI) and ubiquitous connectivity, digital infrastructure is becoming as fundamental to civilization as energy, transportation, and water. The systems that power modern life — cloud computing, financial markets, logistics, communications, research, and others — now depend on vast and growing computational capacity. In many respects, humanity is beginning to organize itself around continuous access to data, compute, and machine intelligence. Training and operating advanced AI systems requires enormous amounts of energy, cooling, materials, and physical security. As these demands accelerate, terrestrial infrastructure alone may become increasingly strained by the competing pressures of cost, land use, grid capacity, resilience, and environmental impact. **The question is no longer only how to build bigger data centers on Earth, but whether some part of humanity's future industrial and computational base should ultimately expand beyond it.** In this context, space data centers emerge as part of a broader proposition: that the growth of human industry need not remain bound entirely to Earth's surface. As access to space improves, it becomes possible to imagine a gradual migration of selected heavy industries and digital infrastructure into space. If energy generation, large-scale manufacturing, resource processing, and computation can increasingly occur off-world, then Earth itself could be relieved of some of the most carbon-intensive and environmentally disruptive elements of growth. Thus, if space data centers prove technically and commercially viable, they might only be the beginning.

This **white paper delves into the opportunities and challenges of space data centers**, exploring the rationale of the computational shift to space, different possible implementations, key players, and possible future trajectories.

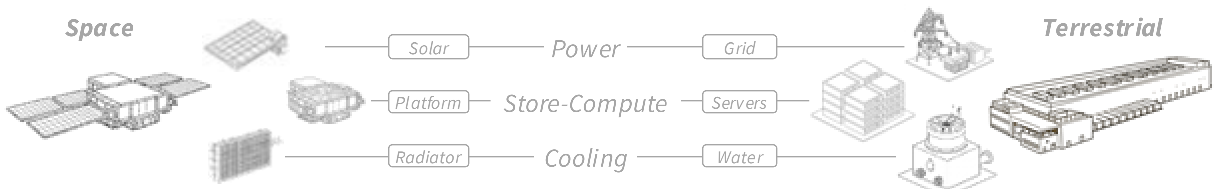

NOVASPACE

Meet The Authors


Carla Filotico
Partner & MD
carla.filotico@nova.space


Maxime Puteaux
Principal
maxime.puteaux@nova.space


Jean-Loup Rondeau
Principal
jean-loup.rondeau@nova.space


Felix Rottmann
Senior Manager
felix.rottmann@nova.space


Lucas Pleney
Manager
lucas.pleney@nova.space


Ivan Iushchenko
Consultant
ivan.iushchenko@nova.space


Claire Fortin
Consultant
claire.fortin@nova.space

2 | Evolutionary logic of the compute shift to space

The development of space data centers is not an isolated event, but the **convergence of two technological paths**: terrestrial computing and space systems. On Earth, computing evolved from early electronics into hyperscale cloud infrastructure and now into larger AI-scale systems, each wave increasing **demand for power**, **cooling**, and **land**. In parallel, space systems evolved from the first satellite into communications, Earth observation, navigation, and today's dense, interconnected **LEO constellations**. The intersection of space systems and AI is already emerging in Earth observation, where edge processing in orbit reduces the volume of data sent back to Earth.

The idea of orbital compute emerges where these trends meet. As terrestrial AI infrastructure pushes against **physical** and **economic constraints**, space systems are becoming **more capable** and **operationally mature**. Space data centers can represent not just another space application, but a **shift toward** a **distributed infrastructure** across Earth and space and toward multi-mission architectures. In turn, space data centers can create a new momentum for space systems.

The push to move computing into space is enabled by **recent developments** in **space systems.** Without the following advancements, achieving the mass, power, autonomy, launch cadence, and communications performance required for meaningful in-space computing would be out of reach:

- **Reusable launchers:** Access to space has become more affordable and frequent, enabling scalable deployment of infrastructure
- **Low Earth Orbit (LEO) constellations:** Massive LEO constellations reflect a step-change in manufacturing capabilities and supporting infrastructure over the past decade. High-speed, reliable optical inter-satellite linked links make it possible to distribute capacity
- **Computing in space:** Hardware is now powerful and resilient enough to be reliable in space and has been showcased by edge computing for EO
- **Power generation & thermal management:** Advancements in solar arrays and thermal systems are making energy generation and heat dissipation for compute-intensive systems feasible

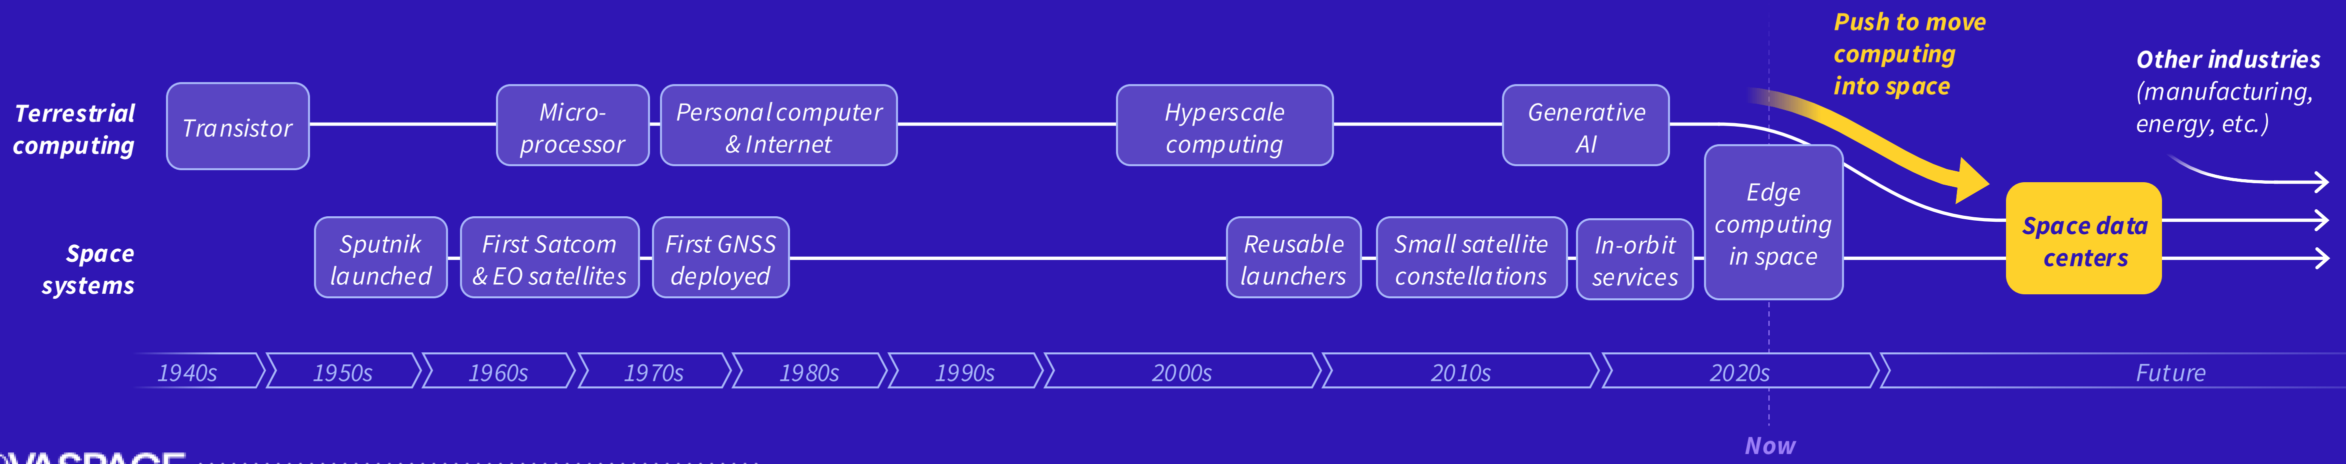

3 | Rationale for space data centers

	Terrestrial pain points	*Space solutions*
Real estate	• *NIMBY effect* • *Land constraints* • *Displacement of locals* • *Noise pollution*	• *Capacity added off-Earth*
Energy & resources	• *Increasing power consumption, stress on grid infrastructure* • *Fresh water usage for cooling*	• *Continuous solar power* • *Radiative cooling without freshwater use*
Regulations & environment	• *Permission bottlenecks* • *Ever-changing local laws* • *High operational CO_2 emissions*	• *Less exposure to local regulations* • *No CO_2 emissions during operations*

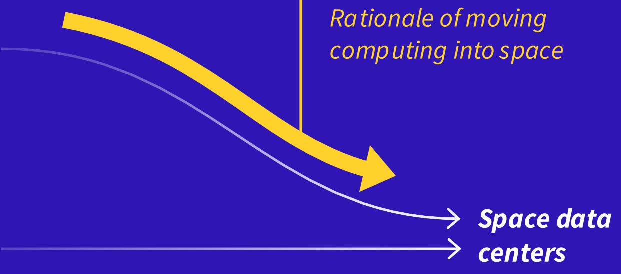

NIMBY = Not In My Backyard

NOVASPACE

Space data centers are gaining attention as **AI increases demand for compute capacity** and puts pressure on the physical footprint and energy grid. As a result, the debate is shifting toward the **practical constraints of scaling** new capacity. In this context, three pain points can be considered:

Real estate – The expansion of terrestrial data centers is often hindered by local opposition over land use, noise, and broader community impact. This pattern is often described as the "**Not In My Backyard**" (NIMBY). In the U.S., some proposals and policy discussions have called for pauses or tighter safeguards around new data center construction. Project delays are already affecting parts of the U.S. data center pipeline. By contrast, space would allow **capacity to be added off-Earth**.

Energy & resources – Global data center electricity demand is expected to approach 1,000 TWh by 2030. **Grid expansion** usually moves much slower than compute deployment, and water use for **cooling** is coming under scrutiny. In California, lawmakers are pushing for limits on water use amid concerns about drought. Space offers near-continuous solar energy and no dependence on water.

Regulations & environment – Local laws, approvals, and shifting regulations **slow terrestrial data centers' development**. In New York, a 2026 bill pauses approvals for data centers while their impact is assessed. **Environmental pressure** is also growing. By 2025, U.S. data centers were estimated to account for 100+ Mt of CO_2 annually. Meanwhile, space solutions are subject to **international law** and **national authorities**, while also removing **CO_2 emissions** during operations.

Nonetheless, these arguments remain debatable. Space data centers do not use land, but LEO is becoming congested, which already hurdles astronomers. Environmental advantage during operations is offset by launch and manufacturing impacts. Energy and water usage concerns can be balanced by available efficiency gains from regular hardware updates, which are much harder with space solutions.

4 | Current and future use cases

Use case	Maturity	Value of performing computations in space	Key benefits
Earth Observation Imagery processing	Very high	Very high	• Latency reduction • More efficient utilization of limited data downlink capacity • Sensor proximity
Defense ISR, secure sovereign storage/ compute	High	High	• Sensor fusion of different data streams • Persistent coverage • Latency reduction
Connectivity Compressing backbone & gateway traffic	Medium	Medium	• Bandwidth efficiency increase and reduction of network congestion • Reduction of backhaul traffic • Increased resilience
AI compute GPU as-a-service, AI inference workloads	Low	Very high	• Bypassing terrestrial real estate, energy, and resource challenges • Possible ability to scale
Niche applications Robotic exploration, in-situ compute capacity, crypto mining	Very low	Medium	• Latency reduction (esp. for distant solar system missions) • In-situ computing

Time to market

Defense and **Earth Observation** stand out today as the **most credible early space data centers** use cases. Both benefit from a strong orbital value proposition, show relatively high maturity, and can generate demand without requiring very large-scale deployment. Defense in particular appears to be one of the clearest near-term entry points, given stronger willingness to pay for mission-critical capabilities.

Connectivity appears to be a more **medium-term opportunity**. While the value proposition is tangible, adoption is likely to depend on broader constellation scale and continued progress in launch economics.

Data computing in space for space is possible in the short/medium-term. In contrast, **AI inference** for users on Earth at scale is extremely compute-intensive and dependent on moving huge amounts of data quickly between chips, storage, and networks. While this use case may offer tremendous potential, the maturity is low and many technical problems and open commercial questions still need to be answered.

Niche applications may also emerge over time, though demand visibility and monetization models are still less established.

5 | Blueprints for possible implementations

As companies seek to establish their own space data centers, they face a spectrum of architectural options, ranging from decentralized and centralized systems to hybrid models in between. Each has its own advantages and disadvantages and deciding which to implement is shaped by the operator's resources and planned use cases. Most of today's proposed implementations focus on Low Earth Orbit (LEO), but other options are also being discussed.

Decentralized architecture

A decentralized architecture comprises a constellation of data centers, or "**nodes**," that can compute independently or communicate amongst each other for larger work demands. As a network, it is highly fault tolerant and offers greater proximity to the user, decreasing latency. It presents an ideal entry point for multi-mission topics to decrease risk. Yet, due to the low computing capacity of individual nodes, large demands require **complex orchestration** and **distributed computing** across nodes, sometimes resulting in less uniform data and inter-node latency. To remain technologically competitive, operators may need to refresh much of the constellation with each generation if decentralized hardware cannot be upgraded in-orbit, creating a less sustainable model. However, a decentralized architecture benefits from low marginal costs as the nodes can be mass manufactured and rapidly deployed, with launch costs a particularly important factor in economic feasibility.

Centralized architecture

A centralized architecture features a single large space data center or a cluster of data centers in orbit. These centers have much **higher computing power** and are ideal for **tightly coupled workloads**; however, the greater distance to many users results in increased latency. Centralized space data centers are comprised of several "modules" that can theoretically be added and removed in-orbit, enabling the operator to scale the data center or conduct hardware updates without having to decommission the base infrastructure, which can have a >10-year lifespan. However, in practice, performing upgrades will be challenging and require an in-orbit servicer, although this maintainable structure also offers sustainability and adaptivity benefits over replacement. Moreover, due to their complexity and size, the centralized approach entail higher manufacturing costs and launch requirements and expose the operator to a single point of failure which would be costly to replace.

Future architectures may also be more nuanced than centralized/decentralized, with the possibility of modular federated systems that scale incrementally and support both distributed and tightly coupled workloads.

No major concerns
Likely manageable
Hard to solve issue

5 | Blueprints for possible implementations

	Decentralized architecture		Centralized architecture	
Investment type	*Lower marginal costs due to mass manufacturing and batch launches, but regular refresh required. Multi-mission applications*	**Phased**	**Upfront**	*Manufacturing and launch costs are high, but less refresh required, so less mass to be launched over time. Can be scaled in phases*
Space system operations	*Operation of a constellation makes it more difficult to track assets and conduct manoeuvres*	**Complex**	**Simpler**	*Operation of a single (or a few) infrastructure facilitates tracking and operation of assets*
Compute & storage operations	*Requires complex orchestration across nodes, distributed computing power*	**Complex**	**Simpler**	*Less data centres reduces the need for orchestration, computing simpler*
Resilience	*Risk is distributed across thousands of satellites, but more vulnerable to space debris due to congested orbits. Risk due to deployment*	**High**	**Medium**	*Single point of failure, but less likely to experience collisions due to less congested orbit. Can be managed by design (opting for a cluster)*
Terrestrial infrastructure	*Complex orchestration of terrestrial infrastructure, multiple ground stations required*	**Complex**	**Simpler**	*Minimal ground stations and orchestration of terrestrial infrastructure required*
Infrastructure deployment and maintenance	*Quickly manufactured, deployed and refreshed due to mass production and batch launches. Hardware cannot be updated in orbit*	**Simpler**	**Complex**	*Complex to manufacture, launch and assemble (technology not active), in-orbit maintenance is operationally challenging*

Distributed computing power increases complexity, but overall feasibility is in line with existing technology and practices

Unproven, critical technological challenges regarding assembly and maintenance. Severity of disadvantages outweigh benefits

6 | Costs of space data centers

Total cost of installing and operating a
1 GW data center for 5 years

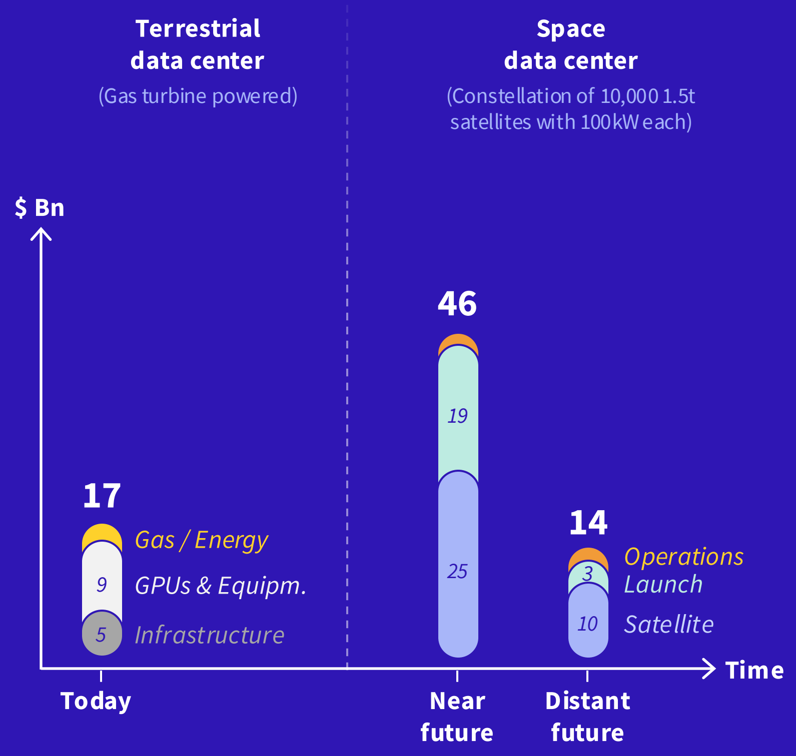

Launch price	$1,000/kg	$150/kg
Satellite cost	$25/W ($2.5m)	$10/W ($1m)

Space data centers for AI training & inference are still far from being cost-competitive with terrestrial alternatives today. When running an optimistic business case using launch costs below today's industry rate and satellite manufacturing costs below what mega constellation are achieving today, space data centers are still nearly three times more expansive than traditional terrestrial data centers: While a terrestrial **1 GW data center** costs around $17 Bn over a timeframe of 5 years, a space data center with the same computational power would cost $46 Bn. The bulk of the cost stems from the cost of the space hardware itself, the rest originates from launching the hypothetical constellation of **10,000 satellites**. After 5 years, the constellation would need to be fully replenished.

To get close to the economics of the terrestrial equivalent, the key cost drivers of satellite manufacturing and launch must come down significantly. In a future where **fully reusable launch vehicles** enable a launch cost erosion by a factor of 10 compared to today's values and orbital data infrastructure/satellite manufacturing at scale pushes down manufacturing cost, space data centers could theoretically become economically viable. However, this simplified view does not consider the immense costs required for the **initial development** of the technology, possible technical hurdles that must be overcome which might impact both cost and time to market, as well as uncertainties regarding **revenue generation**. While customers of orbital data centers might be **willing to pay a premium** due to terrestrial roll-out bottlenecks, orbital systems are unlikely to match the uptime and service level agreements standards of terrestrial hyperscalers.

In order for large-scale AI space data centers to become commercially viable, they must deliver either a compelling premium advantage that justifies the **structurally higher costs** and **lower service reliability** or be able to offer cost parity with terrestrial infrastructure. **Scaling satellite and launch technology** to a level which cost parity is achieved will most likely **take years** if not decades.

However, it must be noted, that smaller-scale space data centers used for defense, Earth observation processing, sovereign compute, and latency-sensitive edge applications might reach commercial feasibility way before AI data centers in space.

7 | Required building blocks

	Building block	Importance	Readiness	Cost impact
Satellites	*Large-scale manufacturing*	Important	Available	**$$$$** *Very high*
	Radiation-hardened components	Crucial enabler	Available	
	Dedicated processing units	Optional	Near-ready	
	Cooling systems	Crucial enabler	Development needed	
	Large solar arrays & power management	Crucial enabler	Near-ready	
	Optical links & network architecture	Important	Fully mature	
	Software & decentralized computing	Important	Near-ready	
Launch	*Super-heavy launch vehicles*	Important	Development needed	**$$$** *High, but can be reduced mid-term*
	Fully reusable launch vehicles	Important	Development needed	
	Launch infrastructure	Crucial enabler	Fully mature	
Operations	*Ground stations & operations*	Crucial enabler	Fully mature	**$$** *Medium*
	Optical ground stations & QKD	Optional	Development needed	
	In-orbit servicing	Optional	Breakthrough needed	
	Debris removal & SSA	Optional	Development needed	

...

Crucial enabler *Important* *Optional*

Breakthrough needed *Development needed* *Near-ready* *Available* *Fully mature*

Space data centers depend on **four main categories**: satellites, launch, and operations. Other factors, such as regulations, affect deployment and scaling, but are not the main constraints at this stage.

The **satellite** category is the largest cost driver. Radiation-hardened components and / or complementary software mitigations, power systems, and cooling are the key enablers of space data centers. Among these, cooling is the least mature and remains one of the challenges. Mass manufacturing, optical links, and decentralized computing are relevant for decentralized architectures. Special processing units are useful, but not essential, as protected conventional processors may be used.

Launch is the second-largest cost driver. Launch infrastructure is already mature, while reusable launch vehicles are becoming the new standard. **Super-heavy** and / or **fully reusable** vehicles are valuable for larger concepts, but they still require development, while smaller projects can be developed without them.

The **operations segment** is smaller in cost and consists of ground and lifecycle. Ground stations are a basic requirement and are already mature. Optical ground stations and quantum key distribution (QKD) can improve throughput and security but are not essential. In-orbit servicing could eventually support upgrades and longer asset life but remains speculative. Debris removal and space situational awareness (SSA) matter for resilience and sustainability than for feasibility.

Overall, satellites remain the largest cost bucket, but the **biggest opportunity** to improve the business case **comes from launch**, where progress in infrastructure, reuse, and eventually super-heavy lift could materially reduce deployment costs.

8 | Key players & projects

At least 13 players are developing or considering space processing solutions, with varying levels of ambition. SpaceX, Blue Origin, and a few others aim for full-scale AI data centers, while players like Kepler and Space Compass focus on data relay with additional processing. Others, like Star.Vision, integrate onboard AI into EO, while Sophia Space and Axiom explore modular compute deployed as ed or standalone systems.

Vertical integration is a key differentiator. Aside from the state-coordinated Chinese CASC ecosystem, SpaceX is the only player close to full integration, while most others have gaps across the value chain.

Launch remains the **main bottleneck**. Many concepts consider large constellations or high-power systems, making capacity a structural constraint. As a result, most players are likely to depend on SpaceX as their primary launch provider, despite it being a competitor. To mitigate this, OpenAI has reportedly explored investing in launch startup Stoke Space.

Demand is another challenge. Only a few projects have internal workloads, while others must rely on external partners to make the business viable.

Compute hardware is also constrained. Even large lack proprietary chip manufacturing, pushing many companies toward NVIDIA as a default supplier, which becomes central to the ecosystem.

Despite these constraints, the space data center **landscape remains highly dynamic**, with frequent new developments reshaping the field.

Operator	Use case	Planned size (n. sats)	Status	Value chain: Compute HW	Value chain: Sat platform	Value chain: Launch	Value chain: Demand
SpaceX	AI compute	1,000,000[1]	FCC filing in Feb 2026	Integrated	Integrated	Integrated	Integrated
Blue Origin	AI compute	51,600[1]	FCC filing in Mar 2026	External	Integrated	Integrated	External: aws
Starcloud	AI compute	88,000[1]	1 sat launched in 2025	External: NVIDIA	Integrated	External	External
Google	AI compute	"A network of satellites"	2 demo sats in 2027	Integrated	External: demos planet	External	Integrated
OpenAI	AI compute	TBD	Exploratory phase[2]	External: NVIDIA	External	External: STOKE	Integrated
Aetherflux	AI compute	TBD	1st sat in Q1 2027	External: NVIDIA	Integrated	External: SpaceX	External
Axiom Space	Compute	TBD	2 nodes launched in 2026	Integrated	External: incl. Kepler	External: SpaceX	External
Sophia	Compute	TBD	$10M seed round in Feb 2026	External: NVIDIA / Integrated	External: incl. APEX	External	External
Kepler	Satcom data relay	TBD	10 sats launched in Jan 2026	External: NVIDIA	Integrated	External: SpaceX	External
Space Compass	EO data relay & compute	TBD	Announcement in Mar 2026	External: Microsoft / Integrated: Partner	External: Swissto12	External	External
ADASpace / Zhejiang Lab	AI compute	2,800	12 launched in May 2025	Integrated	Integrated	External: CASC	External: CASC
Star.Vision	EO data compute	258 OSE sats by 2032	2 sats in May 2026	Integrated	Integrated	External: CASC	Integrated
CASC	AI compute	TBD	Feasibility study in Apr 2026	Integrated	Integrated	Integrated	Integrated

...

Integrated | External

1) According to FCC filings; 2) Sam Altman publicly doubts the idea; HW = Hardware;

9 | Outlook & possible future trajectories

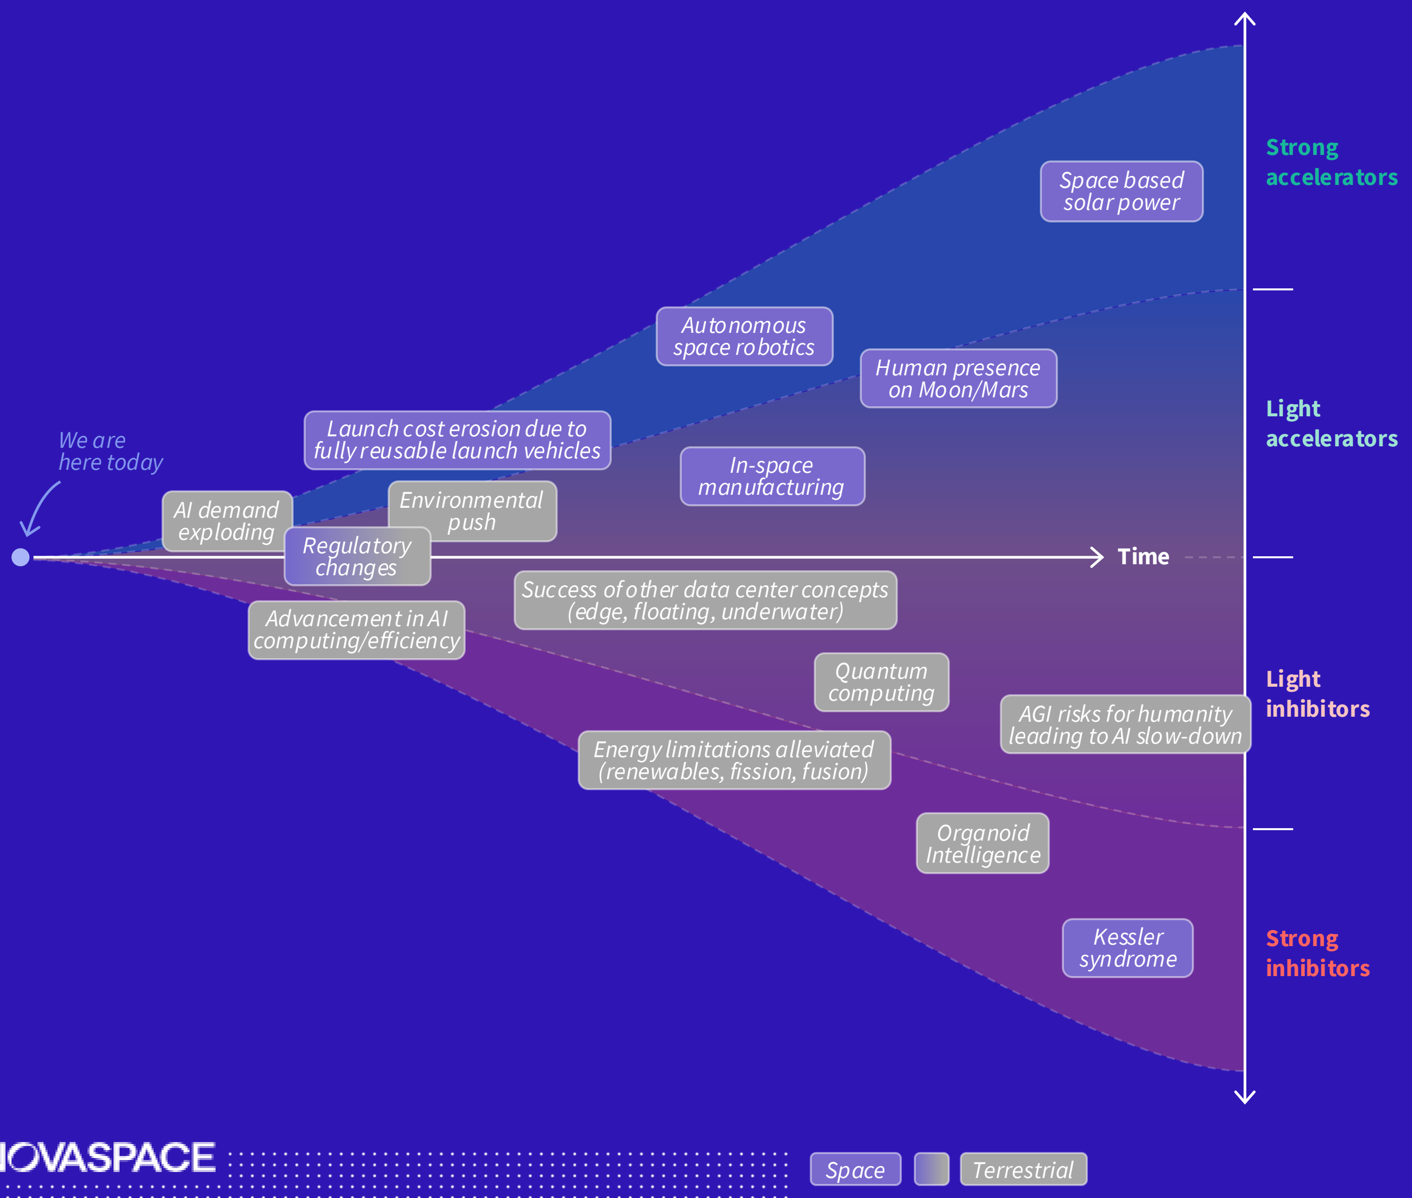

The future of space data centers remains uncertain, shaped by both terrestrial and space-based advancements that could dramatically alter their feasibility. On the one hand, terrestrial innovations such as **breakthroughs in AI computing** (improvement in algorithms, AI-specific chips, novel data center concepts, neuromorphic computing, quantum computing) or the **alleviation of energy limitations** could mitigate some of the advantages currently offered by space data centers, making ground-based infrastructure more competitive. On the other hand, the development of space technologies like **fully reusable launch vehicles** and autonomous space robotics could significantly lower the cost of deploying, and updating space data centers, making them a commercially more viable option.

In the short- to mid-term, **unanswered questions** regarding economics, regulatory aspects as well as the **sustainability** of space data centers and their impacts on the environment, add uncertainties. In the long-term, scenarios like the Kessler syndrome or unforeseen negative impacts of AI on humanity might pose additional risks.

Given these uncertainties the future of space data centers remains unpredictable. Their success will depend on a delicate balance of developments across multiple domains. What is clear, however, is that space data centers will bring about **significant changes** to both the **space and computing industry** if the concept proves successful.

NOVASPACE

Novaspace is a global leader in space consulting and market intelligence with over 40-year legacy of expertise in guiding public and private entities in strategic decision-making. Novaspace offers end-to-end consulting services, from project strategy definition to implementation, providing data-led perspectives on critical issues. Novaspace presents an expanded portfolio of services, featuring combined expertise in management and technology consulting, top-tier executive summits, and market intelligence. Trusted by more than 1,200 clients across 60+ countries, Novaspace operates from offices in Brussels, London, Montreal, Munich, Paris, Singapore, Sydney, Tokyo, Toulouse, Washington D.C., and Bangalore, India.

Follow us

How we can support you

Investors

We provide commercial and technical due diligence for financial and strategic investors. Our advisory team can support in assessing demand, technology maturity, regulatory constraints, competitive positioning, in order to evaluate risk and time investments effectively.

Industry

We can advise your business on market positioning, technology advisory, go-to-market strategy, and building the right partnerships across the in-space infrastructure value chain.

Policy makers & institutions

Shaping government space strategies is one of our key offerings. Applying our socio-economic impact assessment and cost-benefit analysis frameworks to space data centers, we can also help to ensure that public investments yield optimal benefits.