

The Economics of AI in Space



Anirudh Bhoopalam, Ph.D.
Analyst



Karthik Subramanian
Analyst

LUX TAKE

Utilities and data center developers should remain focused on terrestrial infrastructure solutions rather than orbital data centers. While orbital compute may find niche applications in Earth observation, defense, and autonomous space operations, launch economics, thermal management, communications infrastructure, and reliability challenges limit its viability as a large-scale alternative to terrestrial AI infrastructure.



As AI workloads continue to grow, data center developers are facing constraints related to power availability, land acquisition, permitting timelines, and water consumption, among others. In response, they are exploring alternate approaches to increase compute capacity, such as creating distributed data center networks or putting compute capacity underwater. Alongside these routes, a growing number of companies are exploring orbital data centers as a potential alternative to terrestrial infrastructure. Interest in the concept has been fueled both by increasing pressure on terrestrial data center infrastructure and by declining launch costs and rising launch activity, driven largely by commercial providers like SpaceX. The concept places compute hardware in orbit, powered primarily by solar energy and cooled through radiative heat rejection, with proponents arguing that it could bypass many of the constraints facing data center development on Earth.

Interest in orbital computing has accelerated over the past two years. Starcloud demonstrated AI inference and training workloads on an Nvidia H100 GPU operating in orbit. Cowboy Space Corporation is pursuing solar-powered orbital AI infrastructure and large-scale orbital compute networks. Axiom Space has launched orbital data center nodes intended to support cloud computing services in low Earth orbit, while satellite operators and infrastructure providers are exploring in-orbit processing capabilities for communications, Earth observation, and defense applications.

While recent demonstrations and announcements have created momentum around orbital computing, there are questions around the concept's scalability and capability to absorb some of the growing terrestrial compute load. In this brief, we assess whether orbital computing can become a viable alternative to terrestrial infrastructure and where the primary technical and economic constraints are likely to emerge.

Launch economics remain the largest barrier to scaling orbital compute. While AI hardware receives most of the attention, the majority of mass in a large orbital computing platform would come from supporting infrastructure rather than the compute hardware itself. Solar arrays, radiators, structural components, shielding, communications equipment, and power management systems all contribute substantial mass. As a result, overall economics remain highly dependent on launch costs. Concepts like Starcloud assume launch costs on the order of USD 30/kg, whereas current commercial launch pricing remains closer to USD 3,000–USD 3,500/kg based on published Falcon 9 pricing. At scale, this two-order-of-magnitude gap becomes difficult to overcome because the system is increasingly dominated by noncompute mass rather than the graphic processing units (GPUs) themselves.

The energy argument further reinforces this challenge. While orbital data centers offer access to abundant solar energy, capturing and using that energy requires large solar arrays, power systems, and supporting infrastructure that add significant mass. This increases both launch costs and deployment complexity as systems scale.

Thermal management remains a fundamental constraint despite common perceptions that space naturally solves cooling challenges. Space is cold — temperatures in deep space plummeting to as low as 3 K (–270 °C) — but this doesn't mean the cooling problem is solved. In a vacuum, heat cannot be removed by convection and must instead be rejected by radiation. While Starcloud demonstrated that hundreds of watts of GPU heat can be managed in orbit, scaling to megawatt- or gigawatt-class systems requires correspondingly larger radiator systems. While companies like Starcloud are exploring liquid cooling and deployable radiator systems, these approaches don't eliminate the underlying scaling challenge: Larger compute deployments require proportionally larger heat-rejection infrastructure. Even under favorable assumptions, radiative heat rejection is measured in hundreds of watts per square meter, implying that megawatt-scale orbital compute would require thousands of square meters of radiator area, while gigawatt-scale systems would require square kilometers of deployed thermal infrastructure. Since nearly all electrical power consumed by AI compute hardware eventually becomes waste heat, cooling becomes a materials, area, and launch-cost challenge rather than disappearing altogether.

Serving terrestrial workloads requires solving a difficult data-transfer problem between orbit and Earth. Orbital compute can create value when data originate in space, particularly for Earth observation, defense, and autonomous space operations, where processing data in orbit reduces downlink requirements. However, serving terrestrial cloud customers requires moving large volumes of data between orbit and Earth. NASA's TeraByte InfraRed Delivery demonstration achieved optical downlink speeds of up to 200 Gbps, showing that high-bandwidth communications are technically feasible. Yet delivering hyperscale-grade connectivity remains challenging because optical links require line-of-sight connections to ground stations, are affected by cloud cover and atmospheric conditions, and depend on a distributed network of receiving infrastructure to provide continuous coverage. As a result, communications remain a significant constraint for large-scale orbital compute serving terrestrial users.

In 2025, BIC urged the federal government to attract major investments in large-scale biorefineries and biomanufacturing facilities capable of anchoring new value chains around Canadian biomass resources. Canada has begun to make targeted scale-up investments, but these remain project specific rather than part of a broad market-creation policy for biobased chemicals. As in the U.S., the key gap is not the absence of research capability or biomass availability; it is the lack of coordinated demand-pull mechanisms, commercialization financing, and large-scale industrial infrastructure needed to translate Canada's resource base and innovation ecosystem into durable domestic manufacturing capacity.

Reliability becomes more difficult once compute infrastructure is deployed in orbit. Terrestrial data centers benefit from routine maintenance, hardware replacement, and controlled operating environments, while orbital systems must operate in radiation-rich environments with limited opportunities for repair or upgrade. Starcloud-1, for example, is expected to operate for roughly 11 months before deorbiting. Scaling orbital compute would therefore require managing hardware failures, technology refresh cycles, orbital debris risks, and end-of-life disposal across potentially thousands of satellites, increasing operational complexity and cost. The same solar energy and cooling equipment would have a far longer lifetime on earth, meaning capex will go further in terrestrial applications.

Outlook

Orbital data centers have progressed from a theoretical concept to early commercial experimentation, but significant questions around economics, scalability, and operational reliability remain unanswered. The most credible near-term opportunities are in Earth observation, defense, and autonomous space operations, where processing data in orbit can reduce transmission requirements and improve responsiveness. Utilities and data center developers should continue prioritizing terrestrial power and grid infrastructure investments, as orbital compute won't materially offset growing AI-driven demand for data center capacity.



For More Information



VISIT

www.luxresearchinc.com



READ

www.luxresearchinc.com/blog/



CONNECT

www.linkedin.com/company/lux-research



FOLLOW

[@LuxResearch](https://twitter.com/LuxResearch)



CONTACT

questions@luxresearchinc.com



INSPIRE
ILLUMINATE
IGNITE