Space-Based Data Centers: State of the Art, Economics, and Outlook

Jose E. Puente
402T Labs, SL
jepuente@402T.com

Carlos Puente
402T Labs, SL
carlospuenteg@402T.com

Dec 16, 2025

Abstract

Space-based data centers (SBDCs) put compute and storage in orbit and, in some proposals, in cislunar space or on the lunar surface. Recent on-orbit demonstrations have already run server-class hardware—and even GPU-class payloads—long enough to support early "edge cloud" services. Feasibility is no longer the question; unit economics are. This report summarizes the December 2025 state of compute payloads, power, thermal design, and communications; reviews visible commercial and institutional efforts; and uses a quantitative "cost-per-watt" lens to compare orbital and terrestrial deployments. The numbers point to near-term value in in-space analytics and resilience-driven storage, while commodity-scale compute in orbit still hinges on launch cost, spacecraft specific power, and the mass and area required for heat rejection.

Executive Summary

We map the engineering, economics, and operations trade space for space-based data centers (SBDCs) and identify where the concept can pencil out. In most cases, compute itself is not the bottleneck. Power generation, heat rejection, launch logistics, and end-to-end data movement set the ceiling on useful throughput.

- **Dominant constraints:** Specific power (W/kg), radiator area per rejected kW, and reliable link capacity set practical compute density and duty cycle.
- **Economics:** A practical first-order yardstick is *fully-loaded cost per delivered watt of compute*, plus a *cost per delivered bit* for workloads that must exchange data with Earth frequently.
- **Near-term fits:** sovereignty/assurance use cases (tamper resistance and jurisdictional separation), specialized remote sensing processing, and intermittent "burst" compute aligned with orbital power/thermal windows.
- **Key risks:** Radiation-induced faults, thermal transients, launch cadence and insurance, spectrum and regulatory constraints, and single-point failures in power and thermal chains.
- **Outlook:** The fastest way to de-risk the concept is incremental demonstrations that tie power/thermal budgets to measured workload throughput (not just peak TFLOPS).

Acronyms

Acronym	Meaning
SBDC	Space-Based Data Center
COTS	Commercial Off-The-Shelf
LEO / MEO / GEO	Low / Medium / Geostationary Earth Orbit
ISS	International Space Station
GPU	Graphics Processing Unit
LCOE	Levelized Cost of Energy
PCDU	Power Conditioning and Distribution Unit
GNC	Guidance, Navigation, and Control
RF	Radio Frequency
IR	Infrared
ASAT	Anti-satellite (capability/weapon)
BESS	Battery Energy Storage System

Executive Summary

- **Technical feasibility is no longer the primary uncertainty.** Flight results have taken most of the mystery out of running general-purpose compute in space: ISS-class demos and LEO GPU payloads are now public, and several vendors sell on-orbit edge compute modules.

- **Economics remain the gating factor for commodity-scale compute.** First-principles cost models put the constraint where you would expect: dollars per usable watt, driven largely by power-system mass and the thermal hardware needed to reject waste heat.

- **Near-term "best-fit" workloads are (i) in-space analytics and (ii) resilience/security-driven storage.** These segments benefit from proximity to the data source, reduced downlink volume, and physical isolation.

1 Scope, Terminology, and Assumptions

1.1 Scope

This is a technical overview, not investment advice. It focuses on low Earth orbit (LEO) compute payloads and "micro data center" satellites, with only light coverage of lunar and cislunar storage proposals. Unless otherwise stated, the discussion assumes solar-powered spacecraft with radiative heat rejection.

1.2 Terminology

Definition 1 (Space-Based Data Center (SBDC))**.** *An SBDC is a space-deployed system that provides compute and/or storage as an operational service. Implementations range from a single payload in LEO to multi-satellite clusters and (in some proposals) lunar deployments.*

Definition 2 (On-Orbit Edge Compute)**.** *Compute performed in space primarily to process data generated in space (e.g., Earth observation, comms payload analytics), with results downlinked to ground.*

Assumption 1 (Cost Comparisons)**.** *When comparing orbital to terrestrial deployments, the most defensible comparisons are made on* delivered usable electrical power to compute *($/W) and the resulting levelized cost of energy (LCOE), under explicit mass, launch, degradation, replacement, and operations assumptions [1].*

2 Introduction and Motivation

Most discussions of SBDCs start with hardware (solar arrays, radiators, GPUs in vacuum) and only later ask whether the numbers work. A better starting point is the value proposition: *when is a watt (or a flop) in low Earth orbit more valuable than a watt on the surface?* One first-principles model argues that the case rises or falls on delivered usable power cost (and the ability to reject waste heat), not on "big-number" intuitions [1].

Commonly cited motivations include:

- **High sunlight fraction** in selected orbits, improving capacity factor and reducing reliance on terrestrial grids [1, 3].
- **Physical isolation** for high-assurance storage (a form of "air gap" achieved by distance and controlled ground links) [10].
- **In-space processing close to data sources** (Earth observation, communications constellations), reducing downlink burden and latency-to-insight [8, 14].
- **Potential long-run sustainability advantages** (reduced on-Earth cooling water demand; operations powered by solar), contingent on lifecycle impacts dominated by launch and manufacturing [15].

3 Current State of the Art

3.1 Compute Platforms and On-Orbit Demonstrations

Early missions showed that commercial off-the-shelf (COTS) server-class hardware can run in microgravity if the overall system is built to tolerate faults and resets. The HPE Spaceborne Computer program on the ISS remains a widely cited reference point for general-purpose computing in space [4, 5, 6].

Work is now moving toward higher-performance payloads and "micro data-center" satellites. Starcloud, for example, reports an on-orbit demonstration with an NVIDIA H100-class GPU payload as a step toward orbital AI services [7, 8]. In parallel, European efforts (e.g., Unibap with D-Orbit) have demonstrated in-orbit edge computing modules that run applications for analytics directly in space [14, 13].

3.2 Reference Architecture

Figure 1 shows a minimal reference architecture for an SBDC node. The diagram is intentionally technology-agnostic; the key point is that *power, thermal, and communications* subsystems are first-order sizing drivers rather than implementation details.

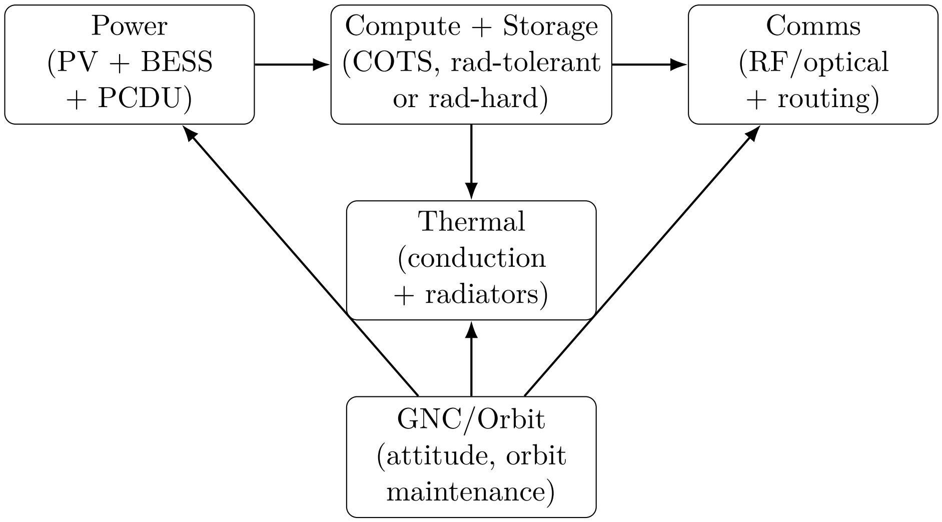

Figure 1: Minimal reference architecture for a space-based data center node.

3.3 Thermal Management in Vacuum

Thermal management is still a gating constraint. Without convection, spacecraft must move heat by conduction to radiators (or radiative structures) and then reject it via radiation; that plumbing drives mass and mechanical complexity [3]. In steady state, radiative rejection can be written as

$$Q_{\text{out}} \approx \epsilon \sigma A \left(T^4 - T_{\text{bg}}^4 \right), \tag{1}$$

where ϵ is effective emissivity, σ the Stefan–Boltzmann constant, A radiating area, and T_{bg} the background temperature. The hard part is not the equation; it is reliably *delivering* heat from dense compute to radiative surfaces without hot spots, fragile interfaces, or unacceptable temperature margins.

Scaling toward multi-kilowatt nodes typically implies large deployable radiator areas and careful control of incident solar and Earth IR loading [3]. Published analyses emphasize that optimistic thermal assumptions can swing mass and cost estimates materially, making thermal modeling a first-order driver in economic assessments [1, 2].

3.4 Power Generation and Storage

Most near-term architectures assume photovoltaic (PV) generation plus batteries. Sun-synchronous and near-terminator orbits can reach very high sunlight fractions, improving capacity factor and reducing the burden on energy storage [1, 3]. A smaller set of concepts extends this to power beaming, where orbital solar collection transmits energy via lasers to other spacecraft or to ground infrastructure [9].

3.5 Communications and Networking

SBDCs depend on high-bandwidth downlinks and, in constellation architectures, inter-satellite optical links. The core challenge is less "can we talk to the ground" and more whether a moving orbital network can deliver data-center-like fabric properties (high bisection bandwidth, predictable latency, and reliable routing) [3]. For many use cases, especially in-space analytics, the standard approach is to transmit *results* rather than raw datasets, reducing downlink requirements [8, 14].

3.6 Reliability, Radiation, and Servicing

Radiation-induced single-event effects and cumulative damage shape both hardware selection and software architecture. Current practice typically combines modest shielding, error correction, redundancy, and an acceptance of some component attrition rather than relying exclusively on fully radiation-hardened computing (which usually trails in performance) [3]. Outside crewed platforms, repairs and upgrades remain non-routine; in practice, reliability comes from autonomous fault management and fleet-level redundancy.

4 Companies and Initiatives

Table 1 lists selected initiatives and summarizes their publicly described positioning. The list is illustrative rather than exhaustive.

Initiative	Approach	Stage (publicly stated)
Starcloud	GPU-centric micro data-center satellites; roadmap toward commercial orbital compute services [7, 8]	Demo in orbit; follow-on planned
Lonestar Data Holdings	Lunar/cislunar storage and disaster-recovery-oriented "off-world data vault" payloads [12, 11]	Early demos; scaling roadmap
Aetherflux "Galactic Brain"	Solar-powered AI compute constellation; ties to space power-beaming concepts [9]	Announced
Cloud Constellation "SpaceBelt"	Encrypted orbital storage marketed as a high-security cloud repository [10]	In development
Unibap + D-Orbit "SpaceCloud"	In-orbit edge compute modules running apps for on-orbit analytics [14, 13]	Demonstrated in orbit
HPE Spaceborne (ISS)	General-purpose computing on the ISS; experiments and operational learnings [5, 6]	Operational / ongoing

Table 1: Selected initiatives in space-based data centers and adjacent on-orbit edge computing.

5 Space vs. Earth Data Centers: Engineering Trade Space

Orbital data centers swap familiar terrestrial constraints (land, water, permitting, and grid interconnect) for harsh-environment constraints (radiation, radiative-only heat rejection, launch mass, and limited physical access). Public analyses generally converge on the same bottom line: small systems look technically plausible, but at commodity scale the outcome is dominated by cost, operations, and networking limits [3, 15].

Dimension	Terrestrial data centers	Space-based data centers
Power	Grid or onsite generation; constrained by interconnect and local supply.	Solar in high-sunlight orbits; high upfront mass/capex; storage depends on orbit [1].
Cooling	Convection and liquid loops; mature designs; can be water intensive.	Radiation-only heat rejection; radiator area/mass dominates at scale [3].
Latency & bandwidth	Fiber-scale bandwidth and low latency; mature switching fabrics.	Uplink/downlink adds latency; optical links promising but operationally complex [3].
Maintenance	Human access, rapid repair, routine upgrades.	Repairs are difficult; autonomy and fleet redundancy are key [3].
Security	Strong cyber controls; physical access risk exists.	Physical isolation potential; new threats (jamming/ASAT) and limited incident response [10].

Table 2: High-level comparison of terrestrial versus orbital data center constraints.

6 An Economic "Cost-per-Watt" Lens

McCalip proposes a concrete comparison metric: the total delivered cost per usable watt of electrical power available for compute, along with the associated levelized cost of energy (LCOE), under explicit assumptions [1]. Under one representative baseline (1 GW target capacity, 5 year analysis period, Starlink-class specific power, and $1,000/kg launch to LEO), the public model yields the illustrative comparison in Table 3 and highlights the magnitude of mass and launch requirements [1].

Scenario (illustrative baseline)	Cost per Watt	LCOE
Orbital solar power for compute (LEO)	$31.2/W	$891/MWh
Terrestrial (on-site CCGT benchmark)	$14.8/W	$398/MWh

Table 3: Illustrative outputs from a configurable first-principles model comparing orbital versus terrestrial power economics for data center capacity. Values reflect one baseline case; sensitivity to launch cost, specific power, sunlight fraction, degradation, and replacement rates is material [1].

6.1 Sensitivity Drivers (Qualitative)

Without asserting additional numeric outputs beyond the published baseline, the dominant sensitivity levers are:

- **Launch cost ($/kg) and packaged specific power (W/kg).** These jointly determine the cost of delivered electrical power to compute.
- **Thermal design margins.** Radiator area and heat transport mass scale nonlinearly with allowable temperatures and radiator emissivity.
- **Replacement rate and degradation.** Radiation environment, component screening, and fault tolerance drive replenishment cost.

- **Network architecture.** Downlink capacity, spectrum constraints, and optical crosslink maturity impact which workloads are viable.

Two strategic takeaways are especially relevant:

1. **Vertical integration is decisive.** If launch, bus, and integration are procured at retail interfaces, the margin stack can prevent cost competitiveness; an actor controlling the full stack is best positioned to explore the frontier [1].

2. **Thermal and mass assumptions are first-order.** Simplifications in heat rejection and coolant loop modeling can change the mass budget and therefore launch cost materially [1, 3].

7 Priority Use Cases

Near-term applications usually share at least one of three properties: they generate data in space, they value physical isolation or extreme resilience, or they can tolerate uplink/downlink latency:

- **On-orbit analytics for Earth observation and communications.** Processing images and signals in orbit reduces downlink volume and can yield faster actionable results [14, 8].

- **High-assurance storage and key management.** "Air-gapped" storage in orbit is marketed as a premium security product for regulated or sovereign datasets [10].

- **Disaster recovery beyond Earth.** Lunar/cislunar storage proposals target long-horizon resilience against terrestrial disruptions [12, 11].

- **AI workloads (conditional).** Training at scale is power and cooling constrained on Earth; orbital concepts claim long-run advantages, but must close the cost-per-watt gap and address networking constraints [3, 1].

- **Blockchain and decentralized services.** Satellite nodes can improve continuity and censorship-resistance properties for some architectures; public demonstrations exist for blockchain payloads in space [16].

8 Risks and Constraints

The following risk categories repeatedly appear in technical and commercial analyses:

Risk area	Why it matters for SBDCs
Thermal scaling	Heat rejection mass/area can dominate at multi-kW scale; thermal interfaces are reliability-critical [3].
Radiation and lifetime	Drives error rates, shielding mass, and replacement cadence; influences total cost of ownership [3].
Debris and threat environment	Constellation operations must manage collision risk and potential adversarial interference.
Networking and spectrum	Downlink bottlenecks can invalidate workloads that require raw data transfer; optical crosslinks may be decisive [3].
Servicing and upgrade model	Lack of routine access shifts value toward autonomy, redundancy, and modular replacement.
Regulatory/export controls	Payload encryption, ground infrastructure, and cross-border operations face compliance constraints (ITU, national regulators).

Table 4: Non-exhaustive risk register for space-based data center deployments.

9 Outlook

The near-term trajectory is likely to be continued demonstrations and small commercial pilots. Higher-performance payloads (including GPUs), improving optical communications, and reusable launch economics sustain interest, while thermal and networking constraints remain fundamental [3]. Feasibility work (e.g., ASCEND) anticipates that meaningful deployment would require major advances in launch systems, in-orbit assembly/servicing, and lifecycle emissions reductions [15].

10 Conclusion

SBDCs have moved past slide decks and into flight demonstrations. The engineering is better understood than it was five years ago, but scaling beyond niche workloads will still require the numbers to close across launch economics, spacecraft specific power, thermal rejection, radiation-driven replacement, and networking. Near term, a hybrid approach is pragmatic: use orbit for in-space analytics and resilience-oriented storage where the advantages are real, and treat orbital commodity compute as an option that only scales once cost-per-watt milestones are met [1, 3].

References

[1] Andrew McCalip, "Economics of Orbital vs Terrestrial Data Centers," 2025. Available: https://andrewmccalip.com/space-datacenters. Accessed: 2025-12-16.

[2] Andrew McCalip, "andrewmccalip/thoughts: Space Datacenters: Orbital vs Terrestrial Economics," GitHub repository, 2025. Available: https://github.com/andrewmccalip/thoughts. Accessed: 2025-12-16.

[3] Space.com, "Data centers in space: Will 2027 really be the year AI goes to orbit?," 2024/2025. Available: https://www.space.com/technology/data-centers-in-space-will-2027-really-be-the-year-ai-goes-to-orbit. Accessed: 2025-12-16.

[4] NASA Spinoff, "Cutting-Edge Computing Goes Spaceborne," 2018. Available: https://spinoff.nasa.gov/Cutting-Edge_Computing_Goes_Spaceborne. Accessed: 2025-12-16.

[5] HPE, "HPE Spaceborne Computing," 2025. Available: https://www.hpe.com/emea_africa/en/compute/hpc/supercomputing/spaceborne.html. Accessed: 2025-12-16.

[6] HPE, "HPE Spaceborne Computer-2 returns to the International Space Station," Press Release, 30 Jan 2024. Available: https://www.hpe.com/us/en/newsroom/press-release/2024/01/hpe-spaceborne-computer-2-returns-to-the-international-space-station.html. Accessed: 2025-12-16.

[7] Data Center Dynamics, "Starcloud-1 satellite reaches space, with Nvidia H100 GPU now operating in orbit," 3 Nov 2025. Available: https://www.datacenterdynamics.com/en/news/starcloud-1-satellite-reaches-space-with-nvidia-h100-gpu-now-operating-in-orbit/. Accessed: 2025-12-16.

[8] Starcloud, "Starcloud-2: In-Space GPU Cluster & Cloud Computing Satellite," 2025. Available: https://www.starcloud.com/starcloud-2. Accessed: 2025-12-16.

[9] Space.com, "Startup announces 'Galactic Brain' project to put AI data centers in orbit," 2024/2025. Available: https://www.space.com/technology/startup-announces-galactic-brain-project-to-put-ai-data-centers-in-orbit. Accessed: 2025-12-16.

[10] Cloud Constellation, "SpaceBelt," 2025. Available: https://spacebelt.com/. Accessed: 2025-12-16.

[11] Data Center Frontier, "Lonestar Data Makes it to the Moon on IM-1 Lunar Lander," 2024. Available: https://www.datacenterfrontier.com/edge-computing/article/33037610/lonestar-data-makes-it-to-the-moon-on-im-1-lunar-lander. Accessed: 2025-12-16.

[12] Reuters, "Lonestar's moonshot: Firm aims to place data center on lunar surface," 21 Jan 2025. Available: https://www.reuters.com/technology/space/lonestars-moonshot-firm-aims-place-data-center-lunar-surface-2025-01-21/. Accessed: 2025-12-16.

[13] Unibap Space Solutions, "In-Orbit Demonstration & Verification," 2025. Available: https://unibap.com/?solutions=software-in-orbit-demonstration. Accessed: 2025-12-16.

[14] Unibap Space Solutions, "D-Orbit Wild Ride," 2022. Available: https://unibap.com/solutions/space-missions/wild-ride-mission/. Accessed: 2025-12-16.

[15] Thales Alenia Space, "Thales Alenia Space reveals results of ASCEND feasibility study on space data centers," 27 Jun 2024. Available: https://www.thalesaleniaspace.com/en/press-releases/thales-alenia-space-reveals-results-ascend-feasibility-study-space-data-centers-0. Accessed: 2025-12-16.

[16] ESA Space Solutions, "SpaceChain," 2020–2025. Available: https://business.esa.int/projects/spacechain. Accessed: 2025-12-16.