The Levelized Cost of Compute

A Framework for Comparing Orbital vs. Terrestrial Data Center Economics

Pranav Myana · December 2025 · x.com/princeofprawns0 · github.com/TheEnergyNerd/orbital-compute-finance

Abstract. Conventional metrics for compute infrastructure—$/GPU-hour, $/rack/month—fail to capture lifecycle costs and operational constraints. This paper introduces the Levelized Cost of Compute (LCOC), a framework adapted from energy economics enabling rigorous comparison across deployment topologies. Under baseline assumptions, orbital compute achieves cost parity with terrestrial infrastructure in the early 2030s.

1. The Input Problem

Data center economics reduce to four constrained inputs: power, cooling, land, and labor. Ground infrastructure faces escalating constraints; orbital deployment substitutes these for a different constraint set with fundamentally different cost trajectories.

Input	Terrestrial	Orbital
Power	Grid queue 3-5 yr, fuel volatility, permitting	Solar 1,361 W/m² continuous, zero fuel, no permits
Cooling	Water rights contested, PUE 1.2-1.4×	Radiative only, no consumables, PUE ≈ 1.0
Land	Zoning, environmental review, escalating costs	LEO unlimited; GEO 1,800 slots
Labor	On-site technicians, remote premium	Autonomous ops, ground control only

2. The LCOC Framework

The Levelized Cost of Compute adapts the LCOE methodology from power markets. LCOC expresses cost in $/delivered GPU-hour at 99.9% SLA, incorporating capital expenditure, operating expense, cost of capital, and utilization constraints.

3. Core Equations

Delivered GPU-hours per year:

$$\mathrm{DeliveredGPUh} = \mathrm{GPU}_{eq} \times 8760 \times \mathrm{SLA} \times \min(1, \mathrm{BW}_{avail}/\mathrm{BW}_{need}) \times \min(1, D/C)$$

Base LCOC:

$$\mathrm{LCOC}_{base} = (\mathrm{CAPEX} \times \mathrm{CRF} + \mathrm{OPEX}) / (\mathrm{GPU}_{eq} \times 8760 \times \mathrm{SLA})$$

Effective LCOC:

$$\mathrm{LCOC}_{eff} = \mathrm{LCOC}_{base} / u_{sell} \text{ where } u_{sell} = \min(\mathrm{BW}_{sell}, D_{sell})$$

Capital Recovery Factor:

$$\mathrm{CRF}(r, n) = r(1+r)^n / ((1+r)^n - 1)$$

Variables: GPU_{eq} = H100-equivalent capacity · SLA = 0.999 · r = WACC (12% orbital, 8% ground) · n = asset life (years) · BW = bandwidth · D/C = demand/capacity

4. Baseline Assumptions

Parameter	Terrestrial	Orbital
WACC	8%	12%
Asset life	5 years	8 years (rad-adjusted)
PUE	1.25 → 1.08	1.0
Energy cost	$0.045/kWh (+2%/yr)	$0 marginal
Launch cost	—	$1,500 → $10/kg
Thermal	Water-cooled	Radiative; droplet 2030+

5. Cost Trajectory Drivers

Ground costs flatten as favorable sites exhaust. Orbital costs decline on multiple fronts:

- **Launch:** $60k/kg (Shuttle) → $1.5k/kg (Falcon 9) → $10-50/kg (Starship)
- **Thermal:** Droplet radiators enable MW-class platforms (vs 250 kW conventional)
- **Power:** Space fission (2035+) enables 5-20 MW; fusion (2045+) enables GW-class
- **Manufacturing:** Volume production drives 30%+ learning rate

6. Key Risks

Thermal. Radiative cooling only. Droplet radiators promising but unproven at scale.

Radiation. SEU requires redundancy (shielding ineffective). TID limits lifetime. MEO harshest.

Bandwidth. Downlink constrains sellable compute. ~100 Tbps global, +35%/yr.

Latency. LEO 4-20ms (competitive). GEO 120ms+. Shell selection application-dependent.

Servicing. On-orbit maintenance requires autonomous vehicles. Costs declining.

7. Conclusion

Under baseline assumptions, orbital compute achieves cost parity in the early 2030s, driven by: (1) launch costs approaching $10/kg, (2) thermal solutions enabling MW-class platforms, (3) tightening ground constraints, and (4) scarcity premiums from demand growth.

Open Source — Model code: github.com/TheEnergyNerd/orbital-compute-finance
Interactive Simulator — astrocompute.dev