

The Orbital Data Hall: Power, Cooling and Servicing Economics

A Decision Framework for Financing Space-Based Compute Capacity

Authored by Chennakeshav Adya

Independent Researcher

October 2026

ABSTRACT

An orbital data hall places computing, storage or data-processing equipment on a spacecraft and connects it to customers through an orbital or ground communications system. Public demonstrations have established limited forms of onboard processing, while larger solar-powered compute constellations remain proposed architectures. The investment question is therefore narrower than whether computing in space is physically conceivable. A sponsor must show that usable compute capacity can be powered, cooled, connected, maintained and sold over a defined life at a price that supports the complete system cost. This paper proposes an integrated diligence framework and a stage-gated financing sequence for that decision.

The analysis draws on Google Research's Project Suncatcher proposal, Axiom Space's reported demonstrations, European space-data-centre work, NASA thermal and radiation guidance, and international satellite-regulation material available through 5 October 2026. These sources describe different use cases and evidence states. None of them establishes a general commercial valuation multiple for an orbital data hall. The paper therefore separates an edge-processing node from a distributed AI-compute constellation, and separates engineering proof from contracted revenue. It traces the dependencies among solar generation, battery and power conversion, waste-heat rejection, communications capacity, radiation tolerance, launch, replacement and end-of-life operations.

A deliberately hypothetical 100 kW compute-payload example shows how capital cost, service availability and contract term affect the annual revenue needed to recover investment. Its figures are teaching assumptions, not estimates for an actual operator or a market benchmark. The paper's principal contribution is a decision system: each financing phase should purchase evidence that changes a particular commercial uncertainty. Equity can fund experiments with bounded losses. Infrastructure-style debt becomes a candidate only after integrated performance, lawful operations, deliverable customer service and contractual risk allocation are independently reviewable. The proposed framework also identifies where an adviser can translate engineering results into a defensible valuation and capital plan.

JEL Classification: G31, G32, L86, O33, Q42

Keywords: orbital data centre, space-based computing, thermal control, satellite servicing, optical communications, infrastructure finance, valuation, data centres

1 Define the asset before valuing it

The expression orbital data hall can describe substantially different assets. A processor aboard a space station may run a small workload close to the data source. A free-flying node can store or pre-process observations from other spacecraft. A tightly connected constellation could, in principle, distribute machine-learning work across many accelerators. The energy requirement, thermal envelope, communications burden and customer contract differ across those designs. An investment memorandum should name the actual service, physical architecture, location and contractual seller before applying a data-centre valuation label.

The first decision is where value is created. If an Earth-observation operator pays to filter raw images before downlink, the relevant benefit may be reduced transmission volume or faster information delivery. If a customer buys general-purpose compute, the sponsor must show that its network can move training or inference data and results at the required speed. A spacecraft that can perform computation is not automatically a competitive cloud facility. A terrestrial data hall's familiar metrics, including powered space and contracted capacity, become meaningful only after the orbital asset can offer an equivalent measurable service.

This paper proposes a full-system investment perimeter. It includes launch, spacecraft bus, solar collection, power management, batteries, compute hardware, storage, thermal control, inter-satellite and ground links, station-keeping, collision avoidance, replacement, insurance where available, and disposal. Ground gateways, software, customer integration and regulatory permissions belong in the same cost and dependency register. Omitting any of these can make a subsystem appear economic while the service cannot be delivered. The perimeter is an author-developed analytical proposal, not a representation that any particular project has obtained the required rights or contracts.

2 Read the present evidence precisely

Google Research described Project Suncatcher in November 2025 as an exploration of solar-powered satellites carrying TPUs and linked by free-space optics [1]. Its publication discusses high-bandwidth links, close formation, radiation testing and a launch-cost scenario. The research also identifies thermal management, high-bandwidth ground communications and on-orbit reliability as remaining engineering challenges. Google said it planned two prototype satellites with Planet for early 2027. The proposal is relevant evidence of serious architectural work; it is not a published commercial operating record for a large orbital compute utility.

Axiom Space reports an orbital data-processing prototype aboard the International Space Station in 2025 and the launch of two dedicated nodes to low-Earth orbit in January 2026 [2]. Its material describes prospective cloud, AI, data-fusion and cybersecurity applications. The public description supports the existence of demonstrator activity as reported by the operator. It does not disclose a complete, independently audited revenue, cost, availability or return series. ESA has separately discussed possible space-based data-centre use cases, including observation-data handling [3]. A European Commission grant description of ASCEND sets out a concept involving modular orbital data centres, solar power and optical communications [4]. A concept, a prototype and a revenue-producing platform are different evidence classes.

Investors should request a dated claim-to-evidence register. For each claimed output, record the configuration tested, environment, duration, instrumentation, independent witness, anomalies and changes since the test. A laboratory optical-link rate should not be inserted into a commercial service model without attenuation, pointing and network-availability assumptions. A short processor demonstration should not be converted into a five-year life estimate. The development plan should specify which test will retire each major uncertainty and what result would cause redesign. This disciplined reading protects capital allocation from the understandable enthusiasm surrounding a visible technical milestone.

3 Choose the commercial workload

An orbital edge node can create value without competing with a terrestrial hyperscale data centre. A satellite constellation may generate sensor data in orbit and need rapid event detection before deciding what to transmit to Earth. The customer might buy processing per observation, stored terabyte-months, reserved compute capacity, or a response-time service. Each price unit creates a different revenue meter. The sponsor should state who owns the source data, what portion can legally be processed or retained in orbit, and whether the output can be independently verified by the customer.

General-purpose AI training creates a different systems problem. Large training jobs move parameters, activations and data across accelerators, and the network becomes part of the compute product. Google Research says its proposed architecture would require inter-satellite links on the order of tens of terabits per second to approach data-centre-scale distributed workloads [1]. That statement describes its intended design requirement. A finance case must translate it into the particular workload's throughput, latency, resilience and ground-ingress needs. A high theoretical optical-link capacity has little commercial value if the contracted job stalls during handover or cannot ingest training data economically.

The market case should use named customer processes rather than a broad total-addressable-market slide. Obtain a technical statement of work that states inputs, outputs, data volume, latency tolerance, security and failure remedies. Compare the proposed orbital service with terrestrial processing at the customer's actual decision point. If a customer values data locality, resilience or reduced downlink, the contract should purchase that attribute explicitly. If the advantage cannot be translated into a measurable service or price, a capital committee should treat it as a hypothesis for customer discovery, not as contracted revenue.

4 Build the physical power balance

Solar arrays generate electricity only under the geometry, degradation and operating assumptions of a specific orbit and design. Power electronics, communications, attitude control, thermal pumps, batteries and compute payload all draw from the same bus. A power budget should show peak and average loads by operating state, including safe mode, eclipse, link acquisition and compute bursts. It should identify battery charge and discharge losses, cell degradation, array orientation and contingency. A quoted panel output under one illumination condition cannot stand in for continuous power delivered to accelerators.

Google's Suncatcher proposal discusses a dawn-dusk sun-synchronous low-Earth orbit to increase sunlight exposure and reduce battery requirements [1]. The orbit is a design choice with operational consequences. It affects thermal exposure, ground contact, collision environment, station-keeping and mission approval. Another project's best orbit may be different. The engineer should supply an hourly energy balance for its own mission profile, not inherit the preferred geometry of a published concept. The financial model should use usable compute power after all non-compute loads and downtime.

The power budget should also identify what happens after a partial failure. A reduced array output might be managed through compute throttling; it might also make a latency contract impossible. Batteries could sustain some operations through eclipse while compute is curtailed. The sponsor must state whether the customer pays for reserved hardware, available compute power or completed work. That contractual choice determines the revenue loss from the same physical event. Finance should use an engineering-controlled source file with versioned assumptions and reconcile it to each term-sheet promise.


5 Close the heat balance before sizing revenue

Almost all electricity consumed by computation becomes heat that must leave the spacecraft. NASA's thermal-control guidance explains that a vacuum has no convective heat transfer to the external environment; external heat rejection relies on radiation, while internal heat can be conducted or moved through active systems [5]. The radiator therefore belongs in the first-order architecture, mass budget and capital estimate. It is not a minor auxiliary item that can be added after the compute payload has been sized. Panel orientation, surface properties, sunlight, Earth infrared exposure and permitted component temperatures affect its performance.

A thermal model should state heat load by mode, cold-plate or chip temperature limits, radiator effective area and temperature, fluid-loop assumptions, pump power, redundancy and degradation. Raising radiator temperature may reduce required area under a simplified radiative relation, but electronics and fluid limits constrain the choice. Larger radiators add mass, deployment complexity and potential collision exposure. Thermal capacity can become the bottleneck even when solar generation appears abundant. The relevant commercial capacity is the amount of compute that can operate continuously within qualified temperature limits, not the nameplate sum of installed chips.

Qualification should cover transients and failures. A high-utilisation AI workload may create a different heat profile from intermittent edge processing. A failed pump, damaged radiator panel, degraded coating or unfavourable spacecraft attitude may require throttling before a complete shutdown. A sponsor should show the control sequence and customer-service consequence for each case. NASA's overview describes multiple passive and active techniques [5]; it does not certify a particular orbital data-hall design. Independent thermal-vacuum tests and a mission-specific operating record should progressively replace modelled heat-rejection assumptions.

Figure 1. Power and heat interfaces in an orbital compute service



Contracted availability requires every interface to operate together.

Author-developed conceptual diagram; each arrow is an engineering and contractual measurement point, not a tested system.

6 Treat links as part of the product

Space-to-ground and inter-satellite communications determine whether compute is usable. The network model should distinguish raw capacity, effective throughput, latency, duty cycle, packet loss, weather sensitivity and handover. It should also identify the path for software updates, telemetry, encryption keys and recovery commands. A service with abundant onboard processing but intermittent or costly data ingress may suit local image filtering and fail as a general cloud platform. The customer's workload map should determine the network architecture and gateway plan.

Optical links can offer high bandwidth, while pointing accuracy, cloud cover for some ground links and route availability need project-specific testing. Radio-frequency links introduce spectrum coordination and interference management. The International Telecommunication Union explains that national administrations license satellite systems within an international regulatory framework and coordinate relevant frequencies [6]. A sponsor should maintain a route and rights matrix for each gateway, uplink, downlink and crosslink. A filing or research paper is not a substitute for an operating authorisation.

Network economics should count terrestrial gateway sites, backhaul, optical terminals, redundancy, spectrum work and customer integration. If a constellation needs close formation to support very high inter-satellite rates, formation control is also a communications cost. Google Research's published close-formation analysis is an architecture proposal and modelling result [1]. The investment model must price the propulsion, navigation, operations and collision-avoidance consequences of its own design. Customer service-level agreements should describe what happens when the orbital compute is healthy but a gateway is unavailable.

7 Protect useful life against radiation and obsolescence

Commercial processors and memory face a radiation environment for which they were generally not designed. NASA describes total-dose and single-event effects that can disturb or damage spacecraft electronics [7]. Google Research reported proton-beam testing of a TPU and identified high-bandwidth memory as a sensitive subsystem in its published work [1]. These are specific tests under stated conditions. They do not establish the lifetime, fault rate or reliability of every accelerator package in a different orbit, shield or workload.

A technical diligence file should identify each critical device, test dose, shielding assumption, expected mission dose, error-correction strategy, spare capacity and failure response. Component failures affect available capacity and customer credits. A model should distinguish a reversible compute reset from permanent loss of a node. It should also account for solar-array degradation, battery ageing, radiator performance and optical-terminal life. Reliability cannot be represented by a single annual percentage chosen to make a financial model work. It is the result of specific failure modes, redundancy and repair or replacement rights.

Obsolescence may arrive before physical failure. A new terrestrial accelerator generation can improve price-performance while an orbital node remains costly to replace. The sponsor should disclose the hardware refresh assumption and the customer's right to demand equivalent performance over a contract term. A modular replacement strategy can preserve value if interfaces, launch slots and servicing capability are real and affordable. A long terminal value based on indefinite reuse of a first-generation payload is difficult to support without a demonstrated refresh path. Technology life and spacecraft life belong in separate valuation schedules.

8 Make servicing a priced option

Servicing can mean software maintenance, remote reconfiguration, replacement of a complete satellite, delivery of a new compute module, refuelling or robotic repair. These activities have different technical and contractual requirements. The most credible baseline may be planned satellite replacement rather than physical repair of a deployed module. A sponsor should specify which servicing action it can actually perform, what hardware interface permits it, how long it takes, and how much lost revenue occurs during the operation. A generic promise of on-orbit maintenance should carry no cost saving until tested.

NASA's OSAM-1 experience is a caution in interpreting servicing roadmaps. NASA announced that the project would be discontinued after technical, cost and schedule challenges and the absence of a committed transition partner [8]. That history does not imply all servicing models are infeasible. It does show why a transaction should separate an established service contract from a hoped-for future capability. The investment case should use a replacement-at-launch baseline and add any servicing benefit as a conditional upside until a provider, interface, mission window and price are supported by evidence.

Serviceability also affects control rights. Who owns the docking interface, software keys and replacement module design? Can a financier step in or appoint a substitute operator? Does a third-party service contract survive insolvency or change of control? The transaction documents should address ownership, technical data access, safety responsibilities and customer continuity. The expected value of servicing depends on these rights as much as on robot performance. A useful diligence exercise prices three cases: no service, scheduled replacement and demonstrated on-orbit intervention.

9 Identify lawful operations and disposal obligations

An orbital compute business requires a licensing route for its spacecraft and communications, and may require additional approvals for gateways, remote sensing, encrypted data or customer-specific activities. The applicable route depends on operator, jurisdiction, orbit, payload and service. ITU coordination addresses one part of a wider national framework [6]. Counsel should verify the actual authorisations, insurance and liability obligations before contracts assume a launch or service date. Cross-border data terms should be assessed against the real location and control of data processing, not a simplified claim that orbit is outside legal systems.

Orbital debris obligations can change design and cost. The US Federal Communications Commission adopted a five-year post-mission disposal rule for satellites subject to its jurisdiction in the relevant low-Earth-orbit region [9]. Other regulators and missions may apply different requirements. A disposal plan needs propulsion or passive-decay analysis, collision-avoidance operations, end-of-life reserves and a response to partial failure. A large compute constellation has more physical objects and replacement cycles than a small experiment. The capital model should include these obligations and the cost of complying with conditions in the issued licence.

Regulatory uncertainty should be modelled as a milestone and possible delay, not hidden in a general contingency line. A gateway site may be ready while frequency coordination is still pending. A satellite may be built before launch permission is complete. A customer may demand data-sovereignty controls whose application to the proposed workflow is unresolved. The project company should maintain a rights register with responsible counsel, status, expected decision, dependency and cost. Construction funding can be staged to the actual permissions obtained.

10 Connect procurement to delivered capacity

An equipment budget starts with the compute payload and expands to the complete bill of materials and services. Space-qualified packaging, shielding, thermal interfaces, deployment mechanisms, navigation, communications and software integration can change the economics of a nominally inexpensive processor. Launch cost depends on wet mass, packing volume, orbit, mission assurance, replacement cadence and the contract's allocation of delay and loss. Ground facilities and customer integration are part of delivering revenue. Finance should reconcile supplier scopes so that no interface is omitted or counted twice.

Google Research modelled a future launch-price path under which space-based data-centre operating and launch costs could become comparable to reported terrestrial energy costs on a per-kilowatt-year basis [1]. The conditional comparison depends on substantial price decline and a particular system design. It should not be treated as a 2026 launch quote, nor as proof that orbital compute services are cost competitive. A sponsor should obtain current binding or indicative terms for its actual mass and mission, then run a flat-launch-price case and a delayed-price-decline case. The customer's willingness to pay also depends on compute and communications quality, not only power cost.

Procurement diligence should identify long-lead components, manufacturing yields, test facilities and reflight options. A fixed-price satellite contract may still exclude a launch failure or design change after qualification. A lower unit cost at high production volume is useful only if the operator can finance and deploy that volume. The model should distinguish prototype cost, first-of-a-kind production cost and repeat-unit cost. Each reduction needs a supplier, learning mechanism and independent test. Without that bridge, the model may combine an expensive prototype's performance claim with a mass-production cost that has never been achieved.

11 Separate capacity economics from hardware counts

Installed accelerator count is a poor revenue denominator if power, cooling or communications constrain simultaneous use. The service model should calculate contracted useful capacity by hour after thermal throttling, power draw, network availability, radiation resets, software maintenance and reserve margin. Workload performance should be benchmarked under the same conditions used in the customer contract. A GPU-hour or accelerator-hour label has limited meaning across hardware generations unless throughput, memory and service quality are specified. The finance team should adopt one definition of billable output and use it consistently in contracts and models.

The operator may sell reserved capacity, completed processing jobs, storage, data reduction or availability. Reserved capacity can improve cash-flow predictability but creates penalties when promised service is unavailable. Per-job pricing shifts utilisation risk to the operator. A minimum-revenue commitment matters only if the counterparty is creditworthy and termination terms preserve the contracted economics. A market study should test actual procurement budgets and migration costs. Expression of interest, research collaboration and paid service contract should be recorded in separate categories.

The result is a contribution-margin bridge: customer invoices less gateway and network cost, replacement allowance, customer support, orbital operations and other service-variable expense. Fixed overhead and capital recovery follow. This bridge can be compared with terrestrial alternatives only for the same workload and service standard. A premium for resilience or data locality should be tied to an explicit customer benefit and procurement process. There is no universal orbital-data-hall multiple or verified market price in the sources reviewed here; project-specific contracts and costs determine value.

12 Test one transparent hypothetical unit

Consider a fictional orbital compute service with 100 kW of installed compute-payload power. For teaching purposes, assume that the payload meets its contracted service definition for 80 percent of the year's 8,760 hours. Assume initial all-in capital expenditure of USD 60 million, fixed annual operating and replacement allowance of USD 5 million, a five-year service life, a 12 percent nominal required return, no taxes, no inflation, no residual value and no construction delay. None of these figures is a market quote, a disclosed company result or a forecast. The example intentionally compresses many real engineering and contractual uncertainties into a small calculation whose assumptions are visible.

The annual billable capacity denominator is 100 kW times 8,760 hours times 80 percent, or 700,800 kW-hours of available installed compute-power capacity. This is a service-capacity accounting unit, not a measure of actual electrical consumption or computational throughput. A five-year annuity factor at 12 percent is approximately 0.27741. Annual capital recovery is about USD 16.64 million; adding the assumed USD 5 million fixed allowance gives about USD 21.64 million. Dividing by the capacity denominator produces a screening revenue requirement of approximately USD 30.89 per available kW-hour of installed compute power. Actual commercial pricing would require a workload benchmark and a contractually meaningful service unit.

The result has no market-comparison value without a matching terrestrial service. It omits customer acquisition, tax, working capital, insurance gaps, financing fees, launch failure, schedule slips, hardware refresh, data-egress charges and many other costs. It assumes that all available capacity can be billed, which is unlikely without a firm take-or-pay contract. Its purpose is to expose the influence of capex, life and uptime on the revenue hurdle. A prospective sponsor should replace every teaching input with a dated engineering estimate, supplier quote or executed customer term and retain the downside version used for an investment decision.

Table 1. Hypothetical annual revenue hurdle per available compute-power capacity-hour

Initial capital	60% availability	80% availability	90% availability
USD 40 million	USD 30.62	USD 22.97	USD 20.42
USD 60 million	USD 41.18	USD 30.89	USD 27.45
USD 80 million	USD 51.74	USD 38.80	USD 34.49

USD per available kW-hour under a five-year life, 12 percent required return and USD 5 million annual fixed allowance. Every input is author-selected for illustration; no actual project is valued.

13 Keep the three critical sensitivities connected

Capital cost, service availability and commercial life are linked. A heavier radiator can raise launch expense while improving sustained throughput. Additional redundancy can raise capex while protecting contracted availability. A shorter life may permit a cheaper component but increase annual capital recovery and refresh frequency. A responsible sensitivity analysis should model those design choices together. It should not use the lowest capex from one architecture, highest uptime from another and longest life from a third. Configuration control is a financial control.

The hypothetical table shows that at USD 60 million initial capital, moving from 80 percent to 60 percent service availability raises the screening hurdle from USD 30.89 to USD 41.18 per available capacity-hour. That arithmetic is mechanical, not an estimate of an achievable service price. The model must also consider whether the customer will pay for reduced output at all. An availability guarantee with service credits can make the cash-flow effect worse than the simple denominator suggests. A lost ground gateway may affect several satellites simultaneously, so correlated outages deserve explicit treatment.

Capital structure can amplify the sensitivity. Debt-service dates do not automatically move when a launch is delayed or a radiator test fails. Sponsor equity may be called before revenue begins; a customer may terminate if the service deadline is missed. The full model should therefore include monthly construction cash flows, launch tranches, commissioning tests, reserve accounts and covenant headroom. A single annual ratio is suitable for screening only. The investment committee should receive a scenario tree showing which technical event produces each financial consequence and who contractually bears it.

14 Fund evidence in sequence

The earliest investment should purchase knowledge, not the appearance of a fully financed utility. A bounded laboratory and thermal-vacuum programme can test power electronics, radiator performance and accelerator faults. An orbital demonstrator can test the integrated bus, communications and processor under actual mission conditions. A paid pilot can test customer work, metering and service support. Each stage needs predefined pass conditions and a capped capital exposure. If a test fails, the sponsor should preserve data, identify the redesign and reprice the next stage before asking for more capital.

Equity and strategic funding can accommodate experiments whose financial return depends on technical learning. Non-recourse infrastructure debt generally requires more predictable contracted cash flows, operating evidence and enforceable rights. This is a financing principle and proposed sequencing, not a statement about any lender's actual appetite. A customer prepayment may help finance deployment if performance obligations and refund protections are negotiated. Government research funding may support a demonstration with public objectives. Each source has a different claim on assets, intellectual property or future revenue; documents should make those claims compatible.

The funding gate should specify an evidence product. For example, a radiator gate could require independently witnessed heat rejection at the sustained compute load and worst-case thermal attitude. A network gate could require measured end-to-end throughput and service availability under representative handovers. A commercial gate could require a paying customer's signed workload specification and observed service metrics. A finance gate could require completed permissions, supplier contracts, a priced replacement plan and customer credit support. The evidence is useful even when the conclusion is to slow or stop deployment.

Figure 2. Proposed financing gates for an orbital data hall

EVIDENCE	POTENTIAL CAPITAL
Component test Heat, power and radiation measurements	Research and strategic capital
Orbital demonstrator Integrated in-orbit operating record	Venture and strategic equity
Paid pilot Customer workload and service meter	Development equity and prepayment
Commercial deployment Permissions, contracts, cost and uptime	Project equity; debt subject to diligence
Repeat system Audited operating and replacement data	Replicable infrastructure capital

Author-developed stage-gate framework. Funding categories are illustrative and do not imply committed financing.

15 Value the right entity and cash flow

An orbital-compute group may contain a technology developer, a spacecraft owner, a network operator, a ground-gateway company and a customer-facing service business. Their cash flows, contracts and risk differ. A sponsor should identify which entity owns the hardware, intellectual property, operating licences, spectrum access, customer data and receivables. The proposed financing vehicle should have enough rights to deliver its promises and give investors meaningful remedies. A project company's discounted cash flow cannot justify a parent-company equity value if the parent has not assigned relevant rights or if customer contracts reside elsewhere.

The valuation should begin with explicit scenarios. A research-stage company may be valued from funded milestones, intellectual-property rights, team capability and potential future markets, with substantial uncertainty and dilution. A paid pilot can add evidence on revenue and service cost. A fleet with repeat contracts can support a cash-flow model if its replacement economics are visible. Comparable-company multiples need similar margins, capital intensity, growth, customer risk and technological maturity. A terrestrial data-centre multiple applied to an orbital prototype without adjustments has no analytical foundation.

Terminal value requires special restraint. Orbital hardware has finite life, and a satellite may need replacement before a customer contract expires. A residual value should reflect feasible transfer, reuse or disposal proceeds after liabilities. A perpetual-growth calculation based on the first generation's gross revenue can conceal refresh capex and licence renewal. The model should show annual new-launch and replacement cohorts, not one undifferentiated opening balance. Investors can then evaluate whether growth consumes more capital than the service generates and whether later cohorts really improve on the first.

16 Structure contracts around measurable failure

Customer contracts should specify the processing service, accepted input, output, latency, security, availability and method for measuring each. They should describe service credits, exclusions, data loss, disaster recovery and termination. An orbital operator may face launch risk, link outages and hardware degradation that are unfamiliar to a customer procurement team. Those risks should be disclosed and allocated in the contract. A price premium for resilience cannot coexist with broad, unpriced disclaimers of service continuity.

Supplier agreements should map onto the customer obligation. If a promised service begins on a fixed date, launch, spacecraft integration, gateway and software contracts need schedules that leave realistic testing margin. A launch provider may limit its liability to transport terms. A processor supplier may warrant only its component. The project company can be left holding the gap. A risk-allocation matrix should identify every interface and the party with practical control. Where no contractor can take a risk, sponsors should budget equity, reserves or customer-side flexibility rather than pretend the risk has disappeared.

Financing documents should respect the same map. Security over a spacecraft, receivables or software has different practical value if regulatory consent or technical keys are needed to exercise it. Direct agreements can provide notice, cure and step-in rights where lawful. The sponsor should price insurance based on actual availability and exclusions, not assume a terrestrial policy can be imported. Specialist legal, engineering and insurance review remains necessary. The table below is an author-developed diligence request, not a description of any existing transaction.

Table 2. Contract and risk interface register

Interface	Evidence request	Economic exposure if unresolved
Launch to commissioning	Manifest, accepted orbit and reflight terms	Delayed revenue and replacement capital
Solar bus to compute	Signed power budget and throttling policy	Lower billable capacity
Compute to radiator	Sustained thermal-vacuum performance	Service credits and hardware damage
Orbital link to ground	Gateway rights and measured throughput	Unusable output or higher backhaul cost
Servicing to customer	Replacement schedule and continuity plan	Contract interruption and refresh capex
Operator to regulator	Licences, spectrum route and disposal plan	Delay, restrictions and end-of-life expense

Proposed diligence questions for a specific project; answers require executed documents and specialist review.

17 Use data tools for traceable decisions

An orbital-compute programme generates engineering, commercial and regulatory records that change at different speeds. A model should connect each thermal test to its hardware revision, each mass budget to a launch quote, each network test to the contracted workload and each permission to the exact spacecraft and gateway. Versioned records allow a committee to see whether a cost estimate refers to the current design. Automated document extraction and anomaly detection can help locate discrepancies when source documents remain available for human review. This is a proposed information architecture, not a claim of measured savings from artificial intelligence.

Data governance is especially important where customers supply sensitive imagery, research or proprietary models. The sponsor should define who can access data in orbit and on the ground, where encryption keys are held, how information is deleted, and what happens when a module is retired. Cross-border legal requirements and export controls require project-specific advice. A marketing claim of sovereignty or security should be converted into a technical architecture, legal opinion and auditable control list. Investors should see the cost of those controls in the operating budget.

The finance model itself should be reproducible. Preserve unit conventions, currency, price date, source, confidence range, scenario owner and approval for every material input. An AI-generated summary cannot establish a radiator's performance, a licence's scope or a customer's binding commitment. Decisions should be traceable to the original measurement or document and the qualified review of that evidence. This standard is commercially valuable because a design change can be evaluated quickly without rebuilding the entire investment story from presentations.

18 Test adjacent transaction opportunities

The first bankable transaction may be a narrow data-processing contract rather than a large orbital campus. A satellite operator with expensive downlink capacity may fund a pilot that processes observations in orbit. A constellation operator may buy a specialised node or lease onboard processing under a defined service level. A communications provider may contribute gateway capacity in return for minimum traffic. A hardware manufacturer may invest strategically to validate a radiation-tolerant product. These are proposed transaction forms; this paper identifies no committed buyer, investor or mandate.

M&A diligence should follow the same physical and contractual boundaries. A buyer of an orbital-compute start-up should ask whether its patents are transferable, whether supplier licences survive a change of control, whether tests relate to the current hardware, and whether customer pilots are paid and renewable. It should inspect cyber controls, spectrum dependencies, space-object ownership, launch commitments and unfunded disposal liabilities. Earn-outs can tie price to technical or revenue milestones, but milestone definitions need

independent measurement and protection against design or budget manipulation. A headline order book without enforceable service delivery rights is a weak basis for consideration.

Joint ventures can align a space operator, processor developer and anchor customer if governance is explicit. The parties should agree which one controls design, funds overruns, owns operating data and resolves conflicts between safety and service level. A customer may seek exclusivity that restricts the market needed to finance the fleet. A supplier may seek ownership of improvements that reduces the project's replacement options. A transaction adviser can model these trade-offs and establish decision rights before a large capital commitment. An attractive technical collaboration becomes an investable enterprise through specified economics and rights.


19 Define stop conditions and downside rights

A strong proposal states when the committee should not proceed. Possible stop conditions include sustained heat rejection below the required operating load, links that cannot meet the contracted workload, unmanageable radiation-induced fault rates, prohibitive replacement cost, unavailable gateway rights, or customer price below the fully loaded revenue hurdle. Each condition needs a quantitative project-specific threshold and a measurement method. This paper cannot set universal numerical thresholds because architectures and customers differ. The sponsor should set them before receiving test results.

The downside plan should describe recoverable assets and obligations. Intellectual property and test data may retain value after a demonstrator fails. Customer prepayments may need to be returned. Spacecraft and hardware may have little resale value once configured for a particular orbit or operator. End-of-life and data-deletion duties can continue after revenue stops. A model that simply assigns zero terminal value may still understate liabilities. Financiers should know who funds disposal, warranty claims and contract termination if the business ceases operations.

Scenario probabilities should be supported by actual engineering evidence and commercial commitments. A committee can approve a small test with a low probability of broad commercial success if the information gained justifies its capped cost. It should not fund a production constellation using a prototype's best-case performance as the base case. The stage-gate map below shows a proposed sequence of decisions. Each gate can result in advancement, redesign, sale of technology, a narrower customer application or closure. These are choices, not a promise that the orbital data-hall thesis will succeed.

Figure 3. Proposed evidence path from orbital prototype to financeable service



Capital increases only when the next uncertainty is retired.

Author-developed decision path; passing one gate does not certify later performance or financing.

20 Commission a decision-ready advisory mandate

An initial advisory mandate can be bounded around a specific investment choice. The proposed deliverables are a system and rights perimeter, a customer-workload and price test, a reconciled power and heat model, a launch and replacement schedule, a source-linked capex estimate, and a stage-gated funding plan. Engineering specialists should certify technical questions and counsel should advise on law and licences. Corporate-finance work should translate their evidence into capital requirements, valuation sensitivities, transaction structure and counterparty strategy. The engagement scope, fee and responsibilities require agreement with an actual client.

The first phase might determine whether to fund a defined orbital demonstration. If the evidence supports it, a second phase could structure customer pilots and strategic equity. A later financing mandate could run a competitive process for commercial deployment, subject to genuine customer contracts and project readiness. These phases need separate go/no-go decisions. A project with a narrow but valuable edge-processing application should not be forced into an enormous general-purpose AI valuation story. Conversely, a compelling constellation design should not be financed on a small edge pilot's revenue without evidence of scale.

Deliverables should be audit-ready. The model must reconcile compute capacity, power draw, waste heat, radiator area, communications throughput, orbit, launch mass, replacement rate and customer service in one controlled configuration. The diligence file should identify verified measurements, supplier estimates, management plans and author assumptions in separate fields. The transaction plan should state what decision each funding round buys and which risks remain unallocated. This process can produce a useful negative decision as well as a mandate to proceed. It offers no promise of a particular client, financing outcome or fee.

Table 3. Proposed investment-committee decision file

Decision	Evidence required	If evidence is absent
Fund thermal and power test	Controlled design, test protocol, capped cost	Rework the test or defer funding
Launch a demonstrator	Flight-qualified interfaces, rights, fault plan	Retain ground-based testing
Sign a paid pilot	Workload, service meter, price and remedy	Conduct customer discovery
Commit deployment equity	Priced system, customer contract, permissions	Stage or reprice the commitment
Consider project debt	Operating record, cash flow, security and reserves	Use appropriate risk capital
Replicate the fleet	Cohort cost and reliability evidence	Redesign or limit scale

Author-developed checklist; a real transaction requires project-specific evidence and professional review.

21 Conclusion

Orbital computation now includes reported prototypes and serious research proposals. Those activities justify rigorous commercial analysis, while they leave a wide gap to a repeatable infrastructure service. An investment decision should start with the actual customer workload, then reconcile the physical energy and heat balances, communications path, hardware life, launch and replacement plan, permissions and contract. Any single impressive subsystem metric becomes valuable only when it contributes to billable, delivered service. The public sources reviewed here do not establish a general orbital data-centre valuation benchmark.

The hypothetical example demonstrates why architecture and finance must be modelled together. Under its fictional USD 60 million capital cost, USD 5 million annual allowance, five-year life, 12 percent return and 80 percent availability, it requires roughly USD 21.64 million annual revenue. That number is not a quoted project cost, attainable tariff or investment recommendation. It is a transparent screening hurdle that a sponsor can replace with verified inputs. Lower availability, shorter hardware life or more frequent replacement can materially change the answer even when an accelerator's peak performance improves.

The proposed way forward is staged evidence. Use early capital to measure thermal, power, link and fault performance. Pay for a customer pilot that tests an actual service and contract. Seek scaled capital only when the full operating perimeter is costed, lawful and independently reviewable. Commercial structure, M&A terms and debt capacity then follow the evidence. A disciplined process may find a narrow, profitable orbital service, a technology asset worth acquiring, or a reason to wait. Each result is more useful than a valuation built on an untested combination of optimistic assumptions.

Frequently Asked Questions

Is an orbital data hall already a proven hyperscale-data-centre asset class?

The cited material describes prototypes, use-case studies and proposed larger architectures. It does not establish a general, independently audited operating or valuation benchmark for hyperscale orbital compute. A specific project needs its own performance and contract evidence.

Why is cooling a central economic issue in space?

Computation produces heat, and a spacecraft must reject that heat within qualified operating limits. In vacuum, external rejection relies on radiation rather than air convection. Radiator size, mass, power use and failure risk therefore affect launch cost and billable compute capacity.

Does nearly continuous sunlight make the power free?

Solar collection still requires arrays, power conversion, storage, deployment, replacement and operation. The usable power depends on mission geometry, degradation and non-compute loads. A project must show delivered payload power over its actual orbit and service schedule.

Can the Google Research launch-cost scenario be used as a current price?

It is a conditional future scenario in Google's published analysis. It is not a binding 2026 launch quote. A project should obtain terms for its actual mass, orbit, timing and replacement needs, then test flat and adverse prices.

Is the USD 30.89 figure a market price for orbital compute?

No. It is the hypothetical screening revenue requirement per available installed compute-power capacity-hour under the paper's stated assumptions. It is not a GPU-hour price, a quote, a forecast or a valuation of an operator.

What should a first customer contract measure?

It should define workload, input and output, throughput or latency, availability, security, metering, service credits and termination. The unit of payment must match the service that can actually be delivered and verified.

When could debt finance be considered?

The framework proposes considering it only after integrated operations, customer cash flows, permissions, replacement economics, security and downside reserves are independently reviewable. Lender appetite and terms require an actual transaction and due diligence.

Is servicing equivalent to simply replacing a failed satellite?

No. Remote software action, module replacement, robotic intervention and complete satellite replacement have different capabilities, costs and rights. The base case should use only the option supported by the project's tested design and contracts.

References

1. Google Research. Exploring a space-based, scalable AI infrastructure system design. 4 November 2025. <https://research.google/blog/exploring-a-space-based-scalable-ai-infrastructure-system-design/>
 2. Axiom Space. Orbital Data Centers: reported demonstrations and deployments. Accessed 5 October 2026. <https://www.axiomspace.com/orbital-data-center>
 3. European Space Agency. Knowledge beyond our planet: space-based data centres. 5 August 2024. https://www.esa.int/Enabling_Support/Preparing_for_the_Future/Discovery_and_Preparation/Knowledge_beyond_our_planet_space-based_data_centres
 4. European Health and Digital Executive Agency. Grant agreements with 45 new space projects, including the ASCEND concept. 20 October 2022. https://hadea.ec.europa.eu/news/hadea-signs-grant-agreements-45-new-space-projects-2022-10-20_en
 5. NASA Small Spacecraft Systems Virtual Institute. State of the Art of Small Spacecraft Technology: Thermal Control. Accessed 5 October 2026. <https://www.nasa.gov/smallsat-institute/sst-soa/thermal-control/>
 6. International Telecommunication Union. Satellite regulation and coordination. Updated September 2026. <https://www.itu.int/en/mediacentre/backgrounders/Pages/Regulation-of-Satellite-Systems.aspx>
 7. NASA Science. How NASA Prepares Spacecraft for the Harsh Radiation of Space. <https://science.nasa.gov/science-research/heliophysics/how-nasa-prepares-spacecraft-for-the-harsh-radiation-of-space/>
 8. NASA. Update on Status of NASA's OSAM-1 Project. 1 March 2024. <https://www.nasa.gov/missions/update-on-status-of-nasas-osam-1-project/>
 9. Federal Communications Commission. FCC Adopts New Five-Year Rule for Deorbiting Satellites. 29 September 2022. <https://docs.fcc.gov/public/attachments/DOC-387720A1.pdf>
- Sources reviewed 5 October 2026. All financial inputs, figures, tables, stages, proposed contracts and decision thresholds in this paper are author-developed hypothetical illustrations unless expressly attributed to a cited source. The paper is educational research and provides no investment, engineering, legal or regulatory opinion on a particular project.

About the Author

Authored by Chennakeshav Adya
Independent Researcher

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEK, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.