Orbital Compute Investment Gates Testing Launch Access Power and Downlink Economics

An Evidence-Based Capital Release Framework for Early Orbital Compute Ventures

Authored by Chennakeshav Adya
Independent Researcher
September 2026

Abstract

Orbital computing is entering an investable but technically unsettled phase. In February 2026, the United States Federal Communications Commission accepted for filing an application for an orbital data-centre system of up to one million satellites. A separate 2026 filing proposed a distributed orbital-compute constellation of up to 100,000 satellites. Acceptance for filing is a procedural step; it does not constitute authorization, technical validation, customer demand or proof of economic viability. NASA and European Space Agency programmes demonstrate useful onboard processing, radiation-tolerant computing and high-rate optical communications at smaller scales. Together, these developments create a financing question: what evidence should justify each successive release of capital to an early orbital-compute venture? [1][2][3][4]

This paper develops a gated investment framework centred on three linked economic systems: launch access, service-available power and data delivery. Launch economics extend beyond a quoted price per kilogram to include contracted cadence, integration, orbit compatibility, mission assurance, schedule exposure and successful commissioning. Power economics begin with end-of-life generation and storage, then deduct platform, thermal, communications and reserve loads before any compute becomes service available. Downlink economics connect optical or radio link capacity with geometry, weather, gateway diversity, protocol overhead, customer acceptance and the cost of delivering a usable output. A venture becomes financeable when these systems are evidenced together rather than presented as independent technical achievements.

The framework uses four capital gates. Gate 1 funds an integrated ground model and regulatory path. Gate 2 funds a flight demonstration that proves sustained workload execution, power balance and an accepted end-to-end data path. Gate 3 funds a customer pilot with repeatable delivery, measured service availability and collected revenue. Gate 4 funds a commercial cluster only after manufacturing yield, contracted launch access, gateway resilience and unit economics have been demonstrated. Each gate specifies required evidence, permitted uses of funds, decision rights and failure responses.

The worked case is wholly hypothetical and describes no identified company. It assumes a pilot of four nodes with 25 kilowatts of usable compute per node, 6,000 kilograms of aggregate delivered mass and a 40-gigabit-per-second nominal optical downlink per node. Its base case applies 15 percent link duty and 60 percent weather-and-geometry availability, producing 3.6 gigabits per second of effective average downlink per node before customer-specific protocol and acceptance adjustments. Illustrative delivered launch costs of USD 3,000, USD 5,000 and USD 8,000 per kilogram produce pilot launch expenditure of USD 18 million, USD 30 million and USD 48 million. These are scenario inputs and do not represent current vendor pricing.

The analysis concludes that investors should release capital against observable conversion events: licensed architecture, controlled mass and power budgets, contracted launch rights, environmental *qualification, flight telemetry, accepted output, customer payment and repeatable replenishment. A technical milestone without a commercial acceptance test can consume capital while leaving revenue unproven. A commercial agreement without demonstrated power and data delivery can create obligations the platform cannot meet. The investment case becomes credible when the same evidence register reconciles engineering configuration, regulatory status, launch manifest, service capacity, customer acceptance and cash collection.*

JEL Classification: G24, G31, G32, L86, L93, O33

Keywords: orbital compute, investment gates, launch access, spacecraft power, optical downlink, unit economics, milestone financing, low Earth orbit, technology financing, space infrastructure

Introduction

Orbital compute may create value when data originates in space, when raw data volumes exceed practical downlink capacity, when rapid onboard decisions matter, or when a customer values processing that is physically separated from terrestrial infrastructure. Earth-observation filtering, sensor fusion, autonomous spacecraft operations and inter-satellite analytics are plausible applications. NASA's High Performance Spaceflight Computing programme addresses autonomy, artificial intelligence, signal processing, data management and fault tolerance. ESA has supported edge-computing demonstrations that process imagery before transmission. These programmes establish technical relevance; commercial service economics require separate evidence. [3][4][5]

The financing challenge is unusually path dependent. A venture may require meaningful expenditure before its integrated system can be tested in orbit. Once hardware is launched, physical rework is difficult or impossible. Schedule slips can strand completed hardware, defer customer acceptance and consume liquidity. A successful payload may still fail commercially if it lacks gateway access, link availability, contractual demand or a repeatable launch and replacement path.

Investors therefore need capital gates that convert uncertainty into defined experiments. The gate architecture should preserve the option to proceed while limiting irreversible exposure. It should also prevent technical enthusiasm from turning planned capacity into reported commercial value before customers have accepted and paid for the service.

1. Define the capital decision

State the decision in operational terms. The investment committee may be asked to fund architecture definition, a laboratory integration model, a hosted payload, a dedicated flight demonstrator, a customer pilot, a first commercial cluster or a full constellation. Each request should identify the evidence it purchases and the decision it enables.

Set the maximum capital at risk before the next gate. Identify the work packages, suppliers, regulatory actions, tests and contracts funded by the tranche. Specify which outcomes permit release, repricing, redesign, delay or termination. A milestone such as “flight ready” is insufficient unless configuration, test criteria and approval authority are defined.

Separate enterprise ambition from the decision at hand. A venture may ultimately propose a large network. The current capital request may only need to prove one workload, one node and one complete delivery path. The investment memorandum should value the funded evidence and the option it creates, while treating unfunded scale as a contingent scenario.

2. Build a four-gate architecture

Gate 1 should establish architecture credibility. Required evidence includes a controlled system design, regulatory roadmap, mass budget, end-of-life power balance, thermal model, representative workload benchmark, communications concept, supplier plan and customer problem definition. The output is a costed flight-demonstration plan with explicit residual risks.

Gate 2 should establish flight credibility. A launched demonstrator must execute the target workload, maintain power and thermal limits, communicate through the intended path, survive representative operating cycles and return auditable telemetry. Flight success should be measured against pre-agreed thresholds rather than a general statement that contact was established.

Gate 3 should establish service credibility. A customer pilot should use representative input data, defined service levels, documented acceptance tests, security controls, invoice logic and a route to cash collection. The venture should reconcile service-available compute and link capacity with the delivered output.

Gate 4 should establish repeatability. Capital for a commercial cluster should depend on manufacturing yield, contracted launch capacity, gateway diversity, customer minimums, replenishment planning, operational staffing and a credible unit-economics model. The board should know how each additional node affects capacity, revenue, working capital and replacement obligations.

3. Use an evidence hierarchy

Rank evidence by decision usefulness. Executed licences, regulator orders, controlled engineering drawings, signed supplier contracts, environmental-test reports, launch agreements, flight telemetry, customer acceptance certificates, invoices and bank receipts provide strong support. Management presentations, vendor discussions and letters of interest provide context and should remain labelled accordingly.

Give each material assumption an owner, source date, confidence level and refresh date. Launch availability, component lead times, optical-ground-station coverage, spectrum status and customer requirements can change before capital is deployed. Expired evidence should trigger re-underwriting rather than automatic roll-forward.

Require traceability across technical and financial models. The spacecraft configuration used in the mass budget should match the configuration priced by suppliers, tested in qualification, reserved for launch and represented in the financial model. A change in shielding, radiator area or optical terminal may affect mass, power, launch integration and service economics simultaneously.

4. Establish the regulatory gate

Review the complete national authorization path, including satellite application, technical exhibits, spectrum requests, earth-station requirements, debris mitigation, launch and re-entry obligations, export controls and customer-data rules. The FCC's February 2026 public notice records acceptance for filing of a proposed orbital data-centre system and seeks comment. The notice does not grant authority to deploy or operate the proposed system. [1]

Map the international spectrum process. ITU procedures for non-geostationary satellite networks can require advance publication or coordination, followed by notification and recording. Filings include orbital and radio parameters, and some systems are subject to deployment milestones and equivalent power-flux-density limits. The financing plan should reflect the actual frequency bands, services and administering jurisdiction. [6][7][8]

Make regulatory progress a measured gate. Evidence can include filing acceptance, completion of information requests, coordination status, licence grant, conditions, earth-station authority and operational reporting readiness. A requested satellite count, frequency range or orbital shell should remain outside the financeable base case until the relevant rights and funded implementation path exist.

5. Underwrite launch access as a service chain

Launch access consists of more than vehicle capacity. It includes payload compatibility, reserved schedule, integration sequence, test requirements, dispenser or separation hardware, transport, range access, licensing, insurance, orbital insertion and commissioning. The venture should identify which party carries each obligation and cost.

Build a launch evidence schedule. A non-binding quotation supports early planning. A signed reservation may establish a date and deposit while retaining displacement or termination rights. A fully executed launch-services agreement, interface-control document, accepted integration plan and funded payment schedule provide stronger evidence. The investment model should reflect the actual contractual position.

Measure delivered cost per commissioned kilogram and per service-available compute kilowatt. Include integration and mission-assurance cost, expected loss, deployment anomalies and commissioning delay. A low quoted price per kilogram can produce poor economics when the chosen orbit, integration constraints or schedule do not support the service design.

6. Test vehicle and orbit compatibility

Reconcile spacecraft dimensions, mass properties, structural loads, acoustic environment, cleanliness, battery state, hazardous materials, electromagnetic compatibility and separation dynamics with the selected vehicle. Confirm that the desired orbital plane, altitude and injection accuracy are available under the contract.

Distinguish dedicated missions, rideshare missions and hosted payloads. A rideshare may reduce launch expenditure while limiting orbit selection and schedule control. A dedicated mission may improve control and increase the capital committed to one event. A hosted payload may accelerate demonstration while introducing dependency on another spacecraft's power, thermal, pointing and operational priorities.

Model orbital transfer where injection does not equal the operating orbit. Propellant, time, radiation exposure and commissioning workload can reduce effective capacity and delay customer acceptance. The launch gate should close only when vehicle, orbit and spacecraft configuration are jointly supported.

7. Convert cadence into a deployment schedule

Public launch activity indicates growing cadence. The United States Space Force reported 93 launches from the Eastern Range in 2024 and described higher cadence as an objective for 2025. This evidence describes range activity; it does not reserve capacity for a particular commercial programme. [9]

Create a node-level manifest with target launch, backup window, provider, vehicle, orbit, integration milestone, regulatory dependency, payment date and customer consequence. Link manufacturing completion to realistic integration cut-offs. Hardware delivered after a cut-off may wait for a later mission and consume storage, insurance and working capital.

Stress provider concentration, vehicle stand-down, launch failure, range closure, shared-ride displacement and supplier delay. Quantify the effect on service commencement, customer obligations and cash runway. Gate 4 should require a credible backup path for any launch dependency capable of delaying a material share of revenue.

8. Set mission-assurance evidence

Define qualification and acceptance programmes before hardware is complete. Tests can include vibration, acoustic, thermal-vacuum, electromagnetic compatibility, deployment, fault injection, radiation and end-to-end communications. Specify margins, configurations, success criteria, waivers and retest rules.

Distinguish qualification articles from flight units. Evidence from a non-representative prototype should not automatically support a flight configuration with different processors, shielding, radiators or optical hardware. Track every deviation and its effect on risk.

Commissioning criteria should include power-positive operation, thermal stability, attitude control, command and telemetry, workload execution, storage integrity, link acquisition and accepted output. Gate 2 should remain open until the complete service chain has been demonstrated for a defined period.

9. Build the end-of-life power budget

NASA's 2026 small-spacecraft technology review describes solar generation, energy storage and power-management considerations across current spacecraft technologies. Available energy depends on array efficiency, orientation, temperature, eclipse, radiation degradation, battery performance and conversion losses. [10][11]

Start with end-of-life generation rather than beginning-of-life nameplate. Deduct attitude control, propulsion, avionics, storage, communications, thermal-control equipment, battery charging and contingency reserve. The residual power envelope supports compute only when temperature and state-of-charge constraints are also satisfied.

Use time-series analysis. Average power can conceal eclipse deficits, peak communications loads and thermal transients. The model should show energy generation, storage, demand and reserve across representative orbits, seasonal conditions, safe mode and degraded hardware cases.

10. Test storage and eclipse resilience

Battery capacity should be modelled across depth of discharge, cycle count, temperature, calendar ageing, radiation and end-of-life margin. The venture should document how compute load is reduced when stored energy approaches reserve limits.

Define recovery from low state of charge. Identify which loads remain active, how the node reacquires attitude and communications, and how long service remains unavailable. Repeated safe-mode events can affect customer service even if the spacecraft survives.

Link energy storage to commercial capacity. Billable compute kilowatt-hours should exclude periods reserved for recovery, battery protection or platform safety. Gate 3 customer acceptance should use measured energy and workload data across representative eclipse cycles.

11. Integrate thermal rejection

Spacecraft reject heat primarily through radiation. NASA's thermal-control review identifies surface area, heat transport, component limits and mission environment as central design factors. High-duty-cycle computing can turn thermal rejection into the binding capacity constraint. [12]

Build a correlated thermal model covering sunlit operation, eclipse, communications peaks, battery charging, degraded surfaces and safe mode. Record radiator area, view factors, emissivity, heat-pipe or pumped-loop performance, conductive paths and allowable component temperatures.

Translate thermal performance into sustained workload. A processor may reach high laboratory throughput for a short burst and require throttling in orbit. Gate 2 should require representative duty-cycle testing and in-orbit telemetry before investors treat peak compute as service-available capacity.

12. Measure compute per service-available joule

Benchmark the actual workload on representative processors, memory, storage, networking and fault-tolerance software. Record useful output per joule rather than relying on accelerator nameplate. Encryption, error correction, redundancy and orchestration consume power and processing capacity.

NASA's HPSC programme illustrates the need for radiation-tolerant multicore processing, power management and fault tolerance in space. Programme progress supports the relevance of advanced flight computing; it does not establish the performance or economics of a specific commercial system. [3][13]

Use a controlled workload suite. Include normal operation, degraded mode, radiation-induced error recovery, model update, data movement and storage scrubbing. Gate 2 should set minimum throughput, energy efficiency, error rate and recovery-time thresholds.

13. Define the complete data path

Map input acquisition, onboard storage, workload execution, inter-satellite transfer, ground downlink, terrestrial backhaul, customer delivery and acceptance. Identify every protocol conversion, encryption boundary, queue and third-party service.

Separate command and control, bulk data transfer and customer output. Each may use different radio or optical links and ground infrastructure. A node can execute a workload successfully and still fail to provide a billable service if the accepted result does not reach the customer within the required window.

Create a data-flow reconciliation. Input volume, processed volume, stored volume, transmitted volume, received volume and accepted volume should balance after documented compression, filtering and protocol overhead. This reconciliation is the foundation of downlink unit economics.

14. Compare optical and radio links

NASA states that optical communications can provide data rates 10 to 100 times higher than traditional radio systems for comparable mass and power. Its Laser Communications Relay Demonstration operates at 1.2 gigabits per second and has tested two-way relay configurations. These results establish demonstrated capability in specific architectures and environments. [14][15]

Optical links require precise acquisition, pointing and tracking. Space-to-ground performance is sensitive to cloud and atmospheric conditions. NASA's LCRD experience highlights the importance of multiple ground stations and weather diversity. Radio links may offer different weather resilience while requiring spectrum rights, antennas and interference management.

Select links by service requirement. High-volume delay-tolerant output may tolerate intermittent optical windows and storage. Time-critical service may require relay paths, gateway diversity or a radio fallback. Gate 1 should compare mass, power, availability, regulatory requirements and cost for the complete path.

15. Model geometry, duty and availability

Nominal link rate is not average delivered throughput. A low-Earth-orbit node has time-limited visibility to each ground station. Pointing, acquisition, weather, scheduled access, congestion and protocol overhead reduce usable capacity.

For each ground station, model elevation mask, pass frequency, pass duration, weather history, optical turbulence, maintenance and network availability. Model simultaneous demand from other nodes and customers. A relay architecture should include inter-satellite route availability and relay capacity.

Calculate effective average throughput as nominal rate multiplied by geometry duty, weather and technical availability, and protocol efficiency. Keep these factors separate so diligence can verify each one. Gate 2 should compare forecast and observed results across enough passes to capture operational variability.

16. Price gateway diversity

Gateway diversity can improve availability and raise fixed cost, integration complexity and regulatory scope. Identify owned, leased and third-party sites. Review land rights, power, fibre, weather, security, maintenance, staffing and permits.

For optical systems, choose geographically dispersed sites with complementary weather patterns where practical. NASA's LCRD uses optical ground stations in California and Hawaii and an RF ground station in New Mexico; the programme has tested atmospheric effects and operational configurations. [16]

Create a resilience matrix. For each site or provider outage, identify lost capacity, rerouting option, customer effect and recovery time. Gate 4 should require sufficient contracted gateway capacity and diversity for the service-level commitments used in the revenue model.

17. Complete spectrum and interference diligence

Radio links require frequency allocation, licensing, coordination and interference analysis appropriate to service, direction, geography and orbit. ITU guidance explains that some non-GSO networks begin with advance publication while others require coordination. National authorizations and earth-station licences remain necessary. [6][7]

Review equivalent power-flux-density limits, out-of-band emissions, antenna patterns, sharing studies, coordination correspondence and operational constraints. A large requested frequency range should not be treated as usable capacity before the regulatory record supports it.

Include spectrum milestones in the capital plan. Architecture funding may precede filing. Flight funding should depend on a credible experimental or operational authority. Commercial funding should depend on rights sufficient for the customer service and deployment schedule.

18. Calculate cost per accepted terabyte

Allocate communications cost to accepted customer output. Include optical and radio terminals, ground stations, leased network capacity, terrestrial backhaul, operations, spectrum and licence costs, maintenance, replacement, data storage and expected retransmission.

Divide by data accepted under the customer contract, not nominal transmitted data. Weather losses, failed passes, protocol overhead, corrupted files and rejected output consume resources without producing accepted service.

Compare downlinking raw data with processing in orbit. If onboard processing reduces data volume by discarding unusable observations or extracting relevant features, the venture may avoid communications cost. The economic benefit equals avoided end-to-end delivery cost and improved decision value, less onboard compute, power, storage and model-management cost.

19. Measure end-to-end latency

Propagation time is one component of latency. Add acquisition wait, queueing, inter-satellite routing, compute, storage, gateway availability, terrestrial backhaul, decryption and customer processing. Report percentile latency and availability rather than a single best observation.

Compare the orbital path with terrestrial alternatives for the specific workload. The relevant benchmark may be raw-data downlink followed by cloud processing, an existing onboard processor, a terrestrial edge facility or delayed batch delivery. Customer value depends on the operational consequence of time saved.

Gate 3 should require a service-level test with representative input and customer acceptance. A technical demonstration of low propagation delay alone does not establish a valuable service.

20. Underwrite the ground segment

The ground segment includes mission operations, telemetry, tracking and command, key management, software deployment, workload scheduling, data receipt, monitoring, customer delivery and incident response. Identify which functions are internal and which depend on service providers.

Test operating coverage, staffing model, cybersecurity, change control, data segregation and disaster recovery. A constellation can increase the volume and velocity of operational decisions. Automation should have human accountability, logged actions and safe fallback modes.

Include ground systems in qualification. End-to-end tests should use production-representative networks, keys, customer interfaces and monitoring. Gate 2 and Gate 3 evidence should show that telemetry and customer delivery can be reconciled without manual reconstruction.

21. Define customer acceptance

For each target workload, specify input source, expected output, accuracy or completeness threshold, delivery window, security requirement, service credits, rejection rights and invoice basis. The technical team and customer should agree the test before launch.

Separate development revenue, milestone payments, reserved capacity, take-or-pay commitments and usage charges. Assess termination, appropriations, launch-delay and force-majeure provisions. A contract can contain substantial conditions before revenue becomes enforceable or collectable.

Gate 3 should require evidence of accepted service and cash collection. Where the customer pilot is funded through research or development expenditure, distinguish that payment from recurring commercial demand.

22. Build unit economics from physical drivers

Revenue may be priced per workload, accepted terabyte, service-available compute hour, reserved capacity or outcome. Costs should include node manufacturing, launch, integration, ground systems, communications, mission operations, insurance, customer support, replacements and corporate overhead.

Calculate contribution margin by service and by node cohort. Include unsuccessful launches, early failures, degraded power, downlink constraints and service credits. Recognize replacement capital as an economic requirement rather than assuming that depreciation alone funds new hardware.

Use cash timing. Launch deposits and long-lead procurement can precede service revenue by years. Customer prepayments may be refundable or performance contingent. Gate 4 should include working-capital and liquidity coverage under delay cases.

23. Gate 1: architecture and regulatory readiness

Gate 1 releases capital for controlled design, regulatory work, supplier engagement, laboratory integration and customer discovery. Minimum evidence should include a version-controlled architecture, mass and power budgets with margin, thermal model, communications options, workload benchmark and preliminary safety analysis.

The venture should have a documented authorization path and responsible advisers. It should identify launch options, interface requirements, critical suppliers and long-lead items. Customer evidence should define a costly problem and acceptance criterion rather than broad interest in space computing.

The gate closes with an independent design review and a costed flight-demonstration plan. Unresolved issues should be assigned owners, budget and decision dates. Capital for flight hardware should remain contingent on closure of the defined critical items.

24. Gate 2: flight demonstration

Gate 2 releases capital for representative flight hardware, qualification, launch, commissioning and an end-to-end technical demonstration. The funded configuration should be traceable to the design reviewed at Gate 1.

Success criteria should include launch survival, deployment, power-positive operation, thermal control, representative workload execution, data integrity, link acquisition, accepted ground receipt and recovery from defined faults. Measurements should cover enough cycles and communication passes to distinguish repeatable performance from one successful event.

Failure responses should be pre-agreed. A recoverable software defect may justify remediation. A structural power or thermal shortfall may require redesign. The investment documents should define information rights, budget control and whether the next tranche can be withheld or repriced.

25. Gate 3: customer pilot

Gate 3 funds a limited service pilot using representative customer data and operational processes. The venture should demonstrate that technical capability can become an accepted and invoiced service.

Measure service-available compute, output quality, effective downlink, latency, availability, incidents, manual interventions and customer acceptance. Reconcile telemetry with invoice quantities. Identify the cost of customer-specific integration and support.

The gate closes when acceptance and payment evidence support a repeatable commercial proposition or when the board concludes that the workload should be changed. A successful technology demonstration without customer adoption should lead to a revised service thesis before scale capital is released.

26. Gate 4: commercial cluster

Gate 4 finances repeatable production and deployment. Required evidence includes manufacturing yield, quality control, supplier capacity, contracted launch slots, gateway capacity, operational tooling, customer minimums, insurance strategy and replacement planning.

Build a cohort model. Each group of nodes should have manufacturing cost, launch date, commissioned capacity, degradation profile, customer allocation, operating cost and expected replacement date. Performance evidence from earlier cohorts should update later procurement and valuation.

Set distribution and additional-debt tests. Cash should support operations, customer obligations and replacement reserves before discretionary distributions. Project debt should depend on contracted and accepted cash flows, robust security and lender step-in rights appropriate to the structure.

27. Match financing to evidence

Early architecture risk is commonly suited to sponsor, venture or strategic equity. Equipment or vendor finance may become possible when hardware configuration and supplier obligations are stable. Customer prepayments can support working capital when refund, performance and security terms are understood.

Project or asset-backed debt requires predictable cash flows and enforceable rights. Spacecraft, spectrum rights, software, contracts and insurance proceeds may have jurisdictional and practical enforcement constraints. Security design requires specialist legal advice.

Use milestone tranches, escrow, conditions precedent and budget controls to align funding with evidence. Avoid financing structures that force deployment before regulatory, technical or customer gates are complete.

28. Design transaction protections

Conditions precedent can require authorization, completed qualification, executed launch arrangements, supplier contracts, customer commitments, insurance and an approved budget. Representations should address ownership, technical data, licences, compliance, contracts and disputes.

Milestone tranches can reduce exposure while creating negotiation risk if criteria are subjective. Use measurable thresholds, named evidence and a defined certifier. Consider cure rights, independent expert determination and consequences of partial achievement.

Information rights should cover engineering changes, test failures, regulatory correspondence, launch status, telemetry, customer acceptance, cash and forecasts. Reserved matters can protect changes to architecture, suppliers, launch provider, spectrum strategy, budget and related-party arrangements.

29. Create one controlled data room

Organize the data room around the investment gates. The regulatory folder should contain filings, correspondence, licences and conditions. The engineering folder should contain configuration, requirements, models, test reports and non-conformances. The commercial folder should contain customer discovery, contracts, acceptance and invoices.

The launch folder should include quotations, reservations, agreements, interface documents, schedule and payment records. The financial folder should connect unit costs, funding uses, cash runway, revenue assumptions and replacement requirements to source evidence.

Maintain a change log. Each material model input should link to a dated source. When configuration changes, update affected mass, power, thermal, communications, schedule and finance records together.

30. Govern post-investment execution

The board should review a compact evidence dashboard. Metrics should cover regulatory milestones, mass and power margin, qualification status, launch readiness, commissioned capacity, service availability, accepted output, customer collections, liquidity and replacement coverage.

Define triggers for escalation. Examples include loss of power margin, missed integration cut-off, failed environmental test, regulator information request, gateway underperformance, customer rejection or runway below the next gate requirement.

Keep decision records. Each capital release should state evidence reviewed, assumptions retained, limitations, dissent and follow-up actions. This discipline protects the venture and investors when the programme evolves.

31. Stress the combined system

Stress variables together. Launch delay may increase storage cost and defer revenue while customer milestones expire. Lower power can reduce compute and increase the time needed to process data, which can increase storage and miss downlink windows. Poor weather can increase retransmission and gateway cost.

Run at least launch, power, communications, customer and liquidity downside cases. Include a combined case with delayed launch, lower usable power, reduced link availability and slower customer ramp. Model the management action available in each case.

Use break-even analysis. Identify the minimum accepted workload volume, contribution margin and node availability required to fund operations and replacement. Gate 4 should not rely on a single optimistic intersection of technical and commercial assumptions.

32. Apply the hypothetical pilot case

The illustrative pilot contains four nodes, each with 25 kilowatts of usable compute and 1,500 kilograms of delivered mass. Aggregate delivered mass is 6,000 kilograms and planned usable compute is 100 kilowatts. A nominal 40-gigabit-per-second optical downlink is assigned to each node.

Illustrative delivered launch costs of USD 3,000, USD 5,000 and USD 8,000 per kilogram produce launch expenditure of USD 18 million, USD 30 million and USD 48 million. These values exclude node manufacturing and are wholly hypothetical. The investment committee should replace them with executed contractual evidence.

The base communications case applies 15 percent link duty and 60 percent combined weather-and-geometry availability. Effective average throughput is therefore 3.6 gigabits per second per node before protocol efficiency and customer rejection. Across four nodes, the gross effective rate is 14.4 gigabits per second. A 75 percent protocol and acceptance factor produces 10.8 gigabits per second of accepted-equivalent throughput, or approximately 116.6 terabytes per day using decimal units.

The arithmetic demonstrates why nominal link rate is an incomplete metric. The effective service depends on ground visibility, weather, terminal performance, network scheduling, protocol overhead and customer acceptance. Each factor should be evidenced separately at Gate 2 and Gate 3.

33. Frame the investment committee decision

Approve Gate 1 when the problem, architecture, regulatory path and evidence plan are credible within a capped budget. Approve Gate 2 when the flight configuration, launch path, qualification plan and success criteria are controlled. Approve Gate 3 when the technical demonstration supports a specific customer acceptance test. Approve Gate 4 when service, manufacturing, launch, gateway and cash evidence support repeatability.

Use conditional approval where residual items have a clear resolution path. State the maximum exposure before each condition is met. Avoid open-ended bridge funding that preserves the programme without purchasing defined evidence.

Record the counterfactual. The committee should understand the value of waiting, using a hosted payload, changing the workload, partnering with a communications provider, acquiring technology or declining the opportunity.

Conclusion

Orbital-compute investment should proceed through evidence gates that join engineering and commercial proof. Launch access must be contractually and technically usable. Power must remain available after end-of-life degradation, platform demand and thermal limits. Downlink must convert nominal rate into accepted customer output under real geometry, weather and operational constraints.

The four-gate framework turns these requirements into capital decisions. Architecture funding buys a controlled design and evidence plan. Flight funding buys an integrated demonstration. Pilot funding buys customer acceptance and cash evidence. Commercial-cluster funding buys repeatability. Each release has a defined purpose, evidence threshold and failure response.

The central discipline is reconciliation. Regulatory status, hardware configuration, launch manifest, power telemetry, link performance, customer acceptance and financial reporting should describe the same system. When that reconciliation is maintained, investors can support technical ambition with controlled exposure and boards can distinguish progress from activity.

Appendix A. Gate 1 diligence checklist

Confirm the customer problem, representative workload, service definition and acceptance criterion. Review architecture, mass budget, end-of-life power balance, thermal model, compute benchmark, communications options, regulatory roadmap, safety concept, supplier map and costed demonstration plan.

Confirm version control, named owners and refresh dates. Require a register of assumptions, limitations and unresolved items. Establish the maximum Gate 1 budget and the independent review needed before flight capital is released.

Appendix B. Gate 2 flight checklist

Confirm controlled flight configuration, qualification and acceptance records, waivers, launch contract, interface-control document, integration schedule, insurance position, commissioning plan and telemetry archive.

Define pass or fail criteria for launch survival, deployment, power, thermal performance, workload throughput, data integrity, communication, accepted receipt and fault recovery. Record the required observation period and approving authority.

Appendix C. Gate 3 customer checklist

Confirm representative customer data, security classification, service-level definition, accuracy or completeness threshold, delivery window, interface, invoice basis, rejection process and payment terms.

Reconcile spacecraft telemetry, ground receipt, customer acceptance, invoice and bank collection. Separate development funding from recurring service demand.

Appendix D. Gate 4 scale checklist

Confirm manufacturing yield, supplier capacity, quality controls, launch reservations, gateway diversity, operating procedures, customer minimums, liquidity, insurance and replacement plan.

Review the cohort model, downside cases, reserve policy, distribution tests, additional-debt limits and board dashboard. Verify that every scale assumption has current evidence.

Appendix E. Hypothetical equations

Delivered launch expenditure equals delivered mass multiplied by delivered cost per kilogram. Service-available compute equals installed compute multiplied by power availability, thermal availability, hardware availability and workload efficiency.

Effective average downlink equals nominal link rate multiplied by geometry duty, weather availability, technical availability and protocol efficiency. Accepted throughput further deducts output rejected under customer criteria.

Contribution margin equals collected service revenue less variable mission operations, communications, customer delivery, service credits and expected replenishment cost. These equations are analytical constructs and require transaction-specific inputs.

Appendix F. Decision figures and tables

Figure 1. Hypothetical pilot launch expenditure by delivered-cost scenario

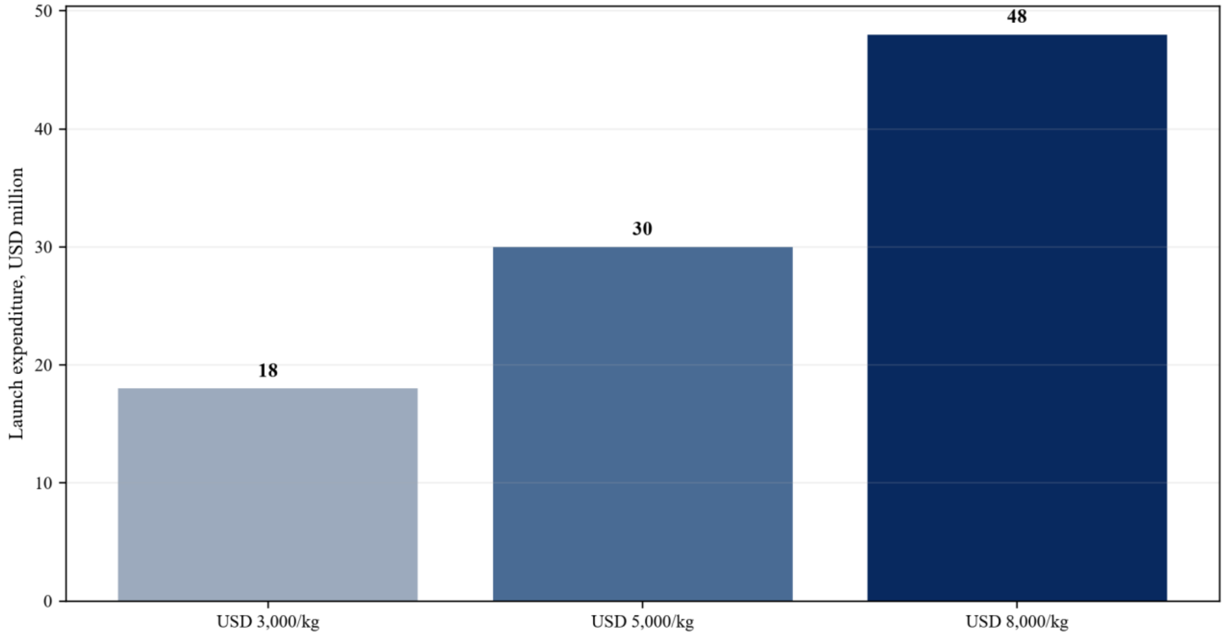

Wholly hypothetical; 6,000 kilograms of aggregate delivered mass.

Table 1. Four-gate capital framework

Gate	Capital purpose	Required conversion event
Architecture	controlled design, regulatory work and ground integration	independently reviewed flight-demonstration plan
Flight	representative hardware, qualification, launch and commissioning	repeatable end-to-end technical performance
Customer pilot	representative service and operational delivery	customer acceptance, invoice and cash evidence
Commercial cluster	repeatable production, launch and gateway capacity	cohort economics and replenishment evidence

Proposed investment framework.

Figure 2. Hypothetical node power allocation

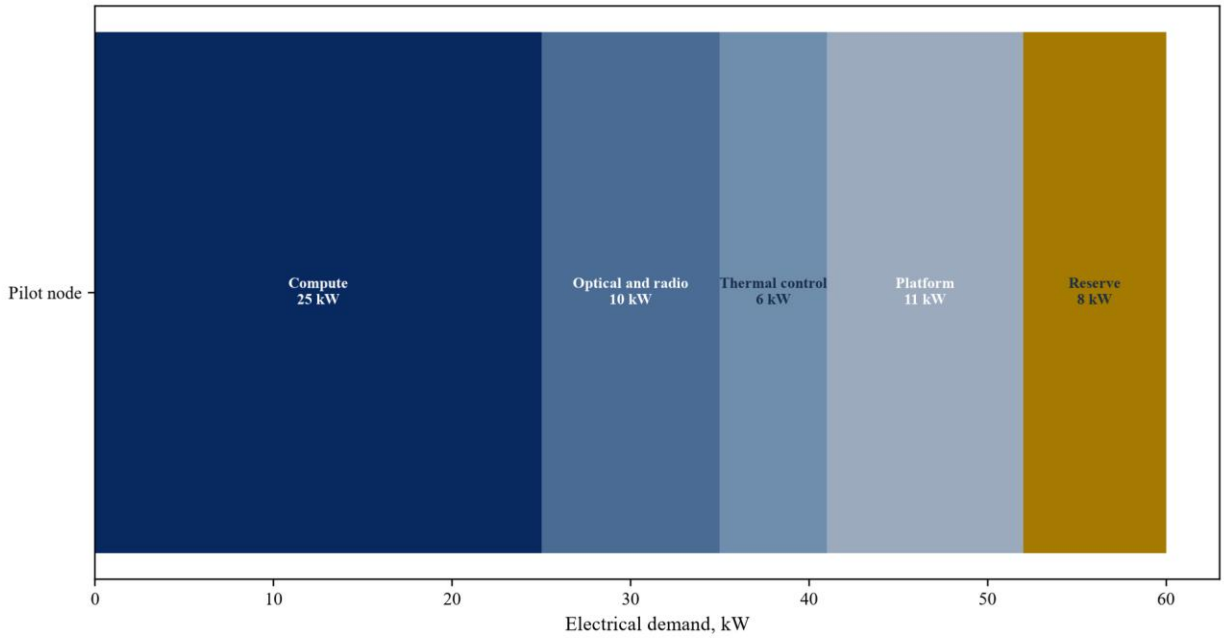

Wholly hypothetical 60-kilowatt total node demand.

Table 2. Launch-access evidence schedule

Evidence state	Example evidence	Investment treatment
Market indication	non-binding quotation and preliminary compatibility	planning range only
Reserved	signed reservation, deposit and target window	schedule evidence subject to rights and conditions
Integrated	accepted interface, test record and integration slot	flight readiness subject to launch and deployment
Commissioned	launch record, deployment and in-orbit acceptance	operating capacity subject to service test

Proposed evidence requirements.

Figure 3. Hypothetical downlink conversion from nominal rate to accepted-equivalent throughput

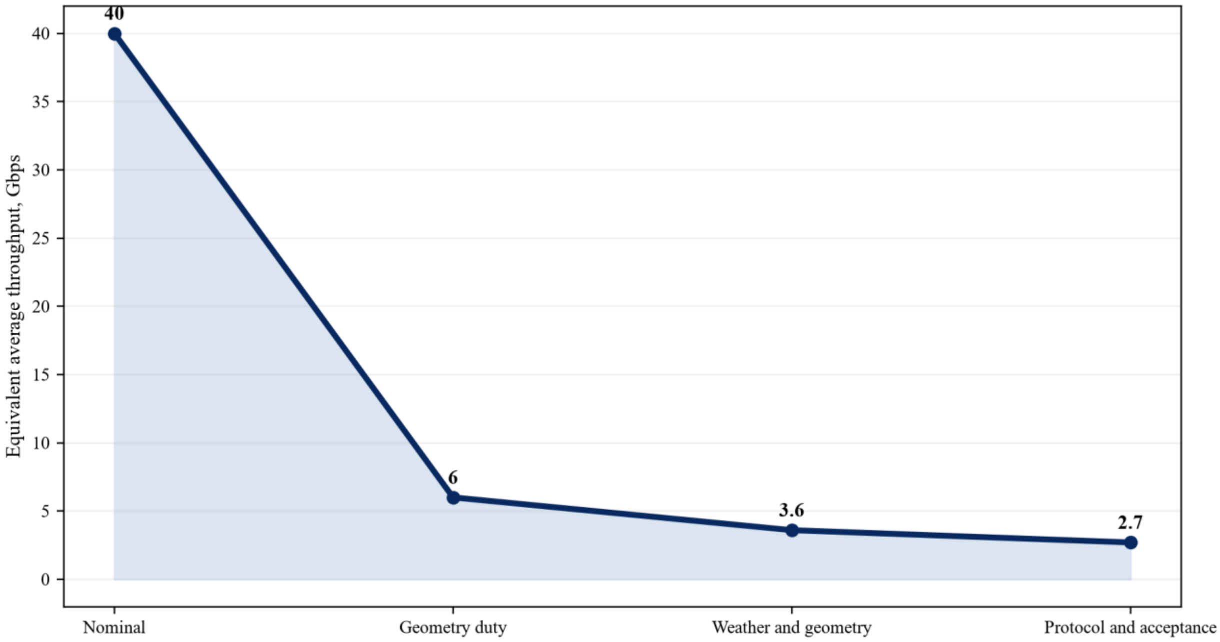

Wholly hypothetical per-node waterfall.

Table 3. Power and communications acceptance tests

Test	Minimum evidence	Commercial link
End-of-life power	correlated orbit and degradation model	service-available energy
Thermal duty cycle	ground and in-orbit temperature telemetry	sustained workload
Workload efficiency	representative output per joule	compute capacity
Link acquisition	repeated acquisition across geometry	delivery availability
Customer receipt	complete accepted output and audit trail	invoice quantity

Proposed framework.

Figure 4. Hypothetical effective downlink sensitivity

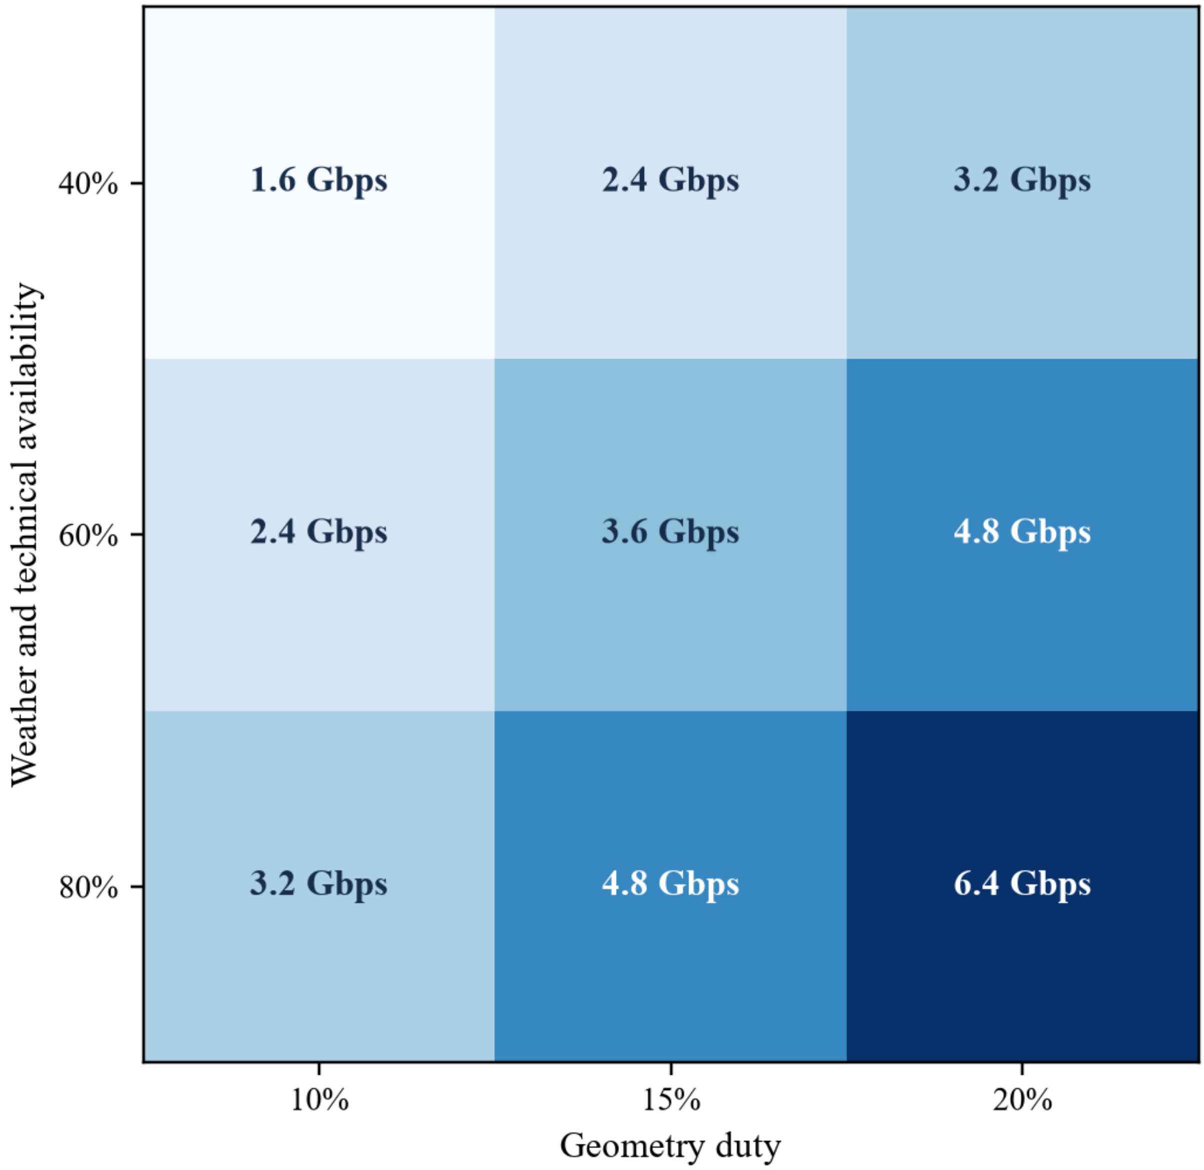

Wholly hypothetical 40-gigabit-per-second nominal link; before protocol adjustment.

Table 4. Hypothetical pilot uses of funds

Use	Central case
Architecture, software and qualification	28
Four flight nodes and integration	84
Delivered launch expenditure	30
Ground and gateway systems	22
Operations, contingency and working capital	36
Total	200

Wholly hypothetical; USD million.

Figure 5. Hypothetical milestone capital release

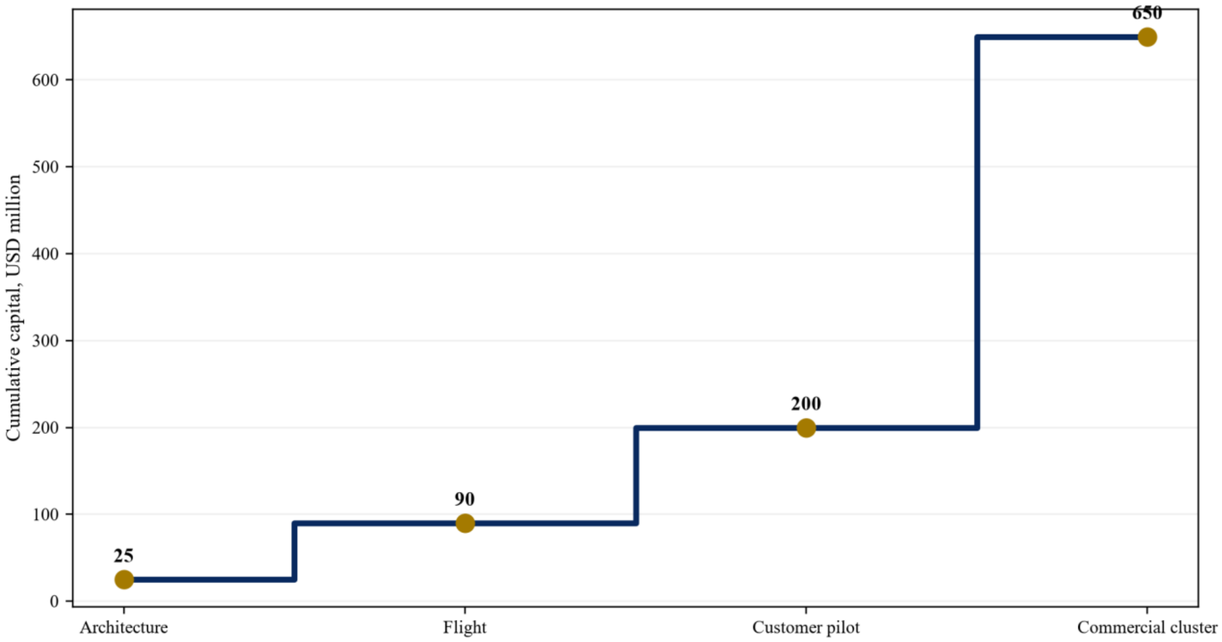

Wholly hypothetical cumulative capital.

Table 5. Illustrative capital release conditions

Gate	Illustrative cumulative capital	Release condition
Architecture	USD 25m	controlled design and approved flight plan
Flight	USD 90m	qualified hardware, launch rights and objective success criteria
Customer pilot	USD 200m	repeatable end-to-end performance and signed pilot
Commercial cluster	USD 650m	accepted service, cash evidence and repeatable deployment plan

Proposed framework with wholly hypothetical amounts.

Figure 6. Orbital compute investment control system

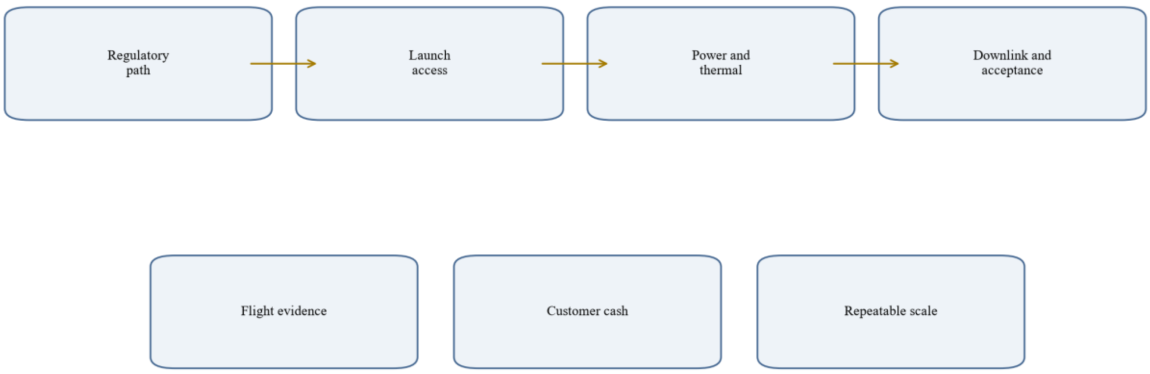

Proposed governance framework.

Table 6. Transaction protection matrix

Risk	Protection	Evidence or trigger
Authorization delay	condition precedent or staged closing	defined licence milestone
Launch displacement	long-stop date and alternative-provider plan	missed integration or launch window
Power or thermal shortfall	milestone tranche and redesign budget	failed acceptance threshold
Downlink underperformance	gateway cure plan and service repricing	measured availability variance
Customer non-acceptance	pilot scope, cure and termination rights	rejected output or unpaid invoice
Replacement underfunding	reserve and distribution test	cohort lifecycle forecast

Terms require transaction-specific legal advice.

Table 7. Board evidence dashboard

Metric	Source	Escalation trigger
Regulatory milestone status	regulator record and counsel report	delay or adverse condition
Mass and power margin	controlled engineering model	margin below approved threshold
Launch readiness and slot	provider record and integration plan	missed cut-off or unreserved capacity
Service-available compute	telemetry and workload results	variance from gate threshold
Accepted downlink throughput	ground and customer records	service-level shortfall
Customer collections	invoices and bank records	overdue or disputed payment
Liquidity to next gate	treasury forecast	runway below approved buffer

Proposed framework.

Frequently asked questions

Q: What makes an orbital-compute milestone investable? A: An investable milestone converts a material uncertainty into objective evidence, has defined success criteria, fits within a capped budget and enables a specific next decision.

Q: Is an accepted satellite application equivalent to authorization? A: No. Acceptance for filing is a procedural step. Investors should review the complete regulatory record and base deployment on rights actually granted and implementable.

Q: How should launch access be measured? A: Measure contracted and compatible launch capacity, integration status, orbit suitability, delivered cost, schedule rights, mission assurance and the probability of successful commissioning.

Q: Why is spacecraft nameplate power insufficient? A: Platform loads, conversion loss, eclipse, degradation, battery policy, thermal limits and reserve reduce the energy available for sustained customer workloads.

Q: How should optical downlink capacity be modelled? A: Multiply nominal rate by geometry duty, weather availability, technical availability and protocol efficiency, then deduct output rejected under customer acceptance criteria.

Q: When should commercial-cluster capital be released? A: Release scale capital after representative flight performance, customer acceptance and payment, manufacturing yield, contracted launch access, gateway capacity and replenishment economics are evidenced.

Q: What does the hypothetical pilot show? A: It shows that a 40-gigabit-per-second nominal link can become 2.7 gigabits per second of accepted-equivalent throughput per node after illustrative duty, availability and protocol factors.

Q: Can project debt fund an early orbital-compute venture? A: Early architecture and flight risk usually require risk capital. Debt becomes more credible when accepted contracted cash flows, enforceable rights, reliable deployment and replacement coverage are demonstrated.

References

[1] Federal Communications Commission, Space Bureau Accepts for Filing SpaceX Application for Orbital Data Centers, DA 26-113, 4 February 2026. https://docs.fcc.gov/public/attachments/DA-26-113A1.pdf

[2] Federal Communications Commission, Orbital Compute Inc. Application SAT-LOA-20260624-00251, 2026 public notice. https://docs.fcc.gov/public/attachments/DOC-424016A1.pdf

[3] National Aeronautics and Space Administration, High Performance Spaceflight Computing. https://www.nasa.gov/game-changing-development-projects/high-performance-spaceflight-computing-hpsc/

[4] European Space Agency Commercialisation Gateway, New Phi-lab in Sweden and the Rise of Edge Computing in Space. https://commercialisation.esa.int/2022/05/new-%CF%86-lab-in-sweden-and-the-rise-of-edge-computing-in-space/

[5] European Space Agency, Data Systems. https://www.esa.int/Enabling_Support/Space_Engineering_Technology/Data_Systems

[6] International Telecommunication Union, Non-geostationary-satellite networks. https://www.itu.int/en/ITU-R/space/support/nonGSO/Pages/default.aspx

[7] International Telecommunication Union, Advance Publication Information. https://www.itu.int/en/ITU-R/space/Pages/API.aspx

[8] International Telecommunication Union, Satellite regulation and coordination. https://www.itu.int/en/mediacentre/backgrounders/Pages/Regulation-of-Satellite-Systems.aspx

[9] United States Space Force, Launch Organization Leans into Modern Mission Assurance and Agile Integration, 2025. https://www.ssc.spaceforce.mil/Newsroom/Article/4170059/space-systems-command-briefing-lanes-phases-and-tranches-ussf-launch-organizati

[10] National Aeronautics and Space Administration, State-of-the-Art of Small Spacecraft Technology 2026. https://www.nasa.gov/smallsat-institute/sst-soa/

[11] National Aeronautics and Space Administration, Small Spacecraft Power Systems. https://www.nasa.gov/smallsat-institute/sst-soa/power-subsystems/

[12] National Aeronautics and Space Administration, Small Spacecraft Thermal Control. https://www.nasa.gov/smallsat-institute/sst-soa/thermal-control/

[13] NASA Goddard Engineering and Technology Directorate, HPSC Radiation-Hardened Multicore Technology. https://etd.gsfc.nasa.gov/our-work/hpsc-transforming-spaceflight-computing-with-radiation-hardened-multicore-technology/

[14] National Aeronautics and Space Administration, Laser Communications Relay Demonstration Overview. https://www.nasa.gov/directorates/stmd/tech-demo-missions-program/laser-communications-relay-demonstration-lcrd-overview/

[15] National Aeronautics and Space Administration, Laser Communications. https://www.nasa.gov/communicating-with-missions/lasercomms/

[16] National Aeronautics and Space Administration, Exploration and Space Communications LCRD. https://www.nasa.gov/goddard/esc/lcrd/

[17] NASA Technical Reports Server, LCRD Experiment Program Characterization and Initial Operations. https://ntrs.nasa.gov/citations/20240001299

[18] National Aeronautics and Space Administration, Small Spacecraft Communications. https://www.nasa.gov/smallsat-institute/sst-soa/communications/

[19] National Aeronautics and Space Administration, Small Spacecraft Avionics. https://www.nasa.gov/smallsat-institute/sst-soa/small-spacecraft-avionics/

[20] National Aeronautics and Space Administration, Small Spacecraft Structures Materials and Mechanisms. https://www.nasa.gov/smallsat-institute/sst-soa/structures-materials-and-mechanisms/

[21] United States Government Accountability Office, Missile Warning Satellites Risks to Capability Delivery, GAO-26-107085. https://www.gao.gov/products/gao-26-107085

[22] National Aeronautics and Space Administration, Orbital Debris Programme Office. https://sma.nasa.gov/sma-disciplines/orbital-debris/orbital-debris-educating-people-on-requirement-changes

[23] National Aeronautics and Space Administration, NPR 8715.6B Requirements for Limiting Orbital Debris. https://nodis3.gsfc.nasa.gov/displayAll.cfm?Internal_ID=N_PR_8715_006B_&page_name=ALL

[24] Federal Communications Commission, Space Innovation; Mitigation of Orbital Debris in the New Space Age, FCC 24-6. https://docs.fcc.gov/public/attachments/FCC-24-6A1.pdf

[25] International Organization for Standardization, ISO 24113 Space systems Space debris mitigation requirements. https://www.iso.org/standard/83474.html

[26] European Space Agency Space Solutions, SpaceCloud Demonstration. https://business.esa.int/projects/scd

[27] IFRS Foundation, IFRS 13 Fair Value Measurement. https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/

[28] IFRS Foundation, IAS 36 Impairment of Assets. https://www.ifrs.org/issued-standards/list-of-standards/ias-36-impairment-of-assets/

[29] IFRS Foundation, IAS 38 Intangible Assets. https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/

[30] NASA Technical Reports Server, State-of-the-Art Small Spacecraft Technology, 2025. https://ntrs.nasa.gov/citations/20250000142

About the Author

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 closure.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, 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 and leads 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.