# Orbital AI Inference for Earth Observation Capturing Value before Downlink

*An Investment Framework for Bandwidth Reduction Decision Latency Customer Acceptance and Lifecycle Economics*

**Authored by Chennakeshav Adya**
Independent Researcher
September 2026

## Abstract

*Earth-observation satellites traditionally collect imagery, store it onboard and transmit selected data to ground stations for processing. Orbital artificial-intelligence inference changes that sequence. A satellite can classify scenes, reject cloud-obscured captures, detect events, compress information and prioritize downlink before the full data set reaches Earth. The investment question is whether this redistribution of computation produces a customer result that is faster, cheaper or more valuable than an alternative ground-processing architecture.*

*Public missions now provide credible technical evidence. The European Space Agency's PhiSat-2 is a six-unit CubeSat with onboard applications for cloud filtering, street-map generation, vessel detection, image compression, wildfire detection and marine-pollution detection. ESA reported in July 2025 that commissioning had completed and the mission had entered its science phase. NASA's Dynamic Targeting work on CogniSAT-6 demonstrated a closed-loop process in which onboard artificial intelligence assessed imagery and changed observation priorities in less than ninety seconds without human intervention. NASA has also described deployment of its Prithvi geospatial foundation model on two in-orbit platforms. These programmes support technical diligence; each commercial case still requires evidence of customer acceptance, operating reliability and willingness to pay. [1][2][3][4][5]*

*This paper provides an investment and transaction framework for orbital inference. It begins with the customer decision, maps the path from observation to accepted output and separates filtering, compression, classification and autonomous targeting. It then connects downlink capacity, latency, model performance, sensor utilization, power, thermal limits, radiation, cyber security, regulation, refresh and replacement to revenue, cost and risk-adjusted value. The framework is designed for founders, strategic customers, investors, lenders and acquirers evaluating a platform rather than a single demonstration.*

*The worked case is wholly hypothetical and describes no identified company. It assumes a twelve-satellite constellation collecting forty-eight terabytes of raw observations per day. Ground and contact constraints permit twelve terabytes per day of baseline downlink. An orbital-inference architecture selects three terabytes of prioritized data and alert products while retaining sampled raw data for validation. Baseline annual revenue is USD 32 million. Faster alerts and improved product availability add USD 20.8 million of accepted annual revenue, while lower transmission and ground-processing requirements avoid USD 6 million of annual cost. Incremental annual operating and lifecycle cost is USD 13 million, producing USD 13.8 million of incremental operating cash contribution before tax and financing. Incremental build and deployment capital is USD 68 million. An illustrative probability-weighted enterprise-value uplift of USD 46 million follows only after adjustments for remaining capital, customer adoption, technical performance and replacement risk. Every amount, rate, probability and value is hypothetical and must be replaced with verified transaction evidence.*

*The analysis finds that onboard inference is most investable where the observation loses value quickly, downlink is constrained, the model can be tested against a stable acceptance protocol and a customer pays for a defined decision product. Cloud screening and compression can create value with relatively simple acceptance tests. Event detection and retasking can create higher value and higher liability. The strongest financing structures release capital against demonstrated stages: representative data performance, hardware qualification, in-orbit commissioning, customer acceptance and repeatable cash collection. Investors should value contracted outputs, measured bandwidth reduction, accepted latency and funded lifecycle controls. Installed compute, model novelty and total imagery collected have no independent value unless they improve an enforceable customer outcome.*

**JEL Classification:** G24, G31, L86, L93, O32, O33

**Keywords:** orbital AI inference, Earth observation, edge computing, onboard processing, downlink economics, satellite data, decision latency, valuation, technology financing, remote sensing

## Introduction

Earth observation creates value after a user receives information that supports a decision. The chain begins with tasking and sensor capture, then continues through calibration, storage, transmission, processing, interpretation and delivery. Delay, cloud cover, congested ground contacts and large data volumes can reduce the useful portion of what a satellite collects. Orbital inference moves selected processing into that chain before downlink.

The economic case varies by workload. A cloud mask can prevent transmission of unusable imagery. A vessel detector can transmit a compact alert and selected image chip. A wildfire model can support rapid notification. An autonomous-targeting model can alter where the satellite observes next. Each use case has a different accuracy threshold, latency value, legal exposure and customer acceptance process.

Capital providers therefore need an evidence system that connects model output to contract, cash and lifecycle cost. This paper supplies that system. It treats orbital artificial intelligence as part of a complete data service, assesses technical maturity at system level and values only outcomes that can be measured and accepted.

## 1. Define the economic product

The product is an accepted information outcome, rather than an algorithm running in orbit. A customer may buy a verified cloud-free image, an event alert, a change-detection layer, a prioritized queue or an instruction that supports another asset. The contract should state what is delivered, when it is delivered, how quality is measured and what happens when the output is wrong or late.

Four operating modes require separate underwriting. Filtering rejects low-value observations. Compression reduces transmitted bits while preserving an agreed level of utility. Inference classifies, detects or estimates something in the scene. Autonomous targeting changes collection behaviour based on an onboard assessment. Technical and commercial risk generally increase across that sequence because later modes place more consequence on the model's action.

The investable perimeter includes the sensor interface, onboard processor, model, flight software, storage, communications, ground systems, customer interface and evidence archive. If material components belong to different companies, transaction diligence should map rights, service dependencies, termination provisions and change-of-control effects.

## 2. Start with the customer decision

An orbital model should be traced to the decision it improves. Emergency managers may decide where to inspect or deploy resources. Maritime users may decide whether to investigate a vessel. Agricultural users may decide whether a crop anomaly merits action. Insurers may decide whether an event requires further evidence. Each decision has a useful time window and a cost of error.

The customer specification should define geography, observation frequency, minimum confidence, tolerated false positives, tolerated missed events, maximum delivery time, required source imagery and human review. These elements convert a technology proposition into a service acceptance protocol.

A broad claim about faster insight is insufficient for valuation. The model should quantify how many decisions occur, how often the orbital path changes the result, what alternative service exists, what the customer pays and how the output is audited. Revenue belongs in the base case only when these elements are supported by an executed contract or comparable evidence.

## 3. Map the observation-to-action chain

The complete chain contains tasking, sensor readiness, capture, calibration, data transfer to the processor, preprocessing, model execution, confidence scoring, storage, prioritization, communications scheduling, downlink, ground validation, product assembly, delivery and customer acceptance. A fast model can coexist with a slow service if any other stage dominates elapsed time.

Investors should record the owner, capacity, failure mode and evidence for every stage. The map should also identify where data is discarded, where a decision becomes irreversible and which raw evidence remains available for later review. These details affect customer remedies, model assurance and dispute resolution.

NASA defines product latency as the time between observation and availability to the user. Its near-real-time LANCE services target availability within three hours for selected products. Orbital inference addresses a narrower part of that chain and can create value where the commercial requirement is shorter or the downlink bottleneck is material. [18][19]

## 4. Quantify decision latency

Latency should be measured from requested observation to accepted customer output. Separate tasking delay, orbital revisit, capture, onboard processing, contact wait, transmission, ground processing, delivery and customer validation. Report median, tail and failure cases because time-sensitive services can lose value in the tail.

NASA's 2025 Dynamic Targeting test used onboard analysis to decide where to point a sensor next, completing the process in less than ninety seconds without human involvement. The result demonstrates a closed technical loop on the tested mission. A commercial underwriting case still needs evidence that the selected target, output quality and response time meet a paying customer's protocol. [3]

Latency value should be expressed as incremental accepted revenue, avoided loss, better asset utilization or reduced operating cost. A minute saved has different value for wildfire confirmation, maritime surveillance, crop analysis and long-horizon mapping. The valuation model should use customer-specific time-value curves rather than a single price per second.

## 5. Measure the downlink constraint

Downlink is a portfolio of contact windows, spectrum, antenna capacity, coding, weather, ground stations, network routes and scheduling rights. Installed radio rate does not equal daily delivered capacity. The model should use achieved throughput after geometry, availability, overhead, contention and service rules.

The ITU Earth Exploration-Satellite Service handbook explains that link choice depends on mission data rate, equipment and link performance, and discusses S, X and Ka-band use. Relevant ITU recommendations address data-downlink systems and spectrum-sharing conditions. These sources support engineering diligence; spectrum access and ground capacity for a specific project require licences and contracts. [12][13][14]

Measure raw bytes collected, bytes retained onboard, bytes selected, bytes transmitted, bytes successfully received and bytes accepted by customers. Orbital inference creates economic bandwidth only when the reduced data stream preserves the evidence the customer needs. A low downlink ratio achieved by discarding useful observations destroys value.

## 6. Separate filtering from inference

Filtering makes a bounded decision about whether data meets a transmission criterion. Cloud detection is a common example because optical imagery obscured by cloud can be unsuitable for some applications. Inference estimates a property or event, such as vessel presence or fire probability. The second task usually has a wider error surface and more consequential customer use.

ESA describes PhiSat-2 applications that discard cloud-obscured images, transform images into street maps, detect vessels, compress images and detect wildfire and marine pollution. This breadth shows that one spacecraft can host distinct workloads. It also reinforces the need for workload-specific evidence, acceptance thresholds and change controls. [1][2][4]

An investment case should avoid aggregating all workloads into one accuracy number. Each model should have a defined population, benchmark, threshold, uncertainty treatment, version and operational envelope. Revenue should follow the subset of workloads that have passed acceptance.

## 7. Define acceptable accuracy

Technical accuracy metrics must connect to the customer decision. Precision matters when false alerts consume scarce response resources. Recall matters when missed events create loss. Localization error matters when the user must act at a particular coordinate. Calibration matters when a confidence score drives an automated threshold.

Acceptance should use representative sensor data, seasons, geographies, atmospheric conditions and operational noise. A curated development set can overstate field performance. Testing should reserve data outside model training and tune thresholds before the contractual test begins.

The service should report performance by relevant cohort. A model may perform differently over ocean, desert, snow, urban areas or low-light scenes. Cohort evidence informs routing: uncertain cases can preserve raw data for ground review, while well-supported cases can receive automated treatment.

## 8. Price false positives and missed events

Model errors have asymmetric commercial effects. A false positive can trigger an unnecessary inspection, consume analyst time or create a misleading customer report. A missed event can delay response, breach a service level or undermine trust. The economic model should assign an expected cost to each error category using customer evidence.

The contract should define whether the output is a screening signal, decision support or an autonomous instruction. It should also define required human review, disclaimers, audit evidence and limits on permitted use. These provisions influence liability and insurance.

Threshold selection is therefore a commercial decision as well as a technical one. A conservative threshold can reduce false alerts while missing more events. A permissive threshold can increase recall while burdening the customer. Product governance should record who approves the threshold and which evidence supports it.

## 9. Build customer acceptance

Customer acceptance converts technical capability into financeable evidence. The protocol should specify a test window, representative locations, reference data, required metrics, latency measurement, sample size, exception handling and sign-off authority. Payment should follow accepted output where possible.

An early contract can separate development, pilot and operating phases. Development revenue may fund integration. Pilot revenue can support representative testing. Operating revenue should begin only after the service passes objective acceptance. Minimum-volume, reservation or availability payments can improve cash visibility when they are enforceable and supported by customer credit.

Retention depends on workflow integration. A fast alert that requires manual reformatting or lacks provenance may not be operationally useful. Diligence should observe the output inside the customer's process and verify renewal, expansion and collected cash.

## 10. Underwrite sensor and data quality

Model performance begins with the sensor. Resolution, spectral bands, signal-to-noise ratio, calibration, pointing, illumination and atmospheric conditions shape the input. A model trained on one sensor or processing level may not transfer to another without validation.

ESA reported that PhiSat-2 entered its science phase after commissioning and provides imagery from a multispectral instrument at approximately five-metre ground sampling distance. ESA also changed the mission approach to provide Level 1 data products for training and application development, alongside onboard processing. That decision illustrates the continuing value of accessible source data and ground experimentation. [4]

The transaction data room should include calibration history, sensor anomalies, preprocessing versions, representative imagery, ground-truth methodology and model lineage. Customer acceptance should identify the processing level it relies upon. Any change in sensor or preprocessing should trigger a defined revalidation process.

## 11. Underwrite compute hardware

The processor must deliver the required workload within mass, power, thermal, storage and radiation constraints. Nameplate operations per second provide limited evidence. Investors need sustained performance on the flight configuration, using the actual model, input size and software stack.

Qualification should cover vibration, thermal vacuum, electromagnetic compatibility, radiation response, memory errors, storage integrity, reset behaviour and fault recovery. Evidence from a commercial processor used elsewhere can inform diligence; integration with the current board, spacecraft and workload still requires testing.

NASA's Small Spacecraft Technology State of the Art reports provide a framework for evaluating avionics and onboard computing options. NASA and commercial demonstrations also show that capable processors can operate in space. Financing should follow configuration-specific evidence and the ability to replace or work around failures. [6][7]

## 12. Model power and thermal limits

Orbital inference competes with payload, communications, attitude control and platform systems for energy. The power model should use orbit-by-orbit generation, eclipse, battery state, conversion loss, degradation, reserve and safe-mode requirements. Average annual consumption can hide critical short-duration constraints.

Compute turns electrical energy into heat. In vacuum, heat must be conducted and radiated. Sustained inference can therefore be limited by temperature even when a processor completes a short benchmark. Thermal-vacuum testing should use representative duty cycles and worst-case environmental assumptions.

Revenue capacity should be expressed as accepted workloads after power and thermal constraints. The model should show which captures receive inference when energy is scarce, how the queue is prioritized and what customer commitments can still be met.

## 13. Account for radiation and corruption

Radiation can produce transient errors, resets, memory corruption and permanent damage. Commercial hardware may tolerate an environment differently from radiation-hardened components. Shielding, redundancy, error-correcting memory, watchdogs, checkpointing and safe-mode recovery can reduce risk while adding mass, power or complexity.

NASA described an intelligent edge-processing system tested aboard the International Space Station and reported that the returned hardware remained functional after months in orbit without observed adverse radiation effects in that unit. This provides relevant demonstration evidence; it is not a general qualification for every orbit, duration or configuration. [8]

The AI service needs end-to-end integrity controls. A corrupted model weight or input can produce a plausible but wrong result. Secure hashes, signed model packages, redundant storage, health checks, known-answer tests and ground comparison can support assurance. The operating model should include investigation, rollback and customer notification.

## 14. Manage model compression

Flight resources can require quantization, pruning, distillation or architecture changes. These techniques may reduce memory, power and execution time. They can also change accuracy, confidence calibration and performance on rare cases. The flight model should therefore be treated as its own controlled product.

Benchmark the ground model and flight model on the same representative test set. Report performance by cohort and resource consumed. A smaller model that completes reliably within the operational envelope can have more commercial value than a larger model that exceeds power, thermal or timing limits.

Compression should remain reversible through configuration management. The data room should preserve source model, training code where available, conversion tools, compiler versions, test evidence and signed flight binary. A transaction buyer should be able to reproduce the accepted build.

## 15. Control model updates

An updatable model can adapt to new geographies, seasons, sensors and customer needs. ESA states that additional PhiSat-2 applications can be uploaded in orbit. Update capability increases the platform's commercial life and also creates cyber, regression and governance risk. [5]

Every update should have an approved purpose, version, test record, rollback path, cryptographic signature and deployment window. The spacecraft should reject unauthorized packages and retain a known-good configuration. Ground operators should be able to prove which version produced a customer output.

Revenue forecasts should distinguish contracted existing models from proposed future applications. A flexible platform supports option value. That option belongs outside the base case until a customer, acceptance protocol and delivery plan are identified.

## 16. Govern training data and provenance

Training data should be traceable to lawful sources, licences and processing steps. Labels should record creator, method, review and uncertainty. Synthetic or augmented data should be identified. Provenance supports model reproduction, dispute analysis and customer trust.

Geospatial data can encode seasonal, geographic and sensor-specific bias. The training set should be compared with the operating population. Gaps should inform sampling, threshold limits and human review. Continuous ground truth can be expensive, so the commercial model should fund validation throughout the service life.

NASA's Prithvi work illustrates the scale of geospatial foundation-model training and the possibility of deploying derived capability in orbit. A foundation model can reduce development time for downstream applications. Each downstream task still requires acceptance on the target sensor, environment and customer decision. [10]

## 17. Protect cyber and command paths

The attack surface includes model development, build systems, update servers, mission control, communications links, onboard storage, processor interfaces, ground stations, customer APIs and operators. A compromised model or command path can alter collection, hide events, leak data or degrade the spacecraft.

NIST's zero-trust architecture emphasizes explicit trust decisions, strong identity and protection of resources rather than reliance on network location. NIST's AI Risk Management Framework organizes AI risk activity around govern, map, measure and manage. These frameworks can structure controls; the satellite system requires implementation evidence. [15][16]

Diligence should test key custody, role separation, signed commands, secure boot, software bill of materials, vulnerability response, logging, anomaly detection, incident recovery and supplier access. Cyber controls should extend to customer delivery and evidence retention.

## 18. Address remote-sensing regulation

Remote-sensing systems can require authorization, operating conditions and continuing supervision. In the United States, NOAA's Office of Space Commerce administers commercial remote-sensing licensing under 15 CFR Part 960. Licence conditions, national-security controls, shutter requirements, data restrictions and foreign ownership issues can affect the service perimeter and a transaction. [20][21]

The legal workstream should identify the satellite registry, licensing jurisdiction, payload operator, data controller, ground-station jurisdictions, customer locations and permitted distribution. It should also review export controls, sanctions, privacy and sector-specific uses. Advice must be specific to the actual structure.

Orbital inference can change regulatory facts because some raw data may never reach the ground, while alerts or derived products cross borders. Contracts and operating procedures should define retention, deletion, regulator access and evidentiary records. A buyer should confirm that change of control and model updates do not invalidate approvals.

## 19. Secure spectrum and ground access

Spectrum rights and ground access determine when the satellite can transmit. The project should maintain a schedule of filings, authorizations, coordination obligations, assigned frequencies, ground-station agreements, service levels and renewal dates. Capacity assumptions should trace to executed rights.

An orbital-inference design may reduce volume and change the traffic profile. It can also increase the value of low-latency contacts because an alert should reach the customer quickly. The network model should test priority service, alternate stations, weather diversity, outages and congestion.

Ground-as-a-service contracts can convert infrastructure into operating cost. They can create concentration and repricing exposure. Diligence should assess provider credit, technical compatibility, data security, termination assistance and the cost of moving to another network.

## 20. Construct the revenue model

Revenue can arise from subscriptions, reserved capacity, accepted scenes, accepted alerts, application licences, analytics, data resale or mission-service agreements. Each stream should have a unit, price, volume, acceptance condition, collection period and cost to serve.

Avoid counting the same value twice. A customer paying for an alert may not also pay for every discarded image. A partner's downstream revenue share should be removed before calculating project cash. Revenue from pilots, grants and integration should be separated from repeatable operating revenue.

The base case should use executed contracts and collected cash where available. Qualified pipeline can support a scenario with explicit probability and conversion timing. General market size, total imagery volume and possible future applications should remain outside the contracted case.

## 21. Model operating and lifecycle cost

Operating cost includes mission control, ground network, cloud systems, model monitoring, customer support, regulatory compliance, insurance, cyber security, data validation and corporate overhead. Lifecycle cost includes software maintenance, processor refresh, constellation replenishment, decommissioning and customer migration.

Onboard inference can reduce ground transfer and processing. It adds flight hardware, integration, qualification, model operations and update controls. The comparison should use total cost per accepted customer result over a full replacement cycle.

Cost estimates should be supported by supplier quotations, employment plans, service contracts and historical operating data. Contingency should reflect maturity and concentration. Currency, inflation and launch-price assumptions require separate sensitivities.

## 22. Compare ground and orbital processing

Ground processing benefits from abundant energy, replaceable hardware, flexible models and direct access to data platforms. Orbital processing benefits from proximity to the sensor, earlier selection and the ability to act before contact. A hybrid design often preserves raw or sampled data while sending urgent derived products first.

The comparison should hold the customer outcome constant. Calculate elapsed time, accepted accuracy, transmitted volume, operating cost, hardware cost, refresh frequency and failure recovery for each architecture. Include the cost of the validation data needed to monitor the orbital model.

Orbital inference should win a workload only when its advantage exceeds added risk and lifecycle cost. Workloads without a binding downlink or time constraint generally retain stronger ground economics.

## 23. Underwrite launch and deployment

Incremental compute affects mass, volume, power, thermal design, integration and test. A hosted payload, dedicated CubeSat or processor on an existing bus has different control and cost. The budget should include engineering, qualification, launch integration, launch, early operations and commissioning.

Launch price alone is not delivered cost. Schedule slippage can delay customer revenue and create storage or requalification cost. A multi-satellite plan also faces cohort risk if the same design flaw affects every unit.

Capital release should follow build completion, environmental qualification, launch integration, successful deployment, processor checkout, model execution, communications and customer acceptance. Insurance proceeds should align with the replacement plan and debt service needs.

## 24. Fund refresh and replacement

Models, processors, sensors and communications can become obsolete at different rates. The economic life should use the shortest binding component or a credible upgrade path. Terminal value should reflect funded replenishment rather than perpetual use of the initial constellation.

A replacement schedule can stagger cohorts to preserve coverage and incorporate improvements. It should also maintain compatibility, validation and customer continuity. Shared launch can reduce cost and introduce schedule dependency.

The financial model should reserve cash before each replacement cycle. Debt maturity, amortization and covenants should fit the demonstrated service life. A platform requiring frequent unfinanced equity is less suitable for operating leverage.

## 25. Allocate service and completion risk

Development, satellite, launch, ground, software and customer parties control different risks. Contracts should allocate design responsibility, interface performance, delay, defects, data loss, cyber incidents, service credits and termination. Liability caps and indemnities should be assessed against plausible loss.

Completion should mean that the system delivers the contracted output, rather than simply reaching orbit. Milestones can include environmental qualification, commissioning, model benchmark, latency test, output provenance, customer acceptance and initial payment.

Sponsor support, holdbacks, warranties, performance security and contingency can cover defined gaps. Their value depends on enforceability and counterparty credit. A financing structure should leave residual risk with the party best able to manage it.

## 26. Define the financing perimeter

Early development suits sponsor equity, grants, strategic capital and customer co-funding. Milestone equity can fund qualification and launch. Asset or project debt becomes more credible after in-orbit performance, customer acceptance and repeatable collection.

The financed entity should control essential intellectual property, spacecraft rights, licences, ground access, customer contracts and operating data. Related-party services need arm's-length terms, continuity rights and lender access. Security should cover assets that can lawfully and practically be transferred.

Revenue concentration, constellation concentration and technology refresh constrain leverage. Covenants can track accepted service, cash collection, coverage, availability, cybersecurity, insurance, replacement reserves and licence compliance.

## 27. Value bandwidth reduction

Bandwidth reduction has three possible values: avoided transmission expense, increased useful data within fixed capacity and earlier delivery of high-priority information. The first can be measured against ground-network charges. The second requires evidence that displaced data had lower value. The third requires customer willingness to pay for priority.

Calculate the ratio of raw data to transmitted data, then adjust for retained utility, validation samples, retransmission and false decisions. A ninety-percent reduction with significant missed events can be worse than a fifty-percent reduction that preserves customer value.

Only incremental cash belongs in valuation. Avoided cost should exclude fixed commitments that remain payable. Added capacity should be monetized only when demand and acceptance exist.

## 28. Value decision-speed improvement

Decision speed can increase revenue, reduce loss or improve resource allocation. Estimate the event population, probability that earlier information changes action, economic consequence, project share of value and contractual capture mechanism.

The model should compare orbital inference with faster ground contacts, more ground stations, inter-satellite links and scheduled priority. The cheapest reliable alternative sets an economic ceiling on the premium.

Evidence can include customer pilots, response logs, renewal, price tests and post-event analysis. Strategic language without a defined decision and payment mechanism should receive no value premium.

## 29. Value sensor-utilization gain

Onboard screening can free storage and downlink, allowing more useful captures. Autonomous targeting can redirect observation toward a detected feature. The gain should be measured as additional accepted products per orbit or per unit of sensor time.

Retasking creates opportunity cost because the satellite may abandon another observation. The scheduler should rank customer priority, contractual penalties, confidence and orbital opportunity. Governance should prevent one workload from degrading higher-value commitments.

Value follows incremental accepted revenue or avoided collection cost. Gross captures, processor utilization and detections remain operating metrics until connected to customer acceptance.

## 30. Apply evidence-weighted valuation

Build the valuation from contracted cash, qualified operating capacity and funded lifecycle. Discounted cash flow should use accepted revenue, probability-adjusted ramp, operating cost, replacement capital, working capital, tax and financing. Comparable transactions can provide a reasonableness check when business model, maturity and contract quality are genuinely comparable.

Stage adjustments should be explicit. A laboratory benchmark, representative hardware test, flight demonstration and paid operating service provide different evidence. The GAO Technology Readiness Assessment Guide emphasizes credible, objective and useful evaluation of maturity before integration into acquisition programmes. [17]

The model should show a bridge from unadjusted present value to value after remaining capital, adoption probability, technical performance, concentration and lifecycle risk. Reviewers can then challenge each adjustment directly.

## 31. Construct the hypothetical case

The hypothetical constellation has twelve satellites and collects forty-eight terabytes of raw data per day. Baseline downlink capacity is twelve terabytes per day. The orbital architecture selects three terabytes per day of priority imagery, image chips, alerts and validation samples. These figures are assumptions and do not describe an operating company.

Baseline annual revenue is USD 32 million. Accepted faster alerts and improved product availability add USD 20.8 million. Avoided transmission and ground-processing cost is USD 6 million. Incremental operating and lifecycle cost is USD 13 million, producing USD 13.8 million of incremental annual operating cash before tax and financing.

Incremental development and deployment capital is USD 68 million. A hypothetical unadjusted present value of USD 118 million is reduced by USD 24 million of remaining capital, USD 30 million of adoption and performance adjustment and USD 18 million of replacement and concentration adjustment. The resulting illustrative enterprise-value uplift is USD 46 million. No external investor should rely on these assumptions.

## 32. Stress the economics

The case should be stressed for lower customer adoption, weaker model performance, delayed launch, lost spacecraft, lower downlink cost, higher ground-station cost, processor degradation, shortened useful life and faster competitor response. Correlated design defects deserve a separate case.

Break-even can be expressed as accepted incremental revenue plus avoided cash cost required to cover operating cost, replacement reserve and return on incremental capital. This makes the investment threshold visible to commercial and engineering teams.

The downside should assume that raw data must be preserved for review, reducing bandwidth savings. It should also test customer refusal to accept autonomous outputs. A recovery plan can move processing to ground, limit service to filtering or replace a cohort.

## 33. Define the investment decision

An investable orbital-inference platform should solve a named customer problem, demonstrate performance on representative flight data, control update and cyber risk, hold required licences, secure communications capacity, deliver objective acceptance and fund replacement. The board should approve the use-case perimeter and the evidence required for each expansion.

The transaction structure should release capital against observable gates. Early capital funds integration and qualification. Deployment capital follows completed tests and launch readiness. Operating leverage follows commissioned service, customer acceptance, collected cash and reserves.

The final decision is a comparison of evidence-weighted value with required capital and downside protection. The model should remain auditable from customer outcome through technical performance to cash flow.

## Conclusion

Orbital AI inference can improve Earth-observation economics when it converts scarce downlink and elapsed time into accepted customer value. Its strongest uses screen unusable data, compress information without losing utility, detect time-sensitive events and retask assets within a validated operating envelope.

Technical demonstrations by ESA and NASA establish credible pathways for onboard processing, uploadable applications and closed-loop targeting. They do not establish the commercial result for a specific company. Investors and counterparties should require configuration-specific tests, customer acceptance, regulatory compliance, cyber controls and funded replacement.

The disciplined valuation unit is the accepted decision product. Bandwidth saved, milliseconds reduced, detections produced and compute installed support that product. Contracted revenue, avoided cash cost and verified lifecycle performance determine value.

## Appendix A. Customer use-case checklist

Identify the customer decision, location, event population, observation frequency, useful time window, output format, minimum confidence, false-positive cost, missed-event cost, human review, raw-data requirement, audit evidence, payment unit, acceptance authority, renewal condition and permitted use. Record the ground-processing alternative and the reason orbital execution changes value.

## Appendix B. Technical evidence checklist

Collect the sensor specification, representative input set, preprocessing lineage, flight processor configuration, model binary, benchmark, environmental tests, power profile, thermal profile, radiation analysis, fault-recovery test, storage integrity, communications plan, ground capacity, commissioning plan and configuration-control record. Link every base-case performance assumption to evidence.

## Appendix C. Model assurance checklist

Document training sources, licences, labels, cohort coverage, reserved test data, model version, threshold, confidence calibration, compression method, signed build, update approval, rollback, drift monitoring, ground comparison, incident response and customer notification. Define the cases routed to human review or raw-data preservation.

## Appendix D. Contract and regulatory checklist

Review satellite and payload licences, registry, spectrum, ground-station rights, export control, sanctions, privacy, data rights, model rights, customer acceptance, service credits, liability, insurance, change of control, termination, step-in, replacement and decommissioning. Obtain jurisdiction-specific legal advice.

## Appendix E. Valuation checklist

Separate contracted revenue, qualified pipeline and future options. Verify unit price, accepted volume, collection, cost to serve, avoided cash cost, launch and integration capital, replacement reserve, concentration, useful life, tax and financing. Present an unadjusted value bridge and explicit technical, adoption and lifecycle adjustments.

## Appendix F. Decision figures and tables

**Figure 1. Hypothetical delivered incremental hardware cost curve**

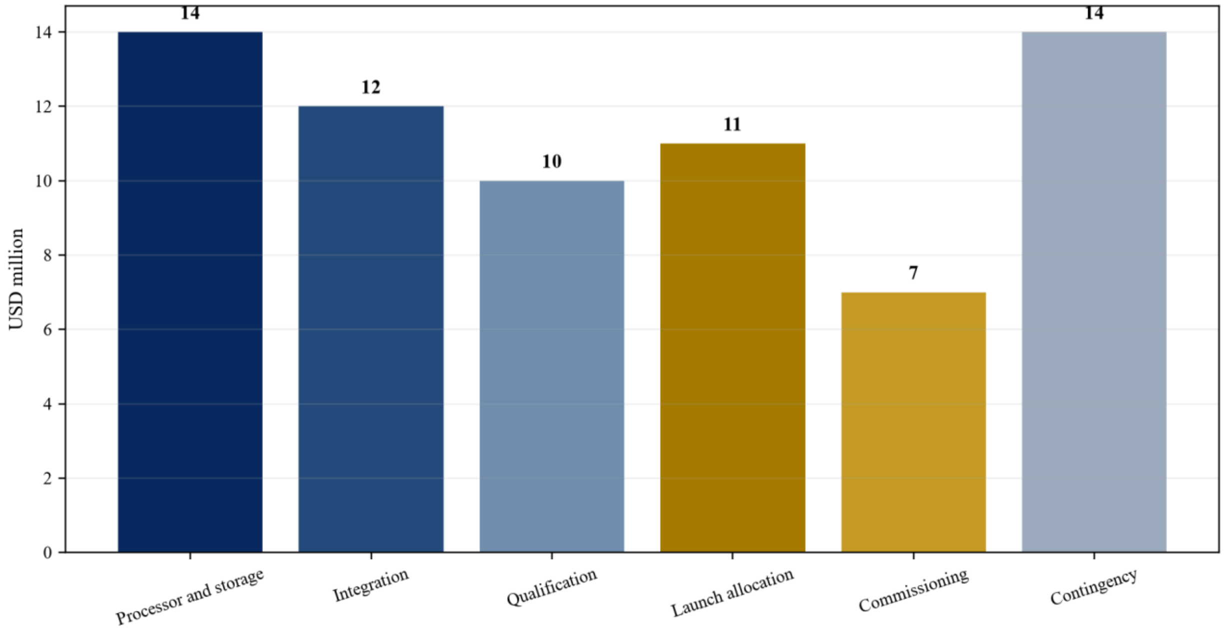


*Wholly hypothetical; USD million for a twelve-satellite constellation.*

**Table 1. Workload and customer-value fit**

| Workload | Primary benefit | Key acceptance measure | Principal risk |
|---|---|---|---|
| Cloud filtering | avoid unusable downlink | retained usable scenes | false rejection |
| Compression | reduce transmitted bits | task utility after reconstruction | lost detail |
| Event detection | faster actionable alert | precision, recall and latency | decision error |
| Autonomous targeting | improve collection | incremental accepted observations | displaced tasking |

*Proposed screening framework.*

**Figure 2. Hypothetical orbital power allocation**

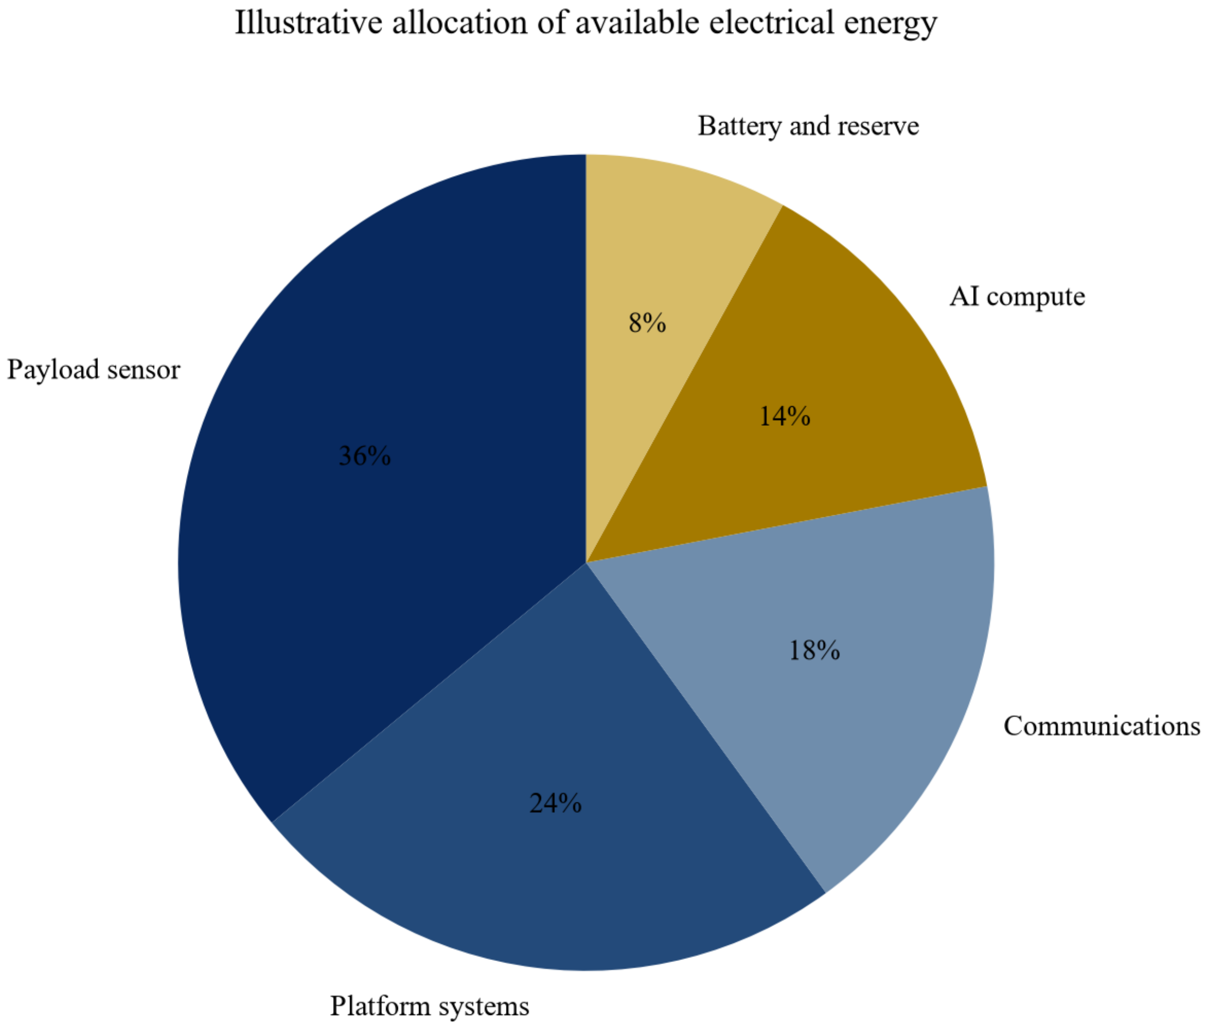


*Wholly hypothetical average available power allocation.*

**Table 2. Evidence maturity and financing gate**

| Evidence stage | Minimum evidence | Indicative capital |
| --- | --- | --- |
| Benchmarked | representative data and frozen metric | sponsor and grant capital |
| Qualified | flight configuration passes environmental tests | milestone equity |
| Commissioned | system operates in orbit | strategic or structured capital |
| Accepted | customer signs objective test and pays | potential operating finance |
| Repeatable | renewal, collection and replacement record | broader debt and acquisition perimeter |

*Proposed capital-release framework.*

**Figure 3. Observation-to-alert latency map**

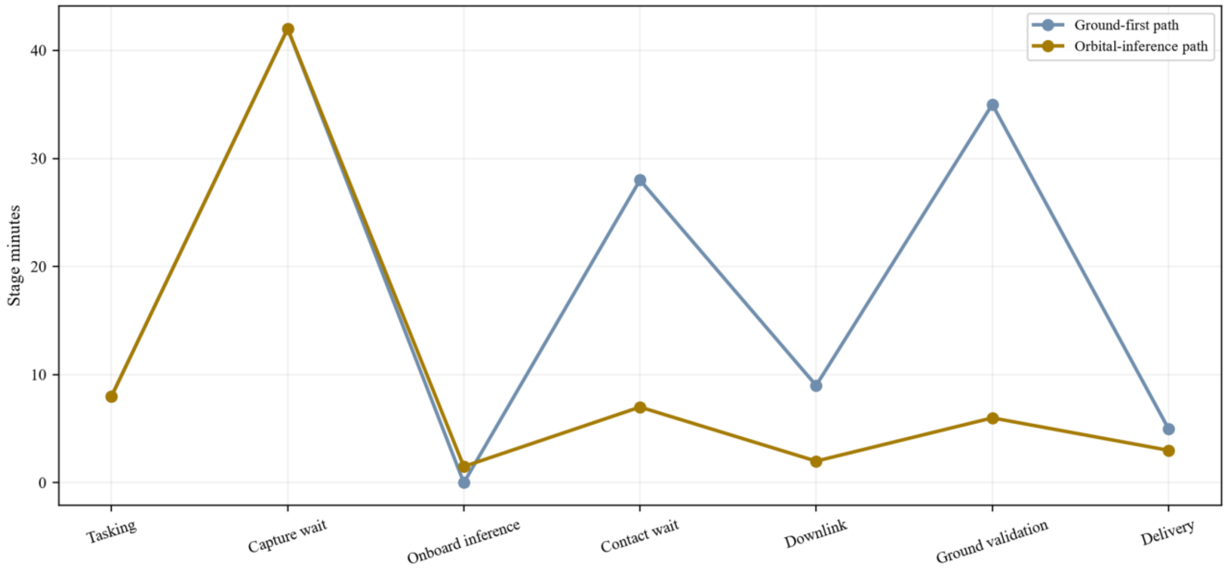


*Wholly hypothetical median minutes; customer validation is included.*

**Table 3. Data-product acceptance matrix**

| Product | Required source evidence | Acceptance metric | Retained review data |
|---|---|---|---|
| Priority image | calibrated scene | delivery time and usability | full selected scene |
| Vessel alert | image chip and coordinates | precision, recall and latency | chip plus sampled negatives |
| Wildfire alert | thermal or optical evidence | event confirmation and time | source window and model record |
| Change layer | paired observations | spatial accuracy and completeness | both observations |

*Proposed contract controls.*

**Figure 4. Hypothetical twelve-satellite refresh schedule**

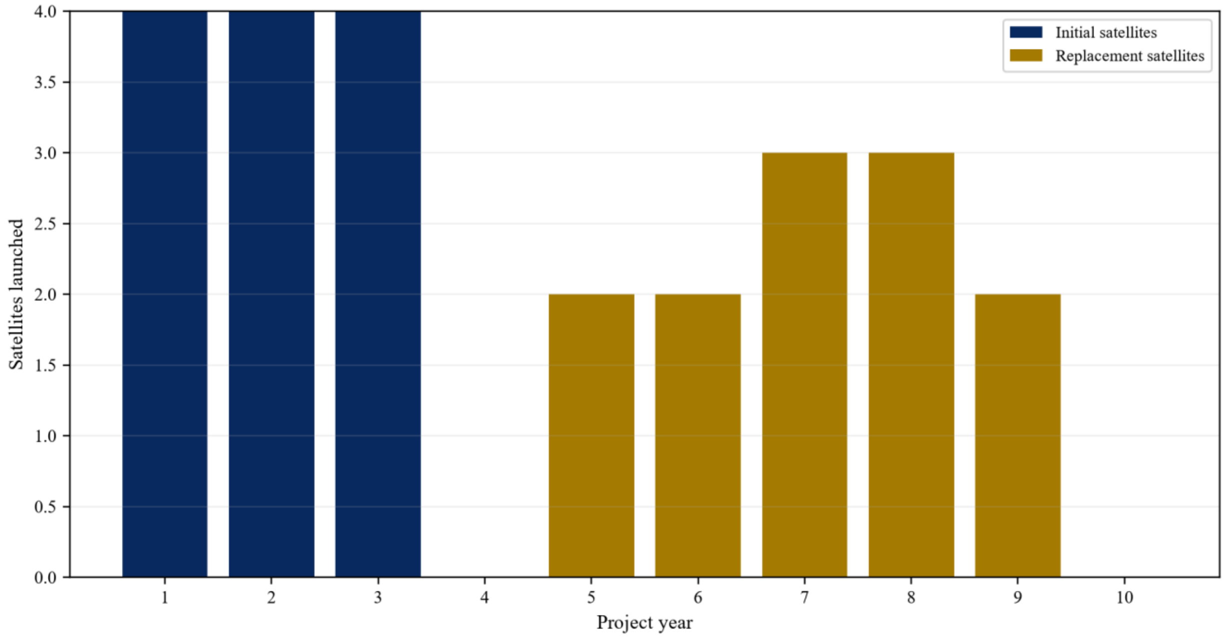


*Wholly hypothetical launches by project year.*

**Table 4. Hypothetical incremental operating economics**

| Metric | Baseline | Orbital-inference increment |
|---|---|---|
| Raw observations, TB per day | 48 | 48 processed locally |
| Downlinked data, TB per day | 12 | 3 selected output |
| Accepted annual revenue | 32 | 20.8 additional |
| Avoided annual cash cost | 0 | 6.0 |
| Incremental annual operating and lifecycle cost | 0 | 13.0 |
| Incremental annual operating cash | 0 | 13.8 |

*Wholly hypothetical; USD million per year unless stated.*

**Figure 5. Hypothetical data-volume waterfall**

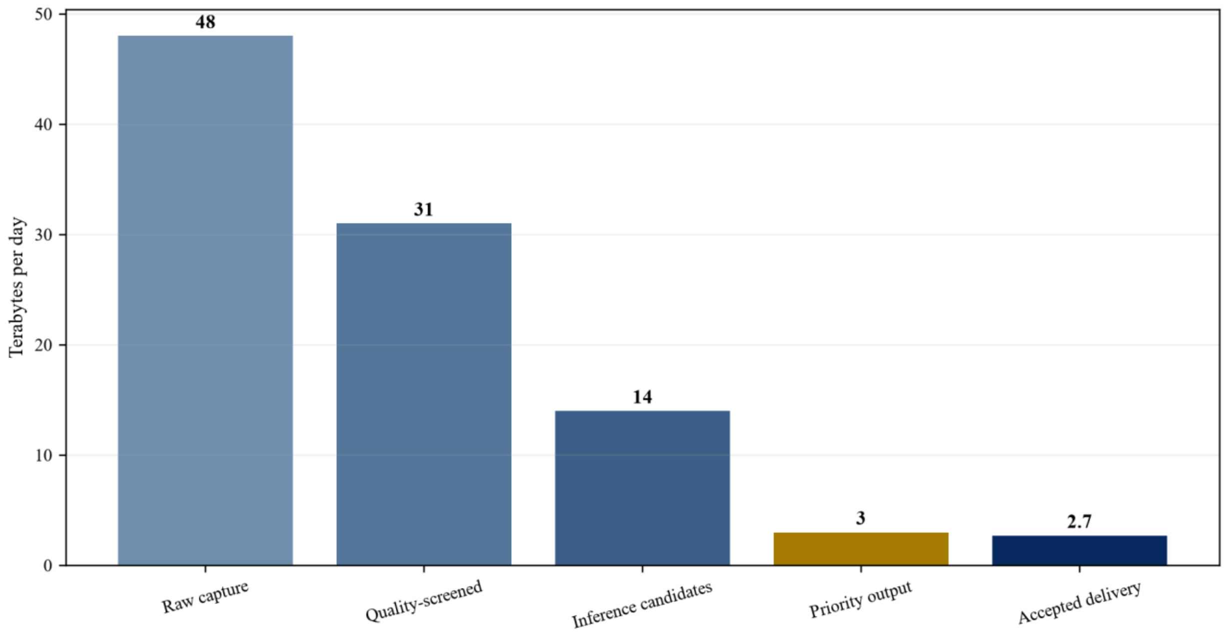


*Wholly hypothetical terabytes per day.*

**Table 5. Model and operating risk controls**

| Risk | Preventive control | Detective control | Recovery |
|---|---|---|---|
| Sensor shift | approved operating envelope | cohort performance | preserve raw data and recalibrate |
| Model regression | frozen acceptance suite | ground comparison | rollback signed version |
| Corruption | secure build and redundant storage | known-answer test | restore known-good package |
| Cyber intrusion | strong identity and signed commands | immutable event logs | isolate, revoke and recover |
| Power constraint | workload budget and priority | battery and thermal telemetry | defer low-value workload |

*Proposed control matrix.*

**Figure 6. Hypothetical enterprise-value bridge**

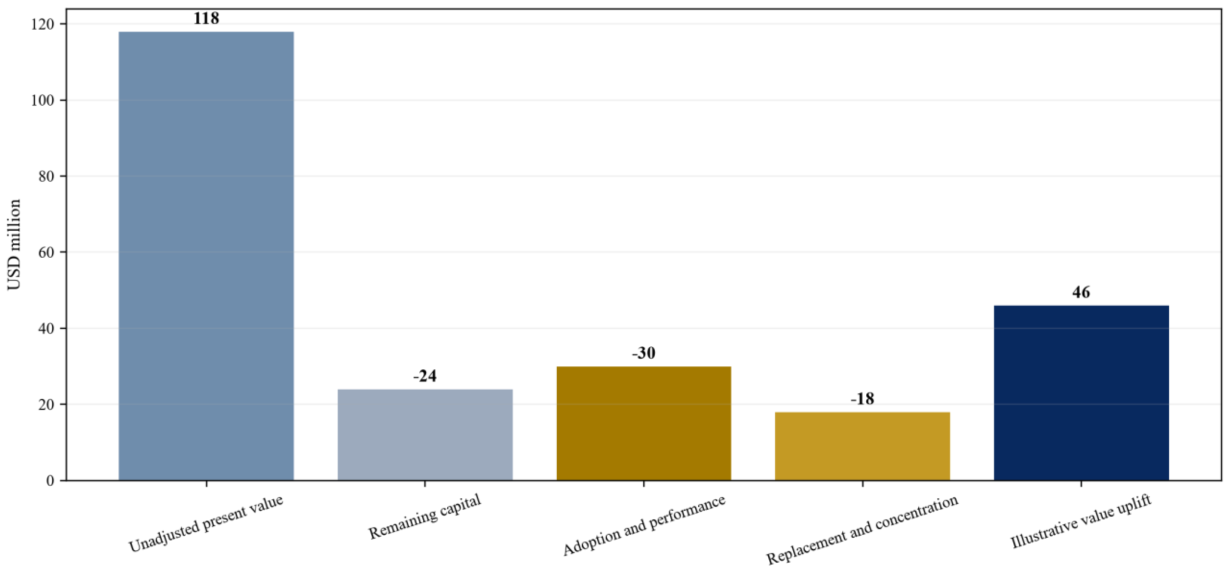


*Wholly hypothetical; USD million.*

**Table 6. Transaction diligence workstreams**

| Workstream | Principal question | Decision evidence |
|---|---|---|
| Customer | who pays for which decision product | executed contract, acceptance and collection |
| Technical | can the flight configuration sustain service | representative tests and telemetry |
| Model | is performance controlled through lifecycle | provenance, benchmark, monitoring and rollback |
| Regulatory | can the system and data service operate | licences, spectrum rights and legal advice |
| Financial | can cash fund operation and replacement | controlled model, reserves and downside |

*Proposed scope.*

**Table 7. Investment scorecard**

| Criterion | Investable evidence | Red line |
|---|---|---|
| Customer outcome | accepted decision product and paid contract | technology without workflow value |
| Downlink | measured constrained capacity and useful reduction | discarded customer evidence |
| Accuracy | representative cohort results | unbounded error or missing benchmark |
| Lifecycle | funded updates and replacement | unsupported terminal value |
| Control | secure version, provenance and recovery | unauthorized or irreproducible model |

*Proposed board framework.*

## Frequently asked questions

Q: What is orbital AI inference? A: It is the execution of a trained model on or near an observing spacecraft so that data can be filtered, compressed, classified or used for targeting before the full data set is downlinked.

Q: When does onboard inference create the strongest economic value? A: The strongest case arises when downlink is constrained, information loses value quickly, the output can be tested objectively and a customer pays for the faster or more selective result.

Q: Does lower data volume automatically create value? A: No. The reduction must preserve the evidence and utility required by the customer. Discarding useful imagery can create missed events, disputes and lost revenue.

Q: How should a customer accept an orbital AI product? A: Use a representative test window, objective metrics, defined latency, retained source evidence, exception rules and authorized sign-off. Payment should follow accepted output where practical.

Q: Can commercial processors be financed before flight? A: Development and milestone capital may support pre-flight work. Operating leverage generally requires configuration-specific qualification, in-orbit commissioning, customer acceptance and repeatable collection.

Q: How should model updates be controlled? A: Each update should have an approved purpose, version, representative test, signed package, deployment record, rollback path and evidence identifying which version produced each output.

Q: What belongs in the orbital-inference valuation? A: Contracted incremental revenue, cash cost avoided, remaining capital, operating cost, replacement, customer adoption, technical performance and concentration. Future applications should remain outside the base case until supported by evidence.

Q: What does the hypothetical case demonstrate? A: It demonstrates how bandwidth reduction, accepted revenue, avoided cost, operating expense, remaining capital and specific risk adjustments can be connected in one auditable decision model. The figures do not describe an identified company.

## References

**[1]** European Space Agency, Introducing PhiSat-2. https://www.esa.int/Applications/Observing_the_Earth/FutureEO/Phsat-2/Introducing_Phsat-2

**[2]** European Space Agency, PhiSat-2 Mission. https://www.esa.int/Applications/Observing_the_Earth/FutureEO/Phsat-2

**[3]** National Aeronautics and Space Administration, How NASA Is Testing AI to Make Earth-Observing Satellites Smarter. https://www.nasa.gov/science-research/earth-science/how-nasa-is-testing-ai-to-make-earth-observing-satellites-smarter/

**[4]** European Space Agency, PhiSat-2 Begins Science Phase for AI Earth Images. https://www.esa.int/Applications/Observing_the_Earth/FutureEO/Phsat-2/Phsat-2_begins_science_phase_for_AI_Earth_images

**[5]** European Space Agency, PhiSat-2 Gets Two New AI Apps. https://www.esa.int/Applications/Observing_the_Earth/Phsat-2/Phsat-2_gets_two_new_AI_apps

**[6]** National Aeronautics and Space Administration, Small Spacecraft Avionics State of the Art 2026. https://www.nasa.gov/wp-content/uploads/2026/05/8-smallsat-avionics-2026-final.pdf

**[7]** National Aeronautics and Space Administration, Technology Tested in Space Is Preparing Us for the Moon and Mars. https://www.nasa.gov/missions/station/technology-tested-in-space-is-preparing-us-for-the-moon-and-mars/

**[8]** National Aeronautics and Space Administration Spinoff, Intelligent Processing at the Edge. https://spinoff.nasa.gov/node/11550

**[9]** National Aeronautics and Space Administration, New Research Launching to Space Station aboard Northrop Grumman's Fifteenth Resupply Mission. https://www.nasa.gov/missions/station/iss-research/new-research-launching-to-space-station-aboard-northrop-grummans-15th-resupply-mission/

**[10]** National Aeronautics and Space Administration, First Geospatial Foundation Model Operates in Orbit. https://science.nasa.gov/science-research/ai-foundation-model-in-orbit/

**[11]** National Aeronautics and Space Administration TechPort, Quantum Machine Learning for Space-Based Earth Observation. https://techport.nasa.gov/projects/106608

**[12]** International Telecommunication Union, Handbook on Earth Exploration-Satellite Service. https://www.itu.int/dms_pub/itu-r/opb/hdb/R-HDB-56-2011-PDF-E.pdf

**[13]** International Telecommunication Union, Recommendation ITU-R SA.2155. https://www.itu.int/rec/R-REC-SA.2155-0-202212-I/en

**[14]** International Telecommunication Union, Recommendation ITU-R RS.2105. https://www.itu.int/rec/R-REC-RS.2105/en

**[15]** National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework 1.0. https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10

**[16]** National Institute of Standards and Technology, SP 800-207 Zero Trust Architecture. https://csrc.nist.gov/pubs/sp/800/207/final

**[17]** United States Government Accountability Office, Technology Readiness Assessment Guide GAO-20-48G. https://www.gao.gov/products/gao-20-48g

**[18]** National Aeronautics and Space Administration Earthdata, EOSDIS Terminology Specification. https://www.earthdata.nasa.gov/s3fs-public/2024-02/ESDIS05100_423-SPEC-002_EOSDIS_Terminology_Spec-5.pdf

**[19]** National Aeronautics and Space Administration Earthdata, Land Atmosphere Near Real-Time Capability for EOS. https://www.earthdata.nasa.gov/data/tools/lance

**[20]** National Oceanic and Atmospheric Administration Office of Space Commerce, United States Commercial Remote Sensing Space Policy. https://www.space.commerce.gov.a.woc.noaa.gov/policy/u-s-commercial-remote-sensing-space-policy/

**[21]** National Oceanic and Atmospheric Administration, 15 CFR Part 960 Licensing of Private Remote Sensing Space Systems. https://www.nesdis.noaa.gov/s3/2021-08/15%20CFR%20Part%20960%20Regs%202020.pdf

**[22]** European Space Agency, PhiSat-1 Artificial Intelligence for Earth Observation. https://www.esa.int/Applications/Observing_the_Earth/FutureEO/Phsat/First_Earth_observation_satellite_with_AI_ready_for_launch

**[23]** European Space Agency, PhiSat-2 Mission Overview. https://challenges.philab.esa.int/wp-content/uploads/2023/02/Phisat-2_Mission_Overview_Web.pdf

**[24]** Consultative Committee for Space Data Systems, Lossless Data Compression Recommended Standard. https://public.ccsds.org/Pubs/121x0b3.pdf

**[25]** National Institute of Standards and Technology, Adversarial Machine Learning Taxonomy and Terminology. https://csrc.nist.gov/pubs/ai/100/2/e2025/final

**[26]** National Aeronautics and Space Administration, NASA Systems Engineering Handbook. https://www.nasa.gov/reference/systems-engineering-handbook/

**[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]** United Nations Office for Outer Space Affairs, Guidelines for the Long-term Sustainability of Outer Space Activities. https://www.unoosa.org/oosa/en/ourwork/topics/long-term-sustainability-of-outer-space-activities.html

## 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.