Earth-Observation M&A: Valuing Data Archives versus Sensor Fleets

A Transaction Framework for Historical Coverage Refresh Capability Customer Workflows Algorithms and Replacement Capital

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Earth-observation businesses combine assets with different economic lives. A historical archive can support change detection, model training and longitudinal analysis long after the originating satellite retires. A sensor fleet creates future refresh, tasking and new spectral coverage but requires replacement capital, launch access and continuing calibration. Algorithms, metadata, ground processing, application programming interfaces and customer workflows can convert imagery into decisions. An acquisition valuation should therefore separate archive cash, refresh capability, algorithm contribution and embedded workflow value before testing the benefits of owning them together.

This paper develops an Archive-Refresh-Workflow Valuation Test for Earth-observation M&A. It measures archive uniqueness through coverage, cadence, resolution, spectral consistency, metadata, rights, accessibility and observable customer use. It measures sensor fleets through usable collection capacity, revisit, latency, tasking, reliability, remaining life and replacement capital. It then tests algorithms and customer integrations through attributable adoption, accuracy, operating leverage, retention and switching cost. The framework prevents the same historical data from appearing simultaneously as an archive premium, model advantage, customer retention benefit and terminal-value uplift.

Primary evidence establishes the transaction context. NASA's Commercial Satellite Data Acquisition programme evaluates and procures commercial Earth-observation data across optical, radar, atmospheric and other modalities, while its discovery tools distinguish archive requests from future tasking. Planet describes a continuously refreshed proprietary dataset and a historical archive containing thousands of observations for many land locations, delivered through subscription and usage-based licences. BlackSky describes the combination of a sensor constellation, AI-enabled software and rapid monitoring workflows. These observations support the framework; they do not determine the value of an unidentified company. [1][2][3][4][5][6]

The worked case is wholly hypothetical and describes no identified company. A seller presents a USD 5.20 billion headline sum of archive, fleet, algorithm and customer values. The transaction model admits USD 1.20 billion of archive cash value, USD 1.65 billion of refresh capability, USD 0.65 billion of algorithm and platform value and USD 1.05 billion of customer-workflow value. It deducts USD 0.45 billion of overlap, USD 0.60 billion of replacement capital, USD 0.20 billion for data-rights and quality limitations and USD 0.15 billion of integration cost. A USD 0.30 billion verified combination benefit produces an illustrative value of USD 3.45 billion. Every amount and operating assumption must be replaced with transaction-specific evidence.

The framework is intended for boards, investors, lenders and transaction teams evaluating acquisitions, carve-outs, minority investments and strategic partnerships in Earth observation and geospatial intelligence. Its central output is a value bridge showing which cash depends on historical data, which

depends on future collection, which arises from software or workflow adoption and how much capital is required to preserve the combined proposition.

JEL Classification: G12, G31, G34, L96, O32, Q56

Keywords: earth observation M&A, satellite data archive valuation, sensor fleet valuation, geospatial analytics, refresh rate, satellite replacement capital, data moat, customer workflow, algorithm valuation, remote sensing transactions

Introduction

An Earth-observation company can continue to monetise historical imagery after individual spacecraft retire. It can also lose customer relevance quickly when refresh, latency, spectral coverage or tasking deteriorates. The archive and fleet therefore complement one another without being interchangeable.

The archive becomes valuable when it is legally usable, searchable, calibrated and connected to customer decisions. The fleet becomes valuable when it can produce required observations at the necessary cadence, quality and cost. Algorithms and workflows create additional value when their contribution can be separated from the underlying data and when customers repeatedly use the output.

This paper treats the target as four linked assets: archive, refresh capability, algorithms and workflows. Each receives an independent cash and evidence test before combination value is admitted.

1. Define the transaction decision

The analysis begins with the acquisition objective, valuation date, perimeter, decision users and intended operating model. Earth-observation value can migrate among data, hardware, software and customer use, so the unit of account must be explicit. The model should identify the legal right, technical dependency, user, cost and cash mechanism for every material component. This permits the team to test whether value survives a change in fleet, platform or ownership.

The evidence set comprises process materials, board mandate, target structure, financial statements and product architecture. Records should be dated, owned and reconciled to the relevant scene, sensor, product, contract and model assumption. Missing support remains a visible diligence gap. The financial model converts the verified record into one transaction question shared by technical, commercial and financial teams. The board should approve the decision and evidence cut-off before modelling value, with every adjustment linked to one ledger and one accountable owner.

2. Map the asset and rights perimeter

Decision usefulness requires a trace from archives, satellites, sensors, ground stations, processing, metadata, software, models, contracts, people and licences to customer action and cash. The trace should distinguish data collected in the past, capability that creates future observations, software that transforms data and workflow adoption that makes the output operational. It should also show where public or third-party substitutes limit exclusivity and pricing power.

Diligence should obtain asset registers, ownership records, data licences, IP schedules and intercompany agreements. The team should test completeness, rights, technical fitness and commercial use. Valuation then measures which assets and rights transfer and which remain dependencies. Decision makers should reconcile legal ownership with operational control; the assumption register should prevent the same observation from appearing as archive value, model advantage, retention benefit and terminal uplift.

3. Separate four value ledgers

The control question concerns archive cash, refresh capability, algorithms and customer workflows. It should be answered through a dated data, operating and customer record. Gross archive size, theoretical

revisit, model demonstrations and aggregate backlog can hide unusable scenes, cloud constraints, degraded sensors, unsupported algorithms and customers who do not rely on the product in production.

Required support includes product revenue, usage, cost, contracts and dependency maps. Exceptions, failed tests, missing rights and deteriorating usage deserve separate review. The model uses the evidence to calculate standalone value components with explicit overlap controls. The investment committee should keep each component separate until the final combination test, retaining deductions for incomplete rights, fleet replacement, remediation and integration.

4. Inventory the archive

The analysis begins with scenes, modalities, dates, geographies, resolution, spectral bands, processing levels and metadata. Earth-observation value can migrate among data, hardware, software and customer use, so the unit of account must be explicit. The model should identify the legal right, technical dependency, user, cost and cash mechanism for every material component. This permits the team to test whether value survives a change in fleet, platform or ownership.

The evidence set comprises catalogues, storage manifests, checksums, provenance and customer entitlements. Records should be dated, owned and reconciled to the relevant scene, sensor, product, contract and model assumption. Missing support remains a visible diligence gap. The financial model converts the verified record into a reproducible inventory of economically usable data. The board should exclude inaccessible duplicate or unsupported holdings, with every adjustment linked to one ledger and one accountable owner.

5. Measure historical coverage

Decision usefulness requires a trace from the geographic and temporal depth required for baselines, trends and change detection to customer action and cash. The trace should distinguish data collected in the past, capability that creates future observations, software that transforms data and workflow adoption that makes the output operational. It should also show where public or third-party substitutes limit exclusivity and pricing power.

Diligence should obtain scene footprints, acquisition dates, cloud masks, quality flags and application demand. The team should test completeness, rights, technical fitness and commercial use. Valuation then measures coverage density by customer use case. Decision makers should value observed usable history rather than gross file volume; the assumption register should prevent the same observation from appearing as archive value, model advantage, retention benefit and terminal uplift.

6. Measure archive uniqueness

The control question concerns the extent to which data cannot be replicated through public sources, competitors or new collection. It should be answered through a dated data, operating and customer record. Gross archive size, theoretical revisit, model demonstrations and aggregate backlog can hide unusable scenes, cloud constraints, degraded sensors, unsupported algorithms and customers who do not rely on the product in production.

Required support includes resolution, timing, modality, rights, event rarity and alternative datasets. Exceptions, failed tests, missing rights and deteriorating usage deserve separate review. The model uses the evidence to calculate incremental willingness to pay and avoided recreation cost. The investment committee should test substitutes before applying an archive premium, retaining deductions for incomplete rights, fleet replacement, remediation and integration.

7. Test data rights and restrictions

The analysis begins with ownership, end-user licence, derivative rights, government restrictions, privacy, export control and redistribution. Earth-observation value can migrate among data, hardware, software and customer use, so the unit of account must be explicit. The model should identify the legal right, technical dependency, user, cost and cash mechanism for every material component. This permits the team to test whether value survives a change in fleet, platform or ownership.

The evidence set comprises contracts, licences, consent records, jurisdiction and customer terms. Records should be dated, owned and reconciled to the relevant scene, sensor, product, contract and model assumption. Missing support remains a visible diligence gap. The financial model converts the verified record into cash that the buyer can lawfully continue and expand. The board should condition value on transferable and usable rights, with every adjustment linked to one ledger and one accountable owner.

8. Test archive quality

Decision usefulness requires a trace from radiometric, geometric, temporal, atmospheric and metadata quality across sensors and years to customer action and cash. The trace should distinguish data collected in the past, capability that creates future observations, software that transforms data and workflow adoption that makes the output operational. It should also show where public or third-party substitutes limit exclusivity and pricing power.

Diligence should obtain quality reports, calibration, validation, error rates and customer acceptance. The team should test completeness, rights, technical fitness and commercial use. Valuation then measures fitness for purpose and remediation cost. Decision makers should retain deductions for inconsistent or unverifiable data; the assumption register should prevent the same observation from appearing as archive value, model advantage, retention benefit and terminal uplift.

9. Test discoverability and access

The control question concerns catalogue completeness, search, APIs, latency, delivery, formats and compute proximity. It should be answered through a dated data, operating and customer record. Gross archive size, theoretical revisit, model demonstrations and aggregate backlog can hide unusable scenes, cloud constraints, degraded sensors, unsupported algorithms and customers who do not rely on the product in production.

Required support includes platform logs, query success, download times, API performance and support tickets. Exceptions, failed tests, missing rights and deteriorating usage deserve separate review. The model uses the evidence to calculate customer use and serving cost. The investment committee should separate valuable content from friction that prevents monetisation, retaining deductions for incomplete rights, fleet replacement, remediation and integration.

10. Measure archive monetisation

The analysis begins with subscriptions, usage, licences, professional services and derived products linked to historical data. Earth-observation value can migrate among data, hardware, software and customer use, so the unit of account must be explicit. The model should identify the legal right, technical dependency, user, cost and cash mechanism for every material component. This permits the team to test whether value survives a change in fleet, platform or ownership.

The evidence set comprises contract terms, product entitlements, queries, downloads, invoices and collections. Records should be dated, owned and reconciled to the relevant scene, sensor, product, contract and model assumption. Missing support remains a visible diligence gap. The financial model converts the

verified record into archive cash after storage, processing and support. The board should attribute revenue through observed product use, with every adjustment linked to one ledger and one accountable owner.

11. Reconstruct the fleet by sensor cohort

Decision usefulness requires a trace from satellite age, sensor type, resolution, swath, spectral bands, capacity, reliability and remaining life to customer action and cash. The trace should distinguish data collected in the past, capability that creates future observations, software that transforms data and workflow adoption that makes the output operational. It should also show where public or third-party substitutes limit exclusivity and pricing power.

Diligence should obtain manufacturing lots, launch records, telemetry, anomalies and calibration history. The team should test completeness, rights, technical fitness and commercial use. Valuation then measures usable collection capacity and replacement timing. Decision makers should model each cohort against customer requirements; the assumption register should prevent the same observation from appearing as archive value, model advantage, retention benefit and terminal uplift.

12. Measure revisit and refresh

The control question concerns the frequency at which a target can be observed under operational and weather constraints. It should be answered through a dated data, operating and customer record. Gross archive size, theoretical revisit, model demonstrations and aggregate backlog can hide unusable scenes, cloud constraints, degraded sensors, unsupported algorithms and customers who do not rely on the product in production.

Required support includes tasking records, achieved collections, coverage, cloud statistics and priority rules. Exceptions, failed tests, missing rights and deteriorating usage deserve separate review. The model uses the evidence to calculate usable refresh by geography and use case. The investment committee should use achieved delivery rather than design orbit claims, retaining deductions for incomplete rights, fleet replacement, remediation and integration.

13. Measure latency and delivery

The analysis begins with time from customer request or event to collection, processing, analytics and usable output. Earth-observation value can migrate among data, hardware, software and customer use, so the unit of account must be explicit. The model should identify the legal right, technical dependency, user, cost and cash mechanism for every material component. This permits the team to test whether value survives a change in fleet, platform or ownership.

The evidence set comprises timestamps, service levels, queue records and customer acceptance. Records should be dated, owned and reconciled to the relevant scene, sensor, product, contract and model assumption. Missing support remains a visible diligence gap. The financial model converts the verified record into time-sensitive value and contract performance. The board should price speed through observed customer economics, with every adjustment linked to one ledger and one accountable owner.

14. Measure tasking optionality

Decision usefulness requires a trace from customer ability to request, prioritise and receive future collections to customer action and cash. The trace should distinguish data collected in the past, capability that creates future observations, software that transforms data and workflow adoption that makes the output operational. It should also show where public or third-party substitutes limit exclusivity and pricing power.

Diligence should obtain tasking rights, queue rules, acceptance rates, conflicts and delivered scenes. The team should test completeness, rights, technical fitness and commercial use. Valuation then measures incremental tasking cash and capacity cost. Decision makers should avoid valuing theoretical tasking that is routinely displaced; the assumption register should prevent the same observation from appearing as archive value, model advantage, retention benefit and terminal uplift.

15. Model fleet reliability

The control question concerns sensor degradation, spacecraft failures, common-mode risk, station-keeping and calibration drift. It should be answered through a dated data, operating and customer record. Gross archive size, theoretical revisit, model demonstrations and aggregate backlog can hide unusable scenes, cloud constraints, degraded sensors, unsupported algorithms and customers who do not rely on the product in production.

Required support includes telemetry, anomaly logs, insurance, fleet availability and failure investigations. Exceptions, failed tests, missing rights and deteriorating usage deserve separate review. The model uses the evidence to calculate capacity loss and expected replacement. The investment committee should stress correlated cohort failures, retaining deductions for incomplete rights, fleet replacement, remediation and integration.

16. Model replacement capital

The analysis begins with design, manufacture, launch, insurance, commissioning and calibration required to sustain refresh. Earth-observation value can migrate among data, hardware, software and customer use, so the unit of account must be explicit. The model should identify the legal right, technical dependency, user, cost and cash mechanism for every material component. This permits the team to test whether value survives a change in fleet, platform or ownership.

The evidence set comprises supplier contracts, production cadence, launch slots, cost history and acceptance. Records should be dated, owned and reconciled to the relevant scene, sensor, product, contract and model assumption. Missing support remains a visible diligence gap. The financial model converts the verified record into sustainable free cash flow after replenishment. The board should fund replacement before applying terminal value, with every adjustment linked to one ledger and one accountable owner.

17. Test launch and ground dependencies

Decision usefulness requires a trace from launch providers, ground stations, cloud, processing, data transport and specialist suppliers to customer action and cash. The trace should distinguish data collected in the past, capability that creates future observations, software that transforms data and workflow adoption that makes the output operational. It should also show where public or third-party substitutes limit exclusivity and pricing power.

Diligence should obtain contracts, capacity, arrears, change rights, lead times and alternatives. The team should test completeness, rights, technical fitness and commercial use. Valuation then measures continuity, cost and concentration exposure. Decision makers should price dependencies through executable continuation plans; the assumption register should prevent the same observation from appearing as archive value, model advantage, retention benefit and terminal uplift.

18. Reconcile archive and refresh

The control question concerns the relationship between historical baselines and new observations that detect change. It should be answered through a dated data, operating and customer record. Gross archive size, theoretical revisit, model demonstrations and aggregate backlog can hide unusable scenes, cloud

constraints, degraded sensors, unsupported algorithms and customers who do not rely on the product in production.

Required support includes customer workflows, model inputs, product design and renewal evidence. Exceptions, failed tests, missing rights and deteriorating usage deserve separate review. The model uses the evidence to calculate combined value without duplicate attribution. The investment committee should identify which product requires both components, retaining deductions for incomplete rights, fleet replacement, remediation and integration.

19. Inventory algorithms and models

The analysis begins with detection, classification, segmentation, forecasting, fusion and tasking capabilities. Earth-observation value can migrate among data, hardware, software and customer use, so the unit of account must be explicit. The model should identify the legal right, technical dependency, user, cost and cash mechanism for every material component. This permits the team to test whether value survives a change in fleet, platform or ownership.

The evidence set comprises model registry, code, training data, benchmarks, versioning and access rights. Records should be dated, owned and reconciled to the relevant scene, sensor, product, contract and model assumption. Missing support remains a visible diligence gap. The financial model converts the verified record into attributable product performance and development cost. The board should distinguish proprietary capability from replaceable tooling, with every adjustment linked to one ledger and one accountable owner.

20. Test model performance

Decision usefulness requires a trace from accuracy, recall, precision, drift, robustness, explainability and mission-specific thresholds to customer action and cash. The trace should distinguish data collected in the past, capability that creates future observations, software that transforms data and workflow adoption that makes the output operational. It should also show where public or third-party substitutes limit exclusivity and pricing power.

Diligence should obtain validation sets, independent tests, customer acceptance and incident records. The team should test completeness, rights, technical fitness and commercial use. Valuation then measures commercial contribution after error and oversight cost. Decision makers should value performance at the customer's operating threshold; the assumption register should prevent the same observation from appearing as archive value, model advantage, retention benefit and terminal uplift.

21. Test training-data dependency

The control question concerns the extent to which models depend on the acquired archive, third-party data or continuing fleet refresh. It should be answered through a dated data, operating and customer record. Gross archive size, theoretical revisit, model demonstrations and aggregate backlog can hide unusable scenes, cloud constraints, degraded sensors, unsupported algorithms and customers who do not rely on the product in production.

Required support includes lineage, licences, feature stores, retraining cadence and model cards. Exceptions, failed tests, missing rights and deteriorating usage deserve separate review. The model uses the evidence to calculate portability and future operating cost. The investment committee should discount models that cannot be retrained lawfully or effectively, retaining deductions for incomplete rights, fleet replacement, remediation and integration.

22. Measure workflow integration

The analysis begins with APIs, alerts, dashboards and outputs embedded in customer decisions and systems. Earth-observation value can migrate among data, hardware, software and customer use, so the unit of account must be explicit. The model should identify the legal right, technical dependency, user, cost and cash mechanism for every material component. This permits the team to test whether value survives a change in fleet, platform or ownership.

The evidence set comprises usage logs, integrations, automation, user roles, support and renewal evidence. Records should be dated, owned and reconciled to the relevant scene, sensor, product, contract and model assumption. Missing support remains a visible diligence gap. The financial model converts the verified record into retention and expansion attributable to workflow adoption. The board should separate genuine switching cost from contract lock-in, with every adjustment linked to one ledger and one accountable owner.

23. Reconstruct customer cohorts

Decision usefulness requires a trace from civil government, defence, agriculture, mapping, energy, insurance, finance and enterprise users to customer action and cash. The trace should distinguish data collected in the past, capability that creates future observations, software that transforms data and workflow adoption that makes the output operational. It should also show where public or third-party substitutes limit exclusivity and pricing power.

Diligence should obtain contracts, usage, price, gross margin, renewal, churn and service requirements. The team should test completeness, rights, technical fitness and commercial use. Valuation then measures cohort contribution and sensitivity to archive or refresh. Decision makers should value each cohort through its actual decision workflow; the assumption register should prevent the same observation from appearing as archive value, model advantage, retention benefit and terminal uplift.

24. Test government and defence contracts

The control question concerns funding, task orders, security, acceptance, assignment, termination and data-rights clauses. It should be answered through a dated data, operating and customer record. Gross archive size, theoretical revisit, model demonstrations and aggregate backlog can hide unusable scenes, cloud constraints, degraded sensors, unsupported algorithms and customers who do not rely on the product in production.

Required support includes executed awards, funded backlog, clearances, invoices and collections. Exceptions, failed tests, missing rights and deteriorating usage deserve separate review. The model uses the evidence to calculate durable cash after change of control. The investment committee should treat unfunded ceilings and non-assignable rights separately, retaining deductions for incomplete rights, fleet replacement, remediation and integration.

25. Measure customer concentration

The analysis begins with reliance on individual agencies, programmes, anchor users and channel partners. Earth-observation value can migrate among data, hardware, software and customer use, so the unit of account must be explicit. The model should identify the legal right, technical dependency, user, cost and cash mechanism for every material component. This permits the team to test whether value survives a change in fleet, platform or ownership.

The evidence set comprises revenue, backlog, receivables, renewals and product use by counterparty. Records should be dated, owned and reconciled to the relevant scene, sensor, product, contract and model assumption. Missing support remains a visible diligence gap. The financial model converts the verified

record into downside cash and bargaining exposure. The board should stress loss or delay of the largest customer, with every adjustment linked to one ledger and one accountable owner.

26. Test storage and compute economics

Decision usefulness requires a trace from object storage, retrieval, processing, model inference, egress and support cost to customer action and cash. The trace should distinguish data collected in the past, capability that creates future observations, software that transforms data and workflow adoption that makes the output operational. It should also show where public or third-party substitutes limit exclusivity and pricing power.

Diligence should obtain cloud bills, unit metrics, architecture and workload patterns. The team should test completeness, rights, technical fitness and commercial use. Valuation then measures gross margin by archive query and refreshed product. Decision makers should include the full serving cost of historical and new data; the assumption register should prevent the same observation from appearing as archive value, model advantage, retention benefit and terminal uplift.

27. Test cybersecurity and provenance

The control question concerns integrity, access, chain of custody, tamper evidence and recovery across data and models. It should be answered through a dated data, operating and customer record. Gross archive size, theoretical revisit, model demonstrations and aggregate backlog can hide unusable scenes, cloud constraints, degraded sensors, unsupported algorithms and customers who do not rely on the product in production.

Required support includes identity controls, logs, hashes, incident tests and customer requirements. Exceptions, failed tests, missing rights and deteriorating usage deserve separate review. The model uses the evidence to calculate trust, remediation and residual risk. The investment committee should require proven provenance for high-consequence use, retaining deductions for incomplete rights, fleet replacement, remediation and integration.

28. Construct sustainable free cash flow

The analysis begins with archive, refresh, algorithm and workflow cash after operations, replacement, tax and working capital. Earth-observation value can migrate among data, hardware, software and customer use, so the unit of account must be explicit. The model should identify the legal right, technical dependency, user, cost and cash mechanism for every material component. This permits the team to test whether value survives a change in fleet, platform or ownership.

The evidence set comprises reconciled accounts, product metrics and fleet plans. Records should be dated, owned and reconciled to the relevant scene, sensor, product, contract and model assumption. Missing support remains a visible diligence gap. The financial model converts the verified record into cash available after preserving the complete service. The board should use the same product state in forecast and terminal value, with every adjustment linked to one ledger and one accountable owner.

29. Build the hypothetical component case

Decision usefulness requires a trace from USD 4.55 billion of admitted archive, refresh, algorithm and workflow value to customer action and cash. The trace should distinguish data collected in the past, capability that creates future observations, software that transforms data and workflow adoption that makes the output operational. It should also show where public or third-party substitutes limit exclusivity and pricing power.

Diligence should obtain illustrative inputs clearly separated from public evidence. The team should test completeness, rights, technical fitness and commercial use. Valuation then measures a component sum before overlap and capital deductions. Decision makers should use the case to demonstrate mechanics rather than represent a company; the assumption register should prevent the same observation from appearing as archive value, model advantage, retention benefit and terminal uplift.

30. Eliminate overlap

The control question concerns duplicate value across archive uniqueness, model training, customer retention and terminal growth. It should be answered through a dated data, operating and customer record. Gross archive size, theoretical revisit, model demonstrations and aggregate backlog can hide unusable scenes, cloud constraints, degraded sensors, unsupported algorithms and customers who do not rely on the product in production.

Required support includes an assumption register linking every benefit to one model location. Exceptions, failed tests, missing rights and deteriorating usage deserve separate review. The model uses the evidence to calculate a USD 0.45 billion deduction in the worked case. The investment committee should reconcile all components before board approval, retaining deductions for incomplete rights, fleet replacement, remediation and integration.

31. Apply replacement and quality deductions

The analysis begins with fleet replenishment, launch, data-rights limits and archive remediation. Earth-observation value can migrate among data, hardware, software and customer use, so the unit of account must be explicit. The model should identify the legal right, technical dependency, user, cost and cash mechanism for every material component. This permits the team to test whether value survives a change in fleet, platform or ownership.

The evidence set comprises engineering plans, rights review, quality tests and executable budgets. Records should be dated, owned and reconciled to the relevant scene, sensor, product, contract and model assumption. Missing support remains a visible diligence gap. The financial model converts the verified record into USD 0.80 billion of deductions in the worked case. The board should release deductions only after verified evidence and funding, with every adjustment linked to one ledger and one accountable owner.

32. Apply integration deductions

Decision usefulness requires a trace from platform migration, data harmonisation, model retraining, customer transition and organisational integration to customer action and cash. The trace should distinguish data collected in the past, capability that creates future observations, software that transforms data and workflow adoption that makes the output operational. It should also show where public or third-party substitutes limit exclusivity and pricing power.

Diligence should obtain workplan, dependencies, costs, owners and acceptance criteria. The team should test completeness, rights, technical fitness and commercial use. Valuation then measures USD 0.15 billion of integration cost. Decision makers should fund the plan and stage value through milestones; the assumption register should prevent the same observation from appearing as archive value, model advantage, retention benefit and terminal uplift.

33. Admit verified combination value

The control question concerns benefits created by joining unique history, future refresh, models and embedded workflows. It should be answered through a dated data, operating and customer record. Gross archive size, theoretical revisit, model demonstrations and aggregate backlog can hide unusable scenes,

cloud constraints, degraded sensors, unsupported algorithms and customers who do not rely on the product in production.

Required support includes cross-sell evidence, product tests, customer commitments and avoided cost. Exceptions, failed tests, missing rights and deteriorating usage deserve separate review. The model uses the evidence to calculate USD 0.30 billion in the worked case. The investment committee should cap synergy by buyer capture and delivery evidence, retaining deductions for incomplete rights, fleet replacement, remediation and integration.

34. Structure the transaction

The analysis begins with price, earn-outs, holdbacks, data-right warranties, fleet milestones, transition and financing. Earth-observation value can migrate among data, hardware, software and customer use, so the unit of account must be explicit. The model should identify the legal right, technical dependency, user, cost and cash mechanism for every material component. This permits the team to test whether value survives a change in fleet, platform or ownership.

The evidence set comprises term sheet, sources and uses, consent map and downside case. Records should be dated, owned and reconciled to the relevant scene, sensor, product, contract and model assumption. Missing support remains a visible diligence gap. The financial model converts the verified record into risk allocation aligned with unresolved evidence. The board should use contingent value for archive quality, refresh and customer retention, with every adjustment linked to one ledger and one accountable owner.

35. Establish post-close governance

Decision usefulness requires a trace from authority over data, fleet, algorithms, customers, capital and integration to customer action and cash. The trace should distinguish data collected in the past, capability that creates future observations, software that transforms data and workflow adoption that makes the output operational. It should also show where public or third-party substitutes limit exclusivity and pricing power.

Diligence should obtain board mandates, dashboards, model governance, assurance and decision logs. The team should test completeness, rights, technical fitness and commercial use. Valuation then measures accountable delivery and early warning. Decision makers should refresh valuation when quality, capacity or retention changes materially; the assumption register should prevent the same observation from appearing as archive value, model advantage, retention benefit and terminal uplift.

36. Define the investment decision

The control question concerns the price, perimeter, funding, conditions and integration actions supported by evidence. It should be answered through a dated data, operating and customer record. Gross archive size, theoretical revisit, model demonstrations and aggregate backlog can hide unusable scenes, cloud constraints, degraded sensors, unsupported algorithms and customers who do not rely on the product in production.

Required support includes the final value bridge, downside liquidity, consent schedule and milestone plan. Exceptions, failed tests, missing rights and deteriorating usage deserve separate review. The model uses the evidence to calculate an auditable decision range. The investment committee should approve, reprice, stage or decline through a dated record, retaining deductions for incomplete rights, fleet replacement, remediation and integration.

Conclusion

Earth-observation M&A requires separate treatment of historical data and future collection. An archive can remain useful after a satellite retires, while its value depends on rights, quality, discoverability and customer use. A sensor fleet produces refresh and tasking, while its value depends on capacity, reliability, remaining life and replacement funding.

The framework adds algorithms and customer workflows as separate ledgers, then removes overlap before admitting combination value. This allows the buyer to identify which cash comes from unique history, which requires continuing refresh, which depends on software and which reflects adoption inside customer decisions.

For a board or investment committee, the decisive test is whether the combined archive, fleet, platform and workflows can sustain customer outcomes and cash after replacement and integration. The answer determines price, conditions, financing and post-close priorities.

Appendix A. Archive register

The minimum register contains dataset, sensor, date, geography, modality, resolution, processing level, quality, metadata, provenance, ownership, licence, customer entitlement, storage location, access method, observed use and serving cost.

Appendix B. Refresh-capability register

For each fleet cohort, record sensor capability, orbit, capacity, achieved revisit, latency, reliability, remaining life, calibration, tasking priority, ground dependency, replacement date, capital, launch and acceptance plan.

Appendix C. Investment checklist

The approval file should contain archive inventory, rights, quality, uniqueness, access, monetisation, fleet cohorts, refresh, latency, tasking, replacement, algorithms, training data, workflows, customer cohorts, contracts, serving cost, security, integration, financing and downside liquidity. Evidence should be dated and owned.

Appendix D. Worked-case figures and tables

Figure 1. Earth-observation transaction value architecture

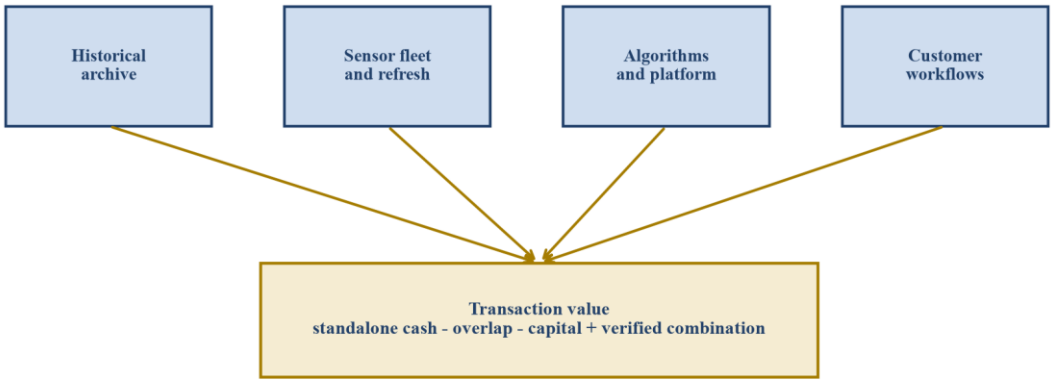

Proposed architecture separating archive, refresh, algorithms and customer workflows.

Figure 2. Hypothetical archive data-moat score

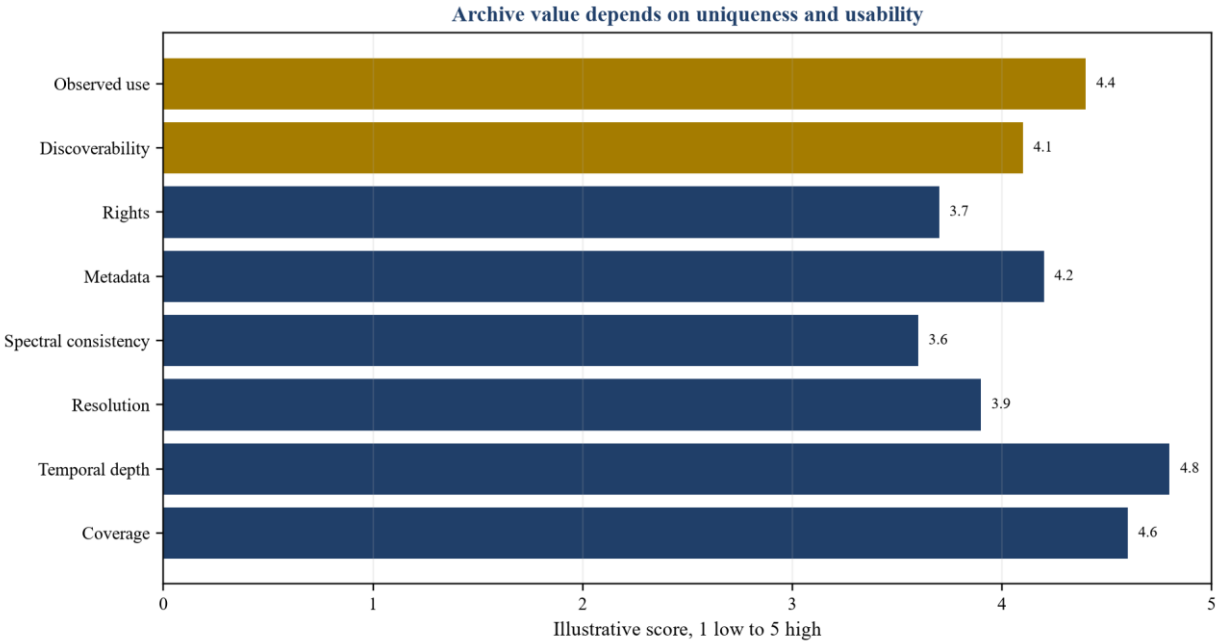

Illustrative assessment of archive uniqueness and usability.

Figure 3. Hypothetical refresh map by customer workflow

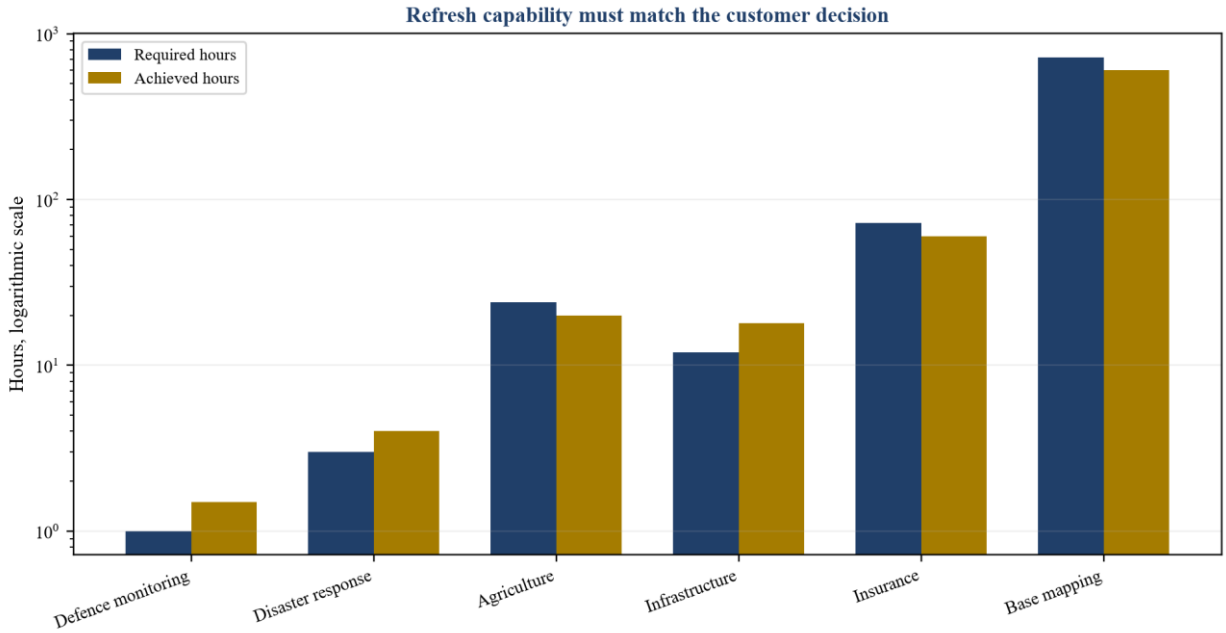

Illustrative required and achieved refresh intervals.

Figure 4. Hypothetical transaction value bridge

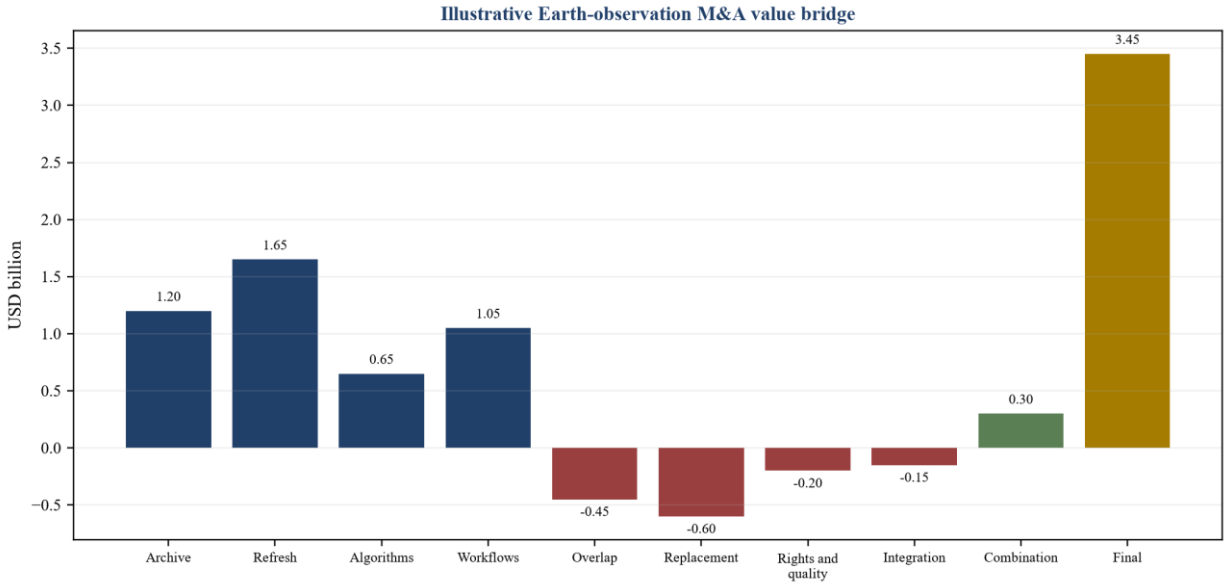

Illustrative bridge from admitted components to final transaction value.

Figure 5. Archive-refresh dependency matrix

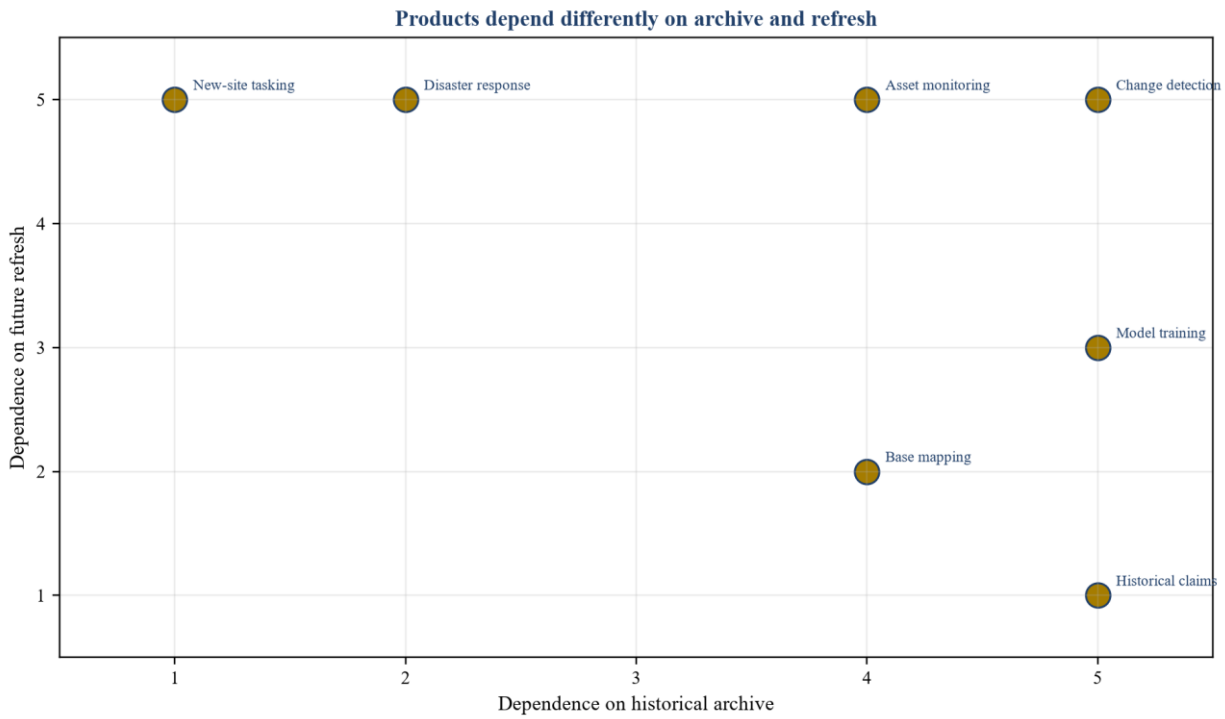

Illustrative dependence of customer products on history and future collection.

Table 1. Four-component valuation ledger

Component	Minimum evidence	Primary cash mechanism
Archive	rights quality uniqueness and observed use	licence usage and derived products
Refresh	achieved collection and replacement plan	tasking subscriptions and capacity
Algorithms	attributable performance and lawful training	analytics and operating leverage
Workflows	embedded use renewal and switching evidence	retention expansion and solution revenue

Proposed valuation treatment.

Table 2. Hypothetical archive profile

Measure	Illustrative result	Valuation use
Unique scenes	2.4 billion	gross inventory before usability tests
Median temporal depth	11 years	longitudinal analysis
Searchable with complete metadata	86 per cent	addressable archive
Rights permitting derived products	72 per cent	monetisable subset
Scenes used by paying customers in 12 months	8 per cent	observed demand

Illustrative measures describing no identified company.

Table 3. Fleet and refresh controls

Dimension	Evidence	Valuation use
Capacity	achieved scenes and area per day	usable supply

Dimension	Evidence	Valuation use
Revisit	achieved intervals by geography	workflow fit
Latency	request to delivered product	time-sensitive value
Reliability	cohort availability and anomalies	expected capacity
Replacement	manufacture launch and calibration	sustaining capital

Proposed diligence fields.

Table 4. Hypothetical value bridge

Step	Amount	Evidence required
Archive cash value	1.20	rights use and contribution
Refresh capability	1.65	fleet capacity and funded replacement
Algorithms and platform	0.65	performance adoption and IP
Customer workflows	1.05	retention integration and contribution
Overlap	minus 0.45	assumption reconciliation
Replacement capital	minus 0.60	executable fleet plan
Data rights and quality	minus 0.20	legal and technical remediation
Integration	minus 0.15	funded delivery plan
Verified combination value	plus 0.30	buyer capture and customer evidence
Final illustrative value	3.45	integrated evidence set

All amounts are illustrative USD billions.

Table 5. Transaction approval gate

Gate	Approval evidence	Stop condition
Archive	inventory rights quality and use	material unusable or restricted data
Fleet	achieved collection and replacement	refresh gap before funded replacement
Algorithms	validated performance and training rights	non-portable or unsupported model
Workflows	production integration and renewal	unverified switching benefit
Economics	serving cost and sustainable cash	margin or liquidity gap
Combination	attributable buyer benefit	duplicated or aspirational synergy

Proposed board control.

Table 6. Transaction protections by unresolved evidence

Risk	Potential protection	Release evidence
Archive completeness	holdback or price adjustment	verified inventory
Data rights	warranty indemnity and condition	transferable licences
Fleet performance	earn-out or milestone	achieved capacity and latency
Replacement	committed funding and covenant	launch and acceptance
Customer retention	earn-out	collected retained revenue
Algorithm performance	staged payment	independent validation

Proposed risk allocation.

Table 7. Post-close value dashboard

Dimension	Core measure	Trigger
Archive	usable scenes queries and contribution	falling usage or rights gap
Refresh	revisit latency and capacity	missed workflow requirement
Fleet	reliability life and replacement	capacity shortfall
Models	accuracy drift and adoption	performance deterioration
Customers	renewals churn expansion and margin	cohort below case
Integration	migrations dependencies and benefits	missed milestone
Capital	cash burn replacement and liquidity	funding shortfall

Proposed monthly decision record.

Frequently asked questions

Q: Why value an Earth-observation archive separately from the fleet? A: Historical data can remain useful after a satellite retires, while future refresh depends on continuing sensor capacity and replacement capital. The cash mechanisms and risks differ.

Q: What makes a data archive valuable? A: Transferable rights, usable quality, unique temporal or geographic coverage, strong metadata, efficient access and observed customer use support value.

Q: How should sensor fleets be valued? A: Use achieved collection, revisit, latency, reliability, remaining life and the capital required to sustain customer refresh requirements.

Q: How should algorithms enter the model? A: Measure attributable performance, adoption, lawful training rights, portability, operating cost and customer willingness to pay.

Q: How should customer workflow value be tested? A: Use production integrations, active users, automated decisions, renewal, expansion and switching evidence after considering contract lock-in.

Q: How is double counting avoided? A: Assign each economic benefit to one ledger and reconcile archive, model, retention, synergy and terminal assumptions before approval.

Q: What belongs in terminal value? A: Sustainable cash after fleet replacement, data serving, model maintenance, customer support, licence compliance and continuing integration obligations.

Q: What is the central board decision? A: The board must decide whether archive cash, refresh capability, software contribution and embedded workflows support the proposed price after overlap, capital and integration.

References

[1] NASA, Commercial Satellite Data Acquisition Programme. https://science.nasa.gov/earth-science/csda/

[2] NASA, Commercial Satellite Data Acquisition available datasets. https://science.nasa.gov/earth-science/csda/commercial-datasets/

[3] NASA, Satellite Data Explorer. https://science.nasa.gov/earth-science/csda/satellite-data-explorer/

[4] NASA, Satellite Data Evaluation. https://science.nasa.gov/earth-science/csda/satellite-data-evaluation/

[5] Planet Labs PBC, 2026 Annual Report. https://www.sec.gov/Archives/edgar/data/1836833/000119312526119957/pl-20260131.htm

[6] BlackSky Technology, 2025 Annual Report. https://www.sec.gov/Archives/edgar/data/1753539/000175353926000032/bksy-20251231.htm

[7] BlackSky Technology, 2026 Annual Shareholder Letter. https://www.sec.gov/Archives/edgar/data/1753539/000175353926000113/a202510k_wrapperfinal.pdf

[8] Satellogic, 2025 Annual Report. https://www.sec.gov/Archives/edgar/data/1874315/000187431526000013/satl-20251231.htm

[9] Spire Global, 2025 Annual Report. https://www.sec.gov/Archives/edgar/data/1816017/000119312526116169/spir-20251231.htm

[10] European Space Agency, Copernicus Contributing Missions commercial data procurement. https://commercialisation.esa.int/opportunities/dynamic-purchasing-system-dps-for-copernicus-contributing-missions-ccm-category-1/

[11] European Commission, Copernicus Data Space Ecosystem. https://dataspace.copernicus.eu/

[12] U.S. Geological Survey, Landsat data access and archive. https://www.usgs.gov/landsat-missions/landsat-data-access

[13] NOAA, Commercial Data Program. https://www.space.commerce.gov/business-with-noaa/commercial-data-program/

[14] NASA, Earth Science Data Systems data standards and practices. https://www.earthdata.nasa.gov/about/esdis/esco/standards-practices

[15] NASA, Satellite Tasking and Archive Request System. https://science.nasa.gov/science-research/earth-science/csda-releases-new-satellite-tasking-and-acquisition-request-system/

[16] NASA, LAADS Distributed Active Archive Center data access. https://ladsweb.modaps.eosdis.nasa.gov/data-access/

[17] Open Geospatial Consortium, standards and specifications. https://www.ogc.org/standards/

[18] U.S. National Institute of Standards and Technology, AI Risk Management Framework. https://www.nist.gov/itl/ai-risk-management-framework

[19] U.S. National Institute of Standards and Technology, Cybersecurity Framework 2.0. https://www.nist.gov/cyberframework

[20] IFRS Foundation, IFRS 3 Business Combinations. https://www.ifrs.org/issued-standards/list-of-standards/ifrs-3-business-combinations/

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

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

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

[24] International Private Equity and Venture Capital Valuation Guidelines. https://www.privateequityvaluation.com/

[25] International Valuation Standards Council, International Valuation Standards. https://www.ivsc.org/standards/

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.