Cyber Insurance for Space Infrastructure: Turning Telemetry into Underwriting Evidence

An Evidence-Led Framework for Cyber Coverage, Aggregation Risk and Transaction Decisions

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Cyber insurance for space infrastructure is difficult to underwrite because an operator's exposure spans spacecraft, ground stations, mission-control software, cloud services, communications links, suppliers, customers and public infrastructure. A questionnaire can identify stated controls. It rarely shows whether those controls operate continuously, whether incidents are detected quickly, whether recovery has been tested or whether several insureds share the same failure path. Operational telemetry can strengthen the evidence base when it is governed, interpreted and connected to loss.

This paper develops an evidence-led framework for translating cyber and operational telemetry into underwriting evidence for commercial satellite operators, hybrid satellite networks and space-service platforms. It separates security signals from insurance conclusions; maps asset, control, incident, recovery and dependency evidence to coverage and pricing decisions; and shows how evidence can inform M&A diligence, representations, policy conditions, limits, retentions and post-close remediation.

The analysis draws on NIST's Cybersecurity Framework, satellite-ground and hybrid-network profiles, enterprise risk guidance and security-control catalogues; ENISA's commercial-space and LEO satellite assessments; NASA space-security practice; regulatory incident-reporting requirements; and insurance-market guidance on cyber aggregation, policy wording and state-backed attacks. These sources establish relevant risk and control concepts. They do not establish an actuarially credible loss rate for any identified operator or constellation.

A wholly hypothetical acquisition illustrates the method. The target operates a communications constellation and supporting ground services. Management proposes USD 180 million of enterprise value and seeks USD 40 million of cyber and technology-errors-and-omissions capacity. The initial submission relies on policies and a point-in-time assessment. A structured evidence pack adds identity, vulnerability, command-path, incident, recovery, supplier and service-continuity measures. The illustrative insurer then differentiates losses that can be modelled, risks requiring sublimits or exclusions, and conditions that require remediation. The buyer uses the same evidence to allocate a USD 12 million cyber remediation reserve, place USD 8 million of consideration in a control-linked escrow and adjust business-interruption assumptions.

The central conclusion is that telemetry becomes underwriting evidence only when its scope, provenance, completeness, interpretation and connection to loss are credible. The most useful programme combines stable exposure definitions, control-performance measures, tested recovery, incident evidence, dependency mapping and clear policy wording. That programme can improve decisions without presenting a risk score as certainty.

JEL Classification: G22, G32, G34, L63, L96, M15, O33

Keywords: cyber insurance, space infrastructure, satellite telemetry, underwriting evidence, aggregation risk, business interruption, space M&A, operational resilience, cyber controls, risk transfer

Introduction

Space businesses depend on digital systems for command, telemetry, tracking, payload scheduling, navigation, customer delivery, billing and regulatory evidence. Those systems can fail through malicious activity, error, software defects, supplier interruption or physical events. The same incident can affect spacecraft operations, ground services, customer commitments, data integrity and public reporting. A conventional cyber application captures only part of that operating reality.

Insurance decisions require evidence about exposure, control and consequence. The underwriter needs to know which services are at risk, which systems support them, how failures are detected, how long recovery takes and how losses could aggregate across customers or portfolios. The insured needs policy language that responds to its actual operating model. An acquirer needs to understand whether insurance transfers a defined risk or leaves a material exposure with the transaction.

Telemetry can improve this evidence base. Identity logs can show whether privileged access is controlled. Vulnerability data can show exposure age and remediation cadence. mission and network telemetry can show service degradation, command-path interruption and recovery performance. Supplier records can show common dependencies. Each dataset also has limitations. Missing coverage, changing definitions, noisy alerts, time gaps and weak governance can create false confidence.

This paper presents a practical framework for boards, acquirers, lenders, insurers and operators. It treats telemetry as one part of an underwriting evidence chain. The method connects evidence to loss scenarios, policy mechanics and transaction value.

1 Define the insurance decision

The first step is to define the decision that evidence must support. The operator may seek first-party cyber cover, technology errors and omissions, business interruption, contingent business interruption, privacy liability, media liability, crime, property, space hull or launch cover. The relevant peril, trigger and loss measure differ across these products. Evidence designed for one decision cannot be assumed to answer another.

The decision statement should specify insured entities, covered operations, policy period, limits, retentions, sublimits and material exclusions. It should identify whether the analysis concerns risk selection, pricing, capacity, wording, renewal, claim preparation or an acquisition. It should also identify the jurisdiction and regulatory perimeter because reporting and coverage obligations vary.

In M&A, the decision extends beyond policy placement. The buyer needs to know whether existing cover survives change of control, whether prior acts are protected, whether known circumstances must be notified, whether tail cover is needed and whether the target's evidence supports transaction representations. The acquisition agreement and insurance programme should use a common exposure definition.

2 Establish the operating perimeter

The operating perimeter should connect legal entities to services, systems and physical infrastructure. A satellite operator can rely on spacecraft buses and payloads, command centres, ground stations, cloud environments, identity services, terrestrial networks, user terminals, mobile-network partners, data providers and specialist suppliers. Some components may be owned. Others are leased, shared or supplied under service agreements.

The perimeter should record which party controls each component, which party observes it and which party bears the financial consequence of failure. Control and loss can sit with different organisations. A ground-station provider may operate the link while the satellite operator owes the customer service credit. A cloud provider may restore infrastructure while the operator must reconstruct mission data and customer delivery.

Hybrid satellite networks require particular care because independently owned components can carry different assurance levels. NIST's hybrid satellite network profile highlights the importance of interfaces and shared responsibilities. An underwriting submission should therefore include responsibility matrices and service dependencies rather than a single asset list.

3 Build the exposure inventory

The exposure inventory links assets and services to financial consequence. Each record should include owner, operator, location, mission, data class, connectivity, customer dependency, criticality, replacement path and recovery objective. Spacecraft and payloads require cohort, software and command characteristics. Ground and cloud systems require environment, identity, configuration and dependency records.

The inventory should be reconciled against observed telemetry. Configuration repositories show what management expects. Network discovery, cloud inventories, endpoint records and mission systems show what is active. Differences should be treated as evidence gaps. An unrecorded gateway can create exposure. A listed system may be retired and should not inflate the declared perimeter.

Financial mapping distinguishes high-value services from technically prominent assets. A low-cost identity provider can sit on the path to substantial revenue. A highly visible spacecraft may have limited near-term cash contribution. Underwriting and transaction analysis should follow the loss path, not replacement cost alone.

4 Separate telemetry from evidence

Telemetry is a recorded observation generated by a system or process. Underwriting evidence is information that has a defined scope, provenance, interpretation and relevance to an insurance decision. The distinction matters because a large volume of telemetry can still provide weak evidence.

An evidence record should identify source system, accountable owner, collection method, time coverage, units, exclusions, retention, access controls and validation. It should state the control or loss question that the record addresses. For example, a multifactor-authentication dashboard may show enrolment. It does not prove that every privileged path enforces the control or that emergency accounts are governed.

Evidence should preserve exceptions. A monthly average can conceal a critical gap during a software release or acquisition cutover. Point-in-time screenshots can omit duration. The evidence pack should provide distributions, thresholds, incidents and reconciliations that allow an underwriter to understand control performance across the insured period.

5 Create an evidence taxonomy

Five evidence classes support a disciplined submission. Exposure evidence defines the assets, services, data and dependencies at risk. Control evidence shows preventive and protective measures. Event evidence shows detections, anomalies, incidents and near misses. Recovery evidence shows restoration and continuity capability. Financial evidence shows how operational disruption becomes insured loss.

Each class needs stable identifiers. A service identifier should link the customer contract, supporting systems, control evidence, incidents, recovery test and cash impact. Without this linkage, the submission

becomes a collection of unrelated metrics. Stable identifiers also support renewal comparisons and post-acquisition integration.

The taxonomy should distinguish reported facts from management assumptions. Observed recovery time from an exercise is evidence. A forecast recovery time for an untested scenario is a management assumption. Both can inform a decision when their status is explicit.

6 Map the threat and loss pathways

The threat analysis should cover malicious access, ransomware, supply-chain compromise, command interference, jamming, spoofing, credential theft, data corruption, denial of service, insider activity and software vulnerability exploitation. It should also cover non-malicious events that can resemble cyber loss, including configuration error, failed deployment and accidental deletion.

The loss pathway translates an event into consequence. A compromised administrator account can enable unauthorised command or disable monitoring. Service may be interrupted while command authority is restored. Customers may claim service credits, cancel contracts or require investigation. Regulators may require notice. The operator may incur forensic, legal, communications, restoration and additional operating costs.

The model should avoid assuming that every cyber event produces physical damage or spacecraft loss. Many events produce data, service, response or liability costs. Conversely, a command-path compromise can have safety or physical consequences that may fall outside a standalone cyber policy. Clear pathways support clear coverage.

7 Measure identity and access control

Identity evidence should cover human, service, machine and emergency identities across corporate, development, ground and mission environments. Useful measures include privileged-account inventory completeness, multifactor-enforcement coverage, dormant-account age, recertification completion, failed authentication, emergency access, credential rotation and separation between development and production.

The evidence should identify control gaps by critical path. Ninety-eight percent multifactor coverage can be inadequate if the excluded two percent includes command administration or signing authority. A risk-weighted measure connects each identity to capability and consequence.

Changes in identity control deserve analysis. A falling number of privileged accounts may reflect remediation or incomplete discovery. A spike in failed authentication may reflect an attack, a new application or a telemetry change. The submission should include explanations and change history so the insurer can distinguish control performance from measurement artefacts.

8 Measure vulnerability and configuration exposure

Vulnerability evidence should identify affected component, exploitability, mission consequence, exposure path, compensating control, remediation owner and age. Raw vulnerability counts have limited value because detection coverage and criticality vary. The useful question is whether material exposure is discovered, prioritised and reduced within the system's operational constraints.

Space systems can have limited patch windows and long component lives. The operator should show how it manages exceptions, tests releases and applies compensating controls. Ground services, internet-facing systems and cloud components may support faster remediation than flight systems. The evidence model should preserve that distinction.

Configuration telemetry should cover approved baselines, deviations, unauthorised change, backup and recovery. Command validation, cryptographic configuration, network segmentation and logging deserve

priority. An underwriter should be able to see whether material exceptions are temporary, approved and monitored.

9 Measure detection performance

Detection evidence should connect monitoring coverage to material attack paths. Measures can include percentage of critical systems forwarding logs, telemetry latency, alert-to-triage time, mean time to detect, false-positive handling, sensor health, clock synchronisation and coverage of privileged and service identities.

Mission telemetry creates additional signals. Unexpected command attempts, authentication failures, link-state changes, anomalous scheduling, configuration drift and unusual data movement can support detection. Those signals need operational context because legitimate mission activity can produce unusual patterns.

Evidence quality matters more than alert volume. A high number of alerts may show broad monitoring or weak tuning. A low number may show effective controls or blind spots. The submission should include coverage, test results and representative incident traces.

10 Measure incident response

Incident evidence should show classification, containment, escalation, decision authority, communications, forensics and lessons learned. Useful measures include time to acknowledge, time to contain, time to restore, preservation of evidence, severity changes and completion of corrective actions.

Tabletop exercises test coordination. Technical exercises test systems and people. The strongest evidence combines both. A response plan should include mission operations, engineering, security, legal, insurance, customer and regulatory roles. The operator should understand when an event becomes a policy notification, regulatory report or material disclosure question.

Incident records should include near misses and control failures. A blocked attack can reveal exposure and control effectiveness. Repeated near misses can indicate concentration or deteriorating conditions. The insurer and insured should agree how sensitive records are shared, protected and used.

11 Measure recovery and continuity

Recovery evidence should demonstrate restoration of the service, not only restoration of files. The operator should test identity, configurations, mission applications, telemetry history, customer interfaces and external dependencies. Exercises should record recovery point, recovery time, manual work, degraded operation and unresolved exceptions.

Space services may require continuity across mission windows. A recovery test should consider whether command and monitoring can continue while the primary site, cloud region or supplier is unavailable. Alternative ground stations, communications paths and operating teams may reduce exposure when they are technically and contractually usable.

Observed recovery should be compared with business-interruption assumptions. If policy waiting periods or financial models assume restoration within twelve hours, an exercise that takes thirty-six hours requires a specific response. That response may include remediation, a larger retention, a lower limit, a sublimit or revised pricing.

12 Quantify service interruption

Business-interruption modelling begins with service units. These can include satellite capacity, accepted imagery, navigation availability, mission contacts, processed data or customer endpoints. The model

should connect unavailable service units to revenue, variable cost, service credits, additional expense, customer churn and restoration cost.

The loss curve should vary by duration. A short interruption may be absorbed by redundancy or service-level tolerance. A longer interruption may miss delivery windows, trigger credits and weaken renewal. Critical government or infrastructure customers can create obligations that differ from commercial contracts.

The model should state which amounts are insured, retained, excluded or recoverable elsewhere. It should avoid treating enterprise value as an insured loss. Enterprise value can be affected through customer retention and future cash flow, while the policy usually responds to defined costs and income loss during a measured period.

13 Identify aggregation risk

Aggregation occurs when one event affects several insured assets, customers or policies. Space infrastructure can share ground stations, cloud providers, identity services, software components, orbital data, communications networks and specialist suppliers. A common vulnerability or provider outage can create correlated loss across operators.

The operator should map first-order and second-order dependencies. A cloud service is a first-order dependency. The cloud service's shared identity, network or software dependency can be second order. Contract terms and technical architecture should be reconciled so the map reflects actual service delivery.

The insurer needs portfolio context that the operator may not possess. The operator can still provide service and supplier identifiers, regions, technologies, fallback paths and maximum foreseeable interruption. Insurers and reinsurers can use this information to control concentration across the book.

14 Address systemic and state-backed events

Cyber attacks can spread beyond one insured or involve state actors. Lloyd's market requirements emphasise clear treatment of war and state-backed cyber attacks, including attribution and territorial consequences. Space infrastructure can be strategically significant, which increases the relevance of these questions.

Policy language should define covered events, exclusions and attribution. The operator should understand how a state-backed attack exclusion interacts with business interruption, physical damage, technology liability and other policies. Silence can create uncertainty rather than protection.

Telemetry can establish timing, affected systems and observed behaviour. It may not establish who directed an attack. Attribution can require intelligence, government assessment and legal interpretation. The evidence framework should therefore separate technical causation from state attribution.

15 Govern telemetry quality

Telemetry used for underwriting should be governed as decision data. The operator should define data owners, schemas, collection health, retention, quality checks, change approval and access. Material metric changes should have version history and explanation.

Completeness should be measured. A dashboard that reports patch cadence for eighty percent of assets should identify the uncovered twenty percent and its criticality. Missing data should not silently become zero. Time gaps should be visible, particularly around incidents and maintenance.

The submission should preserve raw evidence or verifiable extracts sufficient for review. It should avoid providing unnecessary sensitive detail. Aggregated measures can support underwriting, while detailed artefacts can remain in a controlled data room for sampled verification.

16 Protect sensitive evidence

Underwriting evidence can reveal vulnerabilities, architecture, suppliers, customers and incident history. Sharing should follow need, purpose and retention. The parties should agree secure transfer, access controls, permitted use, onward disclosure and deletion.

Legal privilege and regulatory restrictions can affect incident material. Claims and underwriting teams may require different access. The operator should identify information that can be summarised, independently attested or reviewed in a controlled environment.

An evidence-sharing protocol should also address breach of the insurer or broker. Sensitive security evidence becomes part of the risk ecosystem after transfer. The insured should understand where it is stored and which service providers can access it.

17 Build the underwriting evidence pack

The evidence pack should lead with the operating perimeter and loss scenarios. It should then provide control-performance evidence, incident and recovery records, dependency maps and financial models. Policies and certifications provide context. Executed evidence shows performance.

Each finding should carry source, period, scope, owner, interpretation and limitation. Material exceptions should include remediation, funding and due date. The pack should explain changes from the previous period so the underwriter can evaluate direction as well as current state.

A data room can support sampled verification. The insurer may inspect privileged-access reviews, vulnerability exceptions, exercise records, supplier contracts, service maps and loss calculations. The review protocol should preserve confidentiality and operational safety.

18 Connect evidence to policy structure

Evidence can affect risk acceptance, premium, limit, retention, waiting period, sublimit, exclusion, warranty, condition and service. The response should match the identified exposure. A recovery gap may affect the business-interruption waiting period. A shared-provider concentration may justify a dependent-business-interruption sublimit. An ungoverned command path may require remediation before capacity is bound.

Policy wording should reflect the technology stack. Definitions of computer system, data, network interruption, dependent provider, physical damage and service failure should be tested against spacecraft and ground operations. Technology errors and omissions may respond to customer claims that a first-party cyber policy does not.

The operator should model the retained risk after policy terms. A headline limit can overstate practical protection when sublimits, waiting periods, exclusions and proof requirements apply. The board should see gross loss, expected policy response, retained loss and liquidity need separately.

19 Integrate cyber insurance with M&A diligence

An acquirer should review current policies, applications, representations, claims, circumstances, broker correspondence and renewal history. It should compare submitted statements with diligence evidence. A material inconsistency can affect coverage and transaction allocation.

The buyer should identify change-of-control terms, runoff protection, prior-acts dates, claims-made requirements and notice deadlines. Known events should be assessed before closing. The target may need to notify circumstances or purchase extended reporting protection.

The telemetry evidence pack can support both technical diligence and insurance placement. It should not be reused without checking purpose and definitions. Transaction diligence may focus on enterprise-value consequence. Insurance underwriting focuses on covered loss during the policy period.

20 Translate gaps into transaction mechanics

Material gaps should have a transaction response. Options include price adjustment, escrow, indemnity, condition, covenant, specific insurance, retention plan, funded remediation or excluded liability. The response should be linked to evidence and release conditions.

An escrow condition should be objectively testable. Examples include privileged-path multifactor enforcement, a successful command-centre recovery test, closure of defined critical vulnerabilities or executed supplier continuity rights. Broad promises to improve cybersecurity are difficult to value and enforce.

Representations should correspond to the actual evidence. They can address incidents, notifications, access controls, material vulnerabilities, backups, testing, supplier dependencies and policy applications. Technical schedules should be reviewed by engineering, legal, insurance and finance teams.

21 Design post-close evidence continuity

Closing can disrupt telemetry and control. Systems may move, identities may change and monitoring tools may be consolidated. The buyer should preserve evidence across the cutover. Baseline exports, configuration snapshots, access lists and incident records support both operations and claims.

The first one hundred days should recertify access, confirm monitoring coverage, test recovery, validate supplier continuity, reconcile policy statements and close priority exceptions. Changes should be sequenced around mission windows and customer commitments.

The insurer and broker may need notice of material changes. Acquisition integration can alter the declared risk through network connection, identity consolidation, data migration or supplier replacement. Renewal evidence should show which controls remained effective during the transition.

22 Establish continuing assurance

Continuous assurance uses agreed evidence at a useful frequency. Daily telemetry may support operations. Monthly or quarterly summaries may support underwriting and governance. The frequency should reflect volatility and consequence.

Thresholds should trigger review rather than automatic insurance conclusions. A deteriorating patch backlog, repeated privileged-access exception or failed recovery test deserves investigation. Context determines whether the response is remediation, a policy change or an accepted risk.

Independent testing can increase confidence. Penetration tests, recovery observations, control attestations and audit samples should have defined scope. A certification provides assurance only within its stated perimeter and period.

23 Hypothetical transaction case

The hypothetical target operates twelve communications satellites, two mission-control centres and contracted ground-station capacity. It sells managed connectivity to government, mobility and enterprise customers. Management proposes USD 180 million of enterprise value. The buyer seeks USD 40 million of combined cyber and technology-liability capacity for the acquired group.

The initial submission includes policies, a control questionnaire and a recent penetration test. Diligence adds an exposure inventory, service-to-system map and twelve months of control evidence. Privileged

multifactor authentication is strong in corporate and cloud environments. Two legacy ground applications use compensating controls. Vulnerability management meets internal targets for most internet-facing systems, while three flight-support components have aged exceptions tied to limited test windows.

Incident records show two contained credential events and one supplier outage. No unauthorised spacecraft command is observed. Recovery exercises demonstrate restoration of the secondary mission-control environment in eighteen hours against a management objective of eight hours. A shared identity dependency supports both centres. Two ground-station suppliers rely on the same terrestrial carrier in one region.

The insurer's illustrative response keeps the proposed aggregate limit but applies a USD 10 million contingent-provider sublimit, an eighteen-hour business-interruption waiting period and a remediation condition for privileged ground access and identity resilience. The price and terms are management assumptions for the example, not observed market quotes.

The buyer allocates a USD 12 million remediation and continuity reserve. It places USD 8 million of consideration in escrow, released against a tested eight-hour recovery path, elimination of the shared identity failure point and closure or formally accepted treatment of the three aged flight-support exposures. The business plan reduces insured recovery in its central case and holds additional liquidity until evidence supports a revised assumption.

24 Interpret the illustrative economics

The value bridge separates enterprise value from insured loss. The USD 180 million headline value comes from the buyer's commercial model. The USD 12 million reserve funds identity resilience, ground application remediation, recovery engineering and supplier alternatives. The USD 8 million escrow allocates execution risk between buyer and seller. It is part of consideration mechanics rather than an expected claim.

The cyber programme provides defined liquidity for covered response, restoration, interruption and liability. It does not replace operational resilience. The proposed limit is tested against the hypothetical loss scenarios, sublimits and waiting period. Retained loss includes the waiting-period effect, uncovered provider concentration, policy exclusions and amounts above the applicable limit.

The combined decision framework allows the board to see four values: enterprise value, remediation cash, expected insurance response under stated scenarios and residual liquidity need. Each value has a different evidence basis. Combining them into one risk score would reduce transparency.

25 Board decision framework

The board should require a stable perimeter, material loss scenarios, credible control evidence, tested recovery, dependency mapping and policy-response analysis. It should see material limitations and unresolved exceptions. The decision pack should identify which conditions must be met before signing, closing, binding cover and releasing escrow.

The board should also require ownership. Security teams own control evidence. Operations owns mission continuity. Finance owns loss modelling and liquidity. Legal owns wording and notification. Insurance advisers support placement and claims preparation. Transaction leaders connect these workstreams to value and contractual allocation.

Approval should be conditional when evidence is incomplete in a material area. A condition can define the required test, result, funding and consequence. This approach preserves accountability and avoids treating the absence of observed incidents as proof of low risk.

26 Prepare for claims before an incident

Claims readiness begins before loss. The operator should understand notification requirements, consent provisions, panel providers, proof of loss, record preservation and cooperation duties. Incident plans should contain current broker and insurer contacts and distinguish immediate operational action from policy communication.

The evidence architecture should preserve a defensible timeline. Security events, command records, system changes, customer impact, restoration work and cost decisions should use synchronised time where practical. The operator should record who authorised emergency expenditure and why it was reasonable. This record can support operational learning, regulatory reporting and claim preparation.

Cost coding should distinguish forensic response, restoration, replacement, betterment, additional expense, lost income, customer remedies and legal cost. Coverage may differ by category. Finance teams should establish project codes and approval routes that can operate during an incident. Suppliers should provide invoices and work descriptions at a level that supports allocation.

Business-interruption claims require a counterfactual. The operator should document normal service, forecast assumptions, seasonality, customer concentration, mitigation and saved costs. Mission and customer telemetry can help show which services were unavailable and when they returned. The financial model should reconcile to accounting records and avoid counting deferred revenue, lost enterprise value and insured income as if they were the same measure.

27 Connect insurance evidence to financing

Lenders and project-finance providers can depend on cyber resilience when debt service relies on continuous availability, contracted capacity or accepted data delivery. The financing review should identify insurance covenants, minimum cover, lender loss-payee rights, notice requirements and restrictions on policy amendment. A cyber sublimit or waiting period can be material even when the policy's aggregate limit satisfies a headline covenant.

The debt model should test cash available for debt service after a defined interruption and assumed policy response. Timing matters. A valid claim may be paid after debt service is due. Liquidity reserves, revolving facilities and sponsor support can therefore remain necessary. The lender should understand the difference between insured recovery and immediate cash.

Security and account-control arrangements also matter. Insurance proceeds may flow through controlled accounts. The documents should state whether proceeds fund restoration, prepayment or both. An operator can preserve enterprise value by restoring service, while a lender may seek debt reduction after a major loss. The parties should resolve this allocation before an incident.

In an acquisition financing, the lender can use the same evidence pack to evaluate integration risk. Material conditions, escrow and remediation funding should be reflected in sources and uses. The financing case should not assume that cyber cover cures an untested recovery path or a shared provider concentration.

28 Govern underwriting models and external data

Underwriters can combine operator evidence with external scanning, threat intelligence, industry data and portfolio models. These tools can improve consistency and identify exposure outside the questionnaire. They can also create errors when assets are misattributed, services are shared, scans observe a supplier rather than the insured, or historical identifiers remain linked after a transaction.

Material external findings should be validated with the operator. The evidence record should preserve data source, observation date, confidence, affected asset and disposition. A disputed finding should be resolved through targeted verification rather than being silently removed or accepted.

Model governance should define purpose, inputs, limitations, change control and human review. A score used for triage may be unsuitable for pricing or capacity. Portfolio aggregation models require consistent dependency identifiers and clear treatment of missing data. Scenario output should remain distinguishable from observed loss.

Acquirers should ask how an ownership change affects external data and insurer models. Domains, cloud accounts, subsidiaries and service providers may be reassigned. The buyer should ensure that the renewed evidence pack reflects the post-close perimeter and that legacy observations are interpreted in context.

Conclusion

Cyber insurance for space infrastructure requires evidence that follows the service from asset and identity through control, event, recovery and cash consequence. Telemetry can strengthen this chain when collection is complete, governance is credible and interpretation is linked to defined loss scenarios.

The practical programme begins with the operating perimeter. It measures control performance on critical paths, tests restoration of services, maps common dependencies and reconciles insurance language with the technical architecture. It preserves the difference between observed evidence and management assumptions.

For acquisitions, the same evidence can support diligence, price allocation, closing conditions, remediation and insurance placement. The board receives a more disciplined view of enterprise value, insured response and retained liquidity. The result is a decision system that can evolve with the constellation and its risk.

Appendix A Evidence design principles

An evidence measure should be relevant, attributable, complete, timely, comparable and reviewable. Relevance connects it to a control or loss question. Attribution identifies the source and accountable owner. Completeness shows scope and gaps. Timeliness matches the decision period. Comparability preserves definitions across time. Reviewability allows sampled verification.

Measures should be tested for gaming and drift. A target based on closure counts can encourage low-value closure. A target based on average age can hide critical outliers. Risk-weighted distributions and exception records provide better context.

The evidence pack should record version changes. New sensors or asset-discovery improvements can make a metric appear worse because coverage improved. That change can be positive even when the reported count increases.

Appendix B Loss modelling principles

Scenario analysis should identify event, affected services, duration, restoration path, direct costs, customer remedies, liability, policy response and retained liquidity. It should combine dependencies that can fail together. A cloud outage and identity failure may extend recovery beyond either standalone assumption.

Frequency estimates require credible data and careful segmentation. This paper does not provide an actuarial frequency estimate. The hypothetical case demonstrates cash and control decisions under stated events. Insurers should apply their own portfolio data, models, judgement and regulatory requirements.

Stress cases should test long restoration, supplier aggregation, customer concentration and disputed coverage. Management actions should have funding, authority and executable timing.

Appendix C Underwriting data room

The data room can contain the operating perimeter, asset register, architecture, control policies, twelve-month control evidence, incident register, recovery exercises, supplier map, customer service obligations, financial-loss model, current policies, applications, claims correspondence and remediation plans.

Sensitive artefacts should be access controlled. The underwriter can sample source evidence through a review protocol. Requests should be proportionate to the insured decision and avoid operational interference.

The final submission should contain an exception schedule. Each exception records consequence, compensating control, owner, funding, target date and requested insurance treatment.

Appendix D Decision figures and tables

Figure 1. Space cyber underwriting evidence chain

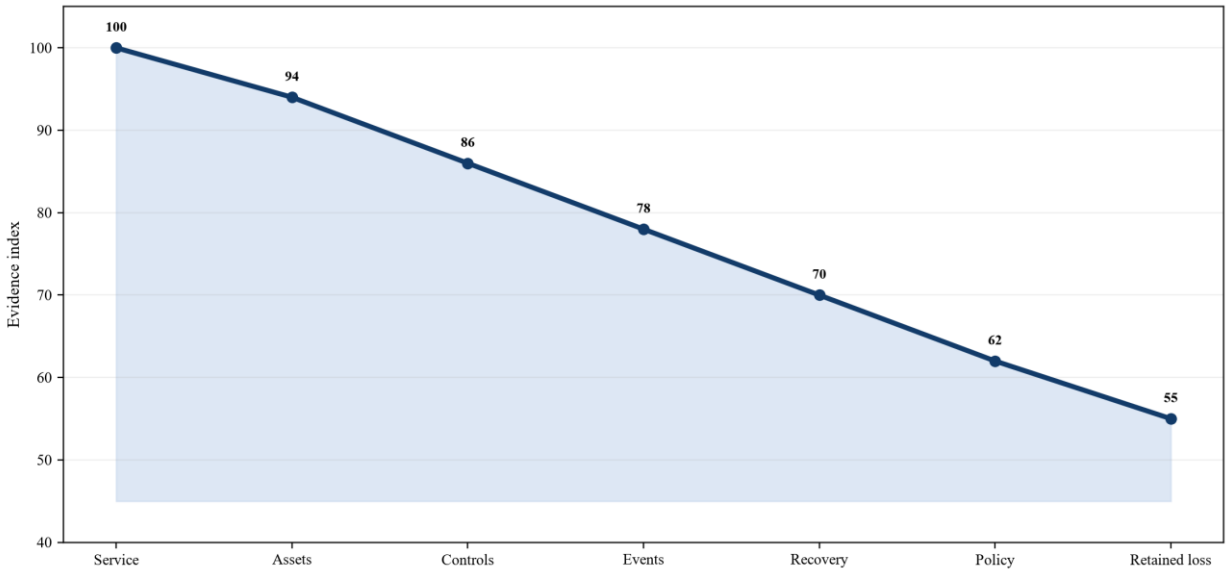

Proposed evidence progression from service perimeter to retained financial exposure.

Figure 2. Hypothetical control-performance profile

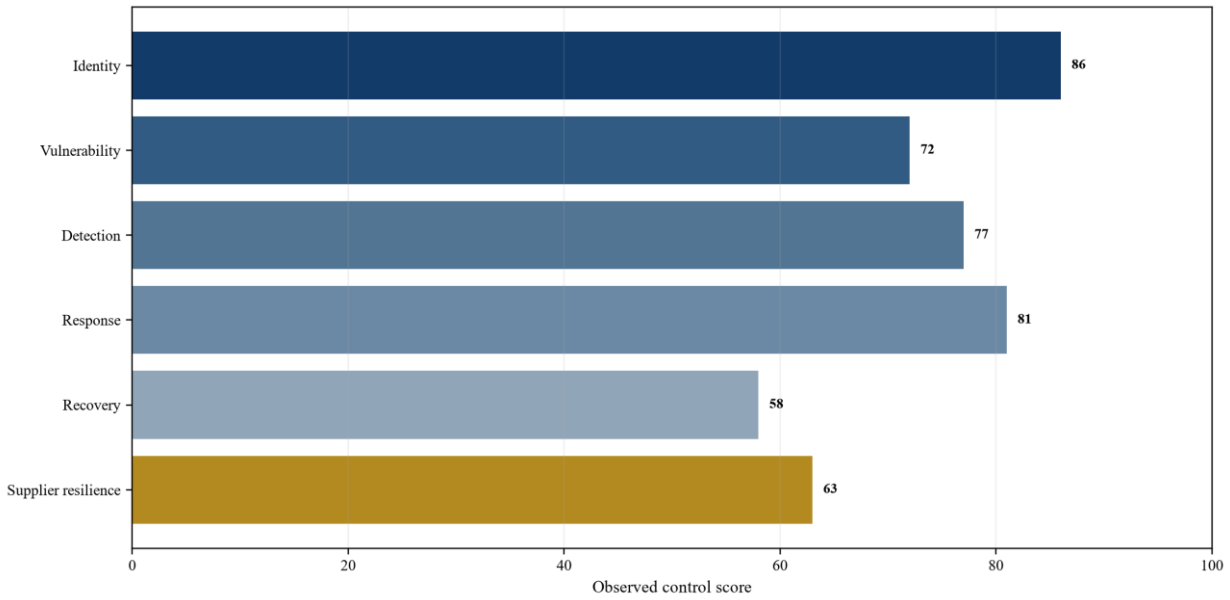

Wholly hypothetical scores; higher values indicate stronger observed performance.

Figure 3. Hypothetical gross and retained interruption loss by duration

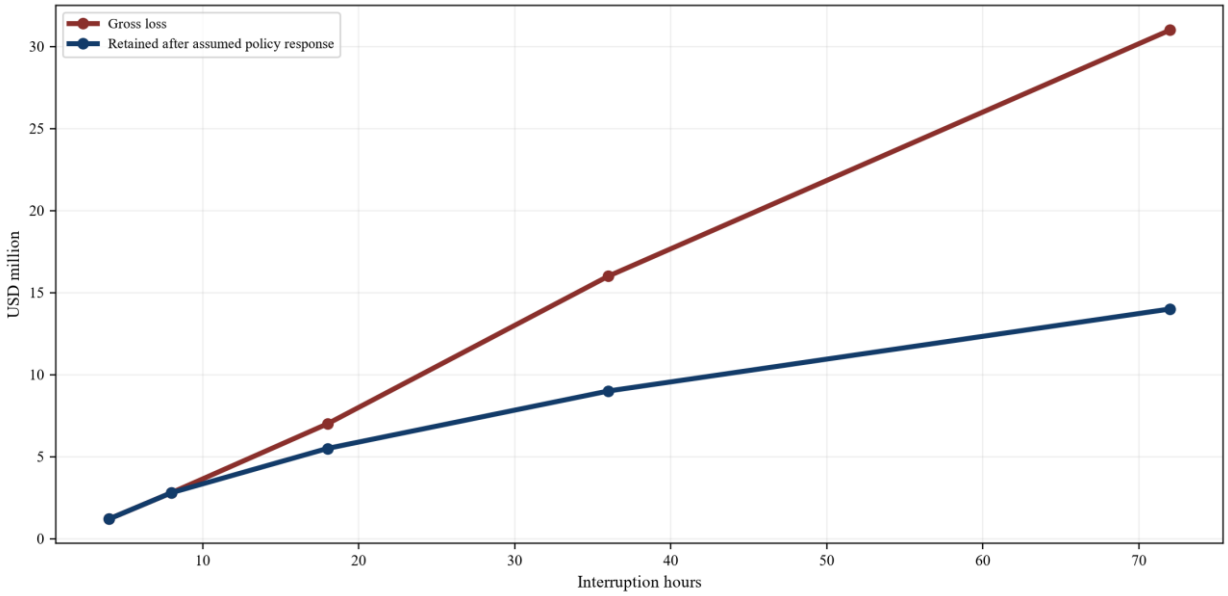

Wholly hypothetical management assumptions; USD million.

Figure 4. Hypothetical transaction and insurance bridge

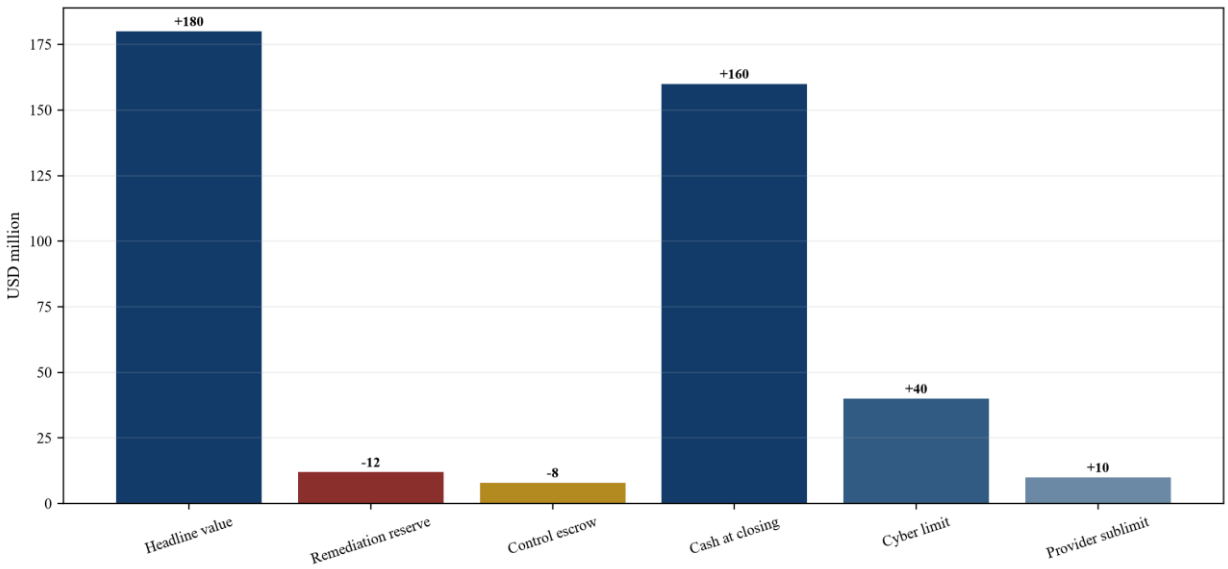

Wholly hypothetical values; USD million.

Figure 5. Illustrative first 100-day evidence programme

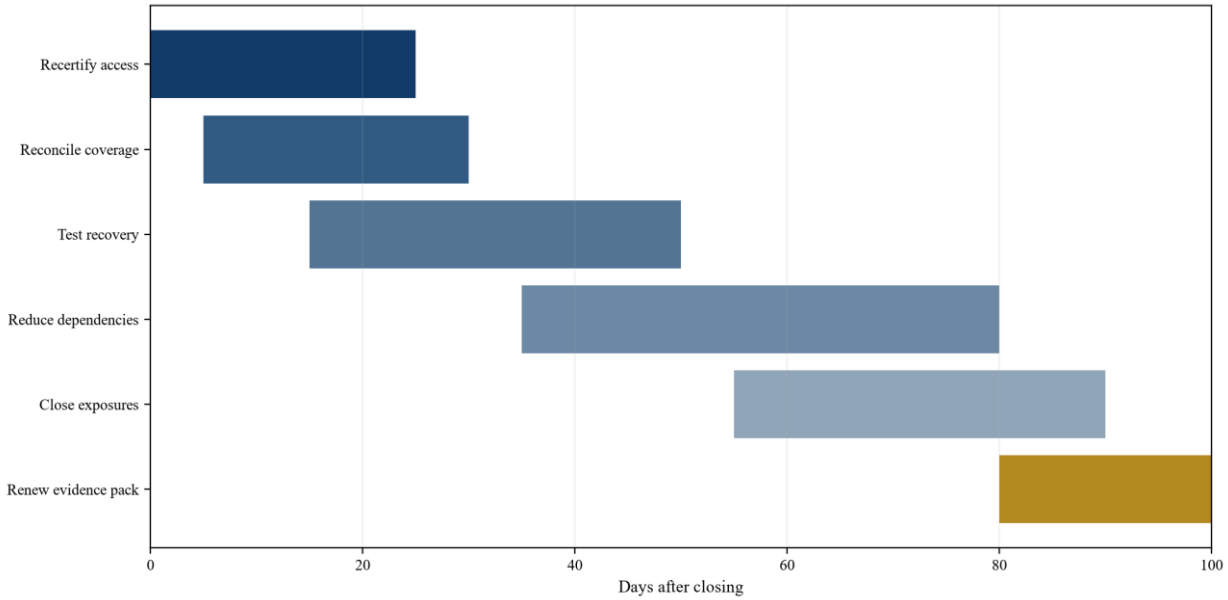

Proposed sequence subject to mission windows and customer obligations.

Table 1. Telemetry-to-underwriting evidence map

Evidence area	Operational measure	Underwriting question	Limitation to disclose
Identity	Privileged paths with enforced multifactor access	Can material control paths be protected and attributed?	Coverage gaps and emergency accounts
Vulnerability	Risk-weighted exposure age	Is material exposure reduced within mission constraints?	Detection scope and patch-window exceptions
Detection	Critical systems with healthy telemetry	Can material activity be identified promptly?	Sensor blind spots and changing definitions
Recovery	Observed service restoration time	Can covered operations resume within assumptions?	Test scope and untested suppliers
Dependencies	Services sharing a provider or control plane	Could one event create correlated loss?	Second-order dependencies and contract opacity

Proposed minimum evidence design.

Table 2. Space cyber loss pathways

Event	Operational consequence	Financial consequence	Coverage question
Privileged compromise	Command or monitoring interruption	Response, restoration and income loss	Definition of computer system and interruption
Supplier outage	Loss of ground or cloud service	Additional expense and dependent interruption	Named and unnamed provider treatment
Data corruption	Reprocessing and delivery delay	Restoration, service credits and customer claim	Data restoration and technology liability
Ransomware	Enterprise and mission-support disruption	Forensics, restoration and interruption	Consent, notification and exclusion terms
State-backed attack	Multi-service or systemic disruption	Large correlated loss	War and state-backed attack wording

Illustrative mapping; coverage depends on actual policy wording.

Table 3. Hypothetical insurance terms

Term	Illustrative value	Evidence driver	Retained exposure
Aggregate cyber and technology limit	USD 40 million	Service and loss scenario analysis	Loss above limit and excluded events
Provider sublimit	USD 10 million	Shared identity and carrier dependencies	Concentrated provider loss above sublimit
Interruption waiting period	18 hours	Observed recovery exercise	Loss during waiting period
Control-linked escrow	USD 8 million	Recovery and identity gaps	Seller release depends on verified cure
Remediation reserve	USD 12 million	Funded post-close programme	Cost overruns and residual risk

Wholly hypothetical management assumptions; not observed market quotes.

Table 4. Evidence quality test

Attribute	Acceptance question	Strong indicator	Weak indicator
Scope	Does the measure cover critical paths?	Risk-weighted coverage reconciled to inventory	Enterprise average without exclusions
Provenance	Can the source and method be verified?	Controlled source with accountable owner	Manual screenshot without lineage
Completeness	Are gaps and outages visible?	Missing data measured and explained	Missing data treated as zero
Comparability	Is the definition stable over time?	Versioned schema and reconciled change	Metric changed without restatement
Relevance	Does it inform loss or policy?	Explicit link to scenario and term	High-volume metric without decision use

Proposed acceptance criteria.

Table 5. M&A risk-to-mechanic mapping

Finding	Transaction consequence	Mechanic	Release evidence
Recovery exceeds plan	Greater interruption and liquidity risk	Escrow and funded remediation	Observed restoration within agreed objective
Shared identity dependency	Correlated control failure	Closing condition or reserve	Independent resilient identity path tested
Aged flight-support exposure	Exploitation and continuity risk	Specific covenant and holdback	Cure or approved compensating control
Inconsistent policy application	Coverage representation risk	Warranty and specific indemnity	Reconciled insurer disclosure
Supplier concentration	Contingent interruption and sublimit risk	Price adjustment or alternative supplier plan	Executed continuity rights and tested fallback

Proposed transaction treatments.

Table 6. First 100-day evidence dashboard

Measure	Day 30 evidence	Day 60 evidence	Day 100 evidence
Identity	Privileged inventory recertified	Legacy paths remediated	Resilient identity path tested

Measure	Day 30 evidence	Day 60 evidence	Day 100 evidence
Detection	Telemetry coverage reconciled	Critical gaps closed	Coverage and latency independently sampled
Recovery	Exercise scope agreed	Integrated test completed	Exceptions funded and accepted
Suppliers	Common dependencies identified	Continuity rights confirmed	Fallback paths tested
Insurance	Statements reconciled	Material changes notified	Renewal-ready evidence pack approved

Proposed milestones.

Table 7. Board decision dashboard

Decision area	Green evidence	Amber condition	Red condition
Perimeter	Services, systems and dependencies reconciled	Minor scoped gaps	Material service lacks ownership or dependency map
Controls	Critical-path performance evidenced	Exceptions funded and time bound	Control assertion unsupported on a critical path
Recovery	Observed restoration meets assumption	Gap protected by liquidity and plan	Critical service recovery untested
Policy	Technical architecture matches wording	Defined sublimits and retained risk	Material ambiguity or disclosure inconsistency
Transaction	Price, escrow and conditions use objective evidence	Residual gap accepted by board	Unfunded material exposure at closing

Proposed decision thresholds.

Frequently asked questions

Q: Can operational telemetry replace a cyber insurance application? A: Telemetry can strengthen an application by showing exposure and control performance over time. The application, policy wording, financial model, incident history and management representations remain necessary. The evidence should be reconciled rather than treated as a substitute.

Q: Which telemetry is most useful to an underwriter? A: Useful evidence covers the material loss pathways. This commonly includes privileged access, vulnerability exposure, monitoring coverage, incident response, recovery performance, service availability and shared dependencies. Its value depends on scope, provenance and interpretation.

Q: Does a high security score prove that a space operator is insurable? A: No. A score can summarise selected measures. It may omit policy terms, systemic dependencies, service consequences, data gaps and untested recovery. Underwriting requires evidence and judgement across the defined risk.

Q: How should satellite telemetry be shared without exposing sensitive operations? A: The parties can use aggregated measures, controlled data-room access, sampled verification and independent attestations. The protocol should define purpose, access, retention, onward disclosure and deletion.

Q: How does cyber insurance affect an acquisition price? A: Insurance can provide liquidity for covered losses. It does not remove remediation cost, waiting-period loss, exclusions, sublimits or operational risk. The buyer should model policy response and retained exposure separately from enterprise value.

Q: Why does aggregation matter for space infrastructure? A: Operators can share cloud regions, identity services, ground stations, terrestrial carriers, software components and suppliers. One event can therefore affect several services or insureds. Dependency mapping supports capacity and sublimit decisions.

Q: What should be placed in a control-linked escrow? A: The escrow should correspond to material and objectively testable gaps. Examples include a recovery-time target, elimination of a critical shared dependency, verified privileged-access control or closure of defined exposures. Release criteria should be agreed before signing.

Q: What should the board see before approving coverage or an acquisition? A: The board should see the operating perimeter, material loss scenarios, control and recovery evidence, dependency aggregation, policy-response analysis, remediation funding, retained liquidity and objective transaction conditions.

References

[1] National Institute of Standards and Technology, Cybersecurity Framework 2.0, 2024. https://www.nist.gov/cyberframework

[2] National Institute of Standards and Technology, SP 1303 Cybersecurity Framework 2.0 Enterprise Risk Management Quick-Start Guide, 2024. https://csrc.nist.gov/pubs/sp/1303/final

[3] National Institute of Standards and Technology, IR 8401 Satellite Ground Segment Applying the Cybersecurity Framework to Satellite Command and Control, 2022. https://csrc.nist.gov/pubs/ir/8401/final

[4] National Institute of Standards and Technology, IR 8441 Cybersecurity Framework Profile for Hybrid Satellite Networks, 2023. https://csrc.nist.gov/pubs/ir/8441/final

[5] National Institute of Standards and Technology, IR 8270 Introduction to Cybersecurity for Commercial Satellite Operations, 2023. https://csrc.nist.gov/pubs/ir/8270/final

[6] National Institute of Standards and Technology, SP 800-53 Revision 5 Security and Privacy Controls for Information Systems and Organizations. https://csrc.nist.gov/pubs/sp/800/53/r5/upd1/final

[7] National Institute of Standards and Technology, SP 800-61 Revision 2 Computer Security Incident Handling Guide. https://csrc.nist.gov/pubs/sp/800/61/r2/final

[8] National Institute of Standards and Technology, SP 800-34 Revision 1 Contingency Planning Guide for Federal Information Systems. https://csrc.nist.gov/pubs/sp/800/34/r1/upd1/final

[9] National Institute of Standards and Technology, IR 8286 Integrating Cybersecurity and Enterprise Risk Management. https://csrc.nist.gov/pubs/ir/8286/final

[10] European Union Agency for Cybersecurity, ENISA Space Threat Landscape 2025. https://www.enisa.europa.eu/publications/enisa-space-threat-landscape-2025

[11] European Union Agency for Cybersecurity, Low Earth Orbit SATCOM Cybersecurity Assessment, 2024. https://www.enisa.europa.eu/publications/low-earth-orbit-leo-satcom-cybersecurity-assessment

[12] European Union Agency for Cybersecurity, NIS2 Directive resources. https://www.enisa.europa.eu/topics/state-of-cybersecurity-in-the-eu/cybersecurity-policies/nis-directive-2

[13] National Aeronautics and Space Administration, Space Security Best Practices Guide. https://www.nasa.gov/general/nasa-issues-new-space-security-best-practices-guide/

[14] National Aeronautics and Space Administration, NASA-STD-1006A Space System Protection Standard. https://standards.nasa.gov/node/2430

[15] United States Office of Space Commerce, Space Policy Directive 5 Cybersecurity Principles for Space Systems. https://space.commerce.gov/president-signs-space-cybersecurity-policy-directive/

[16] Cybersecurity and Infrastructure Security Agency, Cross-Sector Cybersecurity Performance Goals. https://www.cisa.gov/cross-sector-cybersecurity-performance-goals

[17] Cybersecurity and Infrastructure Security Agency, Federal Government Cybersecurity Incident and Vulnerability Response Playbooks. https://www.cisa.gov/news-events/news/federal-government-cybersecurity-incident-and-vulnerability-response-playbooks

[18] United States Securities and Exchange Commission, Cybersecurity Risk Management Strategy Governance and Incident Disclosure, 2023. https://www.sec.gov/rule-release/33-11216

[19] National Association of Insurance Commissioners, Report on the Cybersecurity Insurance Market, 2025. https://content.naic.org/sites/default/files/inline-files/2025_Cybersecurity_Insurance%20Report.pdf

[20] National Association of Insurance Commissioners, Cybersecurity Insurance Topic. https://content.naic.org/insurance-topics/cybersecurity

[21] National Association of Insurance Commissioners, Innovation Cybersecurity and Technology Committee Minutes Packet, 2025. https://content.naic.org/sites/default/files/national_meeting/28_Minutes-HCmte.pdf

[22] Lloyd's, Market Bulletin Y5381 War Exclusions and Cyber-Attack Cover, 2022. https://assets.lloyds.com/media/35926dc8-c885-497b-aed8-6d2f87c1415d/Y5381%20Market%20Bulletin%20-%20Cyber-attack%20exclusions.pdf

[23] Lloyd's, Requirements for Writing War and NCBR Risks. https://www.lloyds.com/market-resources/requirements-and-standards/lloyds-requirements-of-the-council/operating-at-lloyds/performance-management/wr-nbbr-exposures

[24] European Commission, Directive EU 2022 2555 on measures for a high common level of cybersecurity across the Union. https://eur-lex.europa.eu/eli/dir/2022/2555/oj

[25] Consultative Committee for Space Data Systems, Security Working Group Publications. https://public.ccsds.org/Publications/Security.aspx

[26] International Organization for Standardization, ISO IEC 27001 Information Security Management Systems. https://www.iso.org/standard/27001

[27] International Organization for Standardization, ISO IEC 27005 Information Security Risk Management. https://www.iso.org/standard/80585.html

[28] United Kingdom National Cyber Security Centre, Cyber Security Toolkit for Boards. https://www.ncsc.gov.uk/collection/board-toolkit

[29] European Cooperation for Space Standardization, ECSS Space Engineering and Operations Standards. https://ecss.nl/standards/active-standards/

[30] Financial Stability Board, Effective Practices for Cyber Incident Response and Recovery, 2020. https://www.fsb.org/2020/10/effective-practices-for-cyber-incident-response-and-recovery/

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 close.

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

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

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

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

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