76th International Astronautical Congress (IAC), Sydney, Australia, 29 September-3 October 2025.

IAC–25–D5,1,8,x101712

ASTERISK: Space Protocol's quantum-resilient space and air safety framework for trusted operations

Samya Bagchi[a*], Harvey Reed[b], Jonathon Wotton[c], Ruth Stilwell[d], Graeme Hooper[e], Tat-Jun Chin[f], Yasir Latif[g]

[a] *Space Protocol, Adelaide, Australia, samya@spaceprotocol.org*
[b] *MITRE Corporation, USA, hreed@mitre.org*
[c] *MITRE Corporation, Australia, jwotton@mitre.org.au*
[d] *Aerospace Policy Solutions, USA, office@aerospacepolicysolutions.com*
[e] *GPSat Systems, Australia, graemeh@gpsatsys.com.au*
[f] *Australian Institute for Machine Learning, University of Adelaide, tat-jun.chin@adelaide.edu.au*
[g] *Space Protocol, Adelaide, Australia, yasir@spaceprotocol.org*
[*] *Corresponding author*

Abstract

Space operations are entering an era of heightened cyber risk and quantum uncertainty. This paper, authored from the perspective of Space Protocol, presents a mission-centric and forward-looking architecture for **quantum-resilient space system safety**. We propose a hybrid cryptographic model that combines NIST-standard **post-quantum cryptography (PQC)** – including lattice-based key encapsulation (ML-KEM) and digital signatures (ML- DSA), and stateless hash-based signatures (SLH-DSA) – with satellite-delivered **Quantum Key Distribution (QKD)** services. This dual approach fortifies communications against quantum computing threats while providing redundancy if one method is compromised. We describe how this architecture secures multi-modal space sensing networks – spanning electro-optical (EO) telescopes, active and passive RF sensors, and non-Earth- imaging assets – against adversarial attacks, data spoofing, and link compromise. A distributed ledger framework, **"ASTERISK,"** upgraded with PQC, is introduced to ensure data-to-decision traceability and integrity. ASTERISK logs cross-verified telemetry from diverse sensors, offering tamper-resistant audit trails and real-time risk quantification of observations. We illustrate a cross-domain use case where space-based EO/RF telescopes contribute to national air defense by detecting high-altitude drones or aerial objects, fused with cryptographic identity streams (e.g., FAA Remote ID for drones) to distinguish threats from friendly actors. This use case demonstrates end-to-end assurance – from sensor data capture to decision – in a dual-use (civil/military) context. Furthermore, we show how ASTERISK enables **liability-cover-backed data sales** through verifiable and tamper-proof telemetry products: data contributors can monetize high-quality space domain awareness (SDA) data on a marketplace, supported by on-demand liability insurance and smart-contract incentives that reward data integrity and compliance. Throughout, the exposition remains executive and mission-focused, highlighting alignment with sovereign capability goals, emerging cybersecurity mandates, and international safety standards (CCSDS security, NIST PQC guidelines, ITU/ETSI QKD frameworks). We discuss policy implications of an export-compliant dual-use architecture that balances national security interests with cross- border SDA data-sharing. The proposed approach underscores that quantum- resilient cybersecurity and data accountability are foundational to a safe and successful space program in the coming decade. This work is aimed at a cross-section of stakeholders who shape, operate, and safeguard space and air systems, including government space agencies and defense organizations, commercial satellite operators, ground segment providers, sensor network operators, insurers, regulators, and policy makers as well as research and standards communities.

1. Introduction

The attack surface for modern space infrastructure has expanded dramatically. *Sensors* can be spoofed or fed replayed data [1]; *data layers* can be tampered in transit or at rest [2]; *application layers* can compound subtle errors from model drift or malicious code paths; and legacy public-key cryptography (*RSA/ECC*) faces the "harvest-now, decrypt-later" (HNDL) [3] risk where adversaries record encrypted links today for retrospective decryption using future quantum computers. A concrete example highlights the risk: an attacker could tamper with a sensor via a compromised firmware update, causing it to output subtly biased tracking data that distorts conjunction assessments. Without independent verification, such manipulation can silently erode trust in the network. In this work, therefore, we make the case for an end-to-end, mission-centric *quantum-resilient* design that anticipates both classical and quantum-enabled adversaries while pre-

Fig. 1. Threat landscape for space data at different points in the data stack.

serving business viability and standards alignment. Nations worldwide are rapidly advancing their quantum computing and communication capabilities, fueling an emerging technological race parallel to traditional space pursuits. For example, China has demonstrated satellite-enabled QKD links [4], and the United States, Europe, and allied countries are heavily investing in quantum research. This global push toward quantum technology brings the threat horizon closer. Modern space programs depend on an increasingly complex web of satellites, sensors, communication networks, storage, and compute whose *quality and safety* directly impacts mission success. From collision avoidance to real-time Earth observation and military surveillance, trust in space systems' data and resilience against cyber threats is paramount. However, many space assets were launched before cybersecurity was a major concern and therefore *lack basic protections such as encryption or authentication*. As a result, today's space infrastructure remains vulnerable to malicious intrusions – hackers can intercept or spoof satellite signals, disrupt services, or even issue false commands. And this vulnerability will only grow as quantum computing advances: within the next decade, adversaries may develop quantum attacks capable of breaking traditional encryption and compromising satellite communications or control links. Ensuring *quantum-resilient cybersecurity* for space systems has thus become a national priority, with policymakers urging new standards and upgrades for satellites in this "Second Space Race" [5] for cyber-superiority.

Against this backdrop, Space Protocol proposes **ASTERISK: a quantum-resilient architecture for safe space operations** that addresses both current and future threats. This architecture (Fig. 2) is driven by two synergistic pillars: *advanced cryptography* and *distributed trust assurance*. First, we employ a hybrid cryptographic model combining Post-Quantum Cryptography (PQC) – recently standardized by NIST [6] – with Quantum Key Distribution (QKD) via satellites (Sec. 3). By integrating these approaches, we achieve defense-in-depth: PQC algorithms (Sec. 3.1) secure our communications and digital identities against quantum attacks, while QKD (Sec. 3.2) provides information-theoretic secure key exchange backed by quantum physics. This hybrid ensures that even if one scheme is weakened, the other maintains protection. Second, we introduce *a PQC-secured distributed ledger* (Sec. 4) that **logs and verifies every data transaction** from sensor measurement to final decision. This ledger provides end-to-end traceability, accountability, and tamper- evidence, allowing operators to audit how raw sensor inputs are transformed into actionable insights. This traceability is essential for quality assurance in complex multi-party, multi-sensor, multi-algorithm systems where hidden errors or data manipulations could otherwise go undetected.

The ASTERISK architecture is *mission-centric and business-conscious*. It not only hardens space missions against cyber/quantum threats but also creates new opportunities for data sharing and commerce under a safety-first paradigm. We detail how multi-modal space sensing networks (from ground-based telescopes to orbiting sensors) can be secured so that no single false signal can jeopardize a mission. We explore a cross-domain use case (Sec. 5) integrating space and air domain awareness: space-based sensors detecting uncooperative drones or high-altitude objects, combined with authenticated drone identity streams (Remote ID) for national air defense. This scenario highlights the *dual-use capability* of the system – serving both civil space safety and military air security – while remaining *export-compliant* and adhering to international norms. Finally, we discuss the **business model** implications (Sec. 6): a blockchain-backed SDA data marketplace where providers are rewarded for quality data, buyers gain confidence through liability-insured products, and market incentives drive better space safety outcomes. In an industry often lacking accountability, this approach uses economic levers (like insurance discounts for compliant behavior) to encourage cybersecurity and adherence to standards (Sec. 7), ultimately fostering a safer space environment for all stakeholders.

Role of blockchain, AI, and quantum. *Blockchain*—in our case a permissioned, PQC-upgraded ledger — anchors provenance, ordering, and non-repudiation from data capture to decision and post-event audit. *AI* augments human decision making and oversight with three functions at the application layer: causal analysis [7] (why a result), oversight AI (policy and anomaly monitors), and random performance checks [8](continuous verification). *Quantum*

technologies deliver key material with eavesdropper detectability (QKD) [9] and, together with NIST PQC [10–12], enable crypto-agility so that disruptions in one primitive do not cascade into systemic failure.

Intended Audience This work is aimed at a cross-section of stakeholders who shape, operate, and safeguard space and air systems. The primary audience includes government space agencies and defense organizations, for whom quantum-resilient architectures directly support mission assurance, sovereignty, and compliance with emerging cybersecurity mandates. It also speaks to commercial satellite operators, ground segment providers, and sensor networks, who must protect their data integrity while seeking new opportunities for data monetization. In addition, insurers, regulators, and policy makers will find value in the liability-backed marketplace model, which ties data quality to economic incentives and governance. Finally, the work is relevant to the research and standards community, as it aligns with ongoing CCSDS, NIST, ITU, and ETSI initiatives. By addressing this diverse audience, the paper positions ASTERISK not only as a technical framework, but also as an operational and policy enabler for the next era of trusted space operations.

2. Threats to Space Data Integrity and Link Security

Modern space operations increasingly rely on multi-modal sensing networks – a combination of heterogeneous sensors providing a comprehensive picture of the space environment. For example, Space Domain Awareness (SDA) systems fuse data from electro-optical (EO) telescopes, radars and RF sensors (including passive RF receivers and active tracking radars), infrared and infrasound sensors, and even non-Earth imaging telescopes (space-based sensors looking at objects in orbit). By aggregating multiple modalities, operators reduce uncertainty and gain robust situational awareness. However, the very diversity and distribution of these sensors present a security challenge: an adversary can target the weakest link – whether by cyber means or signal interference – to inject false data or degrade the network's trustworthiness. Ensuring *authenticity, integrity, and availability* of sensor data across all modalities is thus critical for the safety and quality of decisions derived from them. Adversarial threats to space sensor networks come in several forms:

2.1 Data Spoofing and Injection:

As demonstrated by recent research, attackers can spoof unprotected satellite data links to inject false measurements. For instance, "Firefly" [13] showed that if an Earth observation satellite's downlink is not cryptographically authenticated, false sensor readings (like phantom wildfire signals in infrared imagery) can be injected by broadcasting fake RF signals to the ground station. The ground systems, trusting anything coming "from the satellite" via RF, may accept these forgeries as real, leading to false alerts or misguided responses. Many legacy EO satellites (NASA, NOAA, etc.) downlink imagery in clear or with outdated encryption, making them vulnerable. Similarly, radar or RF sensors that lack secure authentication channels could be fed with synthetic echoes or fake signal recordings replayed by an adversary. Passive RF sensors, which simply listen for satellite beacons or transponders, are especially susceptible – an attacker can transmit counterfeit signals mimicking a satellite's frequency and modulation, fooling the sensor into reporting a non-existent object or event.

2.2 Link Interception and Tampering:

Without encryption and authentication, sensor data links (whether space-to-ground downlinks, ground network links, or inter-satellite links) can be intercepted and tampered. Attackers might perform man-in-the-middle alterations, modify orbital tracking data, or delay messages. A compromised ground station or relay could alter data in transit, causing subtle errors. Because multiple organizations often share data in SDA (e.g., commercial sensor data feeding into a government system), a malicious actor in the supply chain could corrupt data unless robust end-to-end security is in place.

2.3 Sensor Node Compromise:

If an adversary cyber-infiltrates one of the sensor nodes (e.g., hacking a ground-based telescope's control system or malware on a cubesat's onboard computer), they could manipulate or falsify the data at source. This insider threat is harder to detect unless the system can cross-validate outputs against other independent sources.

2.4 Jamming and Denial of Service:

Although not a data integrity attack *per se*, jamming of sensor communications can create data availability issues or attempt to mask real events. While our focus in this work is on data assurance, we note that robust cryptography (like spread-spectrum modulation with secret keys) can also aid in anti-jamming by making it harder for adversaries to target the signal.

2.5 Attacks against Centralized compute

Centralized compute infrastructures introduce another avenue for adversarial manipulation. If an attacker compromises a central processing hub, they can inject malicious code paths, bias analytic models, or distort outputs, potentially altering collision avoidance or air defense assessments without modifying upstream data. Insider

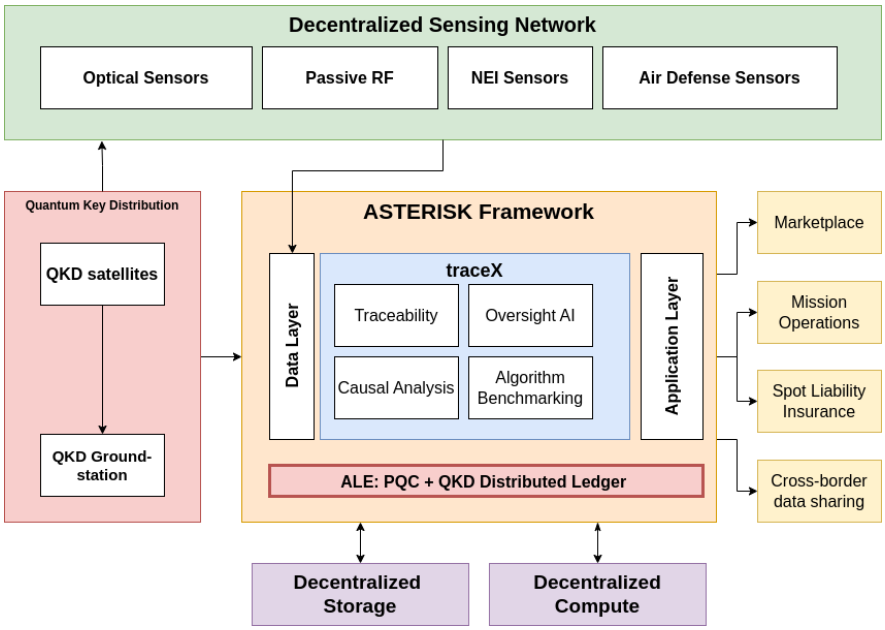

Fig. 2. Space Protocol's ASTERISK framework consumes data from decentralized sensors, is secured by a hybrid QKD + PQC cryptography architecture and enables various applications.

threats and opaque cloud providers amplify this risk, as privileged access or concealed breaches can silently degrade mission assurance.

2.6 Attacks against Centralized storage

Centralized storage architectures present a critical weakness for space operations. A single compromised cloud instance or mission server can enable adversaries to tamper with logs, inject false telemetry, or erase audit trails, undermining trust across the decision chain. Centralized points of failure are also prime targets for ransomware, insider threats, and nation-state campaigns. In a "harvest-now, decrypt-later" context, bulk downloads of centralized archives could be retrospectively decrypted by quantum adversaries, exposing decades of mission data.

3. Hybrid Quantum-Resilient Cryptographic Architecture

Space Protocol's security architecture is built on a **hybrid quantum-safe cryptographic model** that combines **post-quantum cryptography (PQC)** algorithms with **quantum key distribution (QKD)**. This model ensures that all communications and control links in space operations remain secure against both classical and quantum attacks, thereby guaranteeing mission safety even as adversary capabilities evolve.

3.1 NIST PQC Algorithms

The foundation of our model is **Post-Quantum Cryptography** (PQC), consisting of algorithms believed to be secure against quantum computing attacks. In 2022, NIST announced the first standardized PQC algorithms [10–12] after a global competition. Space Protocol adopts the NIST selections for both key encapsulation and digital signatures, specifically:

3.1.1 Module-Lattice Key Encapsulation (ML-KEM):

We employ **CRYSTALS-Kyber**, a lattice-based KEM [10], to perform key exchanges for symmetric encryption. Kyber's security relies on the hardness of lattice problems that quantum algorithms cannot efficiently solve. It provides strong confidentiality for exchanging session keys used in data links and telemetry, tracking & control (TT&C) channels.

3.1.2 Module-Lattice Digital Signatures (ML-DSA):

For authentication of commands, software, and telemetry, we use **CRYSTALS- Dilithium** (and/or the similar FALCON algorithm) as our primary digital signature scheme. These lattice-based signatures [11] ensure that control commands to satellites and data from sensors are provably from legitimate sources. Even a future quantum adversary cannot forge these signatures under current mathematical understanding.

3.1.3 Stateless Hash-based Digital Signatures (SLH-DSA):

We incorporate **SPHINCS+** [12], a stateless hash-based signature scheme, as a backup and for specialized use (such as code signing for firmware updates). Hash-based signatures like SPHINCS+ rely only on the security of cryptographic hash functions, offering a different hardness assumption immune to advances in number theory. Although SPHINCS+ signatures are larger and slower, their conservative design adds diversity to our cryptographic arsenal.

By combining lattice-based and hash-based approaches, our architecture is resilient even if one family of algorithms is later found vulnerable. This reflects a crypto-agile philosophy: multiple primitives can be used in parallel or interchangeably, and the system can be updated to new algorithms in the future. For bulk data encryption, we continue to use symmetric AES-256 [14] (per current standards) for throughput efficiency, as its 256-bit key is considered quantum-safe [15].

Importantly, all cryptographic components are chosen to comply with emerging standards (see Sec. 7 for details). Our cryptographic layer "drops in" to existing space communication stacks, upgrading them for the quantum era without sacrificing interoperability or performance more than necessary.

3.2 Satellite-Delivered QKD Services

In parallel to algorithmic PQC, our architecture leverages Quantum Key Distribution (QKD) delivered via satellites as an advanced service for critical links. QKD uses quantum physics (typically single photons) to distribute encryption keys between distant nodes with *provable eavesdropping detection*. If any adversary tries to intercept the quantum bits, the quantum states are disturbed and the intrusion is noticed – a property guaranteed by the laws of physics, not just computational complexity.

Recent demonstrations underscore QKD's maturity: for instance, China has operational QKD satellite links and completed a 1,000 km secure call by integrating QKD and classical networks [4]. In our design, satellites equipped with QKD transmitters (or ground stations with QKD terminals) can distribute secret keys that will be used to one-time-pad encrypt especially sensitive data or to frequently refresh the encryption keys of classical communications. QKD-generated keys provide information-theoretic security for the data protected by them – even a far future quantum computer cannot retroactively decrypt one-time-pad data if it didn't intercept and alter the key distribution at the time. This is particularly relevant for long-duration confidentiality (e.g., archival science data or strategic military communications that must remain secret for decades).

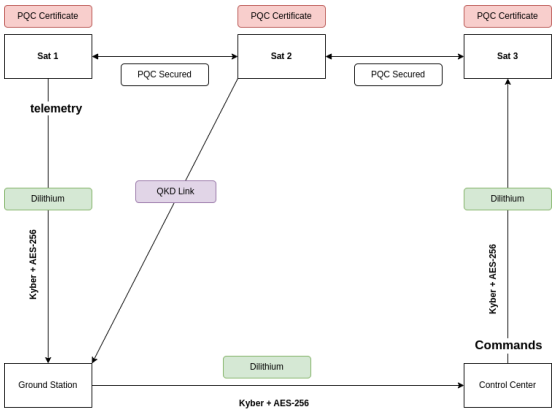

Fig. 3. Hybrid cryptographic architecture

We integrate QKD into a *hybrid key management framework*: PQC and QKD operate in tandem. In practice, this could mean using QKD whenever available (such as between major ground stations and relay satellites) to augment the key supply, and defaulting to PQC key exchange for all other links or as a fallback. In some cases, we use **hybrid keys**, combining a QKD-derived key with a PQC-derived key to encrypt data. This way, an attacker needs to break both methods – an unlikely dual failure – to compromise the key. Hybridization also hedges against uncertainties: QKD has distance and rate limitations and requires specialized hardware, whereas PQC is software-friendly but theoretically could be broken by new quantum algorithms. **By merging their outputs, we achieve a composite key that inherits the strengths of both**. The result is an "end-to-end quantum secure" system wherein communications are safe from immediate threats and protected against future quantum advances.

It is important to note that our QKD framework aligns with international standards (ITU-T and ETSI) to ensure interoperability and trust. For example, we utilize QKD protocols and interfaces recommended by ITU-T SG13/SG17 and ETSI's QKD Industry Specification Group, making our system compatible with emerging global QKD networks. Key management is designed to integrate with existing infrastructure – QKD keys are fed into standard encryption devices (via PQC-protected channels for post-processing) to encrypt data just as any symmetric key would. This seamless integration means mission operators do not need quantum expertise; the QKD service operates behind the scenes, managed by specialized nodes, while the end users simply enjoy stronger keys.

Fig. 3 illustrates our hybrid cryptographic architecture. Space assets (satellites, ground stations, control centers) each possess PQC key pairs and certificates, forming a

web of verifiable identities. Communication sessions begin with a Kyber post-quantum key exchange (or a QKD session if available) to establish shared secrets, which are then used for AES-256 encrypted links. All command uplinks and telemetry downlinks are signed using Dilithium (or SPHINCS+ for specific cases) to guarantee authenticity. QKD- distributed keys can periodically refresh these link keys or provide an additional one-time-pad layer for critical messages. This multi-layered approach ensures no single point of cryptographic failure: even if a new cryptanalytic attack or a flaw in one algorithm is discovered, the system's security would remain intact due to the complementary mechanisms in place. Such robustness is essential for long-term safety of space missions, which often operate for decades and must withstand technological upheavals in the threat landscape.

3.3 Cryptographic Identity and Authentication for Sensors

To counter spoofing and tampering, every sensor or data source in our network is equipped with a cryptographic identity anchored in PQC. Concretely, each sensor platform is provisioned with a unique private-public key pair. The public key, along with metadata about the sensor (owner, location, type), is certified by a trusted authority and distributed to all network participants. This forms a **web of trust** for sensor identities.

Whenever a sensor collects an observation – be it an image, a track, a signal detection – it immediately signs the data with its private key before transmitting. The signature is appended to the data packet or meta-data. Downstream, any receiving station or processing node will verify this signature using the sensor's public key (and PQC verification, which is computationally efficient). Data that fails verification (signature incorrect or unknown identity) is *rejected as untrusted.* In effect, this mechanism means an adversary cannot impersonate a legitimate sensor or alter its data without detection, unless they somehow steal the sensor's private key (which is protected in hardware and further shielded by our PQC+QKD key management). The **"trust anchor" moves from the communication channel to the data itself** – each piece of sensor data carries its own proof of origin and integrity. This counters scenarios Firefly [13]; even if an attacker injects a radio signal, they would lack the sensor's private key to sign false data, so the system would discard that input as forged.

Additionally, all sensor data transmissions are end-to-end encrypted (using symmetric AES-256 keys established via our PQC/QKD hybrid scheme). This prevents adversaries from even reading the sensor data or subtly modifying it in transit. Encryption also thwarts certain spoofing where the attacker aims to overshadow a real signal; if they cannot decrypt and understand the real signal, they cannot seamlessly override it. By using frequent key updates (enabled by QKD's continuous key supply or periodic Kyber exchanges), even a short compromise of one key will have limited impact.

Multi-layer authentication is another aspect: not only is the data signed at source, but any processing step (aggregation, fusion nodes) also signs the outputs it generates [8]. This creates a chain of custody for the data through the network. For example, if a surveillance system merges optical and RF observations into a single track, the fusion algorithm will sign the fused track report. This way, the final outputs have a *traceable pedigree* back to all original sources (we explore this further in Sec. 4).

3.4 Cross-Validation and Consensus Among Sensors

Cryptography aside, a strength of multi-modal networks is the ability to cross-validate data using independent sources. Our architecture leverages this by establishing a **consensus framework** for critical situational awareness decisions. In practice, the ASTERISK framework's ledger (Sec. 4) acts as the medium where observations from different sensors are reported and compared. If one sensor reports an anomaly (e.g., a new object or an event) but three other sensors observing the same volume of space see nothing, the system can flag the single-source report as potentially spurious. Conversely, when multiple sensors corroborate an observation – say an optical telescope sees a glint at a certain coordinate and a passive RF sensor simultaneously detects an unregistered transmission there – the confidence in that event's reality increases. Such a system cannot operate effectively without the ability of actively tasking sensors [16, 17] for validating information when needed.

The consensus algorithm in our system does not necessarily mean majority voting on every bit of data; rather, it is about combining evidence to achieve a *trust score or risk quantification* for each event. For example, ASTERISK might compute a credibility score that considers how many independent sensors confirmed an event, the reliability history of those sensors, and the consistency of the data. This quantitative approach to **data quality** directly ties into safety: operators can set thresholds (only act on high-confidence data) and insurers can price risk (lower premium for highly verified data streams).

Cross-validation also helps detect subtle compromises. If an adversary somehow subverts one sensor (imagine a malicious firmware update on a radar), that sensor might start outputting slightly biased data. Because other sensors will not show the same bias, the discrepancy can be caught by analytics (Section .4) that run on the aggregated ledger data – effectively performing outlier detection. This is analogous to having "multiple witnesses" in

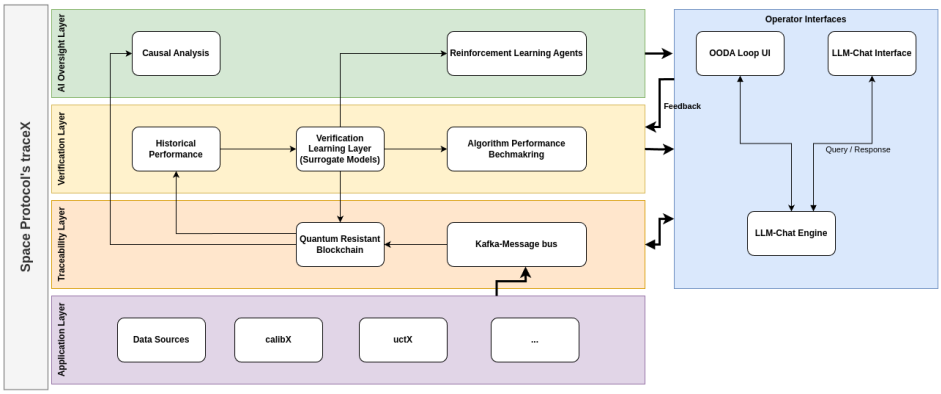

Fig. 4. TraceX [8] enables data and decision provenance along with higher level AI-based causal analysis, benchmarking and verification.

a system: one false witness is exposed by the consensus of truth from others.

3.5 Decentralized Storage

Decentralized storage mitigates risks due to a central storage by distributing data across multiple nodes with cryptographic redundancy and consensus validation. No single actor can unilaterally alter or delete mission-critical information, and tampering at one node becomes detectable through cross-verification. For ASTERISK, this model complements the ledger (Sec. 4) by providing scalable, fault-tolerant repositories where raw sensor outputs, provenance chains, and risk metrics remain verifiable over the long term. Active projects such as IPFS, Filecoin, Storj, Arweave, and Siacoin demonstrate practical paths for resilient, decentralized storage that could be integrated with PQC-secured access controls to strengthen end-to-end assurance.

3.6 Decentralized Compute

Decentralized compute distributes analytic tasks across diverse nodes, with cryptographic consensus ensuring result integrity. Instead of relying on a single data center, ASTERISK cross-validates computations performed by multiple providers, each producing PQC-signed attestations. This reduces trust in any one operator, makes targeted tampering more difficult, and ensures resilience against outages or censorship. It aligns with ASTERISK's guiding principle: traceability not only of data, but of the algorithms and compute pipelines that shape decisions.

Emerging platforms such as Golem Network, iExec, Akash Network, Ankr, and Render Network showcase viable models for decentralized compute. Integrating distributed computing allows analytic results to be tied to verifiable compute provenance, ensuring that mission-critical decisions are both data-verified and computation-verified.

3.7 Quality & Safety in Non-Earth Imaging

A special mention is warranted for non-Earth imaging sensors, such as space-based telescopes tracking satellites or debris. These sensors often operate in challenging conditions (tracking fast-moving objects against star backgrounds) and their data is used for collision avoidance warnings – a safety-critical function. Our security architecture ensures that collision avoidance data (conjunction assessments) cannot be spoofed and are tamper-evident. Each positional observation of a resident space object (RSO) is signed and logged; any maneuver or orbit change broadcast by a satellite (via a trusted channel or the Space Traffic Management service) is authenticated by the satellite's own PQC keys. This prevents scenarios where a fake conjunction alert could be injected to force unnecessary evasive maneuvers, or conversely, a real alert suppressed. It also provides non-repudiation – satellite operators who fail to follow guidelines can be held accountable with evidence (their maneuvers or lack thereof are on record in ASTERISK, tied to their cryptographic identity).

By securing each link and each data point, and ensuring storage and compute authenticity, we create a sensor network that is resilient to adversarial interference, enhancing both the quality (accuracy, reliability) and safety (preventing false decisions) of space operations.

4. ALE: The ASTERISK Ledger

While strong cryptography and sensor protections guard against external threats, complex space systems also demand *holistic traceability and accountability* internally. Space missions often involve automated decision pipelines – from sensor data ingestion to AI algorithms to human operator actions. Ensuring the correct-

ness of each step, and the ability to audit and replay decisions, is vital for safety, especially when things go wrong. To address this, Space Protocol's ASTERISK incorporates a **PQC-upgraded permissioned distributed ledger** named, ALE, designed for data and decision traceability [8] across the space enterprise. This ledger serves as the **single source of truth** that links all inputs, algorithms, and outputs, thereby assuring that decisions are based on verifiable, tamper-resistant data.

4.1 Ledger Overview and PQC Integration

The ASTERISK ledger (ALE) is a custom distributed ledger tailored to space operations. It takes inspiration from blockchain networks (like Ethereum or Hyperledger) but is optimized for handling time-series sensor data, event logs, and analytic results, rather than financial transactions. **Every node in the space network (sensors, satellites, control centers, data centers)** can act as a ledger client, submitting "transactions" that record observations or decisions. These transactions are cryptographically linked (hashed) and validated by a consensus mechanism, creating an *immutable, chronological chain of events*. Unlike public blockchains, ALE is a permissioned ledger: participants are authorized entities (space agencies, commercial operators, defense networks, insurers) with known identities. This ensures sensitive data remains within a trust community while still enabling decentralized validation.

Crucially, ALE is built to be quantum-resilient. All digital signatures, hash functions, and cryptographic protocols within the ledger use PQC algorithms. For instance, when a new block of transactions is proposed by a node, it is signed with that node's Dilithium key rather than an ECDSA key, eliminating vulnerability to future quantum attack on the ledger integrity. Client identity on ALE is likewise tied to their PQC public keys (often the same keys used for sensor data signing). By upgrading the blockchain's crypto primitives to PQC, we avoid making ALE the weak link – a scenario where an attacker with a quantum computer in a decade could fake ledger signatures or alter history.

The consensus algorithm of ALE is a variant of proof-of-authority or Byzantine fault tolerant (BFT) voting among trusted nodes, which is efficient and suitable for our scale (as opposed to energy-intensive proof-of-work). Each block includes a *timestamp, a list of data events, and a hash of the previous block*, forming the chain. Once written, a block (and thus the data records in it) cannot be altered without detection – any change would break the hash linkage and be rejected by the network. This tamper-resistance is at the heart of ASTERISK's assurance: historical data and decisions become *auditable records* that no malicious insider or external hacker can secretly edit.

4.2 Traceability from Data Ingest to Decision Outcome

TraceX [8], built on top of ALE, is engineered to trace the full lifecycle of information in a space system. Consider a scenario in a satellite operations center: multiple algorithms process tracking data after a breakup event to identify the new pieces of debris generated as shown in Fig. 5. With traceX, each step generates a ledger entry:

Step 1: Raw sensor observation: A sensor's signed observation (e.g., "Telescope A spotted Object X at time T in orbit Y") is submitted to the ledger as a transaction. The ledger entry contains the data, sensor ID, timestamp, and the sensor's signature (which is itself verifiable by anyone reading ALE).

Step 2: Fusion and analysis: Suppose an SSA (Space Situational Awareness) software module pulls several such observations to compute the validity of a new object being present. That module, upon producing its result writes a new ledger entry referencing the observation entries it used. TraceX links these through cryptographic pointers or metadata fields (like each entry can list the IDs of ledger records that informed it). This creates a *directed acyclic graph of data provenance* on top of the linear blockchain. In essence, one can traverse backwards ("why did we get this result?") and find all input data that led to it.

Step 3: Decision and action: Finally, a human operator or an automated system decides "Yes, task a new sensor to make further observations" and a command is sent to the possible sensors along the path. This decision event is also recorded: who/what decided, at what time, based on which recommendation. If the command is executed, the sensors might send back an acknowledgment, which again is logged.

If later analysis is needed (for example, if a collision with an tracked debris occurred or a false alarm was caught), investigators can replay the sequence via the ledger: check each algorithm's inputs and outputs, verify that each step was based on authentic data and correct processing. This addresses the "opaque AI/automation" issue where a complex system's output is hard to trust – traceX makes it transparent how a recommendation was formed.

This level of transparency also enables automated oversight and continuous improvement. For example, an **oversight AI** agent could monitor ALE in real time to enforce policies and flag anomalies (as outlined in Sec. 1), while mission controllers or auditors can perform random spot-checks of decision chains for performance verification. In essence, traceX operationalizes the oversight, explainability, and continuous verification principles discussed earlier, turning them into practical capabilities for space operations. Figure 5 depicts an example traceability chain

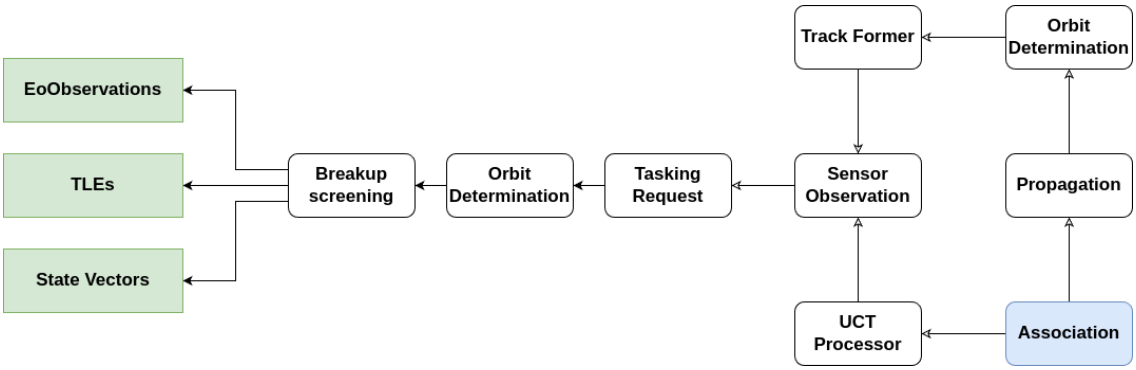

Fig. 5. Traceability: Data and decisions from algorithm interactions are stored on the ledger to provide data provenance. Arrows point to data consumed at each steps, the traceability of data and decision.

from traceX. The left side shows sensor inputs feeding into a chain of processing algorithms, leading to course of action (CoA) decision. The ledger captures two key facets for traceability: state data at each process (instantaneous inputs/outputs) and causal linkages between processes (which output led into which subsequent process). Together, these allow reconstruction of the entire decision flow. This approach is heavily informed by NIST NCCoE's Supply Chain Traceability meta-framework (NIST IR 8536) [18], which Space Protocol has adapted from manufacturing to digital processes. In manufacturing, traceability tracks how each part and subassembly contributes to a final product; here, each data element and algorithmic decision is a "part" that contributes to the final operational decision. By implementing this framework on a secure ledger, we not only log the data but ensure its integrity and availability for future audits. The reader is referred to [8] for a detailed dive into traceX.

4.3 Risk Quantification and Decision Confidence

Beyond traceability, traceX serves as a *risk quantification engine*. Because it aggregates data from multiple sources (Sec. 4.4), it can compute metrics about the reliability and risk of decisions:

Traceability-based confidence score: As each decision entry links to underlying data, traceX can automatically evaluate how well-supported a decision is. For example, a decision to classify an object as hostile could be accompanied by a score indicating it was based on 5 independent sensors and vetted by 3 different algorithms (high confidence), versus one based on a single source (low confidence). These scores help mission controllers and autonomous systems weigh whether to act or seek more data.

Sensor and algorithm performance metrics: Over time, the ledger accrues a historical record of sensor contributions and algorithm outcomes. Using this, traceX can derive statistics such as sensor accuracy rates (how often a sensor's observation was corroborated by others), false alarm rates, latency, etc. Similarly, algorithm performance (how often recommendations were right or led to good outcomes) can be quantified. This feeds into a *reputation system* for network participants – a sensor that consistently provides unique but correct observations builds a high reputation, whereas one that often outputs data no one else can confirm would be assigned a lower trust weighting. TraceX can factor these into real-time consensus (down-weighting data from unreliable sources).

Liability and risk assessment: In safety-critical scenarios, the ledger's data can support quantitative risk assessment. For instance, if a satellite collision warning is issued, ASTERISK, using applications that consume data from ALE, could provide an estimated probability of collision along with confidence bounds, and a "risk index" if no maneuver is taken. If an operator chooses to ignore a warning, that decision (and its rationale) is also logged, which is crucial for accountability and post-event analysis. In Sec. 6 we discuss how such quantification enables insurance mechanisms; basically, by having concrete data and probabilities on record, insurers can underwrite policies (e.g., insuring a satellite against collision if ASTERISK shows due diligence was followed and the residual risk is within certain limits).

From a **quality and safety perspective**, this risk quantification ensures that decisions are made with a clear understanding of uncertainty. It prevents over-confidence in possibly flawed data and encourages obtaining multiple independent sources. Moreover, if something does go wrong, the recorded data allow investigators to determine whether it was simply bad luck (random failure) or a preventable lapse in following the data. This distinction is key to continuous improvement.

4.4 Tamper-Resistant Cross-Verification and Consensus

As mentioned, ASTERISK inherently supports *cross-verification*: multiple parties (observers, validators) see the data and must agree on records. In practice, when a sensor submits a new observation transaction, a set of validator nodes (could be major network hubs or independent third parties) run checks – e.g., does this data fit within physics (no impossible or wildly outlying values), is the signature valid, does it conflict with known truths (like two satellites claiming the same identifier), etc. Only after passing validation is the transaction finalized on the ledger. This distributed validation process itself is a guard against insider threats: a rogue participant cannot unilaterally introduce false data because others will reject it.

Moreover, traceX provides real-time alerts for anomalies. If conflicting data enters (say two sensors give drastically different positions for what should be the same object), the consensus algorithm can mark that in the ledger and notify operators that something is inconsistent. This might indicate a compromised sensor or an ongoing spoofing attack, prompting investigation. By catching such conflicts early and publicly (within the consortium of participants), traceX turns what could be a quiet data corruption into a visible event that stakeholders can respond to.

In summary, traceX is the backbone of data-to-decision assurance: it guarantees the integrity of information, provides transparency of process, and quantifies trust at each step. These qualities directly support the theme of quality and safety – a system with traceX can answer the tough questions ("How do we know the decision was right? Can we trust this data? Who is responsible for this outcome?") with evidence and clarity, which is invaluable for mission assurance, regulatory compliance, and learning from incidents.

5. Cross-Domain Use Case: Space Assets Supporting Air Defense via decentralized sensing

Decentralization plays a pivotal role in security and safety. By distributing sensing across many small nodes, each node can be physically smaller and lower-power, reducing its electromagnetic/visual signature and avoiding *single points of failure* inherent to large, centralized sensors that are easier to detect and target from air or space. One example of such a decentralized, multi-modal, rapidly deployable approach is *Crimson Shock* under development at the SDA TAP Lab (Colorado Springs): a portable kit with multiple sensing modalities that can be fielded at pace and networked into a larger observability fabric. Commercial partners (e.g., ground RF/EO providers) can contribute compact units to this mesh, extending coverage with minimal footprint.

To illustrate the mission and business relevance of our quantum-resilient architecture, we present a cross-domain use case where sensors built for space domain awareness (SDA) are repurposed to detect and track **air-domain** objects such as high-altitude targets. Space-derived observations are fused with cryptographic identity streams (e.g., drone Remote ID) to create a trusted operational picture.

Scenario Background Nation X operates ground-based EO telescopes and passive RF sensors originally deployed for on-orbit monitoring. Facing increased incursions by uncooperative drones and near-space balloons, the operations center retasks some space-facing telescopes toward the atmosphere during predefined windows and retunes passive RF arrays to capture line-of-sight control/telemetry signals. In parallel, authorities mandate Remote ID; cooperative drones broadcast signed identifiers, while adversaries may omit or spoof them. Integrating cryptographically verifiable Remote ID checks with space-derived cues turns SDA assets into a practical air-defense sensor layer.

Architecture in Action An EO telescope (T1) observes a high-altitude object above a critical site; a nearby passive RF node (R1) concurrently detects a 2.4 GHz control-like signal. Both submit *PQC-signed* observations to ALE.

Ledger-linked reports are correlated; a ground Remote ID receiver logs an ID string but the signature is missing/invalid. The fusion node produces a composite track (multi-modal detection, no valid identity) and records the reasoning on chain.

A PQC-secured alert is sent to the battle management system with provenance pointers. Because the assessment is traceable on-ledger, commanders can see the data-to-decision pedigree.

An interceptor is launched; additional RF bearings localize a suspected ground controller. The full incident—from first cue to outcome—is *immutably logged* for accountability and after-action review. The scenario is depicted in Fig. 6.

This scenario shows how *multi-modal decentralization* plus *cryptographic trust* extend air surveillance and harden it against deception. The use of PQC identities means the absence of a valid Remote ID is immediately apparent and auditable; conversely, a valid signed ID could inform a lower-risk response (e.g., monitor rather than neutralize) unless behavior escalates.

6. Marketplace and Business Model for Secure SDA Data

For adoption, enhancing space safety and cybersecurity must align with economic opportunity. Our architecture underpins a **marketplace-driven ecosystem** where a PQC-enabled, permissioned ledger records provenance

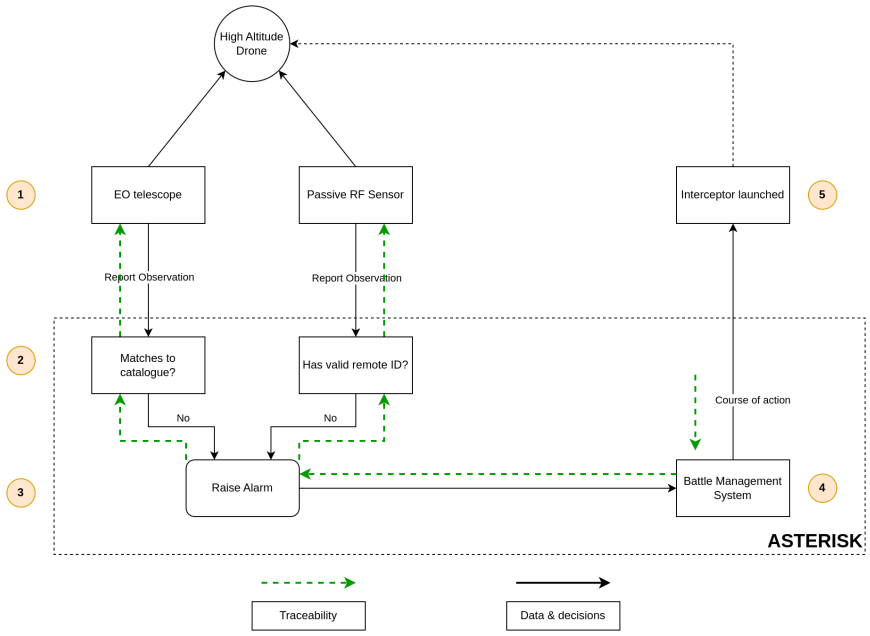

Fig. 6. Cross-Domain application of the ASTERISK framework. In a air defense scenario, detections are stored on ledger, information validated and an alert is raised, leading to a course of action and target interception.

and settlement. The model incentivizes high-quality contributions, introduces **spot liability insurance** to guarantee data quality, and enables cross-border exchange under compliance—creating a durable feedback loop that lifts the quality and safety of space operations.

6.1 SDA Data Marketplace and Incentive Mechanisms

The ASTERISK architecture underpins a blockchain-based marketplace where space surveillance data and analytics can be bought, sold, or shared under transparent and verifiable rules. Within this ecosystem, application providers deliver value-added services such as conjunction risk analyses or traceability reports, while data providers—ranging from satellite operators and ground EO/RF stations to networks of small sensors—contribute observations tied to PQC-anchored identities and reputation scores. Data consumers, including operators, insurers, regulators, defense agencies, and researchers, access this information through tiered permissions, ensuring both security and accountability. Validators and computation providers play a critical role by performing consensus and providing compute and storage, and are rewarded for their contributions. Transactions can be settled in fiat or tokenized payments, with dynamic pricing reflecting the rarity, confidence, and cross-validation level of the data. By lowering participation barriers and rewarding high-quality, verified contributions, this marketplace broadens observational coverage, aligns incentives across stakeholders, and ensures that space safety is reinforced through both technical assurance and economic motivation.

6.2 Spot Liability Insurance and Data Assurance

We integrate transaction-level liability cover so that critical decisions can reliably depend on purchased data. In this model, providers may attach warranted telemetry—such as missed-collision liability up to a defined value—backed either by an insurer or a staked bond. Policies, triggers, and payouts are enforced through insurance smart contracts, with objective evidence drawn directly from the ledger to determine fault. Premiums are underwritten against the historical reliability of data contributors, rewarding high-performing providers with lower costs while penalizing poor performance. This creates natural market differentiation: insured, cross-verified products command a premium price but are more appropriate for safety-critical operations, where quality and accountability are paramount. Flexible coverage options, ranging from per-transaction micro-policies to time-bound windows and capped subscriptions, further adapt the model to varied mission contexts. By embedding insurance at the transaction level, this approach builds trust and assigns accountability up front, replacing slow, adversarial post-event disputes with deterministic and transparent resolution.

6.3 Cross-Border Data Sharing and Export Considerations

Cross-border data sharing in space operations must balance openness with compliance, ensuring that sensitive information is accessible only to authorized partners. In the ASTERISK framework, providers can tag their data with explicit access constraints—such as ITAR or EAR restrictions—so that encrypted datasets are only decrypted for approved parties. Allied consortia can operate within a shared, permissioned ASTERISK instance, forming a secure "data club" that allows vetted partners to collaborate while preserving sovereignty. Nations can maintain sovereign data pools, keeping local copies while selectively sharing subsets under enforceable on-ledger policy controls. To guarantee compliance, immutable audit trails capture every access and export, deterring unauthorized resale and providing verifiable proof of proper handling. This model enables monetization of advanced tracking capabilities while allowing nations and organizations with capability gaps to purchase trusted data instead of replicating full infrastructure stacks, thereby enhancing safety and interoperability without compromising security or sovereignty.

6.4 Economic Viability and Sustainability

The economic model underpinning ASTERISK ensures that enhanced space safety and cybersecurity can be sustained over the long term. Revenues generated through the marketplace provide direct funding for sensor upkeep and upgrades, while integrated insurance mechanisms redistribute losses when rare failures occur, preserving confidence in the system. Dynamic pricing signals highlight where investment is most urgently needed—such as debris removal or new sensor deployments—ensuring that resources flow to areas with the highest safety return. By linking premiums to behavior and reliability, the model incentivizes adherence to best practices and international standards, effectively turning compliance into a competitive advantage. In this way, the ecosystem evolves beyond treating security as a cost center: it becomes a *growth engine* where economic incentives, risk sharing, and technical accountability reinforce one another to foster a cooperative and resilient space-safety economy.

7. Discussion: Sovereign Capability, Standards Alignment, and Policy

We situate the architecture in sovereign and standards contexts and outline an export-compliant path to adoption. This approach creates the opportunity to move beyond initial traceability in shared data to validation, enabling confidence in data from commercial providers and external governments. This supports the sovereign roles of authorization and continuing supervision of space activities by allowing us to move beyond a framework that requires either intrinsic trust in the capability of a third party or access to proprietary process to validate those capabilities. Instead, competence can be demonstrated through comparison to other sources, establishing confidence in data quality based on outputs. In addition, providers can use confidence scores to identify opportunities to improve their data quality. As the commercial space situational awareness industry continues to mature, licensing regimes may become necessary, requiring mechanisms that can validate performance in the context of a regulatory standard.

7.1 Strengthening Sovereign National Space Capability

ASTERISK strengthens sovereign space capability by giving nations an independent, verifiable picture of their space environment. By embedding NIST-standard post-quantum cryptography and optional quantum key distribution into TT&C and data links, the architecture ensures communications remain secure against current cyber threats and future quantum adversaries. Indigenous sensors, anchored in ALE, allow nations to build trusted space domain awareness without over-reliance on external sources. Provenance and risk metrics logged in real time enable decision autonomy, ensuring governments can act quickly and independently during conjunctions or airspace incidents. Moreover, participation in the ASTERISK-enabled marketplace provides economic leverage: nations can monetize their SDA contributions and shape global norms for warranted, liability-backed data, enhancing both sovereignty and influence.

7.2 Alignment with Cybersecurity Mandates and Safety Standards

ASTERISK is designed to align seamlessly with existing cybersecurity and safety standards. Integration with CCSDS protocols ensures that PQC upgrades fit directly into standard telemetry and telecommand stacks, reinforcing long-standing reliability goals. Compliance with NIST and CISA mandates is achieved through authenticated links, encryption at rest and in transit, and supply-chain traceability. International standards bodies such as ITU and ETSI are also supported, enabling interoperability with global telecom networks and quantum communication initiatives like EuroQCI. On the defense side, ASTERISK's ledgered provenance, quantum-safe cryptography, and auditable trails satisfy U.S. Department of Defense supply-chain requirements under CMMC. From a quality perspective, its traceability mechanisms align with ISO 27000 and AS9100 practices, ensuring that data assurance and safety standards are met across civilian and military contexts.

7.3 Export-Compliant Dual-Use Implementation

The architecture is intentionally modular and standards-based to ease licensing and export compliance. By relying on openly published PQC algorithms and ITU/ETSI QKD frameworks rather than proprietary cryptography, ASTERISK avoids regulatory friction and supports transparent evaluation. Permissioned partitions within the ledger allow fine-grained enforcement of access policies, ensuring that restricted data remains within authorized domains while enabling selective sharing with allies. Shared, permissioned "data clubs" can be established among coalitions, providing trusted collaboration with common audit capabilities while respecting national sovereignty. This model not only reduces misuse risk but also provides a clear path for policy advocacy: authenticated, encrypted data sharing demonstrably improves collective safety while minimizing opportunities for exploitation.

7.4 National and International Interest

ASTERISK addresses pressing national and international priorities, from space traffic management to critical infrastructure protection. By providing transparent maneuver and observation logs, along with liability-insured data products, the system incentivizes responsible behavior and discourages negligence in orbit. Its zero-trust design, cryptographic provenance, and continuous audit trails also align with evolving mandates for protecting critical infrastructure against cyber threats. For international stakeholders, ASTERISK offers a framework for building coalitions that balance national sovereignty with global safety, encouraging cooperative norms in space governance while allowing each partner to retain independent control of their assets.

7.5 Adoption path

Practical adoption of ASTERISK has begun with tightly controlled pilot deployments, where strict access policies and limited participation reduce institutional risk while demonstrating operational value. Early benefits—including stronger security, fewer operational incidents, and clearer accountability—build trust among stakeholders and create momentum for broader deployment. As regulatory frameworks increasingly emphasize quantum safety, provenance, and auditability, the policy landscape will further encourage adoption. ASTERISK positions itself as a forward-looking, dual-use architecture that treats safety and security as shared global goods, while still respecting the sovereignty and strategic interests of each participating nation.

8. Space Protocol's offering and the road ahead

ASTERISK is offered as an open-standards, modular software stack: cryptography (NIST PQC + QKD), identities, a PQC-secured ledger, and a smart-contract marketplace. Core cryptographic primitives are openly licensed, while ledger participation follows a permissioned consortium license, allowing governments, operators, and insurers to join without losing sovereignty. TraceX is offered as an open source framework within ASTERISK.

The ASTERISK platform serves a diverse set of stakeholders across the space and air domain awareness ecosystem. This include satellite operators and fleet managers (SpaceX, OneWeb, Amazon Kuiper), who require *trusted conjunction prediction*, *calibration*, and *maneuver compliance* data to safeguard assets worth hundreds of millions. *Government agencies and defense organizations* – both civil regulators and military users – benefit from ASTERISK's dual-use capabilities for space traffic management, debris tracking, and detecting uncooperative aerial or orbital threats. **Insurance providers and underwriters** (Lloyd's of London, AXA XL, Munich Re) represent another core segment, leveraging ASTERISK's liability-backed data and smart-contract insurance mechanisms to offer new, verifiable risk products. Beyond these, launch providers and satellite manufacturers (Mitsubishi Heavy Industries, Interstellar Technologies, Gilmour Space) seek reliable telemetry and insertion assurance to reduce liability during critical mission phases, while commercial SDA/STM firms use the ASTERISK marketplace to resell services with embedded trust, lowering costs and expanding reach. Finally, research institutions and standards bodies gain access to tamper-proof, cross-validated data streams to inform science, policy, and emerging global norms. Together, these segments underscore a broad commercial and strategic demand for decentralized, cryptographically assured sensing and decision infrastructure.

As we look ahead, several avenues for future work and cooperation emerge. On the technical front, ongoing refinement of PQC algorithms will need to be closely tracked. For example, NIST may announce additional standards or find vulnerabilities in current ones (as happened with an early candidate), and our system's modular design is poised to incorporate such updates seamlessly. Similarly, continued developments in QKD—such as higher-rate quantum satellites or quantum repeaters—could expand the reach of our QKD layer, and we plan to integrate those advancements when feasible. On the ASTERISK side, scaling the ledger to handle the full global volume of SDA data and optimizing consensus delays are challenges we intend to tackle. We anticipate collaborating with organizations like MITRE and international SSA data centers to pilot this architecture at scale.

On the policy side, engagement with bodies like the **IAF**, **CCSDS**, **UNOOSA**, and national regulators will be crucial to refine the governance aspects of a shared ledger and marketplace. Open questions remain regarding data ownership, privacy (especially for commercial satellite operators contributing data), and antitrust concerns if such marketplaces become widespread. We propose a multi-stakeholder approach to address these issues: bringing together governments, space companies, insurers, and standards organizations to establish **ethical and legal frameworks** for space data sharing. The safety benefits are clear, but consensus on the "rules of the road" will ensure the solution's longevity.

In conclusion, **Space Protocol's approach aligns with the core theme of "Quality and Safety for a successful space program"** by recognizing that in the high stakes of space operations, *security is safety*. Quality data leads to quality decisions; secure systems prevent disasters. Our quantum-resilient architecture marries cutting-edge cryptographic protection with innovative data governance to create a forward-looking, mission-centric solution. It is an architecture designed not just to survive the next disruption—be it a cyber attack or the quantum leap—but to thrive through it, turning potential perils into opportunities for growth and cooperation. As the space community stands at the cusp of a new era of exploration and commercialization, we offer this blueprint as a timely contribution to ensure that era is underpinned by trust, resilience, and shared responsibility.

[1] B. Cyr, Y. Long, T. Sugawara, and K. Fu, "Position paper: Space system threat models must account for satellite sensor spoofing.," in *SpaceSec*, 2023.

[2] E. Ear, J. L. Remy, A. Feffer, and S. Xu, "Characterizing cyber attacks against space systems with missing data: Framework and case study," in *2023 IEEE Conference on Communications and Network Security (CNS)*, IEEE, 2023, pp. 1–9.

[3] A. T. Olutimehin, S. Joseph, A. J. Ajayi, O. C. Metibemu, A. Y. Balogun, and O. O. Olaniyi, "Future-proofing data: Assessing the feasibility of post-quantum cryptographic algorithms to mitigate 'harvest now, decrypt later'attacks," *Decrypt Later'Attacks (February 17, 2025)*, 2025.

[4] S.-K. Liao *et al.*, "Satellite-to-ground quantum key distribution," *Nature*, vol. 549, no. 7670, pp. 43–47, 2017.

[5] J. Hickman, "Research viewpoint: International relations and the second space race between the united states and china," *Astropolitics*, vol. 17, no. 3, pp. 178–190, 2019.

[6] D. A. Cooper and Q. Dang, "NIST Special Publication 800-208: Recommendation for Stateful Hash-Based Signature Schemes," National Institute of Standards and Technology, Tech. Rep. SP 800-208, Oct. 2020. DOI: 10.6028/NIST.SP.800-208. [Online]. Available: https://doi.org/10.6028/NIST.SP.800-208.

[7] J. Pearl, "Causal inference in statistics: An overview," 2009.

[8] S. Bagchi, L. Pratti, H. Reed, and Y. Latif, "Towards an AI and blockchain enabled Space Battle Management System," AMOS, Space Protocol, 2025.

[9] A. I. Nurhadi and N. R. Syambas, "Quantum key distribution (qkd) protocols: A survey," in *2018 4th International Conference on Wireless and Telematics (ICWT)*, IEEE, 2018, pp. 1–5.

[10] National Institute of Standards and Technology, "FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM)," U.S. Department of Commerce, Tech. Rep. FIPS 203, Jun. 2024. [Online]. Available: https://csrc.nist.gov/pubs/fips/203/final.

[11] National Institute of Standards and Technology, "FIPS 204: Module-Lattice-Based Digital Signature Algorithm (ML-DSA)," U.S. Department of Commerce, Tech. Rep. FIPS 204, Jun. 2024. [Online]. Available: https://csrc.nist.gov/pubs/fips/204/final.

[12] National Institute of Standards and Technology, "FIPS 205: Stateless Hash-Based Digital Signature Algorithm (SLH-DSA)," U.S. Department of Commerce, Tech. Rep. FIPS 205, Jun. 2024. [Online]. Available: https://csrc.nist.gov/pubs/fips/205/final.

[13] E. Salkield, S. Köhler, S. Birnbach, R. Baker, M. Strohmeier, and I. Martinovic, "Firefly: Spoofing earth observation satellite data through radio overshadowing," 2023.

[14] M. J. Dworkin *et al.*, "Advanced encryption standard (aes)," 2001.

[15] S. Rao, D. Mahto, D. K. Yadav, and D. A. Khan, "The aes-256 cryptosystem resists quantum attacks," *Int. J. Adv. Res. Comput. Sci*, vol. 8, no. 3, pp. 404–408, 2017.

[16] Y. Latif, A. Chowdhury, and S. Bagchi, "On-chain validation of tracking data messages (tdm) using distributed deep learning on a proof of stake (pos) blockchain," in *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, 2024, p. 106.

[17] R. Carden, D. Burchett, and H. Reed, "Snare (sensor network autonomous resilient extensible): Decentralized sensor tasking improves sda tactical relevance," in *Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, 2021.

[18] M. Pease, F. Wallace, H. Reed, V. Martin, and S. Granata, "Supply chain traceability: Manufacturing meta-framework," in *NIST Internal Report*, Initial Public Draft, National Institute of Standards and Technology (NIST), Nov. 2024. [Online]. Available: `https://csrc.nist.gov/pubs/ir/8536/ipd`.