Computers & Security 164 (2026) 104848

Contents lists available at ScienceDirect

Computers & Security

journal homepage: www.elsevier.com/locate/cose

Full Length Article

Assessing the attack surface of space organizations: A data-driven analysis

Francesco Casaril *, Letterio Galletta

IMT School for Advanced Studies Lucca, Lucca, Italy

ARTICLE INFO

Keywords:
Space cybersecurity
Space infrastructure
Risk assessment
Space policy

ABSTRACT

The increasing digitalization of the space industry and the rapid expansion of commercial space activities have increased the sector's exposure to cyber threats. As satellite operators and aerospace entities rely on Internet-connected devices (ICDs) for control, communication, and ground-based operations, their attack surface expands accordingly. Despite this growing risk, there remains a lack of standardized methodologies tailored to measuring real-world cybersecurity exposure of ICDs in the space sector. Existing frameworks often overlook the unique characteristics of space infrastructure, including persistent connectivity, long system lifespans, and limited patching opportunities. To address this gap, we propose the Risk Exposure Framework (REF), a methodology to quantify cybersecurity exposure using Internet-facing asset data. REF integrates elements from well-established risk assessment models with targeted analysis of exposed services, known vulnerabilities, and exploit availability. The framework calculates risk through a structured approach that combines Exposure and Likelihood scores based on observable attack surface metrics. Our methodology allows one to compare exposure levels across organizations and supports alignment with sector-specific cybersecurity requirements, and it is adaptable to other critical infrastructure environments where external exposure plays a central role in cyber risk. Unlike general-purpose frameworks, REF directly captures space-specific traits by relying on observable network exposure indicators and by aligning with the principles of attack surface measurement in space environments. REF quantifies the externally observable posture of space organisations, primarily ground-segment and enterprise networks, based on Internet-facing exposure and exploitability. The framework does not model spacecraft constraints, but it can reflect their downstream effects when those constraints manifest at network boundaries. This paper also examines how the REF methodology can support existing cybersecurity policy frameworks and risk assessment strategies in both Europe and the United States.

1. Introduction

Space technologies are vital to modern societies, from navigation and communication to weather forecasting and defense operations, many of the services we rely on depend on satellites and space infrastructure. Recently, the European Union has started treating space not just as a technical field, but as a core part of its security and defense strategy. This approach is expressed in several strategic documents. The "Strategic Compass for Security and Defence" (European External Action Service, 2022), adopted in 2022, was one of the first EU documents to frame space as a contested domain, and calls for stronger tools to detect and respond to threats targeting space infrastructures. The "EU Space Strategy for Security and Defence" (European Commission, 2023) builds upon the previous document by stressing the need to protect space assets and systems from sabotage, espionage, and cyberattacks. More recently, the "White Paper for European Defence - Readiness 2030" (European Commission, 2023) has clearly stated the centrality of space to European security, highlighting again the need for better protection of satellites and ground stations, especially from cyber threats.

These strategic moves come at a time when the threat landscape for space is becoming more serious. The 2022 cyberattack on Viasat during the early days of the war in Ukraine showed how space systems can be disrupted with major consequences (Boschetti et al., 2022; Casaril and Galletta, 2024). Attacks on undersea Internet cables, like those in the Baltic and Mediterranean, have added concerns about Europe's digital backbone (Schaller, 2025). In this context, secure space infrastructure becomes even more important as a backup and as a way to stay connected when other systems fail.

Despite the growing attention to cybersecurity in the space sector, there is still a significant gap in the ability to empirically measure the cyber exposure of space organizations. Current frameworks tend to be qualitative, relying on expert judgment, or focused on threat modeling without incorporating real-world data from Internet-facing systems. At present, there is no standardized methodology for assessing and

* Corresponding author.
E-mail addresses: francesco.casaril@imtlucca.it (F. Casaril), letterio.galletta@imtlucca.it (L. Galletta).

https://doi.org/10.1016/j.cose.2026.104848
Received 15 July 2025; Received in revised form 16 December 2025; Accepted 1 February 2026
Available online 3 February 2026

quantifying this exposure across the space industry. Consequently, policymakers and operators lack actionable metrics to evaluate risk levels, prioritize mitigation efforts, and align with the evolving cybersecurity requirements within the European context.

Given the strategic importance of space technologies in our society, this paper aims to address the above gaps. Specifically, we aim to answer the following research questions, each of which is motivated by a specific driver:

RQ1 How can we effectively measure cyber risk in space-sector organizations in a way that supports both technical and policy decision-making?

The need to address this research question originates from the existing gap between high-level cybersecurity obligations, such as those imposed by EU policies, and specific, automatable technical measures that can be directly verified for effectiveness.

RQ2 Is it possible to measure cyber risk exposure in a clear and repeatable way?

This question stems from the need for a clear and repeatable method to quantify cyber-risk exposure, which supports cross-organizational comparison as internet-facing assets continue to grow and oversight becomes mandatory.

RQ3 How much exposed are space-sector organizations to cyber risks, especially those that come from Internet-connected systems?

This question comes from the evidence indicating that successful intrusions in the space domain often originate from Internet-connected systems within ground and user segments, highlighting the need for an assessment of the resulting exposure at a sector-wide level.

RQ4 Can we connect this kind of measurement to existing European cybersecurity policies and frameworks?

This question is driven by the development of European space cybersecurity policies, particularly the NIS-2 directive and the proposed EU Space Act, which increasingly require continuous risk assessment at both the organizational and sector levels.

To answer these research questions, we introduce the *Risk Exposure Framework (REF)*, a risk assessment methodology that aims to measure two aspects: (*i*) how much is an organization visible to potential attackers (*exposure*), and (*ii*) how likely it is that known vulnerabilities in the organization's attack surface might be exploited (*likelihood*). REF is designed to be flexible and can be applied in different contexts, but in this paper, we focus on space organizations. More precisely, this paper provides the following contributions. To reply to **RQ1** and evaluate the level of exposure of space-sector organizations, we collect data on the Internet-exposed assets of both public and private space entities, identifying connected devices and known vulnerabilities. To determine how to measure cyber risk effectively, we utilize REF to generate a risk score based on observable data. Then, to answer **RQ2** and to test whether our approach is clear and repeatable, we apply REF to a representative group of space organizations and evaluate its consistency across different entities. Finally, to reply to **RQ3** and **RQ4**, we examine whether the framework can connect to existing European policies: we interpret our findings in light of strategic documents such as the EU Space Strategy for Security and Defence and current NIS2 obligations, showing how REF can complement policy-driven approaches to cyber risk governance. Our goal is to offer a simple but useful way to think about cyber risk in space, using a method that is practical, repeatable, and aligned with where European policy is heading. To support reproducibility and further research, a Python implementation of the REF methodology is available in a public GitHub repository (Casaril and Galletta, 2025a).

Note that REF specifically focuses on the ground segment and on supporting digital infrastructure and other Internet-facing assets discoverable through network reconnaissance tools. This ground-segment focus is particularly relevant since several recent cyberattacks against space systems originate through these externally accessible entry points (Viasat, 2022). REF does not directly assess the space segment itself (onboard spacecraft systems, satellite subsystems, or orbital infrastructure), which typically operate in more isolated, proprietary, or air-gapped environments that are not visible through Internet scanning methodologies.

The rest of the paper is organized as follows: Section 2 presents some background information about Internet Connected Device Analysis, space-enabled critical infrastructures, and the relevant literature for our work. Section 3 explains how REF works and the metrics it uses. In Section 5, we apply the REF framework to a group of space-related organizations and illustrate the results of our analysis. Section 6 discusses how REF can be used to fulfill some requirements prescribed by current policies. Finally, Section 7 wraps up the findings and outlines possible future research lines.

2. Background

2.1. Internet-connected device analysis

Internet-connected devices (ICDs) include a great category of devices such as industrial control systems, satellite modems, or energy management systems (Lu and Da Xu, 2018) that are exposed to the public Internet, either by design or misconfiguration. The analysis of these devices is a key activity to understand the evolving cybersecurity posture of modern digital infrastructures.

ICD exposure creates an expanded attack surface that can be monitored and evaluated using a specialized class of search engines. Among the most widely used engines, there are *Shodan* (Matherly, 2015), *Censys* (Durumeric et al., 2015), *ZoomEye* (Li et al., 2020), *Fofa* (Liu et al., 2025), *BinaryEdge* (Daskevics and Nikiforova, 2021), and *Netlas* (Netlas, 2025). Each engine has its own scanning strategies, indexing models, and coverage scope, which affect the consistency and granularity of its results. Below, we briefly review the main features of these platforms to clarify the motivation for our choice:

- Shodan is one of the most used IoT search engines. It systematically scans IP address space and indexes data based on service banners, open ports, software versions, geographic information, and organizational ownership. It integrates vulnerability databases like the National Vulnerability Database (NVD) (Booth et al., 2013), and associates identified services with Common Vulnerabilities and Exposures. It also supports advanced filters (e.g., organization, open ports, and presence of vulnerabilities), allowing advanced targeting.
- Censys focuses on the protocol level and provides highly expressive query capabilities. It collects detailed X.509 certificates, TLS configurations, and other protocol metadata. It offers less direct vulnerability tagging compared to Shodan, but it can fingerprint cryptographic settings and web infrastructure.
- ZoomEye and Fofa are widely used in Asia (Li et al., 2020) and provide extensive metadata tagging and visualization features. These engines often surface device types and service families using machine learning classifiers.
- BinaryEdge offers detailed metadata exports, including packet capture-level summaries, and often identifies honeypots, behavioral patterns, and unusual ports. It can be very useful for correlating IoT device behavior with cloud infrastructure exposure.
- Netlas is a more recent platform that combines banner search, certificate transparency logs, and passive DNS information. The vulnerability integration and query complexity of the tool are still evolving.

Each of these engines contributes to the global reconnaissance landscape, but they are rarely interchangeable. Their coverage is often non-overlapping, meaning that a device visible in Shodan may be completely invisible to Censys or BinaryEdge. For the structured cybersecurity assessment of REF, which involves quantitative risk scoring, we consider

Shodan as the most operationally relevant platform: its direct CVE tagging, historical scan database, and support for Boolean logic queries allow us to track the evolution of exposure across organizations, sectors, or geographic regions. In particular, it can associate each retrieved IP with known vulnerabilities and CVEs, a feature that is fundamental to our REF framework. Furthermore, Shodan's API is reliable, well-documented, and supports advanced filters that help target devices with specific characteristics during the risk quantification at the organizational level. Other platforms, like Censys or BinaryEdge, offer complementary perspectives, but their output is less consistent in tagging vulnerabilities or associating results with organizational identifiers.

To complement Shodan's exposure data with exploit intelligence, we integrate REF with the Vulners API (Tundis et al., 2018). Vulners is a vulnerability intelligence platform that aggregates data from multiple sources, including the National Vulnerability Database (Booth et al., 2013), Exploit-DB (Valea and Oprişa, 2020), Metasploit, and vendor advisories, checking for CVEs against public exploit sources. It also parses information from vendor advisories and security bulletins, including references to known exploitability. Vulners provides access to structured information about vulnerabilities, including their CVSS scores, CWE classifications, publication dates, and, most importantly for our research, exploit availability. Within REF, Vulners is used to query each CVE identified by Shodan to determine whether known exploits exist in the wild. Automating this enrichment process, REF can scale across hundreds of devices and IP addresses, mapping the real-world exploitability. The result is a more actionable measure of risk, grounded in adversarial capability and opportunity.

The ultimate goal of ICD analysis in our research is not merely to catalog devices but to assess their security posture. This includes evaluating the likelihood of exploitation based on exposed services, the impact of potential breaches based on system criticality, and the presence of known software vulnerabilities. As such, ICD analysis is used to perform cyber risk assessment workflows, particularly in sectors like space infrastructure.

This paper focuses on ICDs for methodological reasons, as they constitute one of the only portions of the space ecosystem whose exposure can be empirically measured at scale through externally observable data. In contrast, other risks, such as those that affect the supply chain (Geetha et al., 2024) and vulnerabilities in COTS (Sharmin et al., 2025) or legacy subsystems, typically require privileged access, internal documentation, or forensic analysis that cannot be obtained through open-source intelligence. Second, Internet-exposed services represent the point of lowest-cost adversarial access, and historically constitute the initial foothold for intrusion campaigns across many critical infrastructure sectors (Poirier, 2023). Prior work on attacks against satellite operators and ground stations shows that compromise of externally reachable systems frequently precedes deeper penetration into mission networks (Poirier, 2023). Third, restricting the scope to ICDs ensures reproducibility, as every organization can be evaluated using the same data sources, time windows, and measurement pipeline, allowing REF to provide a consistent quantitative comparison across entities. For these reasons, the analysis in this paper does not attempt to model supply chain dependencies or architectural weaknesses, but it isolates the part of the attack surface that is externally measurable and attacker-observable, which is necessary for constructing a standardized, data-driven exposure metric tailored for the space sector.

While the detailed methodology behind REF and our framework is described in Section 3, the next section describes instead some of the research that has already been carried out in this domain.

2.2. Space-enabled services as interdependent critical functions

Space infrastructure is embedded within critical sectors as a set of enabling functions whose failure degrades terrestrial systems in several ways. Satellite communications provide wide connectivity for operational technology and enterprise backhaul (teleports, gateways, cloud-hosted ground functions), as well as broadcast and broadband services to public services, and mission Telemetry, Tracking, and Command (TT&C). As operators externalise ground functions (commercial teleports and cloud-based ground segment as a service), misconfigurations or common vendor vulnerabilities become increasingly present as points of weakness across otherwise independent organisations. Positioning, Navigation and Timing (PNT/GNSS) provides absolute time and position to power grids, telecommunications, and finance, and anchors aviation and maritime navigation; jamming and spoofing thus propagate beyond navigation to timing-dependent processes when holdover and multi-source strategies are weak (Psiaki and Humphreys, 2016; EASA, 2024; Administration, 2017). Earth Observation (EO) and traffic services (sat-AIS/ADS-B) feed situational awareness and logistics, moving through partner APIs and commercial clouds into emergency response, defence, and supply chains; their integrity and availability constraints translate directly into the quality of downstream public services. These structural ties explain why cyber incidents in the space domain typically materialise first at organisational boundaries (reachable management portals, VPNs, provider APIs) and then radiate into energy, transport, telecoms, finance, and public safety rather than remaining confined to the space sector (European Union Agency for Cybersecurity, 2024a).

2.3. Risk propagation across organisations and sectors: Evidence and how to address it

Cyber incidents in the space domain demonstrate that attacks often affect multiple organizations. The interconnected nature of space requires understanding and measuring each organization's exposure to cyber threats. Quantifying external exposure, specifically the number of systems that are reachable and vulnerable, can help anticipate and mitigate these propagation effects. The following examples illustrate how failures in the space layer can lead to systemic risk and emphasize the importance of measuring exposure to prevent such incidents.

The KA-SAT incident of 24 February 2022 (Casaril and Galletta, 2024) demonstrated how a compromise of the management plane enabled the deployment of a wiper against satellite modems, rendering terminals inoperable. The attack cascaded from the satellite operator to telecommunications providers and ultimately to the energy sector, where Enercon temporarily lost remote monitoring and control of approximately 5800 wind turbines relying on SATCOM backhaul for operations (SentinelLabs, 2022; Reuters, 2022; Viasat, 2022; Institute, 2022). The propagation path in this case ran from a vendor's ground management environment to customer equipment, and onward into geographically dispersed industrial assets.

In 2025, a similar attack targeted the maritime sector, exploiting a supply-chain compromise that disrupted fleets using satellite connectivity. Reports and technical analysis attribute the campaign to the group Lab Dookhtegan (International, 2025). The threat actor attacked an Iranian shipping company by exploiting vulnerabilities in their provider of satellite communications and related services (Fanava Group) (Cyber, 2025). The attackers compromised Fanava's central infrastructure and disabled iDirect Falcon, the network management and control platform that coordinates Very Small Aperture Terminal (VSAT) communications between ships and shore stations. By disrupting Falcon's control processes, used for link allocation, modem authentication, and configuration updates, the attackers were able to wipe storage partitions and sever multiple ship-to-shore links simultaneously. Public reporting indicates that communications were interrupted on more than 60 vessels (39 tankers and 25 cargo ships) in one wave, with two subsequent waves affecting up to 116 vessels overall, extending the impact from maritime communications into associated corporate IT and operational monitoring systems (Cydome, 2025). This incident illustrates how a provider-level supply-chain attack can compromise a satellite service intermediary, rendering it a single point of systemic failure for multiple operators and fleets.

Space-derived positioning and timing services have also been targeted. Aviation regulators warned of persistent GNSS jamming/spoofing near conflict zones, and maritime reporting from the Strait of Hormuz linked dense interference episodes to abnormal tracks and a high-profile collision and fire between tankers near Khor Fakkan on 17 June 2025 (EASA, 2024). Regulators and analysts have emphasized that GNSS interference in constrained waterways can disrupt traffic separation and port scheduling and, when timing distribution depends on satellite signals, can cascade into logistics and other timing-sensitive operations (Administration, 2017; Times, 2025; Reuters, 2025; Psiaki and Humphreys, 2016).

These cyber-incidents demonstrate how Space disruptions can propagate because many terrestrial services rely on shared dependencies for connectivity, data, and timing. The practical way to secure these links is to organize operations so that no single satellite link, one provider's control plane, one timing source, or one ground site can take a whole service with it.

For communications, this means building real diversity and automatic failover across independent paths: where terrestrial backhaul exists, keep it as a live alternative; where it does not, combine more than one satellite constellation and provider, and place ground sites in different locations so a fault or compromise at one teleport or cloud region does not become a sector-wide outage.

Management channels should remain accessible even during periods of primary service degradation. This ensures the ability to reconfigure or quarantine fleets of terminals and remote sites, thereby reducing the impact of modem-fleet "wipe" effects and outages on the provider's infrastructure. These principles are consistent with ground-segment security guidance such as NIST IR 8401 (Lightman et al., 2022), which identifies external management planes and shared service providers as high-impact risks requiring appropriate instrumentation and isolation (Cybersecurity and (CISA), 2024; Innovation & Technology, 2025).

Since many cascading failures originate between the space and user segments, it is essential to strengthen both the ground and supplier layers while ensuring link diversity. Best practices include isolating command and mission operations from enterprise IT, maintaining strict control over remote administration, employing modern authentication mechanisms, and staging signed updates with safe rollback capabilities. These requirements should be included in contracts with ground-segment and cloud-based service providers, together with obligations for transparency regarding software provenance. This way, if there is a compromise at an intermediary level, it cannot silently affect dozens of customers, as we have seen in cases of supply-chain breaches.

Positioning, navigation, and timing need specific protection measures because problems with satellite signals affect more than a single vessel or aircraft; they impact air and sea traffic management, port operations, and any activity that depends on precise timing. The goal is to reduce the effect of disruption: use receivers that combine multiple constellations and frequencies; verify the integrity of the received signals (for example, through Galileo's Open Service Navigation Message Authentication European union agency for the space programme (EUSPA), 2025); detect jamming or spoofing and respond through pre-defined operational actions (CISA, 2022); and maintain accurate timing even during signal loss through holdover techniques (European union aviation safety agency (EASA), 2024). Aviation and maritime guidance on PNT interference also adds the procedural layer: planned alternate routes, backup navigation modes, clear reporting channels, and conservative operational defaults when signal trust is uncertain (U.S. department of transportation, 2024).

Detection and coordination should also match how failures actually spread. Keep a living map of external dependencies, such as teleports, cloud regions, remote-access and identity providers, and terminal vendors. Exercise what happens when the largest nodes in that map fail. REF here can genuinely help identify shared, internet-visible weaknesses across different organisations (such as the same exposed administration portal or the same vulnerable version appearing in many places). Those are the points where one exploit can move fastest across an ecosystem; prioritising them creates the biggest reduction in propagation risk.

Finally, propagation is also a problem of information flow. Incident contacts and formats with satellite and ground providers, as well as aviation and maritime coordinators, shorten the window in which an outage in one place becomes a major disruption. Technical and organizational interconnections can cause failures to spread across different operators and sectors. This highlights the importance of exposure measurement approaches.

Some of the conditions enabling the propagation mentioned in the attacks above can also be examined through externally observable exposure. Because propagation typically originates at publicly reachable endpoints, REF can help quantify the factors that enable such initial accesses: shared, exposed services, recurring vulnerable versions across different organisations, and common third-party entry points that multiple operators depend on. REF highlights structural points where a single compromise can propagate across customers, partners, or sectors by identifying clusters of entities exposing identical software stacks, CVEs, or management interfaces.

The same mechanism applies to space-related hardware that appears within other critical infrastructures. When SATCOM user terminals, GNSS receivers, antenna controllers, or similar components are integrated into the networks of energy, maritime, or telecommunications operators, and are reachable or fingerprintable, REF makes it possible to assess their exposure and the extent to which another sector relies on space-linked systems. The KA-SAT incident demonstrated that externally reachable modem fleets and shared management interfaces acted as the propagation vector through which a provider-level compromise extended into the energy sector. REF can detect whether these shared weaknesses persist over time, indicating when systemic propagation conditions remain in place despite mitigation efforts.

Before delving deeper into the REF framework and its features, the next section discusses some of the major contributions in the fields of space cybersecurity and in the analysis of internet-connected devices.

2.4. Related work

Several recent publications have explored the role of cybersecurity in the space domain, and the increase in scientific production on this topic highlights the relevance and timeliness of the subject.

Pavur and Martinovic (Pavur and Martinovic, 2022) map satellite security into four arenas (RF links, in-orbit platforms, ground segment, and mission/ops), catalog threat actors and attack classes, and propose a research agenda to bridge the gap between policy and technology. The author highlights that the constraints of space systems make IT security insufficient.

Khan et al. (2024) provide a holistic review of space cyber-attack vectors spanning ground, space, communications, user, cloud, and supply chain, with a matched set of mitigations and research needs (space-tailored IDS, zero-trust, crypto/keying, blockchain access control, and regulatory/standards work). Their work presents a useful current landscape of cybersecurity risks and a roadmap.

Manulis et al. (2021) frames how New Space changes the threat landscape: historically, most incidents hit ground and RF links, but the widespread development of large constellations created new vulnerabilities in other segments. The article surveys past incidents, adversary motivations, and likely attack classes, and also walks through enabling technologies and where they open or close attack surfaces. It represents an orientation map tying architectural choices to concrete risks.

Yu et al. (2024) carry out an analysis of nine commodity modems, building a taxonomy across Satellite Communication Interface (SCI), Ground Network Interface (GNI), and hardware, and find numerous issues in standard services (web/Telnet/FTP), unprotected control paths, and hardware/bootloaders, plus ICDs at scale. Their work demonstrates practical exploit chains, and catalogs new categories alongside known ones. It also proposes concrete mitigation (service hardening, signed

updates with rollback, securing debug interfaces) and provides tooling for analysis. Strong, device-level evidence that user-segment weaknesses can enable wide-area impact.

These contributions provide a solid overview of vulnerabilities and their underlying reasons, highlighting how different segments of the space infrastructure serve entry points for attacks. However, their focus differs from our study; they discuss threats and suggest mitigation strategies, but they do not introduce a framework for measuring the risks faced by space organizations through their internet-connected devices. Additionally, while they address vulnerabilities and propose mitigations, they overlook the policy dimension of space cybersecurity; while our work includes aims to bridge cybersecurity requirements to operational frameworks to address them.

Another part of the literature connected to our work concerns the analysis of ICDs and their attack surface. Shodan has been widely studied for its role in identifying and indexing ICDs, including critical industrial control system (ICS) components. Alsmadi et al. (2022) has explored Shodan's ability to reveal vulnerabilities in ICS environments, demonstrating its potential as both a reconnaissance tool for attackers and a valuable resource for cybersecurity professionals. Researchers examined how Shodan indexed programmable logic controllers and proposed service banner manipulation as a mitigation strategy to limit exposure to Shodan queries.

Genge and Enăchescu (2016) went beyond basic device discovery supported by Shodan and proposed ShoVAT, a Shodan-based vulnerability assessment tool capable of automatically identifying vulnerabilities using service-specific banner analysis. Through the testing on 1501 services across 12 different institutions, ShoVAT identified 3922 known vulnerabilities, showing the potential of Shodan-powered tools for cybersecurity assessments.

Similarly, Williams et al. (2017) leveraged Shodan to assess vulnerabilities in consumer IoT devices, collecting a large dataset of indexed devices and analyzing their security posture using Nessus vulnerability scans. Their results revealed a significant number of consumer IoT devices vulnerable to exploits that could compromise user data and privacy, underscoring widespread security concerns associated with the growth of IoT.

Liu et al. (2015) showed that an organization's breach likelihood can be predicted purely from externally observable network properties, achieving around 90% accuracy in forecasting data breaches without any internal data. In the industrial domain, Gourisetti et al. (2021). propose a risk assessment framework for energy delivery control systems that leverages Shodan-derived device fingerprints and publicly known vulnerabilities; their method correlates externally gathered device information with threat intelligence to produce risk-based scores for exposed critical assets.

Harry et al. (2025) developed quantitative attack surface indices by combining Internet-wide scan results with CVE data to measure the "size" and "severity" of an entity's exposure. Their framework analyzes thousands of internet-facing devices (using platforms like Shodan/Censys) and computes metrics such as service diversity and aggregated vulnerability severity, which were found to correlate with certain risk factors (e.g. population served) while revealing gaps (e.g. weak correlation between sheer CVE counts and actual exploit likelihood).

These studies show that data from Internet-exposed devices can support cyber risk assessment, but they usually address the problem in a partial way, for example, by focusing on statistical correlations, single-device scores, or specific domains. The analysis is often limited to counting exposed services or vulnerabilities, and does not clarify how these findings map to standard risk concepts or how they can be combined at the organizational level. The Risk Exposure Framework makes links externally observable services and vulnerabilities to exposure, accounts for the availability of exploits in estimating likelihood, and aggregates these elements into an organization-level score. REF goes from data collection to scoring and interpretation, and provides inputs that can complement existing risk management processes. Our work further differs by grounding the use of REF in European cybersecurity policy requirements, by explicitly comparing it with existing risk assessment frameworks, and by showing how it can be used in a complementary way rather than as a replacement. In this paper, we extend this research line to space cybersecurity by focusing on space-sector organizations and by considering the impact of exposed assets in light of the critical role of satellite and aerospace infrastructure.

Although numerous risk-assessment methodologies exist across IT, OT, and space domains, including general-purpose models and sectoral threat-modelling approaches, such as those we discuss in Section 4, these frameworks do not quantify an organisation's externally observable attack surface. They rely primarily on expert judgement, architectural assumptions, or mission-centric threat scenarios, and therefore cannot capture real-world Internet exposure, persistent vulnerabilities, or the availability of exploits affecting ground-segment systems. The contribution of REF lies in addressing this specific measurement gap. With REF we introduce a reproducible, data-driven methodology to assess the externally visible cyber posture of space organisations.

3. Introducing the risk exposure framework

Our *Risk Exposure Framework* (REF) is a structured, quantitative methodology designed to assess the extent of an organization's exposure to cyber threats and translates it into a quantitative risk assessment. Our goal is to understand how organizations such as aerospace manufacturers, satellite service providers, and space agencies are potentially vulnerable to cyber threats. By defining REF, we also aim to study how this type of assessment can be automated and beneficial to private companies and other actors involved in the space value chain.

3.1. Framework design and risk metrics

REF is designed to be automatic and data-source agnostic, meaning that it can be implemented using any dataset that provides relevant information about exposed digital assets, known vulnerabilities, and a broader threat landscape.

Indeed, it supports different tools or cyber intelligence sources depending on availability and context. Although REF can be implemented using various sources, in Section 5, we use Shodan for identifying publicly exposed services and IPs, and Vulners for retrieving exploit availability and vulnerability severity.

The REF framework is built upon three metrics: *Exposure* (E), *Likelihood* (L), and the *Overall Risk Score* (R). Below, we introduce them.

Exposure (E) aims to measure the breadth and depth of an organization's visible attack surface, considering not only the number of publicly accessible IP addresses but also the concentration of vulnerable services exposed. Specifically, we retrieve the list of accessible IP addresses for the organization under analysis. For each IP address $a \in IP$, we compute the number of vulnerable services it exposes, denoted as $\mathcal{V}(a)$, which is proportional to the number of ports on that host affected by known CVEs. These represent potential entry points for adversaries attempting to compromise the system. The Exposure metric aggregates these values over all IP addresses and normalizes the result by the total number of exposed IPs, capturing both intensity and proportionality of exposure across the organization's infrastructure.

Formally, let IP be the set of accessible IP addresses for a given organization, and let $\mathcal{V} : IP \to \mathbb{N}$ be a function mapping each IP address to the number of vulnerable services it exposes (i.e., ports associated with known CVEs). Then, the Exposure score E is defined as:

$$E(IP) = \frac{1}{\#IP} \sum_{a \in IP} \mathcal{V}(a) \tag{1}$$

where $\#IP$ denotes the total number of publicly exposed IP addresses, and the numerator accumulates the total number of vulnerable services across those addresses. The resulting value reflects the average density of vulnerabilities per host, emphasizing infrastructures where exposure

is not only broad but also concentrated in weak points. This formulation ensures comparability across organizations of different sizes while retaining sensitivity to localized high-risk configurations.

Whereas Exposure focuses on what is achievable, the Likelihood metric (L) estimates the danger that an attacker will succeed in exploiting these weaknesses. Likelihood is composed of multiple factors and is formalized by combining *adversary interest* (A), *exploit availability* (EA), and the *severity of the vulnerabilities* (VSI) themselves.

More precisely, adversary interest (A) is a multiplicative constant that represents the attacker's motivation to target a particular space organization. This might depend on political and economic incentives or the strategic value of the satellite or service at hand. We assigned this value to 1 (low interest) for the reasons we explain in Section 3.4.

Exploit Availability (EA) measures the ratio of public exploits known for the vulnerabilities identified. It is the ratio of total exploits available for an address and the total number of unique vulnerabilities. Formally, let $\mathcal{E}$ be a function that returns the number of exploits available for a vulnerability v on an host a, we define Exploit Availability as a function $EA: IP \to \mathbb{R}$ from a set of IP addresses to a numeric value as follows:

$$EA(IP) = \sum_{a \in IP} \frac{\mathcal{E}(v,a)}{\mathcal{V}(a)} \tag{2}$$

Vulnerability Severity Impact (VSI) represents the exploitation probability through the CVSS v3.0 scores, ensuring that even if exploit code is available, the real impact and complexity of each vulnerability are included. We model $VSI: IP \to \mathbb{R}$ as a function from a set of addresses to a numeric value, and we discuss in detail its calculation in Section 3.3.

The combination of A, EA, and VSI measures how much risky a successful cyberattack may be, given the attacker's motivation, the presence of readily exploitable weaknesses, and the seriousness of each vulnerability. Thus, we defined the Likelihood metric $L: IP \to \mathbb{R}$ as a function given by the following formula:

$$L(IP) = A(EA(IP) + VSI(IP)) \tag{3}$$

Once Exposure (E) and Likelihood (L) are computed, they are combined to produce the *raw Risk Score* (RR). As done above, $R: IP \to \mathbb{R}$ is a function from a list of IP addresses to a real number. We propose two methods for computing it. The first averages Exposure and Likelihood, treating the two dimensions equally:

$$RR(IP) = \frac{E(IP) + L(IP)}{2}$$

The second uses a weighted coefficient C to tune the importance of Exposure versus Likelihood, allowing decision-makers to highlight either how large the attack surface is or how probable exploitation might be. In either case, the result is a single Raw Risk Score that informs analysts about the magnitude of the organization's overall cyber risk posture.

$$RR(IP) = CE(IP) + (1 - C)L(IP)$$

Large space organizations, such as those included in our analysis, can have thousands of exposed IPs and hundreds of potential vulnerabilities. Therefore, Raw Risk Scores can inflate to very large values. To address this, we applied a logarithmic scaling to the Raw Risk Score, obtaining the *Overall Risk Score* (R). Formally, we define $R: IP \to \mathbb{R}$ as follow:

$$R(IP) = \log_2(1 + RR(IP)) \tag{4}$$

The log-based approach preserves the relative ordering of risk values while containing the otherwise unbounded numbers in a more comprehensible range, representing risk scores that are easier to interpret, making it simpler to highlight which organizations or missions deserve priority attention and resources.

Applying the REF framework means computing R above. Our method is flexible enough to account for differences in entity size or complexity, while still allowing detailed analysis of vulnerabilities and exploit feasibility. With the REF, we obtain a clear vision of how many real, exposed attack vectors are present and what impact exploits could cause. REF can be used by security engineers, risk managers, and decision-makers as a practical means of identifying risk and prioritizing defenses in an environment where many components are mission-critical and where reliable operation can be essential to national and commercial interests. In the following sections, we dive deeper into the various components of REF and discuss their applications and limitations.

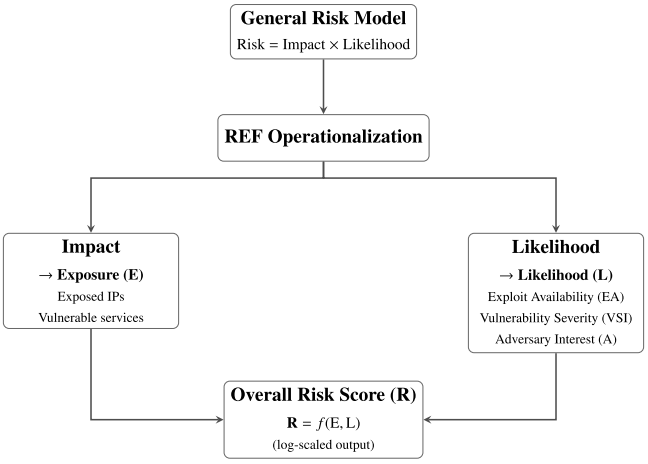

Fig. 1. REF derivation from the general risk model (Risk = Impact × Likelihood).

3.2. Derivation of REF from the general risk model

REF derives directly from the risk model

$$Risk = Impact \times Likelihood$$

a formulation adopted in cybersecurity frameworks such as NIST SP 800-30 and ISO/IEC 27005. In the context of Internet-facing infrastructure, we operationalize Impact through the Exposure (E) metric, which quantifies the breadth and depth of an organization's attack surface, measuring the number of publicly accessible hosts and the concentration of vulnerable services they expose. A higher Exposure score indicates that the potential impact of a successful breach is amplified, as more entry points and vulnerable services translate into greater opportunities for unauthorized access, data exfiltration, or service disruption.

Secondly, we operationalize Likelihood through the Likelihood (L) metric, which estimates the probability that identified vulnerabilities will be successfully exploited. This probability is determined by three factors: adversary interest (A), which reflects the motivation to target specific organizations; exploit availability (EA), which measures the existence of weaponized exploits for discovered CVEs; and vulnerability severity impact (VSI), which incorporates CVSS scores to account for both the complexity of exploitation and the extent of potential damage.

When Computing $Risk = f(E, L)$, E represents impact, and L represents likelihood. REF maintains consistency with the general risk model and grounds each component in observable, quantifiable data drawn from Internet scanning platforms and vulnerability intelligence sources. Fig. 1 illustrates this derivation, showing how the risk formula is mirrored into measurable cybersecurity indicators.

3.3. Assessing Vulnerabilities' severity

The Vulnerability Severity Impact (VSI) is a metric we developed to assess the criticality of an organization's vulnerabilities by evaluating both their frequency and severity, as established by the CVSS v3.0 framework. VSI reflects the real-world impact and ease of exploitation for each vulnerability.

The CVSS v3.0 framework is used for the calculation of VSI and categorizes vulnerabilities based on their exploitability, impact, and environmental factors. This ensures that vulnerabilities that are both severe

and frequently observed within an organization have a significant influence on the likelihood calculation. VSI evaluates how severe an organization's vulnerabilities are, using CVSS v3.0 principles. To calculate VSI, we first calculate the Average CVSS Score of observed CVEs as follows:

$$\text{avgCVSS}(IP) = \frac{1}{\#IP} \sum_{a \in IP} \sum_{v \in \mathcal{V}(a)} \frac{\text{CVSS}(v)}{\#\mathcal{V}(a)}$$

where $\text{CVSS} : V \to \mathbb{R}$ is a function mapping a vulnerability to its CVSS score. Then we define $VSI(IP)$ by cases as follows:

$$VSI(IP) = \begin{cases} 1 & \text{if avgCVSS}(IP) < 5.0 \\ 3 & \text{if } 5.0 \leq \text{avgCVSS}(IP) \leq 8.0 \\ 5 & \text{if avgCVSS}(IP) > 8.0 \end{cases}$$

In our approach, the CVSS v3.0 framework is used to evaluate vulnerabilities by considering both how easy they are to exploit and how damaging they could be. The VSI score captures this by averaging the CVSS scores of an organization's most common vulnerabilities. In this research, VSI values were derived using observed Common Vulnerabilities and Exposures extracted from public internet-facing services, combined with corresponding CVSS v3.0 scores from Shodan and Vulners. The Integration of VSI into our broader Risk Exposure Framework, gives a better visibility to organizations on their security posture. A higher VSI score means an organization has vulnerabilities that are both frequently observed and highly exploitable, making it an interesting target for cyberattacks.

In defining the thresholds of VSI, we adopted the CVSS severity bands, balancing sensitivity and robustness. CVSS outlines qualitative severity bands (*Low*, *Medium*, *High*, and *Critical*) with numeric boundaries (for example, High is defined as a score between 7.0 and 8.9) that serve as reference points. Specifically, we adopted these thresholds to ensure consistency with established practices. However, we also understand that vulnerabilities may tend to cluster in mid-range scores, so it is important that our thresholds do not overly compress the higher severity categories. For this reason, we selected thresholds that maintain clear distinctions among more severe vulnerabilities and also avoid the inflation of the highest category due to small differences in CVSS scores. We use standard CVSS brackets but interpret them with sensitivity to our domain. REF incorporates VSI solely as a component of the Likelihood metric, rather than as a standalone classification for assessing risk. Its role is to calibrate the contribution of exploitability within the overall metric. Finally, we are aware that CVSS has its limitations, such as inconsistencies among evaluators, a lack of contextual awareness, and a static approach to evolving threat conditions (Howland, 2022). For this reason, any thresholding process must acknowledge the need for some degree of interpretive discretion.

3.4. Assessing adversary interest in the space sector

Adversary interest (A) is a critical component of the Likelihood (Eq. (2)) calculation, as it directly reflects the motivation attackers have in targeting a particular sector. Intuitively, a higher adversary interest suggests that cyber threats against the sector are more persistent.

Accurately estimating this metric is very challenging, since an attacker's interest towards a given target is influenced by a wide range of factors, including economic, political, and geopolitical considerations. Defining a precise, objective, and quantifiable metric that incorporates all those aspects is not feasible in practice. Therefore, we must adopt an under-approximation strategy, which by its nature comes with certain limitations. In this paper, we under-approximate this measure using data on historical incidents and trends, taking their frequency into account. Specifically, to objectively estimate the adversary's interest, we analyzed the number of cyberattacks against the space sector compared to other critical infrastructure sectors. We considered the following data for our historical investigations and comparison:

- The ENISA Threat Landscape Report 2024 (European Union Agency for Cybersecurity, 2024a).

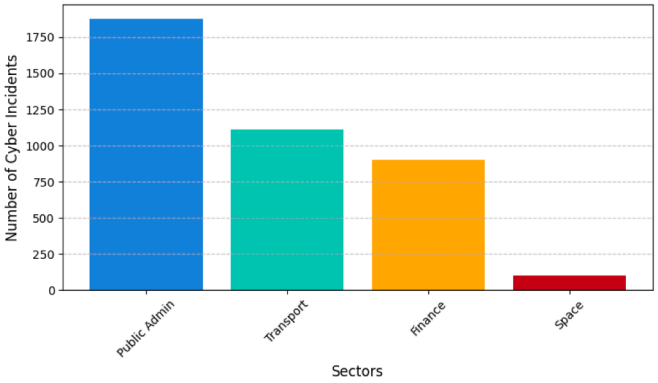

Fig. 2. ENISA Threat Landscape: Attacks against critical sectors 2023–2024.

- The European Repository of Cyber Incidents (EuRepoC) (European Repository of Cyber Incidents (EuRepoC), 2025) in the period 2000–2025. This repository is maintained by an independent research consortium dedicated to better understanding the cyber threat environment in the European Union and beyond.
- Threat intelligence from Cyberinflight (Cyberinflight, 2022), which is one of the founding companies of the EU Space ISAC and a leader in CTI in Europe. We considered a threat intelligence report from 2022.

By analyzing this data, we assign a score to the adversary's interest in the space sector.

Here, we let the adversary interest (A) score range from 1 (low) to 3 (high). The idea is that A should take the value 3 (high interest) for those sectors that are among the top targets for cyberattacks, with a significant share of total incidents. The value 2 (moderate interest) should be given to those sectors that experience moderate but consistent targeting, accounting for around 5–10% of incidents. The value 1 (low interest) should be assigned to those sectors that are rarely targeted, with a share of less than 3% of total incidents.

We use our historical data as follows. First, we analyze the data from ENISA Threat Landscape (period July 2023 – June 2024), shown in Fig. 2. We observe that approximately 100 cyber incidents targeted the space sector, while the number of incidents across all sectors was estimated at 9,900. This means that the space sector accounted for about 1.01% of all cyber incidents within the timeframe considered.

To determine how significant this percentage is, we compare it to the following high-target sectors:

- Public Administration: 1,880 incidents (19% of total)
- Transport sector: 1,110 incidents (11% of total)
- Finance: 900 incidents (9% of total)
- Space: 100 incidents (1.01% of total)

The analysis clearly suggests that the space sector is not yet a primary target for attackers when compared to sectors such as public administration, finance, and transport.

Then, to supplement the previous data, we examine historical records from CyberInflights 2023 threat reports, which document more than 60 cyberattacks against the space sector within 2023, as shown in Fig. 3. The results of the plot align closely with our observations above and indicate that attacks against space assets have remained relatively infrequent so far when compared to other high-risk sectors.

Finally, to achieve a broader perspective, we analyze data from EuRepoC, shown in Fig. 4, which provides the following insights into cyber incidents targeting critical infrastructure:

- Space Sector: 18 incidents (0.51% of all reported cyber incidents globally), representing 1.3% of all incidents targeting critical infrastructure.

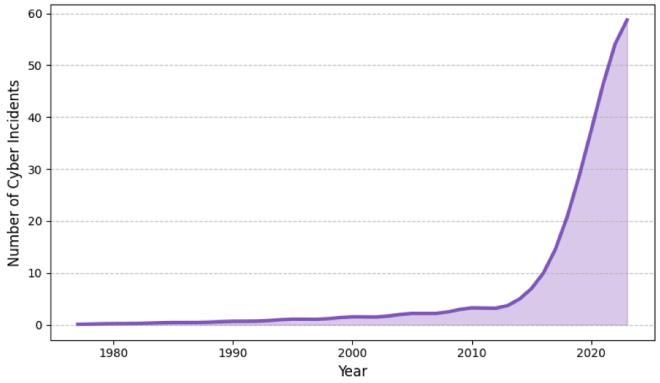

Fig. 3. Evolution of cyberattacks on Space (1977–2023).

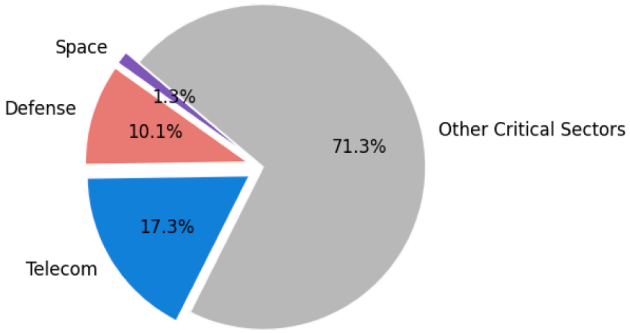

Fig. 4. Cyber incidents in critical infrastructure (2000–2025 EuRepoC).

- Defense Sector: 141 incidents (4.03% of all incidents globally), representing 9.8% of all incidents targeting critical infrastructure.
- Telecommunications: 243 incidents (6.94% of all incidents globally), representing 16.9% of all incidents targeting critical infrastructure.

To assign a score to the space sector, we make the following considerations. The space sector represents 1.01% of incidents (2023–2024 data), 1.3% of critical infrastructure attacks (2000–2025 EuRepoC data), and historically low attack volumes (CyberInfilights 2022 data). This percentage is significantly lower than high-priority targets like public administration (19%), defense (9.8%), and telecommunications (16.9%). Thus, based on our scoring model, the space sector falls into the low interest ($A = 1$) category.

As said above, we are approximating the attack interest through the frequency of past incidents. However, while cyberattacks on the space sector are relatively less frequent compared to other critical infrastructure sectors, this may be related to the nature of the threat actors involved rather than the sector's lack of strategic importance or attractiveness, or to the classification of such attacks. Space assets and networks are widely targeted by state-sponsored hackers and advanced persistent threats (APTs) (Hutchins, 2016; Olszewski, 2018; Young et al., 2025). These types of cyber operations are quite different from the more frequent financially motivated attacks observed in sectors like public administration, finance, healthcare, or retail, often requiring months or even years of reconnaissance, exploitation, and operational execution. The complexity and resource demands of these operations explain why they are less observed if compared to financially driven attacks, which can be cheaply automated. It is only by assigning the same value and characteristics to the threat model across sectors that we conclude that while space assets are targeted by highly capable adversaries, these attacks are less frequent due to the nature of their execution, justifying an $A = 1$ classification.

Another important consideration that should be included in our discussion is the recent upward trend in attacks against the space sector (Pražák, 2021). The past few years have seen a notable increase in cyber incidents affecting space infrastructure, suggesting a growing adversary interest, but the overall volume remains significantly lower than in sectors such as telecommunications, defense, and public administration.

Additionally, we have to consider that many cyber incidents affecting space assets might not be explicitly categorized as space-targeted attacks in existing cybersecurity databases. Space systems are an integral part of telecommunications, energy, and defense networks, and cyber intrusions into these sectors could have originated through vulnerabilities in space-based infrastructure. This interconnectivity creates a blurred boundary in incident classification, and some attacks exploiting space assets as entry points may be recorded under other categories, such as energy or communications. Therefore, the actual impact of cyber threats leveraging space infrastructure is likely underestimated.

In conclusion, the $A = 1$ classification does not imply that the space sector is not a relevant target, but rather that the frequency of cyberattacks – when considering all attackers on the same level – remains low due to the nature of the adversaries involved and the resources required for such operations. The observed trends suggest that this classification may need to be revised in the future as cyber threats against space systems continue to evolve.

3.5. Data collection: Identifying exposed space assets

Our methodology for quantifying cybersecurity exposure within space-sector organizations combines Shodan for host and service discovery with Vulners for exploit intelligence.

To provide a clear overview of the methodology, Fig. 5 illustrates the step-by-step process of the Risk Exposure Framework. The framework follows a linear workflow from initial organization identification through data collection, vulnerability analysis, and final risk score computation.

The first step concerns the identification of space manufacturers, satellite operators, space agencies, and other potentially interesting entities operating in the space sector. This stage involved consulting public lists (European union agency for the space programme, 2025) or referencing known space companies to compile an initial list of relevant organizations. Despite our focus on the space sector, some organizations included in our list are involved in vertical sectors such as defense and aerospace. In Table 1, we summarize some of the information related to the companies we identified in our analysis. The information on the organization name is crucial in our analysis to identify its Internet attack surface. Indeed, during data collection, we leverage Shodan's `org:''<OrganizationName>''` filter to isolate ICDs that appear to be associated with each targeted entity's domain or IP block.

After mapping an organization's potential footprint, we refine our search with the `has_vuln:''true''` filter to query hosts labeled by Shodan as containing known vulnerabilities (CVEs). Shodan tags these vulnerabilities based on observed service banners or software fingerprints. This filter narrows our list to IP addresses that simultaneously belong to the organization's network space and host at least one exposed service carrying a known CVE. For each IP address returned by the query above, we retrieve host data using the Shodan API. This data includes open ports, service banners, potential CPE (Common Platform Enumeration) strings, and CVE references. This host-level information is used for linking each service or port to its reported vulnerabilities and for listing unique characteristics such as version strings or protocols.

From the data retrieved from Shodan, we analyze each IP and store the following information for each of them that allows us to compute the functions $\mathcal{O}$, $\mathcal{V}$ used in the definition of Exposure and Likelihood in Section 3.1:

- The number of vulnerable ports on that IP (ports that Shodan specifically tied to at least one CVE).
- The total number of vulnerabilities detected for that IP.

Table 1
The companies and organizations analyzed in our study. For each company, we report its main domain of activity, its estimated revenue, the main region of operations, and whether it produces dual-use technology and is a government service provider.

Organisation	Domain	Estimated Revenue	Primary Region	Dual-use	Government Service Provider
AIRBUS	Aerospace, Defence, Space	€69B (Airbus, 2025)	Europe	Yes	Yes
Telespazio	Satellite Services, Space Operations	0.7 (Telespazio, 2024)	Europe	Yes	Yes
Viasat	Satellite Communications	€4.28B (Viasat Inc., 2024)	USA	Yes	Yes
SpaceX	Launch Services, Satellite Internet	€8.3B (Sacra, 2024)	USA	Yes	Yes
ESA - Villafranca Satellite Tracking Station	Satellite Tracking / Space Operations	N/A	Europe	Yes	Yes
GMV Aerospace and Defence S.A.	Aerospace, Defence, Space Software	€0.3 B (GMV, 2024)	Europe	Yes	Yes
Gilat Telecom Ltd	Satellite Communications	€0.3B (Gilat satellite networks, 2024)	Middle East	Yes	Yes
Eutelsat	Satellite Broadcasting	€1.5B	Europe	Yes	Yes

Algorithm 1: Risk exposure framework (REF) for space organizations.

Input: Organization name *org_name*
Output: Risk metrics and vulnerability report
`// Step 1: Data Collection`
Initialize Shodan and Vulners APIs;
$shodanResults \leftarrow Shodan.query(org : org\_name \wedge has\_vuln : true)$
`// Step 2: Vulnerability Analysis`
$hostInfo \leftarrow \emptyset$; `// Info on vulnerable hosts`
`// Collect data for each host`
foreach *host h in shodanResults* **do**
  Extract CVEs, CVSS scores, CPEs, and vulnerable ports;
  Store vulnerability data for host h in $shodanResult[h]$;
`// Step 3: Exploit Intelligence`
Collect unique CVEs from all hosts;
Query Vulners API for exploit availability (parallel);
Build the function $\mathcal{V}$ associating each host with its exploit counts;
`// Step 4: Risk Calculation`
$IP \leftarrow \{hostInfo[h].ip$ for each host $h\}$;
`// Exposure score`
Compute $E(IP) = \frac{1}{\#IP} \sum_{a \in IP} V(a)$;
Compute $EA(IP)$ using exploit availability ratio;
Compute $VSI(IP)$ from average CVSS scores;
`// Likelihood`
Compute $L(IP) = A \times (EA(IP) + VSI(IP))$;
`// Raw Risk`
Compute $R_R(IP) = C \cdot E(IP) + (1 - C) \cdot L(IP)$;
`// REF score`
Compute $R(IP) = \log_2(1 + R_R(IP))$
`// Step 5: Report Generation`
Generate summary with risk scores and metrics;
Rank top vulnerable IPs by exposure and exploitability;
Identify most common CPEs and CVEs;
return *Risk assessment report*;

- A list of those vulnerabilities and any relevant details—like CVSS scores—when available.
- The set of CPE strings, if reported, indicating software or operating systems in use.

Each discovered CVE across all IPs is used to perform a query to Vulners to retrieve how many public exploits, namely, proofs of concept, exploit scripts, or modules, exist for that specific CVE. The result is used to build the mapping $\mathcal{E}$ we used in the definition of Exploit Availability (EA) metric of Section 3.1. This data allows for estimating how readily attackers can capitalize on the discovered vulnerabilities. Once for each IP, we have the number of open vulnerable ports, the total number of vulnerabilities, and exploit information, we compute the total count of vulnerable IPs belonging to the organization, the global set of unique CVEs across all hosts, and the overall average CVSS score as described in Section 3.3.

With the per-IP data assembled, we apply the Exposure formula and derive the Likelihood from the aggregated exploit information and the average CVSS score. This calculation yields a measure of how probable it is that these exposed assets could be successfully compromised if targeted.

We then calculate a raw risk value by weighting Exposure and Likelihood, and then apply a logarithmic transformation to maintain interpretability. The final result is direct evidence of publicly visible attack vectors and existing exploit modules, going from a simple port-scan perspective into a more comprehensive threat picture.

Having established the theoretical and mathematical foundations of the Risk Exposure Framework, we operationalize it in Algorithm 1. The pseudocode provides a high-level procedural overview of the REF methodology, delineating how data about organizations' attack surface is systematically collected and processed to derive quantitative risk assessments.

3.6. What REF sees in the space domain

REF focuses on Internet-exposed assets. Therefore, it captures ground-segment and enterprise risks and offers limited visibility on deliberately isolated links or flight software. However, space-specific characteristics often leave observable signals at network boundaries:

- *Isolated environments.* When segmentation and cross-domain controls are effective, externally visible exposure is small or null; conversely, misconfigurations and reachable remote-access paths (e.g., VPN gateways, web management, cloud ground services) appear in REF's exposure and banner data, precisely the pattern ENISA highlights in "Scenario 3 of its space threat assessment: network intrusion due to lack of security protocols and misconfiguration." (European Union Agency for Cybersecurity, 2024a)
- *Limited compute / RT constraints.* Space-qualified/rad-hard processors and deterministic RTOS constraints push conservative change policies and cryptographic agility on adjacent systems; at the boundary, this often shows up as older protocol stacks or slower remediation. REF does not measure on-board resources, but it does register persistent, externally reachable vulnerable versions and weak

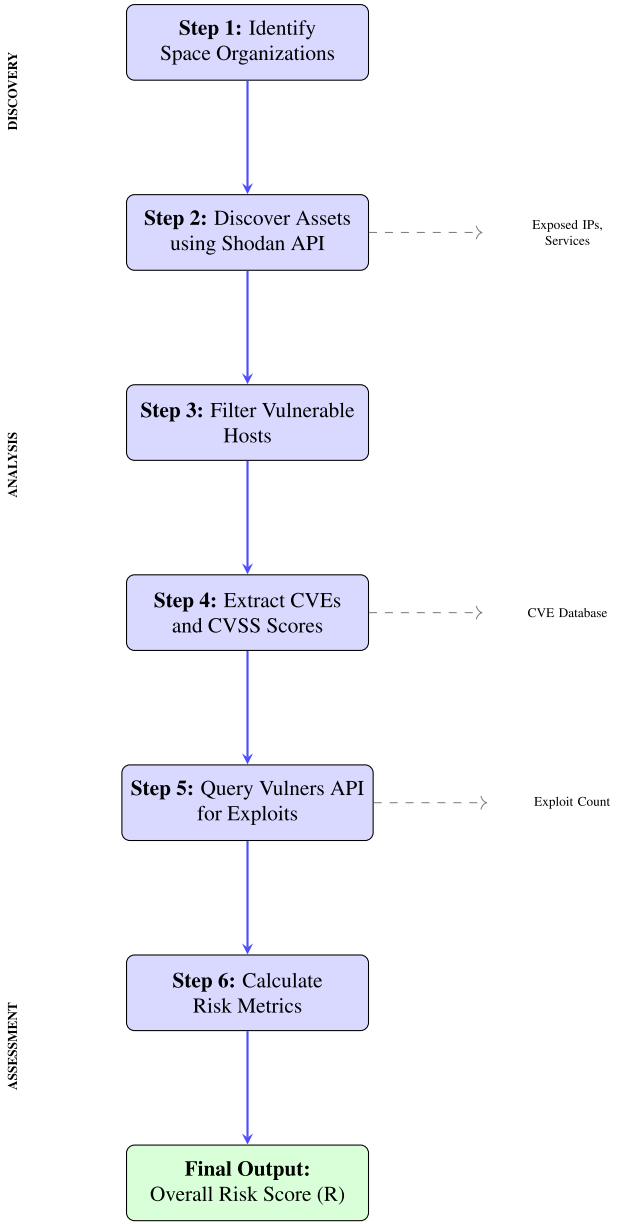

Fig. 5. REF methodology workflow.

protocol fingerprints that can be a downstream effect of such constraints (NASA, 2024; Goodwill, 2024; George et al., 2018; Varadharajan et al., 2024).

- *Firmware/update limitations.* Tight windows and signed over-the-air (OTA) requirements can lengthen patch latency for equipment interfacing with space systems. REF cannot attest to firmware assurance, but when those processes slow remediation, it will observe longer-lived external vulnerabilities on boundary services (Park et al., 2025; Lightman, 2022).

In conclusion, REF reports what a remote adversary can see. It cannot prove air gaps, evaluate flight software integrity, or inspect internal trust zones: those require a mission-internal assessment. Our results and policy mapping already utilize REF in this "first-line triage" role (European Union Agency for Cybersecurity, 2024a).

4. Integrating REF into existing risk assessment methodologies

This section illustrates how REF can be incorporated into some of the existing risk assessment frameworks and processes for space systems. Specifically, we will examine the Notional Risk Score developed by the Aerospace Corporation, the Cybersecurity Risk Assessment for Space Systems created by Honeywell, a case study developed by ENISA, and general-purpose frameworks for or other quantitative risk models such as the NIST SP 800-30 (Risk Assessment Guide) (Initiative, 2012), the FAIR (Factor Analysis of Information Risk) (Jones, 2006). While other frameworks exist, including those developed by NIST, we chose to focus on these due to their operational practicality and relevance. We show that REF can enhance different stages of any risk assessment methodology related to space or other organizations.

4.1. REF and SPARTA NRS

Many ICT risk assessment methodologies often fall short in addressing the dynamic and complex nature of cyber threats in space systems. To bridge this gap, industry and policymakers have requested threat-based risk assessment approaches tailored to space. These methodologies prioritize the identification and analysis of specific threats, such as signal jamming, spoofing, DDoS, or unauthorized access to satellite control systems, to inform the development of targeted mitigation strategies. Focusing on the most probable and impactful threats should allow organizations to allocate resources more effectively and enhance their defensive measures.

One of these frameworks is *Space Attack Research and Tactic Analysis* (SPARTA), developed by the *Aerospace Corporation* (Corporation, 2025), an American nonprofit that operates a federally funded research and development center (FFRDC), and provides technical guidance and advice on all aspects of space missions to military, civil, and commercial customers.

The SPARTA framework defines a taxonomy of potential cyber threats that are specific to spacecraft and space missions. Its primary goal is to support organizations in identifying, understanding, and addressing vulnerabilities across different stages of an attack. The framework organizes threats into high-level tactics, such as *ST0001 - Reconnaissance*, which refers to the phase where threat actors gather information to support future operations, or *ST0002 - Resource Development* that represents the phase where they try to establish resources to support operations. Each tactic is associated with multiple techniques, which describe how an adversary achieves a tactical objective. For instance, a technique might involve exploiting trusted relationships to gain initial access. These techniques can be further broken down into sub-techniques, which represent more specific or granular behaviors. Sub-techniques are variations of the parent technique and often reflect a particular implementation of a threat action, such as compromising mission collaborators (e.g., academic or international partners) to achieve initial access. This hierarchical structure helps organizations trace the progression of potential attacks and design targeted countermeasures accordingly. To quantify and manage cyber risks in space systems, the Aerospace Corporation introduced the *Notional Risk Score* (NRS) within the SPARTA framework. The NRS should provide a metric to evaluate the potential impact and likelihood of cyber threats, assigning numerical values to various threat scenarios. The NRS can be considered a general-purpose framework for its adaptability to assess the cybersecurity risks associated with space systems. It assigns risk scores based on system criticality, likelihood of exploitation, and potential impact. These scores are calculated through a subjective analysis that considers expert assessments of adversary capabilities, system vulnerabilities, and operational priorities. Three elements determine the likelihood of an attack: (*i*) adversary motivation, which is linked to system criticality and assumes that more critical assets attract more attention from threat actors; (*ii*) exploitation difficulty, which considers the complexity of executing an attack technique; and (*iii*) adversary capability, categorized into seven tiers ranging from script kiddies to the most sophisticated state actors. The likelihood assessments are combined with an impact evaluation, which considers potential mission disruption, data integrity loss, or availability concerns. The final risk score is placed on a 5×5 risk

Table 2
Integration of Honeywell risk assessment steps with REF contributions. For each step and activity of the Honeywell risk assessment, we describe how the REF elements can contribute to or complete them.

Risk Assessment Step	Honeywell Risk Assessment Activities	REF Contribution	REF Elements
Security Architecture Creation	Establish system boundaries, identify assets, vulnerabilities, and entry points.	Provides identification and quantification of external attack surfaces and exposed assets.	Exposure (E), Vulnerable IPs, Ports
Definition of Security Risk Elements	Identify risk elements; assign Severity, Occurrence, and Prevention values.	Supplies empirical metrics for exploit availability, vulnerability severity, and adversary interest.	Exploit Availability (EA), VSI, CVSS
Security Model Development	Create security ontologies and threat trees based on identified vulnerabilities.	Offers precise data-driven inputs (CVEs, exploitability) for constructing detailed threat scenarios.	CVEs, Exploit Data (Shodan/Vulners)
Risk Scenario Identification	Aggregate vulnerabilities into structured threat scenarios.	Enables realistic and granular risk scenarios informed by real-world vulnerability and exploit data.	Likelihood (L), Real-world Exploits
Determining Risk Acceptability	Evaluate risk scenarios against acceptance criteria.	Provides quantitative scoring for better-informed risk decisions and clearer acceptability analysis.	Risk Score (R), Risk Scaling (Log)
Continuous Risk Monitoring	Periodic assessment and updates based on new vulnerabilities or threats.	Enables continuous, automated monitoring and updating of the external cybersecurity posture.	Real-time data feeds (e.g., from Shodan)

matrix, where impact and likelihood scores are assigned values from 1 to 5, producing a risk score that falls into a *low*, *medium*, or *high-risk* category.

In practical terms, organizations can employ REF to rapidly identify which Internet-accessible assets are most at risk and derive a specialized Exposure (E) metric. That data can then be fed into an organization's broader application of NRS by becoming part of the "threat likelihood" factor specifically for externally reachable software or systems. Therefore, NRS might be applied to multiple mission segments, and REF provides verified details on real-world, network-facing vulnerabilities. Integrating the two can increase overall risk assessment efficiency, ensuring that the security posture of a specific entity is properly informed by actual vulnerabilities visible on the public internet. While REF focuses on a narrower slice of infrastructure, it can offer an enriched perspective on how easily such vulnerabilities might be exploited, insight that is often lacking in purely theoretical frameworks.

4.2. The REF as part of Honeywell's security risk assessment

In order to give a more comprehensive view of how REF can be integrated into the risk assessment process of space organisations, we considered an additional framework employed by the industry, developed to be applied to space systems (Vessels et al., 2019). We summarized the contributions that REF can give to Honeywell's Risk Assessment in Table 2.

Honeywell's Security Risk Assessment is designed for evaluating security risks in safety-critical systems, including aerospace and space systems. This methodology focuses on identifying, assessing, and managing cybersecurity risks through a clearly defined set of procedures. The first step entails the creation of a comprehensive security architecture, which involves establishing system boundaries, identifying key assets, attackers, vulnerabilities, and possible points of entry. The second step comprehend the definition of security risk elements — assigning quantitative values based on severity, occurrence likelihood, and prevention capability. The methodology involves constructing detailed security models such as threat trees and security ontologies, which help to visualize and analyze the relationships and potential impact of threats. The last step regards the definition of the acceptability of risks according to predefined criteria, guiding decisions on necessary mitigation measures.

In this context, one area where REF provides substantial value is in supporting the creation of security architecture. Traditional risk assessment methodologies require a clear understanding of an organization's attack surface, including vulnerabilities and potential points of entry. REF automates the discovery and quantification of Internet-facing assets and vulnerabilities, directly informing the security architecture and providing empirical evidence on entry points. The integration of REF can help risk assessment teams base their architectural decisions on current and actionable intelligence rather than assumptions or generic risk profiles.

Furthermore, REF enhances the process of defining security risk elements. In various methodologies, risk elements — such as attackers, access vectors, and vulnerabilities — are identified and assigned severity, occurrence, and prevention values based on expert judgment. REF can complement and refine this expert judgment by contributing quantitative, empirically derived metrics related to exposure and exploit availability.

Additionally, REF can help refine security risk models. Honeywell's risk assessment uses threat modeling techniques such as threat trees and security ontologies to represent possible attack paths and outcomes. REF provides concrete, data-driven input for these models, identifying actual vulnerable ports, associated CVEs, and the existence of known exploits, which can be used as inputs to threat trees and scenario analyses. A REF augmented risk analysis can better depict real-world threat scenarios and identify previously unrecognized attack vectors.

4.3. REF and ENISA space threat landscape report

In March 2025, ENISA (European Union Agency for Cybersecurity, 2024b) published the Space Threat Landscape report to drive cybersecurity efforts across the European space sector. The Report outlines threats and realistic attack scenarios that space operators and ground segment entities might face, supporting the development of mitigation strategies. The importance of this report lies in its sector-specific focus, as it addresses the specific challenges that come with managing space assets, both in orbit and on the ground.

One of the critical scenarios highlighted by ENISA is "Scenario 3: Network Intrusion Due to a Lack of Security Protocols and Misconfiguration." This scenario demonstrates how weaknesses in network configuration, such as default settings, the use of unsecured protocols, or missing access controls, can open the door to unauthorized intrusions. The ground segment, a complex web of mission control centers, data relay stations, and tracking infrastructure, is particularly vulnerable to such issues, especially when operational services are inadvertently exposed. REF can support risk evaluation in this specific scenario. Even if it cannot be used as a full forensic audit or to emulate an intrusion detection system, it can flag areas of concern by assessing exposure and estimating the likelihood of exploitation based on known vulnerability data. This process can be useful in the early stages of risk assessment, where broad

visibility into the organization's digital footprint is necessary to identify possible weak spots. Applied to Scenario 3, REF helps identify where misconfigurations might be leading to increased risk, without needing access to internal configurations or administrative policies.

With Scenario 3 ENISA wants to highlight that default settings, insufficient segmentation of networks, and the use of outdated or unencrypted protocols are common among ground operators especially in legacy or multi-tenant environments. REF aligns with this threat landscape because it focuses on externally observable data, such as exposed IP services using weak protocols or devices with known CVEs, which ENISA describes as prime enablers of such intrusions. In this context, REF does not replace internal configuration audits, but identifies publicly visible risk indicators that ENISA flags as early-stage vulnerabilities in cyber kill chains. For instance, if a ground station's control interface is exposed over an unsecured protocol, something ENISA considers as plausible in shared infrastructure scenarios, REF can flag this exposure even without privileged access to the system. Furthermore, REF assigns a dynamic risk score based on exposure and exploitation likelihood that mirrors ENISA's recommendation for continuous risk posture assessments, especially given how quickly threat actors can pivot from reconnaissance to exploitation. ENISA also calls for proactive vulnerability management and improved network hygiene, and here too, REF can provide a scalable first line of analysis by identifying likely misconfigurations, insecure remote access points, and threat-exposed services that deserve in-depth review. The data-agnostic architecture of REF also addresses ENISA's observation that space operators rely on a fragmented toolset and diverse supply chains, offering a platform-neutral method of initial risk triage. In conclusion, REF operationalizes several of ENISA's recommendations: exposure monitoring, risk quantification, and prioritization of mitigation efforts, within a framework capable of automated and continuous external scanning.

4.4. Contextualizing REF within the existing risk assessment landscape

This section analyzes other general-purpose security frameworks and methodologies to better understand how REF fits into them. More specifically, we consider the following:

- NIST SP 800-30 (Risk Assessment Guide) (Initiative, 2012) provides a comprehensive methodology for conducting risk assessments within information systems. This framework employs a structured approach that identifies threat sources, threat events, vulnerabilities, likelihood determinations, impact analysis, and risk determination. It operates primarily through qualitative assessments supported by semi-quantitative scales (from Very Low to Very High), requiring extensive organizational knowledge and expert judgment to evaluate risks across the full system lifecycle.
- FAIR (Factor Analysis of Information Risk) (Jones, 2006) represents a quantitative risk assessment model that attempts to standardize risk terminology and measurement. FAIR decomposes risk into discrete factors: threat event frequency (contact frequency × probability of action) and loss magnitude (primary and secondary losses). Contact represents the frequency at which a threat agent encounters a vulnerable asset, while action refers to the likelihood that the agent, having made contact, decides to attack. FAIR uses probabilistic methods and Monte Carlo simulations to generate risk estimates in monetary terms, making it particularly valuable for business decision-making and cyber insurance contexts.
- CVSS (Common Vulnerability Scoring System) provides a standardized method for rating IT vulnerabilities. Version 3.0 calculates scores based on three metric groups: Base (exploitability and impact characteristics), Temporal (current exploit techniques and remediation levels), and Environmental (modified base metrics considering local requirements). While CVSS excels at individual vulnerability assessment, it does not inherently provide organizational or system-level risk aggregation.
- ENISA Risk Assessment Frameworks encompass multiple specialized matrices designed for European critical infrastructure operators. ENISA has developed several complementary assessment frameworks, including the 5G Security Controls Matrix (ENISA, 2023), the Interoperable EU Risk Management Framework, and the PETs Control Matrix (ENISA, 2016). Each framework addresses specific technological domains while maintaining alignment with EU cybersecurity directives and supporting standardized risk evaluation across member states. The 5G Security Controls Matrix consolidates technical and non-technical controls relevant for 5G network security, mapping various technical controls to an established EU supervisory framework for telecoms. The Interoperable EU Risk Management Framework evaluates interoperability features of risk management frameworks using a four-level scale and provides a methodology for assessing potential interoperability among different frameworks. The PETs Control Matrix includes systematic assessment criteria for privacy-enhancing technologies, distinguishing between generic criteria (applicable to all tools) and specific criteria (assessing particular functionalities).

REF represents a distinct niche within this ecosystem of risk assessment frameworks, functioning as an automated, data-driven exposure quantification tool rather than a comprehensive risk management system. Unlike the frameworks mentioned above, REF specifically targets the observable external attack surface, providing empirical measurements that can feed into broader risk assessment processes. Table 3 compares our framework with the other mentioned above, focusing on their main characteristics.

Rather than replacing the above frameworks, REF should be positioned as a complementary tool that enhances their implementation through empirical data injection at specific assessment phases. Organizations implementing NIST SP 800-30 can leverage REF outputs as empirical inputs for likelihood estimates, particularly during the vulnerability identification and threat-vulnerability pairing phases, transforming what would typically be expert judgment into data-backed assessments. Within FAIR models, REF's exposure metrics can inform Contact Frequency calculations with actual reconnaissance data, while its exploit availability statistics support Probability of Action estimates with real-world exploit prevalence rather than theoretical assessments. REF's aggregated vulnerability view complements individual CVSS scores by providing organizational context for prioritizing patch management efforts, helping security teams understand not just which vulnerabilities are severe, but which severe vulnerabilities are actually exposed and exploitable in their infrastructure. For organizations using ENISA risk matrices, REF enables dynamic updates to likelihood assessments through continuous measurement, transforming static annual assessments into living risk documents that reflect the current threat landscape.

The radar chart in Fig. 6 compares the four major risk assessment frameworks discussed above with REF across eight critical capability dimensions. Each framework displays a distinct capability profile, revealing strategic positioning rather than direct competition. The chart compares the REF to other established risk assessment frameworks across eight capability axes. The comparison is based on a structured qualitative-quantitative evaluation; the methodology builds on established practices of comparative evaluation of risk assessment frameworks (Singh and Joshi, 2018; Jouini and Rabai, 2016).

Scores were assigned by mapping each framework's published features, scope, and methodological constructs to the eight axes defined in the figure. For each dimension, a definition and observable criterion were established (e.g., presence of automation, use of external data, and financial quantification). We examined each framework to verify whether these characteristics were explicitly implemented, partially addressed, or not addressed at all. Each axis was rated on a discrete 1–5 scale, where 1 means minimal or implicit capability and 5 denotes full integration. The axes represent eight dimensions of cybersecurity assessment capability. Automation Level reflects the degree to which data

Table 3
Comparative analysis of risk assessment frameworks.

Framework	Scope	Methodology	Automation	Data Requirements	Output Type	Space-Specific	Integration with REF
REF	External attack surface	Quantitative (E × L)	Fully automated	Internet scan data, CVE/exploit DBs	Numerical risk score	Adaptable	N/A (baseline)
NIST SP 800-30	Full system lifecycle	Qualitative/ Semi-quantitative	Manual	Internal documentation, expert judgment	Risk levels (VL-VH)	No	REF provides data for Steps 2-2, 2–3
FAIR	Information risk (monetary)	Quantitative probabilistic	Partially automated	Historical loss data, threat intelligence	Financial risk ($)	No	REF informs Contact Frequency, Vulnerability
CVSS	Individual vulnerabilities	Quantitative scoring	Manual/Semi-automated	Vulnerability characteristics	Severity score (0–10)	No	REF aggregates CVSS scores
ENISA Matrices	Sector-specific threats	Qualitative (5×5 matrix)	Manual	Threat scenarios, impact assessment	Risk categories	Sector-dependent	REF quantifies likelihood dimension
SPARTA NRS	Space mission threats	Semi-quantitative	Manual	Adversary analysis, system criticality	Risk matrix placement	Yes	REF supplies exposure metrics
Honeywell SRA	Safety-critical systems	Model-based (threat trees)	Partially automated	System architecture, threat models	Risk acceptability	Yes (aerospace)	REF populates asset inventory

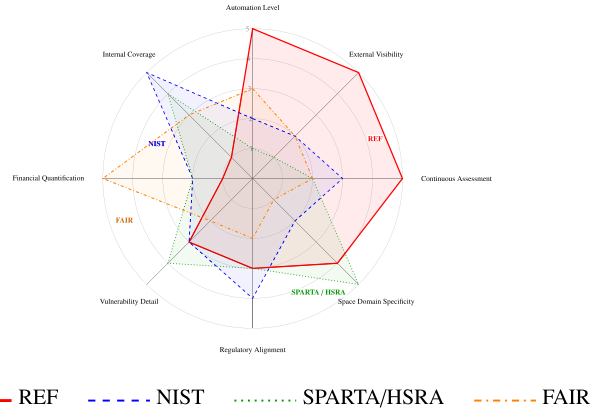

Fig. 6. Radar comparison across eight capability axes.

collection and risk computation occur without manual intervention. External Visibility measures reliance on externally observable data such as exposed IPs or services. Continuous Assessment captures whether the framework supports periodic or real-time reassessment of risk. Space Domain Specificity indicates explicit adaptation to the space domain. Regulatory Alignment evaluates consistency with formal risk management standards (e.g., NIS2, NIST SP 800-30). Vulnerability Detail expresses the depth of vulnerability modeling, including CVE/CVSS integration. Financial Quantification refers to the framework's capacity to express outcomes in economic or loss-expectancy terms. Internal Coverage assesses inclusion of organizational, procedural, and human-factor controls beyond the technical perimeter.

Our analysis results in the following main takeaways:

- *REF's Distinctive Strengths*: REF demonstrates good performance in three interconnected areas: Automation Level (5/5), External Visibility (5/5), and Continuous Assessment (5/5). REF achieves moderate scores (3/5), Regulatory Alignment, and Vulnerability Detail, indicating its ability to support specialized applications while maintaining broad applicability.
- *Complementary Framework Profiles*: NIST RMF excels in Internal Coverage (5/5) and Regulatory Alignment (4/5) while showing limited automation and external visibility capabilities. SPARTA/Honeywell frameworks achieve maximum Space Domain Specificity (5/5) and strong Vulnerability Detail (4/5), reflecting their specialized focus on space-specific threats and attack scenarios. FAIR's strength lies in Financial Quantification (5/5).
- *Strategic Integration Implications*: The non-overlapping capability profiles suggest that these frameworks function as complementary rather than competing solutions. REF's capabilities can feed quantitative data into NIST's comprehensive internal risk management processes, while SPARTA/Honeywell's space-specific expertise can provide contextual interpretation of REF's vulnerability findings. FAIR's financial quantification can translate REF's technical risk scores into economic impact assessments for decision-making.
- *Coverage Gaps and Synergies*: Systematic coverage gaps emerge when frameworks are used independently. REF's minimal Internal Coverage (1/5) and Financial Quantification (1/5) scores indicate clear boundaries for its application scope. However, these limitations become strengths in integrated approaches where REF provides automated external assessment while other frameworks address internal risks, sector-specific threats, and economic impacts.

The comparison supports the integration strategy we defined above and demonstrates that effective cyber risk management requires the coordinated deployment of specialized frameworks. In the following section, we show in detail how REF can be operationalized in practice.

4.5. REF operationalized

This section presents the example of an operational playbook to show how organizations can act on REF outputs in concrete, repeatable ways. A playbook, in this context, is a structured procedure that translates REF's empirical measurements into operational decisions, responsibilities, and evidence. The purpose is to demonstrate that REF can be a decision-support tool that guides remediation and compliance.

Playbook A aims to operationalize a weekly REF run into vulnerability management. The REF outputs are used to prioritize exposed assets, create remediation tasks with fixed deadlines (service-level agreements, SLAs), and generate audit-ready evidence of risk reduction. Table 4 encodes the prioritization policy: each row describes the condition under which an exposed asset is classified (e.g., public exploit available, high CVSS severity, weak management protocols), the action required, the deadline, and the accountable owner.

In the first row (P1 - Critical), the policy addresses assets whose exposure is either actively exploitable (EA > 0), or carries a high technical severity (CVSS ≥ 8.0), or is accessible by an administrative/control interface over a weak protocol. The required action is immediate patching or removal of the exposure. The SLA is 72 hours, reflecting the urgent threat level. The asset owner is accountable, and in the meantime, the SOC should place the IP:port on the SIEM watchlist and enforce blocking via IDS/EDR.

In the "Containment override" row, the trigger is any Internet-exposed control plane or clear-text management protocol (such as Telnet or HTTP admin). In this case, the threshold is not a CVSS score alone, but a weak control mechanism providing attacker access. The required action is de-exposure followed by a scheduled protocol replacement; the SLA is 24 hours because the risk of compromise via exposed control interfaces is very high. Interim controls include network blocks and protocol-specific IDS signatures. The treatment here aligns with prioriti-

Table 4
Playbook A — Prioritization policy for weekly REF runs. Each row specifies how an exposed asset is classified, what must be done, and within what deadline.

Priority tier	Trigger condition (REF evidence)	Required action	Deadline (SLA)	Accountable role	Interim SOC control
P1 — Critical	Exploit available (EA > 0) **or** CVSS ≥ 8.0 **or** exposed admin/control over weak protocol	Immediate patch or de-expose; emergency change permitted	72 hours	Asset owner	Add IP:port to SIEM watchlist; enforce IDS/EDR blocking
P2 — High	5.0 ≤ CVSS < 8.0, EA = 0	Patch or harden configuration; document exceptions	7 days	Asset owner	Targeted anomaly detection; monitor exposure persistence
P3 — Medium	CVSS < 5.0, EA = 0	Patch during maintenance cycle	30 days	Asset owner	Baseline monitoring
Containment override	Any Internet-exposed control plane or clear-text management protocol (e.g., Telnet, HTTP admin)	Immediate de-exposure; schedule protocol replacement	24 hours	Asset owner	Temporary network blocks; protocol-specific IDS signatures

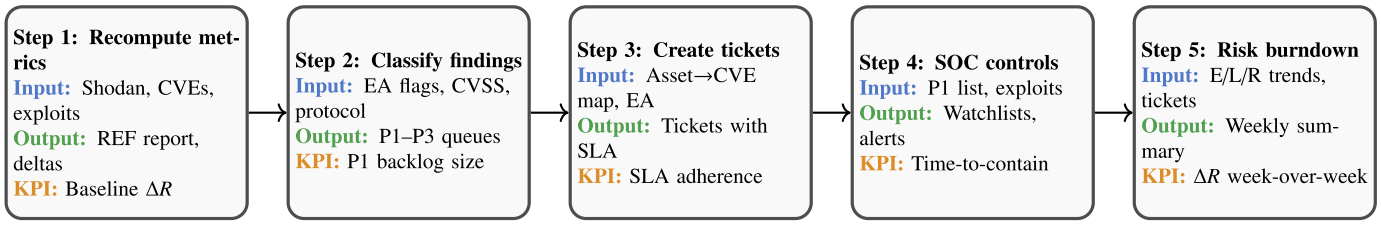

Fig. 7. Playbook A — Weekly REF loop. A horizontal process pipeline: each step has standardised **Inputs**, **Outputs**, and **KPIs**.

sation frameworks that stress exploitability, exposure context, and business risk over severity alone.

Fig. 7 then lays out the weekly loop: five ordered steps from measurement to remediation and reporting. The aim is to make the process explicit for security personnel, demonstrating how weekly REF runs drive prioritization, remediation, SOC monitoring, and measurable risk reduction.

In Step 1 REF recomputes all exposure and risk metrics using Shodan data, CVEs, and exploit information. This produces an updated REF report and a baseline change in the overall risk score (ΔR). Step 2 classifies the results: assets are grouped into priority tiers (P1–P3) depending on exploit availability, CVSS score, or weak protocols.

In Step 3, these classified findings are converted into remediation tickets, each with a clear SLA and assigned owner. Step 4 runs in parallel within the SOC: critical assets (P1) are added to watchlists, and temporary defences such as IDS rules or EDR blocking are applied while fixes are underway. Finally, Step 5 collects all progress data, how many vulnerabilities were closed, how fast containment happened, and produces the weekly risk burndown report.

5. Experimental results and discussion

In this section, we describe how we applied the methodology defined in Section 3 to conduct a risk assessment based on the exposure of the eight companies of Table 1. The code for reproducing our experiments is available in the online repository (Casaril and Galletta, 2025a). However, in the description of our results below, we anonymize the name of the organization for security reasons. We then present and discuss the results of our analysis of the major CVEs identified across the various organizations, highlighting key findings and summarizing the main takeaways. Finally, we validate the methodology on synthetic data and discuss the threat to validity of our results.

5.1. Identification of vulnerable space organizations

Our assessment provides insights into the security posture of the space industry and organizations in the sector at large. Our analysis considers eight leading organizations operating globally, resulting in a combined total of 99,957 publicly exposed IP addresses, among which 4102 were identified as vulnerable, having known associated CVEs. This represents 4.1% of vulnerable hosts over the total exposed ones. The combination of the Shodan and Vulners APIs gives us a snapshot of the industry's current cybersecurity posture, underlining widespread and organization-specific vulnerabilities. A comprehensive summary of the quantitative results from this analysis is available in Table 5. To contextualize our analysis, we present the ranges and trends observed in the main elements of REF. Exposure scores, which reflect the scale and visibility of ICDs, span from as low as 8.22 (Company 1) to nearly 96.48 (Company 5), with a mean value of approximately 37.78. This indicates substantial variability in the online footprint of different organizations, from tightly segmented infrastructures to those with highly exposed configurations. Likelihood scores, which estimate the probability of exploitation based on CVE severity and exploit availability, vary between 6.84 and 7.40, with a mean of 7.17. The distribution remains relatively narrow, suggesting that most organizations share a comparable level of latent vulnerability once exposure is established. The average CVSS score for the vulnerabilities found ranges from 6.67 (Company 7) to 7.08 (Company 3), reinforcing the trend that most exposures involve high-severity vulnerabilities according to standardized scoring. The Overall Risk Score ranges from a low of 3.13 (Company 1) to a high of 5.72 (Company 5), with an average around 4.48. Entities such as Company 5 (5.72), Company 2 (4.87), and Company 8 (4.56) fall into the higher risk bracket, combining moderate-scale exposure with a notable concentration of high-severity vulnerabilities and available exploits. Conversely, Company 1 shows the lowest Overall Risk Score, reflecting a relatively modest footprint and reduced vulnerability context. Looking specifically at the examined organizations, Company 5 demonstrates the highest Overall Risk Score, with a value of 5.72. Despite having only 163 publicly exposed IPs, 72 of these were found to be vulnerable, with a total of 547 unique CVEs and 417 available exploits. This leads to an exceptionally high Exposure Score of 96.48, by far the largest in our sample. The average CVSS score associated with these vulnerabilities is 6.77, further underscoring their severity. This entity shows that even compact infrastructures when characterized by concentrated exposure and poor mitigation, can carry critical levels of cyber risk. Company 2 follows with a scaled risk score of 4.87. It has

Table 5
Cybersecurity risk metrics for space organizations under test sorted by the value of the overall risk score.

Organization	Exposed IPs	Vuln. IPs	CVE Count	Exploits	Exposure Score	L Score	Avg. CVSS	Overall Risk
Company 5	163	72	547	417	96.48	6.84	6.77	5.72
Company 2	1042	73	699	717	49.32	6.99	6.77	4.87
Company 8	2000	34	725	590	37.96	7.37	7.03	4.56
Company 6	44	15	70	101	29.62	7.36	7.07	4.28
Company 4	62,598	2533	560	671	25.28	7.10	6.84	4.10
Company 7	5420	125	810	784	25.11	6.99	6.67	4.09
Company 3	25,537	881	572	662	24.25	7.40	7.08	4.07
Company 1	1253	269	233	380	8.22	7.33	6.87	3.13

a moderate number of exposed IPs (1,042), among which 73 are vulnerable, linked to a striking total of 699 CVEs and 717 known exploits. The Exposure Score of 49.32 is second highest overall, indicating that the organization's digital footprint, although not the largest, is heavily loaded with accessible and potentially exploitable services. The average CVSS score is again high (6.77), reinforcing the critical nature of the exposure. The Likelihood Score for this organization is also significant, contributing to its high risk positioning. Company 8 ranks third, with a total of 2000 exposed IPs and only 34 of them vulnerable. However, those 34 IPs correspond to 725 unique CVEs and 590 known exploits. This density of vulnerabilities results in an Exposure Score of 37.96 and a Likelihood Score of 7.37, one of the highest in our dataset. Combined with an average CVSS score of 7.03, also above the mean, this positions the organization in a high-risk tier despite a seemingly small vulnerable surface. Company 6 presents an illustrative case of risk intensification in smaller infrastructures. With only 44 exposed IPs and 15 vulnerable ones, it would normally be classified as low-exposure. However, the 70 CVEs identified, and 101 exploit matches raise the Exposure Score to 29.62. The organization's Likelihood Score of 7.36 and average CVSS of 7.07, among the highest across the dataset, result in a Risk Score of 4.28. This confirms that compact infrastructures are not inherently safer and may pose comparable levels of risk when vulnerabilities are severe and easily exploitable. Company 4 holds a unique position due to its large-scale online presence. It records the highest number of exposed IPs (62,598) and vulnerable IPs (2,533), with 560 unique CVEs and 671 available exploits. However, its Exposure Score is relatively moderate at 25.28. The Likelihood Score of 7.10 and average CVSS of 6.84 still contribute to a notable Overall Risk Score of 4.10. This organization's vast scale inherently raises the stakes, as a small percentage of vulnerable assets can translate into a significant attack surface. Company 7 shows an Overall Risk Score of 4.09, with 5420 exposed IPs and 125 vulnerable. These are associated with 810 CVEs and 784 exploits, generating an Exposure Score of 25.11. The CVSS average is slightly lower at 6.67, but the concentration of exploits and elevated Likelihood Score of 6.99 contribute to its risk profile.

Company 3, despite its relatively large footprint (25,537 exposed IPs), registers a lower Risk Score of 4.07. With 572 CVEs and 662 matching exploits, the Exposure Score is 24.25. Notably, Company 3 records the highest Likelihood Score in the entire dataset (7.40) and the highest average CVSS score (7.08), pointing to particularly exploitable vulnerabilities. Finally, Company 1 registers the lowest Overall Risk Score at 3.13. It has 1253 exposed IPs and 269 identified as vulnerable, linked to 233 CVEs and 380 exploits. Although the CVSS average is a mid-range 6.87, the relatively low Exposure Score of 8.22 keeps the risk profile limited. This analysis represents a snapshot of the cybersecurity state of specific space organizations but also highlights broader implications for the industry at large. The variability in cybersecurity risk profiles across entities reflects the inconsistent maturity levels in vulnerability management practices and cybersecurity governance. Organizations with expansive and complex digital infrastructures, such as Company 4 and Company 3, face distinct challenges in continuously mapping and securing their digital assets, leading to significantly elevated exposure scores and higher exploitability. Smaller entities with fewer digital assets, like Company 6 or Company 5, face risks that are fewer in number but potentially equally impactful, emphasizing the need for tailored, context-specific cybersecurity approaches.

Table 6
Top observed CPEs by organization and occurrence count.

Organization	CPE	Count
Company 1	`cpe:/a:f5:nginx`	81
Company 1	`cpe:/a:sendmail:sendmail:8.15.2%2f8.15.2`	80
Company 2	`cpe:/o:canonical:ubuntu_linux`	18
Company 2	`cpe:/o:linux:linux_kernel`	16
Company 3	`cpe:/a:f5:nginx:1.5.1`	15
Company 3	`cpe:/a:getbootstrap:bootstrap`	14
Company 4	`cpe:/a:jquery:jquery:3.5.1`	22
Company 4	`cpe:/a:openbsd:openssh:9.3`	21
Company 5	`cpe:/a:apache:http_server:2.4.6`	40
Company 5	`cpe:/a:openssl:openssl:1.0.2k`	34
Company 6	`cpe:/a:f5:nginx:1.18.0`	25
Company 6	`cpe:/o:canonical:ubuntu_linux`	25
Company 7	`cpe:/a:openbsd:openssh:7.4`	25
Company 7	`cpe:/o:canonical:ubuntu_linux`	19
Company 8	`cpe:/a:getbootstrap:bootstrap`	6
Company 8	`cpe:/a:f5:nginx:1.10.3`	4

5.2. Analysis of exposed services and critical attack surfaces

A substantial part of our analysis involves examining the most prevalent exposed services and vulnerabilities. In practice, we extract the most observed Common Platform Enumerations (CPEs) and CVEs from Shodan data. CPEs are structured naming schemes for information technology systems, software, and packages. Through this process, we identified technological assets and vulnerabilities that are particularly widespread, highlighting the most critical areas needing attention in the snapshot of companies we considered. Given the size of our sample, we expect that such findings reflect weaknesses present in the whole space sector.

By aggregating the results globally across all organizations, we identify some clear patterns regarding commonly exposed technologies and their unpatched vulnerabilities. These results are summarized in Table 6.

The table displays the scan results for the ten most common CPEs identifiers across all organizations, along with the number of separate hosts associated with each identifier. A CPE string follows the NIST format `part:vendor:product[:version]`. The initial letter indicates whether the entry represents an application (a), operating system (o), or hardware (h). The next two fields specify the supplier and product, with an optional fourth field that indicates the version. For example, the identifier `cpe:/a:f5:nginx` refers to the Nginx web server application maintained by F5, regardless of the version. In contrast, the identifier `cpe:/o:canonical:ubuntu_linux` encompasses every release of Canonical's Ubuntu operating system.

Among the most frequently observed software were servers such as Apache HTTP Server, Nginx, and OpenSSH implementations, as well as widely-used operating systems like Ubuntu Linux, FreeBSD, and Microsoft Windows. But also, client-side frameworks and libraries such as jQuery and Bootstrap showed a significant presence across multiple organizations.

Table 7
Top CVEs by organization, including frequency, severity, and type.

Company	CVE	Occurrence	CVSS Score	CVE Description
Company 1	CVE-2013-2220	81	7.5	Buffer overflow / PHP
Company 1	CVE-2024-11233	81	9.8	Buffer overflow / PHP
Company 2	CVE-2009-0796	22	5.0	Cross-site scripting (XSS) / Apache HTTP Server
Company 2	CVE-2013-4365	49	7.5	Heap-based buffer overflow / Apache HTTP Server
Company 3	CVE-2020-11022	22	6.1	jQuery / Arbitrary code execution
Company 3	CVE-2020-11023	22	6.1	jQuery / Arbitrary code execution
Company 4	CVE-2008-3844	35	6.8	Externally introduced modification / Red Hat Enterprise Linux
Company 4	CVE-2023-51767	35	7.8	OpenSSH / Authentication bypass
Company 5	CVE-2009-0796	22	5.0	Cross-site scripting (XSS) / Apache HTTP Server
Company 5	CVE-2013-4365	49	7.5	Heap-based buffer overflow / Apache HTTP Server
Company 6	CVE-2009-0796	22	5.0	Cross-site scripting (XSS) / Apache HTTP Server
Company 6	CVE-2013-2765	8	4.3	Apache HTTP Server / DoS
Company 7	CVE-2008-3844	42	6.8	Externally introduced modification / Red Hat Enterprise Linux
Company 7	CVE-2021-36368	42	7.2	OS
Company 8	CVE-2013-2765	11	4.3	Apache HTTP Server / DoS
Company 8	CVE-2013-4365	69	7.5	Heap-based buffer overflow / Apache HTTP Server

Table 7 summarizes the most frequent CVEs detected by our analysis: we observe a consistent presence of outdated HTTP servers and their associated vulnerabilities. For instance, Company 5 notably shows a heavy dependence on legacy software stacks, particularly older Apache HTTP servers (e.g., version 2.4.6 and 2.4.37) coupled with outdated OpenSSL versions (1.0.2k). Specifically, the Apache vulnerability CVE-2013-4365 appeared 49 times in Company 5 systems. This vulnerability also appears 11 times in Company 8. These vulnerabilities enable attackers to perform denial-of-service (DoS) attacks or execute arbitrary code. CVE-2024-11233 appears 81 times within Company1ı́s infrastructure alone, making it one of the most frequent critical vulnerabilities in the dataset. CVE-2024-11233 and CVE-2024-11234 are particularly recent and critical, associated with PHP version 8.2, indicating that some organizations, such as Company 1, continue to use vulnerable and potentially unpatched web frameworks.

Open-source libraries and client-side frameworks such as jQuery represent another significant risk area. Notably, vulnerabilities CVE-2020-11022, CVE-2020-11023, CVE-2019-11358, and CVE-2015-9251 appeared 22 times each. These vulnerabilities, prominently found in the digital infrastructures of Company 4, Company 3, and Company 7, primarily enable Cross-Site Scripting (XSS) and client-side injection attacks, potentially compromising web interfaces or allowing attackers to exfiltrate sensitive operational data. The presence of vulnerable versions of jQuery and Bootstrap suggests insufficient monitoring of software supply chains and inadequate update processes, meaning that automated dependency management tools and regular software composition analysis are still scarcely present.

Vulnerabilities related to cryptographic libraries, particularly outdated versions of OpenSSL are also widespread. For instance, vulnerabilities CVE-2023-48795, CVE-2023-38408, and CVE-2007-2768 were detected 42 times in Company 7 systems. Their presence underscores the neglected updates of security-critical components.

In addition, Table 7 shows a high prevalence of old versions of Nginx (notably, 1.18.0 and 1.10.3) in companies like Company 6 and Company 7. While Nginx itself is widely regarded as secure, outdated versions expose companies to known exploits. Specifically, CVEs associated with older Nginx versions, like CVE-2021-36368, emphasize the urgent need for the sector to implement strict patch management policies.

From the broader industry perspective, these frequently recurring vulnerabilities and exposed services represent significant attack surfaces, increasing the potential impact of cyber incidents. The space sector's unique operational context intensifies the risk implications. Compromise or disruption to satellite tracking stations, telemetry data centers, or satellite communication networks could lead to severe operational, economic, and potential safety implications. Specifically, for tracking stations, reliance on vulnerable systems exposes critical satellite operations and sensitive telemetry data, thereby risking operational continuity and strategic assets. Mitigation strategies must combine different approaches, targeting not just technical remediation but also organizational processes and cybersecurity culture. The organizations should first implement a robust and automated patch management lifecycle, ensuring swift deployment of security patches across all software stacks and prioritizing systems that directly interface with satellite control or telemetry operations. Second, they should implement continuous vulnerability management programs, integrating automated vulnerability scans and dependency tracking, in order to ensure organizations maintain real-time visibility into their cybersecurity posture. The ongoing presence of critical vulnerabilities in key software components highlights significant cybersecurity gaps within the space industry. Addressing these vulnerabilities through timely patch management, continuous vulnerability monitoring, and proactive security measures is the only way forward. These targeted cybersecurity strategies will greatly reduce the attack surface, thereby enhancing the industry's overall security posture and resilience.

5.3. Result validation

To validate our REF methodology, we perform two evaluations. The first consists of applying REF on a synthetic dataset to study how the REF score varies under the variations of the variables. The second evaluation is a sensitivity analysis based on a tornado diagram. In Table 8, we summarize all the formulas used in REF for clarity and reference during validation.

5.3.1. Synthetic dataset

We generated and analyzed a synthetic dataset (Casaril and Galletta, 2025a), designed to test how the metrics respond under controlled variations of key variables. Through this process, we aim to empirically demonstrate the consistency, sensitivity, and reliability of our approach. The dataset of the synthetic organization, named *SpaceTes*, is structured into 4 separate scenarios (identified with letters from 'A' to 'D'), and is designed to isolate and examine the impact of specific variables on the Exposure Score and subsequently on the Overall Risk Score.

In Scenario A, we vary the average number of vulnerabilities per IP, keeping the total exposed IPs and vulnerable IPs constant. The results in Fig. 8 show again a linear growth in Exposure, as expected.

In Scenario B (Fig. 9), while we let every host expose ten vulnerable services, we progressively add more vulnerable IP addresses, letting the total address-space size grow. Exposure rises with the square of the vulnerable-IP tally. The red dashed curve in Fig. 9 therefore climbs steeply: doubling the vulnerable surface from 20 to 40 hosts quadruples Exposure, and moving from 50 to 100 hosts multiplies it by four once again. The plot confirms that even modest expansion of an exploitable footprint drives a non-linear surge in Exposure, underscoring

Table 8
REF metrics and corresponding formulas.

Metric	Formula	Ref.
Exposure (E)	$E(IP) = \frac{1}{\#IP} \sum_{a \in IP} \mathcal{V}(a)$	(1)
Exploit Availability (EA)	$EA(IP) = \sum_{a \in IP} \frac{\mathcal{E}(v,a)}{\mathcal{V}(a)}$	(2)
Vulnerability Severity Impact (VSI)	$VSI(IP) = \begin{cases} 1, & \text{if avgCVSS} < 5.0 \\ 3, & 5.0 \leq \text{avgCVSS} \leq 8.0 \\ 5, & \text{if avgCVSS} > 8.0 \end{cases}$	Section 3.3
Likelihood (L)	$L(IP) = A\,[EA(IP) + VSI(IP)]$	(2)
Raw Risk (RR)	$RR(IP) = \frac{E(IP) + L(IP)}{2}$ $RR(IP) = C\,E(IP) + (1 - C)L(IP)$	–
Overall Risk (R)	$R(IP) = \log_2[1 + RR(IP)]$	(4)

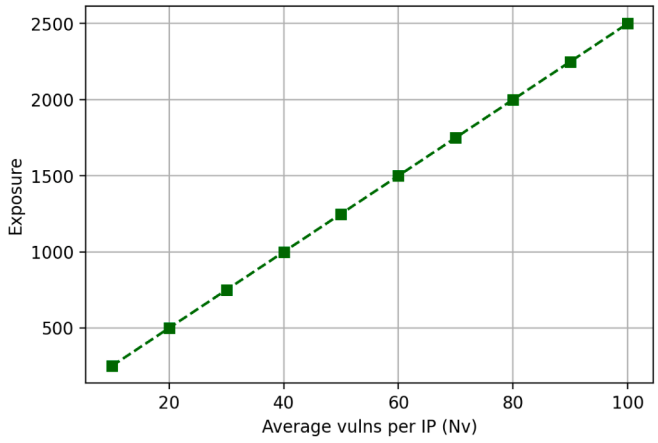

Fig. 8. Scenario A: Exposure vs. avg. N of vulnerabilities per IP.

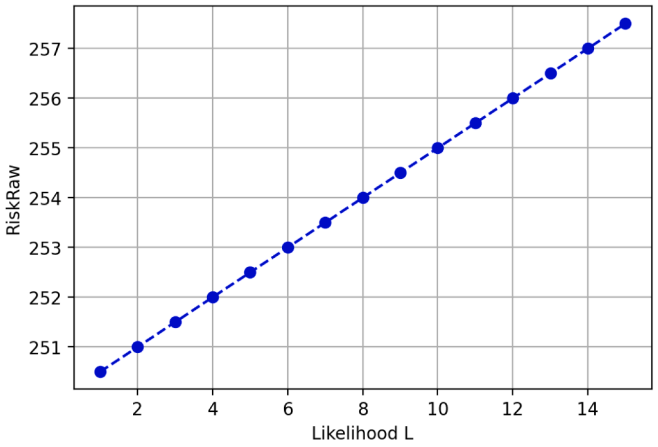

Fig. 10. Scenario C: Risk raw vs likelihood.

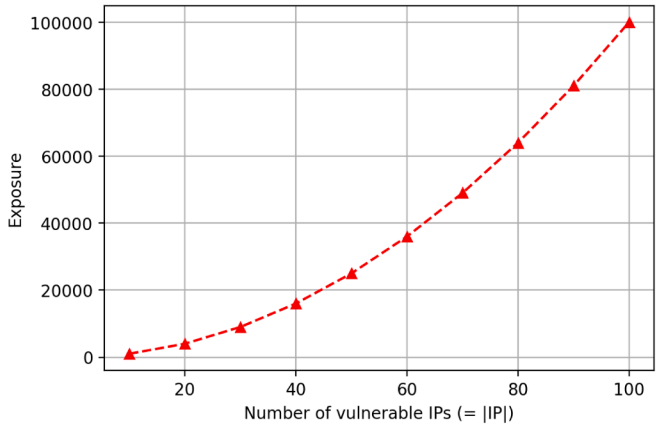

Fig. 9. Scenario B: exposure vs. TotalVulnIPs.

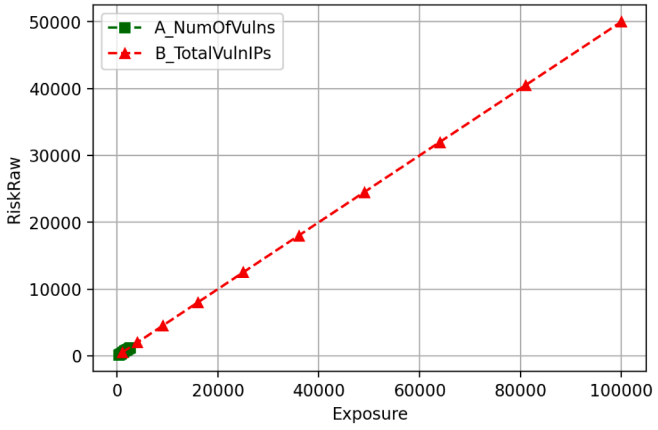

Fig. 11. Risk raw vs exposure.

how quickly overall cyber-risk balloons when a network is allowed to sprawl without hardening each additional node.

Finally, in Scenario C, we examined the variations in the likelihood (L). Here, we keep the technical attack surface constant, the exposure is fixed at 500, and we increase only the likelihood score from 1 to 15. Because the raw-risk formula is a weighted formula, every unit step in L adds exactly 0.5 to the total. The blue dashed line in Fig. 10, therefore, increases with perfect linearity, moving from about 250.5 when $L = 1$ to roughly 257.5 when $L = 15$. This confirms that, once the vulnerable footprint has been set, our metric remains proportionally sensitive to shifting threat likelihoods.

Lastly, we over-laid the results from the two experiments in which Exposure changes, Scenario A (varying the number of vulnerabilities per IP) and Scenario B (varying the count of vulnerable addresses), and plotted the corresponding Risk Raw (RR) values against the resulting Exposure scores. As shown in Fig. 11, the points from both scenarios fall on the same straight line: regardless of whether Exposure grows because each host becomes more vulnerable or because the vulnerable footprint becomes "wider" (more affected IPs), the raw risk rises in exact proportion to the Exposure term in the formula. Scenario C is omitted from this overlay because its Exposure is held constant while Likelihood is swept. The joint plot confirms that our metric responds to any of the technical drivers in a linear way.

5.3.2. *Sensitivity analysis*

To validate the robustness of our methodology and identify the weight of our metrics on the Overall Risk Score, we conducted a sensitivity analysis using the tornado diagram approach. With this analysis

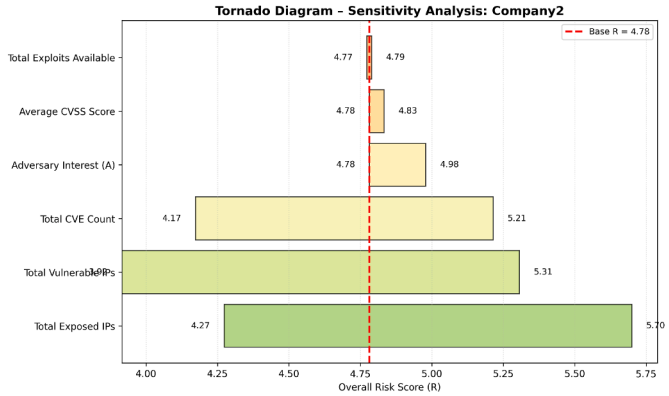

Fig. 12. Sensitivity analysis company 2.

we aim to demonstrate how REF responds to variations in input parameters, establishing confidence in the REF's reliability and discussing the key drivers of cybersecurity risk assessment outcomes.

The tornado diagram method evaluates the impact of each input parameter on the output by varying one parameter at a time while holding all others constant at their baseline values. For each parameter, we calculated the Overall Risk Score using lower bound (10th percentile), expected value (50th percentile), and upper bound (90th percentile) estimates. This one-at-a-time (OAT) approach allows us to isolate and quantify the individual contribution of each parameter to overall risk score variability, following established sensitivity analysis practices in risk engineering (Borgonovo and Plischke, 2016).

We selected three representative organizations from our dataset: Company 2 (high risk, R = 4.87), Company 3 (moderate risk, R = 4.07), and Company 1 (lowest risk, R = 3.13). These organizations represent different vulnerability profiles, through which we can assess whether sensitivity patterns remain consistent. For each organization, we varied the following six key REF input parameters from the observed baseline value: Total Exposed IPs: ±50%; Total Vulnerable IPs: ±50%; Total CVE Count: ±40%; Total Exploits Available: ±30%; Average CVSS Score: varied between 5.0 (medium severity threshold) and 9.0 (critical severity), reflecting the range observed in our dataset; Adversary Interest (A): varied between 1 (low targeting motivation) and 3 (high strategic value).

Fig. 12 shows the tornado diagrams for Company 2 (the other diagrams are on our online repository (Casaril and Galletta, 2025a), with parameters sorted by decreasing influence (widest bar at bottom, following standard tornado diagram approach). The horizontal bars illustrate the range of risk scores achievable by varying each parameter between its lower and upper bounds, while the vertical red dashed line indicates the baseline risk score calculated using observed values. The width of each bar directly represents that parameter's sensitivity: wider bars indicate parameters that, when varied, produce larger swings in the Overall Risk Score. The sensitivity analysis reveals consistent patterns across all three organizations, providing useful information about REF's behavior and validating our design decisions.

The sensitivity analysis highlights several consistent patterns across all organizations. The size of the external attack surface emerges as the most influential factor: variations in the number of exposed IPs cause the largest absolute changes in the Overall Risk Score, confirming that the span of an organization's externally visible infrastructure is the main driver of cyber risk exposure. Closely following, the concentration of vulnerable hosts is almost equally critical. Increases in the number of unpatched or misconfigured systems amplify risk. Parameters related to individual vulnerability characteristics, such as the total CVE count, average CVSS score, and the number of available exploits, have a moderate but consistent influence. These factors matter, yet their impact remains secondary to the structural properties of the attack surface itself. Adversary interest is modeled uniformly as A = 1 in our main analysis. Sensitivity tests show that variations of A between 1 and 3 affect the resulting risk score, demonstrating that refining the estimation of adversary intent, using sector-specific threat intelligence or geopolitical indicators, could improve the accuracy of REF, especially for organizations operating in contested domains or managing strategically valuable infrastructure.

This sensitivity analysis demonstrates that REF responds predictably and proportionally to input variations, with no evidence of unstable behavior. The most influential parameters, attack surface size and vulnerable host concentration, align with established cybersecurity principles, and the framework maintains consistent sensitivity patterns across organizations of different scales and risk profiles, supporting generalizability.

Finally, our analysis validates REF's design choices, particularly the emphasis on attack surface breadth and vulnerable host concentration, and demonstrates that the framework responds predictably to realistic parameter variations.

5.4. Threats to validity

REF offers a structured and reproducible way to assess cybersecurity risk in space organizations, but its validity is subject to several important considerations. These limitations do not invalidate our findings, but they do highlight areas where further refinement is needed.

First, our data collection is based on Shodan and on its filter `org:''...''` tag. This filter is extremely useful for attributing assets to specific entities, but it does not guarantee that the assets identified are part of the core or mission-critical infrastructure. Organizations may operate services under different registered names, third-party domains, or through partners and subsidiaries that may not be captured by this filter. Not all assets captured by the org filter may actually belong to the organization, especially in cases of dynamic IP assignment or misattributed metadata. This can create uncertainty in the completeness and precision of the dataset. However, although the tag may not accurately reflect the true boundaries of an organization, it still provides a consistent and reproducible view of exposed services. Allowing exposure levels to be compared across many entities in the space sector.

Shodan operates through periodic scans that provide only partial, time-bound snapshots of the public Internet. The exposure data is therefore ephemeral: an IP found vulnerable today may be patched tomorrow, or entirely removed from service. Longitudinal sampling over several months can mitigate volatility and reveal persistent misconfigurations or the regular emergence of the same weaknesses, which is a valuable signal of systemic risk.

Another limitation lies in the simplification used when assigning adversary interest. In the current implementation, the adversary interest factor is set uniformly to 1 across all organizations. This assumes an equal level of attractiveness to threat actors, which does not reflect real-world geopolitical or economic dynamics. Some organizations, such as those operating in defense, launch capabilities, or dual-use technologies, may be significantly more targeted than others, especially in certain countries.

Additionally, the framework does not assess the intent, capability, or behavior of specific threat actors. The presence of a known vulnerability does not imply imminent exploitation, and the existence of an exploit does not confirm attacker interest. Without telemetry data, incident logs, or threat intelligence integration, the REF approach remains at the level of potential risk rather than realized threats.

The methodology also presumes a consistent link between vulnerability exposure and organizational impact. A small number of exposed, yet highly sensitive, assets could pose a greater risk than hundreds of low-impact devices. REF does not incorporate contextual factors such as asset criticality, system isolation, or compensating controls, all of which play a significant role in actual risk realization. By highlighting the quantity and severity of internet-facing vulnerabilities, our methodology provides asset owners with essential data that they can easily

cross-reference with their own criticality inventories and compensating controls. REF is not designed to replace internal risk models; instead, it enhances them with externally verifiable evidence.

Finally, the vulnerability scoring itself relies on public CVSS data, which is a generalized and often conservative estimate. CVSS scores may be outdated, incomplete, or misaligned with the specific configurations used in operational systems. However, it remains the most widely adopted severity measure, keeping our results comparable with vulnerability management practices already established across the space industry.

In summary, while the REF is a scalable and repeatable approach to quantifying exposure and likelihood, its outputs should be interpreted in light of these limitations. Future work may improve the framework's validity by integrating more contextual signals, refining attribution methods, and incorporating live threat intelligence data.

6. Policy and security implications

REF is primarily a tool for assessing the cyber exposure of organizations based on a series of observable data, but it can also support the implementation of broader cybersecurity strategies and regulatory frameworks for critical infrastructure protection. Governments and institutions around the world are adopting more strict cybersecurity policies to address the evolving threat landscape in the space domain, and this requires new operational tools that can assist in translating these policies into measurable actions. In this respect, REF contributes by providing a method to assess external risk exposure, making it a building block within comprehensive, multi-layered cybersecurity policy implementation.

In this section, we discuss how REF can support some policy provisions of the EU's NIS 2 Directive, which establishes harmonized requirements for risk management, incident reporting, and supervision across critical sectors, including space, and of the USA Executive Order Strengthening and Promoting Innovation in the Nation's Cybersecurity.

6.1. REF for NIS2 cybersecurity requirements

Analyzing the NIS 2 Directive makes it clear how REF aligns with its emphasis on risk management and visibility over the digital infrastructure of essential and important entities. In particular, REF can assist in fulfilling the requirements outlined in Article 21 of the Directive, which obliges organizations to adopt technical and organizational cybersecurity risk-management measures. REF mechanisms help continuously monitor and evaluate the exposure of an organization's network-facing assets. This complements policy requirements for entities to maintain updated inventories of vulnerabilities, to assess the severity of known weaknesses, and to identify systems that are directly exposed to the Internet.

REF can also contribute to operational cybersecurity obligations under Articles 23 and 24, which mandate timely incident notification and coordinated vulnerability disclosure. Helping organizations identify exploitable vulnerabilities in real time, REF can support earlier detection of potential entry points for breaches or risky configurations that might otherwise go unnoticed until exploited. Moreover, REF can also support the implementation of sectoral cybersecurity policies and coordinated risk assessments as required by Article 13. Since REF is designed to scale across multiple organizations and rely on online available data, its results can be aggregated and compared to reveal trends across different entities or sectors. Supervisory authorities and national competent authorities could use REF-based insights to perform comparative exposure analysis or prioritize oversight and audit activities.

However, REF is not a comprehensive cybersecurity solution, and it does not replace the broader suite of measures required by NIS 2. For example, it does not account for internal controls, organizational governance, supply chain risks, staff training, physical security, or encryption standards. It cannot address strategic, legal, or contractual dimensions of risk, nor does it include mechanisms for managing incident response, crisis communication, or business continuity. In short, REF does not satisfy the full scope of Articles 21–24 on its own.

Instead, our methodology should be understood as a practical tool that complements broader compliance and risk management strategies, an effective early-stage mechanism for identifying and quantifying externally visible risks. It can be integrated into wider risk assessments, such as those required by Article 20, to ensure that organizations begin their cybersecurity work from a clear and empirical understanding of where they are most exposed. Alongside asset classification, threat modeling, and penetration testing, REF supports an organization's ability to comply with the NIS 2 Directive in a data-informed way.

6.2. REF supporting the european space act

The EU Commission released the EU Space Act proposal on 25 June 2025 (Commission, 2025). The Act aims to consolidate previous strategies into a unified legal framework for all space activities within the European Union.

The Act is built on three key pillars: safety, resilience, and sustainability. Safety is prioritized through improved space traffic management (STM), enhanced traceability of space objects, and a well-defined strategy for managing space debris. Resilience is achieved by implementing binding cybersecurity requirements to better protect European space assets and infrastructure. Sustainability is promoted by making environmental impact assessments mandatory for space missions and encouraging compliance and innovation in areas such as in-orbit servicing technologies and satellite end-of-life management.

The Act introduces a sector-specific cybersecurity framework tailored to space infrastructure, addressing gaps left by existing legislation such as the NIS2 Directive, which sets cybersecurity requirements for medium and large-sized public electronic communications networks and certain ground-segment operators, but it does not cover key areas of the space sector. This includes spacecraft, Earth Observation and Space Situational Awareness satellites, micro-operators, launch services from outside the European Union, and critically, Union-owned constellations.

The proposal calls for a dedicated risk management regime for space operators. Under Article 75, the Act identifies itself as *lex specialis* over NIS2: for any operator already classified as an *"essential"* or *"important"* entity, the space-specific requirements outlined in the Act take precedence over the more general provisions of Article 21 of NIS2. Overall, this means that operators will be required to conduct thorough risk assessments, implement security measures proportional to the criticality of their missions, and maintain continuity protocols across all operational phases (Articles 78–88). A new Union Space Resilience Network (EUSRN) will facilitate coordinated responses to major cyber incidents and further support harmonization across member states (Article 94) (Commission, 2025).

Article 78 obliges operators to identify, reassess, and treat vulnerabilities throughout the lifecycle of space missions, and it empowers the Commission to issue delegated acts that standardise threat-modelling methods and ensure comparability of risk assessments (Commission, 2025). Our REF framework can help operators to fulfill such obligations: indeed, it generates the evidence, such as delegated acts, that are needed: a time-stamped map of exposed services, the CVEs linked to each host, exploit maturity, and a composite Risk score.

Asset governance under Article 80 depends on inventories that are drawn up by individual space missions and include the origin and current physical location of assets, including the identification of a cloud-based service.

Again, our REF framework can support these activities through its discovery phase, which collects some of this data, pulling ASN, geolocation, and hosting-provider data.

Article 88 makes Threat-Led Penetration Testing mandatory before launch and every three years thereafter, with the scope to be chosen considering the risk assessment referred to in Article 78(2)

(Commission, 2025). Our REF framework can support organizations with the requirements of this article. Indeed, the ranked list of Internet-facing hosts produced by REF lets security teams transform that legal wording into a concrete test plan: start with the endpoints sitting in the highest-risk decile, verify whether the theoretical exploit paths can be demonstrated in practice, and document the results against REF's baseline so improvements show up as a measurable downward risk.

Article 92 introduces a formal supply-chain risk-management framework and instructs operators to inventory "critical assets of non-Union origin" and collect their dependence on those suppliers (Commission, 2025). Our REF framework can narrow that task by flagging which external libraries, SaaS consoles, or third-party services account for the largest share of exploitable CVEs in the mission, helping procurement teams channel their due diligence and their contract clauses towards the vendors with less exposure. These provisions, even if just part of a proposal, point toward a continuous, evidence-based risk quantification from best practice to legal requirement. REF can deliver and support some of these measurements that Chapter II calls for.

6.3. Applying REF to the cybersecurity oversight of U.S. space infrastructure

By examining other recent policies where REF can contribute, we identified an interesting initiative in the U.S. that represents one of the first policies explicitly demanding and designing some cybersecurity requirements for the space sector. In January 2025, before the end of his mandate, the 46th President of the United States issued the "Executive Order on Strengthening and Promoting Innovation in the Nation's Cybersecurity". The Order aims to strengthen the cybersecurity posture of the United States across multiple critical infrastructure sectors, including space. This document follows up on prior cybersecurity strategies and executive orders, expanding regulatory oversight, strengthening security standards, and emphasizing the role of government agencies and private sector stakeholders in reducing cyber threats.

The U.S. government recently responded to the increasing cyber threats to national security by implementing stricter cybersecurity frameworks, including Executive Order 14,028 (Executive Office of the President, 2025) (Improving the Nation's Cybersecurity) and the National Cybersecurity Strategy. However, "*the rapid expansion of digital infrastructure, the emergence of new attack vectors, and the growing importance of space-based assets* (Executive Office of the President, 2025)" needed further policy advancements, leading to this 2025 Executive Order.

A section of the Executive Order focuses specifically on the space sector, acknowledging the rising cybersecurity risks associated with satellite networks, space-ground communication systems, and federal space programs. The document explicitly states that "as cybersecurity threats to space systems increase, these systems and their supporting digital infrastructure must be designed to adapt to evolving cybersecurity threats and operate in contested environments", underscoring the need for continuous assessment, rigorous vulnerability management, and secure software development in all aspects of space operations.

Among the primary requirements set forth in the Executive Order for the space sector, the following stand out:

1. *Mandatory Cybersecurity Assessments*: The order requires federal space agencies and their contractors to perform continuous cybersecurity assessments of their space-based infrastructure.
2. *Enhanced Procurement Security*: The directive compels NASA, the Department of Commerce, and the Department of the Interior to review and update cybersecurity requirements within the Federal Acquisition Regulation (FAR) to ensure that space-related contracts enforce stricter cybersecurity controls. This measure aims to prevent vendors with inadequate cybersecurity postures from supplying critical space technologies.
3. *Protection of Satellite Communications and Space-Ground Links*: A major aspect of the Order is the security of satellite communications. Indeed, it requires "all new civil space systems to incorporate risk-based, tiered cybersecurity requirements to protect command and control functions, detect and mitigate anomalous activity, and ensure secure software and hardware development practices."
4. *Creation of a Space Cybersecurity Oversight Framework*: The National Cyber Director is tasked with submitting a comprehensive study on space ground systems managed by Federal Civilian Executive Branch (FCEB) agencies, including an inventory of systems and recommendations to improve cyber defenses. This initiative seeks to enhance oversight, provide cybersecurity benchmarks for space agencies, and promote cross-agency cooperation in cybersecurity efforts.
5. *Strengthening Secure Software Development in Space Systems*: Given the increasing reliance on commercial and open-source software in space operations, the Executive Order mandates that software providers follow strict cybersecurity guidelines and demonstrate compliance with secure software development standards. The order states that "software providers must also address how software is delivered and the security of the software itself, with the Federal Government identifying a coordinated set of practical and effective security practices for space system software development."

Our REF methodology closely aligns with several cybersecurity concerns outlined in the Executive Order. Specifically, it addresses the need for continuous assessments, testing, and exercises to verify the cybersecurity capabilities of federal space systems and their supporting infrastructure. REF can be adopted to evaluate the external attack surface of space-related assets, particularly focusing on ground segment infrastructure and Internet-connected components. REF can therefore aid in fulfilling the continuous assessment mandate by identifying externally exposed and potentially vulnerable services across space system operators.

Additionally, the Executive Order mandates that agencies review and update cybersecurity requirements in civil space contracts, using a risk-based, tiered approach (Section 3(e)(i)) (Executive Office of the President, 2025). REF could be part of this process as a monitoring mechanism to support the implementation and enforcement of such requirements. For instance, REF-derived metrics could be used to validate that Internet-facing components of space systems comply with security baselines over time. It can therefore be used as part of a verification layer in contract compliance monitoring. REF's methodology could inform risk tier assignments as it is able to identify systems with a high number of exploitable services or critical CVEs.

However, REF should not be considered as a standalone solution for fulfilling policy requirements. Instead, it offers a data-driven monitoring component that can be integrated into a wider risk assessment and cybersecurity management process. In particular, it can be the first layer for detecting exposure and identifying organizations or systems that may need deeper inspection and additional safeguards. Its applicability to ground components of space infrastructure makes it a practical addition to the cyber resilience toolkit outlined in the policies described in this section and in others that will come, such as the European Space Act (Casaril and Galletta, 2025b).

7. Conclusion and future research

Our framework stems from the need to quantify cyber risk in the space sector (**RQ1**) and to do so in a clear and repeatable way (**RQ2**) from observable data on the external attack surface of space sector organizations (**RQ3**). With the Risk Exposure Framework, we built an indicator that combines Exposure (size of the exposed infrastructure) and Likelihood (probability that vulnerabilities are actually exploited), thus mapping to the requirements of major cybersecurity policies (NIS 2 in the EU, Executive Order 14,144 in the US) (**RQ4**).

Addressing **RQ1**, this paper measured cyber risk through the Risk Exposure Framework (REF), which combines data on exposed Internet assets and known vulnerabilities into a quantifiable indicator. Our approach responds to the need for a measurable link between high-level policy requirements and concrete technical assessments. Translating exposure and vulnerability data into a single metric, REF can allow organisations to monitor the effectiveness of cybersecurity measures and to support both technical and governance decisions.

RQ2 focused on repetability and clarity: REF proved both clear and repeatable, as the same data collection and computation process was applied across several space-sector organisations, allowing for direct comparison of their risk levels. The resulting distributions showed consistent patterns of exposure and exploitability, confirming that the method can be reused for longitudinal studies. The reproducibility of the metric and the transparency of its components address the need for a standardised way to track cyber risk across organisations.

Through **RQ3**, we aimed to investigate the level of exposure of the space sector. The results showed that space organisations remain significantly exposed to risks originating from Internet-connected systems. Out of roughly 100,000 identified assets, about 4% contained known vulnerabilities, with risk levels varying widely depending on how those vulnerabilities were concentrated. As previous attacks have demonstrated, the most critical exposure originates from the ground and user segments, confirming their likely role as entry points in space-related incidents.

RQ4 inquired about the possibility of connecting security measures to policy requirements. We demonstrated how REF can be directly related to European cybersecurity policies, as REF's indicators correspond to the continuous risk assessment and vulnerability management obligations defined under NIS2, and to the monitoring and oversight provisions foreseen in the proposed EU Space Act. Our framework, therefore, provides a concrete technical foundation to support regulatory compliance and sector-level supervision, and can easily link empirical evidence to policy objectives.

From our experimental results, we learned several key lessons. First, the high number of unpatched CVEs highlights how velocity and consistency in patching outrank perimeter breadth; large operators accumulated extreme Exposure scores, probably due to update cycles lagging behind service deployment, whereas limited infrastructures moved into high-risk territory only when critical patches remained outstanding. Second, data shows how the weakest technical layer is still the basic web and cryptographic stacks: Apache HTTP Server, Nginx, OpenSSL, jQuery, demonstrating that dependency management rather than exotic satellite-specific code is the leading attack vector. Third, small does not equate to safe: entities with a few dozen public hosts nevertheless posted high Overall Risk scores, because exploit availability magnified each remaining flaw.

REF should be understood as a complementary tool within a comprehensive cybersecurity assessment strategy for space, rather than a standalone solution. While REF excels at identifying and quantifying ground-segment exposure, where most cyberattacks originate, it does not replace other essential security assessment methodologies such as onboard system audits, hardware security testing, or mission-critical system penetration testing. Organizations implementing REF should integrate its findings with broader security assessments that cover space-segment vulnerabilities, insider threat analysis, and mission-specific risk factors that cannot be captured through external network reconnaissance.

The framework's ground-segment focus aligns with current threat landscapes and policy requirements; however, space operators should ensure that REF-based assessments are part of a comprehensive security program that addresses all mission segments and operational phases.

7.1. Future research directions

Future research should build upon the methodology proposed in this study. One option involves integrating additional tools for identifying and characterizing ICDs beyond Shodan. Platforms such as Censys, ZoomEye, or GreyNoise could be combined to provide broader visibility into exposed space infrastructure and validate findings across different data sources. Additionally, vulnerability data could be enriched with threat intelligence feeds, including indicators of compromise (IOCs), active exploit tracking, or ransomware group targeting patterns, which would support the estimation of likelihood.

Asset criticality is one measure that is not completely addressed in REF. This would enhance the risk model by weighting exposures not only by technical severity but also by their potential mission impact, helping prioritize remediation across infrastructures. Another underexplored domain is the temporal analysis of exposure trends, which would assess how organizational risk evolves over time and how effectively vulnerabilities are being mitigated. While offering insights into the cybersecurity exposure of space sector organizations, our approach has several important limitations. First, it relies solely on publicly accessible information, meaning that only internet-facing systems are captured. Internal infrastructure, assets protected by firewalls, VPNs, or hidden through filtering techniques, are excluded from analysis. As a result, the risk scores produced here reflect only a subset of the true attack surface, what is externally visible, and may underestimate risks tied to internal system vulnerabilities or insider threats. Future work should therefore widen the observation perimeter by fusing external scan results with data drawn from internal vulnerability scanners, EDR agents and industrial inventories, so that the model also captures risks tied to non-internet-facing assets. Second, the framework's dependence on CVE identifiers, while useful for standardization and comparison, overlooks operational context. The methodology does not consider the business or mission criticality of the exposed systems, nor whether compensating controls, segmentation, or isolation mechanisms are in place. Two hosts with the same vulnerability may pose vastly different risks depending on how they are deployed and secured. Similarly, this approach cannot account for custom or proprietary configurations that might either mitigate or amplify the vulnerability impact. A second research track should therefore aim to link business or mission-criticality to every host by linking CVEs to impact scores derived from business-impact assessments.

A further limitation lies in how we infer exploitability. The presence of an exploit in public databases such as Vulners is a necessary but insufficient condition for actual threat activity. The framework does not detect ongoing attacks, zero-day vulnerabilities, and it assumes that all exploits are equally likely to be leveraged, which may not reflect the strategies of sophisticated threat actors. Exploitability estimates can be refined by combining real-time threat-intelligence feeds, measurements of exploit weaponisation in the wild and machine-learning models that rank the effort required to abuse each vulnerability.

In addition, the reliability of service-level data, especially from Shodan, requires careful consideration. Shodan may label a host as "vulnerable" even if no specific open port is visibly tied to a CVE, due to how its scanning engine interprets software banners and metadata. For example, a host running an outdated operating system may be marked as vulnerable even when only non-vulnerable ports are open. Shodan's scans are periodic and port-limited: not all ports are checked at every scan cycle, and not all responses yield full banner data. As a result, an IP might be listed with vulnerabilities but show no active services, or vice versa. In this regard, the reliance on banner-based services such as Shodan can be mitigated by cross-checking multiple scanners or similar tools.

CRediT authorship contribution statement

Francesco Casaril: Writing – original draft, Validation, Methodology, Investigation, Formal analysis, Conceptualization; **Letterio Galletta:** Writing – review & editing, Validation, Supervision, Methodology.

Data availability

Data will be made available on request.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We thank the anonymous reviewers for their careful and helpful comments that helped us improve the paper. This work was supported by Project SEcurity and RIghts in the CyberSpace (SERICS) PE0000014, funded by the European Union NextGenerationEU under the National Recovery and Resilience Plan M4C2 I1.3., CUP: D67G22000340001.

References

Administration, U.S.M., 2017. 2017-005a black sea GPS interference advisory. Accessed 2025-09-27. https://www.maritime.dot.gov/msci/2017-005a-black-sea-gps-interference.

Airbus, 2025. Airbus reports full-year (FY) 2024 results. Accessed: 2025-03-30. https://www.airbus.com/en/newsroom/press-releases/2025-02-airbus-reports-full-year-fy-2024-results.

Alsmadi, I., Dwekat, Z., Cantu, R., Al-Ahmad, B., 2022. Vulnerability assessment of industrial systems using shodan. Cluster Comput. 25 (3), 1563–1573.

Booth, H., Rike, D., Witte, G.A., 2013. The national vulnerability database (NVD): overview. https://nvd.nist.gov/.

Borgonovo, E., Plischke, E., 2016. Sensitivity analysis: a review of recent advances. Eur. J. Oper. Res. 248 (3), 869–887.

Boschetti, N., Gordon, N.G., Falco, G., 2022. Space cybersecurity lessons learned from the viasat cyberattack. In: ASCEND 2022, p. 4380.

Casaril, F., Galletta, L., 2024. Securing satcom user segment: a study on cybersecurity challenges in view of IRIS2. Comput. Secur. 140, 103799.

Casaril, F., Galletta, L., 2025a. Assessing the attack surface of space organizations: a data-driven analysis (Supplementary material). https://github.com/Fracas20/REF-Risk-Exposure-Framework.

Casaril, F., Galletta, L., 2025b. Space cybersecurity governance: assessing policies and frameworks in view of the future european space legislation. J. Cybersecur. 11 (1), tyaf013.

CISA, 2022. Resilient positioning, navigation, and timing (PNT) conformance framework, version 2.0. Technical Report. Department of Homeland Security (DHS) and CISA. https://www.dhs.gov/sites/default/files/2022-05/22_0531_st_resilient_pnt_conformance_framework_v2.0.pdf.

Commission, E., 2025. Proposal for a regulation of the European parliament and of the council on the safety, Resilience and sustainability of space activities in the union. COM(2025) 335 final; 2025/0335 (COD). Accessed 12 July 2025. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025PC0335.

Corporation, T.A., 2025. The aerospace corporation. Accessed: March 13, 2025. https://aerospace.org/.

Cyber, I., 2025. Lab dookhtegan cyberattack on iranian oil tankers traced to supply chain compromise of Fanava's infrastructure. Accessed 2025-09-27. https://industrialcyber.co/supply-chain-security/lab-dookhtegan-cyberattack-on-iranian-oil-tankers-traced-to-supply-chain-compromise-of-fanavas-infrastructure/.

Cyberinflight, 2022. Threat intelligence report - space sector. Technical Report. Cyberinflight. Accessed: 2025-03-30. https://www.cyberinflight.com/.

Cybersecurity and (CISA), I. S.A., 2024. Recommendations to space system operators for improving cybersecurity. https://www.cisa.gov/sites/default/files/2024-06/Recommendations%20to%20Space%20System%20Operators%20for%20Improving%20Cybersecurity%20(508).pdf.

Cydome, 2025. Second wave of cyberattacks on iranian-linked vessels: technical analysis. Accessed 2025-09-27. https://www.cydome.io/blog/second-wave-iranian-vessels-cyberattack.

Daskevics, A., Nikiforova, A., 2021. ShobeVODSDT: shodan and binary edge based vulnerable open data sources detection tool or what internet of things search engines know about you. In: 2021 Second International Conference on Intelligent Data Science Technologies and Applications (IDSTA). IEEE, pp. 38–45.

Durumeric, Z., Adrian, D., Mirian, A., Bailey, M., Halderman, J.A., 2015. A search engine backed by internet-wide scanning. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, pp. 542–553.

EASA, 2024. Safety information bulletin: increased probability of GNSS problems near conflict zones. Accessed 2025-09-27. https://www.easa.europa.eu/.

ENISA, E. U. A. f.C., 2016. PETs Controls Matrix - A Systematic Approach for Assessing Online and Mobile Privacy Tools. Technical Report. European Union Agency for Cybersecurity. Heraklion, Greece. Publication Date: December 20, 2016. https://www.enisa.europa.eu/publications/pets-controls-matrix/pets-controls-matrix-a-systematic-approach-for-assessing-online-and-mobile-privacy-tools.

ENISA, E. U. A. f.C., 2023. 5G Security Controls Matrix. Technical Report. European Union Agency for Cybersecurity. Publication date: May 24, 2023. https://www.enisa.europa.eu/publications/5g-security-controls-matrix.

European commission, 2023. Eu space strategy for security and defence – white paper. Accessed: 2025-04-23. https://defence-industry-space.ec.europa.eu/eu-space/eu-space-strategy-security-and-defence_en.

European external action service, 2022. Strategic compass for security and defence. Accessed: 2025-04-23. https://www.eeas.europa.eu/eeas/strategic-compass-security-and-defence-1_en.

European repository of cyber incidents (EuRepoC), 2025. Cyber incident data framework (2000–2025). https://eurepoc.eu/. Accessed: 2025-03-30.

European union agency for cybersecurity (ENISA), 2024a. ENISA threat landscape 2024. Technical Report. ENISA. Accessed: 2025-03-30. https://www.enisa.europa.eu/publications/enisa-threat-landscape-2024.

European union agency for cybersecurity (ENISA), 2024b. Space threat landscape. Technical Report. ENISA. Accessed: 2025-03-30. https://www.enisa.europa.eu/publications/space-threat-landscape.

European union agency for the space programme, 2025. Euspa industry directory. https://www.euspa.europa.eu/opportunities/fundamental-elements/industry-directory. Accessed: 2025-03-13.

European union agency for the space programme (EUSPA), 2025. Galileo to be the first GNSS to offer authentication service worldwide with launch of OSNMA. https://www.euspa.europa.eu/pressroom/press-releases/galileo-be-first-gnss-offer-authentication-service-worldwide-launch-osnma.

European union aviation safety agency (EASA), 2024. Easa updates safety information bulletin on gnss jamming/spoofing. https://www.euspa.europa.eu/en/newsroom-and-events/news/easa-updates-safety-information-bulletin-global-navigation-satellite.

Executive Office of the President, 2025. Executive order 14144 of January 16, 2025: strengthening and promoting innovation in the nation's cybersecurity. Federal Register, 90 FR 6755–6771. Signed by President Joseph R. Biden Jr.; published January 17, 2025. https://www.federalregister.gov/documents/2025/01/17/2025-01470/strengthening-and-promoting-innovation-in-the-nations-cybersecurity.

Geetha, A., Keerthika, V., Raj, D.M.D., 2024. Cybersecurity kill chain in outer space and cyberspace security. In: Cyber Space and Outer Space Security. River Publishers, pp. 81–95.

Genge, B., Enăchescu, C., 2016. ShoVAT: shodan-based vulnerability assessment tool for internet-facing services. Secur. Commun. Netw. 9 (15), 2696–2714.

George, A.A.D., et al., 2018. Onboard processing with hybrid and reconfigurable computing on small satellites. In: NSF SHREC. https://www.nsf-shrec.org/...

Gilat satellite networks, 2024. Gilat presents fourth quarter and full year 2023 results. Accessed: 2025-03-30. https://www.gilat.com/pressreleases/gilat-presents-fourth-quarter-and-full-year-2023-results/#:~:text=Adi%20Sfadia%2C%20Gilat's%20CEO%2C%20commented,up%2044%25%20over%20last%20year

GMV, 2024. Annual report 2023. Technical Report. GMV. Accessed: 2025-03-30. https://www.gmv.com/sites/default/files/content/file/2024/07/31/114/annual_report_2023_3.pdf.

Goodwill, J., 2024. Current technology in space: briefing on rad-hard compute constraints. Technical Report. NASA. https://ntrs.nasa.gov/api/citations/20240001139/downloads/Current%20Technology%20in%20Space%20v4%20Briefing.pdf/.

Gourisetti, S. N.G., Touhiduzzaman, M., Ashley, T.D., Pal, S., McKenzie, P.L., 2021. Cybersecurity Risk Assessment Framework for Externally Exposed Energy Delivery Systems. Technical Report. Pacific Northwest National Lab.(PNNL), Richland, WA (United States).

Harry, C., Sivan-Sevilla, I., McDermott, M., 2025. Measuring the size and severity of the integrated cyber attack surface across US county governments. J. Cybersecur. 11 (1), tyae032.

Howland, H., 2022. CVSS: ubiquitous and broken. Digital Threats 4 (1). https://doi.org/10.1145/3491263

Hutchins, R., 2016. Cyber defense of space assets. Tufts School of Engineering Medford, MA, USA, 1–18.

Initiative, J. T. F.T., 2012. Guide for Conducting Risk Assessments (NIST Special Publication 800-30 Revision 1). Technical Report NIST SP 800-30 Rev. 1. National Institute of Standards and Technology. Gaithersburg, MD. https://doi.org/10.6028/NIST.SP.800-30r1

Innovation &. Technology, U.K. D. f.S., 2025. Cyber risks of cloud computing in the ground segment of the space sector. Technical Report. UK Department for Science, Innovation and Technology (DSIT). Expert insights on risks in Ground Segment as a Service (GSaaS). https://www.gov.uk/government/publications/cyber-risks-of-cloud-computing-in-the-ground-segment-of-the-space-sector.

Institute, C., 2022. Case study: viasat attack. Accessed 2025-09-27. https://cyberconflicts.cyberpeaceinstitute.org/law-and-policy/cases/viasat.

International, I., 2025. Hackers disrupt communications of iranian tankers and cargo ships. Accessed 2025-09-27. https://www.iranintl.com/en/202509xxxx/hackers-disrupt-communications-iranian-tankers-cargo-ships.

Jones, J.A., 2006. An introduction to factor analysis of information risk (FAIR). Norwich Univ. J. Inf. Assur. (NUJIA) 2 (1) 1–12.

Jouini, M., Rabai, L. B.A., 2016. Comparative study of information security risk assessment models for cloud computing systems. Procedia Comput. Sci. 83, 1084–1089.

Khan, S.K., Shiwakoti, N., Diro, A., Molla, A., Gondal, I., Warren, M., 2024. Space cybersecurity challenges, mitigation techniques, anticipated readiness, and future directions. Int. J. Crit. Infrastruct. Prot. 47, 100724.

Li, R., Shen, M., Yu, H., Li, C., Duan, P., Zhu, L., 2020. A survey on cyberspace search engines. In: Cyber Security: 17th China Annual Conference, CNCERT 2020, Beijing, China, August 12, 2020, Revised Selected Papers 17. Springer Singapore, pp. 206–214.

Lightman, S., 2022. Applying the NIST CSF to the Satellite Ground Segment. Technical Report. NIST. https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=935449.

Lightman, S., et al., 2022. Satellite Ground Segment: Applying the NIST Cybersecurity Framework. Technical Report NIST IR 8401. National Institute of Standards and Technology. https://csrc.nist.gov/pubs/ir/8401/final.

Liu, Y., Sarabi, A., Zhang, J., Naghizadeh, P., Karir, M., Bailey, M., Liu, M., 2015. Cloudy with a chance of breach: forecasting cyber security incidents. In: 24th USENIX Security Symposium (USENIX Security 15), pp. 1009–1024.

Liu, Y., Yuan, Y., Zhu, Y., Hu, L., Wang, L., 2025. Research on design and implementation of an intelligent network asset search system based on LLM agent and FOFA. In: Proceedings of the 2025 4th International Conference on Cyber Security, Artificial Intelligence and the Digital Economy, pp. 483–488.

Lu, Y., Da Xu, L., 2018. Internet of things (IoT) cybersecurity research: a review of current research topics. IEEE Internet Things J. 6 (2), 2103–2115.

Manulis, M., Bridges, C.P., Harrison, R., Sekar, V., Davis, A., 2021. Cyber security in new space: analysis of threats, key enabling technologies and challenges. Int. J. Inf. Secur. 20, 287–311. https://doi.org/10.1007/s10207-020-00503-w

Matherly, J., 2015. Complete guide to shodan. Shodan, LLC (2016-02-25) 1.

NASA, 2024. Small spacecraft technology state of the art 2024. Technical Report. NASA. https://www.nasa.gov/smallsat-institute/sst-soa/.

Netlas, 2025. Netlas.IO – search engine for internet infrastructure. Accessed: 2025-07-08. https://netlas.io/.

Olszewski, B., 2018. Advanced persistent threats as a manifestation of states' military activity in cyber space. Sci. J. Mil. Univ. Land Forces 50, 57–71.

Park, C.Y., et al., 2025. Secure and lightweight firmware over-the-air update for constrained devices. Electronics. 8, 1583.

Pavur, J., Martinovic, I., 2022. Building a launchpad for satellite cyber-security research: lessons from 60 years of spaceflight. J. Cybersecur. 8 (1), tyac008.

Poirier, C., 2023. Breaking the Final Frontier: Cyber Risks and Threats to Space Systems. Report. Center for Security Studies (CSS), ETH Zurich. https://ethz.ch/content/dam/ethz/special-interest/gess/cis/center-for-securities-studies/pdfs/breaking-the-final-frontier-cyberdefense-report.pdf.

Pražák, J., 2021. Space cyber threats and need for enhanced resilience of space assets. In: European Conference on Cyber Warfare and Security. Academic Conferences International Limited, pp. 542–XIV.

Psiaki, M.L., Humphreys, T.E., 2016. GNSS spoofing and detection. Proc. IEEE 104 (6), 1258–1270. https://doi.org/10.1109/JPROC.2016.2526658

Reuters, 2022. Satellite outage knocks out control of enercon wind turbines. Reuters Accessed 2025-09-27. https://www.reuters.com/business/energy/satellite-outage-knocks-out-control-enercon-wind-turbines-2022-02-28/

Reuters, 2025. Two tankers collide near UAE's Khor Fakkan; fire reported. Reuters Accessed 2025-09-27. https://www.reuters.com/.

Sacra, 2024. SpaceX: company financials and overview. Accessed: 2025-03-30. https://sacra.com/c/spacex/#::text=The%20company%20achieved%20profitability%20in,total%20revenue%20of%20%2414.2B

Schaller, C., 2025. Sabotage of submarine cables and pipelines as a use of force and armed attack. Int. Law Stud. 106 (1), 8.

SentinelLabs, 2022. AcidRain: a modem wiper rains down on Europe. Accessed 2025-09-27. https://www.sentinelone.com/labs/acidrain-a-modem-wiper-rains-down-on-europe/.

Sharmin, A., Mahmud, B.U., Nabi, N., Shaima, M., Faruk, M. J.H., 2025. Cyber attacks on space information networks: vulnerabilities, threats, and countermeasures for satellite security. J. Cybersecur. Privacy 5 (3), 76.

Singh, U.K., Joshi, C., 2018. Comparative study of information security risk assessment frameworks. Int. J. Comput. Appl. 2 (8), 2250–1797.

Telespazio, 2024. Company profile. Accessed: 2025-03-30. https://www.telespazio.com/en/company/profile.

Times, F., 2025. GPS interference raises risk of accidents in the strait of hormuz. Financ. Times Accessed 2025-09-27. https://www.ft.com/content/ac3c571f-2c4e-4c57-b582-21e2b27170f8.

Tundis, A., Mazurczyk, W., Mühlhäuser, M., 2018. A review of network vulnerabilities scanning tools: types, capabilities and functioning. In: Proceedings of the 13th International Conference on Availability, Reliability and Security, pp. 1–10.

U.S. department of transportation, 2024. Complementary PNT action plan. https://www.transportation.gov/pnt/complementary-pnt-action-plan.

Valea, O., Oprişa, C., 2020. Towards pentesting automation using the metasploit framework. In: 2020 IEEE 16th International Conference on Intelligent Computer Communication and Processing (ICCP). IEEE, pp. 171–178.

Varadharajan, V., et al., 2024. Security challenges when space merges with cyberspace. Comput. Law Secur. Rev. 67, 101600.

Vessels, L., Heffner, K., Johnson, D., 2019. Cybersecurity risk assessment for space systems. In: 2019 IEEE Space Computing Conference (SCC). IEEE, pp. 11–19.

Viasat, 2022. Ka-sat network cyber attack overview. Accessed 2025-09-27. https://www.viasat.com/perspectives/corporate/2022/ka-sat-network-cyber-attack-overview/.

Viasat Inc., 2024. Viasat annual report FY24. Technical Report. Viasat, Inc. Accessed: 2025-04-06. https://investors.viasat.com/static-files/5e5c4a86-4602-46ad-b17f-e2e2b8158d79.

Williams, R., McMahon, E., Samtani, S., Patton, M., Chen, H., 2017. Identifying vulnerabilities of consumer internet of things (IoT) devices: a scalable approach. In: 2017 IEEE International Conference on Intelligence and Security Informatics (ISI). IEEE, pp. 179–181.

Young, M., Johnson, K., Swope, C., 2025. Space Threat Assessment 2025. Technical Report. Center for Strategic and International Studies (CSIS). Accessed: 5 June 2025. https://csis-website-prod.s3.amazonaws.com/s3fs-public/2025-04/250425_Swope_Space_Threat.pdf?VersionId=orhySgjISemJLjhdQKKes2OVb35jwkU5

Yu, L., Hao, J., Ma, J., Sun, Y., Zhao, Y., Luo, B., 2024. A comprehensive analysis of security vulnerabilities and attacks in satellite modems. In: Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security (CCS '24). ACM, Salt Lake City, UT, USA, p. 20. https://doi.org/10.1145/3658644.3670390