Beyond Obscurity: Developing CubeSat Cybersecurity

Joshua Davis[1], Jill Slay[1,2] and Nalin Arachchilage[3]
[1]Adelaide University, Australia
[2]SmartSat Cooperative Research Centre, Australia
[3]RMIT University, Australia

j.davis@adelaide.edu.au
jill.slay@adelaide.edu.au
nalin.arachchilage@rmit.edu.au

Abstract: As modern space systems, particularly CubeSat constellations, are increasingly embedded in defence, industrial, and civilian operations, fulfilling an increasingly vital role in telecommunications, earth observation, and national security, their vulnerabilities to cyberattacks continue to grow. Despite the advantages of CubeSats; cost-effectiveness, scalability, and enhanced coverage; as satellite providers attempt to maximise payload capabilities, these systems are fundamentally constrained by Size, Weight, and Power (SWaP), imposing significant design trade-offs that limit onboard security. This has contributed to emerging U.S. national security requirements for commercial satellites supporting defence missions to implement real-time on-board intrusion detection and prevention systems forcing cybersecurity for satellite vendors to no longer merely be a desirable design feature, but an explicit mission-assurance and acquisition requirement for space systems used to support national security. This paper examines the growing cybersecurity risks faced by CubeSat constellations, focusing primarily on multi-stage attacks—complex sequences of interrelated behaviours that may appear individually benign but are collectively malicious. Where existing terrestrial intrusion detection systems have succeeded, they lack the distinctive dynamics of satellite networks, and publicly available satellite-oriented datasets to achieve the same success in the space domain. This necessitates the development of bespoke, packet-level datasets for effective cybersecurity modelling. To address the immediate need for enhanced resilience while bespoke datasets are developed, Through the review of existing cybersecurity measures and the integration of machine learning-based solutions, we propose a preliminary cybersecurity framework based on self-supervised learning. This approach uses contrastive learning to produce a nominal operational profile of the satellite's systems. This profile serves as a baseline for an onboard anomaly detection system, aiming to identify deviations that could indicate a cyberattack. By developing adaptive intrusion detection systems (IDS) that prioritise resource conservation, we aim to enhance CubeSat security in future space missions, ensuring both operational integrity and data confidentiality. Furthermore, we propose the development of a minimum cyber resilience capability checklist, accompanied by a standardised testbench. This testbench would provide analogous testing to vacuum and vibration tests common in satellite engineering, covering core, low-cost, high-reward attack vectors such as authentication and encryption vulnerabilities, fuzzing, replay attacks, and Denial of Service (DoS). The results would be delivered to the satellite owner or producer in a simple pass/fail format, identifying which tests failed and recommending standard methods to mitigate the identified vulnerabilities. This framework aims to establish a foundational and verifiable level of cyber resilience for CubeSat missions.

Keywords: Space system cybersecurity, Satellite cybersecurity, On-board intrusion detection systems, Self-supervised learning, CubeSat constellations

1. Introduction

The rapid expansion of the space sector, often termed "New Space", has seen a significant shift from large governmental projects to the small satellite domain, where industry and academia exist as principal drivers. This proliferation is largely driven by CubeSats: compact satellites, typically nanosatellites, comprising modular 10cm³ (1U) units leveraging commercial off-the-shelf (COTS) components, such as Field-Programmable Gate Arrays (FPGAs) and Software-Defined Radios (SDRs). These small, comparatively cost-effective satellites are often deployed in Low Earth Orbit (LEO) constellations, allowing for global coverage and seamless data collection across vast geographical areas (Sweeting, 2018). CubeSats are affordable for academia, offering high sensing resolutions and low latencies, and when designed to be homogeneous; sharing common hardware, software, and mission objectives, constellations, provide industry with greater geographical coverage and redundancy at a cost comparable to larger satellites. This evolving space ecosystem enables CubeSats to be used in a variety of applications critical to national infrastructure, including positioning, navigation, and timing (PNT), for GPS and financial tracking, earth observation for agricultural and climate monitoring, communications for global internet access, and scientific research (Manulis et al., 2021). From the critical services outlined, this makes satellite systems an integral part of critical infrastructure even if only as an enabler.

Despite their growing prominence, CubeSats are inherently vulnerable to cyberattacks that could compromise their functionality, leading to significant risks to the integrity of space missions and national security. The surge

in satellite deployments has inadvertently increased the attack surface of the space domain, making it an increasingly attractive target for cyber adversaries. Of particular concern is the rise of sophisticated multi-stage cyberattacks. These attacks involve an ensemble of attack steps, often chaining multiple single-stage attacks into complex sequences of interrelated behaviours. These steps may appear individually benign and are difficult to detect in resource-rich terrestrial environments, let alone in the resource-constrained, physically inaccessible, and radiation-battered environment of space. With limited power budgets where critical sensors and payloads take precedence, CubeSats are often deployed with limited defence capabilities due to perceived security through obscurity and physical inaccessibility, resulting in an attack surface collectively representative of enterprise networks without the power and processing to match. These attacks are harder to detect than single-stage attacks, as each step may seem innocuous on its own, but forms part of a larger operationally or physically destructive sequence (Gupta & Mishra, 2020), asymmetric to conventional warfare. Furthermore, CubeSats operate in highly resource-constrained environments, with limited computational power, storage, and energy due to their strict Size, Weight, and Power (SWaP) limitations (Weston *et al.*, 2025). These constraints make it difficult to implement traditional cybersecurity defences designed for more powerful systems, necessitating a new approach to CubeSat security. For example, the recent cyber attack on ViaSat's satellite network started with the exploitation of a virtual private network (VPN) vulnerability of several Gateway Earth Stations in ViaSat's ground segment (Boschetti et al, 2022). The adversary penetrated the demilitarised zone (DMZ) and then tunnelled through the trusted management network to the end users' modems. In the final stage of the attack, the deployment of a wiper malware disabled approximately 40,000 modems across multiple countries. This spillover effect outside of Ukraine, whether by mistake or through the selection of intersecting geographic zones, also resulted in one fifth (5,800) of Germany's wind turbines defective for control intervention. This has finally led to the U.S. (Committee on National Security Systems, 2025) requirement that cybersecurity be fully integrated across the lifecycle of space systems used to support national security missions, with real-time on-board intrusion detection and prevention forming part of the minimum defensive capability expected of such platforms. As a result, intrusion detection should be viewed not simply as an optional safeguard, but as a core requirement for mission assurance in security-sensitive satellite systems.

2. Analysis

2.1 Cybersecurity Risks in CubeSat Constellations

Although CubeSats offer many advantages, they are not immune to cybersecurity risks. These risks include traditional cyber threats such as data breaches, denial-of-service (DoS) attacks, and unauthorised access, but more critically, CubeSats are vulnerable to advanced cyberattacks, including multi-stage attacks. A significant aspect of CubeSat cybersecurity concerns the interconnectivity between satellites in a constellation. Although this interconnectedness allows for enhanced system performance, it also means that compromising one satellite can lead to cascading failures across the entire constellation (Martin et al, 2018). For example, in the event of a compromised satellite, the attacker may use it as a launchpad to manipulate communication protocols, inject false data into the system, or disrupt operational capabilities (Pavur et al, 2021).

Ground and user segments typically have a mixture of specialised and general-purpose equipment. "Security through obscurity" refers to the assumption, implemented in some kind of technical practice, that obscured interfaces of specialised equipment deter attackers. In reality, whether for general interest, notoriety or malicious intent, closed undocumented systems are routinely reverse-engineered and the resultant findings circulated. As an example, a voltage fault injection (or "glitching") using a custom modification chip ("mod-chip" for short) was demonstrated on a Starlink User Terminal from SpaceX (Wouters, 2022). There existed a glitching pattern that allowed an attacker to bypass signature verification and subsequently run arbitrary code on the terminal, further bridging the gap between the space and user segment for adversaries. For larger profit margins, vendors of embedded devices often economise on security mechanisms, hard-coding access credentials, using weak cryptographic schemes and failing to patch identified vulnerabilities in a timely manner, if at all (Manulis et al., 2021). Knowing this, attackers often find resource- constrained devices attractive targets. It follows that the ground segment must be locked down, and based on the principles of zero trust, the space segment should operate without any intrinsic trust of the ground and user segments. Furthermore, as illuminated by the Viasat attack where freedom of lateral movement and trust that the ground segment is secure provided the vector necessary for the attack, the principles of defence in depth mandate the space segment to have additional security measures or cyber defences independent of its neighbouring segments.

2.2 Multi-stage Intrusion Detection

Multi-stage attack detection is comparable to classic signature and anomaly-based techniques, albeit with differing terminology and priorities. The primary distinctions are to connect consecutive actions into potential, evolving and current attacks. As shown in (Navarro et al, 2018), multi-stage detection techniques can be classified as similarity-based, causal correlation, structural-based, case-based and mixed approaches. In summary of this research:

Similarity-based: Determine the degree of similarity between alerts to construct the attack scenarios. Following the premise that similar alerts are tied to the same root cause, when sorted by field, such as host IP address, they can be appended, clustered or compared to nominal traffic to detect multi-stage attacks.

Causal correlation: Where causal connections between actions may arise as a result of the latter, causal correlation methods use this as the key to detecting attacks. By establishing which steps facilitate future actions, it is possible to construct attack scenarios and forecast future steps from the causality of the current, allowing a more targeted defence effort.

Structural-based: Working backwards in comparison to prior approaches, structural-based methods utilise known vulnerabilities of the allied system and project current network traffic to construct potential attack paths that would exploit these weaknesses.

Case-based: For well-known and characterised attack scenarios, similar to traditional signature detection, case-based methods compare current network traffic to a database of attack signatures, determining when a specific sequence of network events occur.

Mixed or hybrid: Used to incorporate two or more of the above methods in an attempt to maximise the detection performance while mitigating the drawbacks encountered by differing techniques.

Our research is motivated by the need for space-segment cyber defences that can detect and mitigate multi-stage attacks, which represent a significant threat to the operational integrity of CubeSat constellations and are anticipated to gain command and control of satellite systems (Falco et al., 2021).

2.3 Challenges in CubeSat Cybersecurity

One of the key challenges in securing CubeSat constellations is the resource limitations inherent in these small satellites. CubeSats have limited computational resources, power, and memory, which means that traditional security solutions, such as firewalls, antivirus software, or full intrusion detection systems, cannot be deployed in their entirety. This presents a fundamental challenge for cybersecurity, as it limits the types of defences that can be applied without exhausting the satellite's limited energy supply or reducing its operational efficiency (Manulis et al., 2021). Furthermore, CubeSats operate in a harsh environment that presents additional challenges. These satellites are exposed to space radiation, which can cause hardware malfunctions and data corruption (Ecoffet, 2013). The lack of physical access to CubeSats further complicates security efforts; once a satellite is launched, it cannot be easily repaired or updated in response to emerging threats (Falco, 2020). These factors demand that any cybersecurity solutions be lightweight, adaptive, and capable of functioning in an environment where physical access is not possible. The need for resource efficient solutions is paramount in CubeSat security. For example, CubeSats rely on solar panels for power, and excessive computational tasks may quickly deplete their energy supply, leading to mission failure (Manulis et al., 2021). As such, any cybersecurity mechanisms need to minimise computational overhead and optimise energy consumption without sacrificing the effectiveness of intrusion detection, response, recovery and retaliation.

3. Emerging strategies for CubeSat Resilience

3.1 Machine Learning for Cyber Threat Detection

In response to the unique cybersecurity challenges posed by CubeSats, researchers have increasingly turned to machine learning (ML) techniques for intrusion detection and threat mitigation. ML algorithms have the ability to learn from vast amounts of system data and identify patterns of behaviour that indicate a potential attack (Tedeschi et al, 2022). Anomaly-based intrusion detection systems (IDS), for instance, are capable of flagging irregular behaviour that may indicate the presence of a multi-stage attack. Such attacks often involve subtle changes to system behaviour over time, which can be difficult for traditional IDSs to detect. Supervised learning methods are commonly used for intrusion detection in CubeSats, as they are trained on labelled datasets to identify known attack patterns. However, the lack of labelled data for space-based cyberattacks can be a significant limitation (Li et al., 2020). As a result, unsupervised and semi-supervised learning

techniques are gaining traction. These methods do not require large, labelled datasets, making them particularly useful for detecting novel threats in a dynamic space environment. Another promising approach is the use of deep learning, specifically convolutional neural networks (CNNs), which can learn complex patterns and provide high accuracy in detecting attacks (Fourati and Alouini, 2021) Moreover, CNNs can be designed to operate efficiently in low-resource environments, addressing one of the key challenges faced by CubeSat systems. The ability of deep learning algorithms to process large amounts of telemetry data and identify subtle anomalies makes them a strong candidate for improving CubeSat cybersecurity.

3.2 Data Challenges

The development of effective IDS has historically relied on the availability of comprehensive and realistic datasets. Seminal datasets like KDD99, and its successor NSL-KDD, provided a foundation for much of the early research in the field (Tavallaee et al., 2009). More contemporary datasets, such as UNSW-NB15 (Moustafa & Slay, 2015) and CIC-IDS2017 (Sharafaldin et al., 2018), were created to reflect more modern network traffic and attack vectors. However, a common critique, as noted by Moustafa and Slay, is that the patterns in older datasets like KDD99 are vastly different from those in current network environments, which motivated the creation of UNSW-NB15 (Moustafa & Slay, 2015). This evolutionary trend highlights a core problem: datasets are products of their specific network environments. The traffic characteristics, protocols, and attack surfaces of a corporate LAN are fundamentally different from those of a CubeSat constellation or other LEO satellite constellation such as Starlink, rendering terrestrial datasets a poor proxy for the space domain. Our major challenges therefore include the lack of suitable datasets to develop ML based controls. Where terrestrial security datasets have garnered notable success with a range of machine learning techniques, they lack the distinctive dynamics inherent to satellite networks. These datasets fail to encapsulate the unique dynamics and communication protocols of space systems, such as limited access windows, increased data loss, and latency. Even the recent satellite-oriented STIN dataset (Li et al., 2020) is only publicly released as a flow-level subset (TER20/SAT20) of the original dataset with no nominal satellite traffic available. As a result, there remains a lack of realistic data for on-board intrusion detection systems for space-segment networks. Furthermore, there is a significant class imbalance in existing datasets between benign/single-stage attacks. Central to the methodology of our future work is a data-centric AI process: iterative dataset engineering through scenario-based digital twin simulations, domain-specific feature encapsulation, and data augmentation. This integrated approach preserves theoretical rigor and objectivity while systematically refining data representations to enhance the validity and applicability of intrusion detection models in homogeneous satellite constellations. Though while this area is developed, a preliminary solution is proposed using semi and self-supervised learning.

3.3 Semi-supervised and Self-supervised Learning

While supervised learning has shown great promise in cybersecurity applications, CubeSat systems as mentioned often lack sufficient labelled data for training. To overcome this challenge, semi-supervised learning has been proposed as a viable solution. Existing between supervised and unsupervised learning, as an umbrella term, weakly-supervised learning works with data that may be incorrect (inaccurate supervision), coarsely labelled (inexact supervision) and or partially labelled (incomplete supervision) (Zhou, 2018). Semi-supervised learning, a form of incomplete supervision, is a popular technique utilised to implicitly improve the performance of few explicitly labelled data points through a greater sum of unlabelled data (Song et al., 2023). In semi-supervised learning, the model is trained on partially labelled data, allowing it to make predictions based on incomplete or noisy information (Yang, 2024). Semi-supervised learning works on the principle that similar data yields similar results and it is the purpose of the unlabelled data to aid in discerning which data is similar (Zhou, 2018). This approach is particularly useful for space systems, where attack datasets are scarce and difficult to obtain, and could aid space scenarios where aforementioned data is partially labelled at best.

For example, pseudo-labelling works by taking a small amount of labelled data and a large sum of unlabelled data. The labelled data is used to produce pseudo labels for the unlabelled data and those with a high confidence prediction are treated as ground truth and added to the network [119]. A subsequent neural network can then be trained or fine-tuned using the updated dataset [126] in a supervised manner.

Self-supervised learning, a form of unsupervised learning, is also gaining popularity in CubeSat cybersecurity. Self-supervised learning algorithms generate their own training labels from the data itself, enabling them to learn meaningful representations without relying on external supervision (Yang, 2024). This technique has shown promise in detecting complex, multi-stage attacks that involve multiple phases of exploitation and lateral movement within the system.

Our work preliminarily explores self-supervised contrastive learning, a similarity-based technique, to produce a nominal operational profile for a CubeSat (Liu et al., 2022). This approach creates a baseline of normal behaviour against which telemetry and network traffic can be compared to detect anomalies. Contrastive cosine pair learning can be adapted for satellite cyber security by treating each spacecraft, subsystem, or communication endpoint as part of an attributed network in which normal behaviour is defined by the expected agreement between a node and its local neighbourhood rather than by explicit attack labels. (Liu et al., 2022) argue that normal instances exhibit a consistent matching pattern with neighbouring structure, whereas anomalies emerge when attributes and structure become mismatched. In a satellite setting, this idea can be reframed with cosine similarity by maximising similarity between nominal instance pairs whose embeddings point in similar directions in feature space, while reduced similarity indicates deviation from expected operational state. Such nominal embeddings could be learned from pre-launch validation data and refined during the first few orbital epochs, allowing the system to characterise mission-specific baseline behaviour without requiring explicit attack labels. In this formulation, both system faults and cyberattacks are expected to manifest as lower cosine agreement with the learned nominal state, yielding a continuous measure of system normality rather than a hard class label.

Moreover, the approach could later be fine-tuned for supervised cyber resilience testing by introducing known attack behaviours as deliberately contrasting or opposing pairs, thereby supporting both anomaly detection and downstream attack classification. However, this preliminary method has a core assumption/limitation. It places a heavy reliance on trusted physical and software supply chains to ensure the initial profile is not built on a compromised or malicious foundation. Hence this approach is being proposed as an initial solution to the real-time onboards IDS problem while bespoke space domain datasets continue to mature.

3.4 Defence-in-depth and Zero-trust Architectures

The concept of defence-in-depth has long been a cornerstone of cybersecurity, particularly in complex systems. In a CubeSat constellation, this approach involves implementing multiple layers of security, ensuring that if one layer is breached, others remain intact. Examples of defence-in-depth measures include the use of firewalls, encryption, intrusion detection systems, and secure communication protocols. These layers work together to form a multi-faceted security posture, making it more difficult for an attacker to compromise the entire system. Zero-trust architectures, which operate under the principle of “never trust, always verify,” are also being increasingly applied in CubeSat security. In a zero-trust model, every request for access is authenticated and authorised, regardless of its source. This model assumes that any user or device within the network could be compromised, and as such, access to resources is restricted until the user or device has been properly verified (Kumar & Chatterjee, 2021). Implementing zero-trust principles in CubeSat constellations can significantly reduce the likelihood of lateral movement by attackers, preventing them from exploiting one compromised satellite to launch attacks on others.

3.5 A framework for Verifiable Cyber Resilience

To standardise security practices, we propose the creation of a minimum cyber resilience capability checklist. This checklist would provide a concrete set of security controls and best practices that CubeSat developers and operators should implement (Plotnek, 2022). In line with defence-in-depth principles, this checklist should not rely on any single safeguard, but instead require multiple complementary layers of protection, such as authentication controls, encryption, intrusion detection systems, secure communication protocols, and other mechanisms that ensure if one layer is breached, others remain intact. In parallel, zero-trust principles should also be incorporated, operating under the assumption of “never trust, always verify,” such that every request for access is authenticated and authorised regardless of its source. In a CubeSat constellation, where any user, device, subsystem, or satellite may be compromised, this reduces the likelihood of lateral movement and helps prevent one compromised satellite from being used to launch attacks on others.

In time, this would be accompanied by a testbench for conducting tests analogous to the vacuum and vibration testing performed during hardware qualification. This testbench would include components to cover core low-cost, high-reward attack vectors, with capabilities to test for vulnerabilities in authentication, encryption, fuzzing, replay attacks, and Denial of Service (DoS) scenarios for example. The results of the testbench would be delivered back to the owner or producer in a straightforward pass/fail method, clearly indicating which tests have failed and referencing standard methods to ensure these vulnerabilities are covered. This provides a verifiable and repeatable process for ensuring a baseline of cyber resilience.

Ideally, this process would adhere to the resilience framework taxonomy defined by Plotnek (2022), not only by supporting the detection of threats, but by enabling systems to *react* to threats detected, *survive* adverse effects, *sustain* minimal operational capacity, *recover* to nominal system state, and *adapt* from information gathered. In this sense, the proposed checklist and testbench are not merely mechanisms for identifying vulnerabilities prior to launch, but an initial framework for demonstrating that a system can maintain and restore mission-relevant functionality under adverse cyber conditions. This would provide a foundation to the U.S. requirements through a standardised cyber detection threshold prior to the launch of military or dual use satellites.

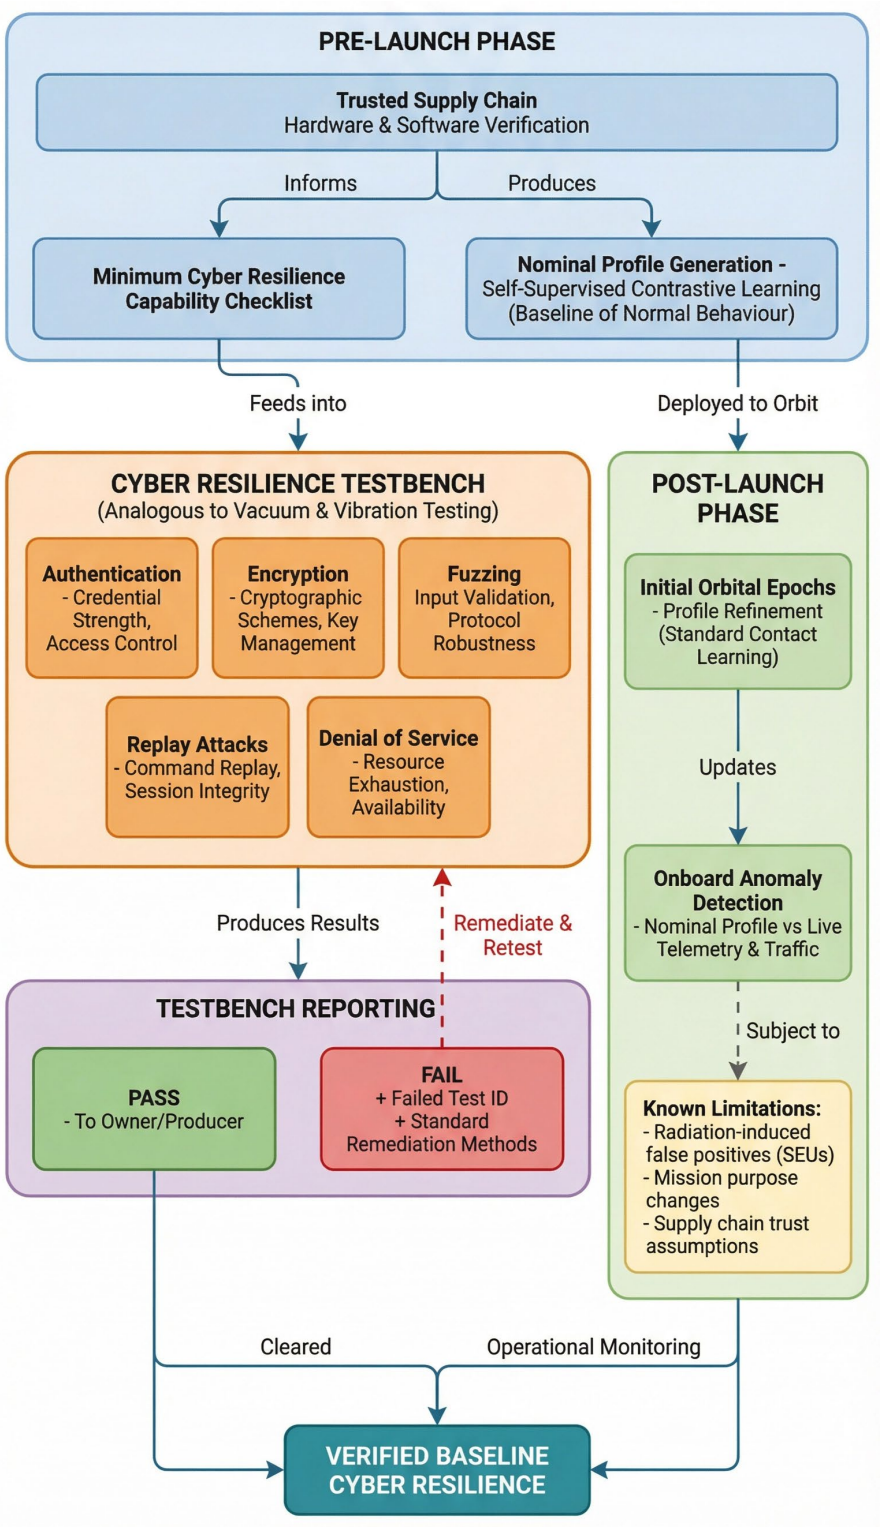

Figure 1: Minimum cyber resilience framework

4. Conclusion

CubeSat constellations are becoming increasingly important, offering low-cost, scalable solutions for global coverage in both civilian and military applications. However, their inherent constraints regarding Size, Weight, and Power (SWaP), combined with the harsh space environment and growing interconnectivity of constellation architectures, make them highly vulnerable to cyberattacks. To address these challenges, this work proposes a framework that combines self-supervised learning for nominal profiling with a structured approach to verifiable cyber resilience. By using contrastive learning to establish a baseline of normal operations, we aim to enable lightweight, onboard anomaly detection as an initial intrusion detection capability suited to resource-constrained platforms. While acknowledging the limitations of this approach, including its reliance on supply chain integrity and the potential for false positives, we propose post-launch refinement of the nominal profile alongside the development of bespoke, space domain-specific cyber datasets. This is complemented by a proposed minimum cyber resilience capability checklist and an accompanying testbench to standardise and verify security postures against common attack vectors, supporting a resilience process that can initially detect threats, supporting future downstream segments to react to adverse conditions, sustain minimal operational capacity, recover toward a nominal system state, and adapt from information gathered. There is not yet a widely applicable IDS solution that can be used across vendors and equipment; accordingly, this work represents a preliminary step toward a standardised cyber resilience requirement for detection and verification. In doing so, it aligns with emerging U.S. national security expectations that cybersecurity be integrated across the full lifecycle of relevant space systems, including the implementation of real-time onboard intrusion detection and prevention capabilities, and may help inform both downstream system design and operator understanding of cyber posture before and after launch.

Ethics Declaration: This research did not involve human participants, human data, or animal subjects. Ethical clearance was therefore not required for this study.

Artificial intelligence Declaration: Generative artificial intelligence tools (ChatGPT, Claude) were used in a limited capacity to assist with language refinement, structural editing, and clarity of expression. All technical content, analysis, interpretations, references and conclusions were developed by the author/s, who retain full responsibility for the accuracy, originality, and integrity of the work.

References

Boschetti, N., Gordon, N. and Falco, G. (2022) 'Space cybersecurity lessons learned from the ViaSat cyberattack', in ASCEND 2022. DOI: https://doi.org/10.2514/6.2022-4380.

Committee on National Security Systems (2025) National cybersecurity instruction for space systems used to support national security missions. CNSSI No. 1200, August. Available at: https://www.mayerbrown.com/-/media/files/perspectives-events/publications/2025/12/cnssi-1200.pdf?rev=b887f8cf41dc4e29913da872329bf5a1&hash=2E34079EA288E6D3C17AB53E1C60EEB3.

Ecoffet, R. (2013) 'Overview of in-orbit radiation induced spacecraft anomalies', IEEE Transactions on Nuclear Science, 60(3), pp. 1791-1815. DOI: https://doi.org/10.1109/TNS.2013.2262002.

Falco, G. (2020) 'When satellites attack: Satellite-to-satellite cyber attack, defense and resilience', in ASCEND 2020. DOI: https://doi.org/10.2514/6.2020-4014.

Falco, G., Viswanathan, A. and Santangelo, A. (2021) 'CubeSat security attack tree analysis', in 2021 IEEE 8th International Conference on Space Mission Challenges for Information Technology (SMC-IT), pp. 68-76. DOI: https://doi.org/10.1109/SMC-IT51442.2021.00016.

Fourati, F. and Alouini, M.-S. (2021) 'Artificial intelligence for satellite communication: A review', Intelligent and Converged Networks, 2(3), pp. 213-243. DOI: https://doi.org/10.23919/ICN.2021.0015.

Housen-Couriel, D. (2016) 'Cybersecurity threats to satellite communications: Towards a typology of state actor responses', Acta Astronautica, 128, pp. 409-415. DOI: https://doi.org/10.1016/j.actaastro.2016.07.041.

Li, K., Zhou, H., Tu, Z., Wang, W. and Zhang, H. (2020) 'Distributed network intrusion detection system in satellite-terrestrial integrated networks using federated learning', IEEE Access, 8, pp. 214852-214865. DOI: https://doi.org/10.1109/ACCESS.2020.3041641.

Liu, Y., Li, Z., Pan, S., Gong, C., Zhou, C. and Karypis, G. (2022) 'Anomaly detection on attributed networks via contrastive self-supervised learning', IEEE Transactions on Neural Networks and Learning Systems, 33(6), pp. 2378-2392. DOI: https://doi.org/10.1109/TNNLS.2021.3068344.

Manulis, M., Bridges, C.P., Harrison, R., Sekar, V. and Davis, A. (2021) 'Cyber security in New Space', International Journal of Information Security, 20, pp. 287-311. DOI: https://doi.org/10.1007/s10207-020-00503-w.

Moustafa, N. and Slay, J. (2015) 'UNSW-NB15: A comprehensive data set for network intrusion detection systems (UNSW-NB15 network data set)', in 2015 Military Communications and Information Systems Conference (MilCIS), pp. 1-6. DOI: https://doi.org/10.1109/MilCIS.2015.7348942.

Navarro, J., Deruyver, A. and Parrend, P. (2018) 'A systematic survey on multi-step attack detection', Computers & Security, 76, pp. 214-249. DOI: https://doi.org/10.1016/j.cose.2018.03.001.

Pavur, J., Strohmeier, M., Lenders, V. and Martinovic, I. (2021) 'In the same boat: On small satellites, big rockets, and cyber trust', in 2021 13th International Conference on Cyber Conflict (CyCon), pp. 151-169. DOI: https://doi.org/10.23919/CyCon51939.2021.9468300.

Plotnek, J.J. (2022) 'A threat-driven resilience assessment framework and security ontology for space systems', Thesis, University of South Australia. Available at: https://www.researchgate.net/publication/370102679_A_Threat-Driven_Resilience_Assessment_Framework_and_Security_Ontology_for_Space_Systems.

Sharafaldin, I., Habibi Lashkari, A. and Ghorbani, A.A. (2018) 'Toward generating a new intrusion detection dataset and intrusion traffic characterization', in Proceedings of the 4th International Conference on Information Systems Security and Privacy (ICISSP), pp. 108-116. DOI: https://doi.org/10.5220/0006639801080116.

Song, Z., Yang, X., Xu, Z. and King, I. (2023) 'Graph-based semi-supervised learning: A comprehensive review', IEEE Transactions on Neural Networks and Learning Systems, 34(11), pp. 8174-8194. DOI: https://doi.org/10.1109/TNNLS.2022.3155478.

Sweeting, M.N. (2018) 'Modern small satellites-changing the economics of space', Proceedings of the IEEE, 106(3), pp. 343-361. DOI: https://doi.org/10.1109/JPROC.2018.2806218.

Tavallaee, M., Bagheri, E., Lu, W. and Ghorbani, A.A. (2009) 'A detailed analysis of the KDD CUP 99 data set', in 2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications, pp. 1-6. DOI: https://doi.org/10.1109/CISDA.2009.5356528.

Tedeschi, P., Sciancalepore, S. and Di Pietro, R. (2022) 'Satellite-based communications security: A survey of threats, solutions, and research challenges', Computer Networks, 216, p. 109246. DOI: https://doi.org/10.1016/j.comnet.2022.109246.

Weston, S.V., Burkhard, C.D., Stupl, J.M., Ticknor, R.L., Yost, B.D., Austin, R.A., Galchenko, P., Newman, L.K. and Santos Soto, L. (2025) State-of-the-art small spacecraft technology. NASA/TP—20250000142 (2024 edition). NASA. Available at: http://www.nasa.gov/smallsat-institute/sst-soa.

Whang, S.E., Roh, Y., Song, H. and Lee, J.-G. (2023) 'Data collection and quality challenges in deep learning: A data-centric AI perspective', The VLDB Journal, 32(4), pp. 791-813. DOI: https://doi.org/10.1007/s00778-022-00775-9.

Wouters, L. (2022) 'Glitched on Earth by humans: A black-box security evaluation of the SpaceX Starlink user terminal', Black Hat USA 2022. Available at: https://i.blackhat.com/USA-22/Wednesday/US-22-Wouters-Glitched-On-Earth.pdf.

Yang, J., Zhou, K., Li, Y. and Liu, Z. (2024) 'Generalized out-of-distribution detection: A survey', International Journal of Computer Vision, 132, pp. 5635-5662. DOI: https://doi.org/10.1007/s11263-024-02117-4.

Yang, X., Song, Z., King, I. and Xu, Z. (2023) 'A survey on deep semi-supervised learning', IEEE Transactions on Knowledge and Data Engineering, 35(9), pp. 8934-8954. DOI: https://doi.org/10.1109/TKDE.2022.3220219.

Zhou, Z.-H. (2018) 'A brief introduction to weakly supervised learning', National Science Review, 5(1), pp. 44-53. DOI: https://doi.org/10.1093/nsr/nwx106.