18th International Conference on Space Operations, Montreal, Canada, 26 - 30 May 2025.

SpaceOps-2025, ID # 194

Cybersecurity for Space Ground Control segments, the Zero Impact approach

Daniel San Miguel Reyero[a], Aurora María Salvador Gutiérrez[b]

[a] GMV, Spain dmiguel@gmv.com
[b] GMV, Spain amsalvador@gmv.com

Abstract

Space ground segments operate under strict validation and isolation requirements that render traditional cybersecurity methods, such as EDR/XDR agent-based monitoring, continuous patching or active vulnerability scanning unfeasible. This paper presents a "Zero Impact" Passive Risk assessment and Cyber Monitoring methodology and tools, specifically designed for these environments. The approach enables non-intrusive, agent-less monitoring of vulnerabilities, configuration changes, and security status using automated, read-only tools that preserve system integrity and validation/certification.

Validated in a representative chain of a constellation ground segment, the solution provides reliable, daily risk assessments with minimal operational overhead and practically reduces to zero false positives and negatives compared to conventional passive vulnerability detection based in the CVE/CPE-based techniques or active scanners. This methodology offers a practical and effective alternative for improving cybersecurity in mission-critical, highly regulated space systems, without compromising operational continuity.

Keywords: Cyber Monitoring, vulnerabilities management, passive vulnerability detection, active vulnerability detection, Cybersecurity, Security-by-design

Acronyms/Abbreviations

CPE: Common Platform Enumeration
EDR: Endpoint Detection and Response
NVD: National Vulnerability Database
NIDS: network intrusion detection systems
NIST: National Institute of Standards and Technology (USA)
SIEM: Security information and event management (SIEM)
XDR: eXtended Detection and Response

1. Introduction

In the realm of space operations, the cybersecurity of ground control segments represents a unique and highly constrained challenge, demanding specialized approaches distinct from those employed in conventional Information and Communication Technology (ICT) systems. Due to the critical nature of space infrastructure, widely used cybersecurity practices such as agent-based monitoring, continuous patching, and Extended Detection and Response (XDR) frameworks often pose unacceptable risks to system stability and operational continuity.

To address these constraints, GMV has developed a viable alternative for space operations with a "Zero Impact" Cyber Monitoring approach tailored specifically for the Space Ground Segment. This methodology emphasizes non-intrusive, agent-less monitoring tools capable of delivering continuous inventory assessments and vulnerability management without compromising system validation or operations.

Through this approach, GMV achieves comprehensive visibility into software assets, system configurations, and security status while ensuring compliance with stringent accreditation requirements.

This paper details the Zero Impact model, evaluates its effectiveness against traditional vulnerability detection techniques, and presents empirical results derived from enterprise-grade Linux environments. The findings

demonstrate the feasibility and advantages of this approach in maintaining robust cybersecurity in mission-critical, highly regulated operational contexts.

1.1 Restrictions on deployments in ground segments

The deployment of infrastructures and software components within the ground segments of aerospace systems, while differing according to the specific organizations and their respective approved methodologies, generally maintains strict control over the validation of the involved segments and components. Moreover, there are significant constraints on the inclusion of changes or updates to the deployed software, with the overarching objective of maximizing system availability by preventing the execution of software that has not undergone the required authorized validation processes.

Accordingly, the following restrictions are commonly enforced:

1. Only software that has undergone a formal validation process is authorized for deployment.
2. Validation of potential interactions between the software and other software/hardware components is required.
3. Any software update mandates either a full or incremental validation of the affected software or component.
4. An airgap to prevent access from external networks, including the internet, is mandatory.

These requirements pose significant challenges to security management, as they severely limit, or completely prevent, the ongoing patching of vulnerabilities. This is due to the inability to continuously validate patchs or patched software. As a result, traditional patch management cycles are reduced or eliminated, leading to an increased risk associated with the presence of a large number of unpatched vulnerabilities. By way of illustration, Red Hat disclosed more than 2,280 vulnerabilities in 2023 alone [1], equating to over 50 vulnerabilities per month in a typical installation or more than 2,000 over a three-year period for a single hosts.

In the absence of the possibility to patch systems and this eliminate potential attack vectors, an alternative strategy involves the use of protection, monitoring, and response tools—commonly referred to as EDR/XDR systems. These aim to block attacks at the operating system level, detect them early, and respond in a way that prevents their success and/or propagation. This approach reduces both the impact of any intrusion and its spread to other components of the infrastructure. Unfortunately, commercial EDR/XDR systems typically require the installation of agents on monitored devices, internet connectivity for updates and access to vendor databases, and the frequent application of security and functionality patches. As previously discussed, this approach is fundamentally incompatible with the strict validation standards mandated in the aerospace sector.

1.2 Risk Identification

Following the discussion in the previous section, the traditional approach adopted by the ground segment of the aerospace sector is to accept the inherent risks associated with both the presence of system vulnerabilities and the lack of visibility into system activities—whether these are authorized operations, human errors, misuse by administrative or operational personnel, or, in the worst-case scenario, security incidents.

However, in order to assume any risk, the first essential step is risk identification. In this context, it is necessary to determine both the vulnerabilities present in the system and any security incidents that may occur.

The following sections analyses the process of identifying existing vulnerabilities along with their associated criticality, as well as the procedure for detecting system changes in order to distinguish between authorized actions (and assess their impact) and unauthorized modifications, along with the corresponding security risks. In both cases, the specific constraints of these systems, as previously outlined, must be taken into careful consideration.

1.2.1 Vulnerability identification

The most widely adopted methodology for vulnerability identification is based on the formal registry maintained by cybersecurity researchers within the National Vulnerability Database (NVD), which is published and managed by the National Institute of Standards and Technology (NIST). This database catalogs software vulnerabilities by linking them to a unique identifier known as a CPE (Common Platform Enumeration), which specifies the software's manufacturer, product, platform, and version.

Unfortunately, this database exhibits several shortcomings, particularly in terms of incomplete coverage of affected software (i.e., incomplete or missing CPE entries). This issue is especially prevalent when third-party software is embedded within operating systems, software distributions, or appliance firmware. Such scenarios are common with open-source or licensed software that is reused across multiple Linux distributions and/or derivative products. As a result of this reuse, the following situations frequently occur:

- Modification of the software package name, which makes it nearly impossible to identify the software using the CPE information provided in the NVD.

- Fragmentation of the software into multiple packages, typically by separating libraries, many of which do not retain the original software name, thus preventing traceability.

In both cases, false negatives are introduced, as the software package names used by vendors do not match or cannot be traced back to the original software packages. This information often remains internal to the entities repackaging or modifying this software.

Additionally, in recent years, to address the instability associated with full software upgrades, enterprise software distributions, both Linux and Windows, have widely adopted the use of targeted security patches. This approach avoids the risks introduced by installing new versions that may include untested features or deprecated functionalities, thereby allowing for safer patching in high-criticality environments.

Consequently, even if a specific software version is listed as vulnerable in public vulnerability databases, it may have been patched by the vendor, thus invalidating the identification process and resulting in a false positive.

Faced with this approach, which involves significant manual effort and unacceptable rates of false positives and false negatives, alternative solutions such as active vulnerability scanners are often considered. These tools require network connectivity to the target systems and hosts credentials for authenticated privileged access in order to identify vulnerable software packages, using a methodology similar to that proposed in this paper, albeit with a more limited scope in terms of supported packages.
While this approach is generally effective at detecting the most relevant and recently disclosed vulnerabilities, its performance tends to decline when addressing older or lower-severity vulnerabilities.
Nevertheless, the use of such scanners is typically not approved in critical production environments, as the high volume of newly introduced plugins, required to support emerging vulnerabilities, may cause unpredictable operational impacts and thus cannot be reliably validated within highly regulated infrastructures.

Finally, the National Institute of Standards and Technology (NIST) has advocated for the use of OpenSCAP as a means of identifying vulnerabilities on target hosts, based on vendor-provided vulnerability definitions and through the local execution of the tool on the systems to be analyzed.
However, despite its high level of accuracy, this technique is not applicable for vulnerability identification in aerospace ground segment environments. This is due to the requirement for deploying dedicated tools that impose substantial resource demands, particularly exceeding 8 GB of RAM and requiring several minutes of execution per hos, rendering it impractical for use in such constrained and highly regulated operational contexts.

1.2.2 Security incidents detection

In the field of cybersecurity, the current industry standard for the detection of security incidents involves the use of various endpoint protection tools, commonly referred to as EDR (Endpoint Detection and Response) or XDR (eXtended Detection and Response) when they include integration with solutions such as SIEMs (Security Information

and Event Management) or the detection of malicious network activities. Unfortunately, this approach is not applicable in satellite and mission control centers, as it requires the installation of specific agents, connectivity with cloud services, and the application of regular updates, all of which are incompatible with the constraints previously outlined [2].

As an alternative, passive network intrusion detection systems (NIDS) may be employed. However, these are often insufficient, as the primary threat in an isolated network, such as those used in aerospace ground segments, is not an external attack via interconnected networks (which are either heavily restricted or non-existent), but rather an insider threat.

Finally, event monitoring solutions such as SIEM platforms offer a degree of visibility into system activity. Nevertheless, they are often inadequate for distinguishing between authorized operations and unauthorized actions carried out by an internal attacker.

Unfortunately, relying solely on NIDS and SIEM solutions is insufficient as a monitoring strategy for critical infrastructures such as the ground segments of space systems. These tools fall short in effectively detecting insider attacks and the exploitation of software systems characterized by a high volume of active vulnerabilities.

2. The Zero impact approach

The proposed objective, is the definition of a new Non-Intrusive Methodology in Ground Segment Environments, given the strict validation and non-connectivity requirements discussed previously, as well as the limitations in risk identification processes.

As a starting point as an alternative non-intrusive approach was implemented in a representative chain within the control center of a satellite constellation. This approach, inspired by white-box security audit methodologies, focused on manual identification of system states to reveal not only the installed software packages but also other key elements such as user accounts, open ports, kernel modules, firewall rules, privilege escalation configurations (e.g., sudo policies), network parameters, and more.

To extract this information, a set of operating system-specific commands and lightweight scripts was defined, relying exclusively on read-only operations and interpreters that are available in the minimal installations of the various operating systems. This ensured that no special execution requirements or system modifications were necessary.

Once these commands and scripts were identified, they underwent a validation process to ensure they posed no operational risk to the systems. Following validation, the scripts were executed in a one-time manner, extracting various system inventories and configuration tables, including:

- Scheduled tasks
- Firmware inventory
- Users and groups
- Network parameters (e.g., interfaces, IP addresses, routes)
- Packaged and unpackaged software
- SSH keys
- Privilege escalation rules
- OS configuration parameters
- Host firewall rules
- Hardening deviations

This collected information enabled two key objectives: passive vulnerability detection (based on the installed software) and the detection of configuration changes, allowing for subsequent analysis to determine whether these changes resulted from authorized actions, unauthorized modifications, human error, software faults, or security incidents.

In terms of vulnerability identification, it was observed that reliance on NVD data led to more than 50% false positives and false negatives, as previously discussed. Consequently, an alternative approach was adopted, based on processing vulnerability data directly from Linux distribution vendors. This approach proved to be significantly more accurate and reliable. In fact, it aligns with current recommendations by NIST, which advocates the use of the OpenSCAP tool.

However, OpenSCAP is not applicable in aerospace contexts, as it requires deployment on monitored hosts, the deployment of vendor vulnerability definitions, and the execution of a resource-intensive process—demanding more than 8 GB of memory and several minutes of processing time per system.
Alternatively, the methodology employed here enabled generation of equivalent results to OpenSCAP, using only the updated list of installed packages and processing it entirely offline, thereby adhering to the operational constraints of aerospace ground infrastructures.

In addition to enabling the generation of a comprehensive infrastructure inventory with highly detailed configuration parameters, the approach also allowed for the detection of low-level changes through delta analysis between successive executions. This made it possible to identify all modifications made to the infrastructure with a high degree of precision.
Recognizing the significant potential of this solution, efforts were undertaken to automate the data extraction process while maintaining a zero-alteration policy toward the monitored systems. To this end, the following components were developed:

- **Task Planner**: Responsible for scheduling, dispatching, executing, retrieving results, and restoring the monitored systems to their initial state on a periodic basis (e.g., every 24 hours).
- **Importer**: Manages the ingestion of updated data and the detection of relevant changes.
- **Workflow Manager**: Oversees delta management through a ticket-based mechanism.
- **Report Generator**: Produces inventory reports and cybersecurity status summaries including the hardening deviations.

This architecture externalized all security-related logic to a system outside the monitored platform, requiring only SSH connectivity, similar to the access granted to operations and administrative personnel. Crucially, it did not introduce any modifications to the monitored systems that might affect their functionality or certification status.

4. Results

Through the implementation of this new approach and the development of supporting tools, it has been possible to maintain a continuously updated cybersecurity status on a daily basis with a zero impact over the standard development and deployment methodology for the ground space deployments. This includes not only a highly reliable assessment of existing vulnerabilities but also enhanced visibility into changes occurring within the monitored systems, providing a valuable complement to traditional SIEM and NIDS solutions.
Furthermore, it has been clearly demonstrated that approaches based on CVE and CPE identifiers generate a significantly high volume of false positives and false negatives. In most cases, the error rate exceeds 50%, meaning that for every correctly identified vulnerability, another is either missed or incorrectly flagged. This results in considerable effort being required for validation, while offering limited reliability in practice.

5. Discussion

Traditional methodologies for the deployment of critical infrastructures, typically characterized by waterfall-based processes and formal validation campaigns, are increasingly being questioned due to their limited responsiveness to change. In response, new agile methodologies, continuous integration practices, and automated validation mechanisms are being introduced. However, until these approaches reach the required levels of reliability to be deployed in high-impact systems, it remains essential to adopt cybersecurity solutions capable of identifying the risks assumed by production infrastructures, as well as detecting potential intrusions—particularly those originating from internal actors, which tend to be the most likely, impactful, and persistent.

6. Conclusions

The proposed methodology, grounded in a zero-impact approach toward existing platforms, combined with the automation of processes through tools specifically designed for the space sector, offers a viable risk assessment providing a precise vulnerability identification, security incidents, unauthorized changes and hardening deviations means of addressing this challenge. It enables the enhancement of cybersecurity posture within current infrastructures without necessitating additional validation of currently deployed components.
The methodology presented, grounded in a zero-impact approach to existing platforms and supported by the automation of processes through tools specifically tailored for the space sector, provides an effective and reliable capabilities of risk assessment. It enables the identification of vulnerabilities, detection of security incidents, recognition of unauthorized changes, and deviations from hardening baselines. Crucially, this approach enhances the

cybersecurity (in terms of reliability of the cyber-status and incident response capabilities) of current infrastructures without requiring additional validation of already deployed components.

References

In the text
[1] Red Hat product security risk report, April 26, 2024 https://www.redhat.com/en/resources/product-security-risk-report-2023 (accessed 13.04.25).
[2] Magic Quadrant for Endpoint Protection Platforms, 23 September 2024 - ID G00808300 https://www.gartner.com/doc/reprints?id=1-2IWARHR9&ct=240924 (accessed 13.04.25).