Security-by-Design at the Telco Edge with OSS: Challenges and Lessons Learned

Carmine Cesarano[1], Alessio Foggia[1], Gianluca Roscigno[2], Luca Andreani[3], Roberto Natella[1]

[1]*Università degli Studi di Napoli Federico II, Naples, Italy*
[2]*System Management S.p.A, Naples, Italy*
[3]*DigitalPlatforms S.p.A, Rome, Italy*

Abstract—**This paper presents our experience, in the context of an industrial R&D project, on securing GENIO, a platform for edge computing on Passive Optical Network (PON) infrastructures, and based on Open-Source Software (OSS). We identify threats and related mitigations through hardening, vulnerability management, digital signatures, and static and dynamic analysis. In particular, we report lessons learned in applying these mitigations using OSS, and share our findings about the maturity and limitations of these security solutions in an industrial context.**

Index Terms—**Edge Computing, Security-by-Design, OSS**

I. INTRODUCTION

Edge computing has emerged as a transformative paradigm in the telecommunications and industrial sectors, enabling low-latency data processing closer to the end-users [1]. Traditionally, edge computing and broadband access networks have operated independently, with edge workloads running on dedicated servers, and Passive Optical Networks (PON) providing high-speed broadband connectivity. However, leveraging PON hardware infrastructure for edge computing presents a promising opportunity to create a high-performance and cost-effective platform for running edge services [2], [3].

Unlike centralized cloud environments, deploying workloads on PON infrastructures introduces security risks. A primary concern is the physical exposure of hardware in uncontrolled environments, increasing the risk of tampering and unauthorized access [4], [5]. Security risks are amplified by the complexity of software architectures, which rely on Commercial Off-the-Shelf (COTS) and Open-Source Software (OSS) components for virtualization, software-defined networking, and orchestration. Software reuse provides flexibility and cost-efficiency, but can also introduce software vulnerabilities and expose to compromised software dependencies [6]. Multi-tenancy introduces more security risks, as different edge applications share the same infrastructure. Maintaining isolation between tenants is critical to prevent escalation of security attacks [7].

The GENIO project [8] is a joint R&D initiative between academic and industry partners, aiming to achieve a secure-by-design edge computing platform integrated with PON networks. Unlike conventional edge models that rely on dedicated servers, GENIO leverages existing PON hardware to host multi-tenant edge services, directly within the telecom infrastructure. This approach can optimize resource utilization and create new revenue opportunities for telecom operators, without the need for additional investments in dedicated servers. One of the main objectives of the GENIO project is to align the platform with security regulations, such as the European Cyber Resilience Act [9] and CE marking certification [10]. This objective shaped the platform by guiding threat mitigations.

Despite the growing emphasis on security-by-design, existing frameworks are not directly applicable for the design of the GENIO platform. Technical standards, e.g., from CISA [11] and NIST [12], provide high-level security guidelines, but these lack practical, actionable implementation details. In real-world deployments, security must be tailored to the specific hardware and software technology and to the operational constraints of heterogeneous industrial environments, thus requiring customized solutions. Academic research has proposed security-by-design methodologies for specific types of systems, such as cloud applications [13], smart grids [14], big data frameworks [15], and 5G networks [16]. However, no blueprint is readily available for security-by-design in PON-based edge computing.

This paper presents our findings from the design of the GENIO platform, covering security risks across hardware, OS, middleware, and applications. We evaluated the maturity and limitations of security mitigations, based on OSS solutions, and documented several challenges due to architectural constraints and software heterogeneity. By aligning security research with deployment realities, this work provides key lessons learned, bridging the gap between theoretical security models and practical implementation in an industrial context.

II. THE GENIO PROJECT

The GENIO project aims to develop a platform that integrates edge computing capabilities into telecom central offices, leveraging PON equipment. The architecture spans three layers: the cloud layer, the edge layer, and the far-edge layer, enabling flexible application deployments based on latency and computation requirements.

Deployment. As shown in Figure 1, the *far-edge layer* includes Optical Network Units (ONUs), which are deployed in residential and business premises, and which connect users to the fiber network. In GENIO, ONUs are equipped with additional low-end computing resources, enabling them to

run applications with ultra-low latency requirements. The *edge layer* includes Optical Line Terminals (OLTs), which are devices located in telecom central offices, and which are repurposed in GENIO to serve as edge computing hubs. Originally designed for PON connectivity management, OLTs are enhanced with additional hardware and software to provide computational and storage resources. This layer is optimized for applications with strict latency and bandwidth requirements, balancing performance and resource availability. For applications with less stringent latency requirements, the *cloud layer* offers high computational and storage resources. The cloud layer also behaves as the orchestration center, managing resources across the edge and far-edge layers, and handling complex tasks that exceed the capabilities of local devices.

Use cases. GENIO is designed to serve both *business users* and *end-users*. Business users, such as service operators and enterprises, are providers of edge applications to be shared as container images on a public registry of the GENIO project. Examples include ML workloads, real-time analytics, IoT data processing (e.g., from smart meters, cameras, and sensors), and telecom network functions. Business users can leverage the GENIO platform through an Infrastructure-as-a-Service (IaaS) model, by leasing computing, storage, and networking resources on the edge to run their applications. End-users, which include individual customers and businesses, interact with the platform via a Software-as-a-Service (SaaS) model, by consuming edge applications provided by business users.

Software Architecture. The GENIO platform a distributed and multi-layered architecture to manage network operations, resource allocation, and application deployment. In particular, OLTs are based on x86 COTS hardware, and they integrate several OSS technologies for software-defined networking (SDN), including ONOS [17], VOLTHA [18] and Open Networking Linux (ONL) [19], to dynamically and efficiently manage PON network resources. OLTs are extended to support secure multi-tenancy for running edge applications, based on OSS virtualization technologies. The physical resources of the OLT are managed using a cluster of virtual machines, managed using the Linux/KVM hypervisor. Edge applications can run in either hard isolation (in dedicated virtual machines) or soft isolation (in containers and network namespaces within the virtual machines), to accomodate different security and performance requirements. Additionally, GENIO uses Kubernetes and Proxmox [20] for orchestrating the virtual machines and containerized applications, allowing scalable, resilient, and efficient workload management based on current network and computational conditions.

III. THREAT MODELING FOR GENIO

Securing the GENIO platform required developing a comprehensive threat model to identify risks across the cloud, edge, and far-edge layers. Using the STRIDE methodology [21], we systematically identified threats, from physical tampering of ONUs and OLTs, to software vulnerabilities and misconfigurations of orchestration services. This process led to the
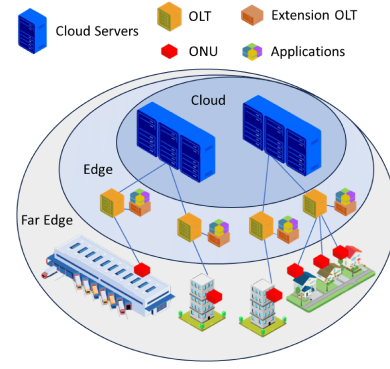

Fig. 1. GENIO deployment across cloud, edge and far-edge layers.

categorization of risks into *Infrastructure-level*, *Middleware-level*, and *Application-level* threats. Figures 2 and 3 provide an overview of the GENIO architecture and a summary of threats and mitigations, respectively.

A. Infrastructure-level Threats

The infrastructure level encompasses hardware components and low-level software, which are the foundation of the GENIO architecture.

T1 Network Attacks The distributed nature of the GENIO architecture creates multiple points of vulnerability for secure data transmission, spanning OLTs, ONUs, inter-OLT links, and cloud interactions. Adversaries can eavesdrop, modify traffic, or impersonate network components at various stages, with *interception and replay attacks* posing a direct threat to data integrity and authenticity. Physical exploits like *downstream hijacking* and *ONU impersonation* target the PON architecture at the hardware and firmware level. Infrastructure-level tampering often involves *physically tapping fiber connections* [22] or modifying device firmware to *siphon traffic* [23].

T2 Code Tampering Attackers can target low-level system components to introduce persistent threats, embedding malware or backdoors into the platform. *Reverse engineering*, *binary untrusted patching/updating*, and *firmware manipulation* are common attack techniques that allow adversaries to manipulate hypervisors, kernels, and system binaries. A successful compromise at this level can provide long-term control over the entire host machine.

T3 Privilege Abuse Misconfigurations in low-level software, such as unrestricted OS accounts, services, and files, can expose the system to *privilege escalation*. Intruders can exploit these flaws to expand their control over the platform, by hijacking administrative functions and achieving persistency. This can facilitate service disruptions and data thefts, posing a severe risk to the security of GENIO operations.

T4 Software Vulnerabilities Unpatched or unknown vulnerabilities in low-level software can be exploited by attackers to gain full access to the host machine, and to break isolation mechanisms. Unfortunately, handling these vulnerabilities can
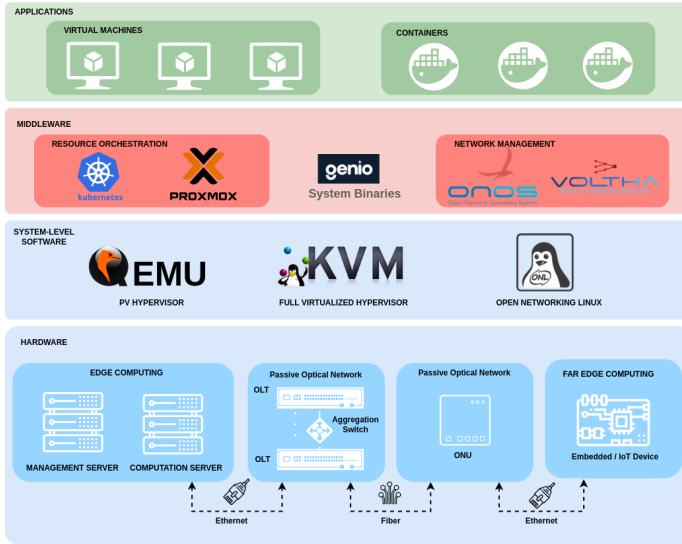
Fig. 2. GENIO architecture.

be quite difficult, since OLTs and ONUs are managed and updated remotely. Moreover, the GENIO platform relies on a custom Linux kernel configuration to support SDN software, requiring continuous vulnerability assessment to secure its custom stack. Any failure to address these vulnerabilities can expose the infrastructure to *kernel exploits* and *container escaping*.

B. Middleware-level Threats

The middleware level in GENIO includes software-defined networking (VOLTHA, ONOS), virtualization and container management (Proxmox, Kubernetes). These components provide powerful interfaces to programmatically manage resources, but also introduce more security challenges.

T5 Privilege Abuse *Misconfigurations* can also apply to middleware, such as overprivileged roles and unrestricted API access. Weak Role-Based Access Control (RBAC) policies can grant excessive permissions, enabling privilege escalation and lateral movement. This risk is often exacerbated by *insecure defaults* in open-source software [24] [25], which prioritizes usability over security by not enabling strict execution policies and strong authentication mechanisms. Without proper hardening, attackers can exploit misconfigurations to manipulate workloads, gain unauthorized access, and disrupt network operations.

T6 Software Vulnerabilities Software vulnerabilities can also arise from flaws in orchestration and network management software, which represent a significant share of the codebase. These weaknesses, such as *bugs in workflows and API implementations* [26], [27] and *vulnerable third-party dependencies* [28], can be exploited to compromise middleware security. These these flaws expose middleware resources to unintended access, allowing adversaries to intercept sensitive data.

C. Application-level Threats

The GENIO platform supports the deployment and execution of applications across its far-edge, edge, and cloud layers. Applications can expose other ones, and the GENIO platform itself, to security attacks.

T7 Vulnerable Applications Since applications are delivered by third-party business users, they can bring additional *application vulnerabilities*. Attackers can exploit such vulnerabilities to gain foothold on a tenant, and pursue malicious actions against users, other tenants, and the underlying platform. Application vulnerabilities arise from the lack of secure software development practices, such as static/dynamic analysis and reuse of insecure components. These issues can expose users to data breaches (e.g., through SQL injection) and injection attacks (e.g., Cross-Site Scripting). Moreover, attackers can gain access through command injection, deserialization, and memory corruption vulnerabilities, which can lead to remote code execution.

T8 Malicious Applications Malicious behaviors can arise both from exploited vulnerabilities (as previously discussed), and from *deliberately malicious applications*. For example, business users can reuse malicious container images from external repositories, which can contain hidden malware or backdoors. These untrusted applications can bypass scrutiny through obfuscation. These applications can execute malicious code to invoke privileged system calls and misusing capabilities (e.g., `CAP_SYS_ADMIN` in Linux containers), in order to escape container restrictions and disrupt the host and neighboring services. Moreover, malicious applications can attack the platform through *resource abuse*, by monopolizing CPU, memory, network, and storage resources, thereby degrading performance and causing service outages for other tenants.

IV. INFRASTRUCTURE-LEVEL MITIGATIONS

A. Mitigating Privilege Abuse

M1 OS environment configurations GENIO ensures a secure ONL Linux configuration, by stripping non-essential components (e.g., unused packages, services, kernel modules) to minimize the attack surface. Security policies are automated using Open-SCAP [29], which enforces SCAP benchmarks to secure SSH configurations, enable NTP synchronization, disable untrusted APT repositories, and protect kernel files. GENIO also aligns to Security Technical Implementation Guides (STIGs) [30], a set of best practices originally developed for Ubuntu and other mainstream Linux distributions, to enforce encryption policies, restrict system access, and secure boot configurations.

M2 OS kernel hardening At the kernel level, memory protections (e.g., `CONFIG_STACKPROTECTOR`) block buffer overflow attacks, while Linux Security Modules (AppArmor/SELinux) [31] restrict privileged system calls. High-risk functionalities like `KEXEC` (runtime kernel replacement) and `KPROBES` (debugging hooks) are disabled. The kernel-hardening-checker tool [32] validates configurations (kconfig, cmdline, sysctl)

Fig. 3. OSS security solutions and standards in GENIO.

against hardened baselines, and speculative execution mitigations (Intel/AMD microcode [33]) address side-channel vulnerabilities like Spectre [34].

> **Lesson 1.** The platform's reliance on Open Networking Linux (ONL) introduced complexities, as ONL lacks formal security guidelines compared to mainstream distributions. The application of STIGs and SCAP benchmarks was required to align with ONL's architecture, demanding iterative adjustments and reviews to balance security, performance, and compatibility.

B. Securing Communication

M3 End-to-End Encryption GENIO safeguards traffic across both Ethernet and PON segments through end-to-end encryption to prevent interception or tampering. At layer 2, the MACsec protocol standardized by *IEEE 802.1AE* [35] encrypts raw Ethernet frames using AES-GCM, providing authentication, confidentiality, and integrity for data on point-to-point Ethernet. In parallel, GENIO follows optical-specific guidelines

such as *ITU-T G.987.3* [36] for GPON, which recommend AES-based payload encryption to defend against fiber taps and low-level tampering in PONs.

M4 Authentication of Nodes GENIO enforces mutual authentication between ONUs and OLTs to verify hardware legitimacy before service provisioning. Certificate-based methods (via PKI) validate device identities, preventing rogue devices from impersonating legitimate infrastructure. Secure key exchange protocols (e.g., TLS 1.3) and secure DNS [37] prevent man-in-the-middle attacks during onboarding and registration, ensuring only trusted devices access network resources.

> **Lesson 2.** Encryption imposes additional engineering efforts and computational resources to enhance the security of the PON network. Implementing secure authentication among heterogeneous hardware (ONUs, OLTs, and cloud systems) demands careful management of certificates. GENIO's alignment with evolving ETSI standards [38] reflects an ongoing effort to maintain interoperability and compliance with industry guidelines.

C. Ensuring Code Integrity

M5 Secure Boot GENIO uses *Secure Boot* and a *Trusted Platform Module* (TPM) to cryptographically verify OS and firmware components before execution. At the earliest stage, the *Shim* [39] bootloader, signed by a recognized certificate authority (e.g., Microsoft), initializes a secure environment before loading the GRUB bootloader. By relying on Shim, the GENIO platform can add custom keys to validate later boot layers, including distribution-specific kernels. Additionally, *Measured Boot* records hashes of critical binaries in TPM Platform Configuration Registers (PCRs) at boot, enabling integrity checks against expected values. Together, these measures help ensure the platform boots from a known-good state and reveal any subsequent compromise.

M6 Secure Storage Beyond TPM-based firmware and OS verification, GENIO protects data at rest using *Linux Unified Key Setup* (LUKS) [40] to encrypt entire partitions with a passphrase. After encryption, the decryption key can be bound to specific PCR values in the TPM. If the measured environment (e.g., the kernel) matches the expected hash chain, the TPM releases the decryption secret; otherwise, access is denied. To automate this, GENIO plans to integrate *Clevis* [41], which seamlessly unwraps the LUKS key at boot when TPM-managed PCRs confirm system integrity. This enables booting without manual passphrase entry, reducing operational overhead in PON settings.

M7 File Integrity Monitoring Even with secure boot, adversaries may attempt to alter system files post-boot. GENIO deploys *Tripwire* [42] for runtime file integrity monitoring (FIM), creating cryptographic baselines of critical system files and alerting administrators to unauthorized changes. Tripwire's configurations and databases are encrypted and signed, with keys protected by the TPM to prevent tampering with the monitoring process.

Lesson 3. Deploying integrity protections in industrial environments faces obstacles. GENIO relies on older OS distributions (ONL Linux, based on Debian 10) to run SDN software and drivers, which lack native support for recent software packages. As a result, manually installing newer dependencies is required, introducing potential conflicts. Libraries required by Clevis for TPM access and automated disk decryption are unavailable, forcing manual passphrase entry at boot, which is impractical for in-field deployments of OLT nodes. Moreover, file monitoring should distinguish between critical resources that should not be mutable (e.g., system binaries and configurations) from mutable ones, to avoid misleading alerts.

D. Mitigating Software Vulnerabilities

M8 Automated Scanning GENIO conducts periodic vulnerability scanning with tools such as OpenSCAP [29], Lynis [43] and Vuls [44] to detect known CVEs across the low-level software, including the Linux kernel, system binaries, and user-space packages. These reports are prioritized based on severity and exploitability, ensuring that critical patches are applied as soon as feasible.

M9 Signed Updates Ensuring the authenticity and integrity of software updates is critical to thwarting supply-chain attacks within GENIO. The platform thus employs different methods tailored to each update scenario. In Debian-based environments, *user-space packages* are distributed via APT, which signs metadata and packages with GPG keys for each repository, and rejects any unverified artifacts. In addition, GENIO employs ONIE [45] for securely delivering ONL kernel updates. Following NIST SP 800-193 [46] guidelines, ONIE images are signed with X.509 certificates, accompanied by a detached signature file that is validated against a locally trusted public key, backed by a TPM. ONIE reboots the system into a minimal environment to apply the update, and fully run this environment by using Secure Boot, reducing potential inference from a compromised OS. Beyond kernel and user-space package updates, GENIO must also distribute additional binaries, such as specialized daemons and custom tools. These are also signed with GENIO's own certificates, which are likewise validated on each target node before installation.

Lesson 4. The maturity of automated scanning solutions facilitated smooth integration into GENIO's custom stack, even if occasional manual tuning is required to handle non-standard paths and configurations in ONL. APT GPG signatures for Debian-based images represent a reliable and straightforward solution to adopt.

V. MIDDLEWARE-LEVEL MITIGATIONS

A. Mitigating Privilege Abuse

M10 Access Control GENIO applies the principle of least privilege across its middleware stack, ensuring each role and service holds only the permissions necessary for legitimate operations. For virtualization and container management, native access control frameworks [47], [48] can mitigate abuses of resources exposed through their APIs. In network management software, including ONOS and VOLTHA, built-in authentication and authorization mechanisms [49] are configured to prevent API misuse. Exposed APIs for OLT and ONT management are strictly restricted to administrative service accounts secured by TLS certificates. On the network-side, GENIO enforces a clearly defined set of capabilities required in production, such as device registration, logical network configuration, and diagnostic logging—while blocking operations that introduce unnecessary privilege risks, such as direct shell access, low-level debugging endpoints, or raw log retrieval.

M11 Security Guideline Compliance GENIO adheres to industry-recognized security standards and continuously audits configurations to maintain compliance. It implements the NSA Kubernetes Hardening Guidance [50], CIS Benchmarks [51], and a suite of community tools, including docker-bench [52], kube-bench [53], kubesec [54], kube-hunter [55], and kubescape [56] to detect misconfigurations in Kubernetes clusters. Additionally, it follows vendor-specific guidelines for SDN-controllers provided by ONOS [57] to address insecure defaults, enforce strong authentication, and detect configuration drift.

Lesson 5. Hardening network management software is straightforward, as required capabilities are well-defined, and unnecessary functions can be blocked without disruption. In contrast, the configuration of RBAC policies for the orchestration platforms is challenging, since they are feature-rich and resource access should be carefully adapted for the workflows of the GENIO platform. Moreover, designers must integrate multiple security guidelines and checker tools, since individual solutions only address a subset of the risks.

B. Mitigating Software Vulnerabilities

M12 Automated Scanning and Patching GENIO integrates multiple sources to track vulnerabilities in middleware components. For Kubernetes, it leverages its official CVE database [58], which provides real-time updates on disclosed vulnerabilities, affected versions, exploitability, impact, and patches. The Kubernetes database offers a structured, programmatically accessible CVE feed for automated monitoring. Other middleware components vary in vulnerability tracking maturity. The Docker runtime publishes security updates [59] as blog-format announcements, making structured extraction difficult. ONOS maintains a structured web interface but is no longer actively updated. Proxmox notifies users only via its web UI. For middleware lacking structured, up-to-date, or programmatically accessible feeds, GENIO relies on the National Vulnerability Database (NVD) APIs to track vulnerabilities. To enhance precision in Kubernetes vulnerability tracking, GENIO integrates the Kubernetes Bill of Materials (KBOM) [60], which catalogs control plane services, node components, and add-ons with their exact versions and images, mapping known vulnerabilities in installed components.

VI. APPLICATION-LEVEL MITIGATIONS

A. Mitigating Software Vulnerabilities

M13 Container Security and SCA A key aspect of securing applications in GENIO is hardening containerized workloads. The platform uses Docker Bench for Security [52] to detect and fix misconfigurations that may introduce vulnerabilities. By enforcing best practices, such as least-privilege execution, restricted volume mounting, and secure networking, GENIO reduces the application attack surface. To address risks from third-party dependencies, GENIO integrates Software Composition Analysis (SCA) with tools like Trivy [61] and OWASP Dependency Check [62], scanning container images, identifying imported packages, and matching versions against CVE databases. These tools provide visibility into vulnerabilities introduced by pre-built components in the application stack.

M13 Static Application Security Testing GENIO applies SAST to detect vulnerabilities in the source code of the application itself. The container filesystem is extracted using Crane [63], allowing further scans to detect quality issues. Java source files are analyzed with SpotBugs [64], identifying issues such as null pointer dereferences, improper resource management, and inefficient exception handling. Pylint [65] serves as the Python counterpart. In addition, GENIO integrates Semgrep [66] and Bandit [67] to detect security vulnerabilities, such as hardcoded credentials, improper input validation, and weak cryptographic functions.

M15 Dynamic Application Security Testing GENIO integrates DAST to uncover runtime vulnerabilities. Specifically, GENIO employs CATS [68], a REST API fuzzer tool, to evaluate OpenAPI-defined endpoints. CATS conducts fuzz testing by injecting malformed, unexpected, and malicious inputs to identify vulnerabilities such as insecure input handling, improper authentication enforcement, and API misconfigurations. Beyond application-level testing, GENIO enforces network security checks when the application is deployed. Nmap [69] verifies TLS enforcement to ensure secure communication and analyzes port configurations, identifying unnecessary open ports that could expose to external threats.

Lesson 7. While SCA and SAST tools are mature and widely available, integrating them into GENIO poses challenges. SCA often flags unused or misidentified dependencies, generating noise and false positives. It also analyzes entire dependencies without linking vulnerabilities to specific functions used by the application, resulting to bloated reports and complicating path prioritization. Lastly, fuzzing containerized applications is feasible only for those exposing standard interfaces, such as REST APIs.

B. Identifying Malicious Applications

M16 Malware Signature GENIO defends against malicious applications using malware signature detection to proactively identify known malicious components before they are deployed or executed. To this end, GENIO utilizes Deepfence YaraHunter [70] to scan container images at rest for indicators of compromise. This tool leverages YARA rules to detect embedded malicious binaries, scripts, or configuration files.

M17 Isolation & Sandboxing To limit the impact of malicious applications, GENIO enforces strict runtime boundaries through isolation and sandboxing. It integrates KubeArmor [71] to restrict container, pod, and VM behavior at the system level using Linux Security Modules (LSMs), blocking unauthorized processes, file access, and suspicious network activity. To strengthen multi-tenancy isolation, GENIO follows the best practices from the PEACH framework [72], which models isolation risks based on interface complexity, tenant separation, and enforcement strength across key dimensions such as privilege, encryption, and authentication.

M18 Runtime Monitoring To enhance visibility into the behavior of running applications, GENIO integrates Falco [73]. Falco monitors system calls in real-time using eBPF and evaluates them against a rich, customizable rule set to detect suspicious behaviors such as unexpected shell execution, unauthorized file access, or unusual network connections. Unlike signature-based scanners or sandboxers, Falco provides deep runtime observability without blocking execution, enabling early detection of post-exploitation activities.

Lesson 8. Our experience with these tools confirms that these techniques are relatively mature and effective in detecting and isolating malicious applications. However, challenges remain in tuning policies and rules to minimize false positives without weakening security. Maintaining performance overheads within acceptable bounds is also a key consideration.

VII. CONCLUSION

Securing an edge computing platform demands to carefully address threats across all layers of the system. The GENIO project identifies and addresses key risks spanning infrastructure, middleware, and applications. Mitigation strategies were deployed using open-source tools and best practices, revealing both the strengths and limits of current security solutions.

ACKNOWLEDGMENTS

References

[1] E. Ahmed, A. Ahmed, I. Yaqoob, J. Shuja, A. Gani, M. Imran, and M. Shoaib, "Bringing computation closer toward the user network: Is edge computing the solution?" *IEEE Communications Magazine*, 2017.

[2] TIM, "Edge Cloud Computing," https://www.gruppotim.it/it/newsroom/notiziario-tecnico-tim/2022/n1-2022/cap03-edge-cloud-computing.html, 2022.

[3] Telefonica, "Telefónica Open Access and Edge Computing," https://www.telefonica.com/es/wp-content/uploads/sites/4/2021/02/whitepaper-telefonica-opa-mec-feb-2019.pdf, 2019.

[4] Y. Xiao, Y. Jia, C. Liu, X. Cheng, J. Yu, and W. Lv, "Edge computing security: State of the art and challenges," *Proceedings of the IEEE*, 2019.

[5] K. Gai, Y. Ding, A. Wang, L. Zhu, K.-K. R. Choo, Q. Zhang, and Z. Wang, "Attacking the edge-of-things: A physical attack perspective," *IEEE Internet of Things Journal*, 2021.

[6] P. Ladisa, H. Plate, M. Martinez, and O. Barais, "SoK: Taxonomy of attacks on open-source software supply chains," in *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2023.

[7] V. Varadharajan and U. Tupakula, "Securing services in networked cloud infrastructures," *IEEE Transactions on Cloud computing*, 2016.

[8] System Management SpA, "GENIO Project," https://sysmanagement.it/genio/, 2023.

[9] European Commission, "Cyber Resilience Act," https://digital-strategy.ec.europa.eu/en/policies/cyber-resilience-act, 2024.

[10] ——, "CE marking," https://single-market-economy.ec.europa.eu/single-market/ce-marking_en, 2024.

[11] CISA, "CISA Home," https://www.cisa.gov/, 2025.

[12] NIST, "NIST," https://www.nist.gov/, 2025.

[13] V. Casola, A. De Benedictis, M. Rak, and U. Villano, "A novel Security-by-Design methodology: Modeling and assessing security by SLAs with a quantitative approach," *Journal of Systems and Software*, 2020.

[14] H. Aranha, M. Masi, T. Pavleska, and G. P. Sellitto, "Enabling security-by-design in smart grids: An architecture-based approach," in *2019 15th European Dependable Computing Conference (EDCC)*. IEEE, 2019.

[15] F. M. Awaysheh, M. N. Aladwan, M. Alazab, S. Alawadi, J. C. Cabaleiro, and T. F. Pena, "Security by design for big data frameworks over cloud computing," *IEEE Transactions on Engineering Management*, 2021.

[16] A. Dutta and E. Hammad, "5G security challenges and opportunities: A system approach," in *2020 IEEE 3rd 5G world forum (5GWF)*. IEEE.

[17] Open Networking Foundation, "Open Network Operating System (ONOS)," https://opennetworking.org/onos/.

[18] ——, "Virtual OLT Hardware Abstraction (VOLTHA)," https://www.voltha.org/.

[19] Open Compute Project, "Open Networking Linux (ONL)," http://opennetlinux.org/.

[20] Proxmox Server Solutions GmbH, "Proxmox," https://www.proxmox.com/, 2025.

[21] OWASP, "Threat-Modeling-Process," https://owasp.org/www-community/Threat_Modeling_Process#stride, 2025.

[22] M. Z. Iqbal, H. Fathallah, and N. Belhadj, "Optical fiber tapping: Methods and precautions," in *8th international conference on high-capacity optical networks and emerging technologies*. IEEE, 2011.

[23] Kelly Jackson Higgins, "How Attackers Siphon Data In Targeted, APT Attacks," https://www.darkreading.com/cyberattacks-data-breaches/how-attackers-siphon-data-in-targeted-apt-attacks, 2011.

[24] Kiuwan, "The Top 15 Open-Source Software Security Risks," https://www.kiuwan.com/blog/open-source-software-security-risks/, 2024.

[25] CWE, "CWE-1188: Initialization of a Resource with an Insecure Default," https://cwe.mitre.org/data/definitions/1188.html, 2025.

[26] X. Ju, L. Soares, K. G. Shin, K. D. Ryu, and D. Da Silva, "On fault resilience of OpenStack," in *Proceedings of the 4th Annual Symposium on Cloud Computing (SOCC'13)*, 2013.

[27] D. Cotroneo, L. De Simone, P. Liguori, R. Natella, and N. Bidokhti, "How bad can a bug get? an empirical analysis of software failures in the OpenStack cloud computing platform," in *Proceedings of the 27th ACM European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2019.

[28] CWE, "CWE-1395: Dependency on Vulnerable Third-Party Component," https://cwe.mitre.org/data/definitions/1395.html, 2025.

[29] OpeSCAP, "OpenSCAP portal," https://www.open-scap.org/, 2025.

[30] STIGs, "Security Technical Implementation Guides," https://public.cyber.mil/stigs/, 2025.

[31] Linux Security Modules, "Linux Secuirty Modules," https://www.kernel.org/doc/html/v4.16/admin-guide/LSM/index.html, 2025.

[32] Kernel Hardening Checker, "Kernel Hardening Checker," https://archlinux.org/packages/extra/any/kernel-hardening-checker/, 2025.

[33] SCWorld, "Intel and AMD chips still vulnerable to Spectre flaw," https://www.scworld.com/news/intel-and-amd-chips-still-vulnerable-to-spectre-flaw, 2024.

[34] P. Kocher, J. Horn, A. Fogh, , D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, and Y. Yarom, "Spectre attacks: Exploiting speculative execution," in *40th IEEE Symposium on Security and Privacy (S&P'19)*, 2019.

[35] IEEE, "802.1AE: MAC Security (MACsec)," https://1.ieee802.org/security/802-1ae/, 2025.

[36] ITU, "ITU-T - G.987.3," https://standards.globalspec.com/std/14459900/g-987-3-amd-2, 2025.

[37] The Internet Society, "RFC 4033 - DNS Security," https://datatracker.ietf.org/doc/html/rfc4033, 2025.

[38] ETSI, "ETSI TS 103 962," https://www.etsi.org/deliver/etsi_ts/103900_103999/103962/01.01.01_60/ts_103962v010101p.pdf, 2023.

[39] SHIM, "UEFI BOOTLOADER," https://github.com/rhboot/shim.

[40] LUKS, "Linux Unified Key Setup," https://www.redhat.com/en/blog/disk-encryption-luks, 2025.

[41] Clevis, "Clevis," https://wiki.archlinux.org/title/Clevis, 2025.

[42] Tripwire, "Tripwire Open Source," https://github.com/Tripwire/tripwire-open-source, 2025.

[43] Lynis, "Lynis," https://cisofy.com/lynis/, 2025.

[44] Future Corp, "Vuls: Vulnerability Scanner," https://github.com/future-architect/vuls.

[45] ONIE, "Open Network Install Environment," https://opencomputeproject.github.io/onie/, 2025.

[46] NIST, "Platform Firmware Resiliency Guidelines," https://csrc.nist.gov/pubs/sp/800/193/final, 2025.

[47] Kubernetes, "RBAC Kubernetes," https://kubernetes.io/docs/reference/access-authn-authz/rbac/, 2025.

[48] Proxmox Server Solutions GmbH, "Proxmox Access Control Framework," https://pve.proxmox.com/proxmox/pve-access-control, 2025.

[49] ONOS, "ONOS Security," https://wiki.onosproject.org/display/ONOS10/ONOS+Security%3A+Security-mode+ONOS, 2025.

[50] CISA, "NSA Kubernetes Hardening Guidance," https://www.cisa.gov/news-events/alerts/2022/03/15/updated-kubernetes-hardening-guide, 2025.

[51] CIS, "CIS Benchmarks List," https://www.cisecurity.org/cis-benchmarks, 2025.

[52] Docker Bench Security, "Docker Bench Security," https://hub.docker.com/r/docker/docker-bench-security, 2025.

[53] Kube-bench, "Kube-bench," https://github.com/aquasecurity/kube-bench, 2025.

[54] Kubesec, "Kubesec," https://kubesec.io/, 2025.

[55] Kube Hunter, "Introducing kube-hunter: an Open Source Tool for Discovering Security Issues in Kubernetes Clusters," https://github.com/aquasecurity/kube-hunter, 2025.

[56] Kubescape, "Kubescape," https://kubescape.io/.

[57] Open Network Operating System, "Onos security & performance analysis (report no. 1)," https://hal.science/hal-03188701v1.

[58] Kubernetes, "Official CVE Feed," https://kubernetes.io/docs/reference/issues-security/official-cve-feed/, 2025.

[59] Docker, "Docker Forum," https://forums.docker.com/t/docker-images-and-security-updates/134504, 2025.

[60] KBOM, "KBOM," https://github.com/rad-security/kbom.

[61] Trivy, "Trivy," https://trivy.dev/latest/, 2025.

[62] OWASP, "OWASP Dependency-Check," https://owasp.org/www-project-dependency-check/, 2025.

[63] Crane, "Crane," https://github.com/google/go-containerregistry/tree/main/cmd/crane, 2025.

[64] Spotbugs, "Spotbugs," https://spotbugs.github.io/, 2025.

[65] Pylint, "Pylint," https://www.pylint.org/, 2025.

[66] Semgrep, "Semgrep," https://github.com/semgrep/semgrep, 2025.

[67] Bandit, "Bandit," https://bandit.readthedocs.io/en/latest/, 2025.

[68] CATS, "CATS," https://github.com/Endava/cats, 2025.

[69] nmap, "nmap," https://nmap.org/, 2025.

[70] YaraHunter, "YaraHunter," https://community.deepfence.io/docs/yarahunter/, 2025.

[71] KubeArmor, "KubeArmor," https://kubearmor.io/, 2025.

[72] Peach, "Peach," https://www.peach.wiz.io/, 2025.

[73] Falco, "Falco," https://falco.org/, 2025.