PROCEEDINGS OF SCIENCE

A comprehensive initiative to enhance the security posture of open-source software

Marica Antonacci,[a,*] **Vincenzo Ciaschini,**[b] **Giacinto Donvito**[a] **and Barbara Martelli**[b]

[a]*Istituto Nazionale di Fisica Nucleare (INFN),*
Via E. Orabona 4, 70125 Bari, Italy

[b]*Istituto Nazionale di Fisica Nucleare,*
Viale Berti Pichat 6/2, 40127 Bologna, Italy

E-mail: marica.antonacci@ba.infn.it, vincenzo.ciaschini@cnaf.infn.it, giacinto.donvito@ba.infn.it, barbara.martelli@cnaf.infn.it

In the dynamic landscape of digital security, safeguarding information assets stands as a paramount concern for organizations. This paper presents a comprehensive initiative undertaken by INFN, a prominent player in research, to bolster the security posture of its open-source components within the DataCloud production middleware. Central to this initiative is the recognition of the pivotal role security plays in the software development lifecycle (SDLC). The paper outlines INFN's strategic approach to align with industry standards such as OWASP SAMM and ISO/IEC 27002 frameworks. Through collaboration and proactive measures, INFN aims to establish virtuous processes aimed at enhancing security governance, self-assessment, continuous monitoring, and timely responses to emerging vulnerabilities. The ultimate goal is to cultivate a more secure and resilient software ecosystem tailored to scientific data analysis.

International Symposium on Grids and Clouds (ISGC2024)
24 -29 March, 2024
Academia Sinica Computing Centre (ASGC), Institute of Physics, Academia Sinica Taipei, Taiwan

[*] Speaker

PoS(ISGC2024)012

1. Introduction

In today's digital landscape, safeguarding information assets has become critical for organizations of all sizes and industries. Information security is not just a technical concern but a fundamental aspect of organizational resilience and reputation protection. The consequences of inadequate information security can range from financial losses and reputational damage to regulatory non-compliance and legal repercussions.

This is particularly crucial in sectors such as research and academia, where the integrity and confidentiality of data are pivotal for trust and credibility.

As a key actor in the research domain, INFN is deeply committed to safeguarding critical information. This commitment is underscored by its extensive involvement in distributed computing infrastructures and different research projects, including those concerning health and other confidential data.

In fact, since 2017 INFN is running a highly secure, ISO/IEC 27001 certified computing infrastructure dedicated to research communities dealing with personal and particular data which need to be managed according to GDPR regulation. Such data include genomics and omics, medical images (MRI, PET, TC) and Electronic Health Records (EHR) shared among research hospitals and organizations to conduct multi-centered studies. As an example, it is worth mentioning some of the main projects INFN is participating in:

- Spoke 8 "In Silico Medicine and Omics Data" in the context of the ICSC National Research Centre for High Performance Computing, Big Data and Quantum Computing [2]
- HBD (Health Big Data) [3]
- DARE (DigitAl lifelong pREvention) [4]
- Elixir [5] Joint Research Unit In the context of these projects, INFN brings his expertise in developing high-performance and high-throughput distributed computing platforms and has the responsibility to develop a GDPR-compliant federated cloud platform enabling research communities to access relevant applications in IaaS, PaaS and SaaS deployment models and to share data in a secure and lawful way.

Recognizing the critical importance of protecting such information, INFN is willing to prioritize the implementation of strong security measures to ensure the integrity and confidentiality of entrusted data.

We are currently addressing these challenges as part of our software development activities in the framework of the DataCloud project. The DataCloud project is a strategic initiative aimed at creating a comprehensive portfolio of IaaS and PaaS cloud services [1], [6]. These services help to accelerate the deployment of innovative solutions by offering seamless access to distributed computing and storage resources.

The main goals of the DataCloud project are twofold. First, it aims to empower researchers by allowing them to use cutting-edge technologies and resources for their projects. Second, the project aims to promote collaboration and knowledge exchange within the research community. By providing a shared platform for infrastructure and tools, DataCloud promotes collaboration among researchers, enabling them to share resources, expertise, and insights.

The software developed within the DataCloud project stands out for its unique focus on meeting the needs of research communities. It provides tailored solutions and cloud-native applications designed specifically to support the unique requirements of research groups.

In the DataCloud project, we embrace a co-design approach, facilitating collaboration between researchers and developers to customize solutions according to users' needs. This collaborative process ensures that the development is shaped by the domain expertise of the users, making their requirements and concerns integral to the design of solutions. By actively involving users in the development process, DataCloud ensures that the resulting solutions are finely tuned to address the specific challenges and requirements encountered in research projects, enhancing usability, efficiency, and user satisfaction.

2. Frameworks for security enhancement

In recent years, the rapid pace of software development has often prioritized swift feature releases over security considerations. Recognizing this historical oversight, we are committed to restoring security as a fundamental aspect of our development lifecycle.

Our current strategic security initiative is focused on evaluating and enhancing the security posture of the open-source components required to build the DataCloud middleware.

Our approach entails a collaborative effort among development leaders, designated security champions, and key stakeholders to establish robust governance frameworks and processes. This initiative encompasses tasks aimed at defining security standards, policies, assigning responsibilities, implementing security training, and fostering awareness programs.

To guide our efforts, we have chosen to adopt the OWASP SAMM (Software Assurance Maturity Model) framework [7] and adhere to the ISO 27K international standard [8]. We believe that the synergy between these frameworks forms the cornerstone of a resilient and unified security posture.

OWASP SAMM is a well-established framework tailored to fortify software security, providing organizations with a structured approach to assess, enhance, and benchmark their security practices. Organized into five security functions (See Figure 1), SAMM offers maturity levels that steer organizations towards a more robust software security posture: it supports the complete software lifecycle and is technology and process agnostic.

Conversely, ISO/IEC 27001 and 27002 focus on information security management, furnishing a comprehensive array of controls and guidelines for effective security governance.

By leveraging both frameworks, we align our software security practices with SAMM's principles while implementing the information security controls outlined in ISO/IEC 27002.

In addition to these frameworks, we are exploring other valuable resources tailored for open-source projects.

The OpenSSF (Open Source Security Foundation) [9] is a collaborative initiative under the Linux Foundation's umbrella, dedicated to enhancing the security of open-source software. Through collaboration with industry stakeholders, it promotes the adoption of secure development practices and mitigates security risks in open-source software.

The OpenSSF Scorecard [10] offers a comprehensive assessment tool to evaluate a project's security posture across various dimensions. Similarly, the OpenSSF Best Practices Program Badge

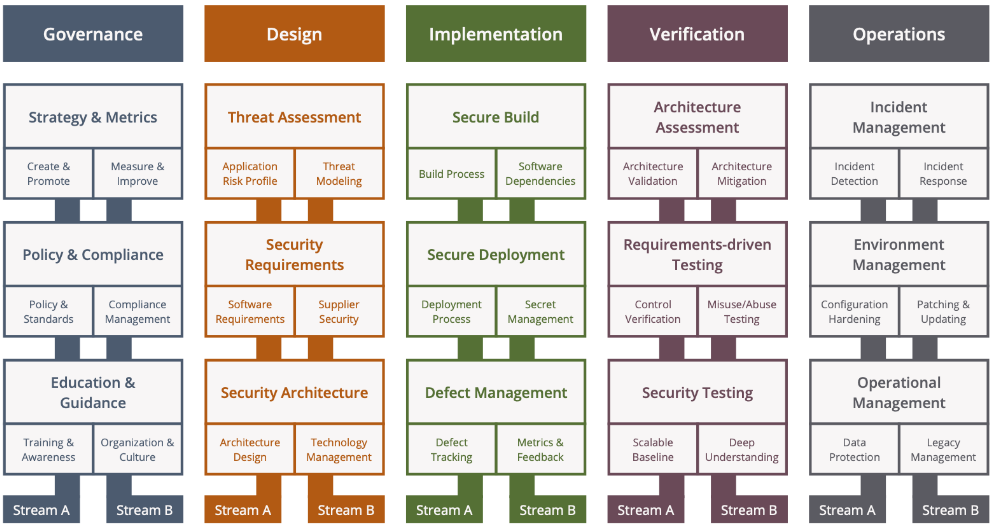

Figure 1: Software Assurance Maturity Model (SAMM) Framework - version 2

[11] recognizes projects adhering to security best practices outlined by the Open Source Security Foundation.

Additionally, the Cloud Native Computing Foundation (CNCF) Security Technical Advisory Group (STAG) [12] plays a pivotal role in advancing security within cloud-native environments. As an integral part of the CNCF ecosystem, the STAG is responsible for identifying security challenges and developing best practices, guidelines, and recommendations to address them effectively. Comprising industry experts, security practitioners, and thought leaders, the STAG collaborates closely with CNCF project maintainers, end-users, and the broader community to promote security awareness, drive innovation, and ensure the integrity and resilience of cloud-native applications and infrastructure.

These resources offer specialized focus on open-source projects and cloud-native security. By leveraging the OpenSSF Best Practices Program Badge and Scorecard, projects can demonstrate their commitment to security and enhance credibility within the open-source community.

3. Our strategy towards a robust security posture

We have developed a detailed work plan with the following objectives:

1. Establishing secure code development processes: our primary goal is to establish formal processes that prioritize risk management in code development. This involves creating robust procedures for secure coding practices.

2. Aligning existing software with secure development strategy: we aim to align our current software projects with the newly established secure development strategy to ensure consistency and adherence to security standards.

To achieve these objectives, we have defined three work packages:

- WP1: Security Policy Development and Training.

  In this package, we are focused on defining, reviewing, and effectively communicating our security policies and standards. It is crucial that these policies align with the ISO/IEC 27002 framework, ensuring compliance with recognized industry standards. Additionally, we are committed to develop a comprehensive security training program for our team members.

- WP2: Security Self-Assessment and Supply Chain Analysis

  This package involves conducting a Security Self-Assessment to identify potential vulnerabilities, assess adherence to security policies, and promote continuous improvement in our security posture. Each project team will engage in a reflective process, analyzing current practices, documenting findings, and developing actionable plans to address weaknesses. This technical introspection will focus on practical insights, risk prioritization, and the implementation of effective corrective actions. Additionally, we are undertaking a supply chain analysis, leveraging the Software Bill of Materials (SBOM) framework, to gain insights into the composition of our software supply chain and identify any associated security risks with third-party components.

- WP3: Ongoing Software Quality Enhancement and Vulnerability Response

  The objective of this package is to establish and maintain ongoing processes aimed at enhancing software quality, updating dependencies promptly, and responding swiftly to emerging vulnerabilities. Security is not a one-time task; hence, we emphasize continuous consideration of security throughout our processes. This proactive approach aims to create a dynamic and adaptive security posture, fostering resilience against evolving threats throughout the software's lifecycle.

  Furthermore, our approach emphasizes the systematic collection of metrics and the provision of insightful feedback to stakeholders. We will generate a regular (e.g., monthly) report tailored for a suitable audience, incorporating key security metrics, vulnerability assessments, and remediation efforts. This report serves as a valuable input for refining our security strategy, guiding improvements in training programs and security verification activities.

  Additionally, we believe in sharing the most prominent or interesting technical details about security defects, along with the strategies employed to address them, with other teams. This knowledge-sharing will occur through regular meetings or dedicated forums, fostering cross-team collaboration and collective learning. By transparently sharing insights and experiences, we will empower teams to collectively enhance their understanding of security best practices and contribute to the overall resilience of our software ecosystem.

3.1 Evolving from DevOps to DevSecOps: strengthening security in development

By shifting security lift, organizations can proactively identify and mitigate potential vulnerabilities, reducing the risk of security breaches and minimizing costs associated with addressing issues later in the development process.

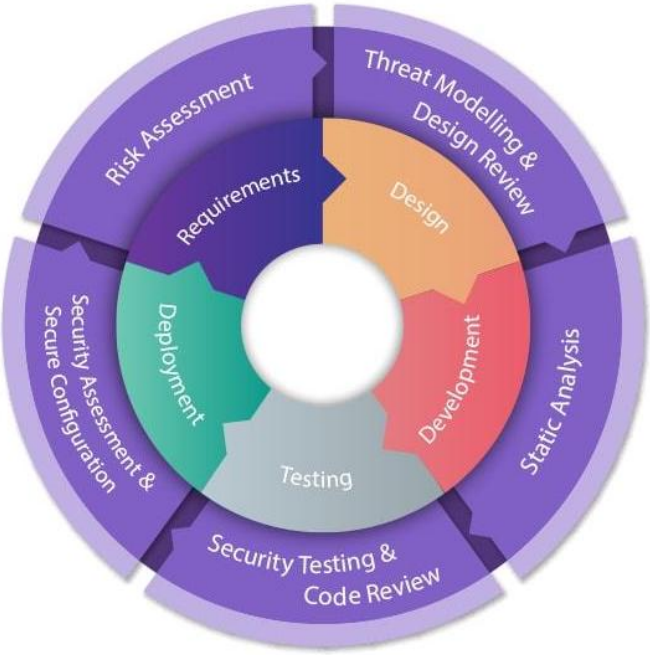
Risk Assessment
Threat Modelling & Design Review
Static Analysis
Security Testing & Code Review
Security Assessment & Secure Configuration
Requirements
Design
Development
Testing
Deployment

Figure 2: Secure Software Development Life Cycle [13]

We aim to implement security guardrails at each stage of our software development cycle.

Security requirements are integral to our software development projects. It is important to conduct thorough risk assessments to identify potential security threats and vulnerabilities within the application, determine the types of sensitive information handled by the application, identify relevant legal and regulatory requirements, and analyze each use case from a security perspective.

During the design stage, it is crucial to conduct threat modeling exercises to systematically identify potential threats and vulnerabilities within the application. We define a comprehensive security testing strategy to validate the effectiveness of security controls and mechanisms implemented during the design phase.

During the implementation stage, it is essential to enforce secure coding practices, conduct regular code reviews, and implement security controls identified during the design phase to minimize security risks.

Finally, it is vital to conduct security testing to validate the effectiveness of implemented security measures and identify any remaining vulnerabilities before deployment.

We are shifting from a pure DevOps [14] approach to DevSecOps [15], which will allow us to automate a wide range of security tasks throughout the SDLC.

DevSecOps embodies the ethos of continuous security testing, automated code analysis, and infrastructure fortification. This proactive approach ensures that security concerns are proactively addressed from the outset, enabling the swift and secure delivery of software without compromising quality or reliability.

Within our DevSecOps automated pipeline (see Figure 3), we harness cutting-edge development tools and methodologies to uphold these principles. Our journey begins with integrated development environments (IDEs) such as Visual Studio Code, enhanced with devcontainers [16]. These configurations come pre-loaded with essential security and software quality assurance (SQA) tools like linters and formatters, empowering developers to address security issues in real-time during code creation.

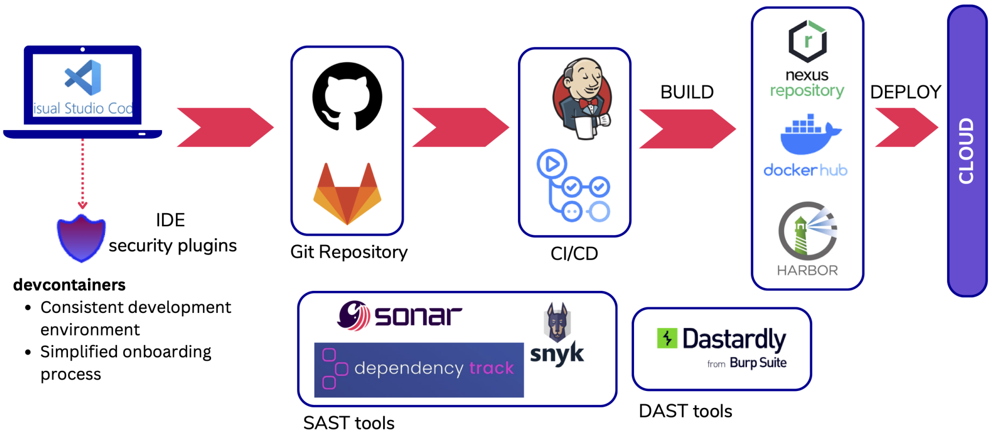

Figure 3: CI/CD pipeline integrating security tools like SAST and DAST.

Our code repositories, hosted on platforms like GitHub or GitLab, are managed through meticulously structured pipelines orchestrated by tools like Jenkins [17] or GitHub Actions [18]. These pipelines execute a battery of tests, encompassing static code analysis via SAST (Static Application Security Testing) tools like SonarQube [19], and dynamic code analysis using DAST (Dynamic Application Security Testing) tools such as Dastardly from Burp Suite [20].

Furthermore, the integration of Dependency-Track from OWASP [21], an intelligent Component Analysis platform, serves to vigilantly monitor and manage dependencies, while Snyk [22] aids in identifying and rectifying vulnerabilities within third-party dependencies. Dependency-Track generates a Software Bill of Materials (SBOM) in CycloneDX format [23], seamlessly uploaded via API from our CI/CD system. The resulting artifacts from the build process are securely stored in a Nexus repository [24], hosted on premise, with container images safeguarded in registries like Harbor [25] (private registry) or Docker Hub (public registry).

Prior to deployment into production environments, rigorous testing in pre-production environments ensures the software's stability and security. This holistic DevSecOps approach empowers us to deliver software with confidence, knowing it meets stringent security standards.

While tools are indispensable, they're only part of the equation. Without a foundational understanding of security principles, information gleaned from security tools may seem cryptic to developers. Thus, cultivating a comprehensive awareness of information security is imperative.

Security must be ingrained as a collective responsibility, transcending individual roles. It's about fostering a culture where every team member comprehends their role in upholding security standards, fostering collaboration, and promoting diligence in addressing potential security vulnerabilities.

In this context, security champions play a pivotal role. They serve as advocates for security awareness, fostering collaboration across cross-functional teams. Moreover, they ensure that

DevOps practices don't solely prioritize speed and efficiency, but also integrate robust security measures at every stage of the development lifecycle. Through their guidance, we ensure that security remains a paramount consideration in our software development endeavors.

4. Conclusions

This paper presents a comprehensive strategy for enhancing information security within the context of the INFN DataCloud project.

By prioritizing security measures and adopting frameworks such as OWASP SAMM and ISO/IEC 27001/27002, we aim to fortify our software development practices and safeguard critical information assets.

The transition from DevOps to DevSecOps underscores the importance of integrating security throughout the development lifecycle, with a focus on automation, collaboration, and continuous improvement. Moreover, the involvement of security champions and the promotion of security awareness among team members are crucial elements in fostering a culture of security within the organization. By adhering to these principles, we seeks to contribute to the creation of a secure software ecosystem that ensures the integrity, confidentiality, and resilience of research data and infrastructure.

References

[1] INFN Cloud | Cloud Services for Research [accessed 2024 Apr 12]. https://www.cloud.infn.it/

[2] ICSC Spoke8 - In Silico Medicine and Omics Data https://www.supercomputing-icsc.it/spoke-8-in-silico-medicine-omics-data/

[3] Health Big Data https://www.healthbigdata.it/

[4] DigitAl lifelong pREvention https://www.fondazionedare.it/

[5] Elixir Europe https://elixir-europe.org/

[6] M. Antonacci and D. Salomoni, *Leveraging TOSCA orchestration to enable fully automated cloud-based research environments on federated heterogeneous e-infrastructures*, Proceedings of the International Symposium on Grids & Clouds (ISGC) 2023 in conjunction with HEPiX Spring 2023 Workshop (ISGC&HEPiX2023), Volume 434, Virtual Research Environment (VRE), October 25, 2023, DOI: 10.22323/1.434.0020, Full text: https://pos.sissa.it/434/020/pdf.

[7] OWASP SAMM Version 2 [accessed 2024 Apr 12]. https://drive.google.com/file/d/1cI3Qzfrly_X89z7StLWI5p_Jfqs0-OZv/view?usp=sharing

[8] ISO/IEC 27000 family Information security management [accessed 2024 Apr 12]. https://www.iso.org/standard/iso-iec-27000-family

[9] Open Source Security Foundation – Linux Foundation Projects [accessed 2024 Apr 12]. https://openssf.org/

[10] Build better security habits, one test at a time [accessed 2024 Apr 12]. https://securityscorecards.dev/

[11] OpenSSF Best Practices Badge Program [accessed 2024 Apr 12]. https://www.bestpractices.dev/en

[12] CNCF Security Technical Advisory Group [accessed 2024 Apr 12]. https://github.com/cncf/tag-security

[13] Secure Software Development Lifecycle (SSDLC) [accessed 2024 Apr 12]. https://snyk.io/learn/secure-sdlc/

[14] DevOps [accessed 2024 Apr 12]. https://devopedia.org/devops

[15] Parveen Bhandari. What is DevSecOps? | The Ultimate Guide [accessed 2024 Apr 12]. https://www.xenonstack.com/insights/what-is-devsecops

[16] Development Containers [accessed 2024 Apr 12]. https://containers.dev/

[17] Jenkins website [accessed 2024 Apr 12]. https://www.jenkins.io/

[18] GitHub Actions documentation [accessed 2024 Apr 12]. https://docs.github.com/en/actions

[19] SonarQube [accessed 2024 Apr 12]. https://www.sonarsource.com/products/sonarqube/

[20] Dastardly, from Burp Suite [accessed 2024 Apr 12]. https://portswigger.net/burp/dastardly

[21] OWASP Dependency-Track [accessed 2024 Apr 12]. https://owasp.org/www-project-dependency-track/

[22] Snyk website [accessed 2024 Apr 12]. https://snyk.io/

[23] OWASP CycloneDX Software Bill of Materials (SBOM) Standard [accessed 2024 Apr 12]. https://cyclonedx.org/

[24] Sonatype® Nexus Repository OSS [accessed 2024 Apr 12]. https://www.sonatype.com/products/sonatype-nexus-repository

[25] Harbor [accessed 2024 Apr 12]. https://goharbor.io/