Contribution ID: **129** Type: **Oral Presentation**

AI-Assisted DevSecOps for Secure Software Delivery in Earth Observation Projects

Wednesday, 18 March 2026 16:00 (22 minutes)

Modern Earth Observation (EO) platforms integrate diverse distributed components and scientific workflows across heterogeneous cloud environments. Ensuring software security, maintainability and rapid delivery within such complex systems represents a major operational challenge. To address this, we developed an AI-assisted DevSecOps framework that augments continuous integration and deployment (CI/CD) pipelines with Large Language Model (LLM) capabilities for automated vulnerability detection, remediation and testing.

The framework extends a GitLab–Dagger–FluxCD toolchain with an agentic AI layer composed of three coordinated components: (i) container-level vulnerability analysis using Trivy [1] combined with LLM-generated hardening suggestions; (ii) source-level security and quality inspection using SonarQube [2] enriched with LLM-based corrective patches; and (iii) automated enhancement of unit tests through LLM-guided generation and refinement. All inference is performed on-premises on an NVIDIA GPU-accelerated cloud instance, where a local deployment of LM Studio [3] exposes an OpenAI-compatible interface secured via NGINX, ensuring full confidentiality of code and security findings.

This approach enables autonomous remediation cycles: vulnerabilities detected during CI trigger the generation of secure Dockerfile updates, code fixes or new test cases, which are validated and reintegrated into the workflow. The solution has demonstrated a substantial reduction in remediation time, improved test coverage and increased consistency of security practices across distributed platform components. The system is deployed using Dagger [4] for reproducible pipeline logic and FluxCD [5] for GitOps-based cluster reconciliation.

Future developments include extending support to additional programming languages, integrating unified security insights across all platform services, optimising on-prem GPU inference performance and exploring policy-driven hardening and runtime anomaly detection techniques.

This work demonstrates how AI-enhanced DevSecOps can strengthen the reliability and security of scientific platforms while maintaining full data sovereignty, aligning with ISGC's focus on AI-enabled scientific workflows and advanced operational practices.

References:

[1] https://aquasecurity.github.io/trivy/
[2] https://www.sonarsource.com/products/sonarqube/
[3] https://lmstudio.ai/
[4] https://dagger.io/
[5] https://fluxcd.io/

Primary authors: PISA, Claudio; Dr FORNARI, Federico (ECMWF); ANTONACCI, Marica (ECMWF (on leave from INFN)); ALBUGHDADI, Mohanad (ECMWF); KAPROL, Tolga (ECMWF); BAOUSIS, Vasileios (ECMWF)

Presenter: Dr FORNARI, Federico (ECMWF)

Session Classification: Artificial Intelligence (AI) - II

Track Classification: Track 10: Artificial Intelligence (AI)