

THE MOSAICO CHRONICLES



Giorgos Pitsiladis 

Project Manager, Research & Innovation Development
at Netcompany

Continuous integration: The backbone of collaborative AI ecosystems

Who we are & our role in MOSAICO

At Netcompany, we specialise in designing and delivering large-scale digital platforms that integrate complex systems into production-ready solutions. In the context of the MOSAICO Project, we act as platform integrators, responsible for ensuring seamless integration and interoperability across diverse AI agents, collaboration frameworks, agent repositories and governance services.

The MOSAICO vision of collaborative AI-agent ecosystems introduces significant technical challenges, including:

- Multi-agent coordination across heterogeneous agents.
- Standardised communication through tailor-made MOSAICO protocols.
- Agent classification and benchmarking.
- Policy enforcement across distributed components.

Our contribution focuses on delivering this integration layer and connecting all core building blocks of MOSAICO, from agent repositories and orchestration engines to governance mechanisms, into a fully functional and deployable platform. This platform is designed to be demonstrated in real-world use cases and made available as open source to the broader AI community.

The integration challenge: How do we bring diverse agents and protocols into a single DevOps platform?

A core challenge in MOSAICO is how to integrate a wide variety of AI agents and services into a single, coherent platform. To address this, MOSAICO adopts a DevOps-driven approach where each agent is treated as an independently versioned, containerised service integrated through a shared CI/CD pipeline. This abstracts implementation differences and enables consistent deployment and operation across heterogeneous components.

Interoperability is ensured through the standardisation of both agent-to-agent and agent-to-service communication. Each agent exposes well-defined APIs based on MOSAICO's tailor-made protocols, which are version-controlled and continuously tested to prevent breaking changes and maintain system-wide compatibility.

Integration quality is enforced through automated pipelines executing integration tests and ensuring that updates to individual agents do not disrupt orchestration workflows or other components. As a result, integration becomes a continuous and controlled process. Security is embedded into the platform through a unified identity and access model, combined with network isolation mechanisms (e.g. VPN), allowing agents to operate across different environments including on-premise and cloud-based infrastructures without compromising system integrity. Finally, centralised monitoring (logs, health status, resource usage) provides a unified view of agent behavior, enabling efficient debugging, performance optimisation and overall system reliability.

Continuous integration (CI/CD): How the infrastructure supports the rapid evolution and testing of AI agents

In MOSAICO, continuous integration and delivery (CI/CD) enables the rapid evolution of AI agents while maintaining system stability. Every change to an agent automatically triggers a standardised pipeline that builds, tests, and prepares it for deployment, ensuring fast iteration within a consistent integration process. These pipelines handle core tasks such as validation, containerisation, testing and versioning, making each agent release traceable. A shared pipeline structure enforces consistent quality checks across all agents, regardless of their implementation.

Beyond standard testing, MOSAICO extends CI/CD to cover specific needs related to AI:

- **Automated evaluations** on curated datasets and regression checks monitor both model performance and non-functional aspects such as latency and resource usage, ensuring controlled evolution of agent behavior.
- The process supports deployment across diverse environments, including **cloud and on-premises infrastructures**, without fragmenting the development lifecycle.
- Before release, new versions are validated through **staged rollout strategies**, allowing them to run alongside existing ones and ensuring compatibility within multi-agent workflows.

Overall, CI/CD acts as a key enabler of MOSAICO continuous innovation through balancing rapid experimentation with system reliability.

From lab to industry: The process of deploying the platform across our 4 distinct pilot sites

A key objective of MOSAICO is to ensure that innovations developed in controlled environments can be reliably transferred to real-world settings. This is achieved through a structured deployment process that promotes AI agents and deploys the MOSAICO platform from the central integration and deployment environment to four diverse use cases:

- Generation of energy-efficient 3D content and code for Mixed Reality (MR) devices.
- Software for personalised investment recommendations in the banking and finance sector.
- Development assurance processes for safety-critical software implementing Remaining Useful Life (RUL) estimation in aerospace.
- Reuse and co-evolution in the development of an IoT catalogue.

Once an agent has successfully passed integration and system-level validation, it is marked as ready for pilot deployment and introduced into a staging environment that mirrors real operational conditions. This enables early validation under realistic constraints before full rollout. All pilot sites rely on a common baseline of platform components, ensuring consistency across deployments, while applying site-specific configurations such as data sources, policies and operational thresholds to adapt agent behavior to local technical, regulatory, and business requirements. In some cases, additional specialised solution agents are developed to better align with pilot needs.

Deployment follows an incremental adoption strategy. Agents are first introduced in a monitoring or decision-support mode, allowing stakeholders to evaluate their outputs without affecting operations. This is followed by user validation and feedback collection, combining qualitative insights with performance metrics. Only after this validation phase do agents transition to more autonomous roles, operating under clearly defined guardrails.

Throughout this process, continuous monitoring and feedback loops are maintained across all pilot sites. Insights gathered during deployment are fed back into the core platform, enabling iterative improvements and supporting the onboarding of new agents or capabilities. This structured approach ensures that the transition from lab to industry is both technically robust and aligned with real-world operational and business needs.

Impact: What this will change for the final users?

For final users, the main change brought by MOSAICO is that AI agents will not be used as isolated commercial assistants, but as continuously integrated components within a controlled software engineering platform. Through the CI/CD process, every agent update is automatically built, tested, containerised, validated and prepared for deployment, ensuring that new AI capabilities can evolve quickly without compromising the stability of the overall multi-agent system or the end-user experience. This is a key difference from commercially available AI tools, where users often receive outputs without full control over versioning, integration quality, deployment environment, traceability or system-wide compatibility.

- **Software engineers and developers** using the MOSAICO platform will benefit from faster and safer adoption of AI agents inside their existing engineering workflows. Instead of manually testing whether an agent update works in practice, developers will be supported by pipelines that automatically validate changes, run checks and detect breaking changes before deployment.
- **DevOps teams and platform integrators** will gain a repeatable process for integrating heterogeneous agents into a single operational environment. CI/CD will support source control, code versioning, automated build-test-deploy processes, containerisation, Docker image management, logging, health monitoring and deployment automation. This makes MOSAICO easier to scale and maintain compared to disconnected commercial AI tools.

- **Use case providers** will be able to access MOSAICO either online, in a cloud sandbox, or through integration with their own IT systems, or through local deployment on their own workstations. This flexibility allows the platform to be tested and adopted under different operational, security and infrastructure constraints.
- **Open-source developers and future adopters** will benefit from a reusable CI/CD-based integration model aligned with Eclipse community practices. This helps MOSAICO become more than a project-specific platform: it becomes a reusable way to onboard, package and deploy AI agents in a transparent and sustainable ecosystem.



FOLLOW US!