Alessandro Pomponio

cv@alessandropomponio.me | alessandropomponio.me | linkedin.com/in/alessandropomponio |
github.com/alessandropomponio | medium.com/@AlessandroPomponio

Experience

Research Software Engineer April 2022 – Present
IBM Research *Dublin, IE*

- Became a reference for the lab's software development standards, introducing DevSecOps practices, tightening style rules and reviews, cutting CI test time by 50% via parallel testing, and adding end-to-end coverage with Testcontainers. Maintainer of open-source projects developed by my team.
- Co-administered two bare-metal OpenShift clusters for 100+ researchers. Designed a policy-driven governance approach using Kyverno, Argo CD, and Kueue to automate best practices and improve the admin/researcher experience.

Education

Alma Mater Studiorum Università di Bologna Bologna, IT
Master of Science in Computer Engineering, with honors *March 2022*

Projects

Accelerated Discovery Orchestrator (ado) | *Python, Ray, Pydantic, FastAPI, Typer* Aug. 2025 – Present

- Extensible framework for executing benchmarks and computational experiments at scale.
- Built a "human-first" CLI to ensure strong validation and smooth user experience. Evolved it to support agent-driven development flows (experiment development, execution, and analysis of results).
- Implemented end-to-end tests against real containerized dependencies (using Testcontainers) for both the ado core package and the Typer-based CLI using pytest.
- Built a DevSecOps pipeline and additional automation for dependency updates and changelog generation.

Simulation Toolkit for Scientific Discovery (ST4SD) | *Python, Flask, JavaScript, Vue.js* Dec. 2022 – Present

- Toolkit for developing, running, and sharing of virtual experiments.
- Created a Web UI for accessing, running, and monitoring ST4SD virtual experiments.
- Developed a thin Flask-based REST API for accessing backend services.
- Mentored 6 junior engineers through issues, design reviews, and PR feedback.

Talks & Community

Taming the Wild West of Research Computing: How Policies Saved Us a Thousand Headaches
· *Kube.fm Podcast* *August 2025*
· *Cloud Native Dublin* *August 2025*
· *KubeCon + CloudNativeCon NA* *November 2025*

- Presented how to use a GitOps approach for self-service namespace management, automating the delivery of group memberships, quotas, and more via Argo CD and Kustomize (bases + per-project overlays, ApplicationSets).
- Explained how to use Kyverno policies to restrict the use of interactive pods wasting GPU resources and automatically keep CPU-only pods off GPU nodes by adding affinity rules, preventing common anti-patterns.
- Described using Kueue to set up an HPC-style system familiar to researchers to guarantee fair-share GPU scheduling by defining GPU flavors and queues (including an "any GPU" queue with borrowing).

Patents

Database query optimization based on analytics: US12353414B2 (Granted)
Reaction-aware adaptive intervention in a monitored area: US20250054376A1 (Published)
Automating a configuration of an infrastructure for cloud applications: US20240356807A1 (Published)

Technical Skills

Cloud / Platform: Docker, Kubernetes/OpenShift, Kyverno, Argo CD, Kueue, IBM Cloud
Languages: Python, SQL, JavaScript, HTML/CSS, Go
Frameworks: Pydantic, Typer, Ray, pytest, FastAPI, Flask