

# Isac Pasianotto



## EDUCATION

<b>Ph.D.</b>	University of Trieste, <b>Applied Data Science and Artificial Intelligence,</b> Trieste, Italy	2024 – present
	<ul style="list-style-type: none"><li>• Research project: "Converged computing: environments and technologies for hybrid HPC and Kubernetes infrastructure"</li></ul>	
<b>Master's Degree</b>	University of Trieste, <b>Data Science and Scientific Computing,</b> Trieste, Italy	2022 – 2024
	<ul style="list-style-type: none"><li>• Final grade: 110/110 cum laude</li></ul>	
<b>Bachelor's Degree</b>	University of Trieste, <b>Statistics and Computer Science for Business, Finance, and Insurance,</b> Trieste, Italy	2019 – 2022
	<ul style="list-style-type: none"><li>• Final grade: 100/110</li></ul>	

## EXPERIENCE

<b>Operations Assistant during Ph.D.,</b> Science Park – Trieste, Italy	2024 – present
<ul style="list-style-type: none"><li>• Management of the production Kubernetes cluster</li><li>• Designed the architecture to serve a multi-user LLM on a GPU-enabled Kubernetes cluster running on-premise</li><li>• Contributed to the maintenance of the internal distributed storage system based on Ceph</li><li>• Assisted system administrators in on-premise cluster management</li></ul>	
<b>Internship,</b> Science Park – Trieste, Italy	2023 – 2024
<ul style="list-style-type: none"><li>• Developed a solution for user authentication and authorization in a cloud environment</li><li>• Implemented a solution based on LDAP and Kerberos to authenticate IdM users inside Kubernetes pods</li><li>• Main technologies used: LDAP, Kerberos, Kubernetes, Vagrant, Libvirt, Docker, Podman, Ansible</li></ul>	

## PUBLICATIONS

<b>Seeking Cost-Optimal Infrastructure Size for Distributed File Systems: A Ceph Case Study</b>	March 2025
N. Tosato, I. Pasianotto, R. Lot, S. Cozzini	
SC Workshops '25: Proceedings of the SC '25 Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis	

## SCHOOLS

<b>Advanced School on Applied Machine Learning</b>	June 2024
International Centre for Theoretical Physics (ICTP), Trieste, Italy	
Presented a poster titled "Gaussian processes for time-series data analysis: case study for a large dataset"	

## SKILLS

**Programming Languages:** Proficient in C, Python, and Bash, with strong experience in scientific computing and systems programming. Working knowledge of C++, R, and Java. Basic familiarity with Rust, Julia, and front-end development (HTML, CSS, JavaScript)

**Parallel & Distributed Computing:** Experienced in high-performance computing (HPC) environments, including parallel programming with MPI, OpenMP, OpenACC, and CUDA. Skilled in building distributed workflows using Python tools such as Dask and PyTorch, and familiar with scheduling and resource management systems like SLURM

**Cloud & Infrastructure:** Experienced in designing, deploying, and maintaining Kubernetes clusters in both bare-metal and virtualized environments. Familiar with distributed storage solutions (e.g., Ceph) and identity management systems. Knowledgeable in disaster recovery, backup planning, and network configuration in virtualized infrastructures

**Containerization & Virtualization:** Skilled in container-based development and deployment workflows using Docker and Podman. Experienced with Vagrant and virtualization platforms (Libvirt, Proxmox, and QEMU) for creating reproducible environments. Proficient in building infrastructure automation pipelines that integrate containers, virtual machines, and orchestration layers

**DevOps & Automation:** Strong background in Infrastructure as Code (IaC) and configuration management with Ansible and OpenTofu/Terraform. Highly familiar with Helm for Kubernetes deployments and GitOps practices

**OS & Environment:** Proficient in Linux system administration, shell scripting, and command-line tools. Experienced with Git-based workflows and remote system management. Skilled in interacting with HPC schedulers and debugging performance issues in parallel environments

**Identity & Access Management:** Experienced in deploying and customizing IAM platforms such as Authentik, including custom enrollment flows and secure provisioning of SSH keys. Working knowledge of FreeIPA and its integration with containerized and cloud environments

**Monitoring, Logging & Tooling:** Familiar with Unix/Linux monitoring and troubleshooting tools (htop, iftop, journalctl, etc.), log aggregation, and basic observability practices. Comfortable with advanced text editors (Vim, Emacs) and terminal multiplexers (Tmux)

**Infrastructure Orchestration & Testing:** Experience in building reproducible testbeds and simulation environments using Vagrant, Ansible, and Kubernetes. Skilled in provisioning multi-VM environments for complex systems (HPC clusters, Ceph storage, FreeIPA) and integrating them into automated CI pipelines

## PERSONAL PROJECTS

**STENCIL** 2023 – present

Comprehensive, ready-to-use framework to spawn an entire virtual cluster that faithfully mimics the behavior of a modern cluster

- Modular approach for both deployment and provisioning stages
- Integrated a Kubernetes cluster, an Identity Management System, a distributed filesystem, and SLURM as a job scheduler
- Completely declarative: built on OpenTofu and Ansible

**dask-bench** 2023 – 2024

Dask-based benchmarking tool for comparing performance across different schedulers on identical hardware.

- Compatible with SLURM and Kubernetes scheduling systems
- Used to evaluate Kubernetes overhead on identical hardware
- Based on the concept of Weak Scaling, executing various workloads with arrays and dataframes

**Federated Learning** 2024 – 2024

Simulated federated learning environment using MRI datasets.

- Scalable simulation both vertically (per-node resources) and horizontally (number of nodes)
- Uses an adapted AlexNet architecture for image classification

## CONTRIBUTIONS

**freeipa-scim** 2024 – 2025

Project to provide a System for Cross-domain Identity Management (SCIM) interface in front of FreeIPA to synchronize identities with Authentik

- Participated in the development of the code
- Managed the production deployment on Kubernetes clusters

#### **slurm-helm**

2023 – 2024

Helm chart to deploy a SLURM server (slurmctld and database) on Kubernetes

- Contributed to templating the chart
- Integrated an IPA client-like configuration based on LDAP+Kerberos to query a FreeIPA server for users' UID and GID

### CERTIFICATIONS

#### **NVIDIA Deep Learning Institute**

October 2024

Data Parallelism: How to Train Deep Learning Models on Multiple GPUs

### LANGUAGES

**English:** Proficient in both written and spoken English

**Italian:** Native speaker

### TEACHING

**Foundations of HPC:** Assistant lecturer for the course taught at the Master's degree in Data Science and Scientific Computing at the University of Trieste

**Scientific Programming Method:** 18-hour module at the Master in Data Management and Curation (MDMC) organized by SISSA and Science Park

**Cloud Computing, 2024:** Delivered lectures in the Advanced Cloud Computing course on CNI plugins and multi-node communication in Kubernetes clusters

### VOLUNTEERING

**AI2S:** Member of the Artificial Intelligence Student Society (AI2S), a student association active in the Trieste area

#### **Data Processing Authorization:**

I authorize the processing of my personal data in accordance with the EU GDPR. Re

