

Nicola Marvulli Site Reliability Engineer

✉ nicola.marvulli7ce@gmail.com 📞 (+39) 3519401871 📍 Via Ippolito Nievo 12, Rome, Italy 🇮🇹 Italian 📅 15/06/1998
🌐 <https://linkedin.com/in/nicola-marvulli-904270136> 🏠 <https://github.com/n1CkS4x0>

Professional Experience

Agile Lab

DevOps & Site Reliability Engineer 10/2022 | Rome, Italy

As a consultant for the Presidency of the Council of Ministers of the Italian Government. My key contributions include:

- **Designed and deployed a cloud-native platform** integrating HashiCorp Vault, Cert-Manager, and a full Observability stack (OpenTelemetry, Prometheus, Grafana, Jaeger, ELK) — improving visibility and reliability for 50+ engineers.
- **Led the migration of a 2PB data ecosystem from Cloudera and VMware Tanzu Kubernetes Grid to a complete on-premise stack with SUSE Rancher and VastData**, enhancing scalability and reducing infrastructure costs by 30%.
- **Built CI/CD pipelines (Jenkins, GitLab CI, Terraform) enabling 20+ monthly releases** across Java, Python, and Scala microservices with automated artifact versioning, reducing deployment time by 75%.
- **Architected a disaster recovery framework** for Kubernetes using Velero, MinIO, and Veeam, automating backups and ensuring high availability.
- **Delivered 24/7 on-call support and incident response**, reducing mean-time-to-recovery (MTTR) within strict SLA targets.
- **Developed Apache Airflow DAGs** to orchestrate data ingestion and model training pipelines, integrating REST APIs for dynamic configuration and job monitoring, reducing manual intervention by 80%.
- **Mentored junior engineers in automation, reliability, and system design**, improving team performance and onboarding efficiency.

Big Data Architect Intern

04/2022 – 09/2022 | Turin, Italy

During the internship, research was conducted for a master's thesis:

- **Designed a Big Data AI pipeline** reducing cloud service costs by optimizing distributed processing and ML workflows.
- **Integrated Apache Spark, Ray on Kubernetes, and MLflow**, accelerating model training and tracking metrics for reproducible experiments.
- **Automated infrastructure provisioning with Terraform and GitOps**, reducing deployment time by 60%.
- **Implemented CloudWatch monitoring for real-time performance** insights and resource optimization.

Education

Master's Degree in Computer Science (curriculum Artificial Intelligence)

10/2020 – 12/2022 | Bari, Italy

University of Bari Aldo Moro

Score: 110/110 cum laude

Bachelor's Degree in Computer Science (curriculum Software Engineering)

10/2017 – 10/2020 | Bari, Italia

University of Bari Aldo Moro

Skills

Programming Languages: Java, Scala, Python, C++, Rust, Go, Bash

DevOps/SRE Technologies:: Docker, Kubernetes, VMware Tanzu Kubernetes Grid, SUSE Rancher, Terraform, Helm, Jenkins, Gitlab CI, HashiCorp Vault, Cert-manager, Prometheus, Grafana, Alertmanager, Jaeger, Elasticsearch, Kibana, OpenTelemetry, KEDA, Kyverno, Nginx, Velero, Istio, ArgoCD

Big Data: Cloudera Data Platform, VastData

Cloud Providers: Amazon Web Services, Google Cloud Platform

Certifications

Kubestronaut

CKA: Certified Kubernetes Administrator

CKAD: Certified Kubernetes Application Developer

CKS: Certified Kubernetes Security Specialist

KCNA: Kubernetes and Cloud Native Associate

KCSA: Kubernetes and Cloud Native Security Associate

VAST Fundamentals

Operating The VAST Management System

CyberChallengeIT

Publications

Real-Time Age Estimation from Facial Images Using YOLO and EfficientNet

10/2021

Balancing Data Within Incremental Semi-supervised Fuzzy Clustering for Credit Card Fraud Detection

06/2021

Languages

Italian, English, Spanish