

Michael Morello

Developer and (distributed) system engineer, France

michael.morello@gmail.com | [LinkedIn](#) | [Twitter](#)

Currently

Working on the best Kubernetes/OpenShift operator for Elasticsearch.

Specialized in

Kubernetes *

Advanced operator development on Kubernetes using Golang and the K8S controller runtime.
Advanced cluster administration (including OpenShift 4)
Somewhat a K8S project contributor: k/k#83895 and the relevant PR k/k#83945

Linux 🐧

Automating things with Ansible (and a little bit of Bash)
Always keep a close eye on security issues (reporter of CVE-2020-7010)
15+ years of Java

Research interests

Distributed systems (CAP and FLP theorems, Google papers ...)
Computing languages (Java, C, Go)
IT Security
Machine learning

Speaker

Meetup ElasticFR #74 : Autoscaling de vos déploiements k8s avec les métriques elastic et autres [French]
Devoxx France 2018 : My Applications in production with Kubernetes [French]
Kubernetes Paris Meetup Juin 2016 [French]

Education

2003

Master in Distributed applications and artificial intelligence

Work History

2019

Principle Software Engineer @ Elastic

Software engineer on the Elastic Cloud on Kubernetes operator:
Key features development in Go
Continuous integration and release process using Jenkins
User support and advanced debugging
Kubernetes cluster management on GKE, AKS, EKS and OpenShift 4 (I'm also a very huge fan of Kind)

Side project:

Custom implementation of the Kubernetes Custom Metrics API for Elasticsearch.

CNAMTS is the French National Health Insurance, it guarantees healthcare access for nearly 55 million people in France

2015

Product Owner and Architect of a Private PaaS cloud platform @ CNAMTS

Lead engineer for cluster management using Openshift, Kubernetes, Docker and Ceph as the storage platform on baremetal and virtual infrastructures.
Automate most of the cluster maintenance tasks with Ansible
Leverage agile and CI/CD methodologies in order to continuously deploy PHP, Java and NodeJS apps in production

2014

Hadoop and BigData Engineer @ CNAMTS

Design and deploy Hadoop clusters :
Hardware and operating system choice (diskless CentOS customized distribution)
Designed processing system (HBase, Hive + TEZ, Spark, YARN, Solr Cloud on HDFS)
Kerberos with FreeIPA
Implemented metrics collection with HBase and OpenTSDB + Grafana to visualize data with nice dashboards
Developed a Spark program for web server log classification using the kmean algorithm

2009

Java and JEE Software Engineer @ CNAMTS

Build a highly available configuration directory with Apache Zookeeper and Apache Curator
Create simple but highly reliable file transfer tool with Netty, Google Protocol Buffers and Google Guice.
This tool is today deployed on hundreds of machines in order to transfer files of any size.
Lead Java engineer on a JEE (Oracle Weblogic) project

2004

Sopra Steria

Multiple positions as a developer (mostly in C and Java) in the domains of Banking & Finance, Telecom, and Healthcare.