

# Giuseppe Lavagetto

Site Reliability Engineer / Backend Developer

I am a Developer turned Site Reliability Engineer with more than 10 years of experience managing high-traffic websites.

Via F.Ozanam 53/B  
00152 Roma (RM) Italy  
+39 3922984018  
lavagetto@gmail.com

## EXPERIENCE

### Wikimedia Foundation, remote — Principal SRE

APR 2014 - PRESENT

During my tenure, I've been involved in many projects. Here are some highlights.

- **Multi-datacenter setup:** moving the Wikimedia infrastructure from working in a single datacenter to work in multiple ones, and automating all the procedures for managing pooling/depooling and moving services across datacenters.
- **Microservices architecture:** I was the first proposer of the migration to Kubernetes, and I drove the definition of quality standards and base requirements of microservices. Addition of service-mesh-like components in the stack. System for managing the build and security updates of docker images.
- **Software migrations:** I've been part of our efforts to move first from PHP 5 to HHVM, and then from HHVM to PHP7. I've managed transitions between major versions of various other parts of the stack (etcd, etc.)
- **SRE code quality:** contributed to the internal puppet style guide and of the corresponding linting plugin. Also set up maintained the CI for the puppet and the helm charts repositories.
- **Technical Committee:** I've been a component of the Wikimedia Technical Committee, the council of engineers who had decision/mentoring power over technical and architectural decisions for 4 years.
- **SRE onboarding:** I have defined the procedure for onboarding new hires onto any of the SRE teams, including a full set of training videos created by various SREs, with clear instructions for mentors on how to optimally organize the first 3 months of work of a new SRE. This is particularly important in a fully remote team.
- **MediaWiki on k8s:** I am the tech lead on a multi-team project to port the Wikimedia installation of MediaWiki to work on Kubernetes, completing the migration of our services infrastructure. This migration required reimagining how we serve MediaWiki.

## SKILLS

Linux  
System Administration  
Configuration Management  
Containers  
Applications design and Architecture  
Troubleshooting in production  
Kubernetes  
Monitoring  
Tracing

Development: Python, Go, PHP, Ruby, bash. Basic knowledge of javascript, C/C++

## Talks

**Kubecon EU 2018** - Building a Kubernetes on Bare-Metal Cluster

**IDI 2019 Keynote** - Microservices: when, why and how? Lessons learned at Wikimedia

**Envoycon 2020** - Improving performance of RPC with Envoy at Wikimedia

**IDI 2020** - Testing in production for fun and profit

## **Autistici.org, remote — admin**

MARCH 2010 - present

I am part of the autistici.org collective, providing blogging/email/IM services that respect the privacy of our users, aimed at helping activists avoid state control.

## **Exepdia.com, Rome — Head of Technical Operations, Venere.com**

MAY 2011 - MARCH 2014

I was hired as a senior Operations engineer, and was moved to lead the team shortly afterwards. Some highlight of my team during my tenure:

- **Drove adoption of MySQL** for the availability engine, and of NoSQL databases for the hotel catalog, resulting in a 70% reduction in latencies for the customers, a 90% reduction in software licensing costs, and less hardware used
- **Introduced devops practices:** configuring infrastructure-as-code with Puppet, got the Operations engineers to work daily with development teams to ease lines of communication and collaboration
- **Introduced RED-like metrics** for all services, improved reliability together with the development teams to get to 99.97% availability

## **Immobiliare.it, Rome — Sr. Software engineer**

MARCH 2008 - APR 2011

I was hired as technical employee #3 during the startup phase of what would become the largest real estate classifieds site in Italy.

- Built an infrastructure from the ground up for a website that would go on and grow 30% every trimester in terms of traffic and data size
- Helped debug and fix backend performance problems across the stack
- Wrote a REST API for the mobile application and took it to production in 1 week.
- Added observability and debugging tools to all the applicational stack, to allow developers to profile their code and find bad SQL queries before they hit production
- Lead the team that built a comprehensive asynchronous classified imports system, that would ingest about 1M classifieds/hour in many different formats, both from an API and ftp uploads.

## **Palermo University, Palermo — Post Doctoral Researcher**

JULY 2004 - FEB 2008

I worked for 4 years as a postdoctoral researcher in applied and theoretical astrophysics, participating in more than 30 publications on peer reviewed international journals.

## **LANGUAGES**

Italian - Native speaker

English - Fluent both in writing and speaking.

## EDUCATION

### Palermo University, Palermo — *Ph.D. in Physics*

JUNE 2001 - JUNE 2004

Graduated cum laude with a dissertation on the effects of General Relativity on the evolution of binary systems.

### Parma University, Parma — *M. Sc. in Physics*

NOV 1995 - NOV 2000

Graduated in theoretical physics (cum laude) with a dissertation on gravitational waves emitted by a binary system.

## SOFTWARE PROJECTS



Github Artic Code Vault Contributor

Over my career, I contributed patches to, or created many software projects. Check out my github profile (linked above). Here are a few of the projects I contributed to that I feel are most interesting:

- [Conftool](#) - a tool to manage a k-v set of entities, with associated tags and a validation schema. Used at wikimedia to control most of our automation via etcd
- [PyBal](#) - Wikimedia's L4 load-balancer
- [Purged](#) - Our daemon for managing cache invalidation on our CDN. I wrote the code to move invalidation consumption to kafka.
- [spicerack](#) - Wikimedia's own automation system

When I'm not working, I pride myself with writing silly projects. A couple I particularly like:

- [Memeoid](#) - the only FLOSS memegenerator supporting animated GIFs
- [Xkcli](#) - a cli tool (that I use in IRC/slack bots) to search the corpus of XKCD for a specific term.

## BLOG POSTS

- [Impact of using HTTP connection pooling for PHP applications at scale](#) - Blog post about the effect of introducing a service proxy for all http calls from a large PHP application
- [Some insight on the inner workings of noblogs](#) - An article introducing how a cheaply distributed installation of wordpress for noblogs.org works.