

MANLIO POLTRONIERI

<http://manl.io>

EXPERIENCE

Guardian News and Media

Software Developer

Jan 2017 - present

London, UK

- Development, debug and maintenance of our real-time analytics app (think Google Analytics, but completely tailored for the newsroom). It's a rather complex, highly scalable microservices architecture that tracks all the traffic to the theguardian.com, as well as integrating third-party apps (Facebook, Twitter etc) providing custom features for the newsroom, and feeding our Data Lake solution. The stack is mostly Scala/Play!, AWS, Elasticsearch.
- I had a short spell with the Data Science team where I helped building a new recommendation engine for theguardian.com. The project failed miserably and it was an eye-opening experience in why projects don't succeed.

Guardian News and Media

Associate Software Developer

Sep 2015 - Jan 2017

London, UK

- I was on the Off Platform and Content API team working on the API that serves both The Guardian's website and mobile offering along with internal applications and off platform services.
- Develop and maintain the backend API that serves billions of requests per month; breaking down the old Monolith into microservices, internal tools to monitor the publication flow.
- Design and implementation of a serverless system based on AWS Lambda to provide monitoring of podcasts consumption by mining our CDN logs and integrate it with our in-house real-time analytics solution.
- Implementation of The Guardian's Alexa Skill for Amazon Echo using NodeJS

Planet Earth

Backpacker

2014 - 2015

South East Asia

- Solo motorbiking trip throughout Asia. Relevant skills include eating with chopsticks, fixing a puncture, sleeping everywhere, developing an appreciation for the world and its diversity, being kind to strangers.

SKILLS

- I have experience designing, building, debugging and maintaining complex microservices architectures deployed on AWS. That said, I know that sometimes the good old monolith works just fine.
- Experience with Scala using the Play! Framework, Akka and actor-based (concurrent) programming, SQL and NoSql database solutions, especially DynamoDB and Elasticsearch.
- Experience with the AWS stack, specifically with EC2, AutoScaling, DynamoDB, Kinesis, S3, Lambdas and server-less architectures, Cloudwatch, Cloudformation etc.
- Experience hacking internal tools, web scrapers and web apps with Python, Flask, Sqlite, Bash or just any other tool which is right for the job.
- Understanding of Data Structures, Algorithms, Computational Complexity, OOP, Functional Programming.

- Other buzzwords include TDD which I don't believe in anyway because strong typing <3, solid Linux fundamentals, relational databases (design, normalization etc), SQL, some Javascript/NodeJS skills that I'd rather not use too much ;), version control with Git, \LaTeX which is now mostly forgotten and university courseworks using C, C++ and Prolog - I'd love to find any real world use case for Prolog, but I suspect there are none.

PERSONAL PROJECTS

Where is Who is hiring? hiring? <http://whereis-whoishiring-hiring.me>
Open-source wrapper for the monthly Hacker News' *Who is hiring?* thread which sorts the job openings by geographical location. It ranked #2 on HN and it was featured on Hacker Newsletter #247. Stack: Flask, Python, BeautifulSoup, Nginx, MySQL, Bash.

manlio <http://manlio.io>
My personal website with programming notes, small projects and other stuff that doesn't fit in a resume. It's self-hosted on a tiny Digital Ocean VM running Debian and Nginx.

EDUCATION

Imperial College London	2014
MSc, Computing Science	
Università degli Studi di Milano	2012
BA, Philosophy	

CONTACTS

Website: <http://manlio.io>
Github: <https://github.com/oilnam>
Email: manlio.poltronieri@gmail.com

Updated: February 2018