

Mark Litwintschik

Date of birth: July 2nd, 1982.

I hold both a British and a Canadian passport.

mark@marksblogg.com

Education

University of California, San Diego

April 2016

Hadoop Platform & App Framework Certificate

University of Amsterdam

May 2016

Basic Statistics Certificate

Johns Hopkins University

May 2016

R Programming Certificate

École Polytechnique Fédérale de Lausanne

June 2016

Functional Programming in Scala Certificate

Udacity

August 2016

Deep Learning with Google's Tensorflow

University College London

March 2008

Estonian Language Certificate

Technical Skills

Backend, Frontend & Version Control

Python	6 Years	Expert
Django, Flask	6 Years	Expert
Tornado	4 Years	Expert
Scala	1 Year	Good
HTML	22 Years	Expert
CSS	21 Years	Expert
CoffeeScript, JADE, SASS, SCSS	2 Year	V. Good
Yeoman, Grunt, Bower	2 Year	V. Good
Backbone.js	3 Years	V. Good
AngularJS	2 Years	V. Good
jQuery	10 Years	Expert
Bash, Coreutils	15 Years	Expert
Git	9 Years	Expert

Data-related

BigQuery	1 Year	Expert
Elasticsearch	6 Years	Expert
Flume	1.5 Years	Expert
Hadoop	6 Years	Expert
Kafka, Kinesis	2 Years	Expert
Memcache	9 Years	Expert
MySQL, MariaDB	16 Years	Expert
Periscope	2 Year	Expert
Postgres	7 Years	Expert
R, RStudio	1.5 Years	V. Good
Redis	7 Years	Expert
Redshift	2 Years	Expert
Solr	1 Year	V. Good
Spark	2.5 Years	Expert

Continuous Integration & Deployment

Ansible	2 Years	Expert
Chef	6 Months	Good
Fabric	6 Years	Expert
Jenkins	3 Years	Expert
Vagrant	2 Years	V. Good

Message Queues & Job Systems

Amazon SQS	1 Year	V. Good
Celery	5 Years	Expert
RabbitMQ	5 Years	Expert
ZeroMQ (ZMQ)	2 Years	V. Good

Cloud, Hosting, Deployments & Monitoring

Apache	16 Years	Expert
Nginx	8 Years	Expert
Gunicorn	6 Year	Expert
HAProxy	1 Year	Good
Varnish	1 Year	Good
Amazon AWS (S3, EC2, Cloudfront, etc..)	9 Years	Expert
Google App Engine, Compute Engine	3 Years	Expert
LXC & Docker	2 Years	Good
Nagios	6 Years	V. Good

Testing

Unit Testing, mock, coverage	6 Years	Expert
------------------------------	---------	--------

Work Experience

Big Data Consultant; Python / HTML5 / Cloud Contractor

September 2011 – Present

London, UK and Dublin, Ireland

My clients, end-clients and partners include Bank of America Merrill Lynch, British Telecom, Ford, Google, the Guardian, News International, Nokia, The Royal Mail, TransferWise and Vertu.

I specialise in Hadoop, Spark, Amazon EMR, Kinesis, Kafka, RDS, Redshift, SQS, Machine Learning at scale and a whole host of other big data technologies and frameworks. I've run clusters of 100s of nodes working on petabyte datasets. I've trained models with 97.5% predictive accuracy.

In the first half of 2016 I've been helping a client with their Kafka and Zookeeper Cluster on EC2 which is used to stream telemetry from professional tennis matches including data from the Wimbledon Tournament.

Throughout 2015 I worked with a client to:

- Tokenise data from a 1.2 petabyte dataset comprised mostly of unstructured and semi-structured data using MRJob on Amazon Elastic MapReduce (EMR)
- Load that data into their Redshift data warehouse
- Load additional data from Amazon Kinesis via S3
- Setup ETL jobs with Airflow (a data pipeline management package that was born at AirBNB)
- Analyse their data (we compared using Looker, Periscope, MicroStrategy and Tableau on top of R- and iPython-based workflows)
- Train models using supervised machine learning algorithms (we compared various algorithm implementations in Spark's MLlib, R, scikit-learn and AmazonML)
- Serve model predictions via a Flask-based API running on EC2 spot instances behind Amazon ELB.

In a 2014 project for Google I setup clusters of machines on Google's Cloud Compute and App Engine platforms to crunch and serve out large amounts of data with Hadoop while centralising logging of all services via Flume to an Elasticsearch and Kibana installation. There is a data pipeline that does enrichment and makes sure data reaches Google's audience in a reliable way by avoiding single points of failure.

When working hands-on I do test-driven development and subscribe to agile methodologies. Django, Tornado, Flask, Backbone.js, AngularJS and various other frameworks make frequent appearances in my work.

Fullstack Web Developer

Bloomberg

All of 2009 & July 2010 – September 2011

London, UK

I was a senior, full-stack developer for bnef.com, I was in charge of delivery of a large volume of transactional emails and I integrated Adobe LiveCycle DRM system into our platform.

Chief Technology Officer

iVoucher

January 2010 – July 2010

London, UK

VC-funded start-up focused on distributing discount codes on behalf of retailers. I built a two-way SMS systems and web-based vendor-facing tools.

Technology Director

Kieon

September 2008 – January 2009

London, UK / Bangalore, India

I managed sixty developers in Bangalore. We delivered technical solutions for the end-clients of several marketing agencies in London.

Fullstack Web Developer; Technical Manager

blue barracuda marketing ltd.

March 2004 – September 2008

London, UK

I delivered several technical solutions for a wide range of end-clients. The LAMP stack was prominent in most of our solutions.

My end-clients included: Angels, Austin Reed, BAA, Barchester, BlackBerry, Buy Cosmetics, Danone, ELC, Fat Face, Financial Times, Goldsmiths, Hamleys, Harrods, ISBA, ITV, Krispy Kreme, Laithwaites, LOVEFiLM, Nectar, NMA, Noble House Media, Novus, Office World, Pizza Hut, PricewaterhouseCoopers, Royal Bank of Scotland, Staples, T-Mobile, TGI Fridays, The Banker, UKTV, Williams Formula 1, Xerox and Visit London.

Fullstack Web Developer

COMBridge AG

September 2002 - March 2004

Bremen / Wolfsburg / Hannover, Germany

I developed web applications for my company to support their SAP projects with Volkswagen and Audi.

Interests

I enjoy Formula 1, photography and mountain biking.