

Mark Litwintschik

British / Canadian Dual-Citizen

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 & Cloud

Python, Django, Flask, Celery & RabbitMQ	> 5 Years	Expert
Apache, Nginx	> 10 Years	Expert
Scala	1 Year	Good
HTML, CSS, JavaScript	> 10 Years	Expert
Git	> 10 Years	Expert
AWS (S3, EC2, Cloudfront, etc..)	> 10 Years	Expert
Google App Engine	3 Years	Expert
Ansible	3 Years	Expert
Jenkins	3 Years	Expert

Data & Deep Learning-related

Airflow	3 Years	Expert
AWS Athena	6 Months	Expert
AWS EMR	4 Years	Expert
Azure w/ HDInsight	2 Years	V. Good
BigQuery	2 Years	Expert
ClickHouse	2 Years	V. Good
Elasticsearch	> 5 Years	Expert
Flume	2 Years	Expert
GCloud Dataproc	1 Year	Expert
Hadoop (YARN, HDFS & Hive on Cloudera & HortonWorks)	> 5 Years	Expert
Jupyter Notebook	3 Years	Expert
Kafka, Kinesis	4 Years	Expert
kdb+/q	1 Year	Good
MapD	1 Year	V. Good
Memcache	> 10 Years	Expert
MS SQL Server	2.5 Years	Expert
MySQL, MariaDB	> 10 Years	Expert
Pandas	2 Years	Expert
Periscope	2 Years	Expert
PostgreSQL	> 5 Years	Expert
PrestoDB	3 Years	Expert
R, RStudio	2 Years	V. Good
Redis	> 10 Years	Expert
Redshift	4 Years	Expert
Solr	1 Year	V. Good
Spark	4 Years	Expert
Tableau	2 Years	V. Good
TensorFlow	1 Year	V. Good
Vertica	1 Year	Good
ZooKeeper	> 5 Years	Expert

Work Experience

Big Data Consultant

2011 – Present

London, UK

I specialise in big data, machine learning and deep learning platform engineering. I've setup 200-node clusters, worked with datasets measured in petabytes and built data pipelines handling 10s of billions of records per day. I've delivered a large number projects and POCs on AWS, Google Cloud and Azure for several global brands. My work has included setting up encryption as rest, Kerberos integrations and training teams of data scientists on best practices.

Over the past few years I've completed work for two airlines, two UK supermarket chains, the largest Investment Bank in Switzerland and France's biggest mobile ad buyer. Some of my more well-known clients, end-clients and partners include Bank of America Merrill Lynch, British Telecom, Ford, Google, the Guardian, LeoVegas, News UK, Nokia, The Royal Mail, TransferWise and UBS.

I also run a popular big data tech blog where I benchmark various databases against a dataset of 1.1 billion records. My benchmarks have been reviewed by the likes of Google¹, IBM², Intel³ and Nvidia⁴.

¹ <http://bit.ly/29o94bN> ² <https://ibm.co/2ixsPqc> ³ <http://intel.ly/2vvUPeU> ⁴ <http://bit.ly/2ozsZPI>

Fullstack Web Developer

Bloomberg

2009 - 2011

London, UK

I was a senior full-stack web developer for Bloomberg New Energy Finance. I also built a high-volume transactional email service and integrated Adobe LiveCycle DRM system into the firm's systems.

Fullstack Web Developer; Technical Manager

blue barracuda marketing ltd.

2004 - 2008

London, UK

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

Our 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.