

Ian Pointer

Data/ML/Deep Learning Engineer. Experienced with Operations and Development at Scale

Skills

Languages: Ruby, Scala, Python, Java, Elixir, Rust, Go, JavaScript, C, Perl, SQL, R, Objective-C

Experienced with: Apache Spark, Apache Hadoop, Apache Kafka, Apache Storm, Apache Beam, Tensorflow, Keras, PyTorch, Apache Apex, Apache Flink, Apache Hive, Cassandra, Elasticsearch, Logstash, Kibana, Oozie, Graphite, Zeppelin, Redis, AWS, Google Compute Engine / Dataflow, Memcache, Docker, Mesos/Marathon/Chronos, MySQL, PostgreSQL, LDAP, Apache HTTPD, Nginx, Samba, Active Directory, SQL Server, Jenkins, Git, Subversion, Hyperic, Nagios, Puppet, Chef, Ansible, CouchDB, HTML5/CSS3, Node.JS

Experience

2016-Present **Senior Solutions Architect / Engineering Manager for Kogentix, Inc** (Remote)

Deep Learning projects for clients including object recognition, segmentation, and classification across various domains.

Apache Hadoop ecosystem (Spark, Kafka, HBase) cloud-based data warehousing and analytical projects for multiple global financial companies.

2013-2016 **Lead Consultant for Mammoth Data** (Durham, NC)

Working on a wide variety of projects for clients large and small, including designing, implementing, and maintaining massive-scale AWS and Google Cloud Platform deployments, data pipelines involving Kafka, Storm, and Spark, along with analysis such as recommendation engines, natural language parsing, and anomaly detection.

Architected and developed an ETL application for a real-time advertisement bidding company using Apache Spark for processing incoming data for further analysis. (5bn records / day)

Architecting, developing, tuning and troubleshooting Spark Streaming applications.

Developed and maintained Kafka/Storm topologies for ETL and in-flight processing at scale.

Designed and implemented an orchestration engine in Ruby revolving around OpenStack/Git/Jenkins for a large email marketing company that sends over 250m emails per day. The engine was responsible for deploying the organization's applications across dev, stage, and production, complete with pluggable APIs for Docker support and alternate cloud providers, along with capabilities for online scaling and self-healing of application groups.

Designed and implemented a multi-region AWS cloud setup using CoreOS, Docker, and Ansible for a company allowing them to increase availability and scalability for their Rails/Redis/MongoDB application stack, serving millions of API requests per day. Profiling and improving Rails applications using Ruby and OS-level profiling tools.

Created financial modeling simulations, streaming solutions, natural language processing (NLP) pipelines, and recommendation engines on Spark clusters, deploying them on AWS on YARN and Mesos.

Architected and deployed data pipeline for a music festival using Apache Kafka, Spark, and Mesos, delivering real-time analysis from disparate data sources. Constructed collaborative

filtering solution for suggesting other events to attendees during the festival, along with sentiment analysis of social media.

Committer on various Open Source projects including spark-testing-base and Ansible.

- 2013 **DevOps Engineer for Reverbnation** (Durham, NC)
DevOps in a high-speed, web-scale Ruby-on-Rails environment (> 3.5m users with 1m unique logins per day). Worked on the migration from Rails 2.3 to Rails 3 and helped improve the transactional mail system to increase throughput and reduce spam rates as well as moving to a new cloud provider and diagnosing and improving Redis/MySQL performance issues.
- 2012-2013 **Systems Administrator for Rho, Inc** (Durham, NC)
Maintaining the company's fleet of CentOS Tomcat application servers running on VMware vSphere. Migrating old hardware to new virtual machines and general troubleshooting & tuning of Java applications. Developed a prototype iOS questionnaire app for internal testing, and working on visualizing data using d3.js. Rolling out automated deployment of servers with Puppet & vSphere.
- 2011-2012 **Developer for Open Software Integrators** (Durham, NC)
Part of a team contracting for various Fortune 500 companies all over the USA, specializing in web front-end development (including a Single-Sign-On project to support one of the largest video game releases of 2011) and DevOps, carrying out troubleshooting and development roles in technologies such as Spring, tcServer/Tomcat, RHEL/CentOS, and HTML5+JavaScript+jQuery.
Worked on the launch of a multi-server platform that receives 1m+ hits a week, rebuilding the front-end system using a combination of HTML5+JavaScript and Spring, while also writing tools and using packages such as Hyperic to obtain performance metrics.
- 2005-2011 **Administrator and Developer for d'Overbroeck's College Ltd** (Oxford, UK)
Managed a team of four people, responsible for over 500 users (teachers, administrative staff, and students - including many non-English speakers) across six sites in Oxford, reporting to the IT Manager and the Principal. Developed various bespoke applications and databases, all running on Linux hardware, including a marketing reports system, printer usage tracking software, a ticket-request system, and a web-based academic grading and reporting application complete with PDF output and email notification to parents, utilizing technologies such as HTML5/CSS3, JavaScript, Ruby on Rails, Sinatra, CouchDB, Redis, MySQL, and SQL Server.
- 2003-2005 **Technical Journalist** (Oxford, UK)
Writing articles for Linux Journal and Linux Magazine on the subject of coding with Mono and Emulation in Linux. Contributor to Open Source Press Books on the subject of DVD Authoring.
- 2002-2003 **TA for University of North Carolina at Chapel Hill** (Chapel Hill, NC)
TA in Computer Science department.
- 2000-2002 **System Administrator for Oxford Brookes University** (Oxford, UK)
Part of a four-person team responsible for over 15,000 users. Maintained Solaris/ Windows NT/Windows 2000/OpenVMS systems. Extended and maintained university's Perl-based webmail application running on a Solaris server. Responsible for keeping Solaris and Windows machines fully patched. Solved network, application, and server issues where appropriate. Responsible for Active Directory maintenance and replication, and DNS for the whole university. Wrote utilities in Perl and Python for diagnostic purposes.

Education

- 1997-2000 *BSc(Hons) in Computer Science* (1st Class Degree with Honours, Manchester, United Kingdom).

Publications

2016-Present *Infoworld*
2013 *Instant Zepto.js (Packt Press)*
2005 *Open Source Video (Open Source Press)*

Talks / Conferences

2016 *Life In The Trenches With Spark*
2015 *Profiling Ruby*
Apache Spark & Storm - Where And When?

Activities and interests

Hobbies Amateur chocolatier, avid fan of British mid-century architecture, and boardgame-collector.
Activities Organizer of the Triangle Apache Spark Meetup Group, an active Big Data group in the Triangle, NC.

ian@snappishproductions.com • [@carsondial](#) • 919-450-5242 • Durham, NC, USA • [github](#) •