

Philip O'Toole

<http://www.philipotoole.com>

github.com/otoolep

philip.otoole@yahoo.com

I am an experienced Software Engineering Leader, with a passionate interest in software and the teams that write it. I have a solid track record of delivering quality software and shipping products on-time, in many diverse fields, and I have led many software engineering teams as both a Technical Lead and Engineering Manager. I have contributed to open-source software, written for company blogs, mentored junior engineers, given customer presentations, spoken at major industry conferences, and assisted with field testing.

TECHNICAL EXPERTISE

- Design:** SaaS Infrastructure, Cloud Computing, data pipelines, search and analytics systems, distributed systems and database design, Linux-system software, Web Services, REST API development, general service-oriented architecture.
- Software:** In-depth experience with C, C++, Python, Go, bash and assembler, and notable experience with SQL, Java, and node.js. Deep understanding of multithreaded software.
- Systems:** AWS Cloud infrastructure, Apache Kafka, Apache Storm, and both SQL and NoSQL databases. Significant experience with Lucene-based search technology such as elasticsearch.

OPEN SOURCE - Creator of various open-source software including:

- **rqlite**, a distributed relational database, with a SQLite storage engine and Raft as consensus protocol.
- **Ekanite**, a Syslog server, with built-in full-text search for received log data.

PROFESSIONAL OUTREACH

- Presented at aws:reinvent, describing Loggly's use and deployment of Apache Cassandra, Apache Kafka, Apache Storm, and Elasticsearch, in the AWS EC2 cloud.
- Presented at multiple tech Meetups, including San Francisco Go and San Francisco Elasticsearch.
- In addition to my own blog, I have authored multiple technical blog posts, most notably for Riverbed Technology, Loggly and InfluxDB. These posts described key technical details of the systems I designed and implemented, promoted new releases and features, and advocated improved industry practices.
- I have acted as official technical reviewer of various publications such as *ElasticSearch Cookbook Second Edition (2015)* and *Implementing Cloud Design Patterns for AWS (2015)*.

EXPERIENCE

Google, Pittsburgh, PA

2018 – present, Engineering Manager

- Leading a team of developers focused on designing and developing Stackdriver, the DevOps and Observability system for Google Cloud Platform.
- Responsible for technical design, growing the team and developing individual engineers within the team, roadmap and vision, planning, and being ultimately responsible for effective execution.
- Working closely with senior Stackdriver Engineering Management, Product Management, and the broader Stackdriver Product Development team.

Percolate, San Francisco, CA

2016 – 2018, Director of Data Platform Engineering

- Managing and growing two teams of developers responsible for customer-facing analytics, external integrations, service-oriented architecture, and data ingestion, for the Percolate SaaS platform.
- Led the team that completely rebuilt Percolate's next-generation Analytics, which launched in early 2017.
- Responsible for roadmaps, technical design, coding and code reviews, effective execution, hiring, personnel management, as well as working closely with Product Management and the Executive team.
- The technologies in use by the Platform teams include Python, Go, Apache Kafka, Elasticsearch, Django, relational databases such PostgreSQL and MySQL, Docker, AWS EC2, and Dynamo DB.

InfluxData, San Francisco, CA

2014 – 2016, Director of Engineering

- Core developer of InfluxDB, the open-source distributed time series database written in Go.
- Significant contributions to the 0.8 & 0.9 releases, and the new high-performance TSM storage engine.
- Development of data ingestion, distributed consensus system, and query engine. Strong advocate for coherent design, testing, and ensuring the development team is as coherent and effective as possible.

Jut Inc., San Francisco, CA

January 2014 – November 2014, Senior Software Engineer

- Developer of Jut's Big Data analytics solutions. Full stack node.js development, within an Agile-Scrum process. Drove backend development of Jut's AWS-hosted Account & Deployment management application, including design and development of a PostgreSQL-based system for the production site.

Loggly Inc., San Francisco, CA

2012 – 2014, Lead Backend Architect

- Technical Lead of the software infrastructure team that designed and developed Loggly's 2nd Generation Cloud-based multi-tenant Log-analytics SaaS platform, with real-time indexing and search. Responsible for an overall coherent design, implementation, scaling, quality, bug tracking, and prioritization. Reporting directly to the CTO, I worked closely with Product Management and was closely involved with hiring as the team grew to 5 developers. 2nd Generation system went live September 2013.
- Designed and developed the Edge Data Collectors and core data pipeline code, involving work in Java, C++, Clojure, and Python, and built on Apache Kafka, Apache Storm, and Elasticsearch.
- Worked closely with the Operations team, writing deployment guides, release notes, advising on best practices, and addressing service outages.

Riverbed Technology Inc., San Francisco, CA

2007 – 2012, Member of Technical staff, later promoted to Engineering Manager, Cloud Services Development

- Managed a team of 5 developers, responsible for design and development of the Riverbed Cloud Portal, an Enterprise-grade web Portal hosted in the AWS Cloud. The Portal allows Riverbed customers to provision, manage, and license their virtual Riverbed network appliances, and Steelhead Cloud Accelerator Services, in multiple different public and private Clouds. Developed using Python and the Django web framework. The Portal first went live in November 2010.
- One of the lead developers of the Cloud Steelhead and Steelhead Cloud Accelerator, a distributed system developed jointly with Akamai, which significantly optimizes network access to multiple SaaS platforms. Django, Python and C development and implementation of numerous REST APIs. The Cloud Steelhead was released November 2010, and the Cloud Accelerator was released February 2012.
- Member of the team that developed the Riverbed Services Platform (RSP). RSP allows Riverbed's customers to run virtual machines on their Steelhead appliances. Launched in 2008.

TiVo Inc., Alviso, CA

2004 – 2007, Member of the Technical Staff

- Key member of the system software team that developed TiVo's flagship Linux-based Series3 high-end DVRs. Developed API definitions, implemented device drivers and coded kernel & user space code (C, C++) for tuners, demodulators, related MPEG input subsystems and the media pipeline. Performed board bring-up of multiple tuners and encoders. DVRs shipped September 2006 and July 2007.

Skystream Networks Inc., Sunnyvale, CA

2000 – 2004, Embedded Software Engineer

- Designed, developed and debugged high performance embedded software and firmware for Skystream Networks' Mediaplex product, an x86 VxWorks-based MPEG data and IP router. Extensive C and assembler programming. Performed board bring-up with HW engineer. Router was released June 2002.

Nortel Networks, Galway, Ireland

1997 – 2000, Software Engineer

- Developed real-time embedded software for wireless telephone system. Product released Q3 1998.

EDUCATION

1993 – 1997, National University of Ireland, Galway, Ireland

- B.Sc. in Applied Physics and Electronics - graduated summa cum laude.
- Awarded Nortel Networks prize for final year project – *An Experiment in Virtual Reality*.