



Ashwanth Kumar

SOFTWARE ENGINEER · DEVELOPER EVANGELIST

☎ (+91) 96000 87935 | ✉ ashwanthkumar@googlemail.com | 🏠 ashwanthkumar.in | 📺 [ashwanthkumar](https://www.youtube.com/ashwanthkumar)

tl;dr

6+ years of experience · Loves to understand, build and operate distributed systems · Ops · Organizes internal / external hackathons · Talks at local meetups and conferences · OSS Supporter · Blogs at blog.ashwanthkumar.in and indix.com/blog ·

Experience

Indix

Chennai, India

PRINCIPAL ENGINEER - DEVELOPMENT

Oct. 2016 - Present

- Currently working on Product Matching - Saying with a very high confidence that given any two products are same or not.
- Tech Lead in Data Ingestion Team that involves Code / Design Reviews and Mentoring for new members to the team.
- Finder - HTML Archive Storage system for storing over 3 billion HTML documents.
- **Suuchi** - OpenSource toolkit to build distributed data systems. Built for Finder initially but later adopted to other projects internally as well.

Indix

Chennai, India

PRINCIPAL ENGINEER - DEVOPS

Oct. 2015 - Sept. 2016

- Responsible for the entire AWS Infrastructure and cost control.
- Work with tools like - Ansible, Chef, Mesos, Marathon, Go.CD, Hadoop and AWS
- Saved around million dollars on AWS using [matsya](#) and [vamana](#).
- Built an [alerting tool](#) for apps running on Marathon.
- Contributed [gocd-janitor](#) to GoCD to support intelligent artifact cleanup on the CI/CD servers.

Indix

Chennai, India

SOFTWARE ENGINEER

Jun. 2012 - Sept. 2015

- First generation of Indix Data Pipeline which is completely based out of λ architecture.
- Two generations of our web crawler - scaling to millions of pages per day across 1000s of websites.
- Two generations of Scheduler which feeds our crawler with the right urls every day.
- Three generations of Parser. Built a DSL to express transformations in a language agnostic way that scaled the parsers to 1000s of websites.
- From Java based MR jobs to Scalding to building a query tool with a DSL (inspired from SQL and on top of scalding) which enabled engineers to access the data much faster for analysis / debugging.

Mu Sigma

Bangalore, India

BIG DATA INTERN

Jan. 2012 - May. 2012

- Evaluating AWS for migrating to cloud from custom data centers
- Worked on running R code on Storm.

Skills

DevOps Tools AWS, Mesos, Marathon, Docker, Chef, Ansible, Basic Linux Administration

Programming Scala, Java, Go, Node.JS

Big Data Hadoop, HBase, Spark, Zookeeper, Kafka, Samza, RocksDB

Talks

MEETUPS

Apr. 2017 **Suuchi - Distributed System Primitives**, DevDay

[Slides](#) | [Video](#)

Nov. 2016 **Building Distributed RocksDB**, Geeknight Chennai

[Slides](#) | [Video](#)

Apr. 2016 **Running Infrastructure on Spot instances in AWS**, Geeknight Chennai

[Slides](#) | [Video](#)

CONFERENCES

April 2018 **Why we built a distributed system**, DSConf 2018

[Slides](#) | [Video](#)

Nov. 2017 **Using Monoids for Large Scale Aggregates**, Scala.io France

[Slides](#) | [Video](#)

Nov. 2016 **Lessons from managing Hadoop clusters on AWS @indix**, DevopsDays India

[Slides](#) | [Video](#)