

PRAJWAL TULADHAR

praj@infynyx.com | <https://www.tuladhar.dev/>

Work Experience

Squarespace (June 2020 - Present) - Senior Engineer

- Currently working in Scheduling. Checkout <https://www.squarespace.com/scheduling>

RunAds / Publicis (April 2016 - June 2020) - Director Software Engineering

- Lead ETL / Reporting backend for RUN 2.0 (new RUN platform) using Spark for batch processing, MongoDB and Elasticsearch for time series queries; Frontend tools: React.js, and D3.js; Rails as an API backend
- AWS Lambda@Edge and Cloudfront based Ad server with global distribution and edge computing powering existing analytics and DMP pipelines; Currently serving 100M impressions per day
- Led data team for managing ETL / Reporting for PeopleCloud (<http://www.publicisgroupe.com/en/news/press-releases/publicis-groupe-advances-in-its-ambition-to-become-the-market-leader-in-marketing-transformation-and-digital-business-transformation>). Responsible for ingesting and aggregating logs from DoubleClick (DCM), Sizmek, TradeDesk for batch processing and storing as time series database in Mongo and Elasticsearch; Also responsible for hiring and interviewing data engineers
- Move existing ETL pipelines (Ruby) from Hive to Spark (Scala) and Apache Airflow as workflow framework; Cost saving and faster execution
- Boost performance of existing Elasticsearch cluster (25 data nodes); Elasticsearch instrumentation via Datadog / JMX; throttle index creation; Shards management; Automate S3 based backup and ES cluster Snapshot; Manage upgrades (from 1.7x to 2.4.x to 6.4.x)
- Work on migration from non-sharded database to sharded MongoDB cluster; Statistics and reporting component of the system involving massive refactoring (Ruby on Rails)
- Ansible + AWS Cloudformation for cloud provisioning and deployment
- Helped to start code reviews culture and mentoring junior engineers
- Initiated and led on call scheduling for data team using PagerDuty; Moved email based alerting to StatsD (Datadog) and AWS Cloudwatch
- Work with Ops team to move EC2 nodes from classic to VPC which increased security (centralized SSH login only via bastion hosts), reduced data transfer cost; thousands of dollars of AWS cost savings by resolving and refactoring “cloud debt”

- Move / containerized some old and new API services to Amazon ECS; cloudformation and ansible for orchestration

Nomi (Jan 2015 - Feb 2016) - Senior Software Engineer (Acquired by Axper)

- Worked on backend component of the system using Java, Python, HDFS (Hadoop), Storm and PostgreSQL
- Migrate tracks data (wifi visits and sessions) from MongoDB to HDFS cluster
- Hive and Impala for offline batch processing (generating report snapshots)
- Designed and deployed storm pipeline for camera data (visual sensors) and wifi data with real time aggregation backed by Kafka as source sink (Protocol Buffer format) and PostgreSQL (9.4.x) as storage layer
- Build application instrumentation using StatsD and JMX with Datadog as Time Series Database
- JVM tuning using tools like Mission Control (production time profiling and diagnostics tool), JMX Console, heap dump and strace
- Configuration management and deployment of services (rpms) using Puppet + Hiera and little bit of Docker
- Code reviews, pair programming, handling system availability and uptime during monthly on-call rotation

Sailthru (Jan 2011 - Jan 2015) - Senior Software Engineer (Acquired by Campaign Monitor - 2018)

- First engineer to be hired hacking on all sort of stuff from backend to frontend to API using mostly using PHP, Java and MongoDB
- Responsible for working on API clients in different languages like: PHP, Ruby, Java, Node (CoffeeScript), C#, Python. Projects can be found at <https://github.com/sailthru?query=client>
- Distributed logging using Scribe
- Metrics driven development using Graphite, StatsD, Metrics
- Wrote initial integration and unit tests for [Public API](#) and backend systems
- Multiple cross data center migration with cloud (AWS) + on premises datacenter setup
- Helped to scale user profiles from 10 million to 600 million. This was done initially with custom sharding followed by Mongo's default sharding backed by a common abstraction layer in PHP and Java.
- Migrated legacy mailer services with plain text protocol to HTTP Restful service using Dropwizard (0.6.x and 0.7.x)
- Build backend for iOS push notifications and Android push notifications (XMPP)

- JVM Tuning (Garbage Collection tuning, debugging, profiling) for both I/O bound and CPU bound services

Open Source Contributions

- Some minor patches to Apache Spark
<https://github.com/apache/spark/commits?author=infynyx>
- Contributed few [patches](#) to [Netty Project](#)
 - Enable checkstyle to maintain code consistency -
<https://github.com/netty/netty/commit/05850>
 - Ported HashedWheelTimer fixes from Akka Project -
<https://github.com/netty/netty/commit/915cb8b>
- [Facebook Nifty](#)
 - Added interface to expose Channel statistics for metrics collection -
<https://github.com/facebook/nifty/pull/86>

Connected Ventures | College Humor (Aug 2009 - Jan 2011)

- Backend development and scaling for the main site: collegehumor.com using PHP and MySQL; Created backend for link tracking internal analytics system using PHP, MongoDB, Nginx and Scribe (Centralized logging)
- Tools / Technology used: PHP, MySQL, Scribe, Nginx, MongoDB, Memcached

Fli Digital, Inc (Dec 2008 - June 2009)

- Worked as a developer for a social networking site and involved in writing backend for RESTful APIs
- Technologies: PHP, Javascript, PrototypeJS, MySQL, Memcached

Education

- Tribhuvan University, Nepal; Bachelor in Information Management (2008)