

Prateek Agarwal

Passionate about building scalable and fault-tolerant infrastructure.

+91 99366 30955
✉ [prat0318\[at\]gmail.com](mailto:prat0318[at]gmail.com)
📄 prat0318.com
🌐 [prat0318](https://www.linkedin.com/in/prat0318)
🔗 [prat0318](https://github.com/prat0318)

PROFESSIONAL EXPERIENCE (~16 years)

Current - **Uber, Data Infrastructure**, Bangalore, IN.

Jan. 2019 **Senior Staff Sw Engineer** (Jul. 2023 - Now)

Staff Software Engineer, TLM (Jan. 2021 - Jun. 2023)

Senior Software Engineer (Jan. 2019 - Dec. 2020)

- Leading the migration of Uber's Data Lake (~500PB Storage, ~1000K Cores Compute) from On Premise to Cloud without customer disruption (a multi-year effort).
- Responsible for Disaster Recovery Strategy for the complete Data Platform.
- Built and Managed the Storage / HDFS team in India from scratch.
- Responsible for high availability and durability of Uber **HDFS clusters** which physically store ~1.5EB data (One of the biggest in the Industry).
- Working on various Cost Efficiency efforts to keep the clusters running at their maximum utilization.
- Led the **Kafka Security** and performance projects to adhere to the PII, PSD2 and SOX compliance.
- Technologies got to work on: **HDFS, YARN, Airflow, Kafka, ZooKeeper**

Dec. 2018 - **Appdynamics, Common Infrastructure Platform**, Bangalore, IN/San Francisco, USA.

Apr. 2017 **Staff Software Engineer** (Nov. 2017 - Dec. 2018)

Senior Software Engineer (Apr. 2017 - Oct. 2017)

- Part of the *Data Ingestion team* responsible for scaling the backend services on cloud.
- Architected the **Kubernetes deployment** framework to make it multi-tenanted.
- Designed workflows to easily provision new Kubernetes clusters on both On-Premise and AWS.
- Built the Kubernetes deployment framework to break monolith Platform services into micro-services.
- Built Monitoring/Alerting/Logging solutions to efficiently monitor deployments on Kubernetes.
- Technologies got to work on: **Kubernetes, Helm, Kops, Python, Jenkins, Ansible**

Apr. 2017 - **Yelp, Distributed Systems**, San Francisco, USA.

May. 2014 **Software Engineer** (Jan. 2015 - Apr. 2017)

Software Engineer Intern (May 2014 - Aug. 2014)

- Maintained and provided tooling around NoSQL datastores - **Cassandra, Elasticsearch**.
- Built **Zipkin** (Distributed tracing) infrastructure at Yelp to debug service latencies.
- Wrote **pyramid_zipkin, swagger_zipkin** packages to easily integrate Zipkin with existing services.
- Co-authored **bravado** and **swagger-py** packages to integrate python services with Swagger specs.
- Technologies got to work on: **Cassandra, Elasticsearch, Python, Swagger, Zipkin**

Apr. 2014 - **Texas Advanced Computing Center**, Austin TX, USA.

Dec. 2013 **Graduate Research Assistant** (Dec. 2013 - Apr. 2014)

- Added docker integration to web portal digitalrocksportal.org to streamline the deployment process.
- Technologies got to work on: **Python, Django, Docker, Docker-compose, Celery**

Jul. 2013 - **Flipkart, Supply Chain Backend**, Bangalore, IN.

Jun. 2011 **Software Development Engineer II** (Jun. 2012 - Jul. 2013)

Software Development Engineer I (Jun. 2011 - Jul. 2012)

- Part of the Supply chain team which managed Warehouse and Logistics System.
- Built a new Supply chain framework from scratch which scales to 10X compared to the older system.
- Technologies got to work on: **Ruby, Java, Hive, MySQL**

Jun. 2011 - **IBM Software Labs, Informix DB**, Bangalore, IN.

Jul. 2009 **Software Engineer** (Jul. 2009 - Jun. 2011)

- Handled high priority issues for JDBC client driver for IBM Informix database.
- Was awarded for resolving an obscure deadlock situation, making Informix JDBC driver robust.
- Technologies got to work on: **Java, IBM Informix, IBM DB2**

GITHUB PROJECTS

- [prat0318/json_resume](https://github.com/prat0318/json_resume) - Details at prat0318.com/json_resume

- Ruby library to generate resume from JSON to HTML, pdf, md, tex formats.
- Around 1K stars on Github. Front page of Hacker News
- Current resume has been generated using `json_resume`.
- [prat0318/2048-tetris](https://prato318.com/2048-tetris) - Live at [prat0318.com/2048-tetris](https://prato318.com/2048-tetris)
 - Wrote a Tetris variant of 2048 game, 1M+ unique visits
 - Front page of Hacker News, mentioned in TechCrunch.
- [Yelp/pyramid_zipkin](#) - Zipkin integration with Pythonic services.
 - Built `pyramid_zipkin` as a wrapper to bind `Pyramid` based services with Zipkin.
 - Also wrote `swagger_zipkin` as a zipkin decorator over `swagger` clients.
 - Authored an engineering blog post: [Distributed Tracing at Yelp](#).
- **Major contributor to swagger packages** - `swagger-py`, `bravado`, `pyramid-swagger`
 - Wrote `swagger-py` as part of my internship at Yelp - currently used in production.

TALKS/BLOG POSTS

- Uber Engineering Blog [Apr. 2022] : [Securing Kafka Infrastructure at Uber](#)
- OpenZipkin/Distributed Tracing Conf [SF, July 2016] : [How Zipkin works @ Yelp](#)
- Yelp Engineering Blog [Mar. 2016] : [Distributed Tracing at Yelp](#)

EDUCATION

2013-2015 **Master of Science, Computer Science**, *University of Texas at Austin*, GPA: 3.73/4.

2005-2009 **Bachelor of Technology (Hons.), Computer Science**, *Indian Institute of Technology, Kharagpur (IN)*, GPA: 8.91/10.

ACADEMIC ACHIEVEMENTS

- Was among the top 5% in the institute merit list of the undergraduate batch at IIT Kharagpur.
- Secured 383rd Rank in IIT-JEE 2005, an exam conducted by Indian Institute of Technology(s).