

BHASKAR DIVYA

Bengaluru, India | <https://github.com/koldbyte> | <https://linkedin.com/in/bhaskardivya>

EXPERIENCE

Reflektion, Principal Engineer

JUN 2018 - PRESENT

- As a Senior member of the Data Platform team, The core responsibilities and contributions include
 - Designing and developing highly-available and fault-tolerant data pipelines using Kafka and Flink to capture and process click-stream data from retail websites
 - Designing and developing horizontally scalable Golang Microservices to provide data access to other teams and for other processing steps required in the data pipelines with <5 ms average latency using AWS DynamoDB and DAX
 - Building pipelines in Airflow and using the Snowflake warehouse for ETL and generating Analytics related to different entities viewable through the Portal UI
 - Performance testing and Scaling up of the Go Microservices & the data pipelines to handle 5x traffic (~500K events/min) during peak holiday season (Black Friday and Cyber Monday sales) and making it robust by handling out-of-order late events
 - Creating & Managing the horizontally-scalable setup of Airflow - including standardised development and deployment processes for the DAGs created by multiple teams
 - Building and Maintaining Monitoring dashboards for different systems and modules using Grafana and InfluxDB
 - Creating and Maintaining a standard Runbook for the Data platform Team for Production Issues
- Highlights from other major contributions
 - Designed and developed a Real Time Identity Graph system using Flink, Golang microservices and AWS DynamoDB
 - Designed and led the early development of a Cross Category Recommendation dataset generator system with custom DSL as config using ANTLR grammar in Golang
 - Implemented Reliability Patterns like hedged requests to reduce the tail latency of microservices by more than 87% (as described in the Tail of Scale paper)
 - Designed and built a config-based Alerts & Notification System for Customers in Python that uses Analytics data from Snowflake DB for checking various alert conditions
 - Built an internal Error and Metrics reporting infrastructure for standard usage among different teams
 - Migrated Flink Infrastructure from standalone infrastructure to Yarn and then later to Kubernetes
 - Migrated old Kafka deployment to the latest version of Kafka using Ansible and Jenkins
 - Worked on adding support for Product Fitments throughout different modules and creating analytics pipelines for customer reports
 - Built a core pipeline to validate the ETL data process in Snowflake DB on a daily basis using Airflow
 - Led multiple fixes and minor enhancements throughout different modules to decrease the number of High-Urgency Prod alerts by ~75%
 - Initiated effort to document and prioritise Technical Debts in multiple modules that includes cleanup of deprecated codes, updates to tools and technologies used, following best practices and other changes to increase reliability, observability, ease of development & deployments and to decrease costs
 - Initiate the team-level Documentation effort to help onboarding new engineers, improve documentation for different modules and help other teams know about any implicit or explicit dependencies on the Data Platform
 - Mentored new joiners to the Data Platform Team
 - Interviewed 20+ candidates for (Mid to Senior) Data Engineering roles in the Data platform Team
 - Helped the CSM team with Operational tickets from Customers and also set up a runbook for frequent issues/questions from customers
 - Helped the QA team as a go-to person of contact for any clarifications required for both technical or business logic

Sapient, Senior Associate Technology L1

JULY 2015 - MAY 2018

- SDLX : Sapient Data Lake Accelerator
 - Utilised open source tools like Teradata Kylo, Nifi, Spark, Hive, ElasticSearch
 - Built Nifi processors to add data rollback capability and Implemented a new table merge strategy for Change Data
- EnTRAC : Energy Transaction, Reporting and Accounting - Alternate Receipt Points Project
 - Migrated full code base from Java Struts 1.2 to Spring 4.3 in a record time of 60 days using a custom-developed JSP Tags parser reducing development time by ~20x
 - Built a fast and fully responsive Message & Alerts Inbox system with support for various filters, Search and detailed views
 - Created multiple scripts and tools to reduce efforts & increase the Developer/QA productivity by more than 10x
 - Improved performance of critical PL-SQL components by 100x for large inputs to support the automatic market transition on large scale
 - Won multiple "Star of the Month" Awards for excellent performance

Sapient, Junior Associate

JAN 2015 - JULY 2015

- Completed a 3 month classroom training under GMI (Global Markets Institute)
- Developed Equity Trading Simulation System from scratch using Spring & Hibernate in a agile sprint of 2 weeks
- Won "Supernova" Award for excellent performance

EDUCATION

Bachelor of Engineering (Computer Science), Birla Institute of Technology, Mesra, Ranchi

JULY 2011 - JUNE 2015

GPA: 7.98 / 10

- Co-Founder, ACM Student Chapter - Started the first ever Student Chapter in the state of Bihar, conducted various teaching sessions for Data Structures & Algorithms, organised various coding contests and talks by prominent tech leaders which resulted in 300% rise in participation in coding contests and teams getting selected for ACM ICPC regionals
- Represented the college in ACM ICPC Kharagpur Regionals
- Participated in & Won many programming and Web Development Contests under various Tech-Fests

SKILLS

Programming : Golang, Python, Java

Big Data : Apache Flink, Apache Airflow, Apache Kafka, Hadoop

Databases : DynamoDB, Snowflake, MySQL, InfluxDB

Others : AWS, Kubernetes, Docker, Bash, Git, Jenkins

Past worked on: Flutter, Dart, Scala, Apache Spark, Apache Nifi, Teradata Kylo, ElasticSearch, Apache Hive, Sqoop, C++, PL-SQL, Spring, HTML, CSS, JavaScript, ReactJS, VueJS, Apache HBase, Solr, Firebase, Qt C++, Java Swing

Courses & Trainings: Machine Learning (Coursera), 15+ courses on Scala, Big data and DevOps on PluralSight/Lynda, In-house trainings on Hadoop, Spark, Hive and HBase

PET PROJECTS

- Grocery Slot Alert App - Android App to notify when delivery slots become available on BigBasket (a problem when Covid pandemic started). Built in Flutter and Dart.
- Termoid : Linux Terminal as a desktop widget for KDE. Was ranked among Top 15 Plasma 5 addons based on the number of downloads. [KDE Store Link](#)
- CodeBackup : A Java GUI tool to download user submitted codes and problem statements from Coding contest Sites – Codeforces, SPOJ, and Codechef
- MySN (My Search Network) : A Search Engine in PHP & MySQL with Word Stemmer and Boolean search developed for a contest winning 1st prize
- Minesweeper : A clone in Java Swing

OTHER ACHIEVEMENTS

- Popular award - Reflektion's hackathon - Created a chatbot to answer business questions from DWH in natural language.
- Global rank within top 4% and Country rank within top 2.5% on Codechef(alias bhaskardivya) with 4 Stars ratings
- Global rank in top 15% overall and in top 10% in Algorithms track on HackerRank(alias Koldbyte) winning 4 silvers and 1 Bronze medals in contests
- Ranked among top 100 programmers on HackerRank Codesprint contest in India