SIDDHARTHA THOTA

Senior Software Developer — Distributed Systems & Data Architecture — ex-Shopify
(236) 509 7834 ⋄ sddhrthrt@gmail.com

SUMMARY

Senior Software Engineer with 10+ years of architecting high-throughput distributed backend services and full-stack tooling. Proven track record of owning mission-critical systems at Shopify-scale—optimizing APIs for sub-second latency, building resilient microservices in Golang, and streamlining developer workflows. Passionate about product engineering and systems design as an IC.

EXPERIENCE

Shopify *June 2020 - May 2025*
Senior Data Developer *Remote, BC, Canada*

- Architected and delivered the core execution layer of a declarative streaming framework in Flink and Google DataFlow, owning schema migration, data type safety, and incident tooling. Working alongside a principal engineer on the broader platform vision, reduced pipeline onboarding from weeks to days for 20+ downstream teams.
- Architected correctness verification to drive highly resilient deployment cycles that integrate seamlessly with the codebase, and provide feedback within minutes of deployments, replacing hours of hand-written verification tasks.
- Spear-headed SOC 1 compliance and unlocked enterprise adoption by provably reducing inaccuracies in financial analytics by fully re-architecting the reporting backend.
- Owned a multi-tenant Query Federation Service in Golang, orchestrating execution across BigTable, Parquet, and MySQL: reduced tickets from largest customers by 98%. Slashed p99 latency by 40% through predicate pushdown, adaptive shard pruning, and granular result merging.

StackAdapt *Apr 2019 - Mar 2020*
Data Engineer *Toronto, ON, Canada*

- Built demographic profiles for millions of user profiles using billion-scale distributed, real-time data sources.
- Improved and validated user recommendation for advertisement campaigns by using demographic profiles using Spark and billions of data points.
- A-B tested usage of demographic profiles on historical campaign audience data using statistical analysis to build POC for the project.
- Developed data pipeline in Golang to write real-time, distributed data into Aerospike, then into a data lake, processed and analysed using Spark.

Adobe Systems Inc. *Jan 2014 - Aug 2017*
Senior Software Developer *Bangalore, India*

- Pioneered an ML-as-a-service architecture for data preparation, model deployment and scoring (Azure, Docker with support for Python and Scala models). Used Mesos and Airflow with SparkML to offer a platform for other teams across the org, to observe and manage training, validation and deployment of machine learning models.
- Built framework to manage deployment of trained models and automatically generated inference APIs for use by dev teams, flush with UI, best practices and support.
- Built a platform for digital marketing clients to optimize spends on multiple channels, supporting multiple, parallel campaigns, environment variables and revenue prediction.
- Overhauled the Marketing Mix Model and improved the optimizer performance 250x by optimizing data reads and storage ops. Responsible for eliminating ETL bugs using well-tested, robust code.

EDUCATION

University of Toronto, Canada *Sept 2017 - Dec 2018*
MSc. in Applied Computing, GPA 3.925/4.

National Institute of Technology Karnataka, India *July 2010 - March 2014*
BTech. in Information Technology, GPA 8.47/10

TECHNICAL STRENGTHS

Backend & Systems	Golang, Python, Kotlin, Scala
	Postgres/MySQL, MongoDB, BigTable
	Docker, Kubernetes (K8s), Terraform, Linux, Git
Streaming & Data Infra	Kafka, Flink, Spark, BigQuery, Airflow, Apache Beam, DBT
Full-Stack & Tooling	React, JavaScript (ES6), Django, Vim, Shell scripting