

SPANDAN BEMBY

spandanbemby.com ◊ github.com/spandanb ◊ linkedin.com/in/spandan-bemby
spandan.bemby@gmail.com ◊ (650) · 283 · 5376 ◊ Aliso Viejo, CA

ABOUT

I am a highly self-directed software developer with strong design and implementation fundamentals. I thrive working on meaningful problems, with technical depth and broad impact.

SKILLS

Concepts Design, implementation, and evolution of distributed systems optimized for scalability, latency, reliability, consistency/correctness, user-experience, and other higher-level derived metrics. Relational and NoSQL databases. Performance analysis.

Languages Python, C#, SQL

Technologies PostgreSQL, SQL Server, Docker, Kubernetes, AWS, Azure, Google Cloud

EXPERIENCE

Microsoft

Software Engineer

Feb 2020 - Present

Aliso Viejo, CA

- Worked on evolution of platform for running performance related experiments and diagnosing performance issues
- Extended the platform to enable expression of new type of test scenarios, e.g. complex concurrent execution scenarios, which were used to performance benchmark the product (data warehouse SaaS)
- Leading a new effort to modernize the performance platform

Cruise Automation

Data Infrastructure Engineer

May 2017 - Aug 2019

San Francisco, CA

- Worked on end to end management of sensor data, including: storage, transformation, validation, indexing, and retrieval of data optimized for several key use cases
- Specifically, led design and implementation of catalog service which provided RESTful access to metadata about various datasets; and was critical to other derived pipelines, e.g. for model training and performance validation.
- Contributed to the design reviews, prototyping, and implementation of other parts of the data infrastructure stack, e.g. ingestion of real and simulated sensor data into data lake

Insight Data Science

Data Engineering Fellow

September 2016 - December 2016

Palo Alto, CA

- Designed a platform to give runners real time feedback on how their instantaneous speeds compare with other users'.
- Utilized Spark for batch and real time processing, Elasticsearch for geospatial queries, and Kafka for ingestion.

EDUCATION

University of Toronto

M.A.Sc. in Computer Engineering

Thesis: Distributed orchestration of heterogeneous resource types over public and private clouds

September 2016