Sumeet Bansal

LinkedIn: s-bansal Email: sumeetbansal@gmail.com Cell: (408) 250-3097

Education

University of California, San Diego — La Jolla, CA
B.S., Computer Science; Minor in Mathematics. — *Graduated March 2020*

Skills & Qualifications

Languages – Java, Python, SQL, TypeScript/JavaScript, Rust, Shell
Technologies/Frameworks – CI/CD, Docker, Node, React, Cassandra

Experience

Pre-seed startup — New York, NY
CTO, Co-founder — *March 2024 to Present*

Sabbatical — *September 2023 to February 2024*

Palantir — New York, NY
Software Engineer — *May 2020 to September 2023*

- Developed the time series product, which ingests large datasets (up to tens of Terabytes each) containing millions of unique time series and enables users to execute complex queries composed of transforms and aggregations against those time series within seconds.
- The product supports business-critical workflows at global institutions like oil-and-gas supermajors, automotive and aerospace manufacturers, and national health agencies. For example, Airbus instruments its aircraft with thousands of sensors recording measurements at millisecond frequency, for every flight, across their fleet, and uses our time series product to manage that volume of sensor data and plan preventative aircraft maintenance. Clients frequently use time series to monitor the health of their assets in real-time, e.g. British Petroleum and its energy platforms, several national health agencies and their COVID responses.
- I worked extensively across the product stack: multiple Java services, a Rust service, a Python library, and a Typescript/React front-end.
- Owned streaming workflows for time series, a functionality highly-demanded from current and prospective customers. The architecture for streaming workflows barely existed as a skeleton; I shipped tens of thousands of lines of code to production, wrote documentation, added monitoring and alerting, collaborated with early adopters to iron out issues and address pain points, released the feature, and enabled our largest customer to migrate from their custom streaming solution. As clients scaled their streaming usage, I profiled and analyzed runtime to improve stability (e.g. memory usage) and optimize performance.
- Enabled usage-based pricing for time series and unlocked millions of dollars in new contract value by implementing compute reporting for time series services and updating all significant time series clients in the Foundry platform one-by-one to pass along compute attribution to time series.
- Reduced cloud costs by several hundred thousand dollars annually by building the next-generation of time series architecture, which would be more tightly integrated with the Foundry platform and write orders of magnitude less metadata, eliminating the need for dedicated Cassandra clusters on hundreds of customer deployments.
- Contributed to the support rotation, where I responded to pages and investigated and fixed time series issues in the field, such as security vulnerabilities, product bugs, configuration errors, performance issues, and complex system failures. I closed over 1000 pages. While on support, I was responsible for maintaining strict SLAs across dozens of customer deployments.

Software Engineer Intern — *June 2019 to September 2019*

- Worked on a a permissioned registry service to store user-authored code called Functions, and Function metadata and specifications (e.g. input/output data types, documentation, location of Function code). Designed and implemented a search API for published Functions in Java using Cassandra and Elasticsearch.
- Developed a custom TypeScript logger conforming to the internal service specification, e.g. integration with internal authentication service, tracing by Function execution context, support for log rotation.

IBM Watson Health — Milwaukee, WI
Software Developer Intern — *June 2018 to September 2018*

Actiance — Redwood City, CA
Engineering R&D Intern — *June 2017 to September 2017*

Zscaler — San Jose, CA
Emerging Technologies Intern — *June 2015 to August 2015*