

EXPERIENCE

- **Citadel (Fixed Income and Macro)** London, United Kingdom
Senior Software Engineer → Engineering Manager *April 2023 - Present*
 - **Research Data Platform (Storage & API Layer):**
 - * Sole architect and developer of a centralized research data platform, now adopted by **all systematic desks**
 - * Built on **PostgreSQL and Arrow Flight** over gRPC with bidirectional streaming APIs directly from database cursors; deployed on **Kubernetes**
 - * Introduced sharded per-desk architecture for independent scalability, enabling heavy-usage teams to run their own stacks, with built-in permissions and TTLs to protect desk-specific IP and restrict the growth of unused data
 - * Designed and implemented the research storage API, which became the preferred interface for department-wide data access, prompting migration of other storage systems to use this API due to its intuitive design and streamlined workflow
 - **Market Data Infrastructure (Tick Pipeline):**
 - * Sole developer of a tick data ingestion pipeline ingesting data from multiple market venues in real time, processing **30B+ market events/day** across **27M+ unique instruments**
 - * Built a low-latency, high-throughput system using **Kafka, Redis, and C++** with concurrent and lock-free design patterns; **primary expertise is in system design and Python-based data engineering**
 - * Designed for production resilience, scalable for increased market data needs, and seamlessly rolled out without disruptions to existing usecases
 - **Leadership & Platform Transformation:**
 - * **Promoted to Engineering Manager**, leading two teams (6 engineers) focused on data onboarding and data platform infrastructure
 - * Modernized the data onboarding platform, overseeing migration of **5,000+ ETL jobs** and reducing onboarding times from 5+ months to **<30 days**
 - * Architecting and leading the team developing a department-wide data quality and discovery platform, providing validation checks, dataset health monitoring, metadata management, multi-storage support, and data arbitration
- **University of St. Andrews** United Kingdom
Career break for full-time study (details in Education section) *September 2021 - September 2022*
- **Bloomberg LP** London, United Kingdom
Senior Software Engineer *2019 - September 2021*
 - **Taxonomy and Metadata System:** Infrastructure team developing a taxonomy and metadata system to support various asset classes at Bloomberg. Python, GraphQL, Redis, RabbitMQ, Solr, RDF, PostgreSQL, Data Pipeline
 - * **Architected ingestion pipelines** that handle ingestion of metadata for datasets such as commodity index tickers.
 - * Added **support for sharding** to increase the speed of the ingestion pipeline. Split up the initial ingestion into 8 shards that are running distributed across an entire cluster.
 - * Implemented node overlay and concatenation executed on RDF graphs.
 - * Worked on **rearchitecting the ingestion pipeline** from a synchronous pipeline to a file based ingestion.
 - * Implemented services in GraphQL to avoid overfetching and underfetching on the existing Bloomberg specific microservice framework.
 - * Worked on a side project, debugging tool, to help find metadata for given entities regardless of the source of provenience, using a **topological sort algorithm on service APIs and RDF Graphs**.
 - * Team lead backup for a team of 7 engineers.

Software Engineer *2016 - 2019*

 - **Real Time Analytics:** Application team, maintaining a large variety of analytics, integrated in different products across Bloomberg. Processing several billion trades per day. C++, Distributed, Realtime
 - * Implemented Average Volume at Time for non equities and Top Exchanges in the past X minutes.
 - * Implemented a **load testing facility** that is used to simulate the impact of market open on the system to assess performance and capacity. I identified a bottleneck in the system and by fixing it, subscriptions are being processed 3 times faster.

- * **Increased availability** of an analytic by implementing fallback to disk in case the database connection cannot be established. This prevents the 1 minute downtime per week when the database cluster is being rebooted.
- o **Recommendation Disclosures and Trade Ideas:** Python, Typescript, Microservices, Event Sourcing
 - * Implemented asynchronous messaging using RabbitMQ to allow distribution of recommendations. Added support to publish messages, implemented consumers and set up the middleware.
 - * Reduced the codebase by 8% by removing around 100.000 lines of dead or unused code.

Intelligent Software Systems

Timisoara, Romania

Software Engineer

2013 - 2016

- o Built a desktop-to-CRM sync tool (workbooks.com) and contributed features and bug fixes to the FAFSA questionnaire.

EDUCATION

University of St. Andrews

United Kingdom

MSc Statistics

Distinction, Deans' List, grade 17.9/20, September 2022

- o **Modules:** Markov Chains, Bayesian Statistics, Multivariate Analysis, Statistical Modelling using GLM/GAMs, Quantitative risk analysis, Machine Learning
- o **Dissertation:** An Ensemble of HMM and KNN for signal generation in horse betting markets. Grade 19.3/20 - Highest grade in my cohort

University of Sheffield

Remote, United Kingdom

Graduate Certificate in Statistics

Distinction, grade 90/100, July 2021

- o **Modules:** Mathematical Methods for Statistics, Probability and Probability Distributions, Basic Statistics

Politechnica University of Timisoara

Timisoara, Romania

B.S.E in Computer Science

July, 2014

- o **Modules:** Software Engineering, Computer Architecture, Algorithms, Software Development Fundamentals, Design Patterns, Programming Languages, Databases

SKILLS

Coding

- o **attrs-strict <https://github.com/bloomberg/attrs-strict>:** Created and open-sourced a library that allows runtime validation for attrs dataclasses in Python.
- o **bloomberg-bas-middleware:** Maintainer and creator of an inner source library used for Python services at Bloomberg. Created Flask-like middlewares that are used in over 10.000 services running in Python 3 at Bloomberg. Any new service has by default the functionality opted-in.
- o **Python Guild:** Part of a group of 18 engineers leading the direction and adoption of Python at Bloomberg, led the Design Reviews Working group, helping teams to bring projects to a broader audience.

Public speaking for large audiences

- o **"How to migratate services to Python 3"** - Attendance of over 200 engineers.
- o **"Modeling the commodity world"** - Presentation on the architecture of the system I worked on and the different architectural decisions taken in the past years. Attendance of over 150 engineers.

Mentorship

- o Mentored interns and junior/mid engineers at Citadel and Bloomberg, conducted interviews, ran workshops on Python 3 migration, and supported new joiners in ramping up and contributing to inner-source projects.

Architecture and technology:

Large scale distributed systems, Data pipelines, Microservices, GraphQL, PostgreSQL, Redis, Kafka, Solr, RabbitMQ, Event Sourcing, Git, RDF

Data onboarding:

External data ingestion via APIs, SFTP, web scraping and PDF parsing; sources include Bloomberg (PerSecurity/BLPAPI), Refinitiv, S&P; ETL pipeline management at scale

Languages:

Python (primary), C++ (production use), JavaScript/TypeScript (basic), SQL

Mathematics:

Probability theory, Markov chains, Bayesian models, GLMs, statistical inference