

Quentin Bouget

Amsterdam, Netherlands

[LinkedIn](#)

Online version of this document: <https://iamlookingforajob.dev>

Experiences

Software Engineer, IMC, Netherlands

Feb 2021 - Apr 2026

Strategy developer

Developed a vectorized backtesting system used to analyze and trade systematic volatility strategies. The system supported two sources of pricing data: pre-computed and stored in a relational database -- which e.g. allowed researchers to backtest >200 points on SPX's volatility surface over 9 years at a 5min delta-hedged frequency in about 10min -- and on-the-fly pricing using third party historical data and a vectorized Black-Scholes pricer -- which allowed researchers to backtest strategies on FX for which IMC did not have historical data in-house.

Worked for 3 years alongside a systematic volatility trading desk, providing them with technical solutions to their trading visions. Practically this meant: research tooling such as the aforementioned backtesting system, integration with IMC's data and trading infrastructure, technical and operational support, as well as research projects of my own, when appropriate.

Worked for 1.5 years in a more global role, developing and rolling out a general purpose signal framework.

Spent my first 6 months migrating all the trading desks in Amsterdam from bare-metal ipython servers to a jupyterhub instance in k8s, as well as a custom cron service for jupyter notebooks, with a track record of 0 operational incident.

Systems engineer, CEA-DAM DIF, France

October 2016 - Nov 2020

Lead developer

Rewrote [RobinHood](#) -- the most critical data management tool used in my team -- for an increase in performance ranging from 20x to 1000x. ^{ref} RobinHood allows its users to collect any filesystem's metadata and store it in a database, making it easily and efficiently searchable.

Package maintainer

Packaged and maintained [Lustre](#) -- an open source distributed filesystem used by more than half of the top 100 supercomputers, including #1 -- for a private RPM-based distribution, which is deployed on production clusters of thousands of nodes.

Open source contributor

Conceived and implemented a prioritization mechanism for Lustre's HSM (Hierarchical Storage Management) that reduced first access latency, from hours down to seconds.

Improved the Lustre HSM test suite to complete 1.8x faster for successful runs, from ~1h10 down to ~40min. ^{ref}

Designed and built from the ground up [parablos](#) and [mercurio](#) -- respectively, a single-node/local parallelization framework in python based on `multiprocessing`; and the data transfer application built on top of it -- to parallelize data transfers on high bandwidth networks, leading to near optimal network utilization.

Internship director

Managed three interns: two for 3 months in 2018, and one for 6 months in 2019.

Intern, CEA-DAM DIF, France

June - Sept 2015, February - August 2016

Contributions to the Lustre filesystem.

Wrote patches to extend the configuration syntax of a large-scale monitoring tool, making it possible to use nodeset-like notations (eg. `host[0-100]`).

Education

ENSIMAG, Computer Science and Applied Mathematics, Information Systems Engineering, Grenoble

2013 - 2016

Technical Skills

Languages

Python, Bash, C, English, French

API design

I am committed to designing APIs that are idiomatic, intuitive, and user-friendly.

Miscellaneous

Presented progress reports at an international workshop, twice. ¹ ²

Contributed to the organization of said international workshop, twice ¹ ². Lead the transition to a virtual event in 2020 due to the pandemic. ¹