# Brendan MacDonell

brendan@macdonell.net

San Francisco, CA

## Summary

Staff engineer with significant experience leading complex products to launch and beyond. At Sight Machine, I led the team that built Sight Machine's streaming ETL product, plus the data acquisition product that moves customer data into our platform. These products — plus a suite of analytics on the transformed data — are Sight Machine's core offering, driving over $10M in annual recurring revenue.

## Technical Skills

| | |
|---|---|
| **Languages** | Java, C, C++, SQL, Python, OCaml, Javascript, CSS, sh, awk |
| **Technologies** | AWS, GCP, Git, Kubernetes, Linux, perf, PostgreSQL, Prometheus |

## Professional Experience

**Sabbatical** Jul 2024 — Jan 2025

After nine years at Sight Machine, I took six months off to recharge and work on personal projects related to high-performance data processing. Ask me if you'd like to know more!

**Staff Software Engineer** Sight Machine Jul 2019 — Jul 2024

Most recently, I led the Factory BUILD project, creating a streaming ETL product to model manufacturing data and transform it into an analyzable form. In particular, I:

- **Pitched and led implementation of the core UI features for the ETL product**, including a drag-and-drop interface for connecting data flow operators, interactive data exploration, and a macro system for composing reusable groups of operators.
- **Designed and implemented a high-performance distributed streaming ETL engine using Java and RocksDB**, with data flow operators capable of processing 250k wide records (of 1000+ fields) per CPU core per second. The unusual ordering and lateness requirements for manufacturing data prevented us from using an existing framework like Flink — instead, I built a new system from scratch using concepts from Flink, Millwheel, change data capture, and a novel approach to streaming stateful transformations.
- **Led the team to build dozens of types of data flow operators, plus an SDK to let users build their own** — all without having to worry about implementing per-operator support for exactly-once processing, state storage, or handling out-of-order data, as these concerns are handled by the core ETL engine.
- **Created a quality system that eliminated almost all regressions.** Manufacturing customers demand high reliability — a few bugs in a year can lose a customer. We relied on careful code review, unit testing, property-based testing, end-to-end testing, and nightly smoke tests against customers' data pipelines. As a result, we shipped only one (minor) regression between launching in July 2021 and my sabbatical in July 2024.

In addition, I:

- **Pitched and architected Sight Machine's externally-facing SQL interface**, allowing users to access their transformed data in tools like Excel and Tableau.
- **Acted as manager for the Factory BUILD team** from July 2020 until getting headcount to hire a new engineering manager in May 2023.

**Senior Software Engineer** Sight Machine Jul 2016 — Jul 2019

Improving Sight Machine's platform using Python and PostgreSQL. Among other projects, I:

- **Led the design and implementation of Factory CONNECT,** a Python-based product for acquiring data from dozens of different manufacturing systems and products.
- **Designed and implemented a versioned configuration store** to let internal and external data engineers configure data transformations in a user interface, test them on staging systems, and merge them to production. This replaced a workflow with no notion of versioning or merging, which often caused changes to be lost.
- **Designed and built Sight Machine's authentication system**. The system allows different groups of users (e.g. customers, data integrators, and Sight Machine employees) in the same tenant to sign in by federating multiple OpenID Connect providers.

**Software Engineer** Sight Machine Apr 2015 — Jul 2016

Improved the performance, reliability, and scalability of Sight Machine's platform using Python, MongoDB, and PostgreSQL, particularly related to data acquisition.

- Redesigned the protocol used to upload data from factories to Sight Machine's cloud platform. The new implementation improved throughput by 10×, solved message ordering issues, and allowed data to reliably pass through the great firewall of China.
- Rewrote the data acquisition code for Mitsubishi PLCs to fix data corruption issues. The new code was simple enough that solutions engineers began to contribute to it.

**Software Engineer** E la Carte Jun 2014 — Apr 2015

Developed and maintained high-availability web services to manage and monitor 30 000 restaurant tablets using Python, Django, PostgreSQL and Redis. Also carried out several of my own initiatives, including:

- Discovered and fixed a security architecture flaw that made all restaurants' point-of-sale TLS certificates interchangeable.
- Created an automated system to reflash tablets returned due to software defects, cutting hardware expenses by 4% (roughly $300 000.)

**Software Developer** Sheepdog May 2013 — Sep 2013

Wrote software for a variety of clients. For example, I created a Python-based system to import Sharepoint and eRoom sites into a client's new document management system.

**Software Developer** E la Carte Jan 2012 — Jan 2013

Contributed to all several products, from maintaining restaurant tablet software written in C++ to developing web services with Python and PostgreSQL. Among other things, I was responsible for:

- Building an online menu editor, reducing the time to create a menu from weeks to days.
- Integrating the Micros 3700 point-of-sale system for ordering and payments.
- Unifying several divergent point-of-sale integrations into a cohesive Python codebase.

**Software Developer** Sheepdog May 2011 — Sep 2011

Contributed to several projects with Django and PostgreSQL, including a reporting engine for healthcare survey data and a backend for seat-based subscriptions.

**Software Developer Intern** Sheepdog May 2010 — Sep 2010

Worked on gTrax, Sheepdog's time tracking product (built on GWT and AppEngine.)

## Education

**Carleton University** Ottawa, Ontario

*B. Eng. in Software Engineering (4.3 GPA)* Sep 2009 — Apr 2014