

# Tom Mitchell

📍 Seattle 📞 (425) 533-1632 ✉️ [tom@tom.kiwi](mailto:tom@tom.kiwi) 🔗 [tom.kiwi](https://tom.kiwi) 🌐 [tomnz](https://tomnz.com)

Accomplished software engineering leader with 16+ years of industry experience and a strong track record of designing and delivering complex backend systems and infrastructure as an individual contributor, tech lead, and manager. Experienced in all stages of a rapid software development lifecycle, from design through implementation, testing, operability, and maintenance. Passionate about beautiful code, the power of software, and transformative technologies.

## Competencies

Versed in a wide variety of languages and frameworks; fast learner when presented with new platforms.

Strongest current expertise in Python, Go, TypeScript, C, and SQL; utilizing React, Django, Terraform, PostgreSQL, SQLAlchemy, Bazel, Docker, Mesos, Protobuf, etc.

Strong background in complex backend systems, observability and testing, databases (relational and non-relational), team-oriented development practices (source control, bug management, documentation), and managed cloud services (AWS EC2, RDS, etc).

Three years of experience in Tech Lead and Engineering Manager positions.

## Education

### Bachelor of Engineering

2005 – 2008 • *University of Auckland*

Specialized in Software Engineering.

800+ hours of industry experience pre-graduation.

### References

Available upon request.

## Experience

### Staff Software Engineer – Oscar Health

Nov 2015 - Jun 2022 • New York City, Seattle (remote)

Founding member and leader on Oscar's next-generation claims processing engine - critical technological infrastructure for a tech-oriented health insurer. The system processes claims in under 10s (P90), with an industry-leading auto-adjudication rate (no human intervention) of 95%+.

**Staff Tech Lead and Manager (2020-2022):** Returned to Claims System team in leadership role with five engineers.

- Formulated strategic long-term goals, managed progress and milestones, and provided technical leadership in a complex business domain. Business areas covered included authorizations, cost-sharing calculation, benefit determination, accumulator ingestion, repricing, adjustments, and many others.
- Oversaw complex multi-year technical migration from a legacy third-party claim processing system while ensuring business continuity.
- Supported business-wide efforts for +Oscar and Cigna partnerships, Medicare Advantage, and new state market expansions.

**Staff Software Engineer (2019-2020):** Returned to an IC role in Engineering Infrastructure.

- Built an internal tool from scratch that allows users to query any Oscar gRPC service using a React form generated dynamically according to the corresponding Protobuf request schema. The tool is used by engineers and non-technical users alike.
- Contributor to internal infrastructure designed to replicate data from a MySQL or PostgreSQL replication log into various sinks, such as BigQuery or Kafka. This centrally-developed tool replaced 4+ bespoke solutions across the company.

**Staff Tech Lead and Manager (2018-2019):** Founded new Plans and Benefits team to meet identified business need.

- Built a brand new backend team of five engineers in a year and supported professional and personal development for each.
- Designed and executed multi-year technical roadmap for stewarding Plans and Benefits data across Oscar including data models, services, configuration tools, and benefits adjudication business logic. Collaborated with insurance experts in Risk and Data teams.

**Senior Software Engineer (2015-2018):** Founding member of the Claims System team.

- Key decision-maker and engineer for foundational features of the system, such as accumulators, pipeline orchestration, and generic object lookup storage. Utilized a microservices architecture with gRPC services written in Go and Python, connecting to PostgreSQL for transactional storage and Kafka + Confluent for queueing.
- Developed a custom domain-specific language for business logic, which combined Lua with a custom type-checked language for safety. This language is utilized by non-engineers / insurance domain experts to configure the system and write unit tests.

Throughout tenure, was a prolific contributor to development infrastructure, including work outside core projects.

- Designed, implemented, and evangelized a novel storage platform for bitemporal data on top of PostgreSQL, utilized by multiple teams at Oscar. Data source schemas are defined in YAML, which is code-genned into SQL, ORM, and service definitions. The platform manages entity history transparently, while database constraints prevent bitemporal inconsistency.
- Voluntarily championed good engineering culture - through changes to onboarding processes (documentation, introductory sessions, codelabs, development environment), improvements to interviewing (trialing new questions and techniques), writing new technical documentation, presenting tech talks, etc.

📍 Seattle 📞 (425) 533-1632 ✉️ [tom@tom.kiwi](mailto:tom@tom.kiwi) 🔗 [tom.kiwi](https://tom.kiwi) 🌐 [tomnz](https://tomnz.com)

## Experience Cont'd

### Founding Engineer - Ravel Data

*Jun 2022 - Present • Seattle (part-time)*

Recent part-time engagement as sole founding engineer for a web data extraction and competitive insights startup.

- Built out Python codebase from the ground up, including web crawler, data models, supporting scripts, documentation, and engineering infrastructure. Designed to be extensible, scalable, and understandable for easy transition to a new technical cofounder.
- The fully-functional data product has been launched and sold to multiple B2B customers, totalling ARR of \$100k+ to date.
- The crawler is an extensible platform built on top of Scrapy, backed by PostgreSQL. With minimal oversight, it regularly evaluates 9000 identifiers across 100k+ websites, using 400 parallel shards. At this level of parallelization, it can complete a crawl in about 6 hours.
- Terraform is used to provision all AWS resources, utilizing a highly scalable architecture on top of Batch and Step Functions, writing results to RDS. Bazel is configured as the build system for testing, packaging, and deploying code to AWS Lambda and Docker images.
- Low-code internal tooling with Retool supports the configuration and operation of the system by non-technical users without significant engineering investment - allowing exploration of new identifiers to boost sales prospects, and monitoring the status of running jobs.

### Software Development Engineer II - Microsoft

*Oct 2010 - Aug 2015 • Seattle*

Worked on multiple machine learning and infrastructure problems - initially for the Contextual Relevance group in Bing, followed by Cortana's Proactive card ranking.

- Designed and implemented a generalized system for making Cortana personalized ranking features (computed daily from user logs) available online with minimal effort from data scientists, allowing highly nuanced features across multiple user dimensions (e.g., demographics, location). Prior to this system, bespoke deployment solutions were created for each new feature set by engineers.
- Regularly collaborated with individuals in MS Research to bring new research ranking techniques to fruition in the Bing live site. One example was Category-Promoted Search with Paul Bennett, which sought to categorize user intent over time, then reorder web results according to similarity to the user's profile, leading to more relevant ranking for ambiguous queries.
- Migrated two critical personalization features from a C++ monolith into C# with a completely new service paradigm - requiring redesign and rewrite of multiple components.
- Designed and built a feature known as Ranking Stability, which caches personalized query results per user to stabilize Bing's web results between queries. This critical issue required a fast turnaround time on implementation and was successfully deployed within several days, requiring no further maintenance.
- Developed a centralized service to expose user information and models from various parts of the Bing stack in a single easy-to-consume schema for partners.

### Software Engineer - Canary Data Solutions

*Nov 2006 - Oct 2010 • Auckland, New Zealand*

Developed custom data-oriented business software in a small software shop (<10 engineers), owning several small business customer relationships end-to-end.

- Appointed as project lead on two ongoing support contracts, communicating with clients daily, and overseeing new feature development and bug fixes.
- Led account for a high-volume vehicle dismantling operation in Japan, consisting of a public website and multiple internal systems handling everything from vehicle purchase to dismantling, inventory, and customer invoicing. ASP.NET applications interfaced with a central MSSQL database.
- Championed modernization of company engineering practices, including replacing Subversion with Mercurial, and pushing for the adoption of unit tests on all projects.