

Sam Wanuch

Santa Clara, CA | swanuch@gmail.com | swanuch.com | github.com/TheNotChosenOne

Summary

Architect at Adeptmind, growing with the company from an intern among fewer than 10 engineers to setting company standards for system design, AI usage, language choice, and efficiency for more than 50 engineers. Brought multiple systems from drawing board to production products both independently and leading a team.

Experience

Senior Software Architect, Adeptmind

Feb 2022 - Present

- Owned technical feasibility evaluation for company-wide migration to a managed cloud vendor, migrating core search services to cost/performance parity on production load; requiring the vendor to implement fixes and missing features across autoscaling, load distribution, and billing.
- Authored company architecture guidelines; ran reviews on major new system designs and designs referred by the CTO and director of engineering.
- Cut core Python multi-service search p95 latency from 282 to 131ms, reducing CPU usage by 30% and company-wide compute bill by 3%.
- Built a system to automatically connect source code, infrastructure, configs, and metrics using Docker image/dependency analysis and LLMs, allowing AI agents to find cost savings company-wide, reducing compute spend 13% after team leads implemented recommendations.
- Introduced Go and Rust to production, growing Go to 12% of active repos with Rust limited to performance-critical rewrites, by building the first systems in each, teaching language-selection, running intro sessions, and recommending them in design reviews.
- Reduced CockroachDB database from 40TB to 0.5TB cutting storage cost 20x through automated cold storage archiving with resumable crash recovery and checksum validation before delete.

Senior Software Engineer, Adeptmind

Jul 2020 - Feb 2022

- Designed a live in-store inventory and price scraping platform and led development of 290 retailer integrations.
- Built a Go crawl platform around freshness targets for 3.6M pages per day, with hot-swappable per-site processing graphs, runtime-submitted embedded Python processors, and provenance-tracked reprocessing to re-derive link frontiers and extracted data for fast client integration cycle time.
- Built and ran data team interviews with automated per-candidate take-home trial environments; designed and ran distributed systems design interview for general engineering; 3 of 10 hires rose to team lead or higher.

Software Development Engineer, Adeptmind

Jul 2018 - Jul 2020

- Designed and implemented a distributed web crawler with product extraction; led a team of 6 integrating 330 retailers, with automated per-crawl data validation flagging integrations for repair.
- Designed and shipped a multi-tenant search platform for hundreds of Shopify clients with at-least-once processing of idempotent zero-downtime catalog operations.
- Designed and built a catalog ingestion pipeline for a key client's 42M documents, with snapshots ingested into PostgreSQL within 10 minutes and streaming updates reconciled into a point-in-time catalog.
- Made core search service multi-tenant with parallel index queries to cut p50 search latency by 25%.

Software Development Engineer Intern (Co-op), Adeptmind

Sep 2017 - Dec 2017

- Wrote a Python distributed crawler on Kafka with domain-based partitioning and product data extraction.

Software Development Specialist (Part-time Contract), Sentrex Wind Services Apr 2020 - Dec 2020

- Designed and built battery-aware lifecycle interface in C++ for remote solar-powered devices on ARM to load shed hardware and services at low power.

GPU Development Tools Intern (Co-op), NVIDIA Jan 2017 - Apr 2017

- Optimized data processing tool, cutting peak memory 6x and overnight runs to 21 minutes.

Software Developer (Co-op), Big Viking Games May 2016 - Aug 2016

- Rebuilt Java asset processing pipeline with incremental rebuilding and multithreading for 3x faster full rebuilds.

Software Developer (Co-op), Crank Software Sep 2015 - Dec 2015

- Rewrote the embedded XML loader in ANSI C89, cutting its peak memory use 87%, to less than the size of the file it loads.

Associate Programmer (Co-op), Relic Entertainment Jan 2015 - Apr 2015

- Built content creation tools for Company of Heroes 2 that let designers finish common tasks 80% faster.

Skills

Languages: Rust, Go, Python, C++, C

Platforms: Linux, Kubernetes, Docker, Google Cloud, AWS, PostgreSQL

Education

Bachelor of Mathematics, Honours Computer Science, Minor in Combinatorics and Optimization,
Japanese Language 1 Certificate, University of Waterloo

2018

Projects

Demos for some projects at swanuch.com

- Personal web platform: streamed video recording, game server, webnovel to audiobook pipeline; rewritten from Go to Rust with idempotent lease-based task queues over PostgreSQL; operated solo on Kubernetes since 2019.
- Core War engine whose Cranelift-based JIT turns game memory into self-modifying machine code.
- Data model built on wasm PostgreSQL (PGlite) with copy-on-write dataset branching and allowing the same query to run on server or client with transparent delivery of relevant table rows.
- Real-time ARM920T microkernel in C with 0.6 microsecond context switches.
- Live compilation of Rust snippets into a call graph with dynamic rewiring and fused recompilation of hot call chains.
- Rust SAT solver with clause learning and per-variable satisfiability queries for logic puzzles on web/Android.
- Rust port of PBRT v4 for progressive in-browser rendering with per-ray arena allocation.
- Rust ECS game engine without entity IDs for Wasm/WebGPU.
- Visualizer for custom 16-bit CPU with word-level gate stepping, assembler, and peripheral simulation.