

Paul Amonson

Personal Information: Hillsboro, OR | 503-680-9193 | paulamonson@gmail.com

Linked-In: <https://www.linkedin.com/in/paulamonson/>

GitHub: <https://github.com/paul-amonson>

Summary

Senior Software Engineer with 15+ years building backend and distributed systems that other teams depend on, from HPC cluster control systems at Intel to company-wide telemetry infrastructure at Anaconda. Specializes in the connective layer between systems: the schemas, APIs, and data pipelines that let services, teams, and languages work together. Proven record of measurable performance gains, including cutting CI runtime by ~80% and tuning PostgreSQL and query engines for Intel hardware. Pragmatic adopter of AI-assisted engineering, keeping rigorous human review in the loop.

Core Skills

- *Programming:* Java, Python, C++, Swift, TypeScript, Rust (intro), Go (intro)
- *Backend Development:* REST APIs, microservices, middleware, Linux services
- *Telemetry & Observability:* OpenTelemetry schema design, collectors, exporters
- *Distributed Systems:* HPC control systems, cloud optimization
- *CI/CD & DevOps:* GitHub Actions, caching, automated testing, pipelines
- *Architecture & Design:* Software architecture, API design, secure systems
- *Data Engineering:* Data modeling, ETL pipelines, PostgreSQL
- *AI-Augmented Engineering:* Debugging, security scanning, performance tuning
- *Team Collaboration:* Agile, mentorship, cross-functional teamwork

Professional Experience

Anaconda Inc. as a Senior Software Engineer from 4/2025 to 8/2026

- Designed an OpenTelemetry-based telemetry schema and backend pipeline, including storage models, collectors, and exporters.
- Developed multi-language telemetry libraries (Python, TypeScript, Rust) to support a unified observability API across the company.
- Implemented new REST endpoints and integrated SpiceDB authorization logic into the core resource manager API.
- Researched and prototyped a potential migration from Connexion (OpenAPI) to FastAPI, producing technical recommendations.
- Reduced CI runtime from 35–40 minutes to 8 minutes through AI-assisted analysis, test optimization, caching, and refactoring slow test cases.
- Used AI tools to surface bugs, security issues, and architectural improvements while maintaining full human oversight.

Intel Corp as a Software Engineer and Senior Software Engineer up to 10/2024

As a Software Engineer and Senior Software Engineer, I was responsible for designing, developing, and maintaining software applications, ensuring high performance and scalability.

- Optimized PostgreSQL, Velox, and Gluten for Intel architectures, improving cloud performance using flame graphs and custom profiling tools.
- Designed and scaled a Unified Control System for HPC clusters, enabling telemetry ingestion, transformation, RAS event processing, and distributed storage across large node environments.

- Contributed to Ecma standards for low-power data transfer, including technical requirement definition and security evaluations of Intel's implementation.
- Developed and enhanced distributed systems to improve performance, reliability, and scalability for internal cloud and HPC platforms.

Projects

Daily Food Planner (SwiftData / SwiftUI, iOS and macOS)

- Built a multi-platform meal-planning and nutrition application with CloudKit synchronization and offline support.

Shopping Application (SwiftUI)

- Developed two versions: one using CloudKit for multi-device syncing, and one using Redis on a Tailscale private network to support multi-user collaboration

Custom Golang Backend with PostgreSQL

- Built a local-first backend to replace CloudKit, including offline synchronization logic and application migration.
- Leveraged AI systems to accelerate backend architecture design, troubleshooting, and application porting.

Education

Bachelor of Science, Computer Science (Minor in Mathematics) University of Idaho