

# Nathan Baune

Platform & Systems Engineer

Remote • Open to Relocation  
nathanbaune@gmail.com | Portfolio | GitLab | LinkedIn

## Technical Stack

### LANGUAGES

C++20, Go, Python, C#, TypeScript, Dart, SQL

### DATABASES

PostgreSQL (Supabase), SQLite, MongoDB

### COMPUTE

GPU/CPU kernels, CUDA, Metal, Vulkan, SIMD

### BACKEND & INFRA

Node.js, Docker, Linux, WebSockets, AWS Lambda

### FRONTEND & UI

React, Flutter, Electron, Unity, Vite, Skia

### ARCHITECTURE

Control planes, async jobs, runtime schedulers, codegen

## Summary

Platform and systems engineer with 9 years building distributed systems, real-time engines, and research tools and systems. Built software and infrastructure for 9 labs across three institutions; co-founded one startup (Proprio, STTR-funded), founded another (Gothic Grandma). PhD (Washington University School of Medicine) with specialization in computational and systems neuroscience. Currently building MUSE—a 100k+ entity simulation platform—in off-hours.

- End-to-end ownership: architecture, implementation, databases, runtime, observability
- Systems depth: C++ simulation engines, Go backends, GPU compute, real-time processing
- Platform thinking: schema-driven architectures, code generation, control planes
- Production experience: NIH-funded research systems, STTR-funded startup, active clinical use

## Education

PhD — Washington University in St. Louis School of Medicine, specializing in computational neuroscience • 2020

BS — Experimental Psychology, University of Oregon • 2013

## Experience

### LEAD PLATFORM ENGINEER • Precision Neural Engineering Lab, Emory University

2021—Present

Backend systems engineer building research platforms, ML pipelines, and experimental control systems for NIH- & NSF-funded clinical research. Lead rapid prototyping for four active labs at the Emory Rehabilitation Hospital.

- Built **MR.Flow**: GUI orchestrating neuroimaging pipelines with parallel execution optimized to user hardware, automated DICOM-to-BIDS conversion, processing hundreds of subjects with multimodal scans. Setup and deployment reduced from weeks to ~45 minutes. Adopted by multiple labs. Python • PyQt6 • Docker • BIDS
- Built **Epoche**: ML workbench for neurophysiology—feature extraction, grid search, ensemble optimization, cross-model interpretability, and serialization for deployment to real-time C++ execution backends. 18 classifier types + ensembles, database-driven reproducibility. Supporting active R01-funded clinical research. Python • PyQt6 • scikit-learn • TensorFlow • SQLite • MNE-Python
- Built **real-time closed-loop classification system**: ML classification of brain state for closed-loop experiment control, therapy onset, and neurostimulation. <40ms end-to-end latency. Integrated with fMRI for multimodal closed-loop experiments. Python • C++ • MATLAB • Lab Streaming Layer
- Architected **multi-modal integration solutions** coordinating VR, robotics, biosensors, and neuroimaging with sub-millisecond event marking precision. Lead engineer across 3+ labs for NIH-funded research. Python • C++ • MATLAB • Simulink • LSL

## MUSE Living Worlds: Emergent World Simulation Platform

Building an emergent world simulation platform—100k+ entities with biological needs, psychological states, and social relationships running on deterministic, inspectable systems. Primary application is interactive storytelling; narratives emerge from simulation state rather than scripts, with LLMs translating causality into natural language. **Developing in off-hours while employed full-time.**

### PLATFORM ARCHITECTURE

- **C++20 simulation engine:** GPU kernel batching with structure-of-arrays layouts, deterministic execution with explicit scheduling, 100k+ entity scalability, heterogeneous CPU/GPU execution. [C++20](#) · [CUDA](#) · [Metal](#) · [Vulkan](#)
- **Go backend:** PostgreSQL persistence with row-level security, remote-first with retry queues, WebSocket real-time collaboration, async job queues, multi-tenant Auth/RBAC. [Go](#) · [PostgreSQL \(Supabase\)](#) · [JWT](#) · [OAuth](#)
- **PostgreSQL-to-code pipeline:** Visual schema editor (Constructor) enabling non-programmers to design biological/psychological systems → automated C++ kernel generation → compiled execution. [PostgreSQL](#) · [Python](#) · [C++ codegen](#)
- **Multi-surface architecture:** CYPHER (R&D workbench), CLIO (development management with Go backend and Python subprocess capability for LLM-integrated tooling), GLYPH (e-reader), CALLIOPE (marketplace). Cross-platform desktop (macOS, Linux, Windows). [Electron](#) · [React](#) · [Flutter](#) · [Dart](#) · [Python](#) · [Go](#)
- **Semantic platform interface (BABEL):** NLP intent parsing with spaCy, sentence-transformer embeddings for command similarity search (FAISS), multi-model local LLM orchestration (privacy-first, no cloud dependencies), permission-aware tool-use dispatch across all workbenches. [Python](#) · [spaCy](#) · [FAISS](#) · [Ollama](#)

### KEY SYSTEMS DESIGN

- Unified all entity types under single kernel execution model using exclusion-based system masks
- Eliminated per-frame entity filtering via event-driven kernel membership updates
- Designed VRAM-backed detail tables shared across kernels preventing data duplication
- Implemented deterministic stepping, snapshotting, and replay for debugging and time-travel inspection

## Proprio: Wearable Activity-of-Daily-Living Monitoring Platform

Co-invented Apple Watch IMU-based activity tracking system for stroke rehabilitation. Led technical direction, data infrastructure, clinical validation design, and ML pipeline.

- Built real-time streaming: Apple Watch IMU (30Hz 6-axis) → iOS app with packet queue preventing dropped packets → AWS Lambda processing → MongoDB storage → activity classification with <500ms latency. [Swift](#) · [CoreMotion](#) · [AWS Lambda](#) · [MongoDB](#)
- Designed ML pipeline: time-series feature extraction (sliding windows, FFT) → patient-specific model training → clinical deployment with 85-92% accuracy across 12 activity classes. [Python](#) · [scikit-learn](#) · [pandas](#)
- Built data operations: C# pipeline converting MongoDB to Python data structures, automated quality checks, handling 50GB+ sensor data. [C#](#) · [Python](#)
- Deployed in clinical validation studies with multiple stroke patients demonstrating system feasibility and accuracy
- Secured **\$100K STTR grant** from NIH. Published peer-reviewed validation results.