

MATTHEW (GRAY) MARSHALL

Austin, TX

✉ matthew.gray.marshall@gmail.com  github.com/GrayM8  linkedin.com/in/graymarshall graymarshall.dev

Education

The University of Texas at Austin

Expected May 2028

Bachelor of Science in Computer Science

Austin, TX

- **Coursework:** Operating Systems, Computer Architecture, Data Structures, Cloud Computing, Machine Learning

Experience

Longhorn Sim Racing

Mar 2025 – Present

Co-Founder & Chief Technology Officer

Austin, TX

- Architected and deployed a full-stack TypeScript platform (Next.js, React, Supabase, PostgreSQL) serving 170+ members and 60+ active drivers.
- Built scalable event registration and admin systems supporting 28+ competitive events with role-based access control.
- Integrated custom storefront UI with Shopify checkout and transactional email workflows; deployed via Vercel CI/CD.
- Implemented structured SEO (schema markup, sitemaps, metadata), achieving #1 ranking for organization name and 3.7K page views in first 30 days.
- Collaborated cross-functionally on sponsor relations, business planning, and infrastructure scaling decisions.

Longhorn Racing Electric (Formula SAE EV Team)

Sep 2024 – Present

Telemetry Software Engineer & Planning Lead

Austin, TX

- Develop distributed real-time telemetry pipeline ingesting 60–120Hz vehicle sensor streams via MQTT.
- Built React dashboards used by 20+ engineers trackside for live timing, energy analysis, and driver input visualization.
- Integrated PostgreSQL-backed storage on local server infrastructure for replay, analytics, and performance modeling.
- Designed and implemented 20K+ LOC on-car driver display system optimized for clarity and low-latency feedback.
- Owned web server and client interface architecture for telemetry visualization stack.
- Designed replay/test harness enabling faster-than-real-time playback for validation and performance analysis.
- Led architecture whiteboarding and planning phases for distributed telemetry system.

Projects

AgentWorkspaces — Real-Time AI Collaboration Platform | *TypeScript, React, Cloudflare Workers*

- Engineered a real-time multi-user collaboration platform on Cloudflare's edge using Durable Objects and WebSocket Hibernation to maintain stateful, cost-efficient persistent connections per room.
- Designed a custom WebSocket protocol to synchronize chat, shared memory, todos, artifacts, and settings across concurrent clients with debounced batched writes to durable storage.
- Implemented a structured AI tool-calling engine (Llama 3.3 70B) using a two-pass inference architecture and 9 typed function tools enabling deterministic, autonomous state mutation across sessions.
- Architected a full-stack TypeScript monorepo (React 19, Vite, Tailwind) deployed serverlessly via Cloudflare Pages + Workers, eliminating traditional server infrastructure and minimizing cold-start latency.

Operating Systems Kernel (Pintos) | *C*

- Implemented kernel threading, scheduling, synchronization primitives, and process management.
- Built virtual memory subsystem with paging, address translation, and page fault handling.
- Designed concurrent file system with buffering and persistent storage.
- Debugged race conditions and memory faults in multi-threaded kernel environment.

chArm-v3 CPU Pipeline & Cache Simulator | *C*

- Built cycle-accurate 5-stage 64-bit CPU pipeline simulator with hazard detection, forwarding, and branch resolution.
- Implemented write-back, write-allocate cache with LRU replacement and variable memory latency modeling.
- Validated correctness using trace-driven binaries and reference test suites.

Technical Skills

Languages: TypeScript, Python, C, C++, Java, JavaScript

Frontend: React, Next.js, Vite, WebSockets

Backend & Data: PostgreSQL, Supabase, Prisma, Cloudinary, Durable Objects, SQLAlchemy

Systems: Distributed Systems, Concurrency, Virtual Memory, CPU Architecture, Real-Time Data Processing

Infrastructure: AWS EC2, AWS S3, Vercel, Cloudflare, Git, MQTT, GCP