Matthew Davoren

Software Engineer

Ventura, CA 93004 | matthewdavoren@gmail.com | 630-310-7173
[LinkedIn](#) | [GitHub](#) | [www.matthewdavoren.com](#)

Systems-focused software engineer building backend services, workflow automation platforms, and distributed media pipelines. Experienced in containerized deployment, hybrid cloud infrastructure, and building containerized processing pipelines for large-scale media ingestion and video workflows.

PROFESSIONAL AND TECHNICAL SKILLS

Languages: Python, C++, SQL
Backend & Systems: Node.js, Express, Sidekiq, REST APIs
ML Systems: PyTorch, CUDA, LoRA fine-tuning
Databases: PostgreSQL, MySQL, Microsoft SQL Server
Infrastructure & DevOps: Linux (Arch, Ubuntu, CentOS), Docker, Kubernetes, Nginx, Git
Cloud Platforms: AWS (S3, IAM, Lambda), Google Cloud Platform
Additional Languages: Ruby, Java, TypeScript, JavaScript, Bash

WORK EXPERIENCE

General Services Agency, *Ventura County*
Software Engineer - Python, Ruby   *07/2024 - Present*

- Architected an enterprise workflow platform in Rails 8 featuring configurable multi-step routing, status lifecycle management, and event-driven notifications via a Sidekiq/Redis job processing pipeline to digitize hundreds of paper-based forms.
- Engineered a configurable workflow engine that allowed new form types to be defined and deployed without code changes, supporting multi-page layouts, conditional field visibility, role-based access control, and automatic state changes.
- Modeled a polymorphic status tracking and audit system across 5+ workflow types, normalizing heterogeneous status enums into a unified reporting schema with full change history, reassignment tracking and role-based routing rules.
- Designed a hybrid data architecture integrating a legacy MSSQL employee directory with application-specific models across 30+ entities, supporting real-time database lookups, organizational hierarchy resolution and cross-system reliability.
- Containerized production services using multi-stage Docker builds, including asset precompilation, secrets management, and git-tag-based rollback for zero-downtime production releases.
- Built a Python-based ingestion pipeline to normalize and load CSV billing data into Microsoft SQL Server, orchestrated via Make, reducing processing time by 75%.
- Deployed and operated a Linux-based, open-source media server using FFmpeg to support automated transcoding and distribution of thousands of digital assets.

National Football League (NFL), *Los Angeles*
Media Engineer - Java, SQL   *06/2021 - 05/2024*

- Developed Java-based ingest workflows for the NFL's media asset management (MAM) system, automating the intake of video and image assets from multi-feed acquisition pipelines across live game-day and post-production environments.
- Developed Java tooling to analyze object storage usage patterns in Dell ECS, identifying 23+ storage inefficiencies across media asset archives.

- Designed automated metadata normalization and tagging using JSON parsing, filename conventions, and upstream system lookups to ensure compliance with NFL schema standards.
- Created modular, reusable Java workflow components using Spring and Maven for asset normalization, metadata enrichment, and validation.
- Integrated workflows with external APIs and internal databases to enrich assets with contextual data (e.g., game ID, team names, production tags).
- Built SQL-based Grafana dashboards to support game-day operations, improving observability and response time by 50%.
- Debugged Java logs and JSON metadata to resolve critical system errors across video and image pipelines.
- Designed scalable AWS S3 storage tiers with IAM policies and lifecycle rules to optimize cost and performance across hybrid cloud and on-prem environments.

Spotify, *Los Angeles*
Media Manager   *05/2019 - 06/2021*

- Supported and integrated automated ingest workflows for media files and associated JSON metadata, reducing manual intervention by 90%.
- Resolved data transfer issues with external partners by deciphering Linux CentOS error logs.
- Patched and configured fiber channel connections between NAS storage servers and macOS/Windows editing workstations to enable high-throughput access to shared media volumes.
- Integrated OpenDrives storage with Telestream Vantage Transcoding systems to enable automated watch-folder–based transcoding, delivery, and asset migration.

PROJECTS

Cinezoo
ML Systems & Streaming Infrastructure Project   *2024 – Present*

- Designed backend services in Node.js and TypeScript to manage video channels, scheduling, and asynchronous processing of live stream metadata.
- Built a stream ingestion pipeline that samples frames from external HLS/RTMP streams using FFmpeg without hosting the video infrastructure directly.
- Developed a React-based interface for managing channels, stream playback, and internal tooling for media moderation and discovery.
- Containerized backend services with Docker to support reproducible deployment and scalable processing of media workflows.

SupOtto
Animation Background Generation Pipeline (PyTorch + CUDA)   *2025 – Present*

- Built a GPU-accelerated training pipeline using PyTorch and CUDA to fine-tune Stable Diffusion (v1.5) with LoRA adapters for generating stylized cartoon environments.
- Engineered a custom dataset preprocessing pipeline to normalize prompts, captions, and training images for consistent style generation across multiple scene perspectives.
- Containerized the training environment using Docker with CUDA runtime, enabling reproducible GPU training and inference across development environments.

EDUCATION

Ventura College (ongoing)   *MS Computer Science Prep Courses*
Columbia College Chicago   *Bachelor's Degree*