Maximiliano Merced Rios

max@maxrios.dev | github.com/maxrios | maxrios.dev

SKILLS

Rust ◦ Java ◦ Go ◦ Dart ◦ JS/TS ◦ Kubernetes ◦ Systems Programming ◦ System & Database Design
Concurrent Network Architecture ◦ Incident Management ◦ Load Balancing & Traffic Routing

EXPERIENCE

Google   **December 2021 - Present**
Software Engineer

- Serve as primary domain owner and SME for Originality Reports, Classroom's plagiarism detection service, leading incident triage, reliability operations, and onboarding engineers via tech talks and 1:1s.
- Architected and led the transition of Google Classroom to an RBAC-based system.
  - Designed and implemented a backfill service and ACL validation layer to improve the security of 67 billion files, including an incident management strategy to detect and recover from data corruption during migration.
- Instrumented custom logging and observability pipelines to analyze high-volume tenant traffic, identifying bottlenecks and improving the plagiarism service's SLO by 14%.
- Led the end-to-end integration of Microsoft Word support for the plagiarism detector, authoring the design and resolving dependency bugs.
- Directed a cross-functional team of 1 SWE and 2 researchers to migrate legacy auto-feedback services to LLM-based architectures, improving feedback quality from 60% to 85% across a pilot of 17,000 teachers.

PROJECTS

MCS (Chat Application with Custom Load Balancer)   **January 2026 - Present**
github.com/maxrios/mcs (load balancer)

- Built an asynchronous, real-time chat client and server in Rust using Tokio, featuring a custom binary protocol and TLS encryption.
- Engineered a custom Layer 4 TCP load balancer featuring round-robin routing, dynamic service discovery via Redis, and connection and bandwidth rate limiting.

Home Lab/Infrastructure   **May 2024 - Present**

- Architected a secure, multi-node Linux cluster (Ubuntu, TrueNAS) operating entirely behind a private VPN with zero public ingress.
- Configured internal network routing using a reverse proxy to map subdomains to isolated, self-hosted services, including headless GPU-accelerated LLM deployments.
- Implemented a 3-2-1 backup strategy to ensure fault tolerance and data integrity across machines.

SPEAKING

Intersectionality Between AI & Education   **September 18, 2025**
youtu.be/ZL-IW-9SedE   *Dickinson College*

EDUCATION

Dickinson College   **2021**
B.S. in Computer Science