

Erick Ti

hello@erickti.com

erickti.com | github.com/erick-ti | linkedin.com/in/erick-ti

SUMMARY

Software Engineer owning 15+ production systems at CellLink, a high-velocity startup, with end-to-end responsibility spanning production LLM agents, SQL Server backends, and full-stack web applications used by 500+ employees across multi-site operations. Fluent across Python, C#, and TypeScript; focused on AI-powered tooling, ML retrieval systems, large-scale data infrastructure, and shipping pragmatic systems end-to-end - from architecture through self-hosted deployment and operation.

TECHNICAL SKILLS

Languages: Python, TypeScript/JavaScript, SQL, C#

Backend: FastAPI, Django + DRF, Node.js/Express, ASP.NET MVC

Frontend: React, Next.js, TanStack Query, Tailwind CSS, Vitest, Playwright

Databases: PostgreSQL (pgvector/HNSW), Redis, SQL Server

AI/ML: RAG, hybrid retrieve-rerank, LLM tool-calling & orchestration, MCP, vector search, text/audio embeddings, LLM deployment

DevOps: Docker, Kubernetes (K3s), AWS (S3, Timestream), GitHub Actions, GitLab CI, Caddy, systemd, Cloudflare R2, Grafana, Vercel

EXPERIENCE

Software Engineer

May 2024 – Present

CellLink Corporation - San Carlos, CA

- Built a production LLM agent with custom tool-calling orchestration on FastAPI, the Anthropic SDK, and the M365 Agents SDK, enabling 500+ users to query live SQL Server data in natural language from Microsoft Teams; **cut ad-hoc data requests by 80% at \$115/month**
- Built an internal developer Q&A system with semantic vector search and RAG over 170+ Confluence pages using locally hosted LLMs and the Atlassian REST API, **cutting recurring software-team support requests by 20%**
- Architected and maintained 15+ production ASP.NET MVC applications on SQL Server serving 500+ employees across California and Texas sites with **99.9% uptime**
- Optimized critical SQL Server stored procedures through cursor elimination, set-based CTE and CROSS APPLY rewrites, and composite indexing, **reducing p95 execution time from 11s to 500ms** in the product validation system
- Designed 20+ event-driven ETL pipelines (C# file-watcher services, AWS Timestream, MinIO) **processing 50,000+ manufacturing records daily**, then migrated them to a K3s Kubernetes cluster with multi-node failover and a 24-hour soak (stability) test
- Developed 15+ REST APIs for ERP integration and 30+ Grafana dashboards for real-time telemetry, improving operational workflow speed and **reducing manual data entry by 70%**

Software Engineer / Database Intern

Jan 2022 – Sep 2022

Ojibwe People's Dictionary - Minneapolis, MN

- Refactored the Ruby on Rails backend and optimized MySQL reporting queries on the Ojibwe People's Dictionary, an Indigenous language preservation platform, **driving a 44% increase in organic monthly traffic** (24K to 34.6K users)

Software Engineer Intern

May 2022 – Aug 2022

Ecolab Inc. - Saint Paul, MN

- Deployed 3 computer vision models on Azure Custom Vision + IoT Edge to detect packaging defects on a 4,000+ bottle/day production line, **achieving <1% false positive rate and 100% recall** with sub-second per-bottle inference
- Integrated PLC-driven reject logic with Azure Stream Analytics + Azure SQL and surfaced real-time results through Grafana dashboards

PROJECTS

Doppel - Machine Learning Music Recommendation Engine | doppel.erickti.com

2026

- Built a hybrid retrieve-then-rerank music recommendation engine that finds songs with a similar feel: it pulls candidates from Last.fm/ListenBrainz, matches them against MusicBrainz, then reranks by audio similarity using LAION-CLAP audio embeddings and PostgreSQL + pgvector, with Anthropic's Claude used only to explain the results, not rank them
- Designed a lazy, self-growing vector corpus that embeds only candidates surfaced by real queries, **cutting repeat-run latency ~60x (from ~12 min cold to ~12 s warm)**
- Validated retrieval quality with a 19-seed eval harness across 8 genres, benchmarking CLAP audio reranking against cultural-retrieval baselines and HNSW nearest-neighbor search, and shipped a Next.js console visualizing pipeline telemetry

Millennium - Self-Hosted Trading Card Portfolio Tracker | millennium.erickti.com

2026

- Designed and self-hosted a full-stack Django 5.2 + DRF + Next.js app on a Hetzner Linux VPS behind Caddy TLS, with containerized services, pull-based continuous deployment, nightly systemd ETL jobs, and off-box PostgreSQL backups to Cloudflare R2
- Built a confidence-scored CSV import and reconciliation engine that matches collection exports against a **43,000+ printing catalog**, computing per-lot cost basis, market value, and unrealized gain/loss while preventing cross-edition pricing errors
- Shipped a Next.js product UI with interactive valuation views, type-safe OpenAPI-generated TypeScript clients, **873 automated pytest/Vitest/Playwright tests**, and a live ops dashboard surfacing pipeline uptime, heartbeat, and deployed SHA

EDUCATION

University of Minnesota, Twin Cities

Graduated: May 2023

Bachelor of Science in Computer Science | Minor in Information Technology | Minor in Digital Media Studies