

John C. Waters_

Platform Team Lead. I build the internal tools that make development teams move faster. Yes, that includes agents.

Shelby, MT john@johncwaters.com johncwaters.com github.com/johncwaters linkedin.com/in/johncwaters

\$ **cat** experience.md

Platform Team Lead @ ELM Utility Services

May 2025 – Present

- **Our engineering agent platform** (Python, pydantic-ai, Azure DevOps): 753 commits in six months, its own CI, canary runs and a soak phase before cutover. Where it shines: pulling from Datadog and Azure, figuring out what failed, fixing it, and leaving a PR for developers to review.
- **Self-healing email parser**: 811 locate tickets arrive as email (I know), in dozens of inconsistent formats. Instead of an LLM parsing every message, one LLM call per unknown format writes a static extraction profile. Zero AI calls on the hot path of a regulatory workflow.
- **After-hours dispatch automation** (C#, Azure Durable Functions, Datadog): emergency tickets used to route through a paid external call center. Now the system calls and texts live techs in the field directly, and a silent failure pages within minutes instead of going unnoticed.
- **Org AI adoption**: a Claude Code plugin encoding our PR, story, and release standards as skills that chain into complete workflows (Task > Plan > Feature > Test > Commit > PR). Standards apply at the point of work, not on wiki pages nobody reads. Plus an Azure DevOps extension dashboarding PR review health.

Software Developer @ ELM Utility Services

Dec 2022 – May 2025

- Built and operated C#/.NET services, Python tooling, and TypeScript/React frontends on Azure for utility infrastructure workflows.

Earlier at ELM: Data Analyst, then Tier 1 and Tier 2 tech support (2021 – 2022). Freelance software developer since 2022.

\$ **ls** projects/

glissa/ npm, TypeScript, Node.js

I wanted many Claude Code sessions running in parallel without them stepping on each other. Each gets its own git worktree; a Node server with a localhost dashboard runs the show.

card-harbor/ Electron, React

Desktop app for organizing scanned TCG singles: camera-based card identification and automatic inventory sync. Agents reverse-engineered the APIs that made sync possible.

claude-setup/evals/ Python, PostHog

Which context makes a coding agent best at real PostHog tasks: none, llms.txt, MCP, or a skill bundle? An eval harness that answers it, instrumented with PostHog itself.

keeplings/ Flutter, Android and iOS

A gentle habit app where lapses let your companion rest. No streaks, no punishment. Going live as soon as Google approves it.

\$ **grep** skills

Languages: C#/.NET, Python, TypeScript/React, SQL

AI systems: agent architecture, pydantic-ai, Claude Code and skills, MCP, evals, LLM cost and reliability design, routing across models (Claude, GPT, Grok)

Cloud and ops: Azure (Functions, Durable Functions, DevOps, Bicep), CI/CD, Datadog, PostHog, Vercel