

Finn Rutis they / them

☎ 503-267-5052 | ✉ tech@finnrut.is | 🏠 www.finnrut.is | 🎧 Multipixelone | 🌐 Finn Rutis | 📍 NYC

Content producer who can host the podcast, cut the video, and write the how-to, with a homelab full of LLMs to back it up.

June 26, 2026

Tavily

Remote

Dear Hiring Manager,

I'm applying for the DevRel Manager role at Tavily. What caught me was the framing: this isn't traditional engineering evangelism, it's content and community. That lines up with what I've been doing separately and want to do together: producing video and audio end to end, and building LLM tooling in Python that actually talks to APIs and acts on the results. I can write the tutorial, record the walkthrough, and cut the promo clip, because I've done each of those pieces and I'm looking for the place that asks for all three.

In the summer of 2020 I ran on-site video and audio capture for two productions at Oregon Children's Theatre. I set the camera positions and the mics, built a recording setup I could trust, ran the sessions, then cut the full edits and mixed the audio for digital release. Nobody handed me footage; I made the footage, then made it into something. That's the kind of hands-on production ownership your podcast and video work needs.

A few months earlier I cut a virtual choir for Bridgetown Conservatory. Twenty performers recording alone at home, most with no technical setup at all. I walked them through capturing their parts, troubleshooted their phones and laptops, and managed a flood of dropped files and out-of-spec recordings. Then I synced the video, mixed the audio stems, and delivered on a lockdown deadline. The technical part mattered, but the real work was getting twenty people to show up sounding like themselves on tape. That maps pretty directly to guest management, remote recording, and the editing that comes after.

The audience side comes from my current work too. As a Student Ambassador at Molloy University I give campus tours and run welcome events, and the script lands differently for a nervous parent than for a seventeen-year-old. At Salt & Straw I lead a team through the dinner rush, calling the next thing before it has to happen. Both are the same muscle: reading a room and adapting in real time. Hosting a live podcast or running a community meetup feels like an extension of that, not a pivot away from it.

The technical credibility matters because Tavily builds the search API that powers RAG and agent reasoning. I write Python tools that feed LLMs tight system prompts, enforce strict JSON schemas, and parse the output so the next step can run automatically. I've built agent workflows that chain calls and handle structured responses, which means I've used the kind of APIs you're building for. When a developer wires your search into their RAG pipeline, I know what they're trying to solve because I've built with similar tools. I can write the how-to, produce the demo, and answer the follow-up questions from having done the thing.

I'm in New York and set up to work remotely across US hours. Happy to share any of the projects I mentioned, or to talk through what the first few episodes of a Tavily podcast or a tutorial video series could look like.

Sincerely,

Finn Rutis