

CASE STUDY

Secure Multi-Agent Outreach System

Two AI agents research and draft outreach on local hardware — with a structural wall between them and your data.

— THE WORK

I built a two-agent system where one agent researches a contact and finds genuine, relevant angles, and a second drafts personalized outreach in the sender's own voice. Both run on local hardware, so sensitive data never leaves the building. The research agent runs Qwen 3 8B; the drafting agent runs a local Qwen 3 32B that reproduces a sender's writing style with striking accuracy.

The client started on an off-the-shelf agent framework and chose to rebuild it custom, for tighter security, real accountability, and full telemetry. A tiered firewall structurally blocks sensitive facts from ever reaching the drafting model, the agents hold no database credentials, and every action is logged and replayable. It scales by simply adding more containers or hardware.

— TECHNOLOGIES

Ollama

Qwen 3 (8B / 32B)

Docker

Vellum

Supabase


Retool

Python

— WHAT I TOOK AWAY

- ◆ Real AI security is architectural. Agents that hold no keys cannot leak data, even if one is compromised.
- ◆ Build versus buy is a real decision. We started on an off-the-shelf framework and went custom for security, accountability, and telemetry we could stand behind.
- ◆ Local inference plus containers means data sovereignty and easy scaling — add hardware, add throughput.

THE AGENTS, LIVE



The Agent Operations console: Lennie researching, George drafting, and “0 tier-3 filtered” in the live stream – the firewall at work.

THE SECURITY MODEL



No keys at the edge

The agents hold no database credentials. A breach reaches nothing that matters.



A tiered data firewall

Sensitive facts are structurally blocked from ever reaching the drafting model.



A human approves every send

Nothing goes out without a person's sign-off in the console.