

PORTFOLIO · CASE STUDIES & SELECTED WORK

Building the **future,** one product at a time.

Gabe Lalasava — full-stack developer & founder with 20+ years shipping AI-powered SaaS, real-time platforms, and mobile apps from zero to production.

This portfolio showcases selected projects, architecture highlights, and technical deep-dives across products I've founded, led, or made significant contributions to.

82%

SUPPORT CUT
(SENHUB)

70%

BANDWIDTH CUT
(SEDESS)

20+

YEARS DEV

4x

FOUNDER

CASE STUDY

SendHub.ai

AI-powered omnichannel messaging for businesses. Multi-tenant SaaS delivering a unified inbox, intelligent auto-replies, voice AI, and e-commerce integrations on the edge.

82%

REDUCTION IN
CUSTOMER SUPPORT
LOAD

OUTCOME

AI auto-replies, cascading business prompts, and a semantic knowledge base **cut customer support workload by up to 82%** for pilot businesses — letting small teams handle enterprise-scale inbound volume with sub-second response times on the edge.



SendHub.ai

sendhub.ai · Founder & Lead Engineer

PRODUCTION

A **multi-tenant SaaS** giving businesses AI-powered omnichannel messaging — WhatsApp, SMS, voice, email, and embeddable web chat — in a single unified inbox. Supports WhatsApp Business API for business clients, personal WhatsApp via a **Matrix bridge**, plus SMS and voice over **Telnyx and Twilio**. Runs entirely on **Cloudflare Workers** for low-latency global edge delivery.

Edge-Native Architecture

Entire API on Cloudflare Workers + Hono. Zero cold starts, sub-100ms global latency, queues for message delivery, Workers AI for free embeddings.

Multi-Provider AI

Gemini 2.5 (default), Claude Sonnet, GPT-4.1 — selectable per org. Unified aiQuery() interface with cascading prompts (channel > business > org).

Semantic Knowledge Base

Documents chunked and embedded via Cloudflare Workers AI, stored in **pgvector**. ~1ms queries on IVFFlat indexes. Powers RAG for auto-replies.

Dual WhatsApp Strategy

WA Business API for enterprise clients via Meta OAuth Embedded Signup. **Matrix bridge** (mautrix-whatsapp) for personal WhatsApp accounts — one platform, every WA user.

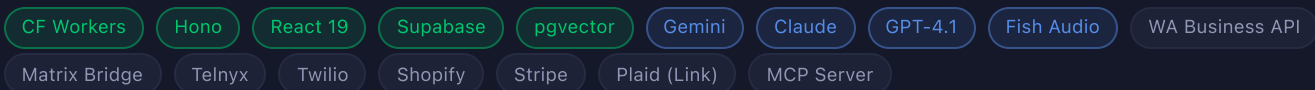
Voice AI + Payments

Fish Audio for low-latency voice replies. **Telnyx + Twilio** for numbers and SMS/voice. **Stripe + Plaid (Link)** for cards and bank-account payment processing in-conversation.

"Client Never Delivers"

Dashboard only inserts messages into Postgres. All delivery handled server-side by the ai-worker — centralized retries, clean separation, no client-side race conditions.

ARCHITECTURE & STACK



CASE STUDY

Seedess

P2P 4K video streaming on BitTorrent. Browser-native WebTorrent delivery turns every viewer into a peer, slashing bandwidth costs without sacrificing quality.

70%

REDUCTION IN
STREAMING BANDWIDTH
COSTS

OUTCOME

By moving video delivery onto a **client-side WebTorrent mesh**, Seedess cuts CDN and origin bandwidth costs by **up to 70%** — every additional viewer becomes a peer, so costs flatten rather than scale linearly with audience size. 4K streams run smoothly even on constrained origins.

S

Seedess

seedess.com · Founder & Lead Engineer

OPEN SOURCE · PRODUCTION

P2P 4K video streaming over BitTorrent, delivered natively in the browser. A **Node.js + Redis** tracker/coordinator serves the swarm, while **client-side WebTorrent** turns each viewer into an active peer. Includes DLNA/UPnP Smart TV casting, a service-worker delivery layer for native `<video>` playback, and a custom React video player tuned for P2P seeking and buffering.

Hybrid BitTorrent + WebTorrent

Native BitTorrent peers seed alongside browser-based WebTorrent peers via WebRTC. Single swarm, two transports — every viewer contributes upload bandwidth back.

Service Worker Delivery

Service worker intercepts `<video>` range requests and serves from the P2P buffer — zero-copy handoff from WebTorrent to native playback with seek support.

4K Optimization

WebTorrent tuned for high-bitrate piece selection, sequential prioritization near the playhead, and adaptive fallback to origin when swarm health drops.

Node.js + Redis Backend

Node.js tracker and swarm coordinator. Redis for peer state, piece availability, and real-time swarm health metrics. Scales horizontally on standard infra.

DLNA / UPnP Casting

Cast P2P streams directly to Smart TVs and media devices via DLNA/UPnP discovery and HTTP bridge — Chromecast-style without Google infra.

Custom React Player

React video player built on native HTML5 media, with P2P-aware buffering indicators, swarm stats, and graceful degradation to HTTP origin.

ARCHITECTURE & STACK

WebTorrent

BitTorrent

WebRTC

Node.js

Redis

Service Workers

DLNA / UPnP

React

HTML5 Video

Open Source

SELECTED WORK

More things I've built.

Additional products I've founded, contributed to, or shipped as lead developer — ranging from collaborative knowledge bases to P2P video streaming and AI extensions.

H

FOUNDER · LEAD

Hardroad

hardroad.app

Cloud SaaS daily action planner with real-time sync via MQTT, offline-first architecture, and a novel relational-integrity framework from Prisma schema all the way into MobX stores on web, iOS, and Android clients.

React · React Native · Swift · MobX · MQTT · Fastify

O

FOUNDER

Offmode.ai

offmode.ai

AI-powered collaborative workspace and knowledge management. Real-time editing with Yjs, vector search on pgvector, role-based access control, and AI chat assistance embedded in the workspace.

Next.js · Yjs · pgvector · OpenAI

N

FOUNDER

NoteOn.space

noteon.space

Live team notes with an AI knowledge base. Rich editing with BlockNote, Excalidraw drawing integration, real-time collaboration via WebRTC + PartyKit, and Stripe billing for team plans.

Next.js · BlockNote · Excalidraw · PartyKit

N

2024

Nexai

[nexai platform](https://nexai.platform)

Conversational AI platform for live support with a developer chatbot API. Achieved the fastest known in-browser semantic search using FlexSearch + semantic cache with RAG over live support docs.

LangChain · Socket.io · TensorFlow.js · Redis

F

2017

Factmata

[Chrome Extension](#)

AI fact-checking Chrome extension that validates statistical accuracy of claims on webpages using NLP. React frontend with Python + HUG API backend. Early AI-meets-browser-extension work.

React · Python · NLP · Chrome Ext

H

2018

Hectare / SellMyLivestock

hectare.com

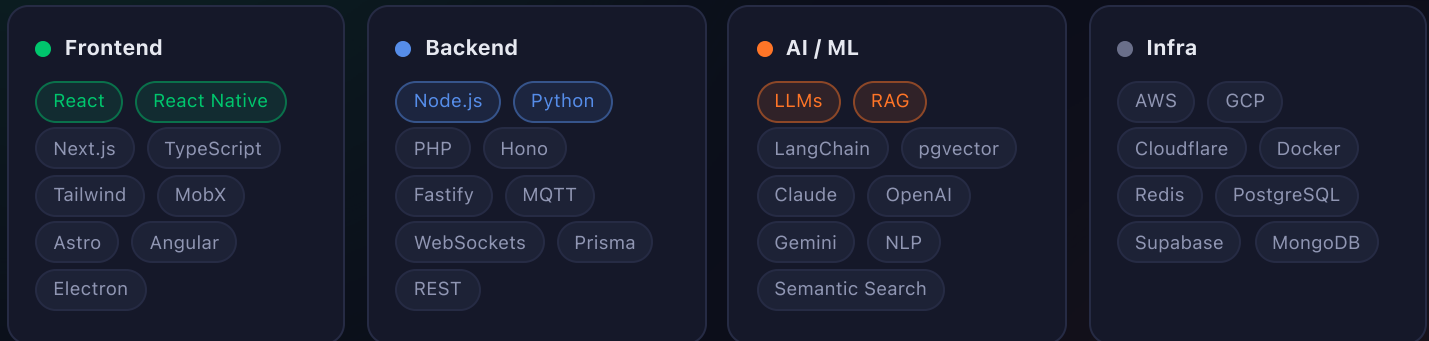
React Native marketplace for Europe's #1 online livestock platform. Algolia-powered semantic search serving ~65,000 farmers on iOS and Android. Cross-platform mobile architecture.

React Native · Algolia · iOS · Android

TECH I WORK WITH

Full-stack, end to end.

Comfortable across the entire stack — from React and React Native on the frontend, to Node.js and Python on the backend, to deploying on AWS, GCP, and Cloudflare.



DEV PROCESS

AI-augmented development.

SendHub.ai and other recent projects were built with extensive use of Claude Code. Not about generating boilerplate — about having an AI pair programmer that understands the full architecture and makes informed decisions across the stack.

CONTEXT-AWARE

Architecture-level understanding

Full codebase context via CLAUDE.md files documenting architecture decisions, conventions, and the multi-tenant model. New features generated with correct RLS patterns, proper error handling, and consistent API design from the start.

SCHEMA + MIGRATIONS

SQL migrations with RLS policies

Denormalized org_id/business_id columns, proper policy functions, and cascade-aware foreign keys — generated following established patterns, catching security gaps before they reach production.

MULTI-SERVICE

Consistency across 6+ services

With Worker, Dashboard, Admin, Shopify App, Matrix Socket, and MCP Server deployed — type alignment across services, shared conventions, and enforcement of rules like "Client Never Delivers".

END-TO-END

Full features in single sessions

Semantic search KB, scheduled messaging, Shopify webhook sync, and email channel integration implemented end-to-end — migration + API + UI — in single sessions, production-ready and consistent with project conventions.

LET'S BUILD SOMETHING TOGETHER

Available for new projects

Looking for a full-stack engineer or founding engineer? I'm open to senior IC, tech lead, and founding roles working on **AI-powered SaaS, real-time platforms, and developer tools**. Reach out at gabe@fijiwebdesign.com.