



Triple Tres (333)

# Ricardo Jaimes Ch.

- > hola@tripl3.dev
- > tripl3.dev
- > github.com/tripl3tr3s
- > linkedin.com/in/tripl3tr3s
- > X: @TripleeeTres
- > Querétaro, México · EST-aligned

## AI / LLM

MCP Protocol    Agentic Loops

LLM Orchestration

MCP Primitives    Multi-Agent

Claude API    Tool Use API

RAG · pgvector    SSE Transport

n8n Flows    Langfuse    Mastra

HITL Patterns    Make · Zapier

## BACKEND / INFRA

TypeScript    Node.js    Python

PostgreSQL    SQLite    Docker

Railway    REST APIs    Git    Linux

## FRONTEND

Next.js    React    Tailwind

TypeScript    shadcn/ui

## EDUCATION & CERTS

Model Context Protocol: Advanced Topics  
Anthropic · Oct 2025

Claude Code in Action  
Anthropic · Sep 2025

Introduction to Model Context Protocol  
Anthropic · Sep 2025

AI Fluency: Framework & Foundations  
Anthropic · 2025

n8n Self-Hosted for Enterprises  
n8n · Nov 2025

Building Scalable Agentic Systems  
DataCamp · 2025

B.A. Graphic Design  
UAA · 2013

## LANGUAGES

Spanish Native

English Professional — technical documentation, proposals, and async collaboration with US-based teams

AI engineer and solo founder focused on replacing manual fiscal/accounting work with **production-grade agent systems**. Shipped a **43-tool MCP server**, multi-tier Claude orchestration, specialized agents and a full fiscal compliance SaaS; all from zero in TypeScript with strict Zod schema validations. Communicates and documents in **English-first environments** as a default.

## EXPERIENCE

### Founder & Lead AI Engineer · **DISAI**

Jun 2025 – Present

- **43-Tool MCP Server:** Architected and deployed a MCP server (TypeScript/Node.js, Railway) covering all three primitives — **Tools** for CFDI 4.0 fiscal ops, **Resources** for live SAT catalog injection (batch-loaded per agent domain), **SSE + HTTP Streamable** dual transport, multi-tenant CSD management. EFOS/EDOS blacklist sync, PAC stamping. **100k+ invoices/min at <2ms latency**. **MCP** **TypeScript** **Railway**
- **3-Tier Agentic Orchestration:** Haiku router (classifies 10 fiscal domains + query intent in ~100ms) → 10 domain **DomainAgents** (Sonnet, Anthropic native tool\_use, real agentic loops up to 6 iterations, parallel tool calls, self-correction on failure) → **Expert Registry** with scoped tool allowlists (5-8/domain) + live MCP Resource injection. Zero hallucinations. **Claude API** **Tool Use** **Multi-Agent**
- **Vector DB & MCP Resource Injection:** pgvector (PostgreSQL 16) for semantic SAT catalog search. MCP Resources pre-loaded into each agent's context window before the first LLM call, eliminating hallucinated catalog codes without RAG query overhead. **pgvector** **PostgreSQL**
- **Self-Hosted AI Infra:** Deployed self-hosted n8n with direct MCP client integration for SAT tool execution. Replaced Zapier/Make — **80% cost reduction**, 400+ integrations, full data ownership. Deployed MVPs to Railway before scaling. **n8n** **Docker**
- **Streaming React Frontend:** Next.js 16 + shadcn/ui chat interface with SSE event pipeline — Activity Log, Thinking Steps, Skeleton bubbles for 30-60s async LLM waits. Compliance dashboard, EFOS monitor, CSD status, driver.js onboarding, per-call savings tracking. **Next.js** **React 19** **Tailwind**
- **Langfuse Observability:** Instrumented the full orchestrator with Langfuse tracing — every LLM call, tool invocation, model selection, token cost, and streaming event tracked end-to-end. Supports cost attribution per session and failure triage without log digging. **Langfuse** **Tracing**
- **Human-in-the-Loop Reviewer Surfaces:** Built compliance dashboard and EFOS monitor as Next.js reviewer interfaces with one-click approval flows, MCP-apps for interactive responses in chat and per-call savings tracking — wired directly into the agent pipeline as the human approval layer. **Next.js** **HITL** **MCP-UI**

### Core Contributor / Lead Developer · **RetailDAO**

Jun 2024 – Jan 2026

- **Financial Analytics Terminal:** Full-stack "Bloomberg-style" dashboard aggregating real-time data from Federal Reserve, on-chain APIs and crypto exchanges via WebSockets. Rate-limiting, memory caching, and graceful degradation for a self-healing low-latency system. **React** **WebSockets** **APIs**
- **Smart Contract Audits & Documentation:** Led internal audits and built a token distribution visualization platform. Mitigated investor risk and drove a **30% increase in token valuation** to all-time highs.

### Crypto Research Analyst · **Coinsider**

Mar 2024 – Sep 2025

- **Data-Driven Intelligence:** Delivered 25+ peer-reviewed analytics reports (1,200+ pages) — SWOT analysis, investor risk models, and macroeconomic forecasting for Web3 ecosystems. Translated dense technical data into content generating **+70k views**.

# Ricardo Jaimes Ch.

## KEY PROJECTS

MCP SDK · TypeScript  
Claude API · Tool Use  
Next.js 16 · React 19  
Langfuse · HITL  
n8n · Docker  
PostgreSQL · pgvector  
Railway · Coolify

## KEY PROJECTS

### SAT-MCP Server

43-tool MCP server for Mexican fiscal compliance. Full primitive set: tools, resources, prompts, elicitations, tasks + 8 MCP-UI mini-apps for in-chat fiscal rendering. SSE + HTTP Streamable. CFDI 4.0 lifecycle, EFOS/EDOS sync, PAC stamping, multi-tenant CSD. 3,960 tests · 91%+ coverage.

TypeScript MCP SDK Node.js  
Railway SQLite

### DISAI\_Conta – 3-Tier Agent System

Haiku router + 10 domain DomainAgents (Sonnet, native tool\_use, 6-iter agentic loops) + Expert Registry with scoped allowlists and live MCP Resource injection. Langfuse full-stack tracing. HITL compliance dashboard. Streaming SSE UI for Mexican SMEs.

Next.js 16 Claude API Tool Use  
Langfuse shadcn/ui

### n8n AI Automation Stack

Self-hosted n8n with direct MCP client integration for SAT tool execution. Specialized sub-agents per CFDI type, 400+ integrations, context bundle architecture. Replaced Zapier/Make — 80% cost reduction.

n8n MCP LLM Agents Docker

### Homelab AI Infrastructure

PostgreSQL 16 + pgvector for semantic SAT catalog search. Self-hosted Gitea, static IPs, Samba share. Powers the MCP Resource injection pipeline in production.

PostgreSQL pgvector Linux Gitea  
Docker

### Mastra + Langfuse Stack Evaluation

Evaluated Mastra framework for HITL patterns, agent workflows, and MCP connections. Identified tradeoffs vs. custom harness: Soniox TTS incompatibility, shallower Langfuse A/B testing depth. Production stack uses a custom orchestrator with pure Langfuse cloud for full observability and cost-attribution control.

Mastra Langfuse HITL Observability MCP