

Martin Gonzalez

AI Engineer · Remote · martin.oscar.gonzalez1@gmail.com · +54 291 4167300 · LinkedIn

SUMMARY

AI engineer who builds production software with AI coding agents, not just around them. I use Claude Code with MCP servers wired into the systems I operate to do live reviews and deploy production automations for a consultancy serving 15+ US clients, from LLM document pipelines to AI voice intake agents. I design the prompts, context strategies and agent skills that make those workflows reliable, and I review every line the agent produces before it ships. That review habit is why I keep finding the bugs a clean test case hides.

SKILLS

AI-assisted engineering: Claude Code daily (configuration, custom skills, project-level instruction files, MCP servers), agentic development workflows, context strategies for multi-system codebases, prompt and system-instruction design, versioning and iteration on agent behavior, critical review of AI-generated code for correctness and security

LLM engineering: Azure OpenAI, OpenAI API, Anthropic Claude API, structured outputs, tool use and function calling, LLM document pipelines (extraction, translation, reasoning), prompt engineering for production, AI voice agents (Retell AI), conversational design

Backend & data: Supabase (PostgreSQL) schema provisioning and queries, REST APIs, webhooks, OAuth2, background jobs and scheduled workflows, JavaScript/Node.js, Python, SQL

Deployment: Vercel (including programmatic project creation via API), Next.js, Docker and Docker Compose, self-hosted infrastructure, Git

Orchestration & integration: n8n, Make, Zapier, event-driven architecture, telephony (Telnyx, RingCentral, CallRail), CRM (GoHighLevel, Docketwise), third-party and government APIs

Infrastructure: Linux, Windows Server, virtualization, backup and disaster recovery, network security, credentials and secrets handling

EXPERIENCE

AI Automation Engineer at Immigration Finder (contract)

Jan 2026 - Present

US consultancy building AI systems for 15+ immigration law firms. I own each build end to end, from scoping to production.

Building with AI coding agents

- Use Claude Code with MCP servers connected to the systems I operate to run live reviews and ship production deployments, including the Visa Bulletin automation and the AI voice intake agents described below.
- Design and version the prompts, system instructions and agent skills that drive those workflows, plus the context strategies that let an agent work across an automation platform, a CRM, telephony and third-party APIs at once.
- Review every AI-generated change before it ships. I don't trust agent output until it has run against real data, which is how I found four bugs in one pipeline that any design review would have passed, including a row-contamination bug where every row after the first inherited the first row's identity.

Production AI systems

- Build LLM pipelines (Azure OpenAI, OpenAI) for translation, structured extraction and reasoning over legal documents. Rebuilt a monitoring pipeline the day its scraper got blocked: pull the source PDF directly, structure 150 rows per issue with GPT-5-mini into structured output, deduplicate monthly, alert. Added an explicit check separating "not published yet" from "the fetch is broken", two failures that return the same empty result.
- Design and operate about 30 AI voice agents (Retell AI) across 15+ client deployments, wired into client telephony and CRM through webhooks, with a reusable implementation guide that standardizes every rollout.
- Built a 27-node case-status engine polling a government API daily and notifying firms and their clients on state changes, verified end to end against a real filing before release.

Shipping to production

- Shipped an internal usage dashboard from empty schema to live in one working session: Supabase data model, an n8n reporting workflow, and a Next.js app deployed on Vercel with per-client access. Fixed a cross-client authentication flaw found during the build, where a shared secret let one client's password unlock another client's dashboard.
- Built a government case-status tool as a Next.js app, with the Vercel project created 100% programmatically via API. Verified without mocks: a real success, a real validation rejection and a real 404, both locally and on the live deployment.
- Audited an existing live automation end to end and found three real production bugs, including one sending recurring notifications regardless of case status, and a fan-out bug where chained SDK calls produced a straight line instead of parallel branches, silently dropping the logging branch.
- Document every build so it can be reproduced and supported by others, and set team conventions for how automations link back to the workflows that created them.

System Administrator at Coopeviajes.com

May 2023 - Jan 2026

Travel agency. Where I started automating, with scripts before platforms.

- Rebuilt email and marketing infrastructure, raising domain sender reputation from 3.7 to 10 with 100% deliverability.
- Implemented Veeam backup infrastructure with recovery verified during real incidents.
- Automated recurring IT operations with Bash and PowerShell.

IT Support at Hospital Interzonal Dr. José Penna

Sep 2019 - Nov 2019

- Implemented a ticketing system that improved incident tracking and resolution times.

IT Assistant at Escuela de Enseñanza Técnica n°2

May 2019 - Sep 2019

- Deployed a Moodle server on Ubuntu Server and redesigned the institution's WiFi.

EDUCATION

Computer Engineering (coursework), Universidad Nacional del Sur

2019 - 2023

Technical Secondary Diploma in IT and Networks, Escuela de Educación Secundaria Técnica n°2

2012 - 2019

LANGUAGES

Spanish (native) · English (professional working proficiency; daily client-facing work with US clients) · Portuguese (elementary)