

MOSES OYEBAMIJI

Career Portfolio | Growth Systems, AI Automation and Commercial Architecture

+234 913 497 7608 | Mosesoyebamiji95@outlook.com | linkedin.com/in/moses-coweb | Lekki, Lagos, Nigeria

POSITIONING

I build growth systems, not campaigns. Every project in this portfolio is a production system I designed, wrote and operate: AI agents running against live APIs, conversion tracking that reconciles to banked revenue, and machine learning models scored against real outcomes rather than demonstrations. The distinction matters. Most marketing leaders adopt AI tools. I build the architecture underneath them, and I carry the commercial number that architecture produces.

Across eleven years these systems have produced over 50,000 leads and more than 12,000 sales for my own ventures and for clients. This document sets out six of them: the problem each solved, how it was built, and what it measurably returned.

\$500,000

Client revenue generated

\$160,000+

Freelance client revenue

10,000+

Marketplace sales

8.6x

Return on ad spend

5.0

Rating, 61 reviews

14

AI agents shipped

SELECTED SYSTEMS

PROJECT 01

LinkedIn Demand Engine for a Professional Services Business

Client. Lhurd Resume, founded by Lynda Hurd. A professional services business operating almost entirely on LinkedIn, with no owned web presence and no capture layer beneath its audience.

Problem. The business had genuine authority and an engaged LinkedIn following, but every lead lived inside someone else's platform. There was no landing surface, no way to capture intent, and no measurable path from attention to revenue. Growth depended entirely on the founder posting and responding personally, which is a ceiling rather than a strategy.

What I built. I recommended moving the business onto owned infrastructure and then built that infrastructure personally: the website, the landing pages, and the funnel connecting paid distribution to conversion. The architecture routes LinkedIn advertising into dedicated landing pages, captures intent there, and returns qualified prospects into LinkedIn messaging where the founder's authority converts them. Paid reach and personal credibility were made to compound rather than compete.

ARCHITECTURE

- LinkedIn advertising driving traffic to purpose built landing pages rather than to a profile or feed
- Capture and qualification layer on owned web infrastructure, designed and built end to end
- Structured handoff from landing page into LinkedIn direct messaging for conversion
- Full funnel instrumentation so lead source, volume and closed revenue could be attributed

OUTCOME

- **Approximately 1,000 qualified leads generated in roughly three weeks** from launch
- **Over 300 sales closed** from that cohort, a conversion rate of roughly 30% from lead to paying customer
- **\$500,000 in total revenue** generated for the client through the system
- A business dependent on organic founder activity converted into one with a measurable, repeatable acquisition channel

PROJECT 02

Marketplace Ranking Strategy: Engineering Review Velocity

Problem. Selling on a third party marketplace where search placement decides everything. New listings face a cold start trap: products without reviews never surface in search, so they never sell, so they never accumulate reviews. Paid placement was the obvious escape and an expensive one.

What I built. I reverse engineered the marketplace's ranking behaviour and identified that review count carried disproportionate weight in search placement. Rather than buy visibility, I engineered it. I generated demand offline, then deliberately routed those buyers through the marketplace listing instead of selling to them directly, forgoing the immediate direct margin to capture a ranked transaction. Each buyer's contact was captured at the point of sale and followed up specifically to convert the purchase into a review.

MECHANISM

- Identified review count as the dominant ranking signal in marketplace search
- Converted offline demand into on platform transactions to seed ranking, accepting lower per unit margin to buy algorithmic position
- Systematic post purchase follow up by phone to drive review conversion rate
- Compounding loop: reviews lifted ranking, ranking produced organic search volume, organic volume produced further reviews

OUTCOME

- **Over 10,000 sales generated** through the marketplace storefront
- Products ranked organically in category search, replacing paid placement with earned position
- Sourcing and importation of jewellery, bags and consumer electronics managed end to end alongside the listing operation

PROJECT 03

AI Assisted Paid Acquisition and Conversion Tracking System

Problem. A live ecommerce operation was running paid social with unreliable measurement. Purchase events fired without revenue values attached, so return on ad spend read as zero, conversions were under reported, and no spending decision could be justified with evidence. Budget was being allocated on instinct.

What I built. A complete acquisition system covering campaign architecture, creative and copy strategy, and the measurement layer beneath both. Server side conversion tracking was implemented through the Conversions API alongside browser pixels, with deduplication so events counted once and revenue values passed correctly. Campaign structure was rebuilt so the ad platform could exit its learning phase rather than resetting on every edit.

ARCHITECTURE

- Browser pixel and server side Conversions API running in parallel with event deduplication
- Purchase events carrying real transaction value and currency, reconciled against banked revenue
- Campaign structure designed around a single consolidated ad set to preserve learning signal
- Automated campaign creation and audit scripts written against the platform's marketing API

OUTCOME

- **8.6x return on ad spend** and **11% click through rate** on a ■25,000 product
- **■2,895 cost per purchase**, measured against confirmed revenue rather than platform estimates
- The methodology packaged as a versioned, reusable system deployable to a new store without rebuilding it

PROJECT 04

Multi Agent AI Operating System for Marketing and Support

Problem. Content production, campaign operations, customer support and social engagement each required human attention that did not scale. Output quality varied by who did the work. Hiring was the obvious answer and the wrong one.

What I built. A library of fourteen specialised AI agents and skills, each owning a defined function, invoked on demand and running against live production systems. These are not prompt templates. Each carries its own operating rules, tool permissions, guardrails and domain knowledge, and several write directly to production infrastructure.

ARCHITECTURE

- Two autonomous agents with scoped tool access, owning paid marketing and customer support respectively
- Twelve invocable skills covering ad creation and auditing, brand voice, and a ten skill suite for social content, engagement analytics and reply handling
- Guardrails encoded as rules, including ownership boundaries preventing agents writing to assets they do not own
- Integration with live REST APIs across advertising platforms, publishing systems and internal reporting

OUTCOME

- Content, campaign and support functions run at a throughput a small team could not otherwise sustain
- Quality became a property of the system rather than of the individual performing the task
- Each capability is portable and deployable into another business without rebuilding

PROJECT 05

Machine Learning Order Risk Model with Shadow Deployment

Problem. In a pay on delivery market, failed deliveries destroy margin. Goods travel, the customer is unreachable or declines, and the cost is absorbed entirely by the business. The operation needed to know which orders were likely to fail before dispatch.

What I built. A risk scoring model running a local large language model to assess each order across address findability, customer responsiveness, order characteristics and confirmation signals, returning a risk band and a recommended action. It was deployed in shadow mode, scoring live orders and logging predictions without acting on them, so accuracy could be measured against real outcomes before any automated decision was permitted.

ARCHITECTURE

- Llama 3.2 served locally, scoring structured order data against weighted risk signals
- Three band output with a specific operational action attached to each band
- Shadow deployment writing timestamped predictions to an append only log, with no writes to the live order system
- Back test methodology excluding in flight orders from the success rate, because counting them inflates it

OUTCOME

- **Just under 2x separation** between risk bands: 28% failure among flagged orders against 15% among cleared
- **64% of failures caught**, representing ■175,000 of protected revenue across the back test
- First shadow run scored 19 in flight orders worth ■525,000 and flagged 6 worth ■150,000 for confirmation before dispatch

A note on rigour. The project documentation states plainly that weights were fitted on the same orders they were tested against, that eleven failures is too small a sample to draw conclusions from, and that an earlier version of the analysis overstated the delivery rate by counting shipped orders as successes. That correction is recorded in the findings rather than quietly removed. I hold my own analysis to the standard I would expect to defend in front of a board.

PROJECT 06

End to End Commerce Platform: Checkout to Settlement to Profit

Problem. Revenue reporting stopped at the storefront. There was no reliable view of what an order actually earned once payment fees, landed product cost, delivery and acquisition spend were accounted for. Cash received was being mistaken for profit earned.

What I built. A complete commercial system spanning the customer facing store, payment capture, order lifecycle management, inventory, and a reporting layer resolving each order to a true profit figure. Built in Python against live transaction data, with an operations dashboard used daily to run the business.

ARCHITECTURE

- Payment gateway integration handling card, bank transfer and mobile payment flows
- Order lifecycle model where every status is explicitly bucketed, so no revenue silently disappears between states
- Per order and per product profit ledger incorporating landed unit cost, payment fees and daily allocated acquisition cost
- State aware inventory with automatic deduction, and an administrative dashboard deployed to a cloud instance

OUTCOME

- **■1,979,000 in gross transaction value** tracked through the system to date
- Cash position and earned profit reported as separate figures, correcting a material misreading of business performance
- A reporting discipline tying marketing activity directly to banked revenue rather than platform attribution

CORE CAPABILITIES DEMONSTRATED

Growth and acquisition: paid social strategy, funnel architecture, marketplace and search ranking strategy, conversion rate optimisation, attribution design, cost per acquisition management, creative and copy strategy

AI and automation: multi agent system design, agent orchestration, prompt engineering, local and hosted LLM deployment, retrieval augmented generation, content and campaign pipeline automation

Technical: Python, REST API integration, payment gateway integration, server side event tracking, data modelling and reporting, dashboard development, cloud deployment

Commercial: full profit and loss ownership, unit economics, pricing and margin strategy, supplier negotiation and international sourcing, inventory and fulfilment operations

INDEPENDENT VERIFICATION

Eleven years of client delivery, publicly documented and independently rated. Across the two largest freelance marketplaces I have generated **more than \$160,000 in client revenue**: over \$100,000 on Upwork and over \$60,000 on Fiverr. That work spans web development, brand development, store design, WordPress builds and marketing delivery for clients across multiple continents.

My Fiverr profile carries a **5.0 rating across 61 client reviews** accumulated over more than five years. Sustaining a perfect rating across that volume and that timespan is itself the credential: it reflects delivery, not marketing.

Verifiable at: [fiverr.com/emmababy](https://www.fiverr.com/emmababy)

CERTIFICATIONS

Certificate of Completion, Claude Code 101 | Anthropic, September 2026 | Credential ID epvqxpp8awim

AI Engineer Core Track: LLM Engineering, RAG, QLoRA, Agents | Udemy, August 2026 | Credential ID UC-0fa02254-fc52-4814-a130-3af78b58c22f

Certificate of Completion, Claude 101 | Anthropic, June 2026 | Credential ID vipjqgpvd6vd

Google Digital Skills for Africa | Google

WHAT THIS EVIDENCES

A growth leader who can architect the system, write the code that proves it works, and carry the commercial number it produces. The measurement discipline is the through line: every claim here traces to a figure reconciled against banked revenue or scored against a real outcome. Where a result is not yet proven, the documentation says so. That is the standard I would bring to a board reporting cadence.