

Mujeeb Ahmad

AI SYSTEMS ENGINEER

AUTONOMOUS AGENT ARCHITECT

CRM & AUTOMATION SYSTEMS BUILDER

+92 302 584 3504

m.ahmad.aidigital@gmail.com

Sahiwal, Punjab, Pakistan

linkedin.com/in/mujeebahmad-246893378

PROFILE

Production-grade AI engineer specializing in **zero-leakage, cloud-native architectures** that bridge no-code platforms (n8n, Make.com, Zapier) with custom Python backends and **multi-agent AI pipelines (LangChain · LangGraph)**. Proven track record designing enterprise RAG systems and Kubernetes (AWS EKS) deployments — not a prototype builder.

PORTFOLIO & LINKS

Enterprise AI Triage (GitHub)

AEC RAG Pipeline (GitHub)

SEC-RAG System (GitHub)

Full Project Portfolio (Drive)

AUTOMATION & CRM

Zapier Webhooks

n8n

Make.com

HubSpot API

GoHighLevel

ATS Integration

LTU Architecture

Idempotent Events

Code by Zapier

AI & LLMS

LangChain

LangGraph

RAG Architectures

Pinecone

Vector Databases

Semantic Search

Vector Embeddings

Agentic Workflows

Claude Sonnet 4.6

Gemini 2.5 Pro

Vertex AI Imagen 4 Ultra

Multi-Agent Design

Prompt Engineering

Self-Healing JSON

PROJECTS & DELIVERED WORK

SEC EDGAR Financial Intelligence RAG System

Live Demo

GitHub

ENTERPRISE RAG

LangChain · LangGraph · Pinecone Vector DB · FastAPI · Docker · Kubernetes (AWS EKS) · CI/CD (GitHub Actions) · React/Vite · Render · Vercel

- Orchestrated a 5-Node LangGraph State Machine:** Designed a stateful multi-agent pipeline (Router → Retrieval → Analyst → Citation → Guardrail) autonomously executing financial analysis, multi-stock comparisons & risk assessments across 50 S&P 500 companies from real SEC 10-K filings.
- Built Enterprise-Grade RAG Pipeline:** Integrated LangChain + Pinecone Vector DB with semantic similarity search, indexing 4,000-char chunks with section tags from SEC EDGAR government filings — zero hallucination architecture.
- Deployed Cloud-Native on Kubernetes (AWS EKS):** Containerized with Docker, configured Horizontal Pod Autoscaling (HPA) to scale 2→10 pods automatically during high-concurrency multi-ticker queries.
- Automated CI/CD Pipeline:** Built GitHub Actions workflows for automated testing, Docker image builds, ECR pushes, and zero-downtime rolling updates to the production EKS cluster.
- Engineered "Zero-Leakage" BYOK Architecture:** Client-side 'Bring Your Own Key' — LLM API keys processed in-memory, never persisted server-side. Enterprise security & compliance ready.
- Resilient Frontend:** Exponential backoff auto-retry (3-attempt, 5s/10s/15s) handles 120s+ async AI reasoning without user-facing connection drops.

AEC Document Intelligence Pipeline

GitHub

ENTERPRISE RAG

Python 3.11 · AWS Lambda · AWS SQS (Fan-Out) · AWS EventBridge · Pinecone · Supabase PostgreSQL · Gemini 2.5 Pro (Vision + Text) · Vertex AI · Instructor · PyMuPDF · PDFPlumber

- Engineered a production-grade 3-tier multimodal RAG pipeline** for the Architecture, Engineering & Construction industry — autonomously extracts structured material specification data from raw commercial construction PDFs with **98.8% accuracy across 84 tags in 75.86 seconds** at under \$0.12 per document.
- Designed event-driven serverless fan-out architecture:** S3 upload → EventBridge → SQS ingestion queue → Ingestion Controller Lambda → parallel Loader + Router invocation → atomic Supabase RPC fan-out to 5 concurrent Standard Worker Lambdas → Vision fallback queue — zero persistent compute, fully idempotent.
- Implemented atomic distributed orchestration:** PostgreSQL try_trigger_orchestrator RPC guarantees exactly-one fan-out even when 100 loader chunks complete simultaneously — prevents thundering-herd SQS storms without any coordinator service.
- Built OCR + Unicode resilience layer:** 10-variant Unicode dash normalization (U+2010–U+FF0D) + OCR confusable substitution (0↔O, 1↔I/L) enables semantic vector matching even on character-corrupted PDF text layers. Zero-hallucination guard explicitly rejects cross-tag data attribution — TER-1 ≠ TERT-1.
- Delivered measurable AEC industry impact:** Replaces 4–8 hours of senior estimator manual work (\$225 avg cost) with a sub-76-second automated extraction — validated

AI Voice Payloads

Instructor + Pydantic

text-embedding-004

Multimodal Vision API

CSI MasterFormat

BACKEND & CLOUD

AWS (EKS/EC2)

AWS Lambda

AWS SQS + DLQ

AWS EventBridge

Fan-Out Architecture

Kubernetes (K8s)

Docker

FastAPI

CI/CD GitHub Actions

Auto-Scaling (HPA)

Redis Cache

Python asyncio

Playwright

Modal.com

Supabase

PostgreSQL

Row-Level Security

JWT / HMAC

PyMuPDF (fitz)

PDFPlumber

E.164 Normalization

Deno/TypeScript

FRONTEND

React 18

TypeScript

Vite

Tailwind CSS

shadcn/ui

Recharts

Framer Motion

Chrome Extension API

APIS & INTEGRATIONS

Slack API

Resend API

Google Calendar

TwelveData

DataForSEO

MT5 / MetaTrader

Telegram Bot API

SEC EDGAR API

EDUCATION & CERTS

AI for Business Strategy

Columbia University · 2024

Social Media Marketing

HubSpot Academy · Feb 2026

SEO Specialization (5 Courses)

UC Davis via Coursera · Feb 2026

Foundations of Business Strategy

Univ. of Virginia, Darden · Feb 2026

Fundamentals of Digital Marketing

Google Digital Garage · Aug 2025

Executive Diploma of CEO

MTF Institute · Aug 2025

Intermediate in Computer Science

Govt. College, Sahiwal · 2021

against Cardozo High School renovation specification (GCS-SIGAL LLC / HartmanCox Architects, Washington DC).

● Autonomous Recruitment AI Pipeline — ATS × CRM × AI Voice Screening

NEW · NDA

Zapier · Code by Zapier (Python) · HubSpot API · Webhooks · AI Voice Agent · US Staffing & Logistics Firm

- Architected a fully autonomous end-to-end recruiting pipeline — ATS webhook intake → Python JSON sanitization → HubSpot CRM sync → AI voice interview → Pass/Fail routing → auto-scheduled calendar booking. **Zero human intervention** from application to booked interview.
- Engineered **Python sanitization engine** (Code by Zapier) that flattens nested JSON arrays, enforces strict **E.164 phone formatting**, normalizes email casing — garbage data never enters the CRM.
- Implemented **LTU (Lookup-Then-Update) architecture** using unique ATS candidate IDs — zero duplicate contacts, non-destructive field updates, full audit trail at every node.
- Built conditional decision engine (Zapier Paths): PASS → pipeline stage updated, SMS + email with calendar link fires instantly; FAIL → tagged, pipeline closed, recruiter notified. Live end-to-end tested — **zero errors across all paths**.

● LeadFlow AI — Intent-Based B2B Prospecting & Lead Scoring Platform

CLIENT

React · TypeScript · Tailwind CSS · Gemini 3.1 Pro · Google Search Grounding · Leaflet Maps

- AI-powered B2B prospecting platform scanning open web & forums for **high-intent prospects** — shifts from cold-calling to hot-intent targeting via Gemini 3.1 Pro + Google Search Tool.
- Lead Scoring algorithm (0–100) grades businesses on custom criteria; built **self-healing JSON parser** handling 8,192-token AI payloads with zero application crash on overflow.
- "Executive AI Brief" auto-generates a 3-sentence customized pitch per lead on demand + mini-CRM dashboard with localStorage persistence.

● Enterprise CRM & Revenue Pipeline Automation

CLIENT

n8n · HubSpot API · Python · Slack API · Resend API · Snitchy Payment Gateway

- "Read-Calculate-Write" pipeline** on n8n — live webhook ingestion → HubSpot CRM lookup → real-time financial aggregation → non-destructive write-back. **100% calculation accuracy, zero leakage**.
- Strict parseFloat()/Number() data integrity layer inside n8n JS nodes eliminates string-concatenation bugs across all donation records.
- Conditional **VIP routing engine** (\$500+ threshold) triggers simultaneous Slack alert + branded email — idempotent webhook handling prevents all duplicate processing.

● BidVeritas — AI Procurement Intelligence SaaS

OWN SAAS

Python · Modal.com · Playwright · Gemini 2.5 Pro · Supabase · ReportLab

- Production B2B SaaS monitoring government procurement portals via stealth browser automation — **25 parallel workers**, anti-captcha bypass, serverless containers on Modal.com.
- 4-phase pipeline: portal scraping → 2-stage AI triage → ReportLab PDF → Resend email dispatch — fully cron-scheduled Mon–Fri. **SHA-256 amendment detection** with race-condition-proof row-level locking.

● ForexFundAI — Institutional-Grade AI Trading Platform

CLIENT

Supabase · Deno/TypeScript · Gemini 2.5 Pro · MT5 · TwelveData · forexfundai.com

- › Complete trading platform: Gemini 2.5 Pro macro bias engine, TwelveData real-time data, MT5 AutoTrade queue with EA polling, retry/backoff logic, JWT auth, row-level security.
- › Automated reporting pipeline to **Telegram (12,000+ subscribers)**, admin RBAC dashboard, branded transactional email — production-hardened 24/7.

● Autonomous 9-Agent AI Decision Pipeline

CLIENT

Python · Claude Sonnet 4.6 · Gemini 2.5 Pro · Supabase · GitHub Actions (hourly cron)

- › ARIA: **3-persona internal AI debate** (Aggressive / Conservative / Neutral) deliberates before any output — "WAIT is valid" prevents forced low-confidence decisions.
- › **Self-learning memory**: Claude Haiku auto-generates banned patterns from last 50 outcomes; LOCKED vs SOFT rule classification — system improves every single run.

● Propace Autofill — AI Document-to-Form Chrome Extension

OWN SAAS

Chrome Extension API · OCR · AI Extraction · 256-bit Encryption · propaceautofill.live

- › **98% OCR accuracy** on passports, visas, IDs — autofills any web form in under 5 seconds. 50+ languages, GDPR compliant, SOC 2 aligned.
- › Deployed to **10,000+ professionals** saving 10+ hours weekly per user. Shipped from zero to production as a complete self-contained SaaS.

● Full-Stack Business Management Platform

CLIENT

React 18 · TypeScript · Supabase · PostgreSQL · ahmadgraphics.live

- › Complete **double-entry accounting system** from scratch — Chart of Accounts, P&L, Balance Sheet, Cash Flow, Bank Reconciliation — automated via PostgreSQL triggers across 20+ RLS-protected tables.
- › Granular RBAC with JSONB permissions, HMAC-signed AI chat tokens, security-definer SQL functions, JSON-LD structured data, 6 Supabase Storage buckets.