

JAEHWAN HA

GenAI Solutions Engineer • Technical BD • Pre-Sales

Seoul, South Korea • +82 10-8453-2114 • gkwoghks96@gmail.com • linkedin.com/in/hajaehwan • github.com/hwan96-ai • hajaehwan.com

SUMMARY

GenAI Solutions Engineer / Technical BD who sells by building the proof. Turns customer discovery into working evidence — self-built RAG/LLM demos and end-to-end PoCs — and carries it through technical Q&A, proposals, quotes, packaging, and pricing into deals; currently owns GTM and B2G strategy for the Sapie-Guardian GenAI security gateway. 3+ years in IT, 2+ focused on GenAI pre-sales, PoC-to-productization delivery, and product PM. Evidence across the full arc — validation: solo end-to-end enterprise callbot PoC, first sales contact to pricing; a won deal: smart-home voice AI, through technical validation and scope control; productization: PM of the PoC that became Sapie-Reco (AWS Marketplace global listing); selection: ~KRW 3.1B (=US\$2.2M) national semiconductor AI-agent program, ~20:1 written and 3:1 final rounds.

CORE STRENGTHS

- **Fast technical validation that de-risks decisions** — design and build RAG/LLM demos and PoCs (ingestion → hybrid retrieval → generation → API/UI → cloud) in 1–2 week cycles, enabling customers to evaluate working systems, not slides
- **Deal design through proposals, packaging & pricing** — translate discovery into PoC scope, proposal storylines, quotes, tiered packaging, and pricing that balance customer budget against delivery and cost risk
- **Customer-facing range** — solution presentations, demos, and technical Q&A from inbound first response to public-sector procurement requirements (GS certification, KISA evaluation tracks, national-program alignment)
- **Judgment on where AI belongs** — repeatedly scoped AI to where it adds value and kept rule-/standards-based components where accuracy is non-negotiable, protecting quality, cost, and delivery timelines

PROFESSIONAL EXPERIENCE

GenAI Solutions Engineer / Technical BD • Saltware (AWS Partner), Seoul *Jan 2026 – Present*

- Own GTM execution for **Sapie-Guardian**, a GenAI security gateway (AI-DLP control layer over existing DLP/proxy infrastructure): customer meetings, demos, technical Q&A, and inbound first response
- Rebuilt sales assets around buyer concerns: separate private-sector and public-sector (B2G) sales kits, a reusable Q&A base from recurring customer questions, and tiered packaging with a lower entry cost
- Reviewed and proposed B2G entry requirements (GS certification, N2SF readiness, KISA fast-track / performance evaluation) and supported certification documentation; run initial technical review of AI SI inbound inquiries, coordinating feature requirements with engineering
- Owned a major enterprise customer's AICC chatbot→callbot PoC **end-to-end, solo** — from first sales response through build, validation, quote, and pricing structure (case study below)
- Helped **win a smart-home (wall-pad) voice-AI deal** by pairing hands-on technology validation with PoC / contract scope control (case study below)
- Contributed proposal direction, technical composition, architecture, and pitch-deck storyline to the **NIPA national semiconductor AI-agent program** — selected (case study below)

AI Service Development (RAG / LLM) • Saltware *Feb 2024 – Dec 2025*

- PM of the **idus hyper-personalization recommendation PoC** — requirements, recommendation-flow design, and technical validation through daily scrums with the customer; productized as **Sapie-Reco**, listed globally on AWS Marketplace (case study below)
- PM of **Sapie-Braille** (accessibility AI): validated the initial “AI braille transcription” request, then re-architected — standards-based braille conversion, AI in document understanding, alt-text, voice, and information agents (below)
- Built document-RAG chatbot PoCs end to end for a global QSR franchise, a large convenience-store franchise, Yeonsung University admissions, and Saltware Marketplace (parsing → indexing → vector DB → RAG → demo UI)

App / IoT Development • Saltware (Smart-farm team) *Feb 2023 – Jan 2024*

- Built a Flutter/Firebase smart-farm app with Raspberry Pi / sensor IoT integration — the hands-on engineering base behind today's fast PoC delivery

SELECTED CASE STUDIES • CUSTOMERS MASKED (NDA) • FULL DETAILS: HAJAEHWAN.COM

AICC Chatbot → Callbot PoC — solo end-to-end: sales → build → validation → pricing *Mar – Apr 2026*

- For a major enterprise customer evaluating an LLM-based callbot upgrade, led the full cycle: first sales response, requirements, PoC build, validation materials, quote, and operating-cost assumptions (M/M and delivery risk reviewed with the engineering lead)
- **Build:** hybrid RAG chatbot on ~250 official support documents/FAQs in ~1 week — FAISS semantic + BM25 keyword retrieval merged via RRF, with lightweight-LLM query routing/rewriting and relevance filtering to keep answers in supported scope — extended in ~1 more week into a callbot-compatible API (conversation-ID multi-turn; 294 pre-cached entries, 28-day TTL; health check & cache statistics)
- **Decision evidence** (per validation materials delivered to the customer): 96.7% single-turn / 95% multi-turn accuracy; all API responses <5s; ~0.1s on cache hit vs. ~5s on miss
- **Business case:** presented two build options (full RAG vs. retrieval-only with customer-side generation) and flagged AWS + LLM usage-fee volatility as a cost risk — restructured pricing to lower the per-channel fee instead of absorbing volatile usage costs, balancing customer cost against company risk

Sapie-Guardian — GenAI Security GTM, Packaging & B2G Strategy

Jun 2026 – Present

- Repositioned messaging from feature lists to buyer risk — private sector: control GenAI data-leak risk while preserving productivity, layered on existing DLP/proxy; public sector: certification, evaluation, and procurement eligibility as purchase enablers
- Designed lower-entry packaging: baseline pattern/keyword controls first, sLLM contextual detection and GPU capacity as expansion options; proposed AWS Marketplace listing and usage-based GPU operation to cut upfront cost
- Extended positioning beyond external-GenAI control to a **pre-ingestion sensitive-data filtering layer** for on-premise manufacturing AI — making semiconductor materials-and-equipment customers a focus vertical in coordination with the NIPA program, supporting pricing and proposal direction

NIPA Semiconductor-Manufacturing AI Agent Program — proposal contribution → selected *Mar – Jun 2026*

- Contributed to pivoting the proposal domain from an initial gaming-industry direction to semiconductor materials/parts manufacturing, after expert feedback that policy fit and industrial impact — not feasibility alone — drive national-program evaluation
- Built the storyline deck-first (“why now → why this industry → field problems → 4 workflows / 21 specialized AI agents → validation & diffusion”) and carried the structure, architecture diagrams, and key visuals into the formal business plan
- **Result:** passed the ~20:1 written round and won the 3:1 final presentation — ~KRW 1.55B year-1 / ~KRW 3.1B (=US\$2.2M) cumulative program secured

Smart-home Wall-pad Voice AI — technical validation & scope control → deal won *Dec 2025 – Mar 2026*

- Benchmarked Korean TTS candidates (gTTS, MeloTTS, Zonos, CosyVoice, Kokoro, Coqui) on language quality, install complexity, runtime environment, and latency behavior — evidence separating “technically possible” from “deliverable in this contract”
- Re-baselined scope against meeting minutes, functional specs, and contract revisions: narrowed the PoC to core device control + LLM validation, expressed sLLM/RAG latency as best-effort targets, and deferred high-risk features to follow-on discussion
- Defined post-delivery support boundaries (contract-based criteria separating defects from new requests), reducing schedule, scope-creep, and free-support risk — **deal won and delivered**

idus Hyper-Personalization Recommendation — first PM → productized as Sapie-Reco *2024 – 2025*

- Reframed the problem as **gift recommendation — the reference person is the recipient, not the buyer:** designed a conversational flow extracting recipient, relationship, occasion, and budget to narrow candidates over turns (LLM for intent, follow-ups, and explanations; structured fields for hard constraints like price band)
- Switched the vector store FAISS→Chroma when metadata-conditioned retrieval (price, category, occasion) proved more decisive than pure similarity — stabilizing candidate quality for cold-start users
- **Outcome:** productized as Sapie-Reco — 2x+ cold-start accuracy (publicly reported), AWS Marketplace global listing, AWS AI Solution Showcase recognition

ADDITIONAL PROJECTS

- **Fair Trade Open-Data AI (“RAG Gyeolsadae”)** — **team lead, Grand Prize:** recruited and led a 4-person team; hybrid BM25 + Chroma retrieval with metadata-scoped indexing (reranker dropped on Docker-size/stability grounds); led RAG and final presentations — Grand Prize, AI Model Development track (1 of 23 teams), Korea Fair Trade Commission · Jul 2025
- **Sapie-Braille — accessibility AI platform, PM/architecture:** kept braille conversion rule-/standards-based where accuracy is non-negotiable; placed AI in a multi-agent supervisor for document understanding, VLM alt-text, voice interface, and information agents → Procurement Encouragement Award, AI Champion Top 100, CES 2026 Innovation Award submission · 2024 – 2025
- **CardFollow — self-initiated internal sales-productivity app:** Excel/CSV-first contact workflow (column mapping, row-level review, event/template/attachment assignment), preflight checks (duplicates, attachments, Graph connectivity), and Outlook draft generation via Microsoft Graph — draft-only by default; pivoted from a multimodal-LLM card-extraction hypothesis (kept as auxiliary) on cost/latency/usability grounds · May 2026

AWARDS & RECOGNITION

- Grand Prize, AI Model Development — Korea Fair Trade Commission “1st Fair Trade Data Competition” (1 of 23 teams; team lead) · Jul 2025
- Selected — NIPA “AI Agent Convergence & Diffusion Support Program” (~KRW 3.1B cumulative; proposal contribution) · May 2026
- AWS AI Solution Showcase — Sapie-Reco (commercial product) selected as an outstanding solution, earning booth sponsorship for AWS AI x Industry Week 2025 (originating PoC PM) · Oct 2025
- Encouragement Award — Public Procurement Service “2025 Korea AI Innovative Procurement Expo” (Sapie-Braille, PM) · Sep 2025
- Top 100 pre-screening — MSIT “2025 AI Champion Competition” (~6:1) (Sapie-Braille, PM) · Aug 2025

CERTIFICATIONS & EDUCATION

- Databricks Certified Generative AI Engineer Associate (Apr 2026) · AWS Certified AI Practitioner + Early Adopter (Jan 2025) · AWS Certified Cloud Practitioner (Dec 2024) · Anthropic Claude Code 101 / Claude 101 (Apr 2026)
- B.Eng., Information & Communication Engineering — Anyang University · Mar 2015 – Feb 2022

SKILLS

- **GenAI · RAG:** RAG architecture, hybrid retrieval (BM25 + vector, RRF), retrieval evaluation, vector DBs (FAISS, Chroma), agent & prompt design, document parsing/extraction, caching, STT/TTS, GenAI security (AI DLP, PII masking, guardrails)
- **Pre-Sales · Technical BD:** customer discovery, solution framing, technical Q&A, PoC scoping & delivery, demos, proposals & architecture, quotes, tiered packaging & pricing strategy, B2G / procurement requirements, sales enablement
- **Prototyping · Cloud · AI-assisted delivery:** Python, FastAPI, React, Next.js, TypeScript, Docker, AWS; SSE streaming, Microsoft Graph; Codex & Claude Code with quality gates, preview-first / dry-run, human-review handoff