



WABS

With AI Be Sustainable

AI Agent Specialist · Since 2025

WABS

WABS specializes in solving real-world business challenges through AI agent technology. We design, develop, and operate AI solutions across logistics, energy, manufacturing, and business automation.

"We see problems first, not technology" — We deeply understand our clients' business challenges and deliver fully working solutions end to end.

Company	WABS Inc.
CEO	Hyunjun Choi
Established	2025-06-04
Address	401-166 A-ho, 4F, 79 Gimpohangang 9-ro, Gimpo-si, Gyeonggi-do, Republic of Korea
Email	wabs@wabs.io
Website	https://www.wabs.io/

VISION & MISSION

With AI, Be Sustainable

Contributing to a sustainable future with AI. We connect the potential of technology to real-world industry problems, creating lasting value rather than one-time solutions.

CORE VALUES

Problem-First

Real problem-solving is our reason for being. We start by deeply understanding our clients' business challenges.

Deep Completion

We take full responsibility until every project works in production. We deliver real services, not prototypes.

Big Bowl

We embrace problems beyond the boundaries of domain, technology, and people — connecting logistics, energy, AI, and automation as one.

Sustainable

We pursue long-lasting value, not quick fixes. Our designs consider operations and maintenance from the start.

OUR PEOPLE

WABS is a team of people who believe humanity can build a sustainable future together with AI.

01 Problem Seekers

We ask 'why' before 'how', defining problems in our clients' language.

02 Finishers

We don't back down from challenges — we prove our worth through working results.

03 Boundary Crossers

We go beyond our own expertise, connecting different domains and technologies.

04 Sustainability Advocates

We work with a long-term perspective, driven by the conviction that AI can contribute to humanity's future.

Solving Business Problems **with AI**

From consulting to development and operations — we define problems and deliver AI solutions.



AI Agent Design & Development

Custom AI agents built on LangGraph and multi-agent architectures to automate your workflows.

LangGraph

Multi-Agent

RAG



Full-Stack SaaS Development

End-to-end SaaS platform development from planning through deployment and operations.

Next.js

Supabase

Docker



Data Visualization & Analytics

Real-time dashboards, interactive 3D visualizations, and BI reports that turn data into insight.

React

Three.js

Recharts



MLOps & GPU Deployment

GPU-based AI inference pipelines deployed automatically through Kubernetes and Triton.

Kubernetes

Triton

Docker



LLM Integration & Agent Consulting

Integrate Claude, GPT, Gemini, and other LLMs into business workflows while optimizing reliability and cost.

Claude API

Prompt Caching

MCP



Synthetic Data & Test Automation

Synthetic data and automated test environments that stay safe and useful inside regulated industries.


LLM-First

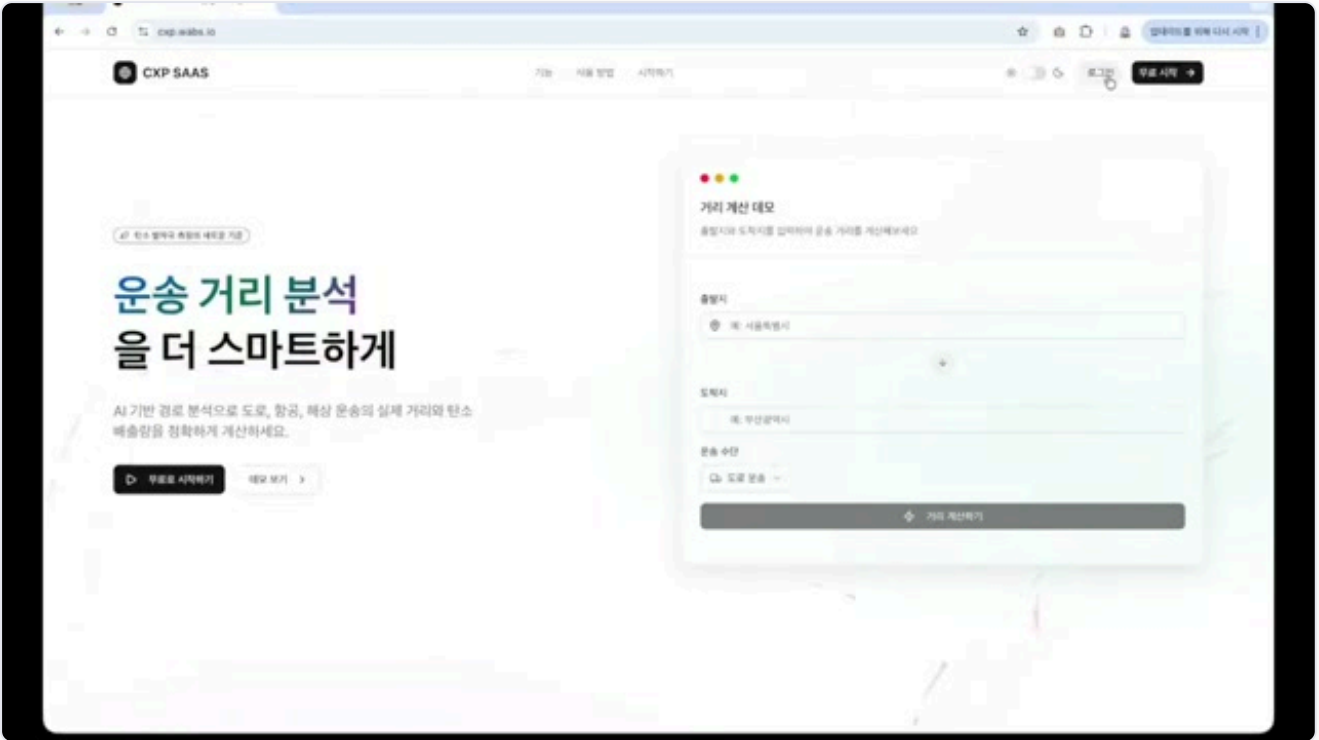
Synthetic Data

QA


CXP SaaS

LCA transport route analysis platform — from CSV upload to AI route analysis, carbon emission calculation, and visual reporting.

<p>TYPE WABS in-house solution</p>	<p>PERIOD 2025.12 - 2026.03</p>	<p>TEAM 2 engineers</p>	 <p>DEMO VIDEO</p>
---	--	------------------------------------	---



Architecture




CHALLENGE

Logistics teams and ESG officers were calculating per-route carbon emissions manually, which made data processing slow and segment-level analysis of multi-modal routes (road, air, sea) effectively impossible. Applying international emission factors consistently was also a persistent challenge.

SOLUTION

We built an automated geocoding and route-optimization engine powered by Google Gemini 2.5 Flash, packaged as a Next.js 16 + Supabase full-stack SaaS. The pipeline covers bulk CSV upload, AI analysis, interactive 3D globe visualization, and Excel/PDF report export end to end. A multi-tenant architecture provides organization/project-level collaboration with role-based access control.

Key Screens



Key Features

AI Route Analysis

Auto-geocodes origins and destinations with Google Gemini 2.5 Flash and analyzes optimal transport routes.

Multi-modal Decomposition

Automatically decomposes road, air, and sea routes into segments, calculating distance and emissions.

3D Globe Visualization

Visualize transport routes and carbon emissions on an interactive Three.js 3D globe.

Multi-tenant SaaS

Organization/project-level collaboration, role-based access control, and bulk CSV upload.

Achievements

90%+

Analysis time reduction

Road·Air·Sea

Supported transport modes

Excel/PDF

Report formats

Technology Stack

Next.js 16

React 19

Tailwind CSS 4

Supabase

Google Gemini 2.5 Flash

Leaflet

Three.js

AWS EC2


Docker

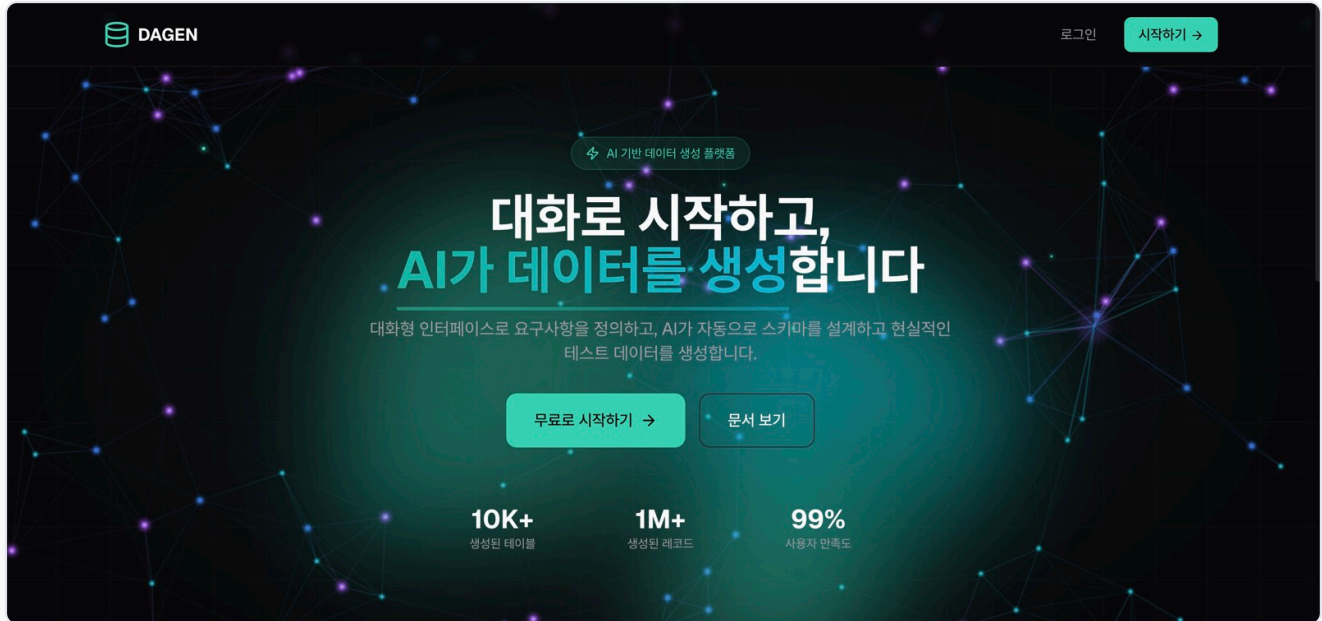
OUTCOME

Reduced route-analysis time by over 90% compared to the manual baseline. Multi-modal routes (road, air, sea) are decomposed into segments automatically with distance and emission calculations. International emission factors are applied consistently to deliver accurate, repeatable carbon-footprint data. Live at cxp.wabs.io.


DAGEN

LLM-first synthetic financial data platform — generates causally consistent records via persona-based Actor and journey simulation.

<p>TYPE</p> <p>WABS in-house solution</p>	<p>PERIOD</p> <p>2025.11 - 2026.04</p>	<p>TEAM</p> <p>2 engineers</p>	 <p>DEMO VIDEO</p>
--	---	---------------------------------------	---



Architecture



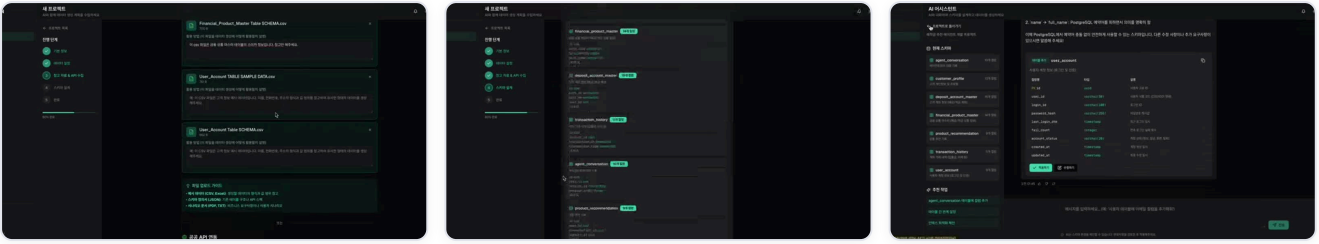
CHALLENGE

Building financial AI agents demands realistic test data that mirrors actual user behavior, yet regulation makes real financial records off-limits. Conventional random or rule-based synthesizers generate columns independently, so foreign-key integrity, date arithmetic, and monetary constraints cannot all be satisfied at once — leaving teams to patch data manually for every demo and test cycle.

SOLUTION

Built on a Claude API-centric LLM-first architecture that combines persona-driven Actor generation with journey-based simulation to guarantee row-level consistency. An LLM-driven schema classifier, event-to-table mapping, and a context pre-generation pipeline resolve business rules and foreign keys automatically. Additional capabilities include automatic ERD image analysis (Vision API), token-cost optimization through Prompt Caching, and SSE-based real-time progress monitoring. Five RDBMS targets (PostgreSQL/MySQL/MSSQL/SQLite/Oracle) and four export formats (CSV/SQL/JSON/Debug Dump) are supported out of the box.

Key Screens



Key Features

LLM-First Synthetic Data

Claude API-centric architecture generates financial synthetic data with row-level causal consistency through Persona and Journey simulation.

Automatic ERD Analysis

Vision API auto-parses ERD images to extract tables, relationships, and constraints, pre-generating context.

5 RDBMS · 4 Export Formats

Supports PostgreSQL/MySQL/MSSQL/SQLite/Oracle and CSV/SQL/JSON/Debug Dump exports.

Real-time Progress Monitoring

SSE-based real-time progress and log streaming with Prompt Caching for token cost optimization.

Achievements

5 engines

Supported RDBMS

4 formats

Export formats

**Row-level
100%**

Data consistency

Technology Stack

Next.js 14

Node.js 20 LTS

Express.js

Prisma

PostgreSQL

Redis / Bull Queue

Anthropic Claude API

Google Gemini API

Socket.io

Tailwind CSS

OUTCOME

Generates synthetic data that follows realistic financial journeys (signup → account opening → deposit → transfer) ready for immediate demo and test use. The LLM consistently enforces date arithmetic such as maturity = open date + term, monetary range constraints, and state-driven logic, removing the overhead of manual data verification. Live as a SaaS at dagen.wabs.io for external use.

Proven Capabilities

A track record of successfully building and delivering AI solutions across diverse industry domains. Each project is presented on its own dedicated page on the following pages.



ENERGY / MANUFACTURING 2025.06 - 2025.12

ENMS

Smart-factory energy monitoring and 3D visualization platform — real-time power surveillance, carbon emissions analytics, and metaverse integration.



AI / DEVOPS 2025.09 - 2026.01

Harness Monitor

AI skill management dashboard — search, visualization, automated configuration, and orchestration in one place.



BUSINESS AUTOMATION 2025.07 - 2025.11

LAMDA

Desktop AI agent for LG CNS in-house marketing analytics — automates BigQuery/GA4 data analysis and site-UX inspection from natural-language queries.



BUSINESS AUTOMATION 2025.01 - 2025.06

Meeting Assistant

GPU-powered AI meeting-minutes automation — audio upload, speaker diarization, STT, AI summarization, and email delivery.



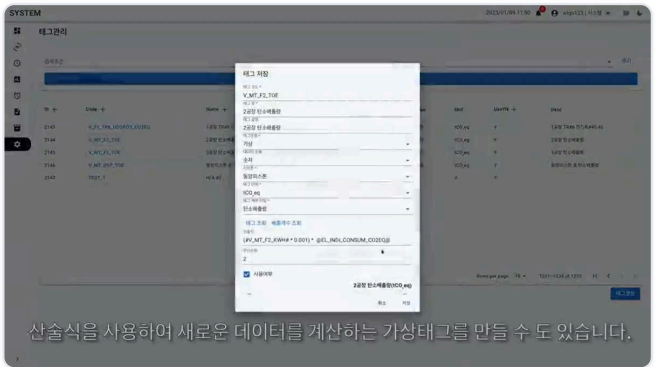
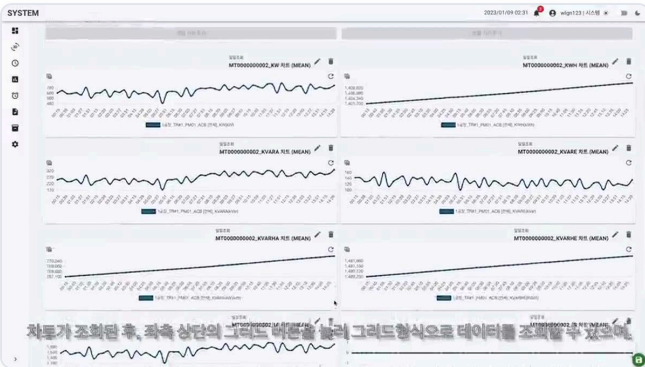
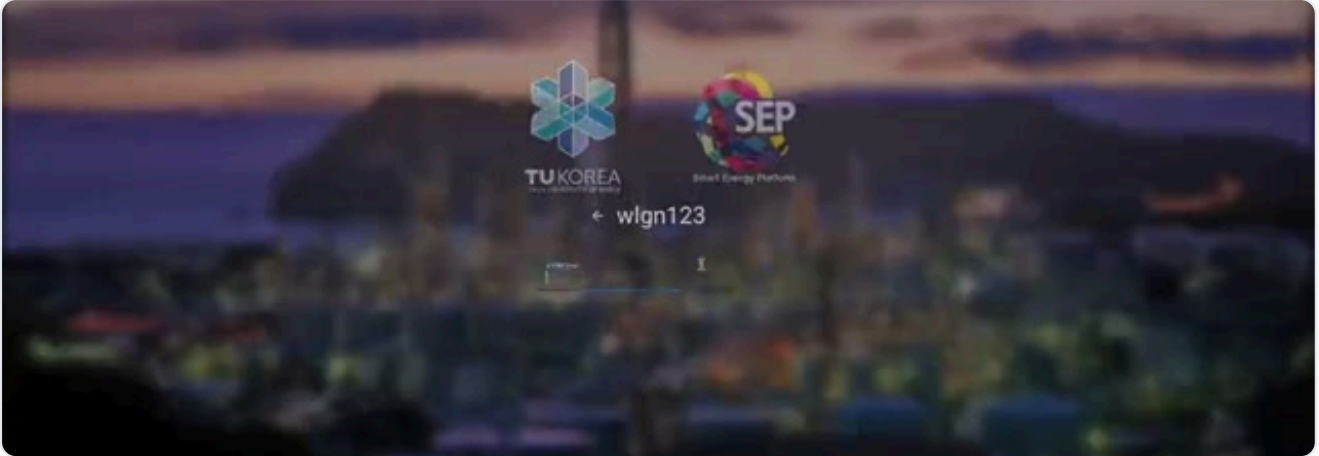
MULTI-AGENT OPTIMIZATION 2025.03 - 2025.09

Multi-AI Agent System

LangGraph-based multi-stage enterprise AI agent system — three-stage pipeline with six domain sub-agents.

ENMS

Smart-factory energy monitoring and 3D visualization platform — real-time power surveillance, carbon emissions analytics, and metaverse integration.



CHALLENGE

Smart factories lacked a system to monitor and manage energy consumption in real time, so energy waste could not be detected proactively.

SOLUTION

We built an IoT-sensor-based real-time energy data pipeline and rendered factory energy flows intuitively through a React dashboard with 3D metaverse...

OUTCOME

Real-time monitoring of plant-wide energy consumption enables instant detection of abnormal patterns.

24/7 Real-time monitoring IoT sensors Data ingestion 3D metaverse Visualization

React

TypeScript

Recharts

Spring Boot

InfluxDB

MariaDB

3D Metaverse



DEMO VIDEO


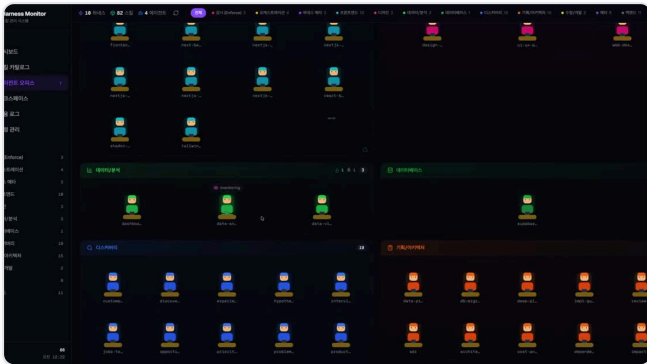
Scan the QR to watch the actual solution in action.
<https://www.wabs.io/portfolio/videos/enms.mp4>

AI / DEVOPS

2025.09 - 2026.01

Harness Monitor

AI skill management dashboard — search, visualization, automated configuration, and orchestration in one place.



CHALLENGE

Managing 62+ AI skills systematically and matching the optimal combination to each project was inefficient by hand.

SOLUTION

We developed a natural-language search and automated skill-composition system powered by Gemini 2.

OUTCOME

Centralized management of 62+ AI skills with natural-language search.

62+ Skills managed **3 stages** Auto-configuration **4-layer** Architecture

Next.js

TypeScript

Prisma

Gemini 2.5 Flash

Docker



DEMO VIDEO

Scan the QR to watch the actual solution in action.

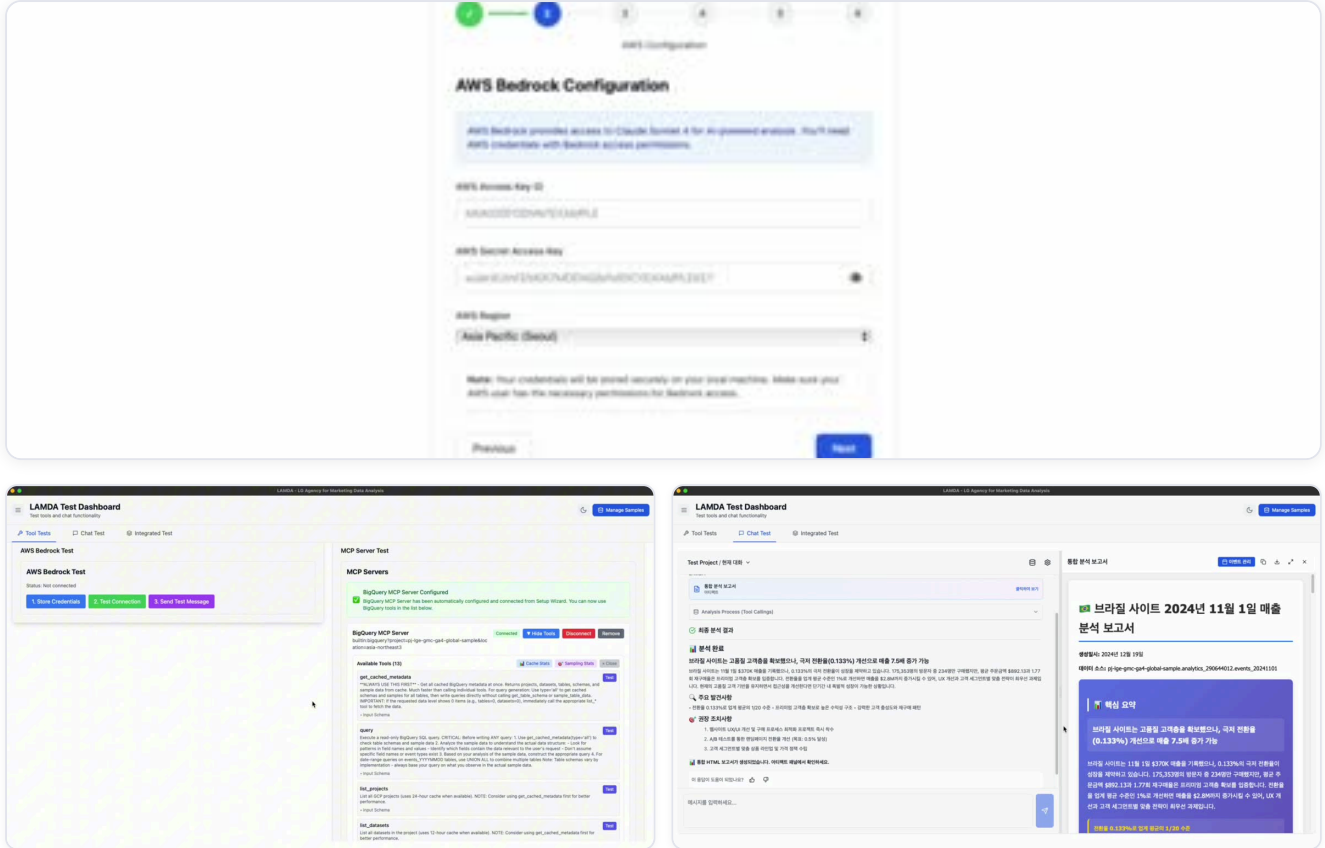
<https://www.wabs.io/portfolio/videos/harness-monitor.mp4>

BUSINESS AUTOMATION

2025.07 - 2025.11

LAMDA

Desktop AI agent for LG CNS in-house marketing analytics — automates BigQuery/GA4 data analysis and site-UX inspection from natural-language queries.



CHALLENGE

The LG CNS marketing team needed to analyze GA4 data in BigQuery alongside live website UX, but no single tool could combine quantitative analysis...

SOLUTION

We anchored the analytics engine on AWS Bedrock Claude Sonnet 4, integrating a BigQuery MCP server with metadata caching, adaptive sampling, and a...

OUTCOME

Delivered to LG CNS as an in-house tool covering the full flow of natural-language question → automated BigQuery generation → validation → chart and...

5 agents Specialist agents **8 formats** Artifact formats **Up to 10 rounds** Multi-turn tool calls

- Tauri 2.x (Rust)
- React 18 + TypeScript
- AWS Bedrock (Claude Sonnet 4)
- Google BigQuery
- GA4 API
- Model Context Protocol (MCP)
- Playwright
- SQLite
- Vite
- Chart.js / Plotly



DEMO VIDEO

Scan the QR to watch the actual solution in action.
<https://www.wabs.io/portfolio/videos/lamda.mp4>

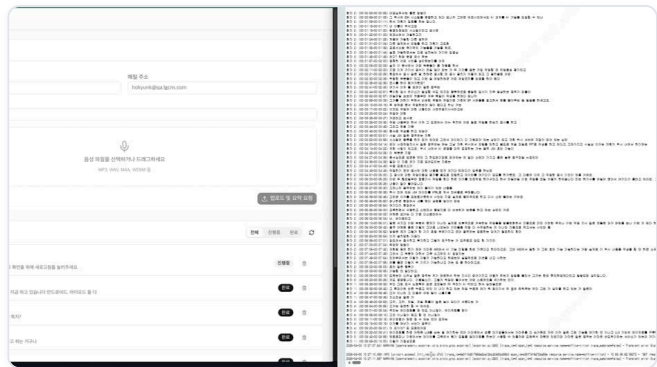
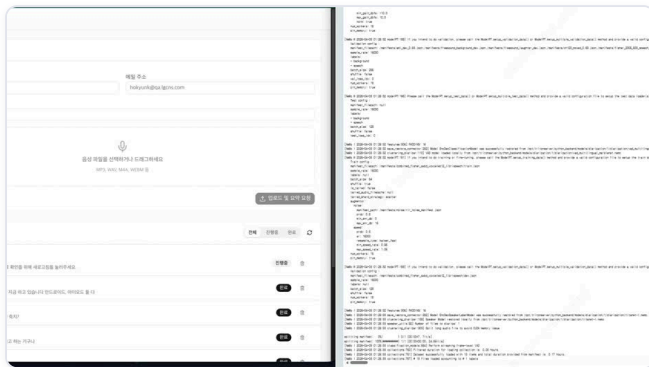
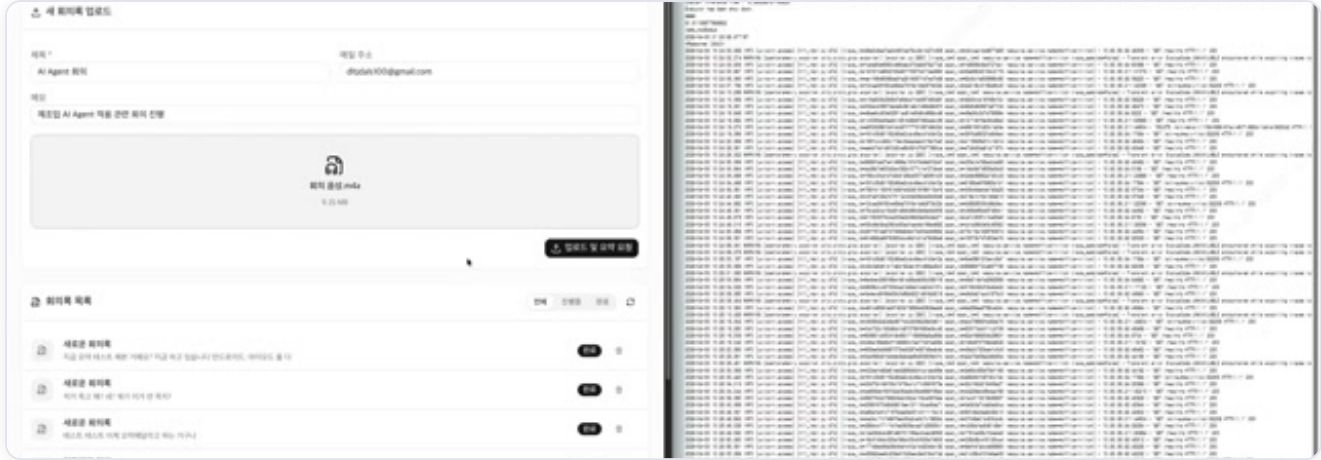


BUSINESS AUTOMATION

2025.01 - 2025.06

Meeting Assistant

GPU-powered AI meeting-minutes automation — audio upload, speaker diarization, STT, AI summarization, and email delivery.



CHALLENGE

Drafting meeting minutes consumed excessive time, and low speaker-identification accuracy forced manual transcription.

SOLUTION

We built a six-stage AI pipeline: audio upload (MinIO + RabbitMQ) → pre-processing (FFmpeg) → speaker diarization (NeMo Diarization + Triton) →...

OUTCOME

Fully automated the path from raw recording to structured minutes.

6-stage AI Pipeline **Private/Hybrid/Cloud** Deployment options **Triton Server** GPU acceleration

FastAPI

NVIDIA Triton

NeMo Diarization

Whisper v3

GPT-4.1

Kubernetes

RabbitMQ

MinIO

MongoDB

Grafana



DEMO VIDEO

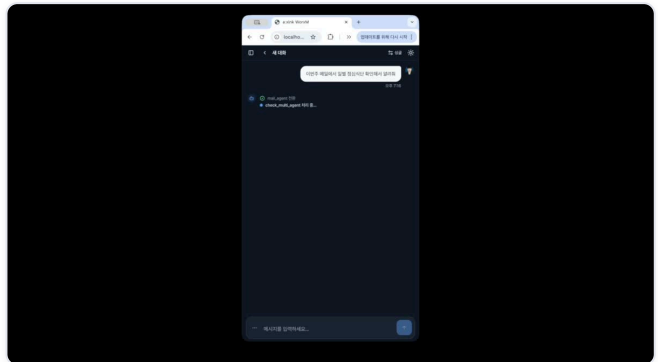
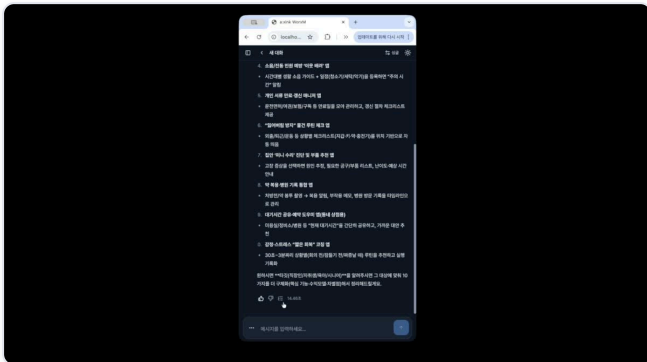
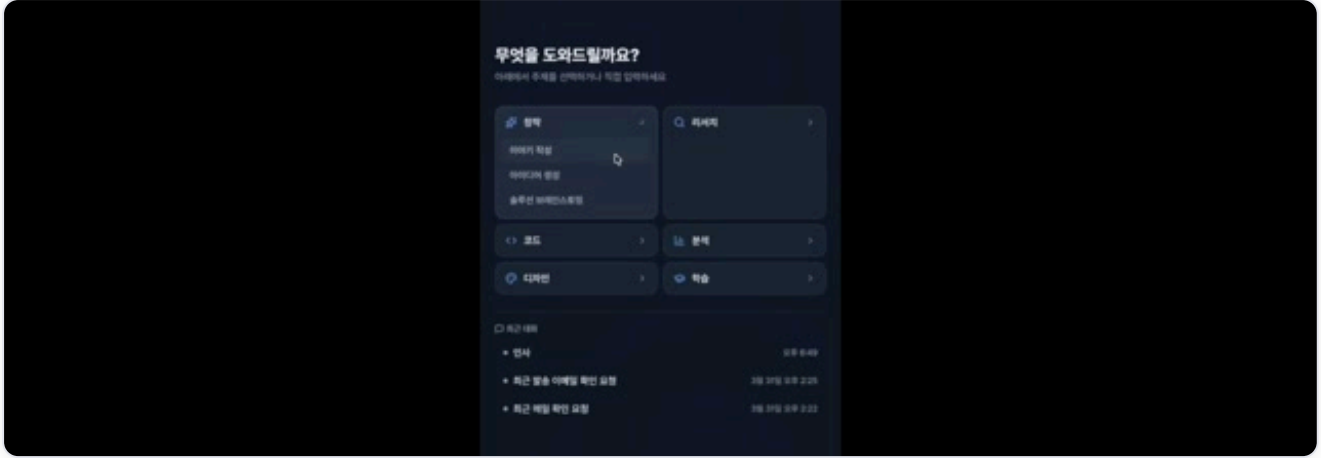
Scan the QR to watch the actual solution in action.

<https://www.wabs.io/portfolio/videos/meeting-assistant.mp4>



Multi-AI Agent System

LangGraph-based multi-stage enterprise AI agent system — three-stage pipeline with six domain sub-agents.



CHALLENGE

Handling cross-domain queries with a single agent was difficult, and building an enterprise-grade multi-agent system ran into technical limits around...

SOLUTION

We designed a three-stage LangGraph pipeline (Foundation → Understanding → Behavior) with each stage independently compilable, testable, and...

OUTCOME

A single POST /chat endpoint supports auth propagation, dynamic stage selection, and streaming.

6 domains Sub-agents **3 stages** Pipeline **3 layers** Memory tiers

- LangGraph
- LangChain
- OpenAI GPT-4o
- FAISS
- FastAPI
- PostgreSQL
- Docker



DEMO VIDEO

Scan the QR to watch the actual solution in action.
<https://www.wabs.io/portfolio/videos/multi-ai-agent.mp4>

Overall Capabilities

7 Projects delivered	6 Domains covered	52+ Technologies	Full- stack AI Team expertise
--------------------------------	-----------------------------	----------------------------	---

Technology Stack

Frontend Next.js 16, React 19, Tailwind CSS 4, React, TypeScript	Backend & Database Supabase, Prisma, Spring Boot, InfluxDB
AI & ML Google Gemini 2.5 Flash, LangGraph, LangChain, OpenAI GPT-4o	Infra & DevOps AWS EC2, Docker, Kubernetes

Milestones

- 2025.06** WABS incorporated — AI agent specialists
- 2025.06-12** Delivered ENMS, Harness Monitor, Multi-AI Agent, Meeting Assistant projects
- 2025.12-** Developed CXP SaaS platform and launched cxp.wabs.io
- 2026.04** 7 projects across 5 domains, 52+ technologies



WABS

With AI Be Sustainable

EMAIL

wabs@wabs.io

WEBSITE

(main) <https://www.wabs.io/>

(reference) <https://www.wabs.io/portfolio>

ADDRESS

401-166 A-ho, 4F, 79 Gimpohangang 9-ro,
Gimpo-si, Gyeonggi-do, Republic of Korea

Build the Future **Together**

Have a problem you'd like to solve with AI?

Let's build a sustainable solution together.

[Get in Touch](#)