



# Philip Vana

## AI / GenAI Engineer

RAG · AWS Bedrock · Retrieval Engineering · LLM Evaluation

## PROFILE

AI / GenAI Engineer with end-to-end ownership of a RAG platform in the public sector — from the ingestion pipeline through retrieval optimization to an evaluated AWS Bedrock deployment. My background as a data engineer (ETL, data warehousing, large-scale data) shapes a system-first, data-driven approach — optimized for reproducibility, measurability, and operations. I think in systems, data flows, and failure modes.

## EXPERIENCE

### Data Engineer · GenAI / RAG Focus

2024 – Present

Deutsche Rentenversicherung Bund, Berlin

- Single-handedly designed and built the **end-to-end RAG pipeline** of an internal GenAI assistant — a self-designed deterministic chunking strategy and an AWS Bedrock Knowledge Base over ~800 documentation pages. Deployed and in user-acceptance testing ahead of production rollout (Aug 2026).
- Optimized the retrieval layer data-driven — **Recall@5 53% → 67%** (78% on expert-validated pilot queries), MRR 0.50 → 0.66 — each lever measured in isolation against a self-built **4-pillar evaluation harness** (537 tests, CI regression gate, automated stakeholder reporting).
- Built an **EU-data-resident, agentic AI development environment** in AWS SageMaker (opensource → AWS Bedrock via a LiteLLM proxy) — credential-free via IAM execution roles, automated single-file bootstrapping, self-healing lifecycle.
- **Data-engineering foundation** in the BI team (SAP HANA / data warehouse): end-to-end ownership of production batch ETL at the billion-record scale — monitoring, metadata, root-cause analysis, production deployments.

### Software Engineer (Internship)

06/2023 – 08/2023

Vaultree Ltd., Dublin

- Object-oriented backend development (Java, Python) in an international, product-oriented team; encryption-focused API applications (PostgreSQL, REST APIs, Vaultree Encryption API).

**Earlier experience (2014–2021):** project & event management (Vienna) and vehicle-inspection engineering (KÜS), prior to transitioning into tech.

## EDUCATION

### IT Specialist, Data & Process Analysis (Dual Vocational Training) · Grade 1.7

2021 – 2024

Deutsche Rentenversicherung Bund, Berlin · German grade scale, 1.0 = best

### Higher Technical Secondary Education (Fachabitur) · Engineering Focus

2012 – 2014

Fachoberschule Technik, Munich

## CONTACT

Berlin, Germany

+49 163 5217112

philip.vana95@gmail.com

philip-vana.com

## CORE SKILLS

### Generative AI / RAG

RAG · AWS Bedrock (Knowledge Bases, Titan, Rerank) · Vector Search · Embeddings · Chunking · Query Expansion · LLM Evaluation (Recall@k, MRR, nDCG) · LLMops

### Cloud & Infrastructure

AWS (Bedrock, SageMaker, S3, Lambda, DynamoDB, Cognito, API Gateway, CloudFront) · IAM · AWS CDK · CodeBuild · LiteLLM

### Data Engineering

ETL/ELT · Data Warehousing · Data Modeling · Batch & Chunk Processing · SAP HANA XSA · Oracle DB · PostgreSQL

### Languages & Tools

Python · SQL · FastAPI · Pytest · Git · Scrum

## LANGUAGES

German native

English C1, fluent