

# Alex Narayan

## Senior Data Platform Engineer

[alex@alexandarnarayan.com](mailto:alex@alexandarnarayan.com) · [gigatexal.blog](https://gigatexal.blog) · [linkedin.com/in/alexandarnarayan](https://linkedin.com/in/alexandarnarayan)

Berlin, Germany · EU Blue Card

## Summary

---

Data platform engineer with 13+ years across database administration, Python development, SRE, and data engineering. Designs and operates production data platforms on Snowflake, AWS, and GCP. Experience includes warehouse and database migrations and GDPR data masking and right-to-be-forgotten processes. Currently hands-on engineering manager of Blacklane's data platform team (~60% engineering).

## Experience

---

### Blacklane GmbH · Berlin

**Engineering Manager - Data Platform** · May 2026 – Present

**Data Platform Engineer** · Oct 2024 – May 2026 · *technical lead for the platform team from Jul 2025*

Hands-on lead (~60% engineering) for the AWS/Snowflake data platform; team of 3 engineers (staff, senior, junior).

#### Architecture and Platform

- Set technical direction and roadmap; own design sign-off and RFCs
- Operate near real-time GPS ingestion (SQS, Kinesis, Data Firehose, S3) with no custom code, built to scale from ~50M records/day today to 500M+ without redesign
- Operate the company's sole Airflow instance (MWAA): ingestion DAGs and dbt-core orchestration via Cosmos
- Maintain AWS infrastructure in Terraform, 80+ Fivetran sources, and a co-maintained legacy Java pricing application
- Hold Snowflake spend growth to 10% for the current contract term despite multiplying data volume, through warehouse right-sizing and tech-debt removal
- Migrated the AWS EMR data warehouse to Snowflake, including the dbt project and related Airflow DAGs

#### Governance and Security

- Own the request workflow (Jira/Confluence service desk) for company-wide GDPR data masking in Snowflake and Omni; review the masking implementation
- Review the design and rollout of the company's first GDPR right-to-be-forgotten process across the data stack
- Govern access through layered RBAC (database roles composed into functional roles, PII denied by default) with Okta SSO onboarding; ~90% of engineers connect AI agents (Claude, MCP, Snowflake CLI) to curated production data
- Own business-critical and regulatory data feeds: airline customer reporting, a regulatory feed to the Riyadh transport authority, and Salesforce reverse ETL

#### AI on Snowflake

- Direct Snowflake Cortex cost and quality optimization: model benchmarking for classification and estimation, AI-function query rewrites, spend monitoring, and published guidance
- Align new AI initiatives with the existing Snowflake platform to avoid duplicate infrastructure

#### Technical Leadership

- Set dbt standards, run dbt upgrades, and deliver Git training for the analytics team
- Provide Terraform and AWS infrastructure support to Data Science
- Coach engineers from junior to staff through pairing, clear expectations, and regular feedback
- Organize the on-call rotation

### Smartbroker AG · Berlin

**Lead Data Engineer** · Nov 2023 – Oct 2024

- Led 4 senior engineers across AWS and GCP
- Designed a central BigQuery warehouse, cutting insight delivery from days to hours
- Owned engineering practices: source control, linting, dev workflows, retrospectives, documentation, mentoring

### Delivery Hero GmbH · Berlin

**Senior Data Engineer** · Jun 2020 – Nov 2023

- Migrated a 5TB+ Redshift warehouse to BigQuery: fact/dimension modeling, partitioning, clustering, Airflow

DAGs

- Contributed to the in-house Google Analytics 4 replacement as an engineer on the platform team
- Established cost-effective BigQuery practices (partition filtering, anti-pattern reviews)
- Built a GitHub Actions CI/CD pipeline for AWS Lambda (HR analytics)
- Redesigned hiring with pair-programming interviews in place of take-home projects, cutting time-to-offer from 1–2 months to 3 weeks
- Raised code documentation coverage to 95%+; authored internal Git/GitHub guides and a [video course](#); trained teams on dataset use

## **Finleap Connect (formerly figo GmbH) · Berlin**

**Site Reliability Engineer** · Oct 2019 – Mar 2020

**Python Software Engineer** · Nov 2018 – Oct 2019

- Migrated a 1TB+ product database to CockroachDB on Kubernetes, with export and validation tooling that saved days of engineering time
- Built a streaming data lake (Pub/Sub, Cloud Functions, BigQuery), increasing reporting capability 5x; built monitoring dashboards and advised on SLAs
- Developed features for the flagship banking API (Python/Flask)

## **CorVel Corporation · Portland, OR**

**Data Warehouse Engineer** · Sep 2017 – Apr 2018

**Database Administrator** · Aug 2013 – Sep 2017

- Owned deployments across 21+ SQL Server databases and 3 environments; tuned queries, secured and performance-tested servers, and shared on-call
- Extended dimension and fact tables and the internal C#/SSIS ETL tool; resolved data-quality issues
- Built Python/C# tooling for DBAs, maintained source control, and trained engineers

## **Skills**

---

**Data platform:** Snowflake (Cortex, RBAC, masking policies, cost management), dbt, Airflow (MWAA, Cosmos), Fivetran, Omni, Streamlit, BigQuery, Redshift, data modeling, ETL/ELT

**Cloud and infrastructure:** AWS (S3, SQS, SNS, Kinesis, Data Firehose, Lambda, EMR), GCP, Terraform, Kubernetes, GitHub Actions, Linux

**Languages:** SQL, Python, Java (maintenance)

**Governance and AI:** GDPR (masking, right to be forgotten), access control, data quality, MCP, AI agents

**Leadership:** technical direction, RFCs, mentoring, hiring, cross-team adoption, training and public speaking

## **Education**

---

**BSc Financial Economics** · Brigham Young University – Idaho · 2009 – 2013