

Personal Overview

Behnam is an experienced Full-Stack and Data Engineer with close to a decade spent designing, building, and tuning data-centred systems. He has a solid record of shaping relational models, optimising queries, and enabling reliable reporting and analytics.

At Siteagle Tech he simplified the firm's asset-tracking store in PostgreSQL/PostGIS by moving to a cleaner dimensional layout, trimming average report times by about 70 percent. Day-to-day he works with fact and dimension tables, manages slowly changing attributes, and keeps storage costs down through smart indexing and partitioning.

Behnam is comfortable across SQL Server, PostgreSQL/PostGIS, MySQL, and Oracle and writes ETL pipelines in Python, .NET Core, SSIS, and Azure Data Factory. He has built feeds that land high-volume IoT telemetry and uses incremental loads with reconciliation checks to keep data accurate. While at Mealzo in the UK, his refactor of the operational-to-warehouse flow on Azure cut Power BI refresh times by roughly 70 percent and halved the overnight load window.

A natural collaborator, Behnam works closely with product managers, analysts, and compliance teams to capture business rules, tighten data-governance processes, and introduce ISO-aligned security controls. He is at home in cloud techs, packages services with Docker and Kubernetes, and follows Agile ways of working. Backed by certifications in advanced SQL and database normalisation, he consistently delivers well-governed, high-performing data solutions in demanding environments

Key skills & Technologies

- Data Modelling & Warehousing
 - Star & Snowflake schema design
 - Slowly Changing Dimensions
 - Fact table optimisation
- ETL & Orchestration
 - SSIS
 - Azure Data Factory
 - Python (Pandas, Airflow)
 - .NET Core
- Databases
 - SQL Server
 - PostgreSQL/PostGIS
 - Oracle
 - Snowflake
- BI & Visualisation
 - Power BI
 - Tableau
- Cloud & DevOps
 - AWS (EC2, S3, Lambda)
 - Azure (VMs, Functions)
 - Docker
 - Kubernetes
 - Git
- Other
 - Data governance & security
 - Query tuning & indexing
 - Agile / Scrum

Education & Qualifications

- Nabi Akram University Master of Science – MS Artificial Intelligence
- University of Applied Science and Technology Bachelor's degree, Computer Software Engineering
- Database Normalization and Design Principals (Oracle University)
- Advanced SQL Programming (Microsoft)

Experience

SiteAgle Tech (Canberra)

Data Warehouse and Full-Stack Engineer

2025 – Present

Responsibilities:

- Redesigned the company's asset-tracking warehouse around a star schema, cutting average report run-time from 18 s to under 5 s.
- Built Python-based ETL jobs that land IoT telemetry in PostgreSQL/PostGIS every two minutes, handling bursts of 75k events/min with zero data loss.
- Worked side-by-side with the BI team to define grain, dimensions and slowly changing attributes, ensuring new dashboards could be delivered without re-work.
- Introduced partitioning and columnar storage for large fact tables, reducing storage costs by 23 %.
- Set up automated data-quality checks and role-based security to satisfy ISO 27001 audit requirements.

Mealzo (UK)

System and Database Analyst

2024 – 2025

Responsibilities:

- Migrated a monolithic OLTP database to a dedicated Azure SQL data-warehouse layer using a snowflake design; nightly load window shrank from 4 h to 65 min.
- Created SSIS and Azure Data Factory pipelines to ingest sales, marketing and third-party logistics feeds.
- Optimised key fact tables with bitmap filtering and clustered column-store indexes, improving Power BI refresh times by 70 %.
- Worked with finance and marketing stakeholders to document business logic and data governance rules. • Prepared capacity-planning forecasts that guided the move to Azure Synapse, saving an estimated \$22k p.a.
- Delivered five green-field data-mart projects for SMEs, each built around star schemas and exposed via RESTful APIs.
- Custom ETL frameworks written in .NET Core and Python (Pandas, SQLAlchemy).
- Implemented incremental load strategies, change-data capture and automated reconciliations to guarantee data integrity.

Self Employed

Full Stack/Database Developer

2020 – 2023

- Delivered five green-field data-mart projects for SMEs, each built around star schemas and exposed via RESTful APIs.
- Custom ETL frameworks written in .NET Core and Python (Pandas, SQLAlchemy).
- Implemented incremental load strategies, change-data capture and automated reconciliations to guarantee data integrity.

Manir

SQL Database Specialist & Software Developer

2016 – 2020

Responsibilities:

- Designed SQL Server data marts for national electricity clients; heavy use of surrogate keys, SCD logic and partitioned tables.
- Built SSRS and Tableau dashboards that reduced manual reporting effort by 80 %.
- Optimised stored procedures and indexes, pushing key billing queries from minutes to seconds.

Selected Project Highlights

ISODCTI – Integrated System of Data Centre Technical Information

- Architected the warehouse layer (SQL Server 2008) using a star schema; provided near real-time reporting on 130+ data centres.

Telecommunication Infrastructure Documentation Portal

- Delivered a web-based document repository (ASP.NET 2008) backed by a dimensional model; improved search accuracy by 40 %.

Bill Management System – Azerbaijan Regional Electricity Company

- Re-engineered legacy reports into a modern Power BI suite fed by a Snowflake-style warehouse, slashing month-end processing from three days to six hours.

Training Management System – National Iranian Oil Refining & Distribution

- Implemented conformed dimensions across HR and training data, enabling cross departmental analytics for the first time.