

Georgi Daskalov

Senior Senior Data and ETL Engineer

SUMMARY

- 15 years of experience in software and data engineering. Expertise in designing and implementing cloud-native data lakehouse solutions using Snowflake, Azure Databricks, and Apache Kafka with CDC and Lambda/Kappa architectures.
- Proficient in Go, Python, SQL, PL/SQL, ETL, dbt, Airflow, Debezium, and CI/CD pipelines with GitLab, enhancing data ingestion and transformation workflows. AWS and Azure cloud services.
- Led development of zero-downtime upgrade features for VMware vSphere, holding patents for innovative software-defined data center upgrade methods.
- Data consulting for Neoshare AG, DSK Bank, OpenTag EOOD, VMware, CODIX S.A and IBM.

TECHNICAL SKILLS

Main Technical Skills	Snowflake (2 yr.), Databricks (1 yr.), SQL, ETL Pipelines, AWS Glue
Programming Languages	Go (2 yr.), Python (2 yr.)
.NET Platform	ASP.NET Core Framework (4 yr.)
Mobile Frameworks and Libraries	Streaming
Data Analysis and Visualization Technologies	Airbyte (1 yr.), Apache Airflow (2 yr.), Databricks (1 yr.), ETL, ETL Pipelines, Fivetran (1 yr.), Lakehouse, OLAP
Databases & Management Systems / ORM	Data Lake, dbt (2 yr.), DWH, Flyway, IBM DB2 (5 yr.), Oracle Database (5 yr.), PL/SQL, PostgreSQL (2 yr.), Snowflake (2 yr.), SQL
Cloud Platforms, Services & Computing	Azure (1 yr.)
Amazon Web Services	AWS Glue, AWS Lambda
Azure Cloud Services	Azure Databricks, Azure MSSQL (3 yr.), Databricks (1 yr.)
Industry Domain Experience	Banking, Gambling & Betting, retail
Methodologies, Paradigms and Patterns	Agile, FDD, microservices, microservices architecture, Scrum, SDLC
Web/App Servers, Middleware	Apache Kafka
	Atlassian Confluence



Collaboration, Task & Issue Tracking	
Scripting and Command Line Interfaces	Batch
Deployment, CI/CD & Administration	CI/CD (2 yr.)
Project Management & Administration	Data Models
Platforms	DBX
Virtualization, Containers and Orchestration	Docker, VMware vSphere
Version Control	Git (2 yr.), GitLab
Blockchain and Decentralized Software	NEAR
Message/Queue/Task Brokers	RabbitMQ (2 yr.)
Other Technical Skills	CDC, Cloud Native, Dagster, Data Model, Data Platform, Data Processing, Data Transformation, Debezium (2 yr.), OLTP, Schema, snapshot

WORK EXPERIENCE

Senior Data Engineer - Neoshare AG (Cloud Native Data Lake House Solution)

Duration: 10 Apr 2025 – Current

Summary:

- Design and build schemas to solve the day-to-day problems of the market and intelligence forecast team
- Designed and built ETL Pipelines, AWS Glue
- Built and deployed a cloud-native Data Lake house solution using Snowflake, implementing CI/CD processes and a data ingestion orchestrator to streamline data workflows

Responsibilities:

- Design and build data schemas for market and intelligence forecasting.
- Build and deploy a cloud-native Data Lake house solution using Snowflake.
- Create and implement GitLab CI/CD processes.
- Create and implement data ingestion processes in collaboration with business stakeholders.
- Create and implement a data ingestion orchestrator.
- Implement schema migration tooling in the delivery process.

Technologies: AWS Glue, Snowflake, GitLab CI/CD, Data Lakehouse, Data ingestion orchestrator



Data Platform Architect - DSK Bank (Retail Banking Data Platform Redesign and ECB Regulatory Reporting Migration)

Duration: 2 Oct 2023 – 7 Apr 2025

Summary:

- Redesign and optimization of retail banking data platforms, including migration of ECB regulatory reporting from on-prem Oracle DWH to Azure Databricks with snapshot and micro-batch data processing
- Implementation of modern data architectures and CI/CD practices in a legacy environment

Responsibilities:

- Deep study and redesign of retail banking data model and analytics platforms.
- Analyze, refactor, and optimize existing business-specific data models.
- Establish business work-related processes for Data Engineering teams.
- Develop, implement, and document data processing information systems.
- Optimize legacy information system infrastructure.
- Research and develop CI/CD philosophy in a non-CI/CD SDLC environment.
- Design architecture and implement CDLH on Azure with Databricks.
- Implement and propagate CDC principles and Lambda/Kappa architectures.
- Re-create ECB regulatory reporting from on-prem Oracle DWH to Azure DBX platform.
- Align cross-team corporate data strategy.

Technologies: Azure, Databricks, Oracle DWH, CDC, Lambda architecture, Kappa architecture, CI/CD

Senior Data Engineer - OpenTag EOOD (Cloud-Based OLAP Platform and Data Streaming Implementation)

Duration: 3 Jan 2022 – 7 Sep 2023 (1 yrs 9 mos)

Summary:

- Development of cloud-based OLAP near real-time platform and OLTP microservices architecture for the sports betting industry
- Migration from on-prem batch processing DWH to modern cloud DWH using streaming data solutions and scalable ETL pipelines

Responsibilities:

- Examine and study sports betting industry data models.
- Implement OLTP microservices architecture and build a cloud-based OLAP near-real-time platform.
- Analyze, redesign, and reintegrate on-prem batch processing DWH to cloud-based DWH using Apache Kafka streaming.
- Design, develop, and implement proprietary scalable ETL pipelines following landing staging data flow philosophy.
- Implement custom data model logic and ETL scheduling with client-specific dependencies.
- Implement Apache Kafka data streaming using CloudSQL and Debezium CDC on Postgres.
- Use Flyway for DWH versioning.
- Implement DBT and Airflow for data transformation and processing of complex datasets.



- Implement custom business data ingestion using Fivetran and AirByte for unstructured data ingestion.

Technologies: Apache Kafka, CloudSQL, Debezium, Postgres, Flyway, DBT, Airflow, Fivetran, AirByte, OLTP microservices, OLAP

Software Developer / MTS 3 - VMware (vCenter Product Data Migration Framework and Zero Downtime Upgrade Features)

Duration: 10 Sep 2018 – 3 Jan 2022

Summary:

- Design, development, integration, and maintenance of a proprietary data migration framework supporting vCenter product upgrades
- Development of zero downtime upgrade features for vSphere product and internal database tools

Responsibilities:

- Design, develop, integrate, and maintain a custom data migration framework for vCenter product upgrades.
- Develop and maintain conceptual, logical, and physical data models of the product.
- Design and develop new features for the vSphere product enabling zero downtime upgrades.
- Design and integrate internal tools such as database health checker and cleanup utilities.
- Provide expertise and support to other product teams on MSSQL, Oracle, DB2, and Postgres platforms.

Technologies: MSSQL, Oracle, DB2, Postgres, Data migration frameworks, vSphere

Oracle Database Developer / Performance and Tuning Expert - CODIX S.A (Database Development and Performance Optimization)

Duration: 10 Nov 2016 – 5 Sep 2018

Summary: Development and optimization of Oracle databases including SQL and PL/SQL code, database tools, and refactoring of SQL applications to improve performance and maintainability.

Responsibilities:

- Database development and design using SQL and PL/SQL.
- Optimize and tune SQL and PL/SQL code for performance.
- Develop database tools to support application needs.
- Analyze, debug, and optimize SQL and PL/SQL queries and packages.
- Design SQL and PL/SQL packages, queries, functions, procedures, and triggers in C and ProC programs.
- Refactor SQL applications for improved performance and maintainability.

Technologies: Oracle, SQL, PL/SQL, C, ProC

Remote Software Support Engineer - IBM

Duration: 11 Aug 2011 – 16 Nov 2015



PROJECTS

11 May 2022 – 13 Feb 2024

Reduced downtime during upgrade of an application hosted in a data center

A method of upgrading an application in a software-defined data center (SDDC) includes:

- deploying, by lifecycle management software executing in the SDDC, a second appliance, a first appliance executing services of the application at a first version, the second appliance having services of the application at a second version, the services in the first appliance being active and the services in the second appliance being inactive;
- expanding, by the lifecycle management software, the state of the first appliance to support both the services at the first version and the services at the second version; replicating, by the lifecycle management software, the state of the first appliance to the second appliance;
- performing, by the lifecycle management software, a switchover to stop the services of the first appliance and start the services of the second appliance; and contracting, by the lifecycle management software, state of the second appliance to remove a portion unused by the services at the second version.

<https://patents.google.com/patent/US20230106877A1/en?q=US20230106877A1>

EDUCATION

- **Bachelor's Degree, Engineer of Electronics**
Technical University of Sofia
15 Sep 2009 – 16 Feb 2015

