



esentri

Modern Data Stack Hands-on
building a cloud data platform with Snowflake

Business Analytics Day 2024

7.5.2024 an der Hochschule Offenburg

Agenda

- Motivation
- Modern Data Stack
- Snowflake as Data Platform
- DEMO
- Lessons Learned
- Discussion

Agenda

- Motivation
- Modern Data Stack
- Snowflake as Data Platform
- DEMO
- Lessons Learned
- Discussion

Cloud Data Platform with WHY

Snowflake enthusiasts



NICOLAS HAUBNER
Senior Data Engineer



FILIP STEPNIAK
Senior Data Engineer

Cloud Data Platform with ❄️ WHY

customer's success story

Status Quo & Pain points (in 2021)

- Little overview of business & KPIs
- Many ad-hoc reports / SQL queries
- Dependence on few people & resources

Requirements

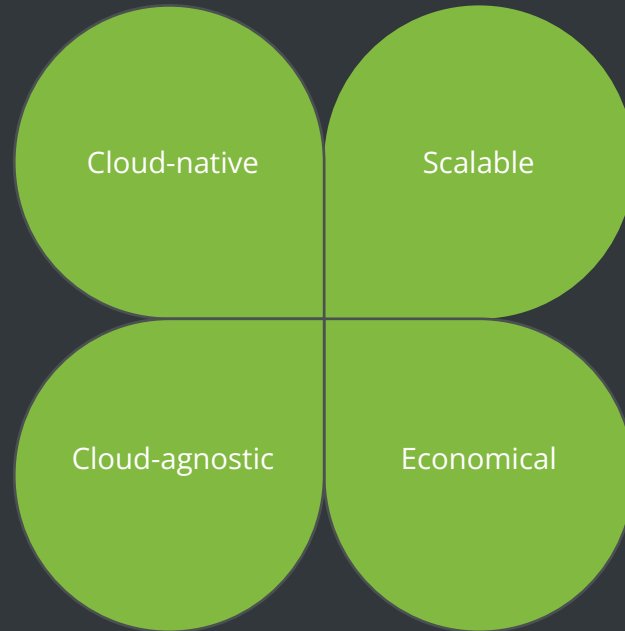
- Self-service analytics / reporting by departments
- Transparency
- "Democratization" of data / analytics
- Automatic updating / daily updated figures

Agenda

- Motivation
- Modern Data Stack
- Snowflake as Data Platform
- DEMO
- Lessons Learned
- Discussion

Modern Data Stack

Modern Data Stack is an ecosystem of open source as well as proprietary tools for data analytics purposes in the cloud.



Modern Data Architecture

Project Architecture



Data Catalogue

Data Lineage

Data Testing

Orchestration

Governance & Orchestration Layer

BI Tools Data Science Workbench AI-Deployment

Consumption Layer

Data Warehouse

Data Lake

Storage & Compute Layer

Integration Tool

Ingestion Layer

Datenbank Rest API Dateien ...

Source Layer



Data Flow

Agenda


- Motivation
- Modern Data Stack
- Snowflake as Data Platform
- DEMO
- Lessons Learned
- Discussion

Snowflake: Quick Facts

Cloud Data Warehouse



- Modern cloud-native data warehouse
- AWS, Azure, or Google Cloud
- Widespread and growing adoption
- Fully managed





Why Snowflake?

Best fit for the project

- Performance & scalability
 - Separation of storage and compute
 - Elastic scale → compute can be scaled up and down automatically
- Cost management
 - Pay as you go
 - Zero-copy clone
- Security & governance
 - End-to-end data encryption in transit and at rest
 - Snowflake Editions (Standard, Enterprise, Business Critical, Virtual Private Snowflake)
- Integrations to many common tools
 - BI and ML → Optimized for analytical workloads
 - dbt & Dagster

Agenda

- Motivation
- Modern Data Stack
- Snowflake as Data Platform
- **DEMO**
- Lessons Learned
- Discussion

Agenda

- Motivation
- Modern Data Stack
- Snowflake as Data Platform
- DEMO
- Lessons Learned
- Discussion

Best Practices

- Infrastructure as Code (IaC)
- Version control
- Automation (CI/CD, data tests, etc.)
- Code reviews
- Isolated environments
- Avoid vendor lock-in by using open technologies where possible

Lessons Learned

- Scalability of Snowflake allows building up the data platform step by step
- Positive experience orchestrating end-to-end with Dagster (especially dbt & Snowflake integration)
- Business requirements don't have to be boring! Using design-thinking methods, we have made the workshops at the customer's side fun.
- Management believes in our performance only if they see the dashboards.
 - Keep all stakeholders involved
 - Release often and in small iterations
- Snowflake is great! But tool selection is not everything → an established development process is also very important → we invite you to work with esentri 😊

OUR DATA & AI TEAM




NICOLAS HAUBNER

Senior Data Engineer

+49 175 5878815

nicolas.haubner@esentri.com

[LinkedIn](#)



FILIP STEPNIAK

Senior Data Engineer

+49 171 1975167

filip.stepniak@esentri.com

[LinkedIn](#)

Agenda

- Motivation
- Modern Data Stack
- Snowflake as Data Platform
- DEMO
- Lessons Learned / Next Steps
- Discussion