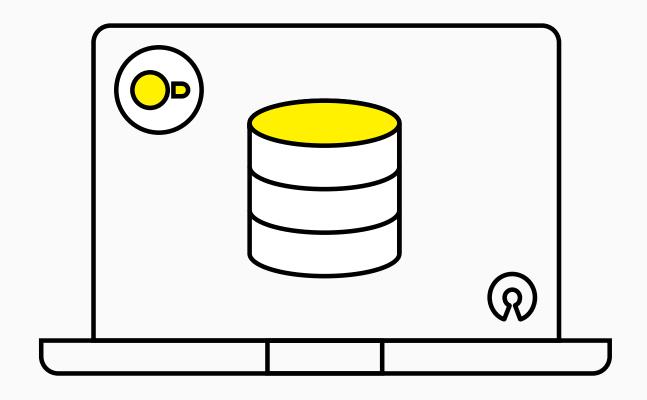


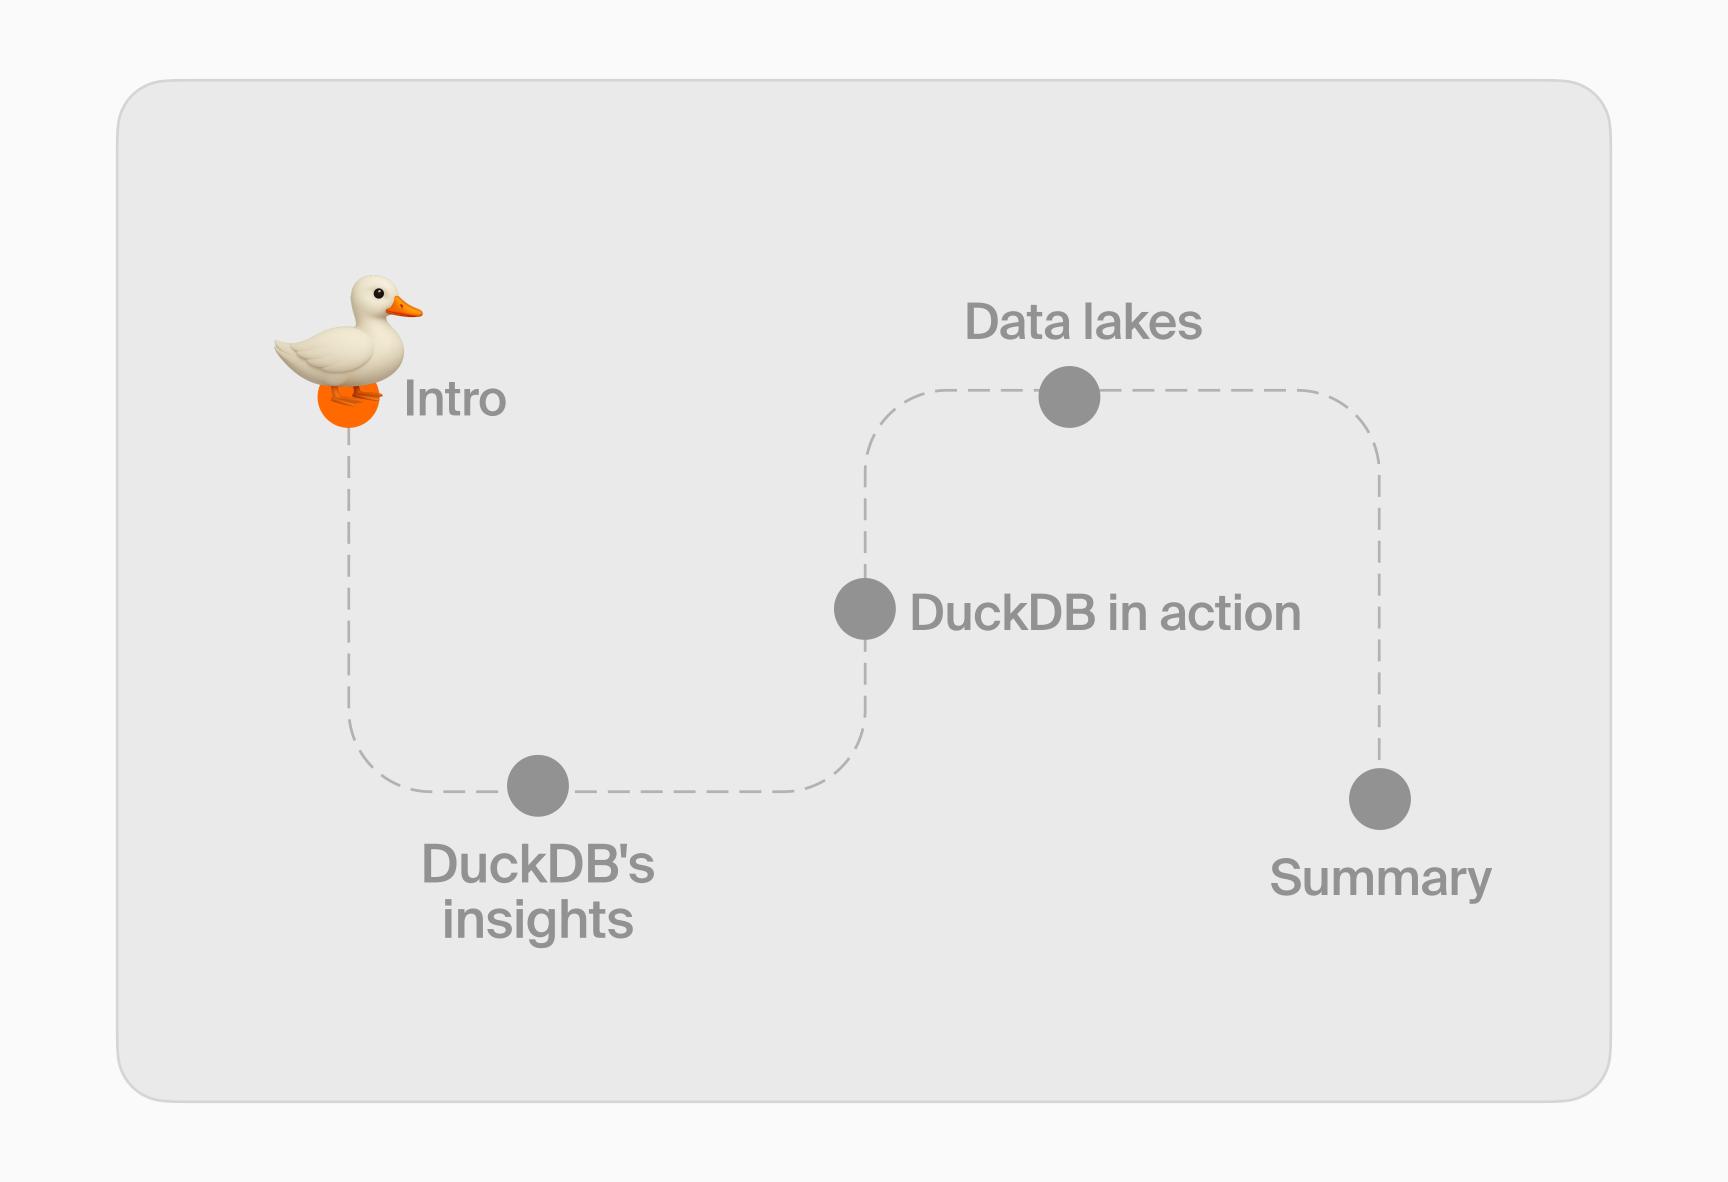
Big Data Unplugged

GÁBOR SZÁRNYAS

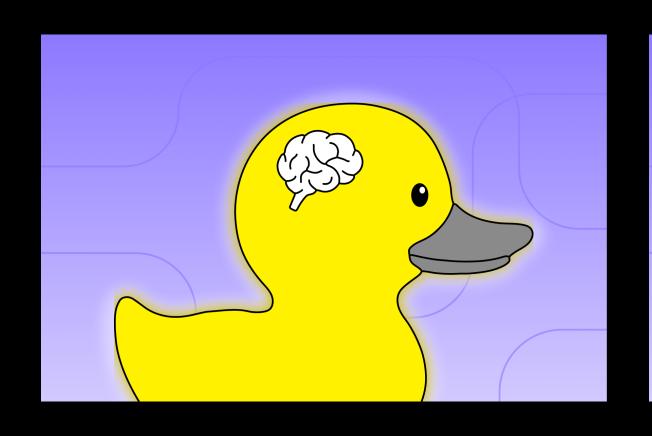


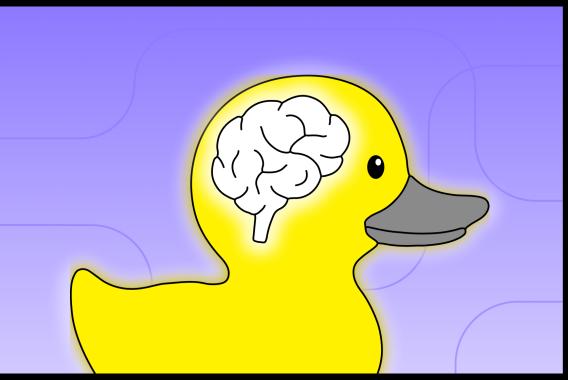


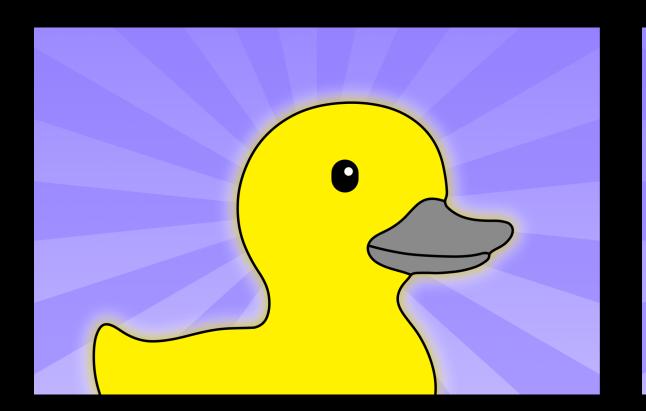
An open-source in-process SQL database for analytics

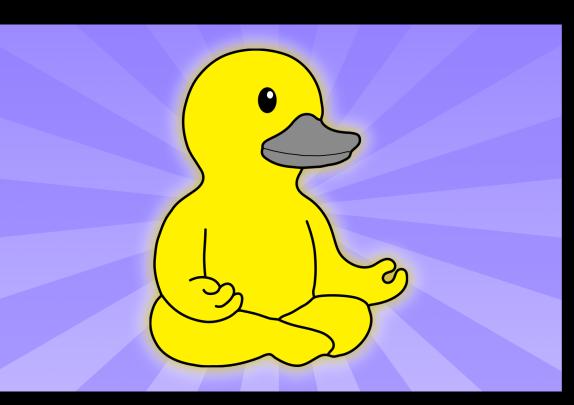


Data processing in increasing levels of enlightenment

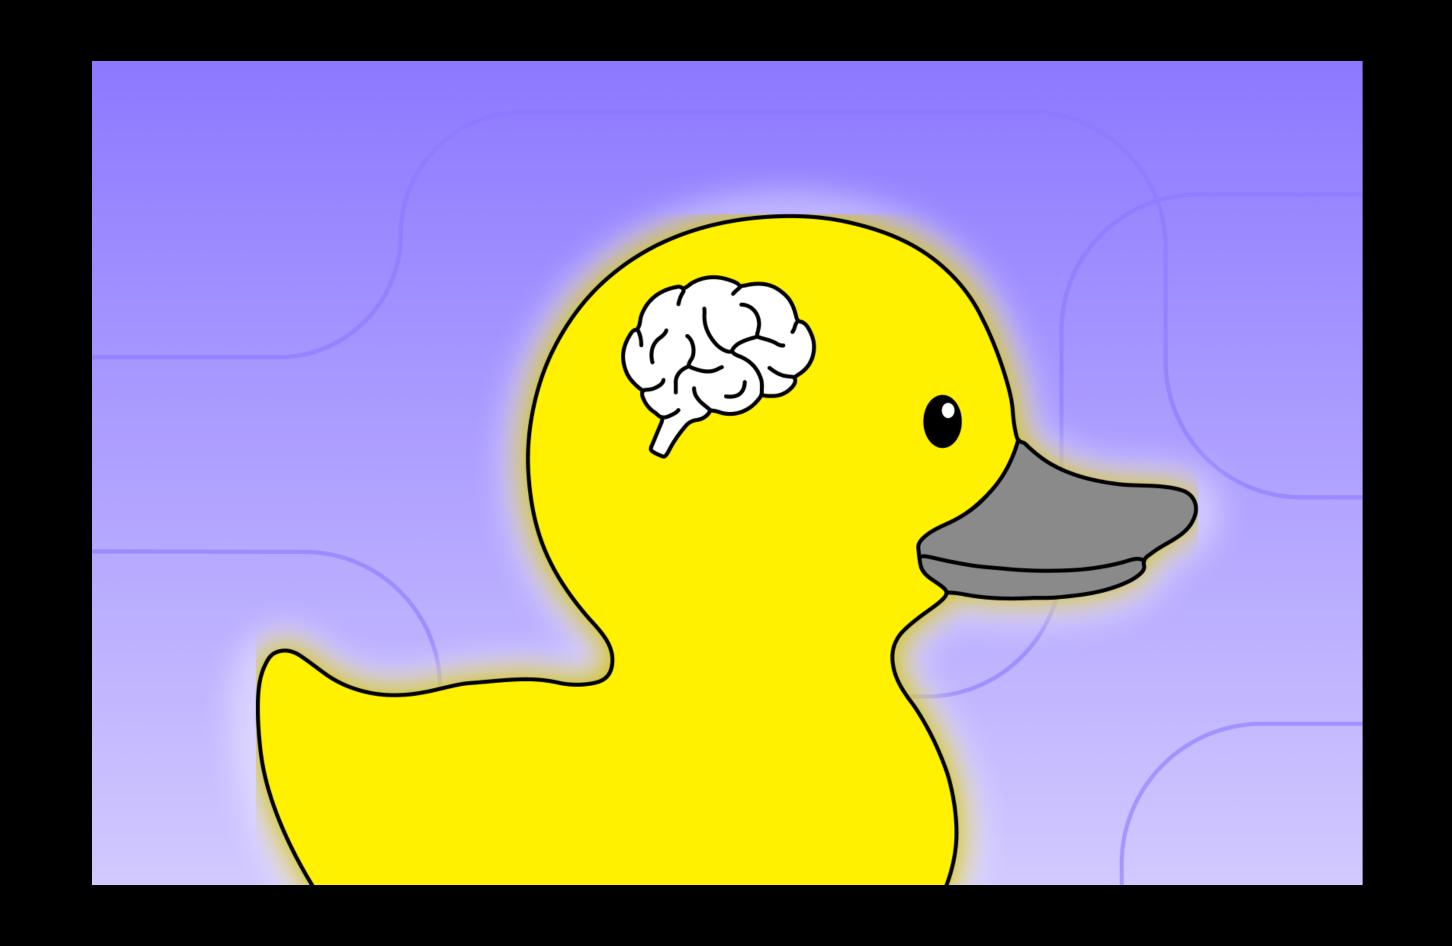






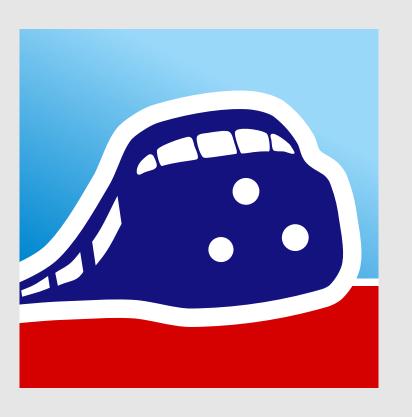


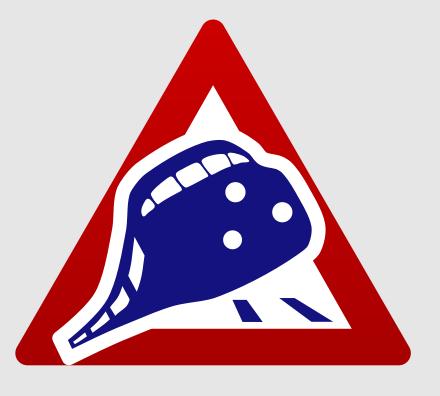
Text files



Dutch railway dataset by "Rijden de Treinen"

Which train stations had the worst delays over the weekends of February 2025?





```
$ wget https://blobs.duckdb.org/trains-2025-feb.csv.gz
$ gunzip trains-2025-feb.csv.gz
$ head -n 2 trains-2025-feb.csv
id,date,service_type,company,train_number,...
15277134,2025-02-01,Intercity,NS,1410,...
$ db_client -u admin -p admin
# CREATE TABLE services (
      id BIGINT, date DATE, service_type VARCHAR, ...
  );
# COPY services
  FROM 'trains-2025-feb.csv'
  (HEADER true, DELIMITER ',');
```

```
# SELECT
    station,
    avg(delay)
FROM services
WHERE strftime('%a', date) IN ('Sat', 'Sun')
GROUP BY station
ORDER BY #2 DESC
LIMIT 3;
```

station	avg(delay)
Siegburg/Bonn	5.58
Emmerich-Elten	3.97
Brussel-Zuid	3.86

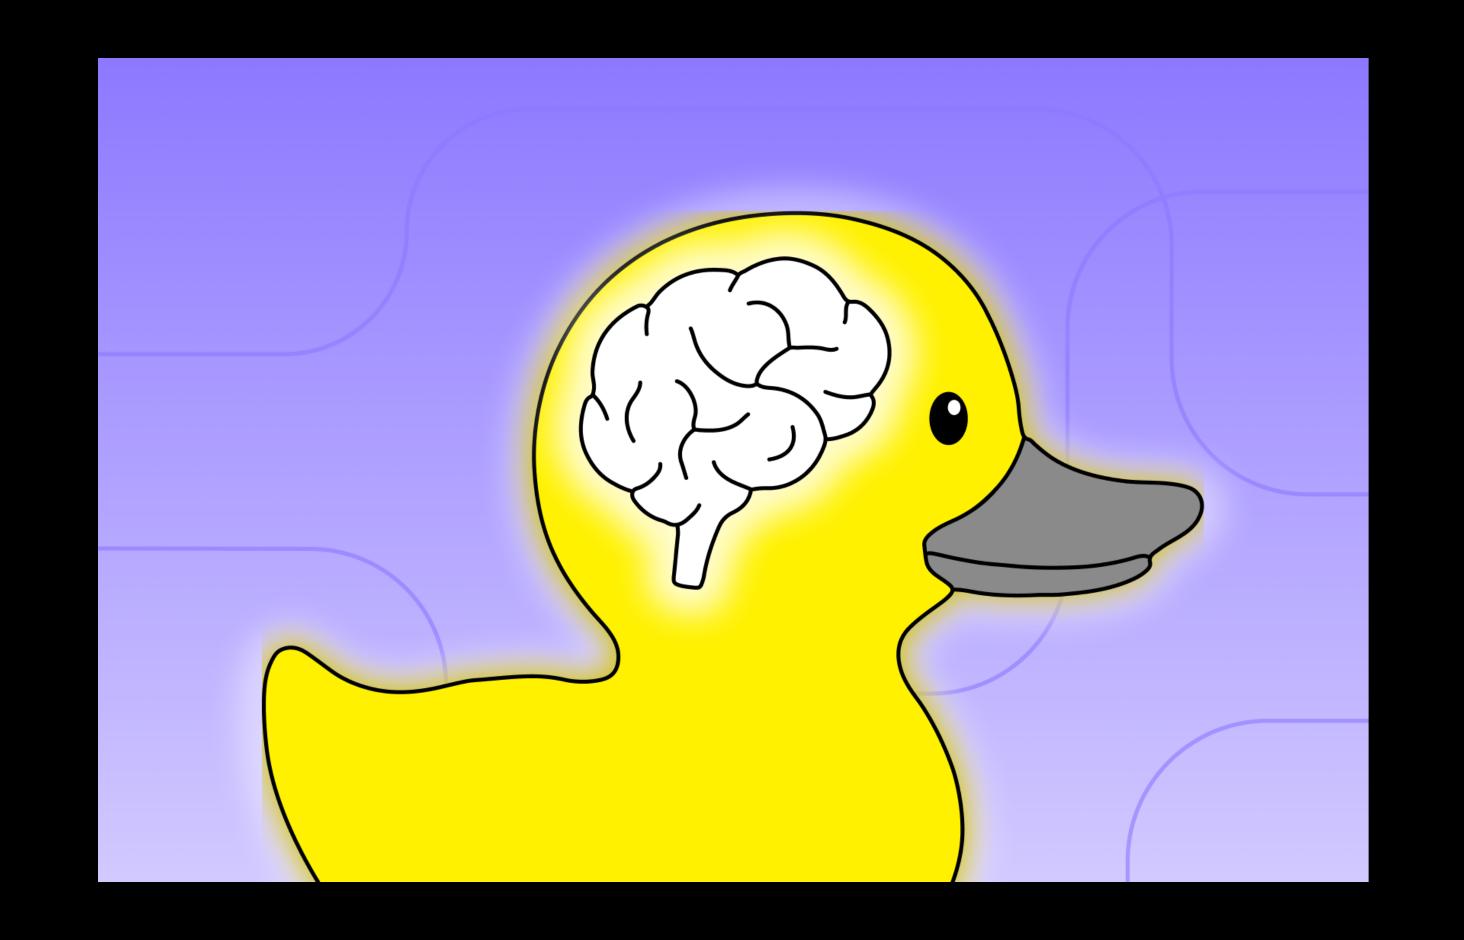
```
$ wget https://blobs.duckdb.org/trains-2025-feb.csv.gz
$ gunzip trains-2025-feb.csv.gz
$ head -n 2 trains-2025-feb.csv
$ db_client -u admin -p admin
# CREATE TABLE services (id BIGINT, date DATE,
service_type VARCHAR, company VARCHAR, train_number
BIGINT, cancelled BOOLEAN, partly_cancelled BOOLEAN,
maximum_delay BIGINT, stop_id BIGINT, station_code
VARCHAR, station VARCHAR, arrival_time TIMESTAMP WITH
TIME ZONE, delay BIGINT, arrival_cancelled BOOLEAN,
departure_time TIMESTAMP WITH TIME ZONE, departure_delay
BIGINT, departure_cancelled BOOLEAN);
# COPY services FROM 'trains-2025-feb.csv' (HEADER true,
DELIMITER ',');
# SELECT station, avg(delay) FROM services WHERE
strftime('%a', date) IN ('Sat', 'Sun') GROUP BY station
ORDER BY #2 DESC LIMIT 3;
```

750+ keystrokes for a simple analysis

Legacy solution

Works fine in DuckDB!

Text files (deluxe)



```
$ duckdb
```

```
D SELECT
    station,
    avg(delay)
FROM 'https://blobs.duckdb.org/trains-2025-feb.csv.gz'
WHERE strftime('%a', date) IN ('Sat', 'Sun')
GROUP BY station
ORDER BY #2 DESC
LIMIT 3;
```

download

load

decompress

detect schema

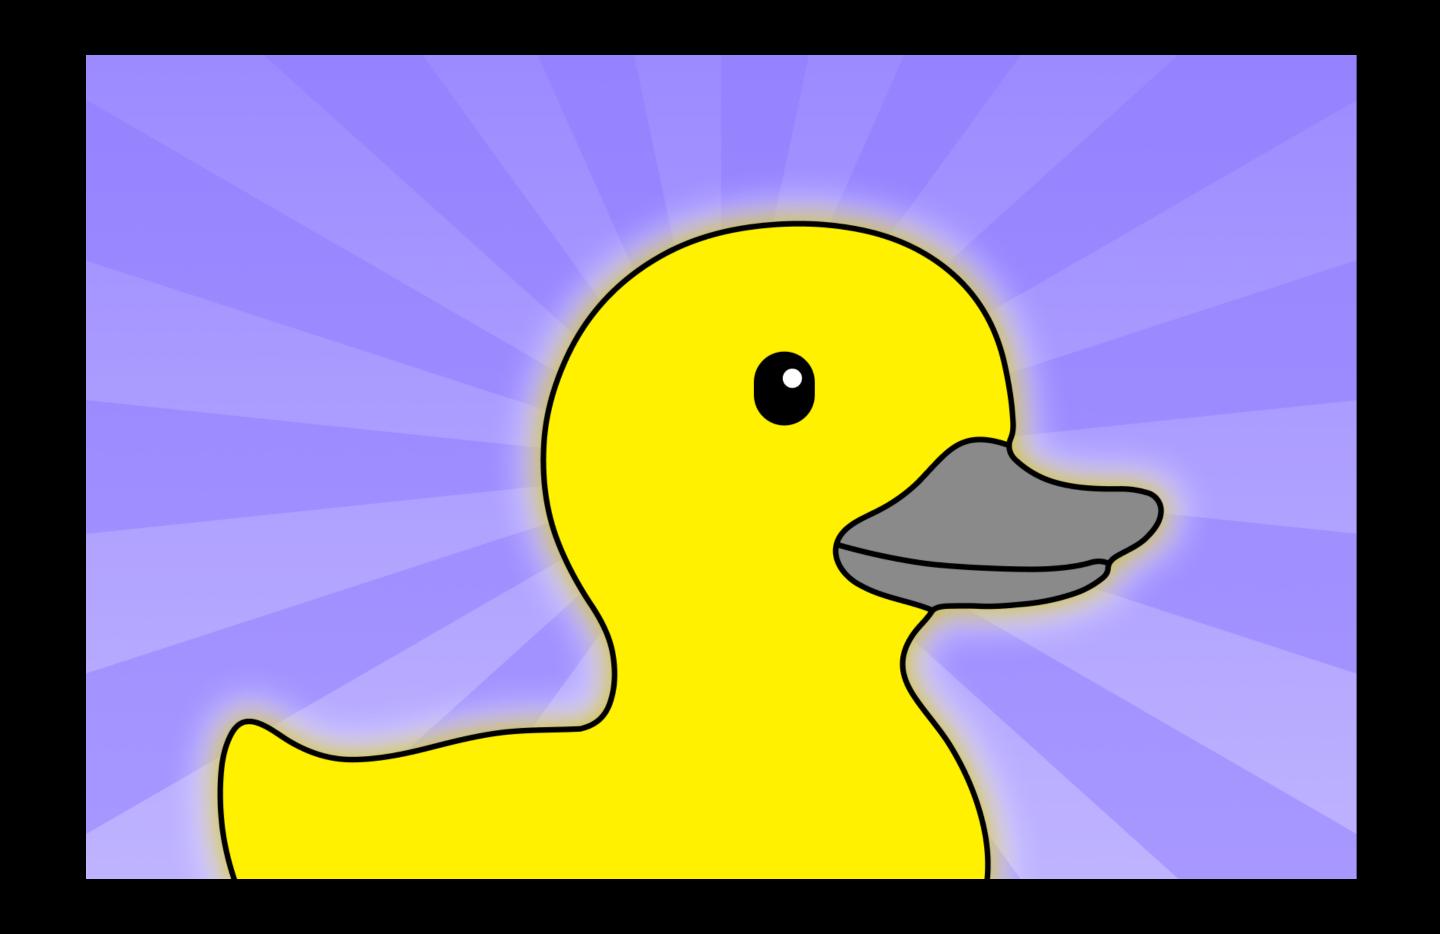
station	avg(delay)	
Siegburg/Bonn	5.58	
Emmerich-Elten	3.97	
Brussel-Zuid	3.86	

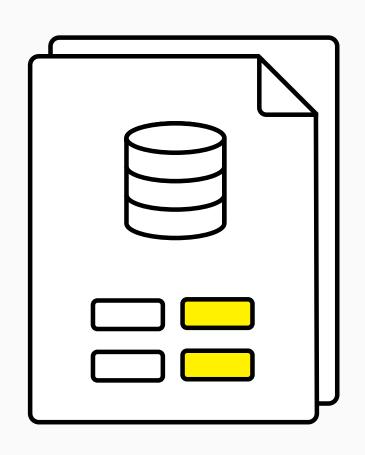
(4 seconds)

200 characters = 70% reduction

Modern SQL

Binary files





Data layout: Row-based or column-based?

Row-based storage

date	station	delay
Feb 15	Delft	0.0
Feb 15	Edam	0.0
Feb 15	Breda	0.0
Feb 15	Delft	1.6

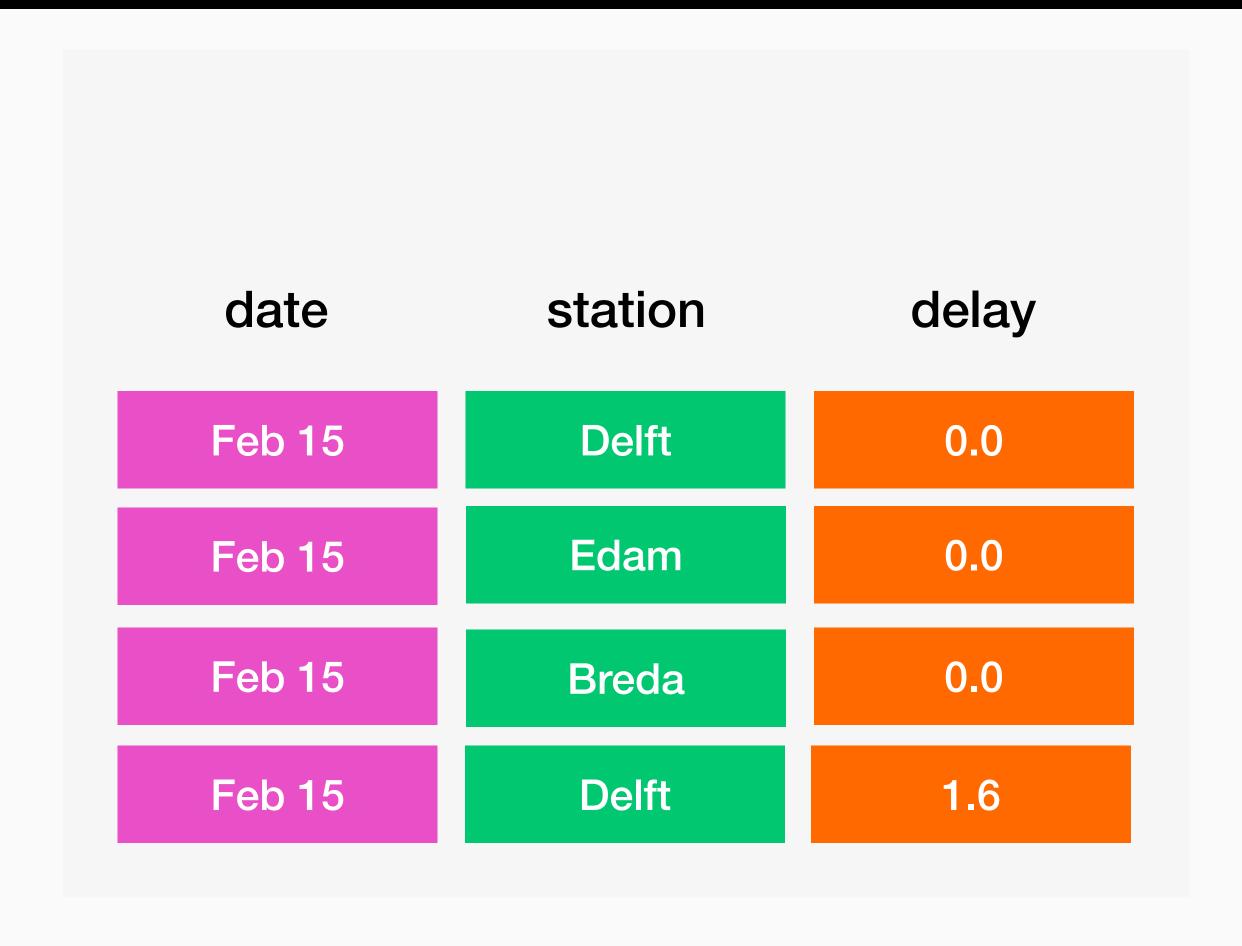
Row-based storage



Row-based storage



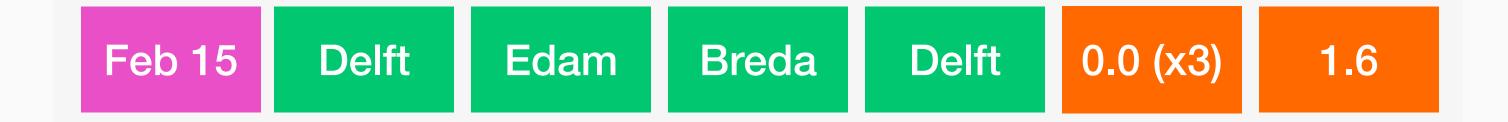
Column-based storage



Column-based storage



Column-based storage



```
$ duckdb
```

```
D SELECT
     station,
     avg(delay)
FROM 'https://blobs.duckdb.org/trains-2025-feb.parquet'
WHERE strftime('%a', date) IN ('Sat', 'Sun')
GROUP BY station
ORDER BY #2 DESC
LIMIT 3;
```

station	avg(delay)
Siegburg/Bonn	5.58
Emmerich-Elten	3.97
Brussel-Zuid	3.86

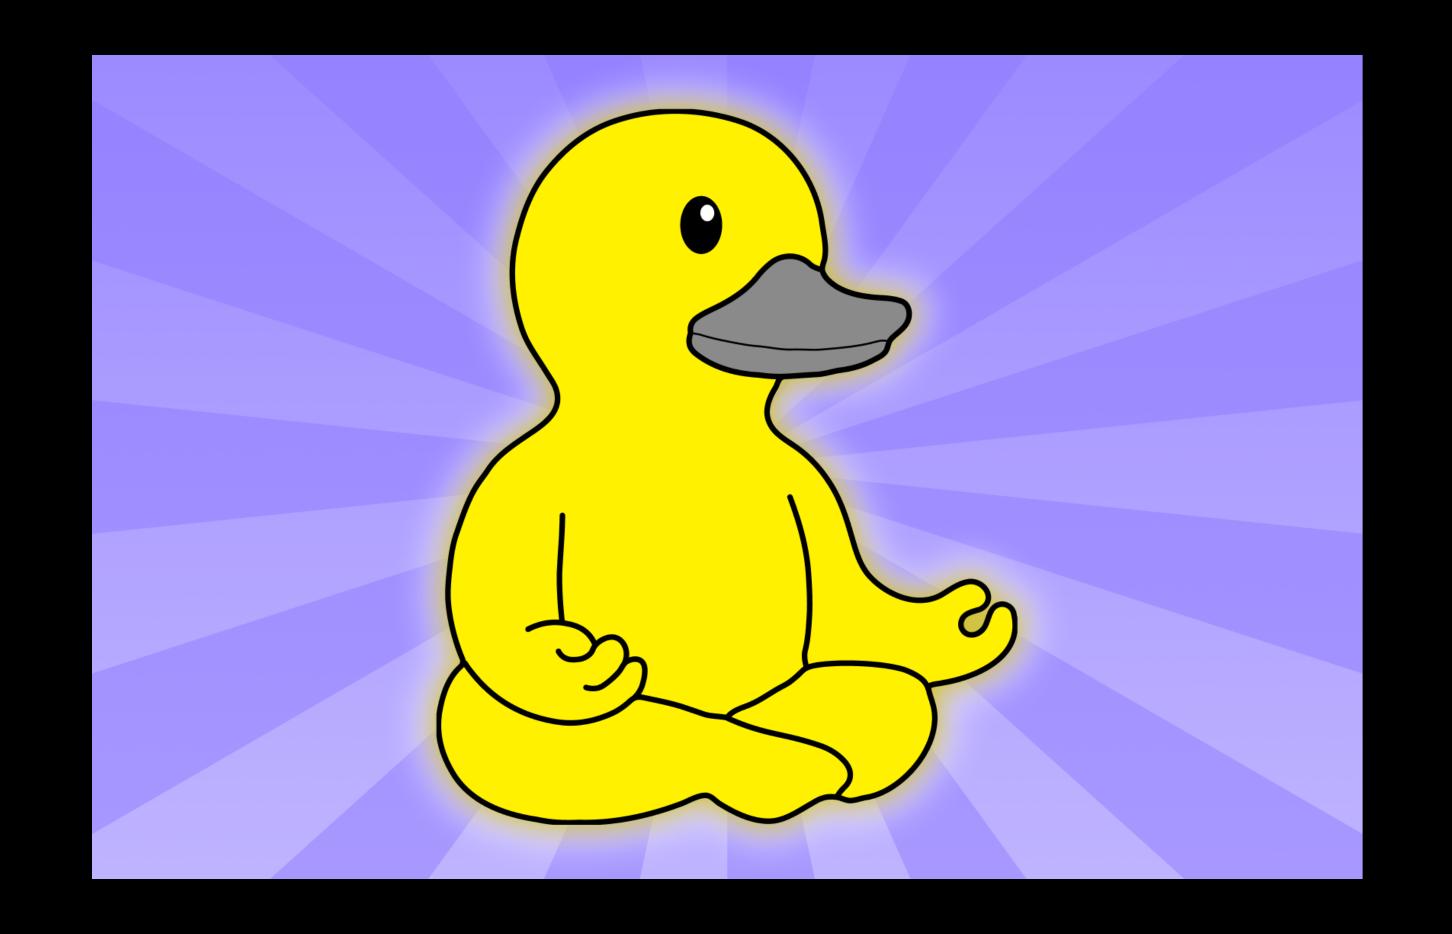
(1 second)

Concise

Safe

Fast

Database



```
$ duckdb train-services.db

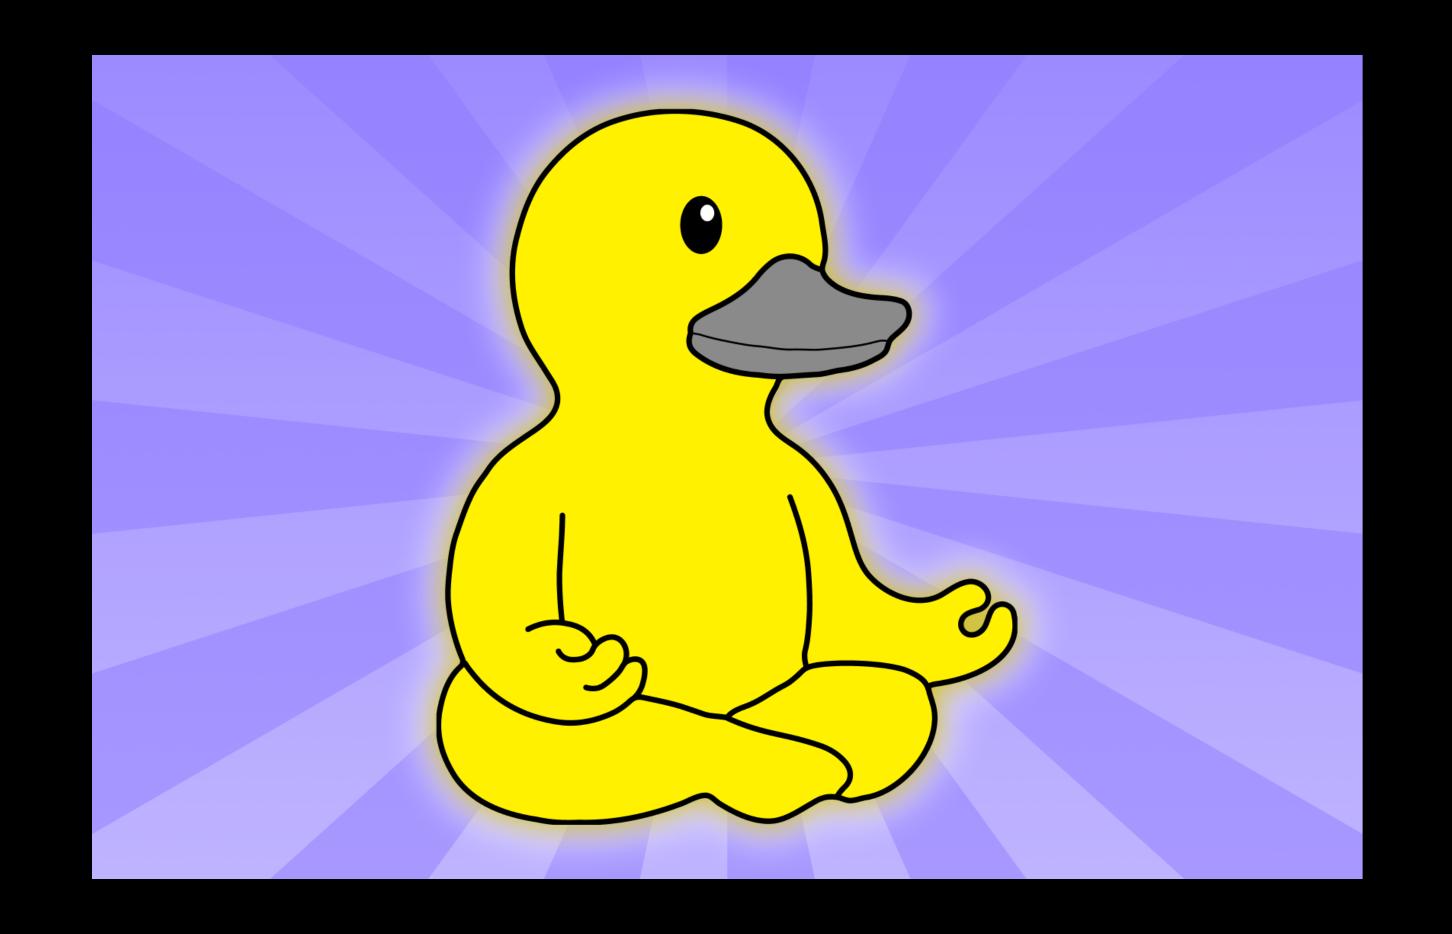
D SELECT
     station,
     avg(delay)
FROM services
WHERE strftime('%a', date) IN ('Sat', 'Sun')
GROUP BY station
ORDER BY #2 DESC
LIMIT 3;
```

station	avg(delay)
Siegburg/Bonn	5.58
Emmerich-Elten	3.97
Brussel-Zuid	3.86

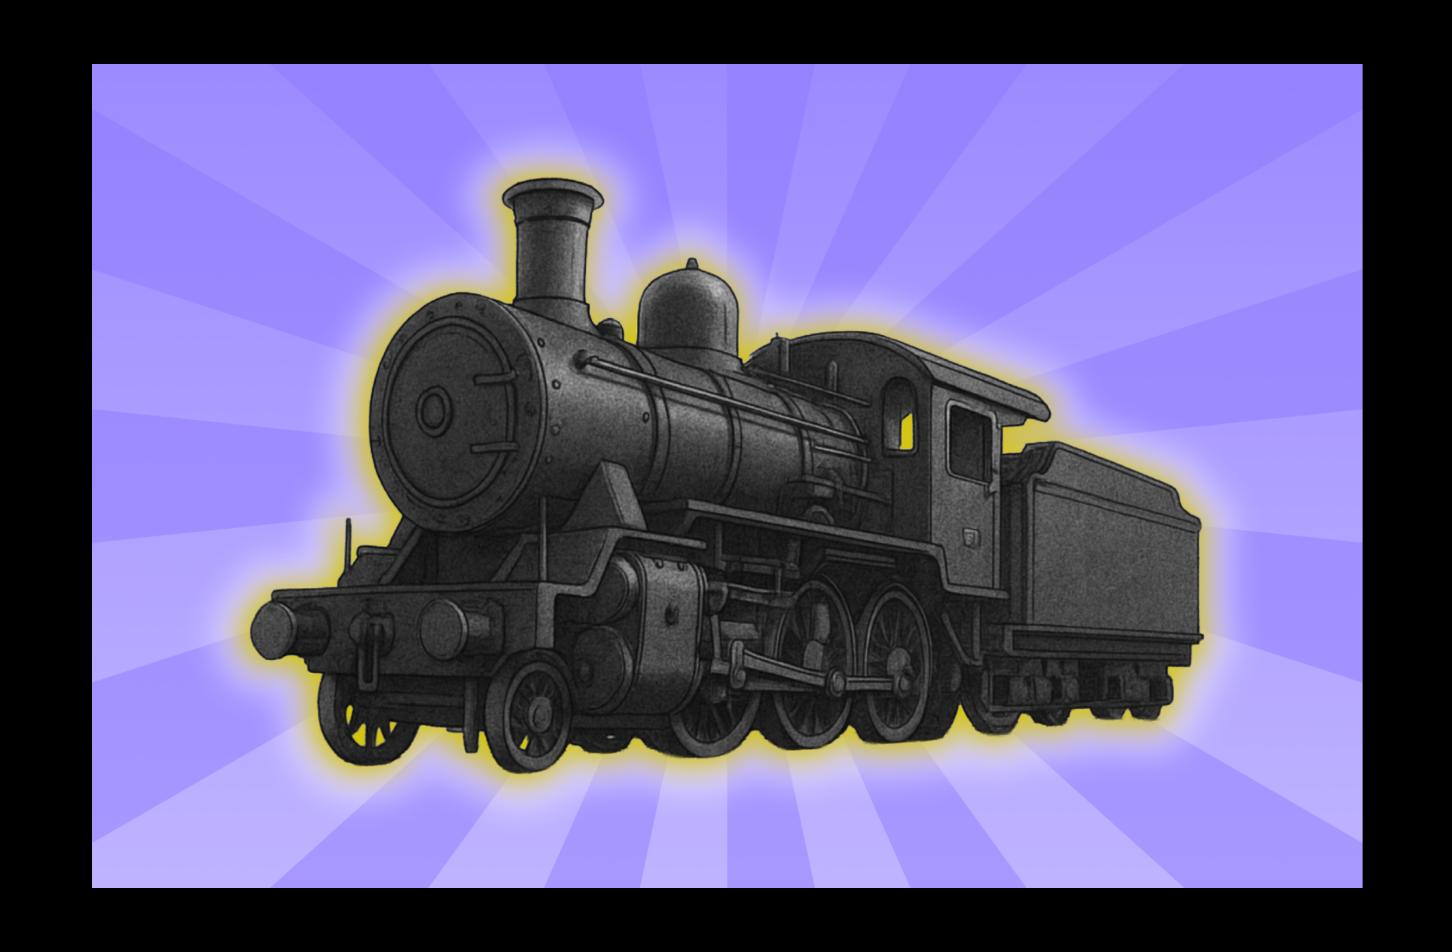
```
(0.01 seconds)
```

Super fast!

Database



Railway operators



The worst delay in Amsterdam Centraal is almost 3 hours!

```
SELECT date, train_number, delay
FROM services
WHERE date = '2025-02-27'
  AND station = 'Amsterdam Centraal'
  AND train_number = 420;
```

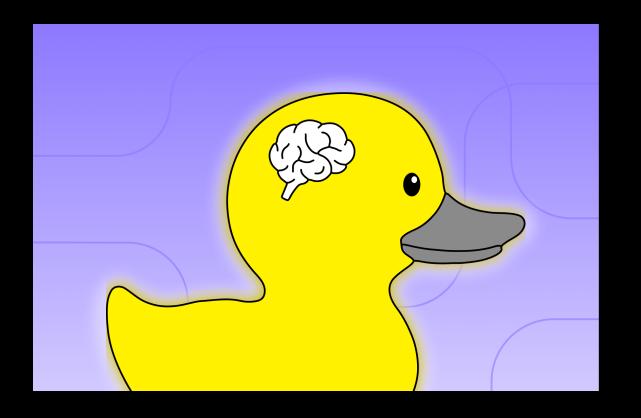
date	train_number	delay
2025-02-27	420	174

A fictional story from a railway operator

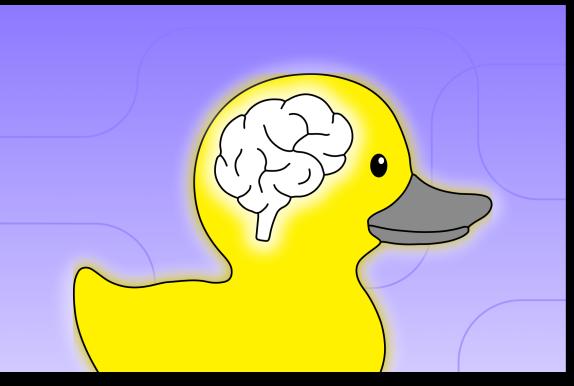
Change the train's number!

```
UPDATE services
SET train_number =
    train_number + (random() * 100 + 1)::INT
WHERE delay > 150;
(23 rows updated)
         train_number
                        delay
  date
            0 rows
```

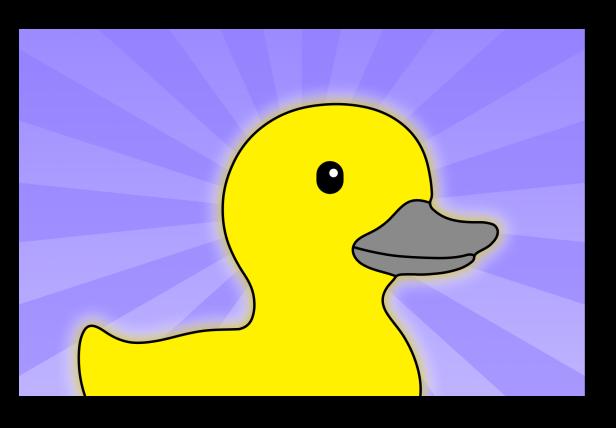




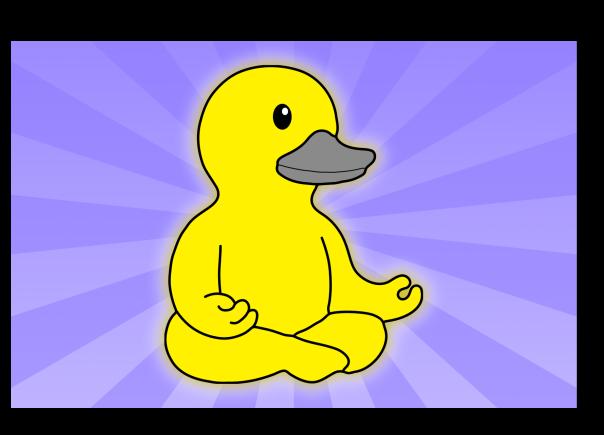
Text files



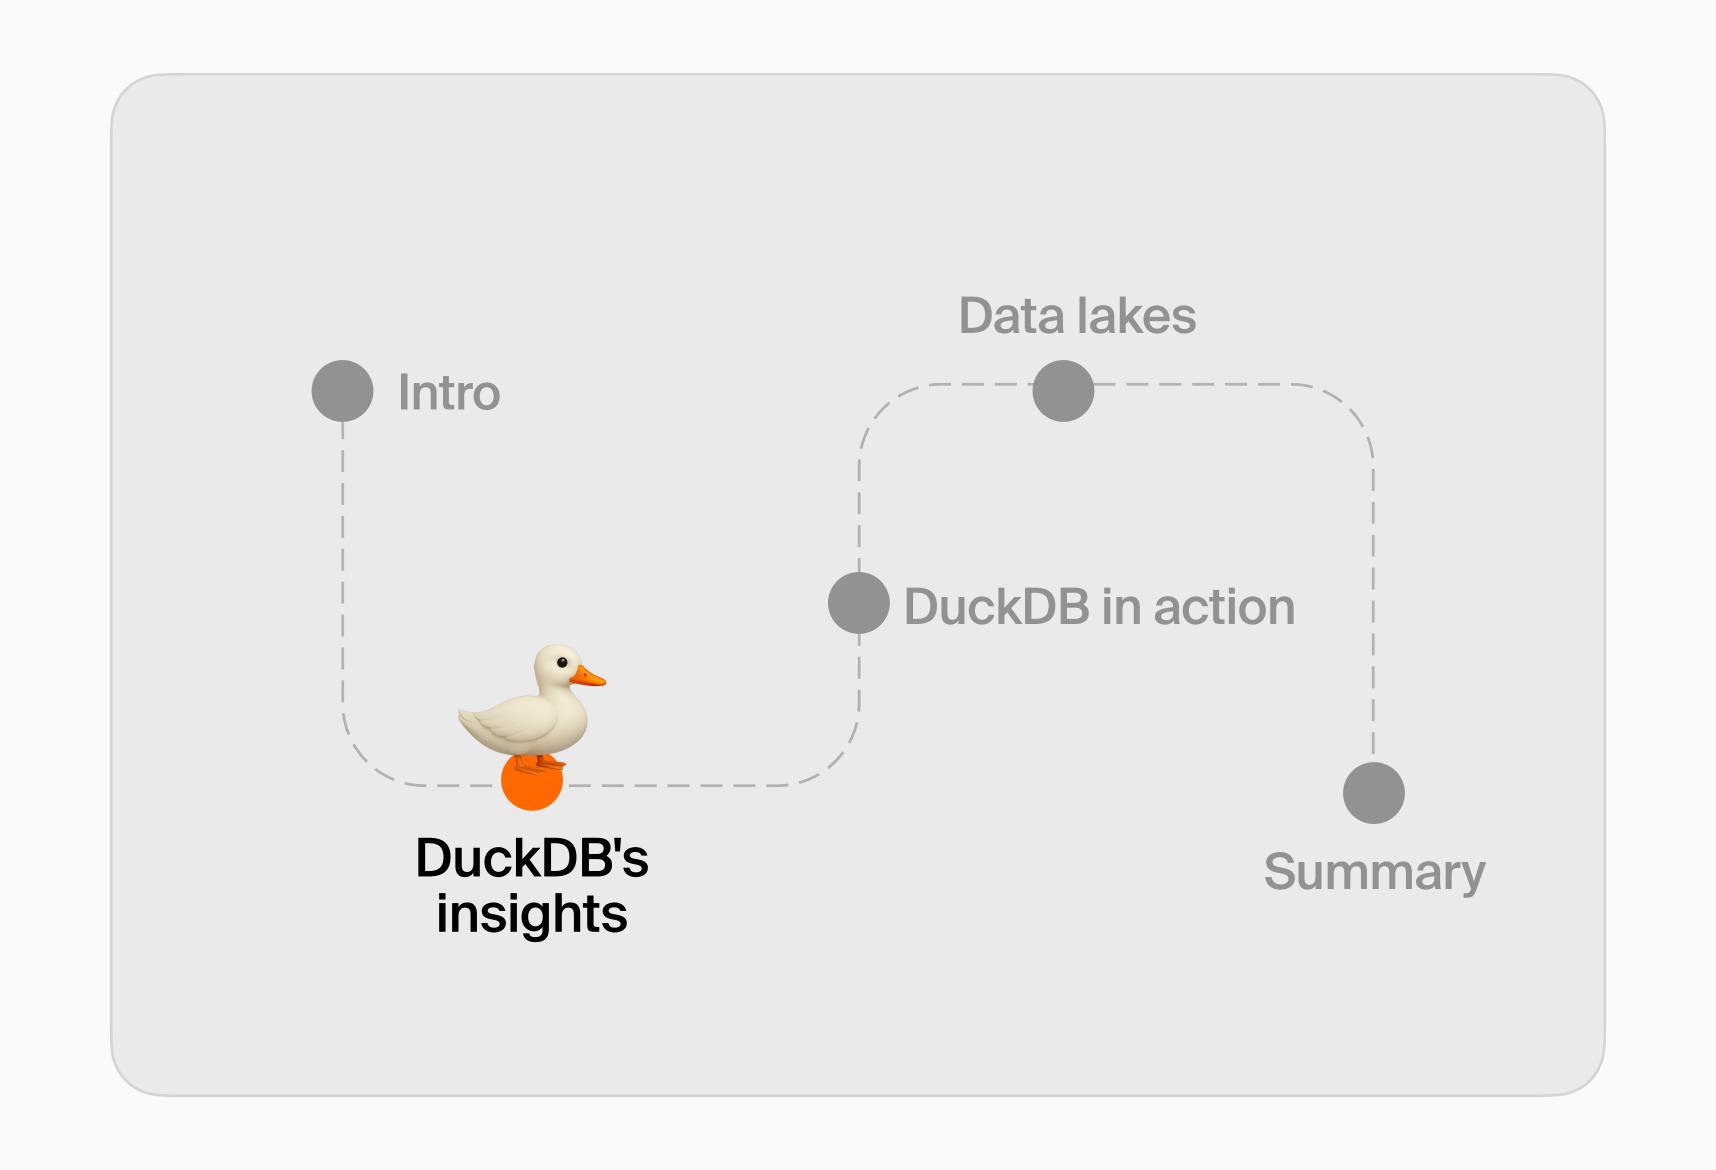
Text files (deluxe)

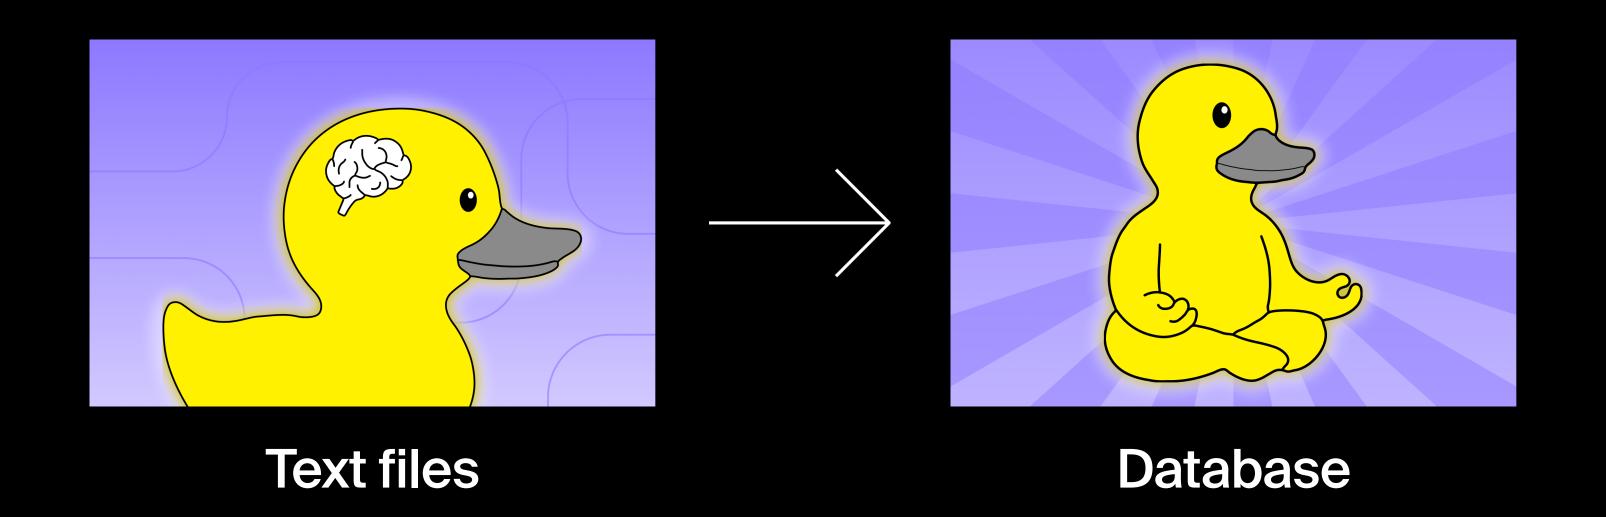


Binary files



Database



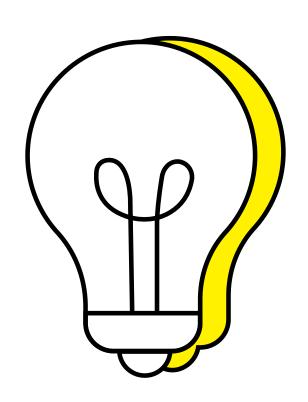


Concise queries

Safe

Fast

Updatable



INSIGHT #1

Databases can target data science workloads

Data science landscape, 2015





?

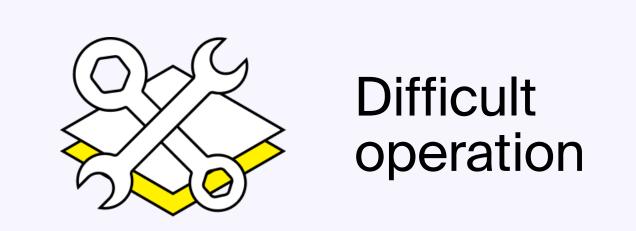
data.tables



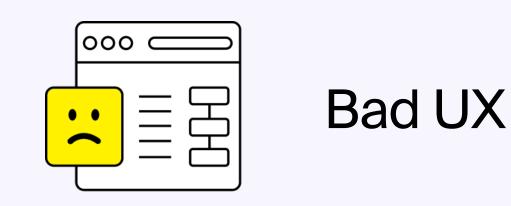
data size

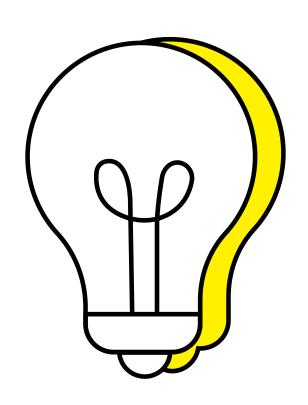
Resistance: Why would I use a database?

```
$ sudo apt install your_favorite_database
$ ... configure for 60 minutes ...
$ sudo service start your_favorite_database
$ wget https://blobs.duckdb.org/trains-2025-feb.csv.gz
$ gunzip trains-2025-feb.csv.gz
$ db_client -u admin -p admin
# CREATE TABLE ...; COPY ...; SELECT ...;
```









INSIGHT #2

Focus on usability

End-to-end performance

setup

define schema

load

write queries

run queries

End-to-end performance

setup

define schema

write queries

load

run queries

End-to-end performance

setup

define schema

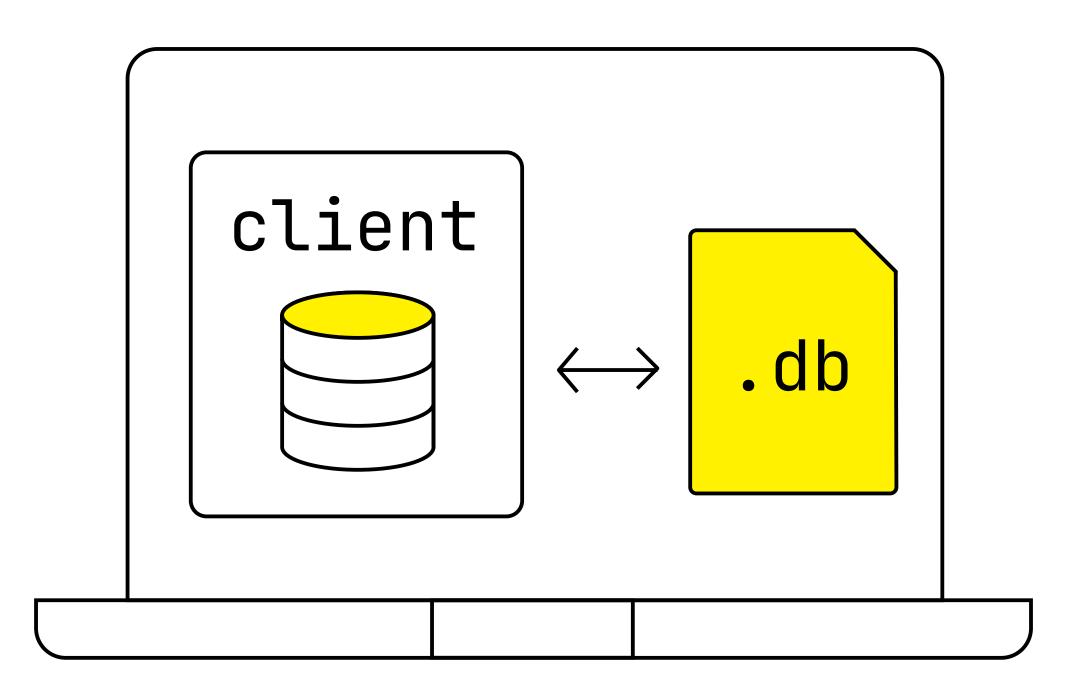
load

write queries

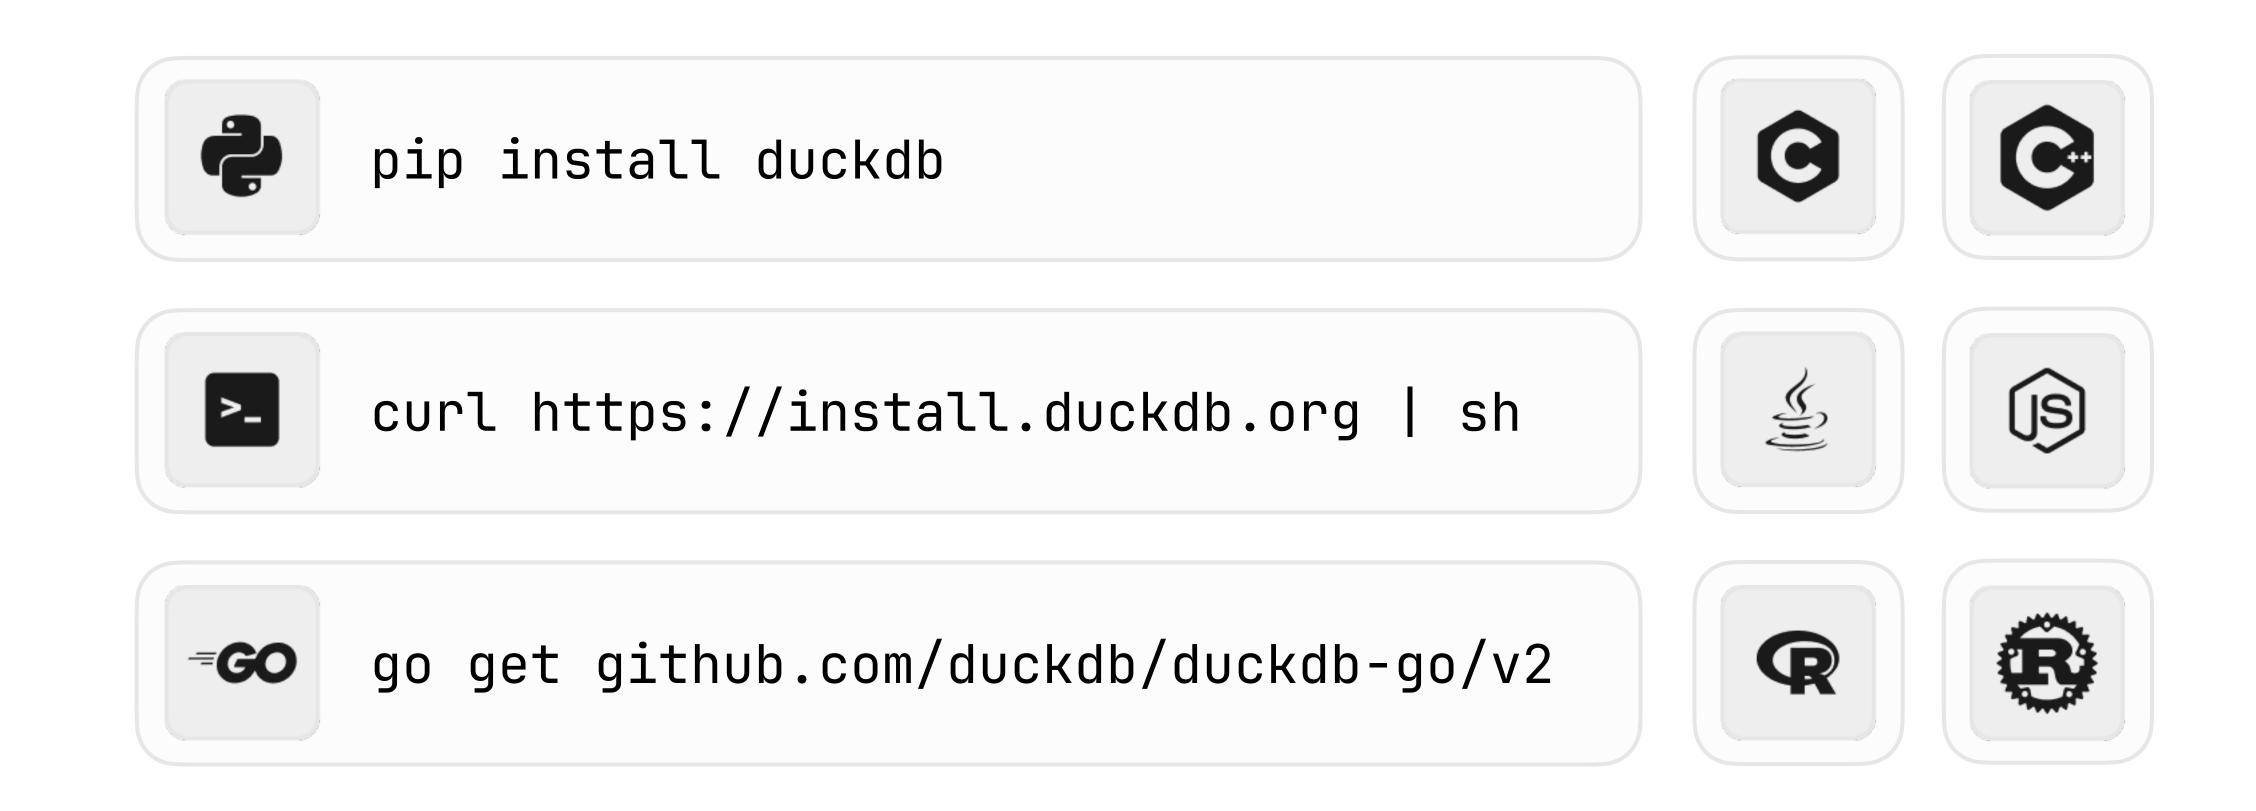
run queries

In-process architecture

setup



In-process architecture



End-to-end performance

setup

define schema

load

write queries

run queries

Friendly SQL

define schema

write queries

```
SELECT
    date, station, avg(delay),
-- min(delay)
FROM 'https://blobs.duckdb.org/trains-2025-feb.csv.gz'
GROUP BY ALL;
```



||||· ClickHouse





2022

2022

2023

2023



2024

amazon REDSHIFT



PostgreSQL

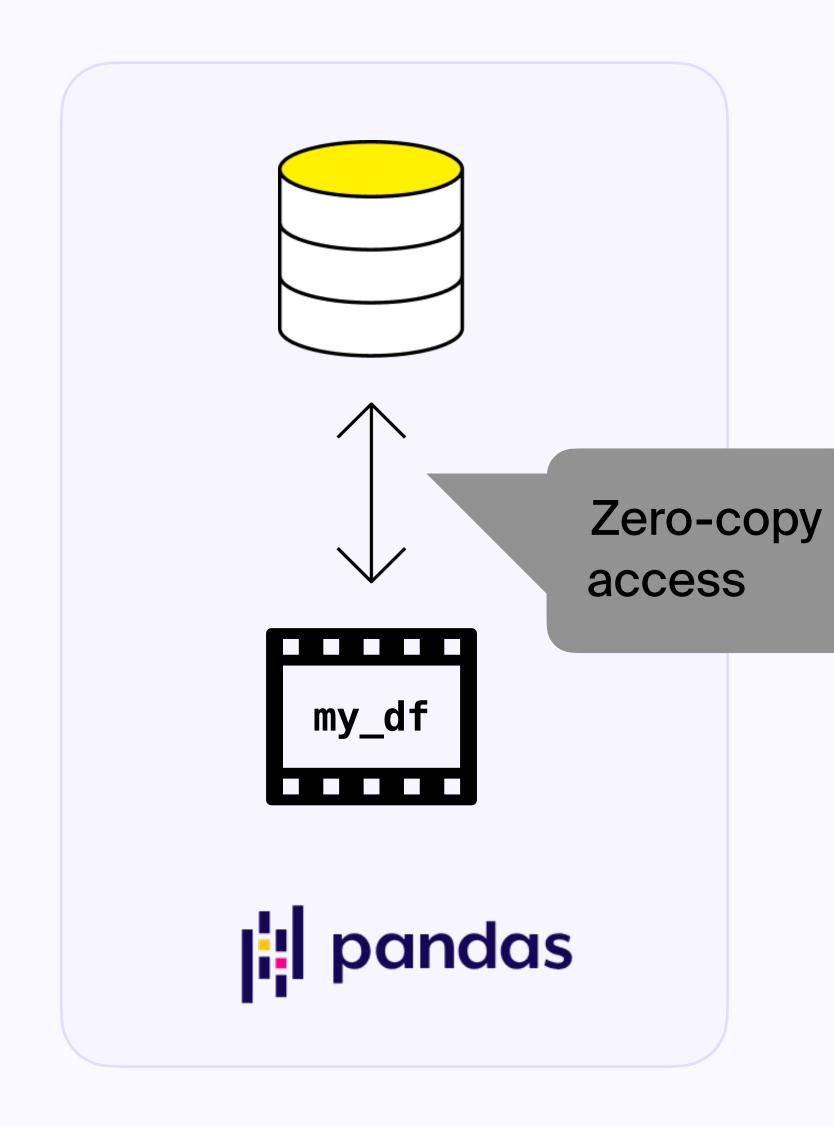
2025

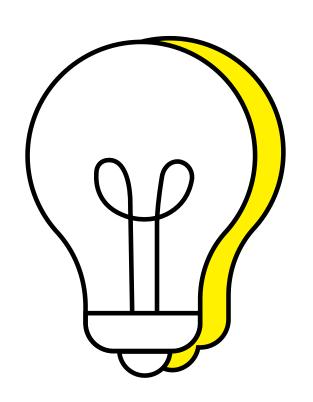
2025

2026

Pandas integration

```
import duckdb
import pandas as pd
my_df = pd.DataFrame \
          .from_dict({'a': [42, 43]})
res = duckdb \
        sql("SELECT avg(a) FROM my_df")
res.df()
                          Replacement scan
# 0 42.5
```

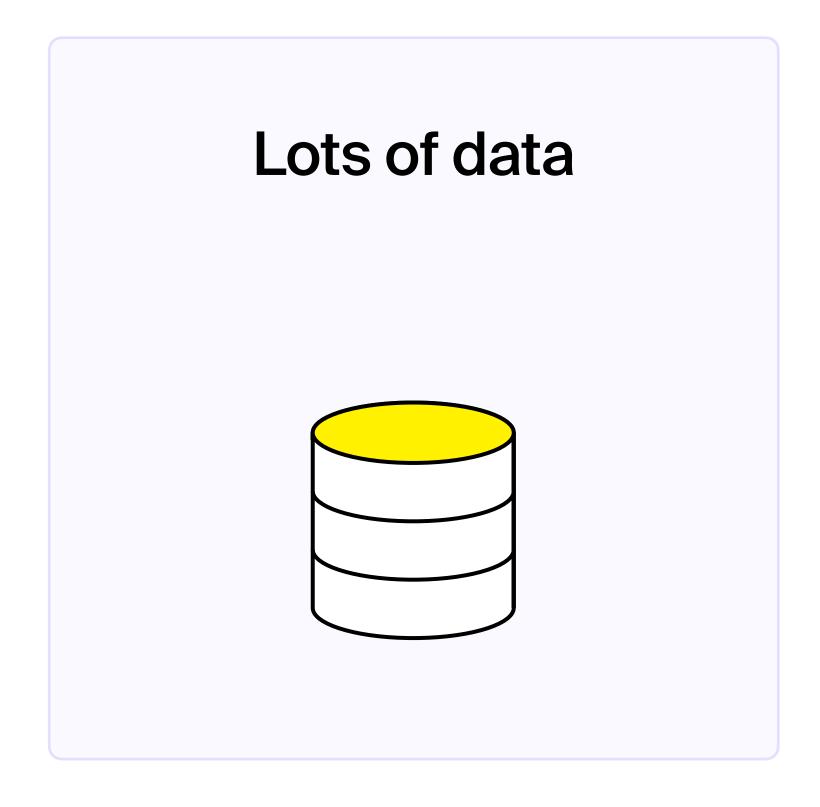


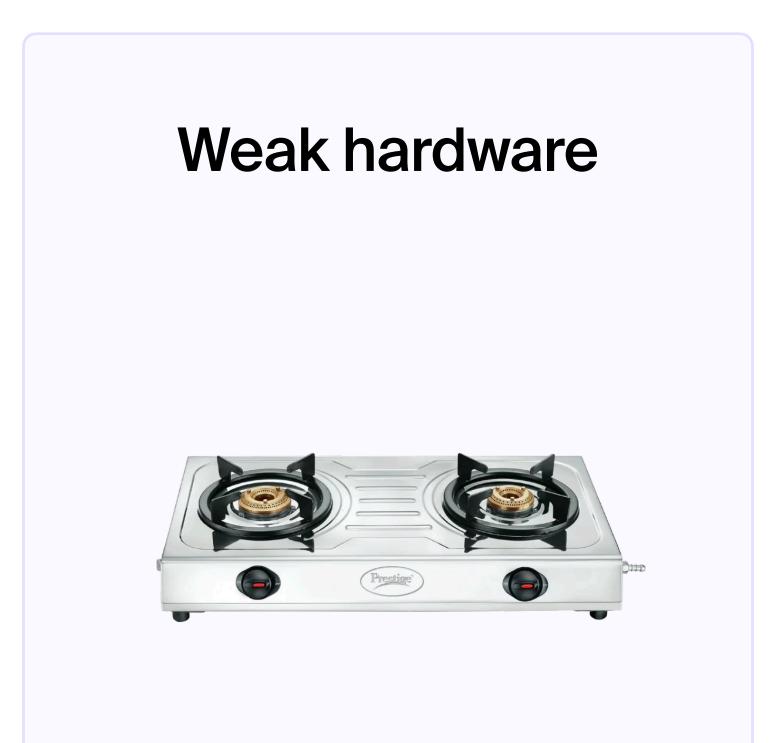


INSIGHT #3

"Big Data" is overrated

Data processing landscape

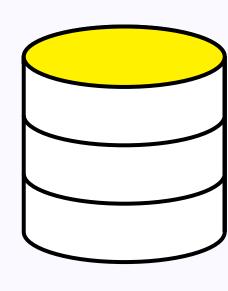




Data processing landscape

Lots of data

but most queries only access a small amount

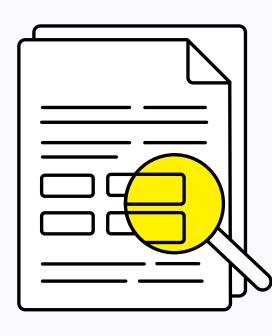


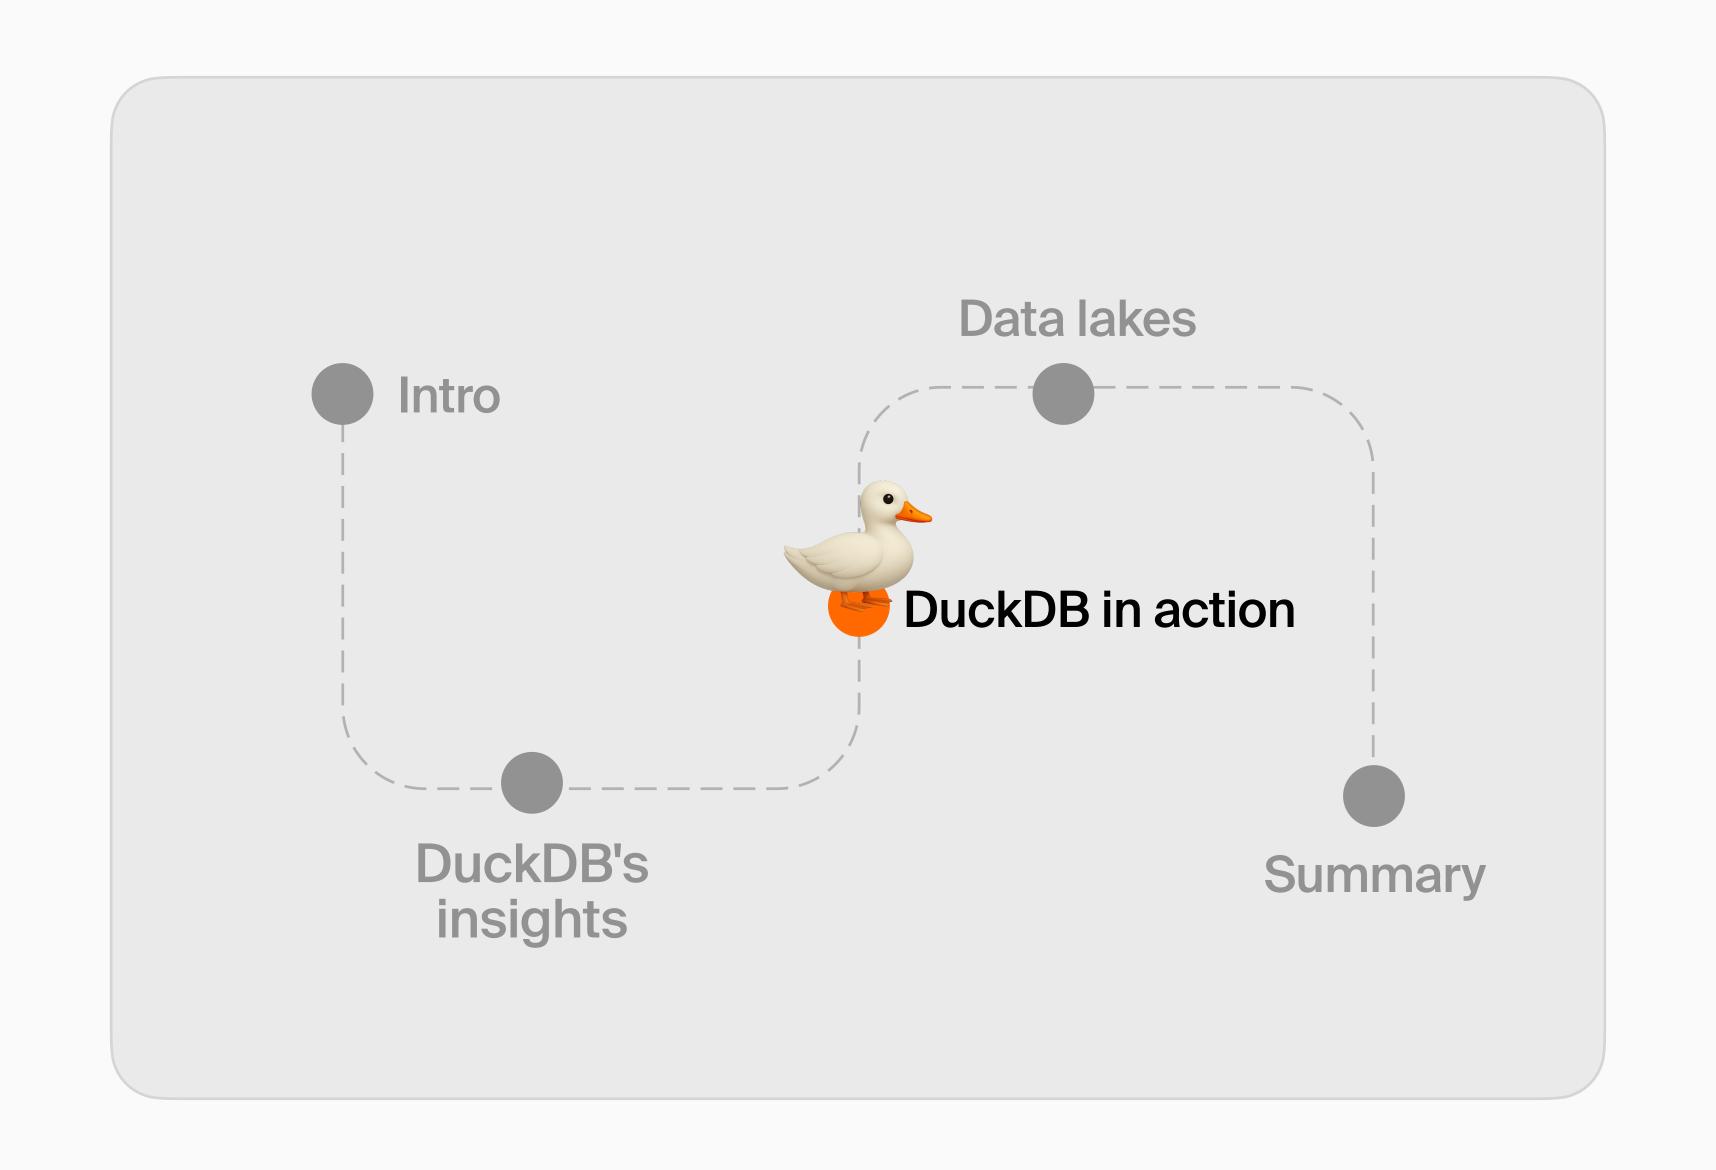
Strong hardware

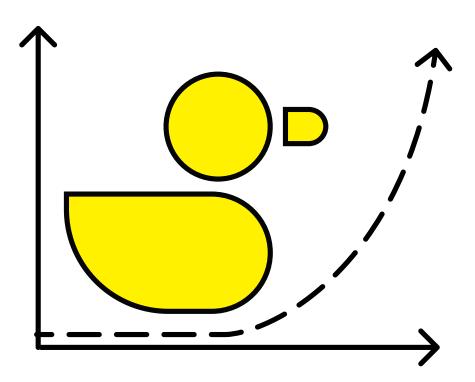


New techniques

columnar storage, lightweight compression, vectorized query execution







Performance

TPC-H benchmark





DVIDIA.



ORACLE

ieit



6 浪潮云

Red Hat





TRANSWARP













vmware



Home About the TPC▼

Benchmarks/Results ▼

Downloads ▼

TPCTC

Miscellaneous ▼

Search

timecho i

Newsletter

Member Login

1,000 GB Results

Rank	Company	System	QphH	Price/kQphH	Watts/KQphH	System Availability	Database	Operating System	Date Submitted
1	Hewlett Packard Enterprise	HPE ProLiant DL345 Gen11 (with AMD EPYC 9375F)	1,531,016	192.37 USD	NR	10/01/25	Microsoft SQL Server 2022 Enterprise Edition 64 bit	Microsoft Windows Server 2022 Standard Edition	10/01/25
2	Hewlett Packard Enterprise	HPE Proliant DL345 Gen11 (with AMD EPYC 9374F Proc	1,257,628	233.57 USD	NR	07/29/25	Microsoft SQL Server 2022 Enterprise Edition 64 bit	Microsoft Windows Server 2022 Standard Edition	07/29/25
3	Hewlett Packard Enterprise	HPE ProLiant DL380 Gen12	1,184,211	263.13 USD	NR	06/02/25	Microsoft SQL Server 2022 Enterprise Edition 64 bit	Microsoft Windows Server Datacenter 2025 Edition	04/23/25

TPC-H unplugged







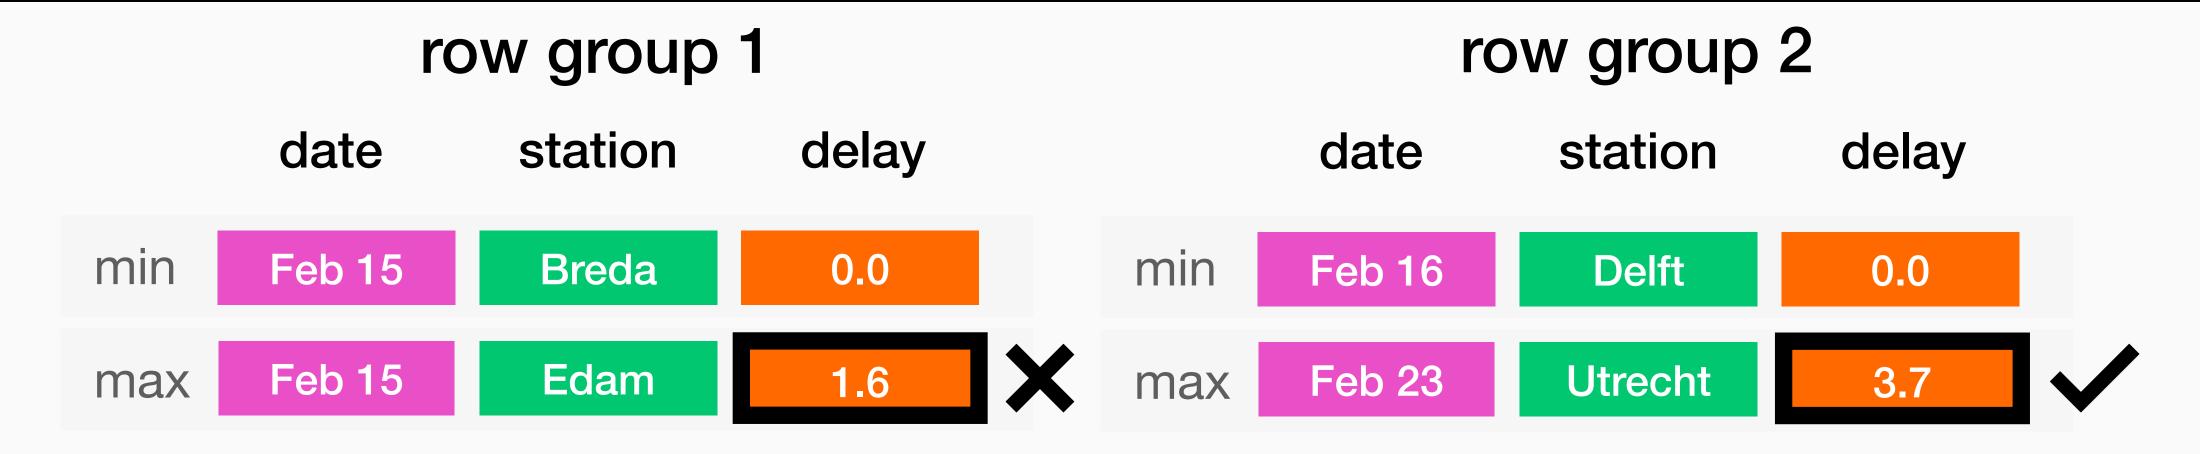
100 GB 1000 GB 10 000 GB

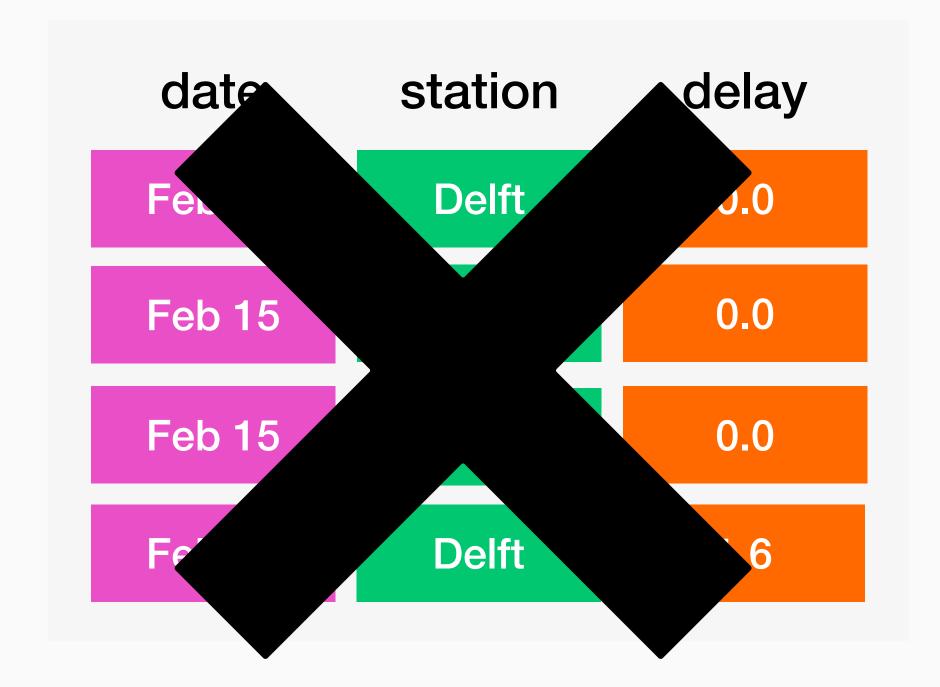


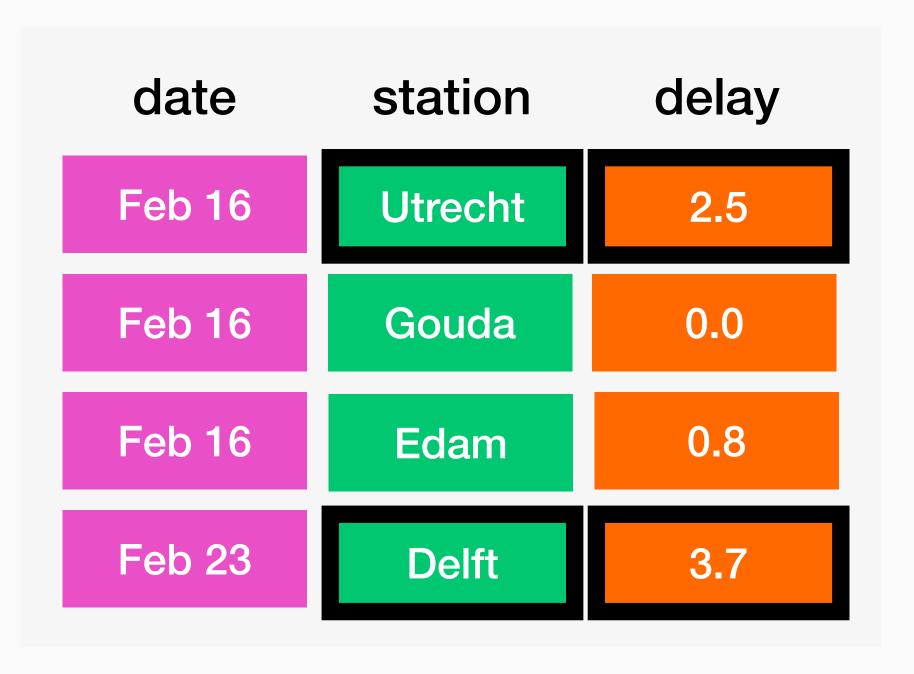
Lightweight indexing

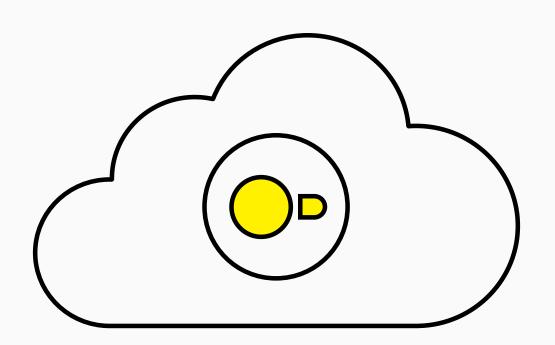


Which stations are affected by delays > 2 min?





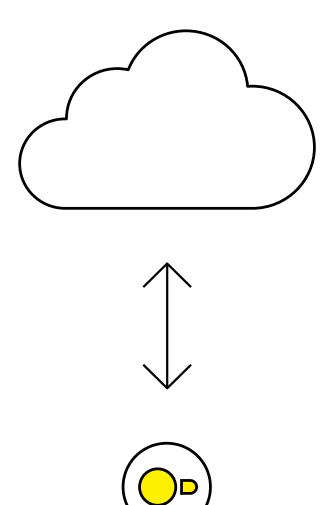




DuckDB and the cloud



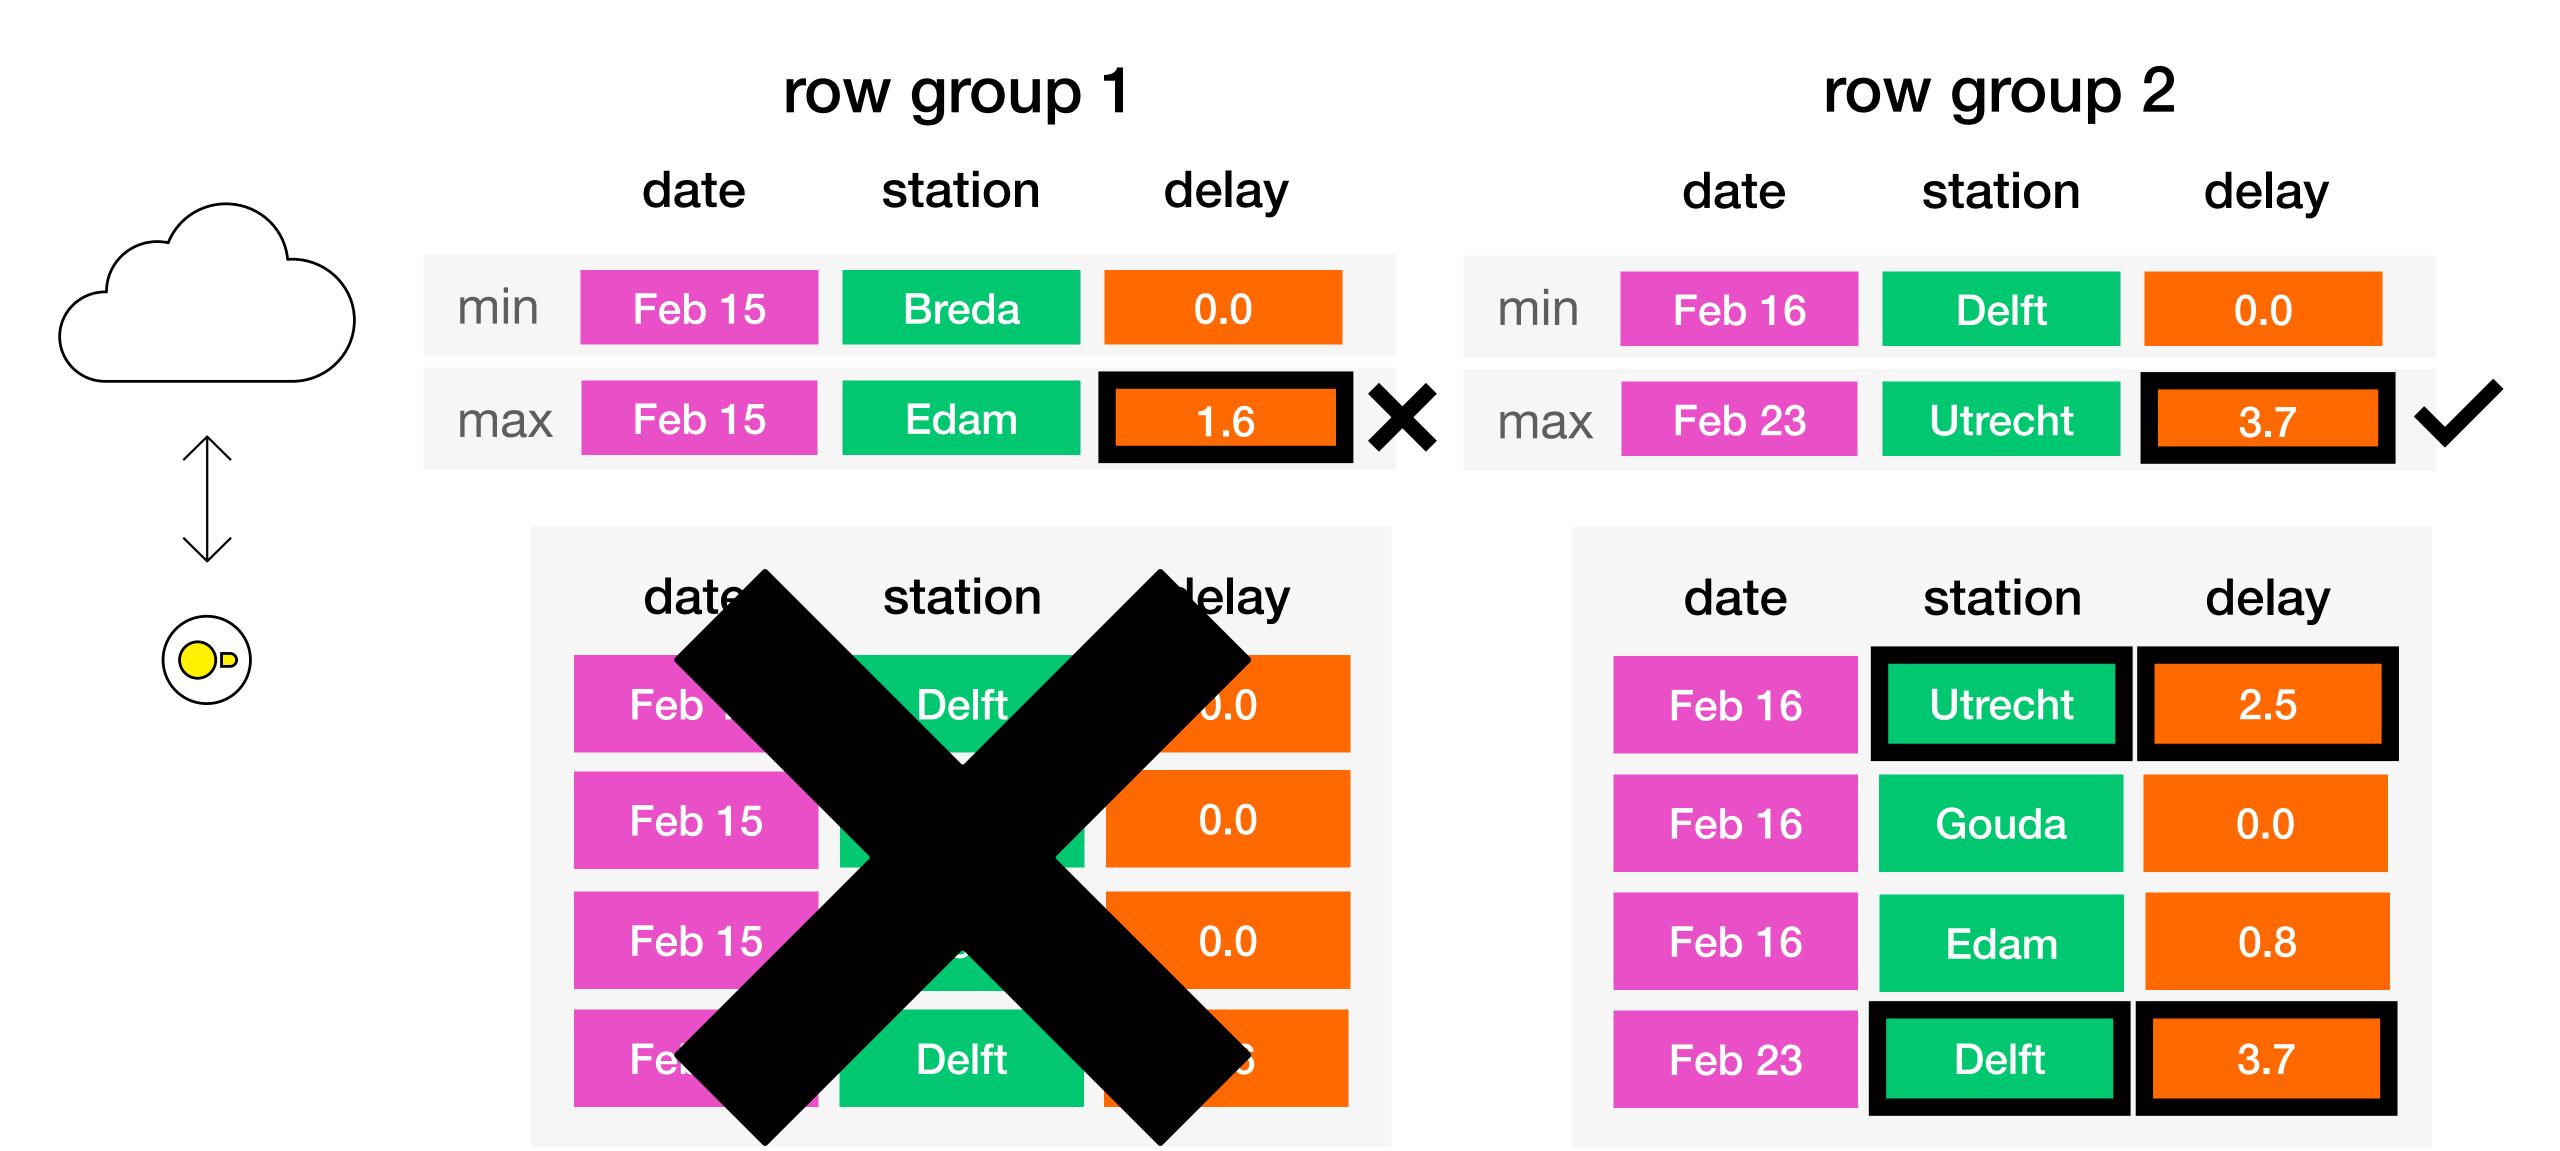
Which stations are affected by delays > 2 min?

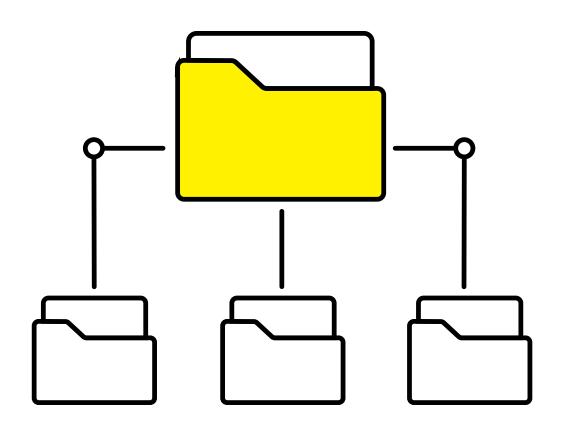


```
SELECT station, delay
FROM 's3://my-bucket/trains-2025-feb.parquet'
WHERE delay > 2;
```



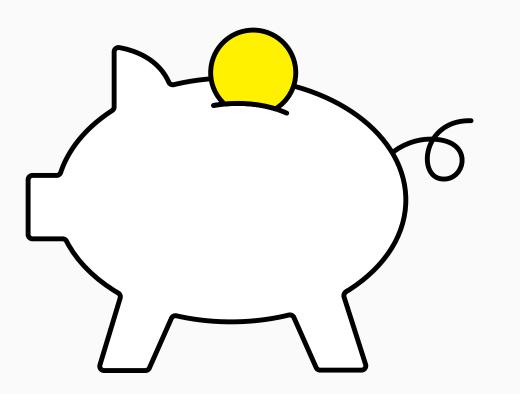
Which stations have delays > 2 min?





Use cases

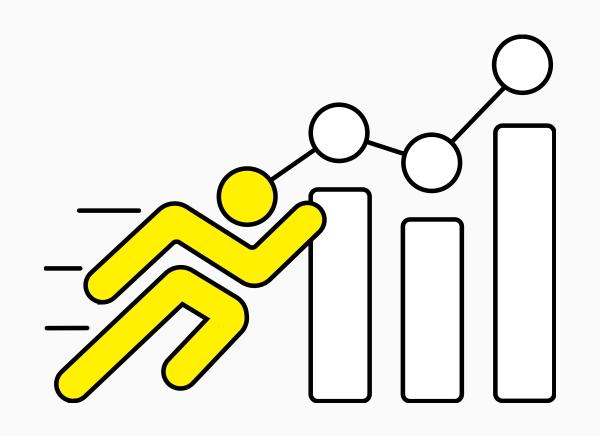
Cost saving



Fast local development: time & egress fees

Replace proprietary systems: license cost Replace distributed systems: hardware cost

Last mile analytics

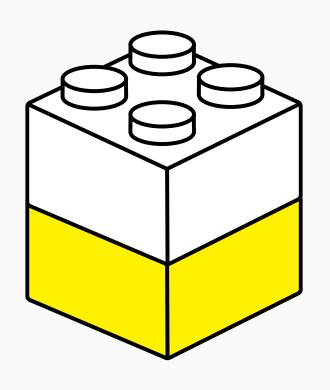


1PB log → pre-process with Spark

Process the 200 GB aggregate DuckDB

Build fast dashboards

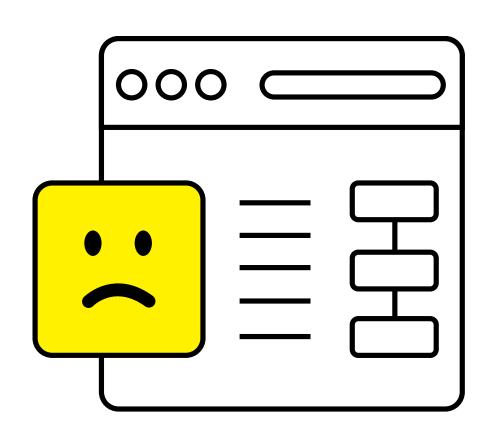
Modernizing codebases



Postgres-compatible SQL dialect

Builds on open formats (CSV, Parquet)

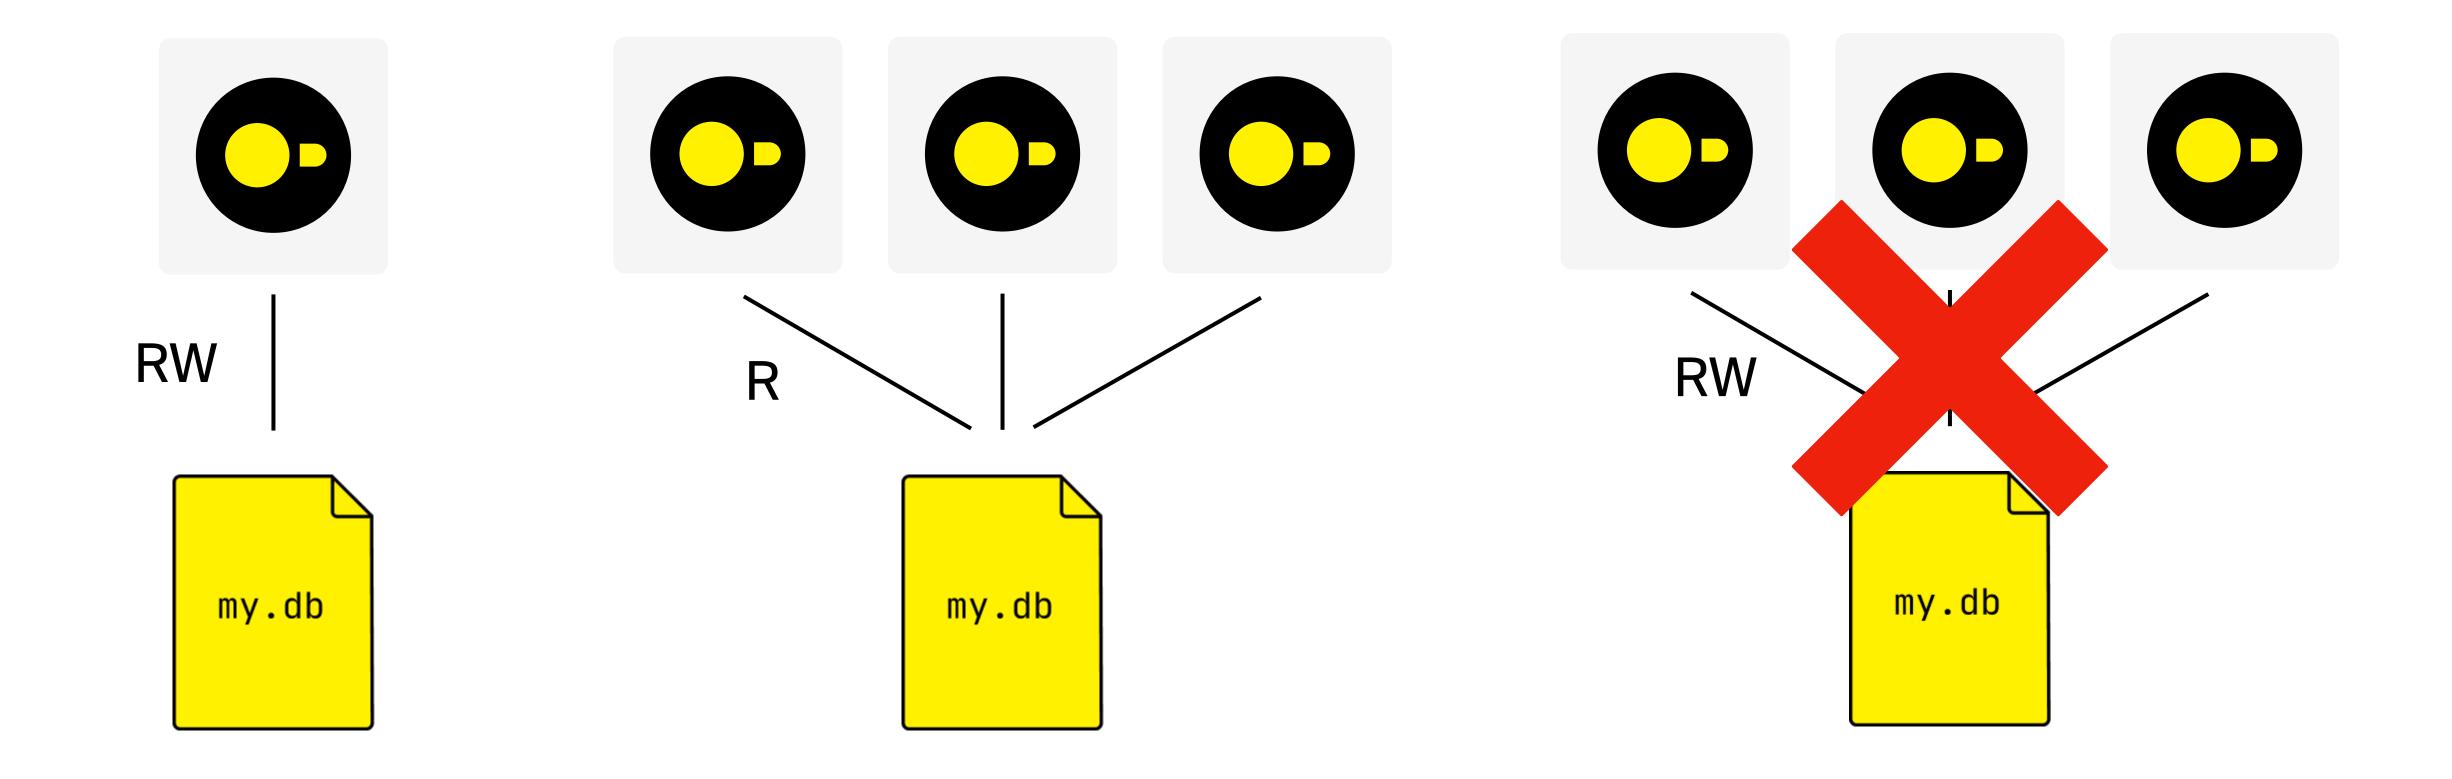
LLMs can refactor code to use DuckDB



Limitations

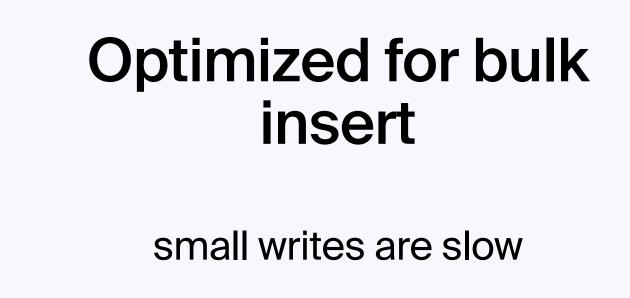
What's not so great about DuckDB?

Single-player experience: just one writer



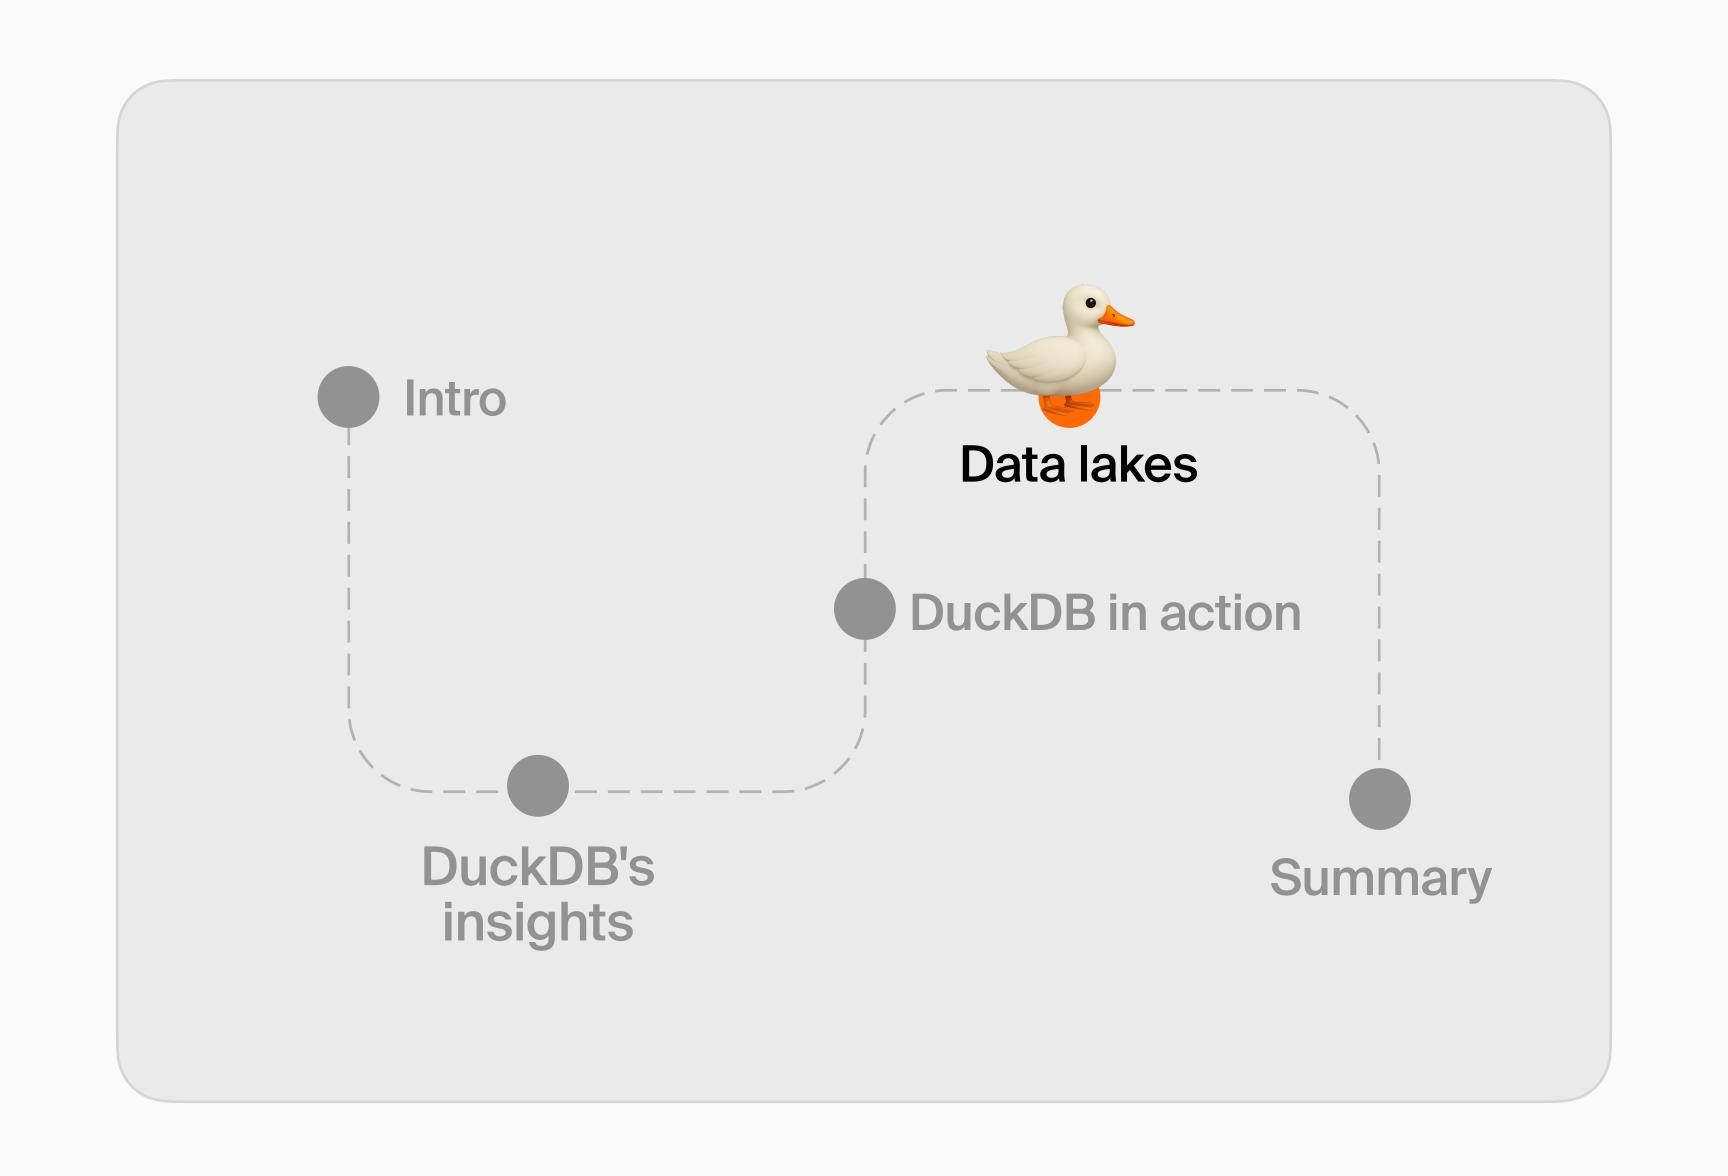
What's not so great about DuckDB?

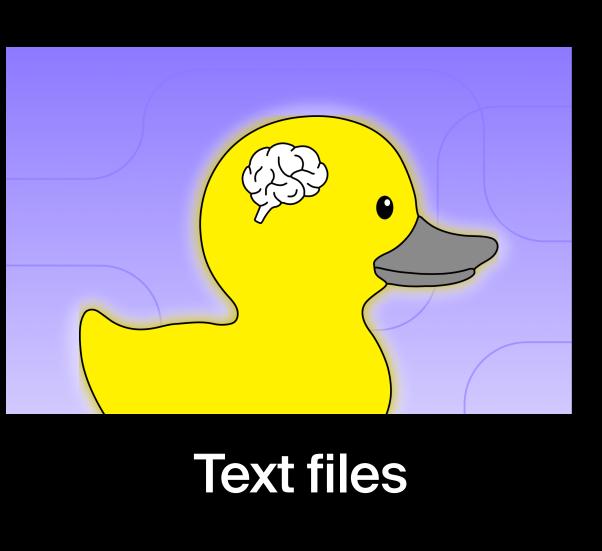
10 TB CSV = 2.7 TB database file

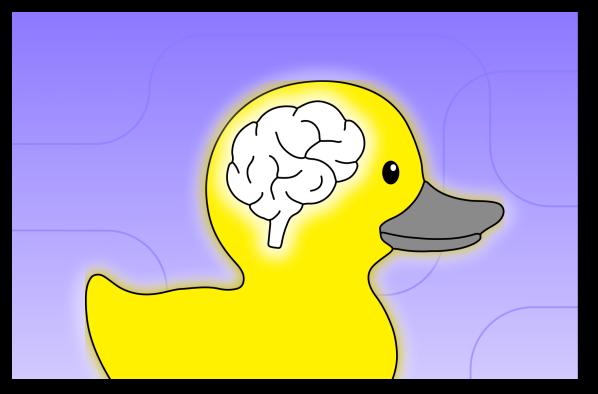


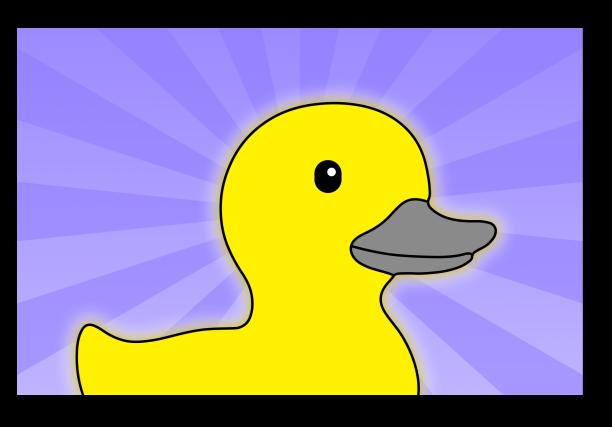


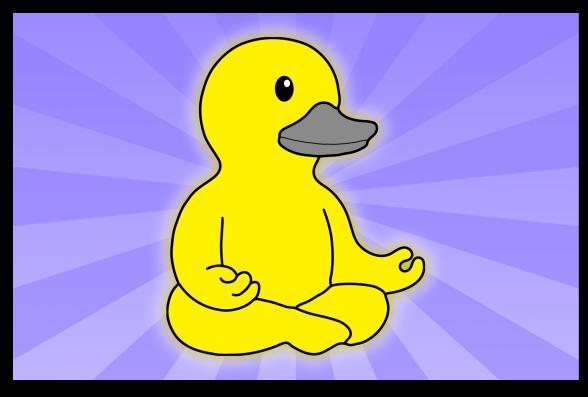








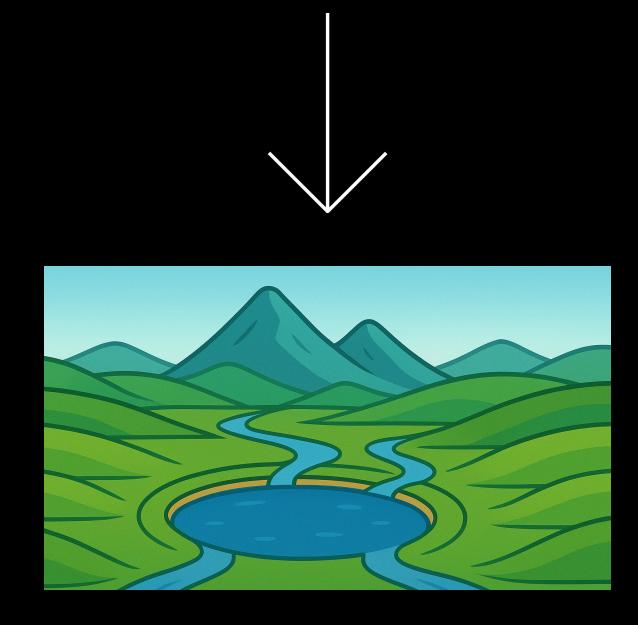




les Text files (deluxe)

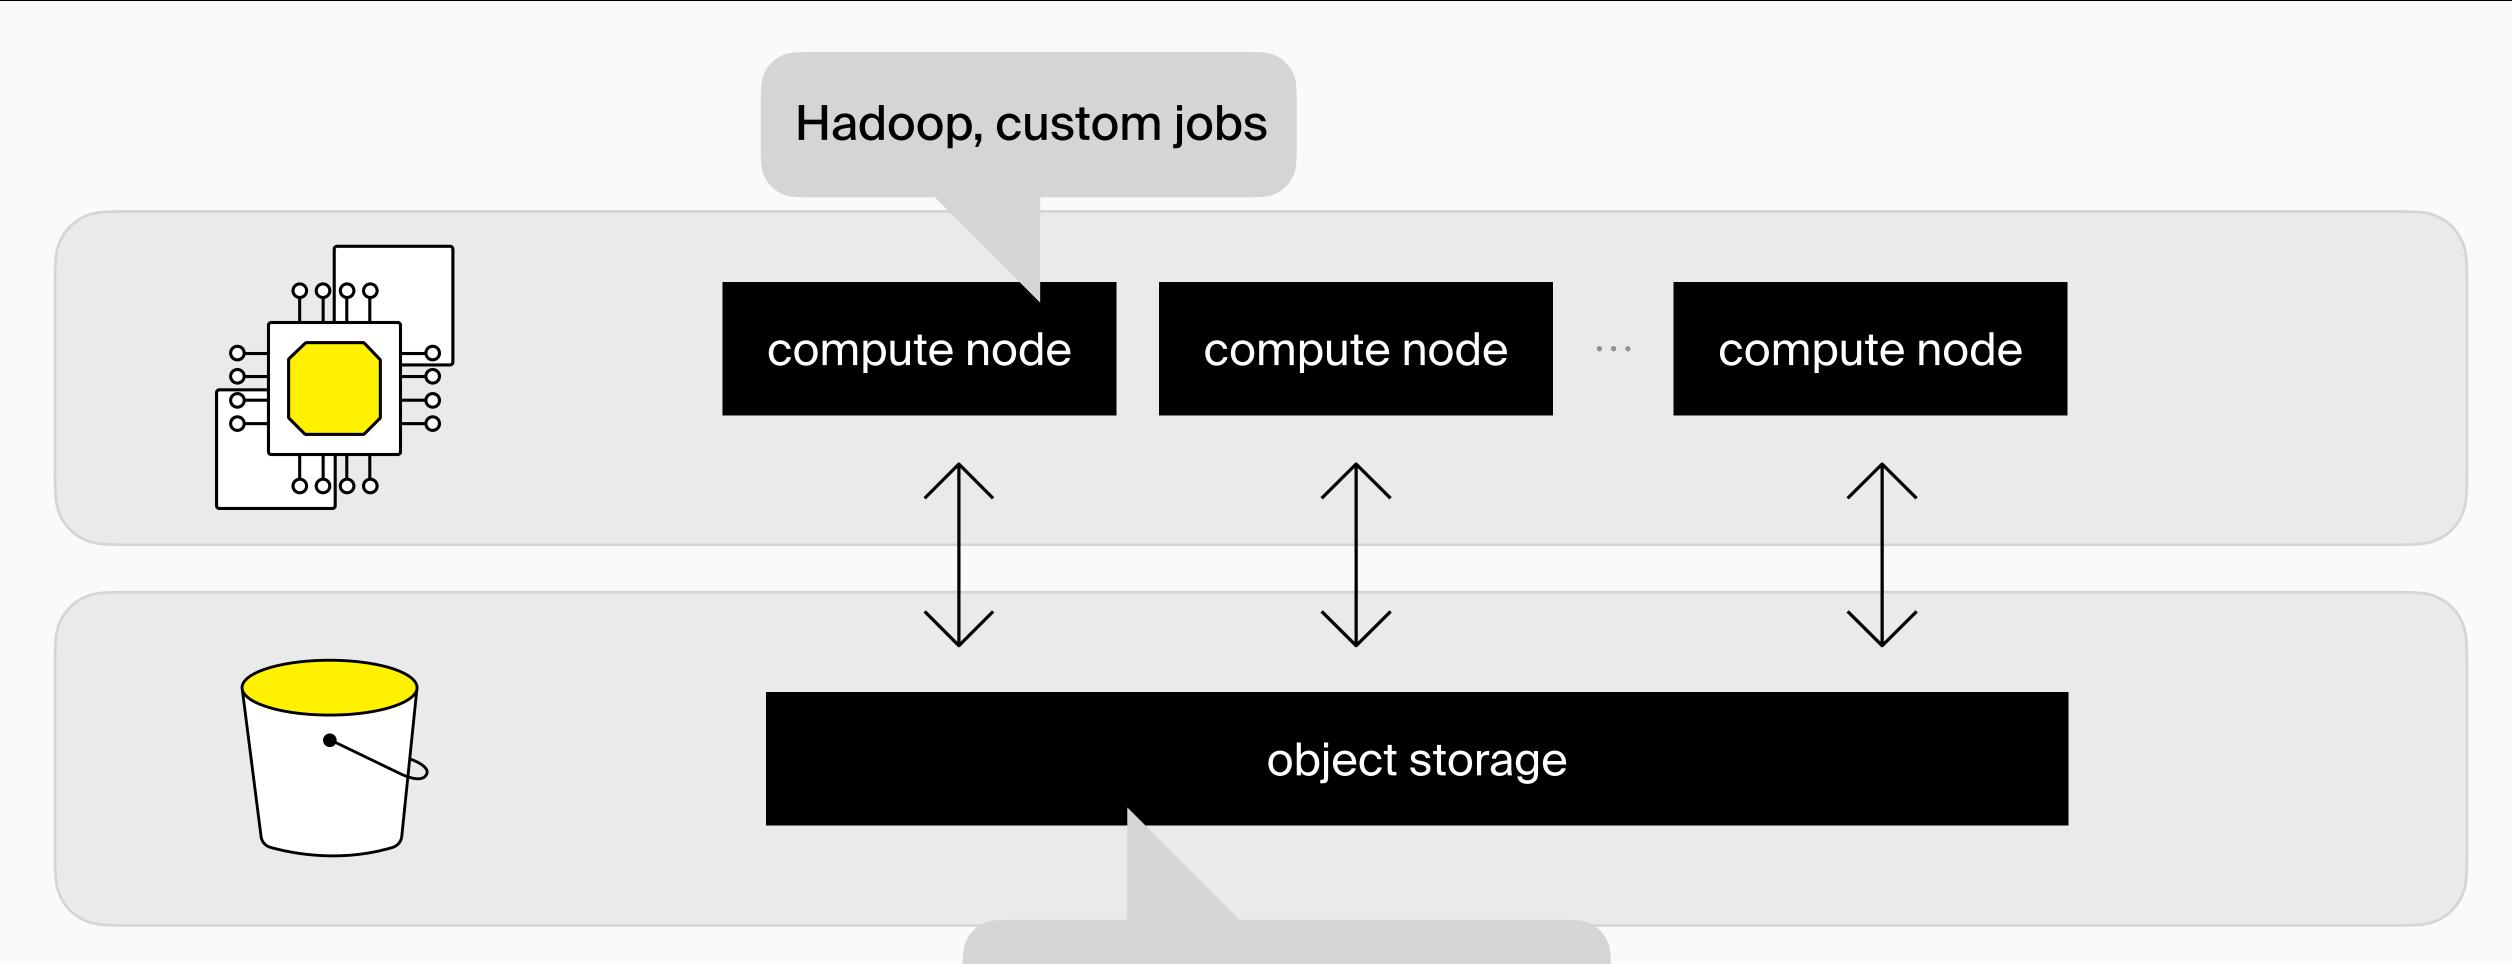
Binary files

Database





Disaggregated storage

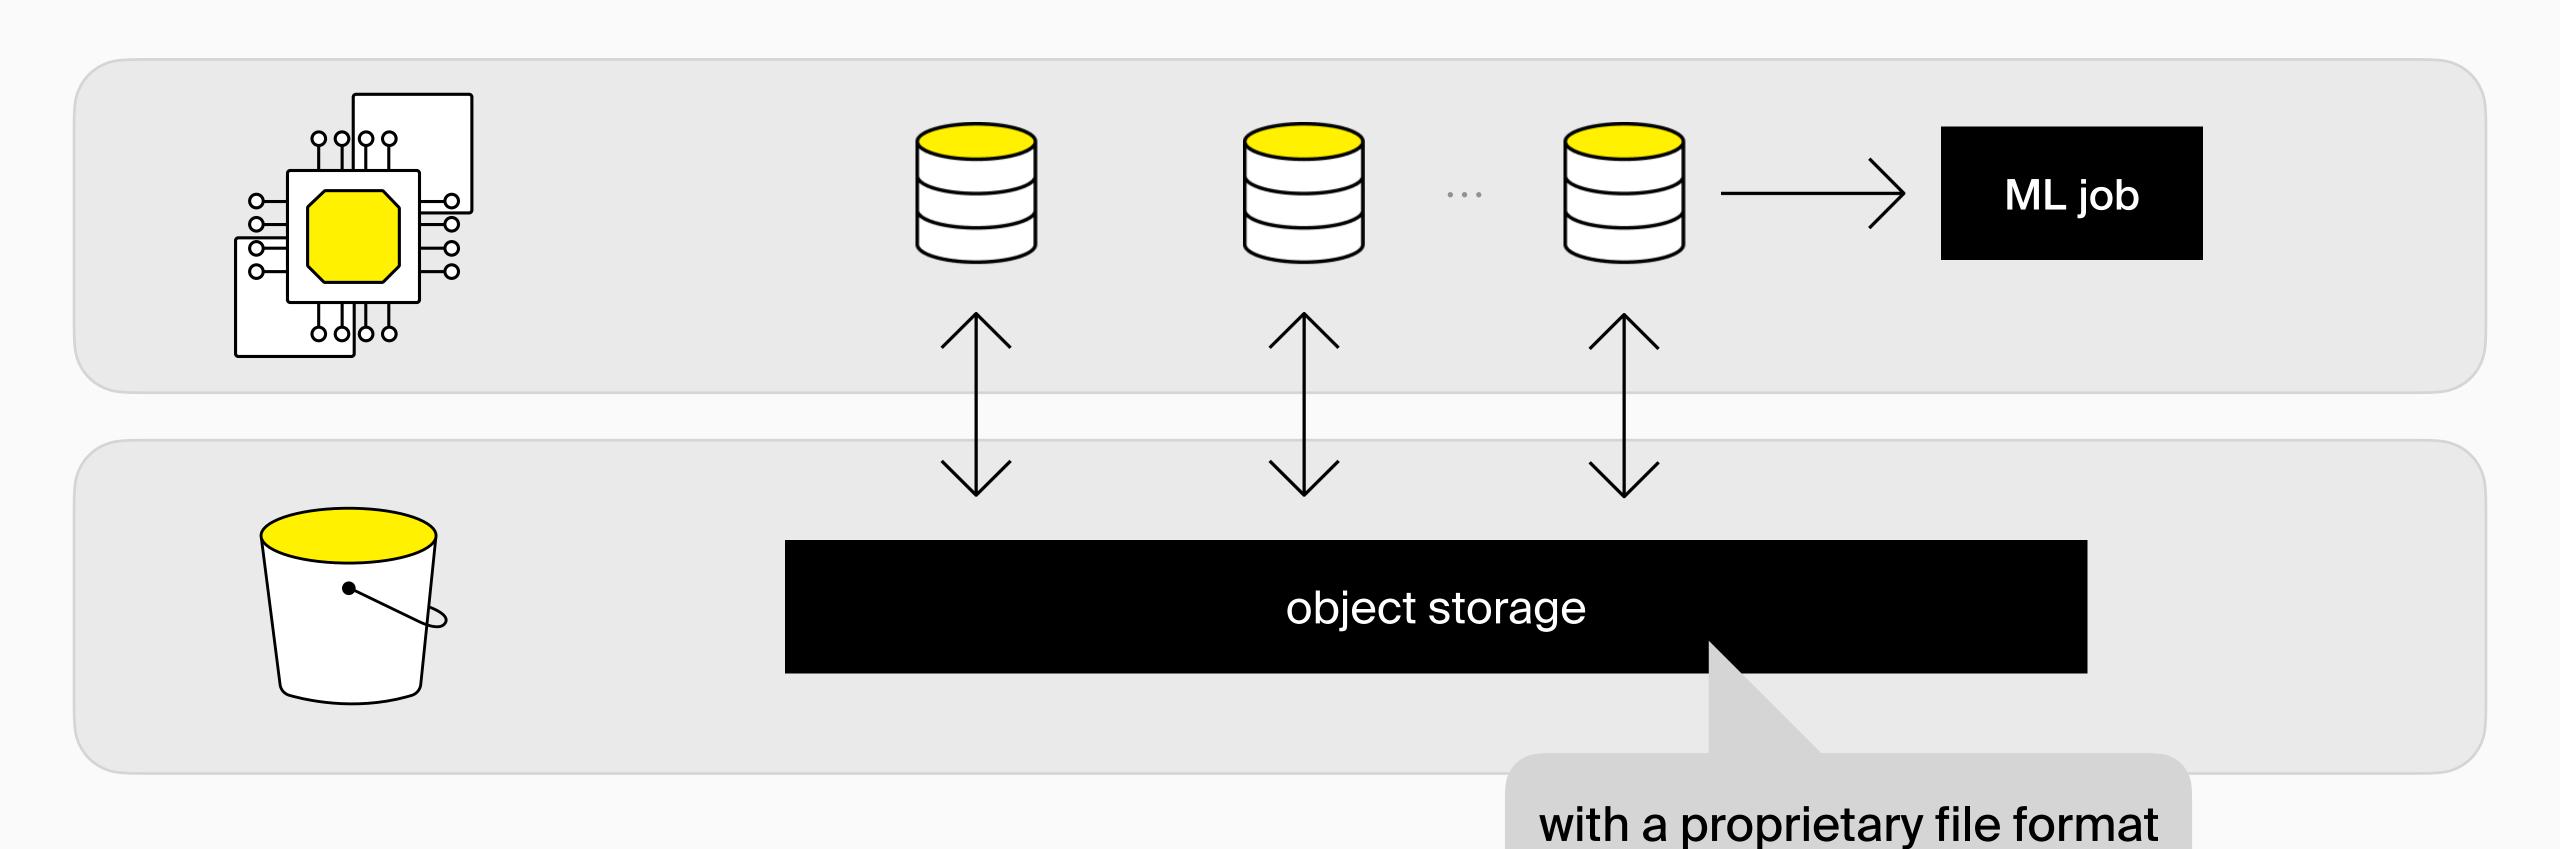


HDFS / S3, immutable files

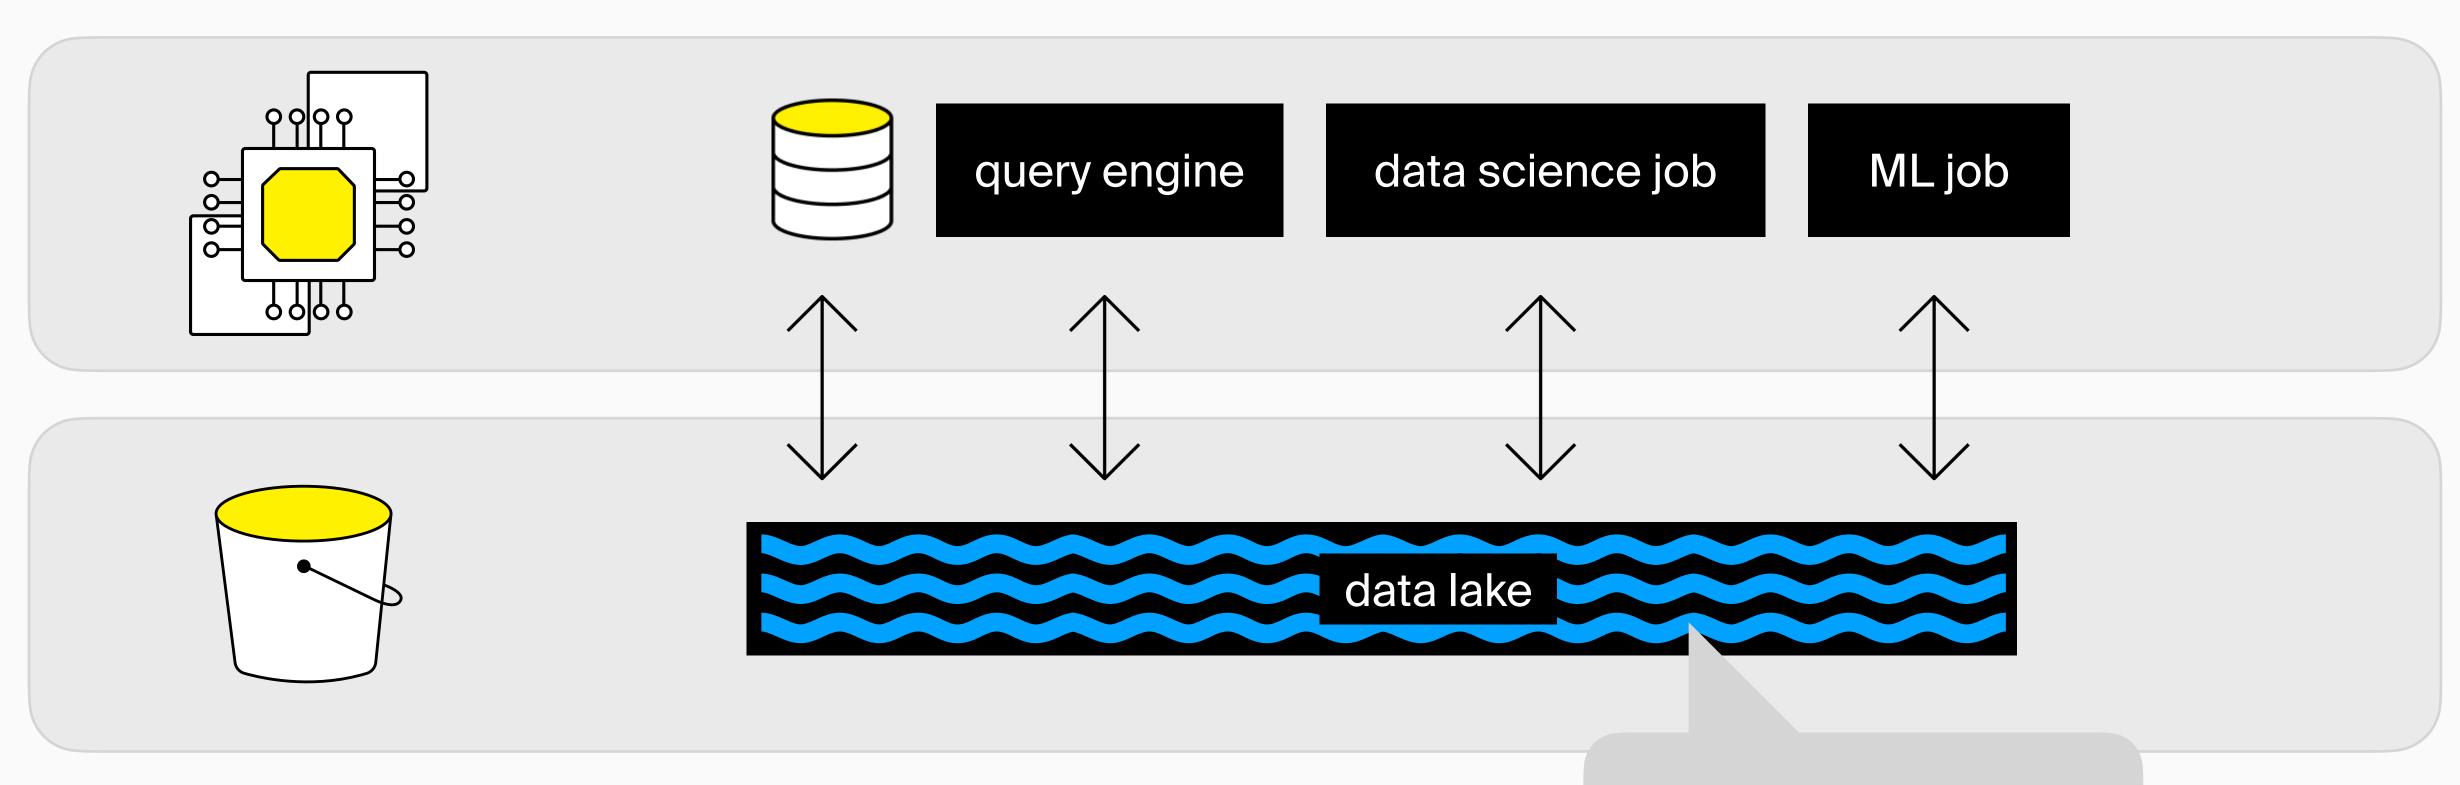
2010

Cloud data warehouses

Google
Big Query
Snowflake



Data lake architecture

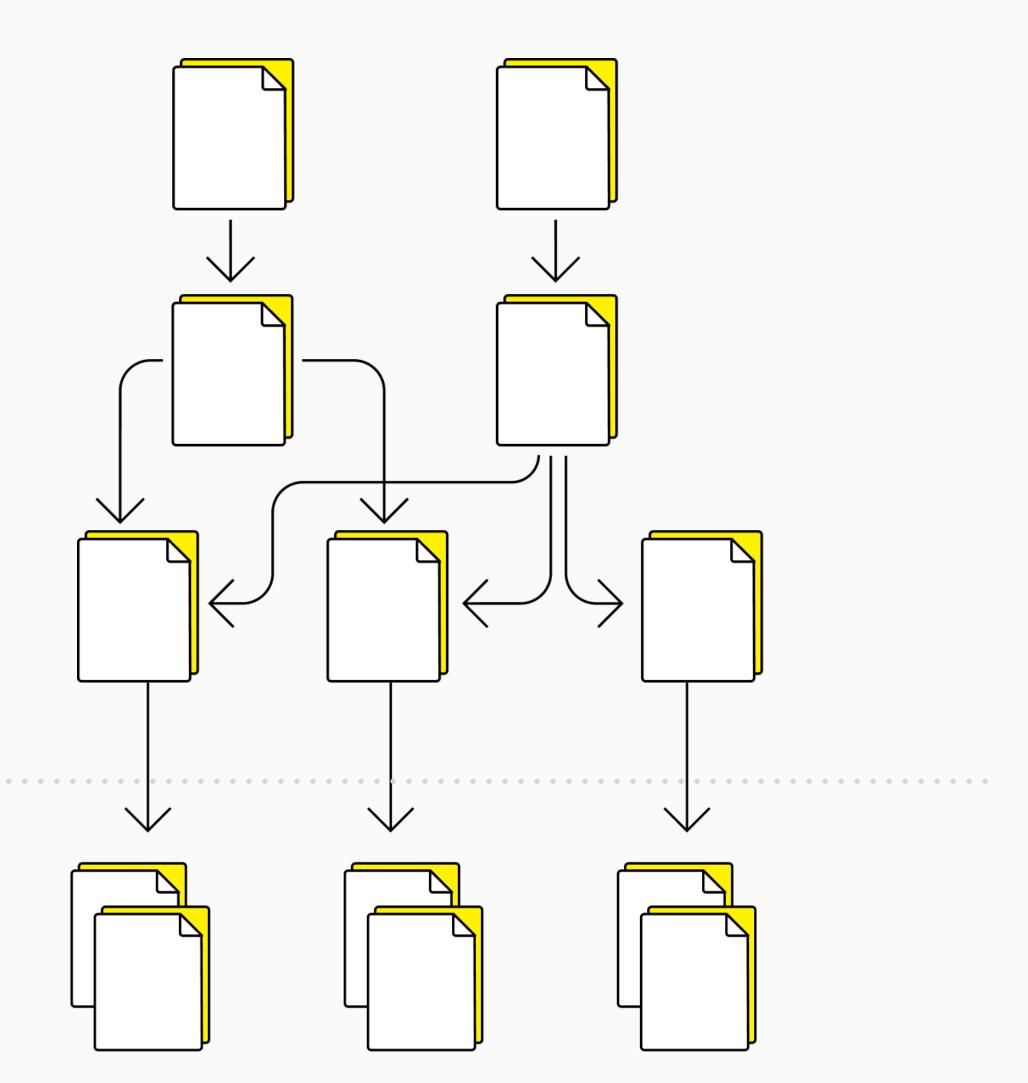


Iceberg, Delta Lake, ...



metadata layer (json and avro)

data layer (parquet)

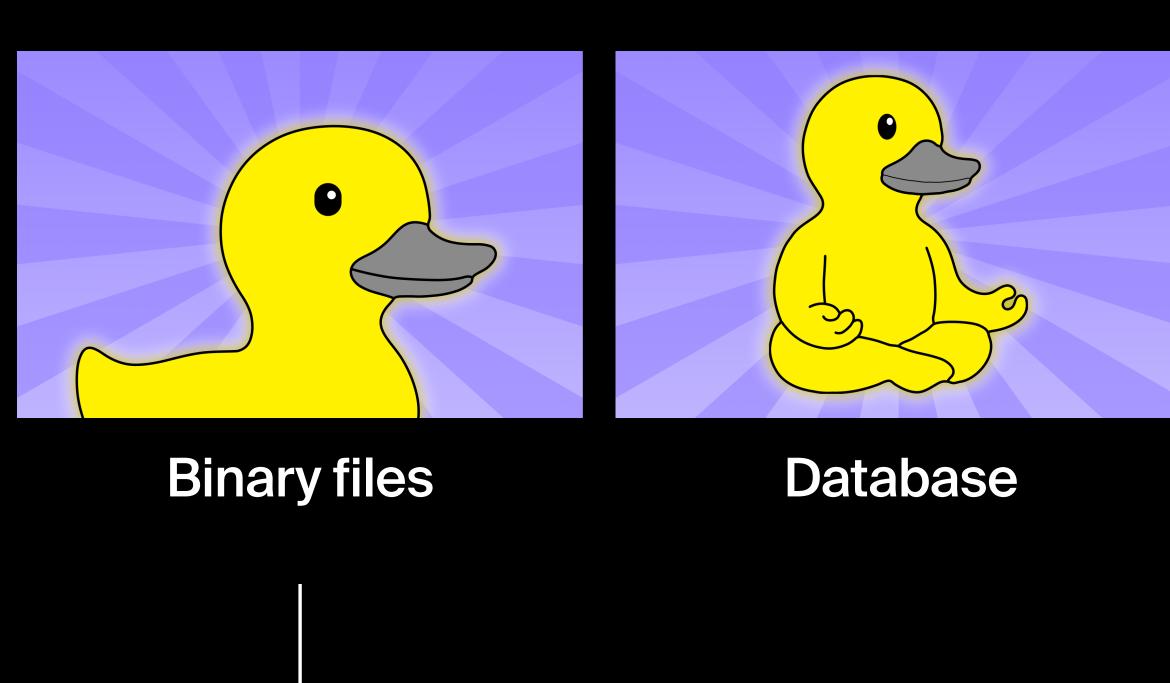


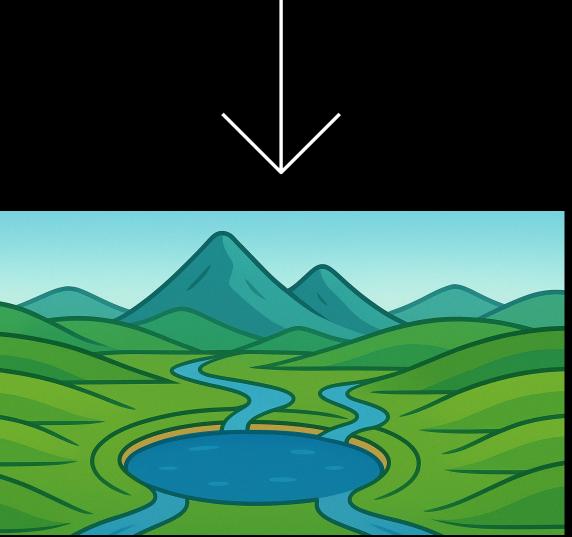


- File-only architecture
- Transactional updates
- Time travel queries
- Schema evolution

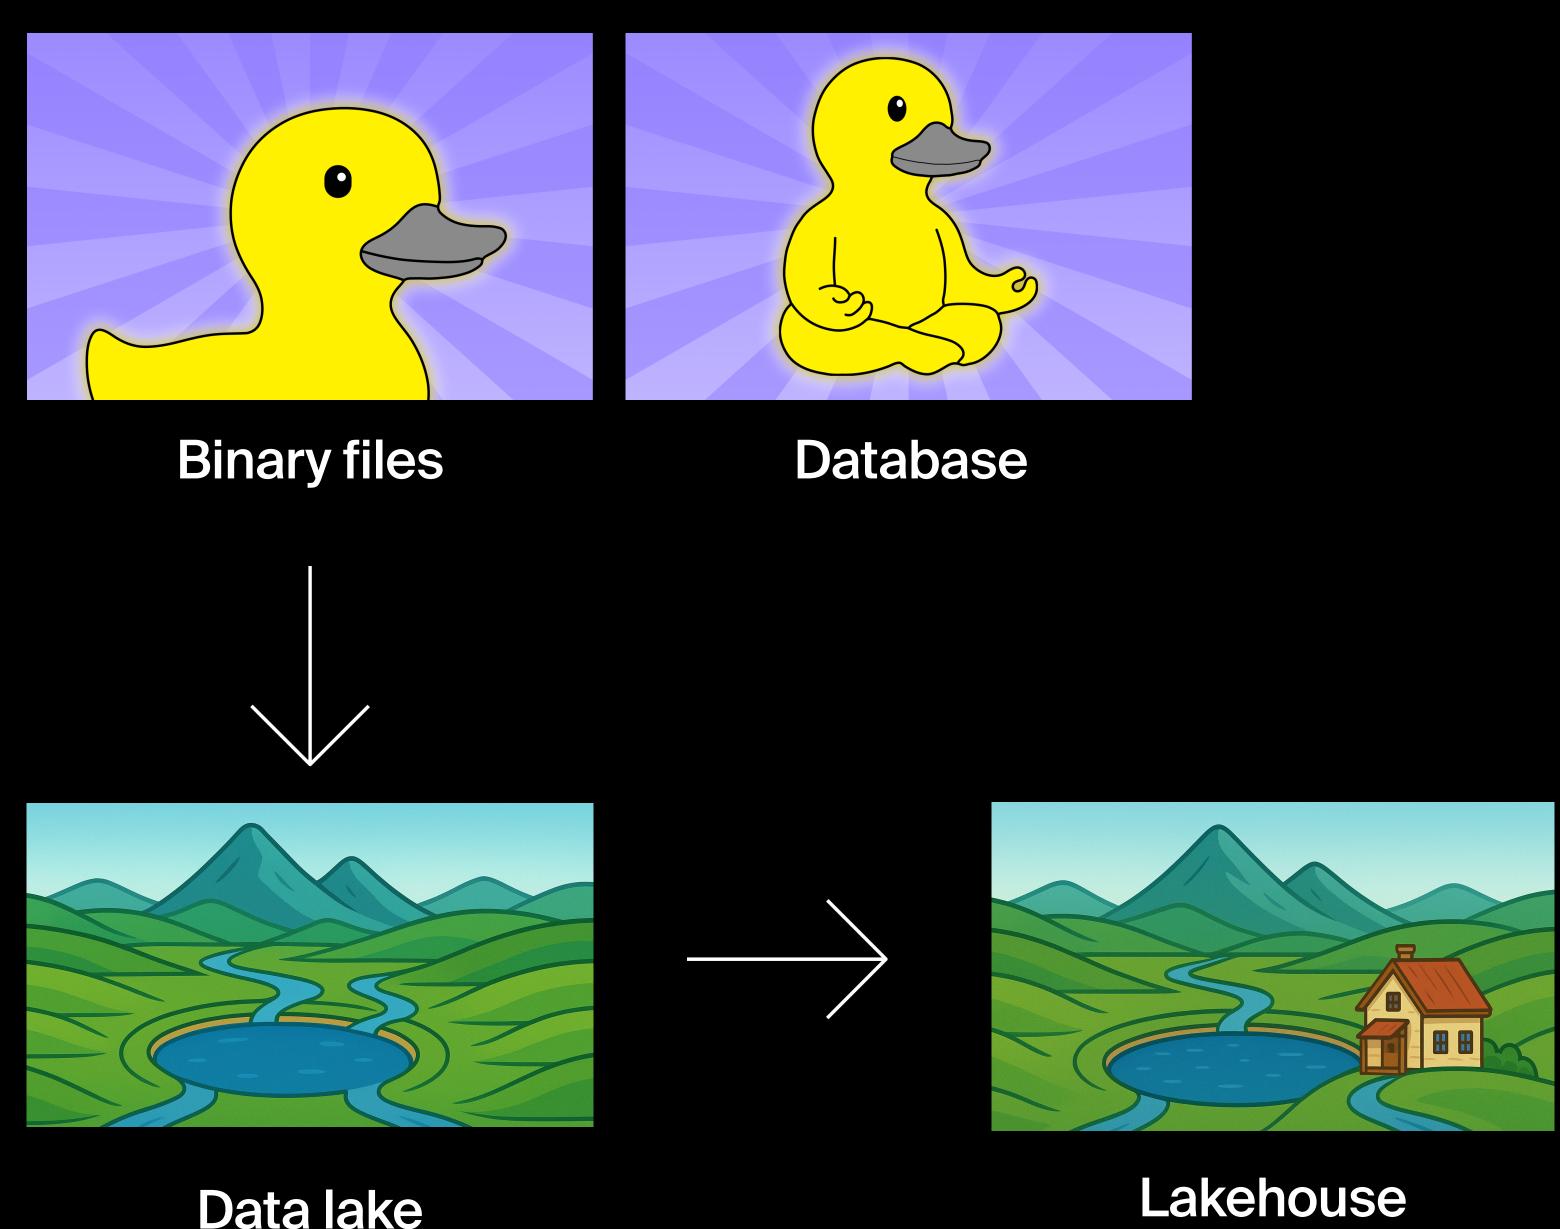


- Performance issues
- "Small files" problem
- Limited to a single table



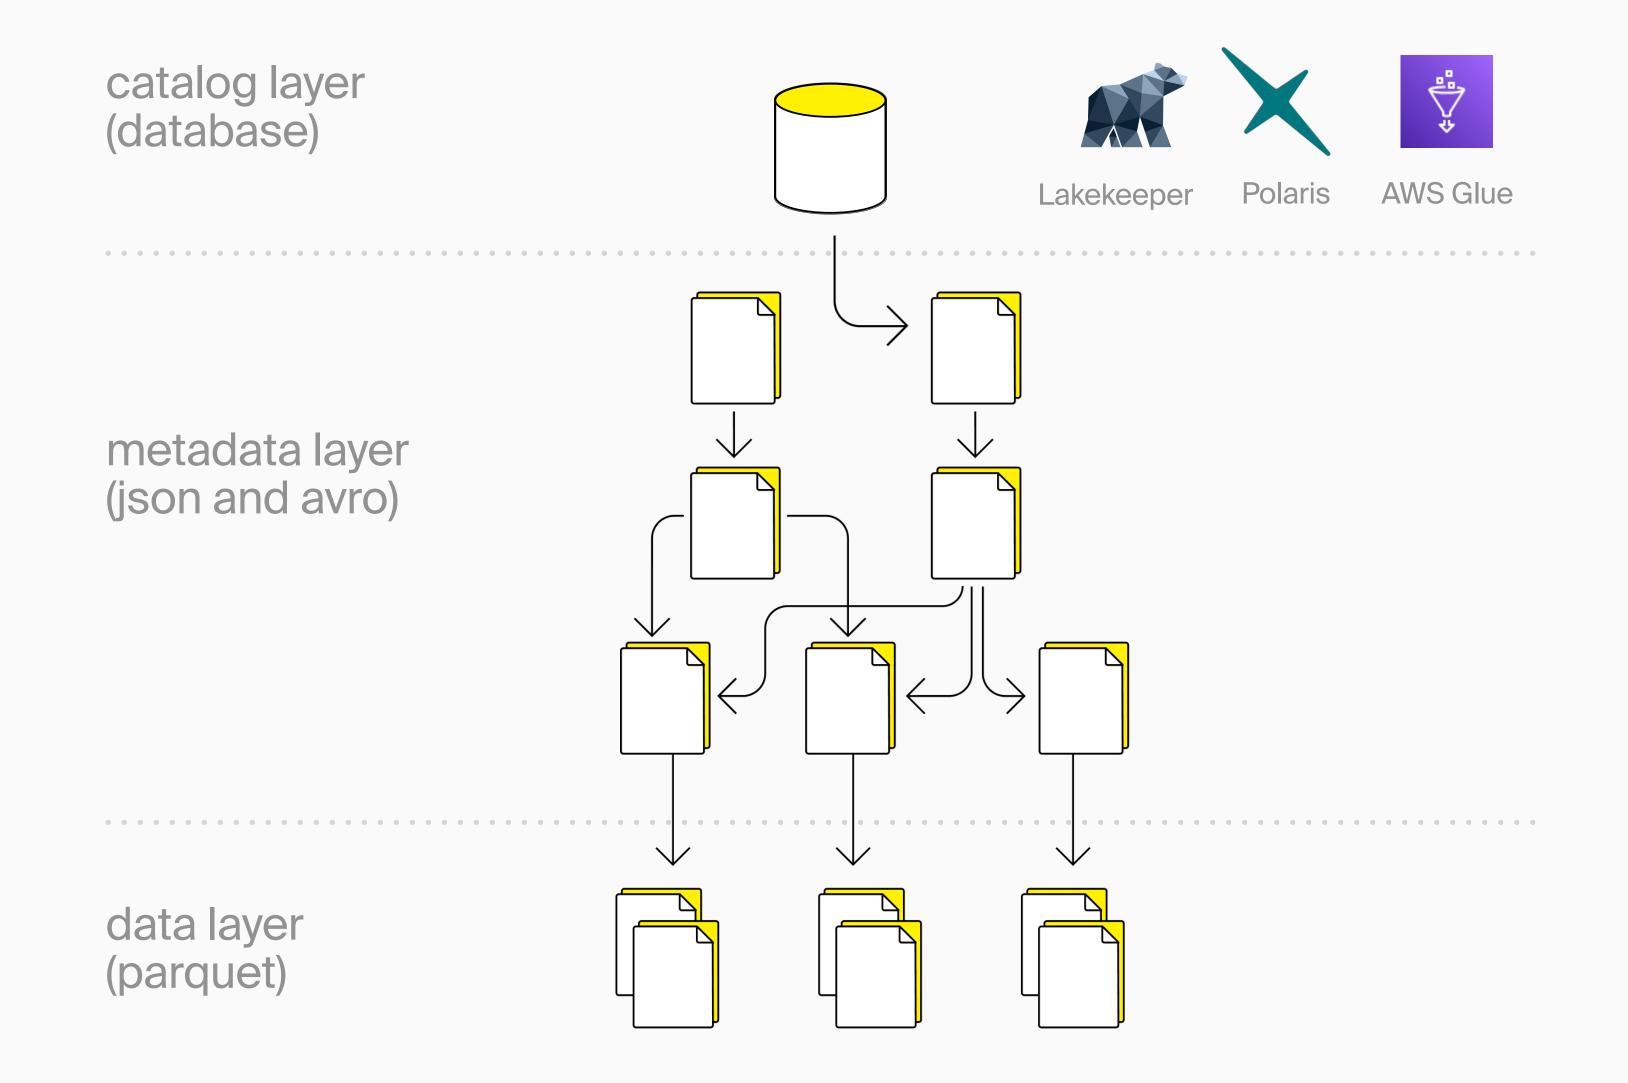


Data lake



Data lake

2020 Lakehouse



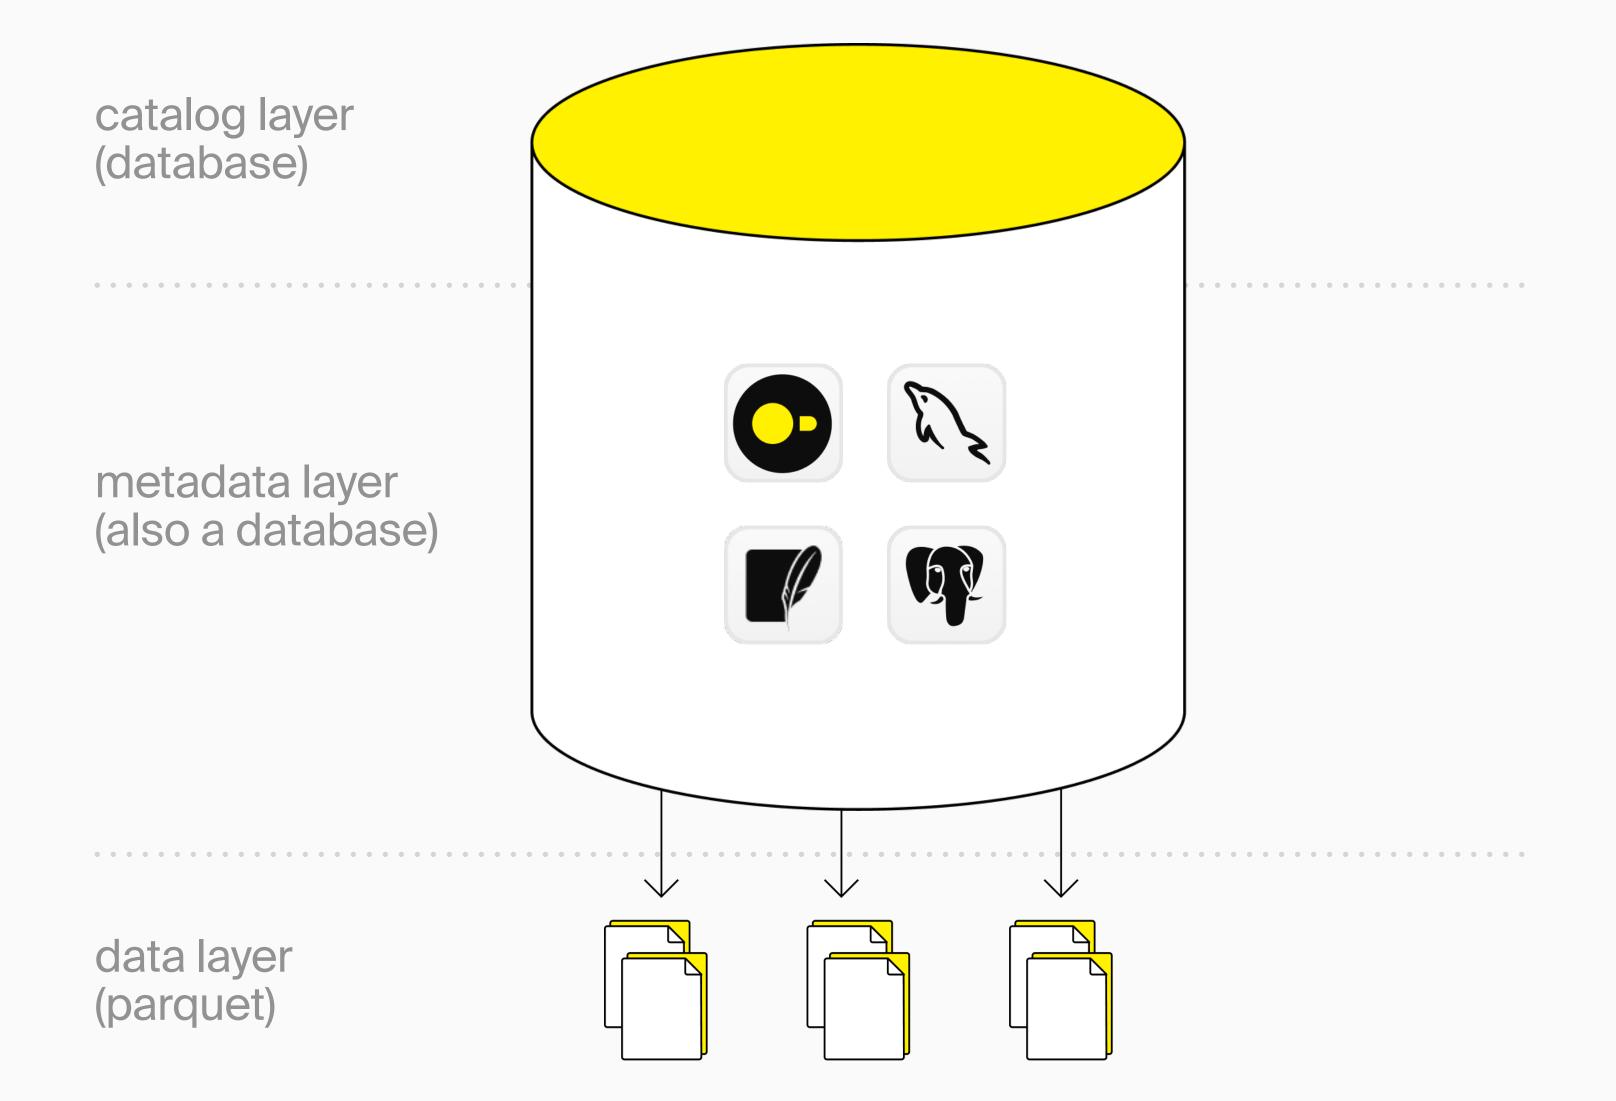


- Cheap object storage
- Small catalog database
- Full database features
- Scalable



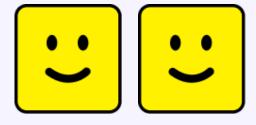
- Performance issues
- Operational complexity

2025 DuckLake





- Cheap object storage
- Small catalog database
- Full database features
- Scalable



- Lightweight snapshots
- Multi-table ACID
- High performance

```
ATTACH 'ducklake:trains.ducklake' AS trains;
USE trains;

CREATE TABLE services AS
FROM 'trains-2025-feb.csv';
```

```
UPDATE services
SET train_number =
   train_number + (random() * 100 + 1)::INT
WHERE delay > 150;
SELECT snapshot_id, changes
FROM ducklake_snapshots('trains');
```

snapshot_id int64	changes map(varchar, varchar[])				
1	{schemas_created=[main]} {tables_created=[main.services], tables_inserted_into=[1]} {tables_inserted_into=[1], tables_deleted_from=[1]}				

Use time travel

```
SELECT date, train_number, delay
FROM services AT (VERSION => 1)
WHERE date = '2025-02-27'
  AND station = 'Amsterdam Centraal'
AND train_number = 420;
```

date	train_number	delay	
2025-02-27	420	174	



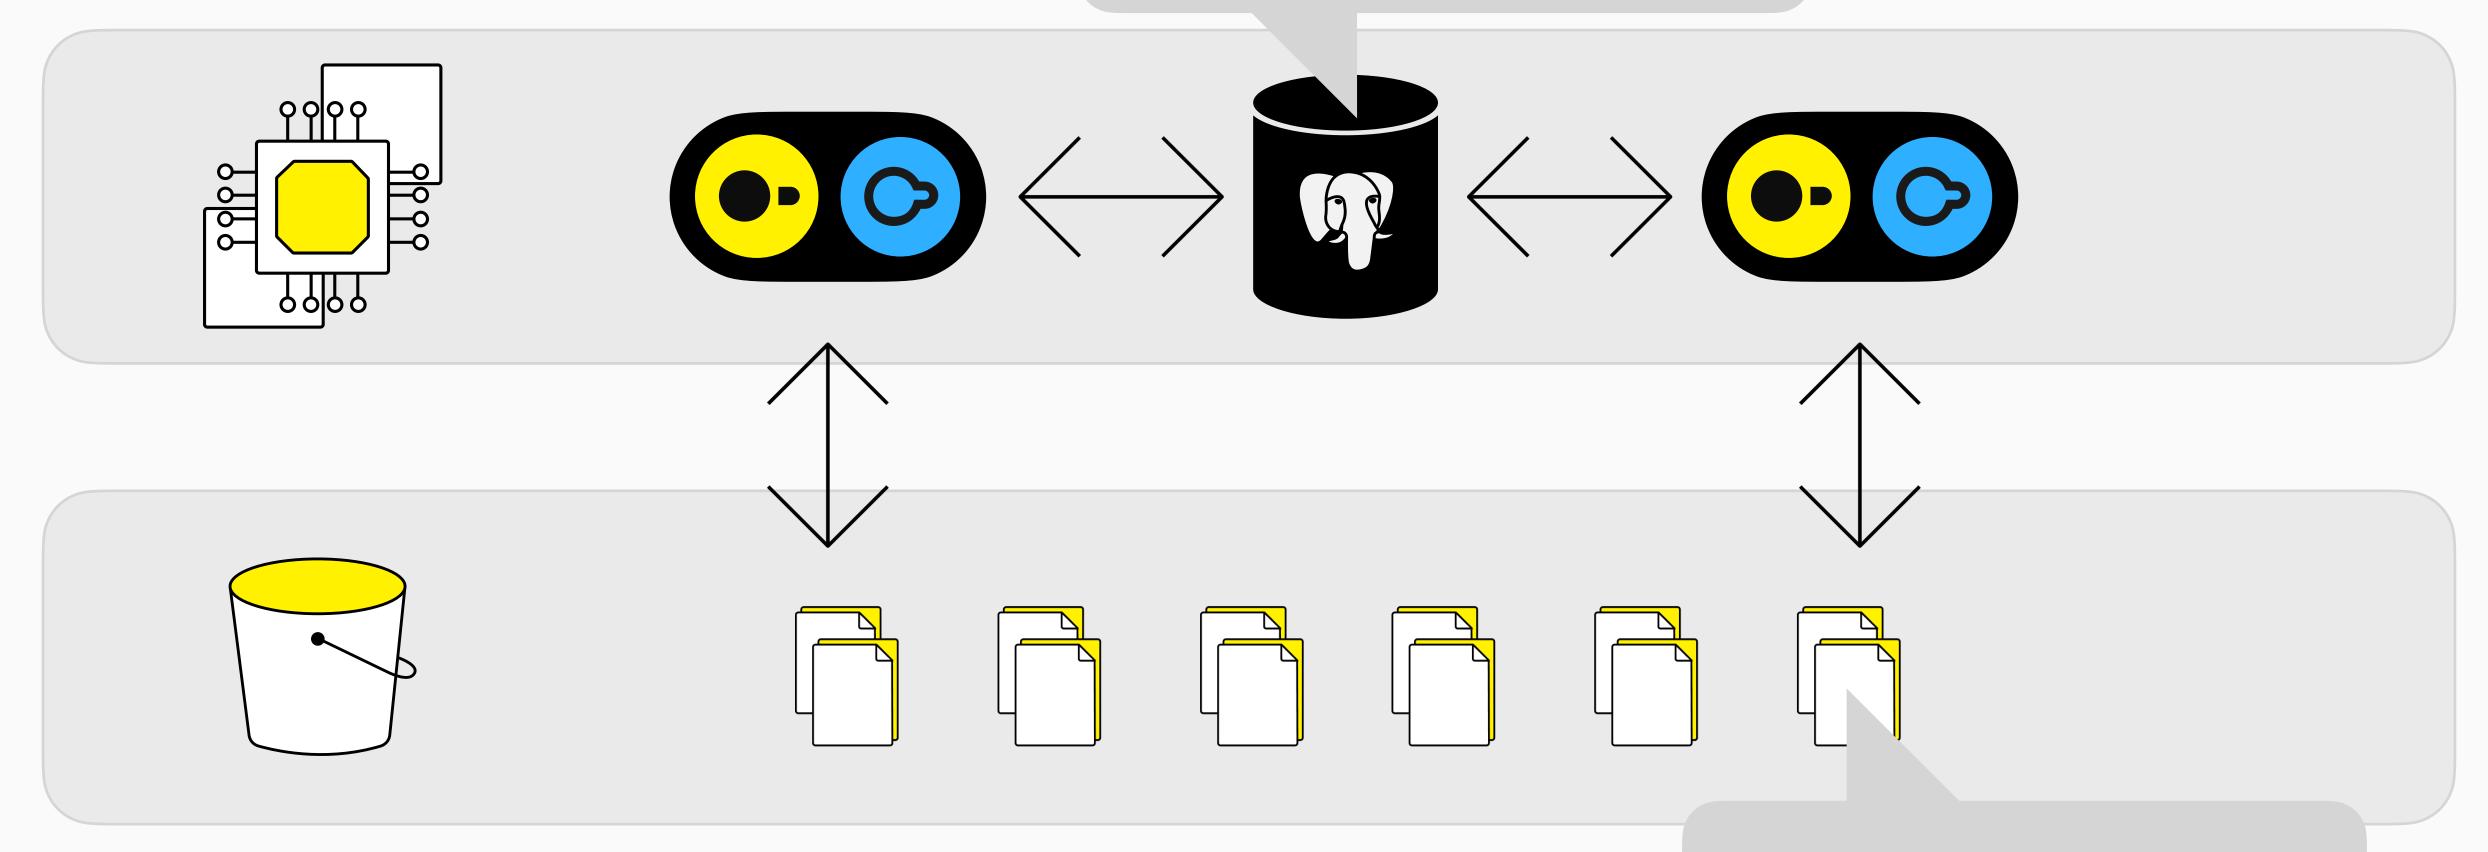
You can look back in time: train company busted

Also great for re-running fixed scripts on old data



Multiplayer DuckDB architecture

Catalog database with inlining



Parquet files with pushdown



Laboratory experiment, do not replicate!

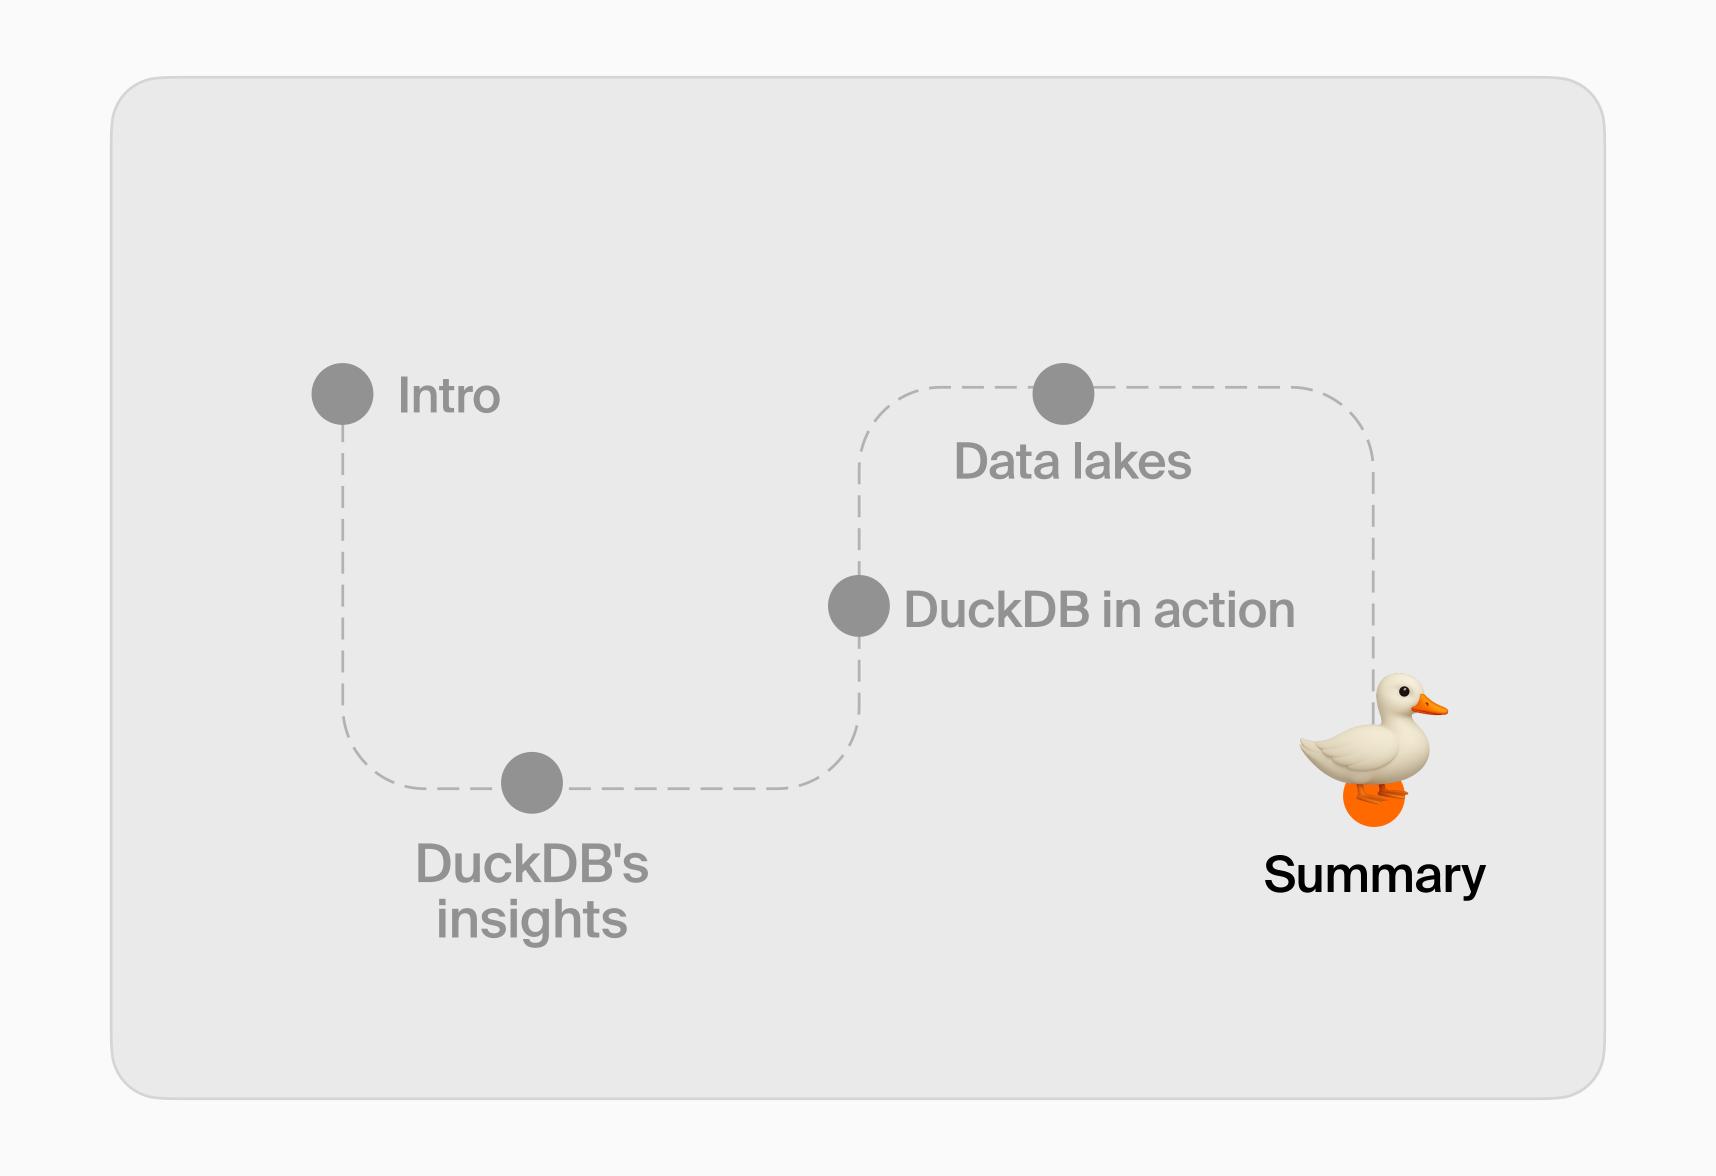
Demineralized water

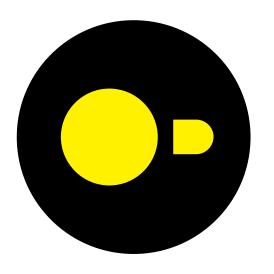
Brand new phone

Limited time underwater

Still a bad idea: Wi-Fi and 5G signals drop while underwater!

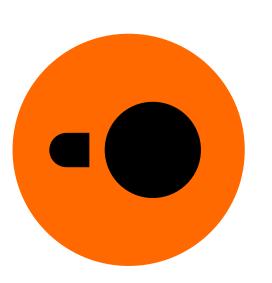
DuckLake on a phone *underwater*





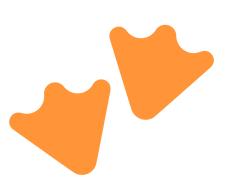
DuckDB Foundation

Intellectual property and TM



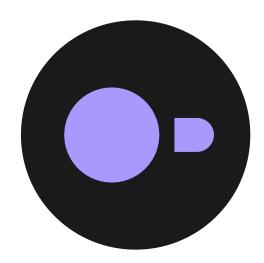
DuckDB Labs

Core DuckDB / Amsterdam



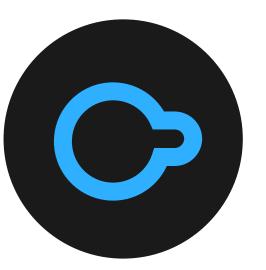
MotherDuck

Cloud data warehouse / Seattle



Community Extensions

Third-party contributions



DuckLake

Lakehouse standard



pg_duckdb

Run DuckDB in Postgres