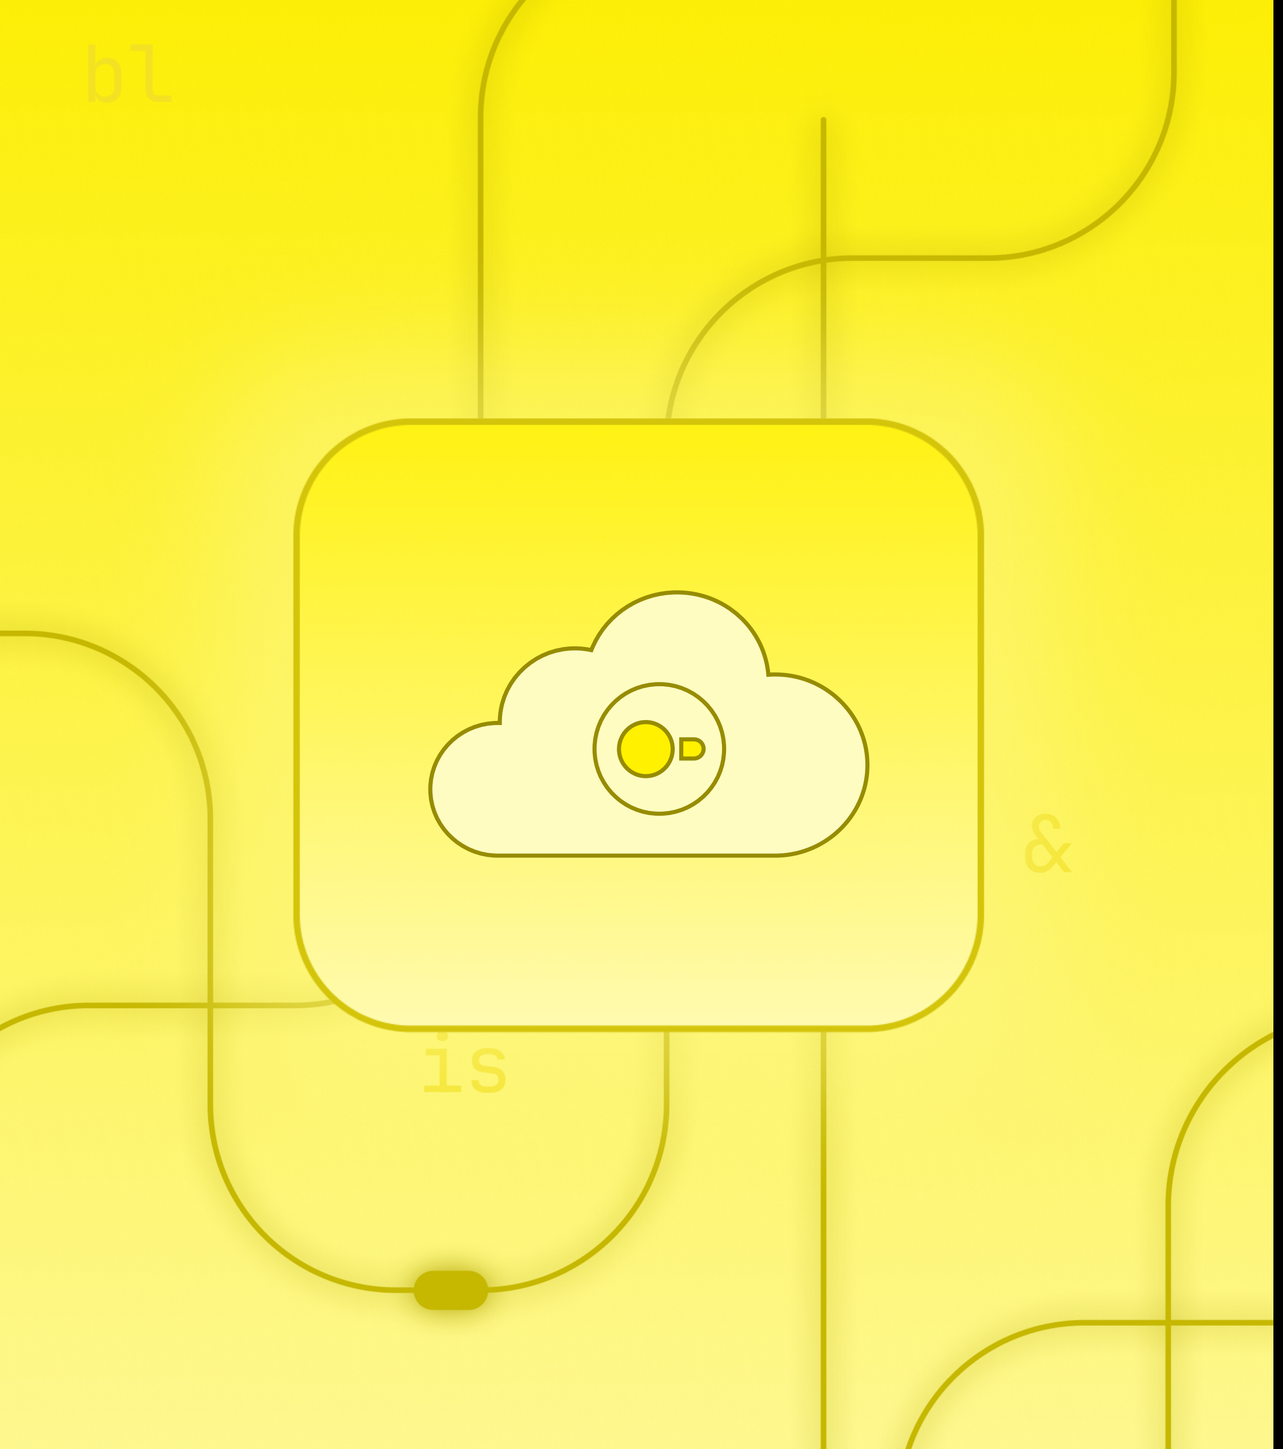

FOSDEM 2026 | DATABASES DEVROOM

DuckDB in the Cloud:

A Simple, Powerful SQL Engine for Your Lakehouse

GÁBOR SZÁRNYAS

DUCKDB LABS

DuckDB = a SQL database for analytics

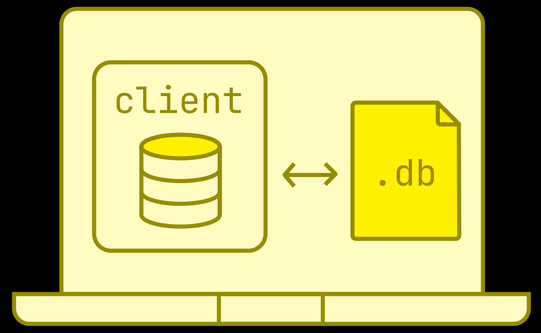


```
$ curl https://install.duckdb.org | sh
$ duckdb my.db

D SELECT *
  FROM '/your/favorite/file.csv';
```

```
$ pip install duckdb
$ python

>>> import duckdb
>>> duckdb.sql("FROM my_data.parquet")
```

```
... and 15 other clients
```

DuckDB deployment models

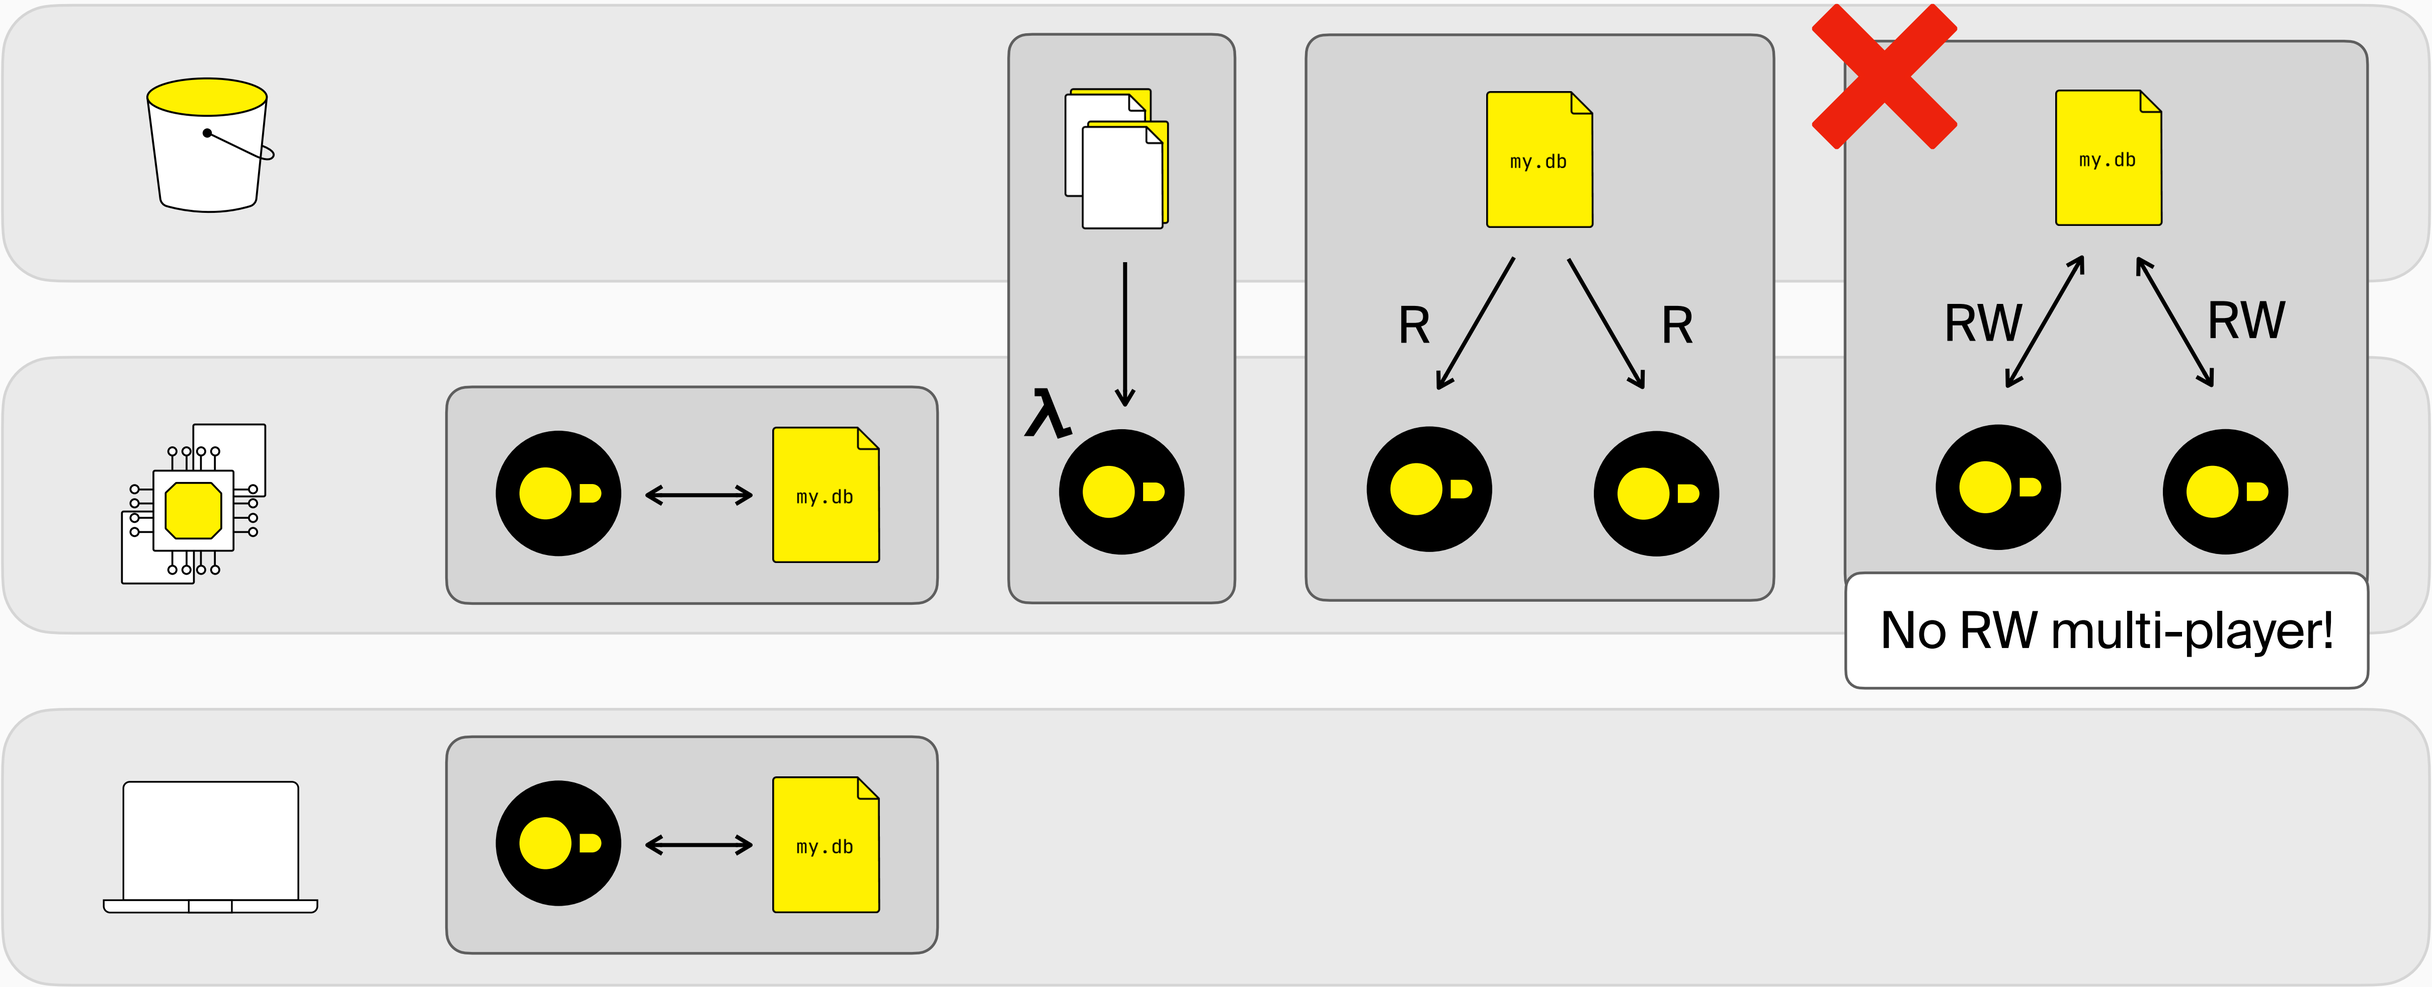

Data lake architecture (~2018)

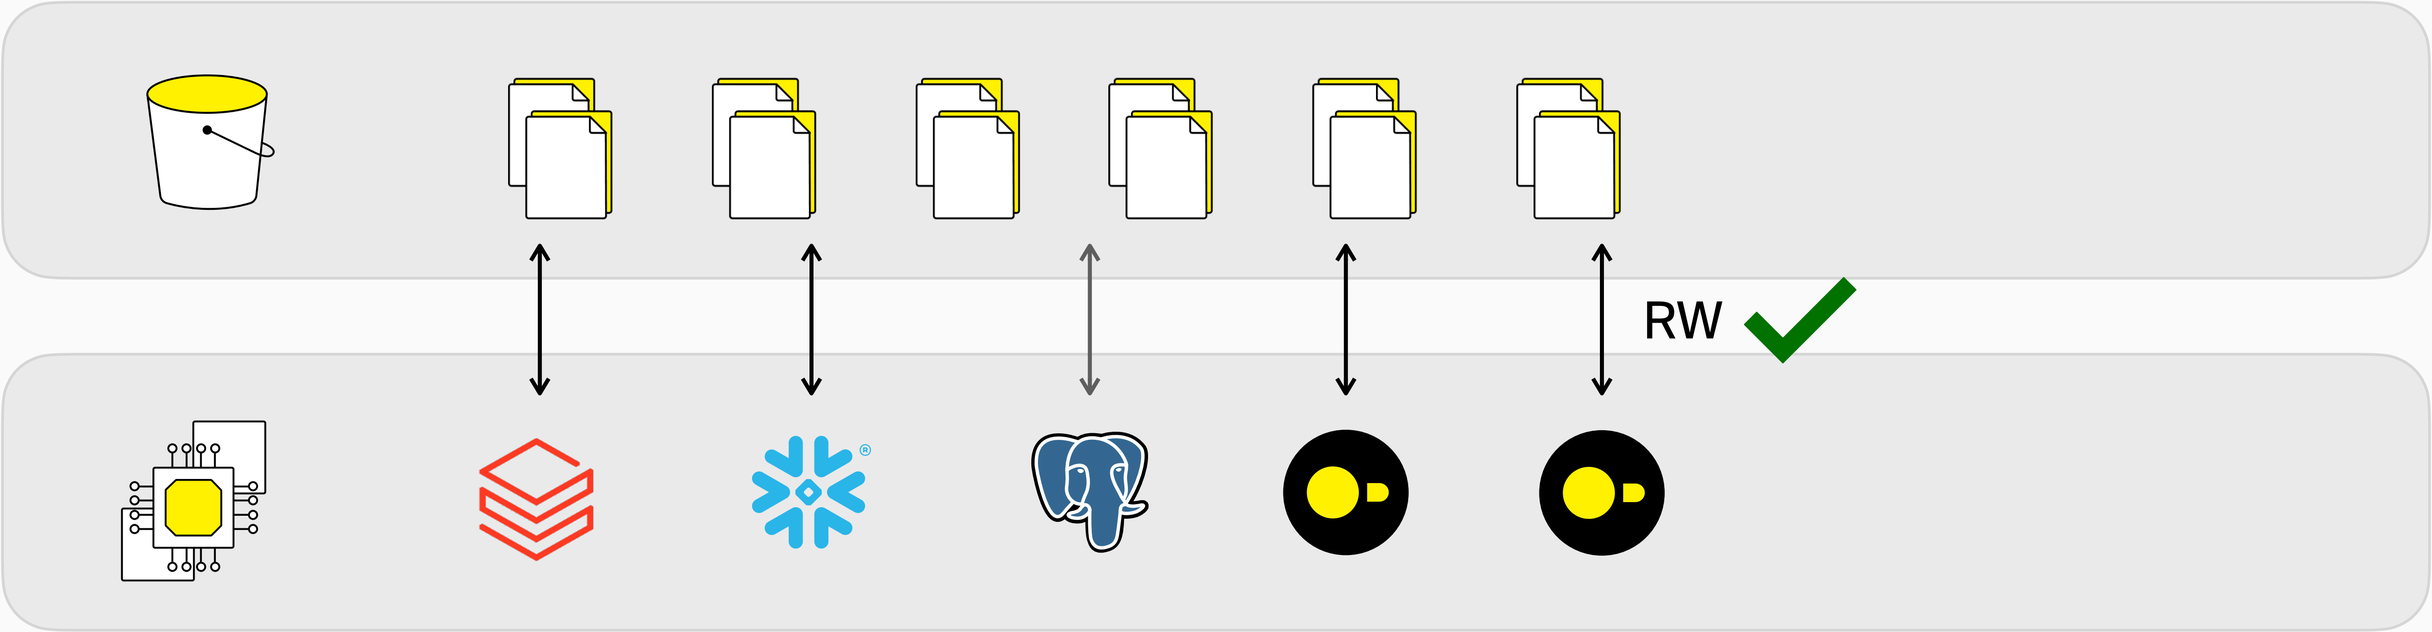

Multi-player & updates on object storage

Limited to a single table

Iceberg in DuckDB

```
$ duckdb

D SELECT *
  FROM iceberg_scan(
      'https://blobs.duckdb.org/data/lineitem_iceberg/metadata/v1.metadata.json'
  );
```

l_orderkey int32	l_partkey int32	l_suppkey int32	…	l_shipmode varchar	l_comment varchar
1	1552	93	…	TRUCK	egular courts abov…
1	674	75	…	MAIL	ly final dependenc…
1	637	38	…	REG AIR	riously. regular, …
1	22	48	…	AIR	lites. fluffily ev…

Lakehouse architecture

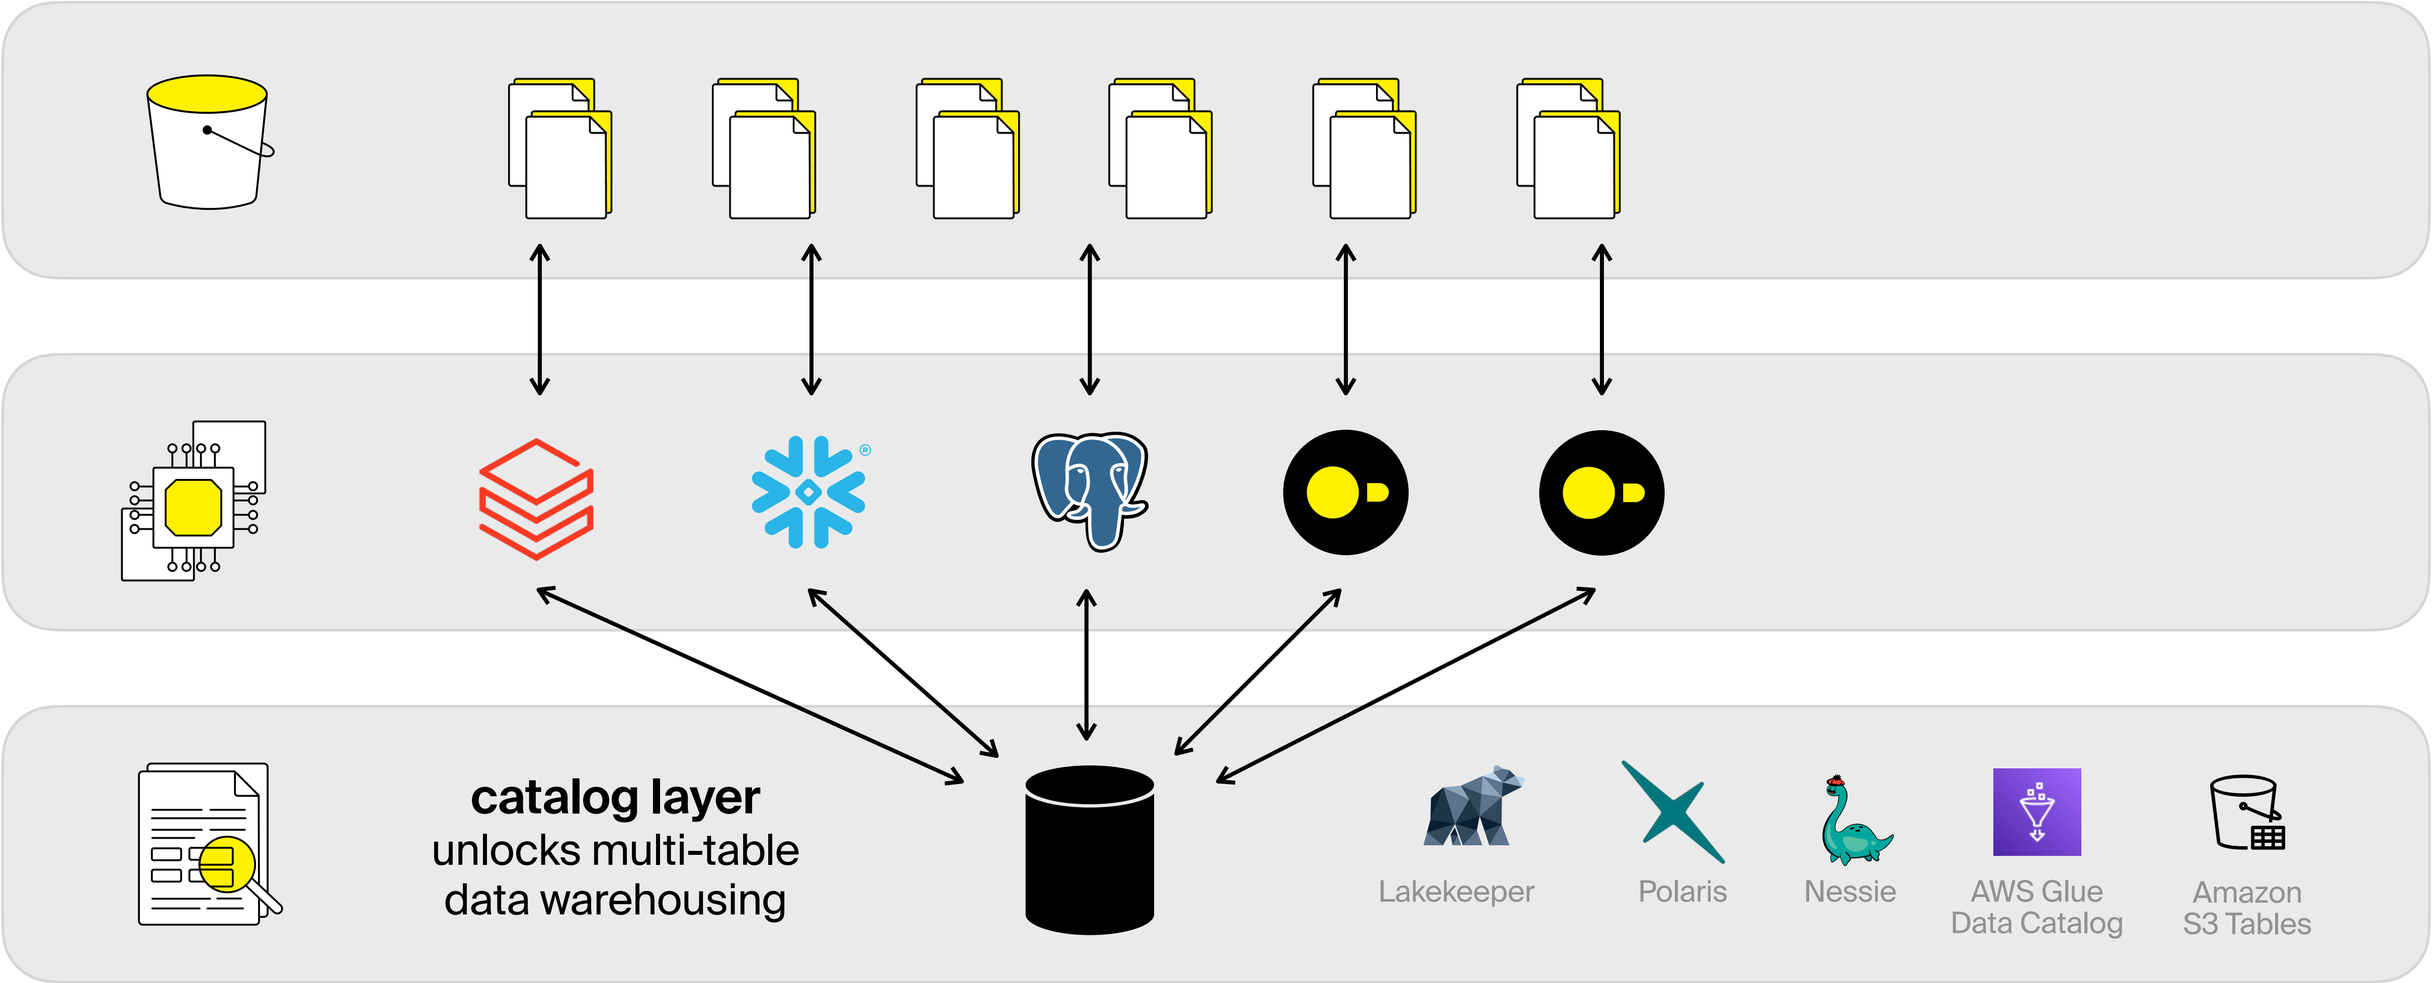

Iceberg REST Catalog support in DuckDB

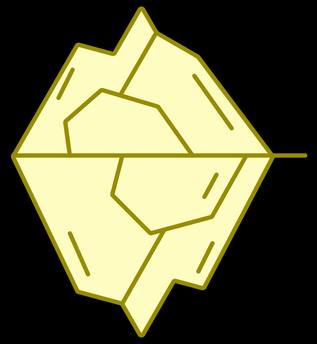

```
D CREATE SECRET (
    TYPE s3,
    PROVIDER credential_chain
  );

D ATTACH 's3_tables_arn' AS s3_tables (
    TYPE iceberg,
    ENDPOINT_TYPE s3_tables
  );

D INSERT INTO s3_tables.my_ns.my_table
    VALUES ('hello world', 42);
```

DuckLake

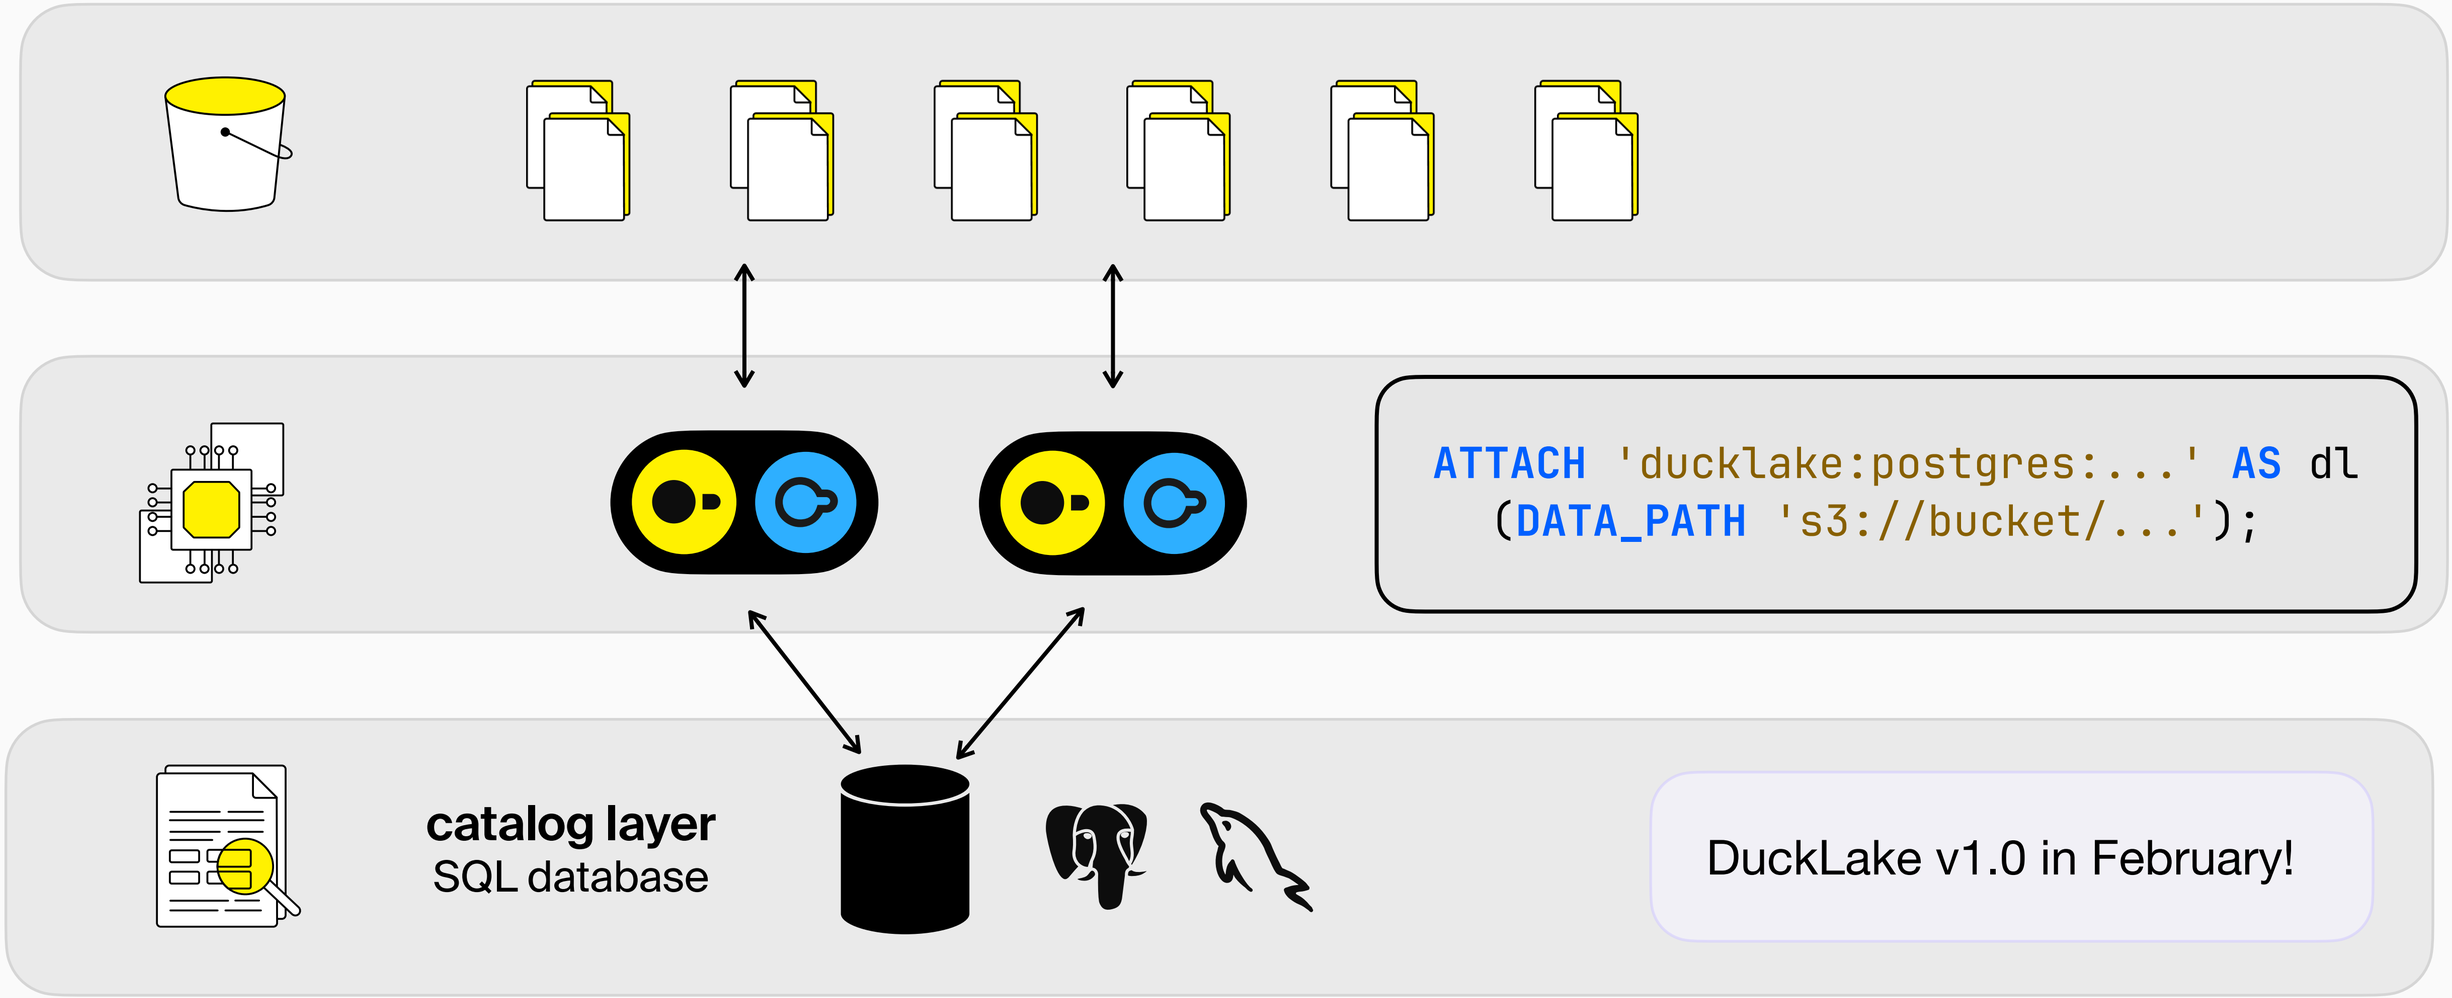


```
ATTACH 'ducklake:postgres:...' AS dl
  (DATA_PATH 's3://bucket/...');
```
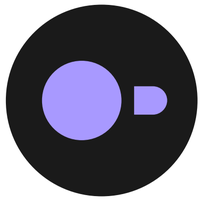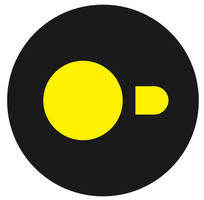
DuckLake v1.0 in February!