

# DuckLake

## The SQL-Powered Lakehouse Format



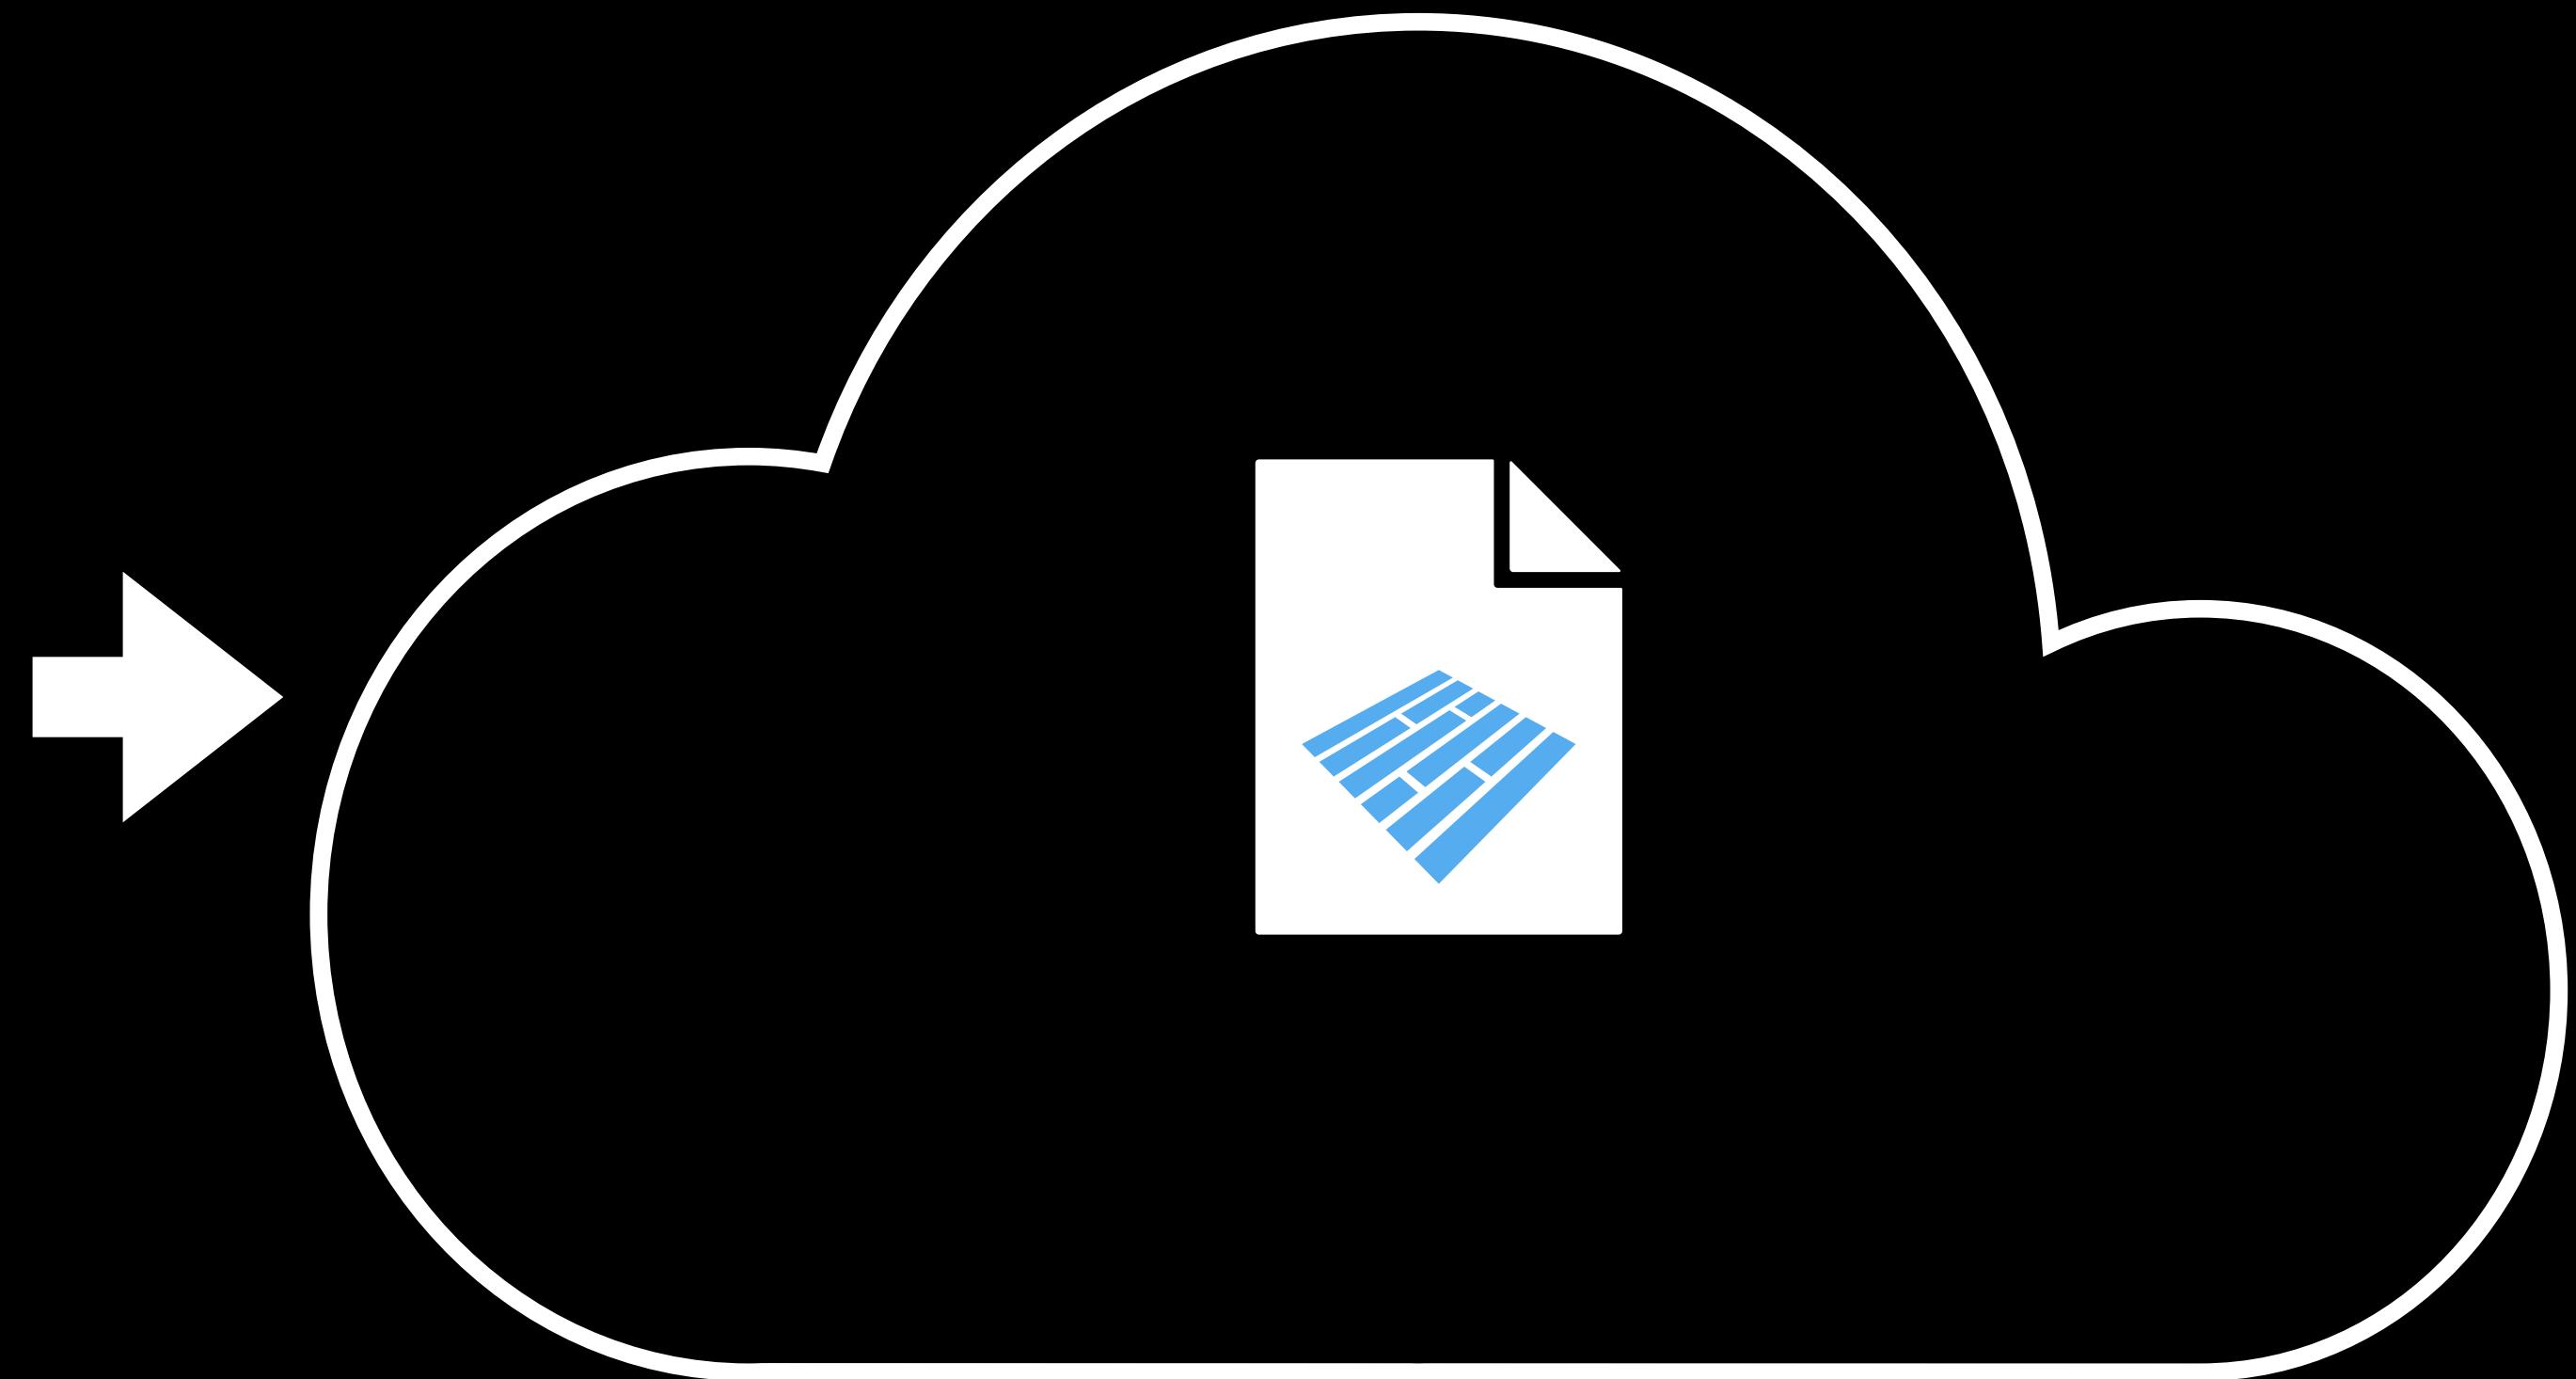


Name	Speed	Tonnage



v1

Name	Speed	Tonnage



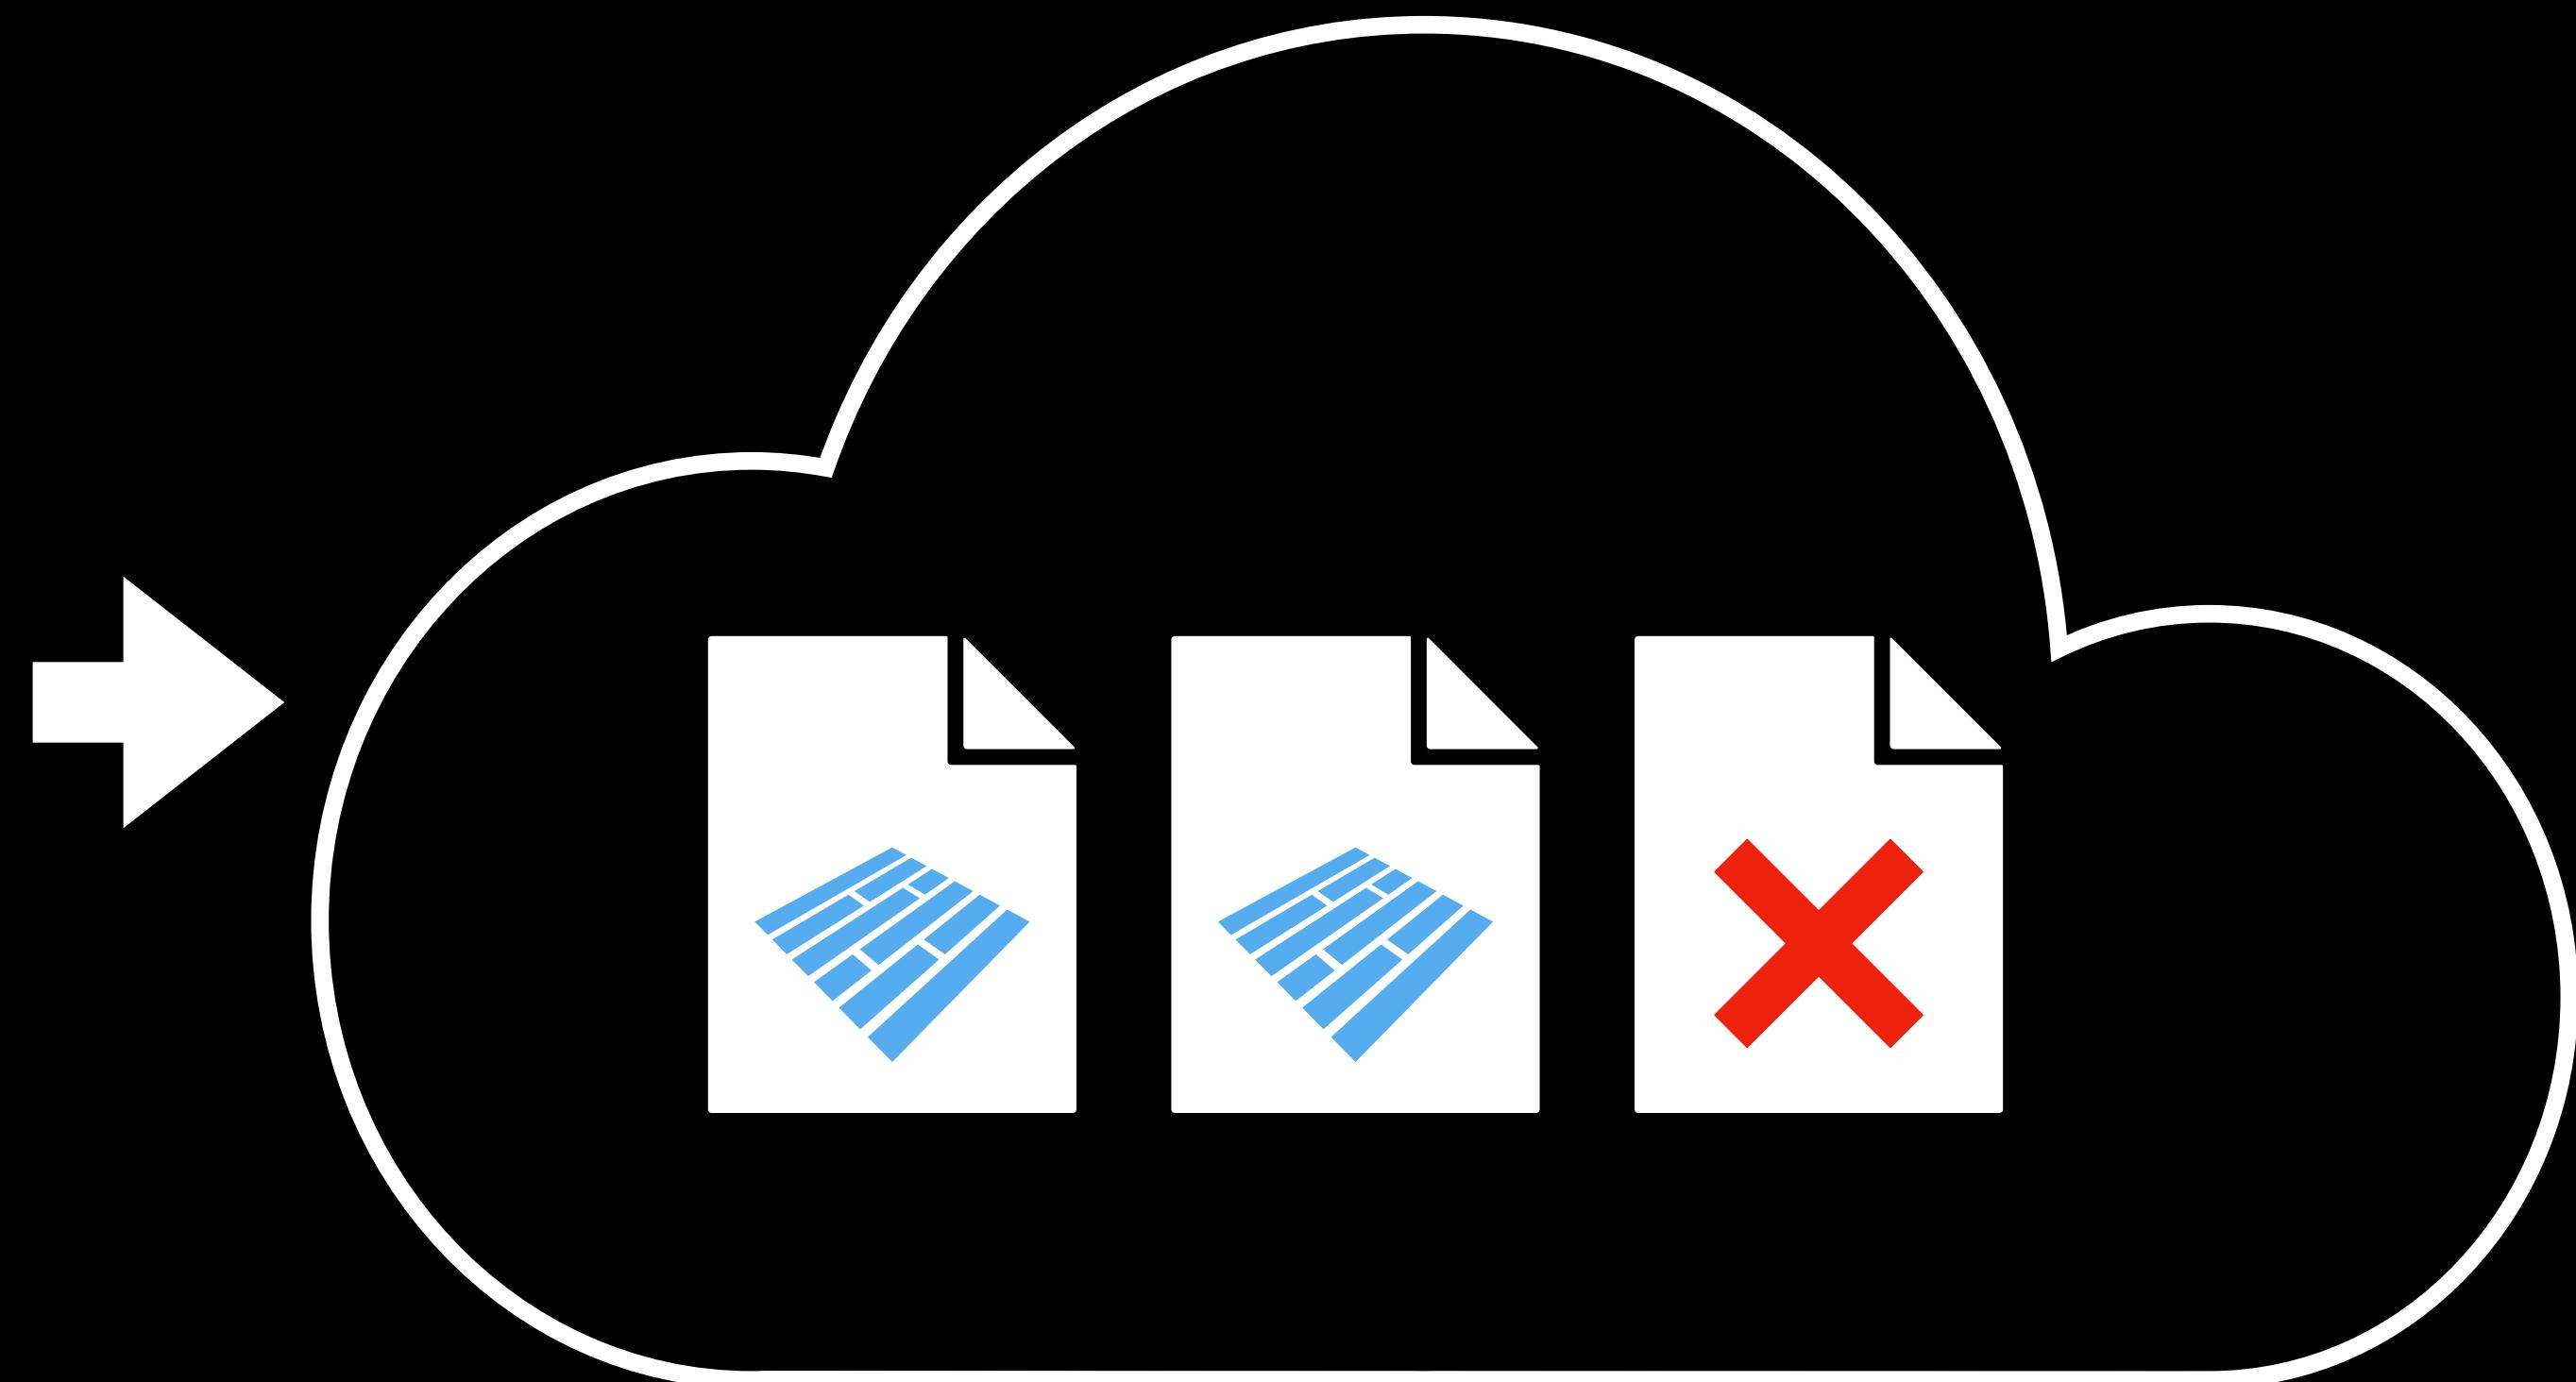
v2

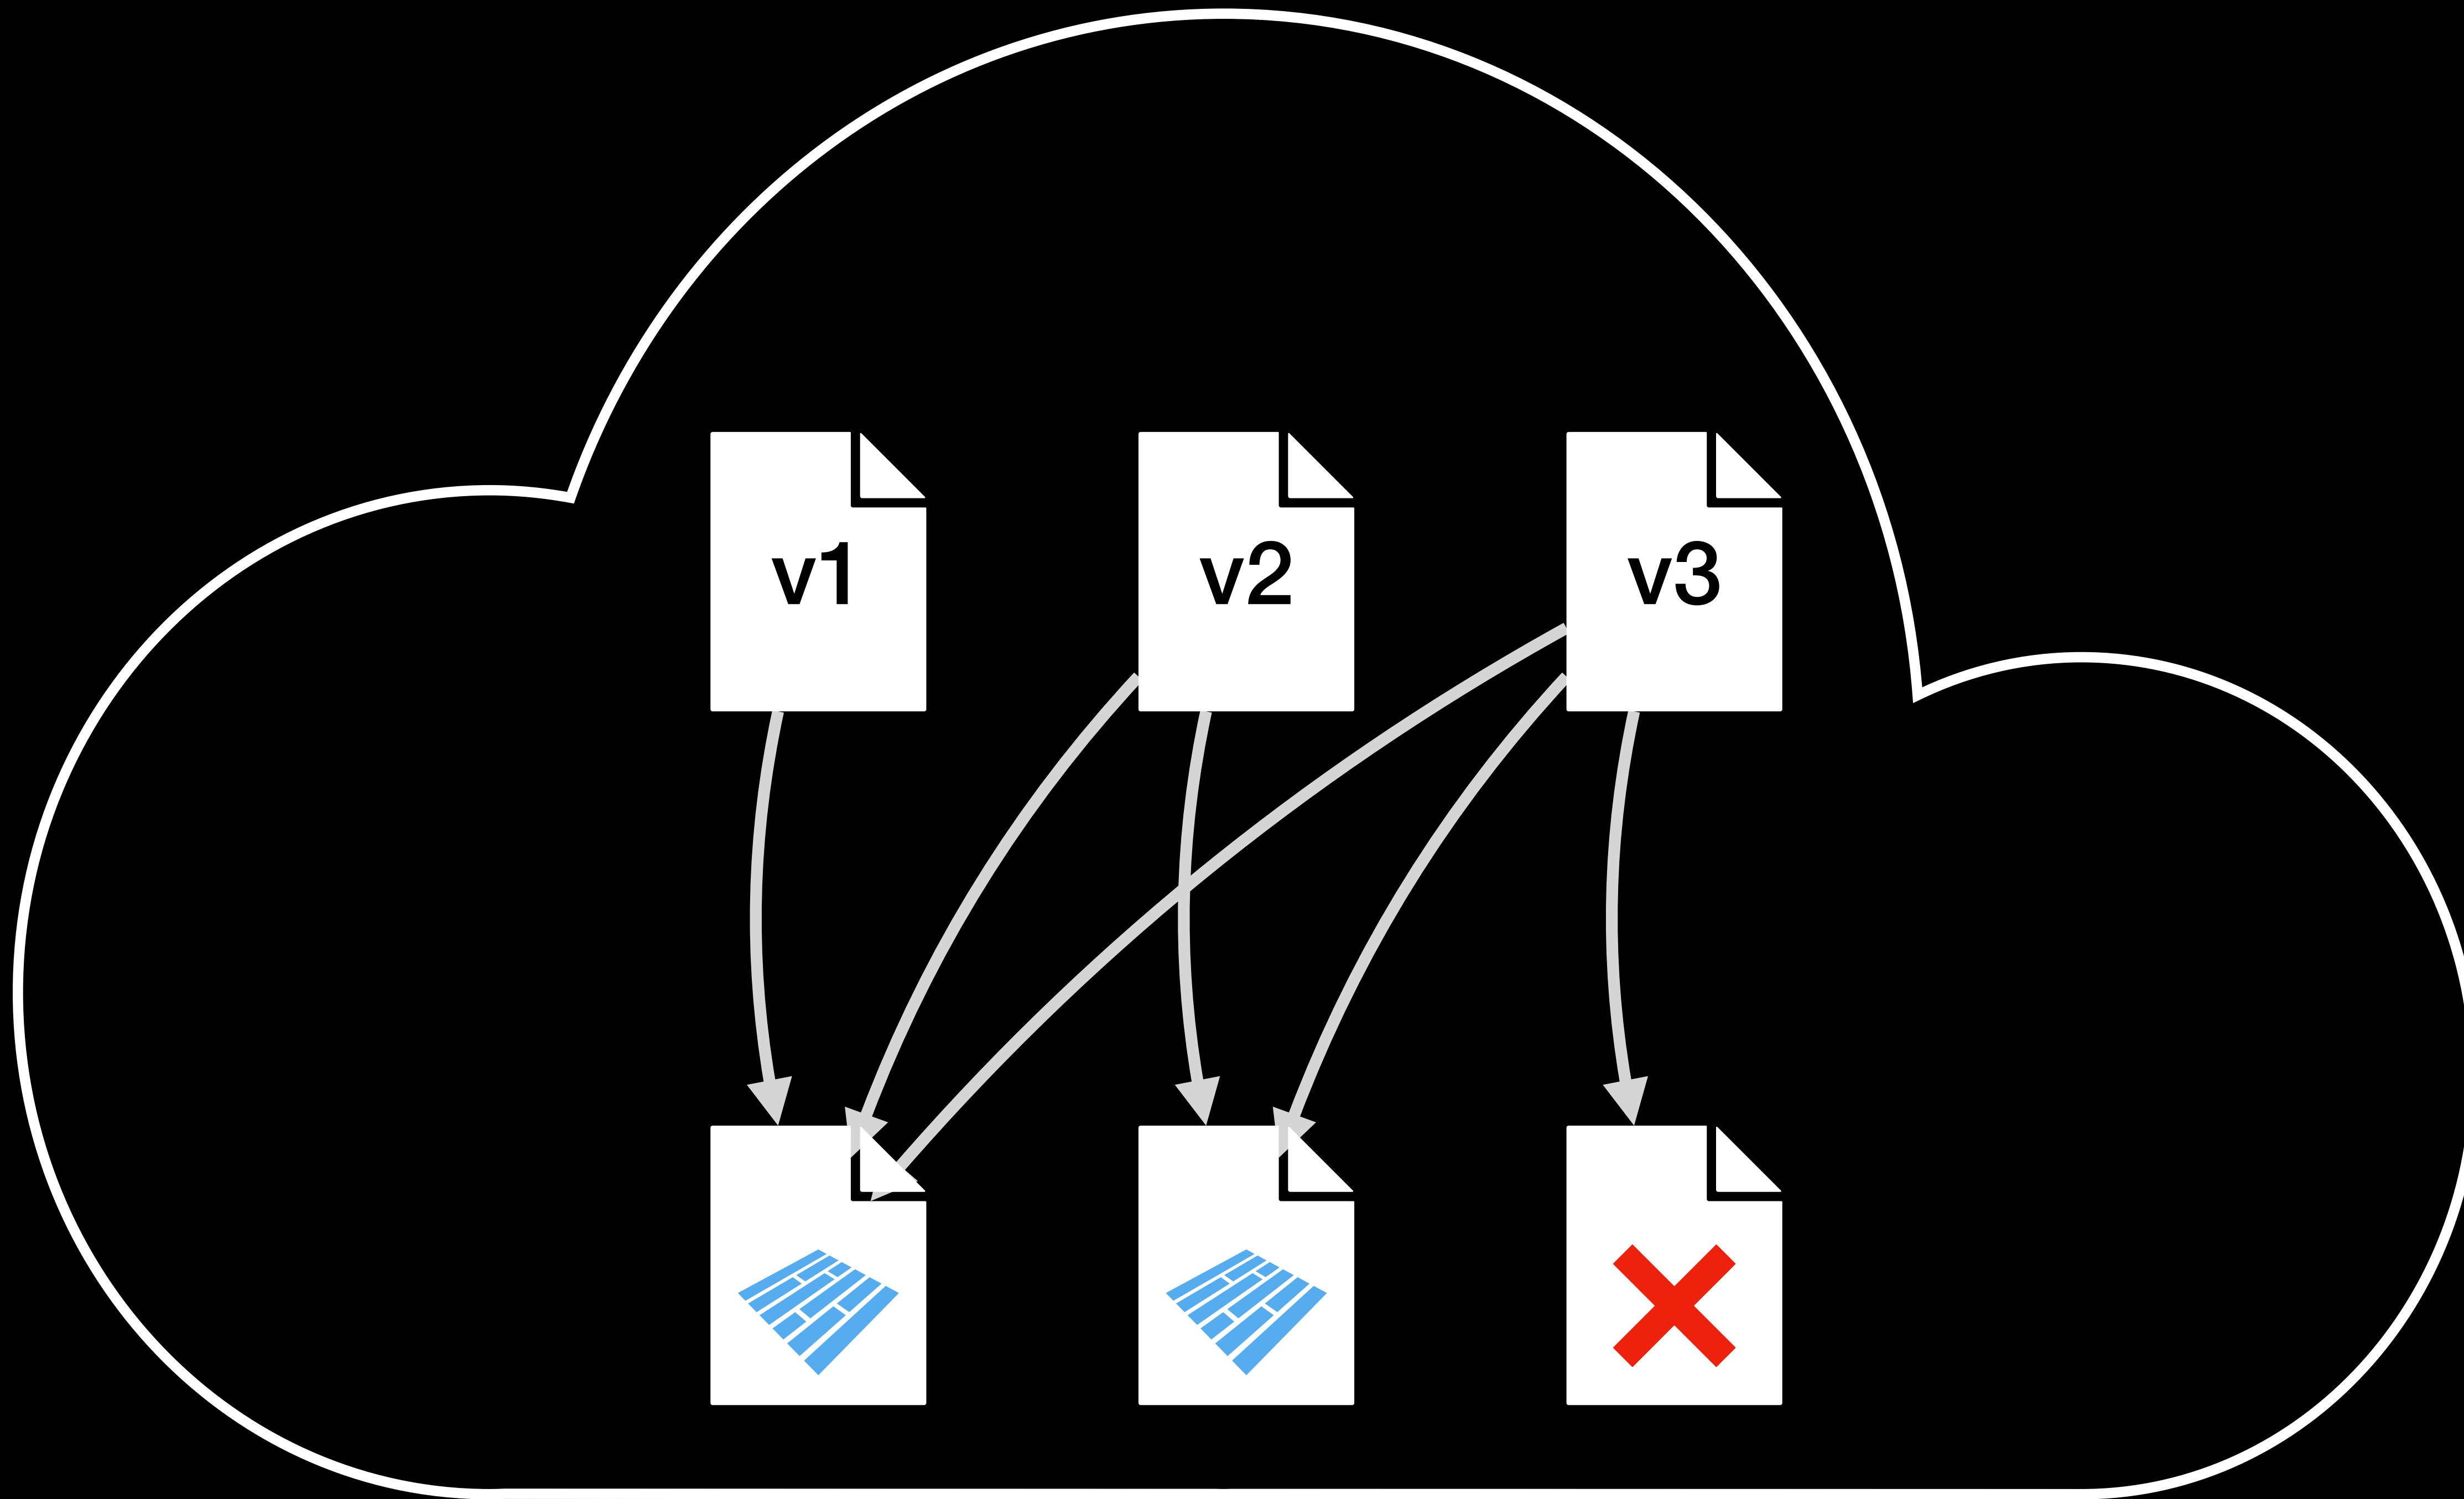
Name	Speed	Tonnage

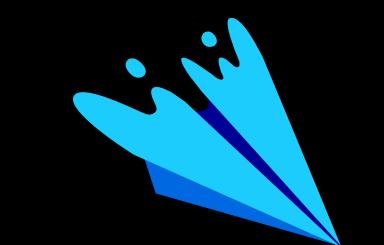
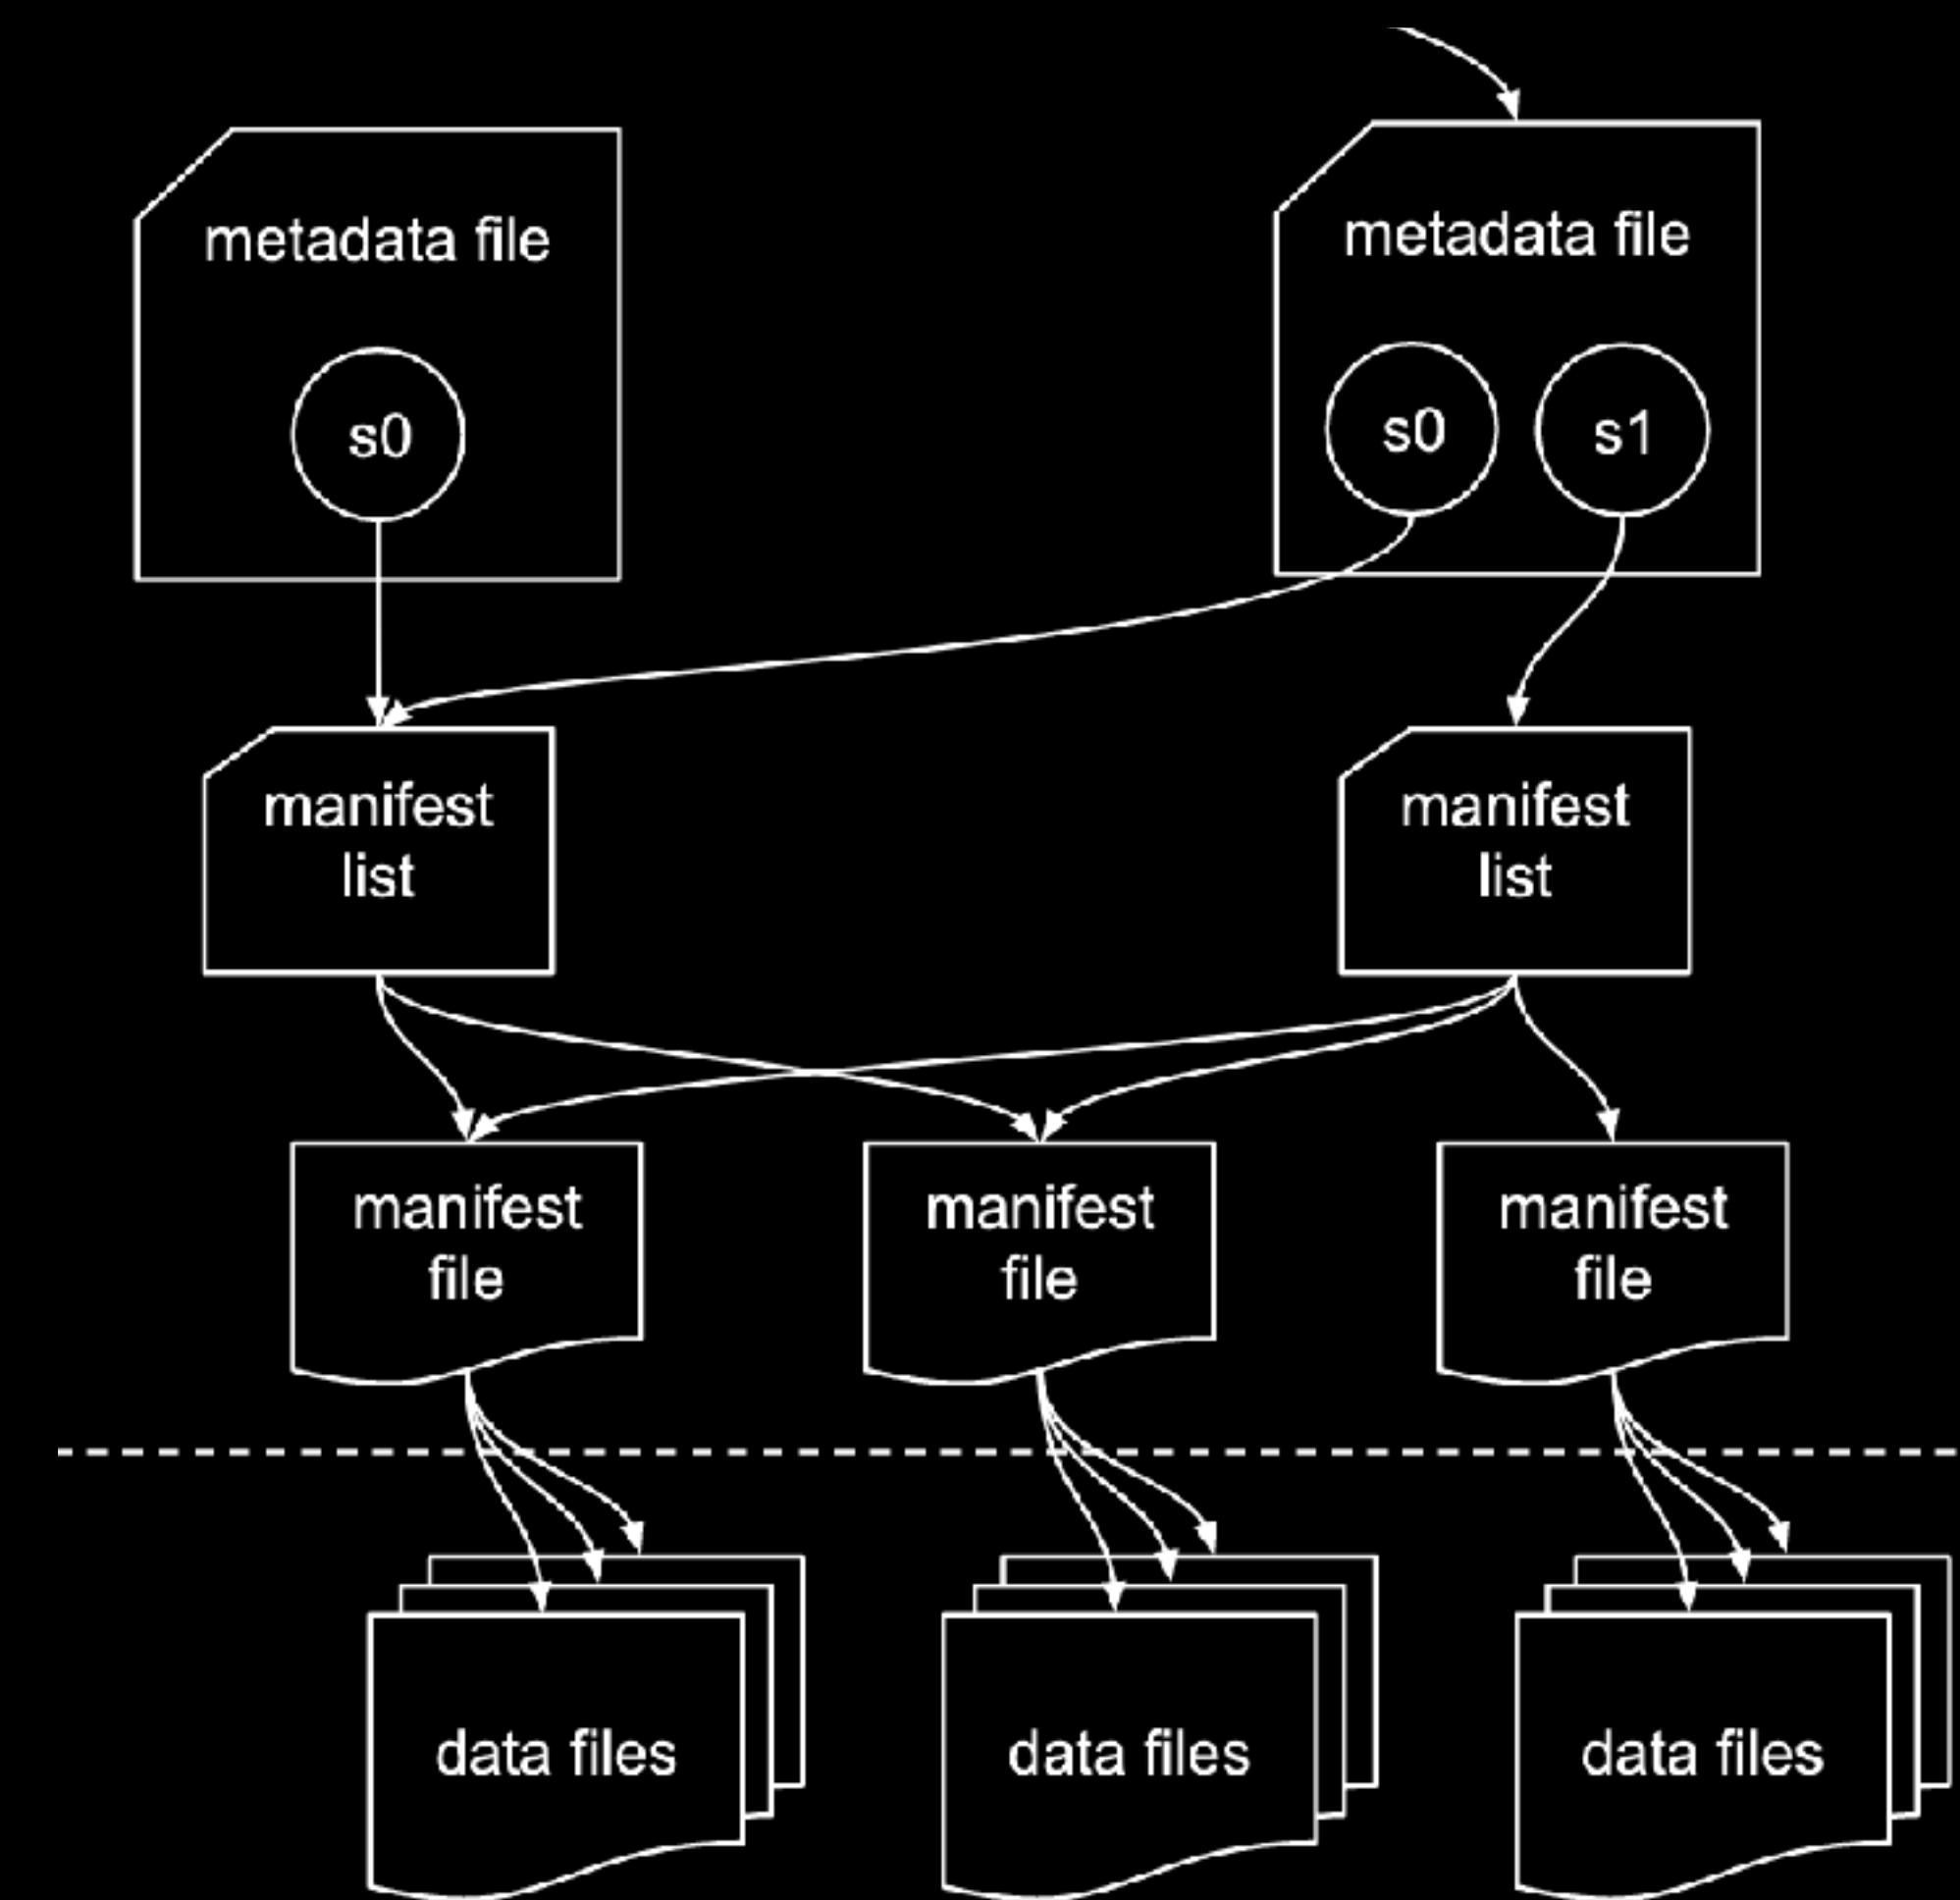


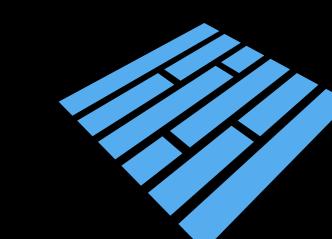
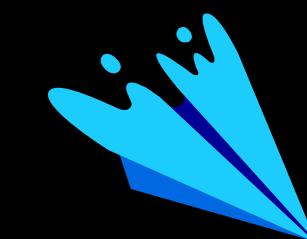
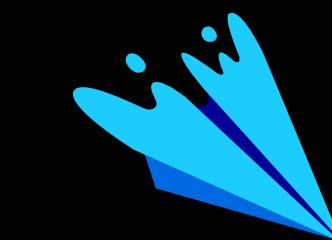
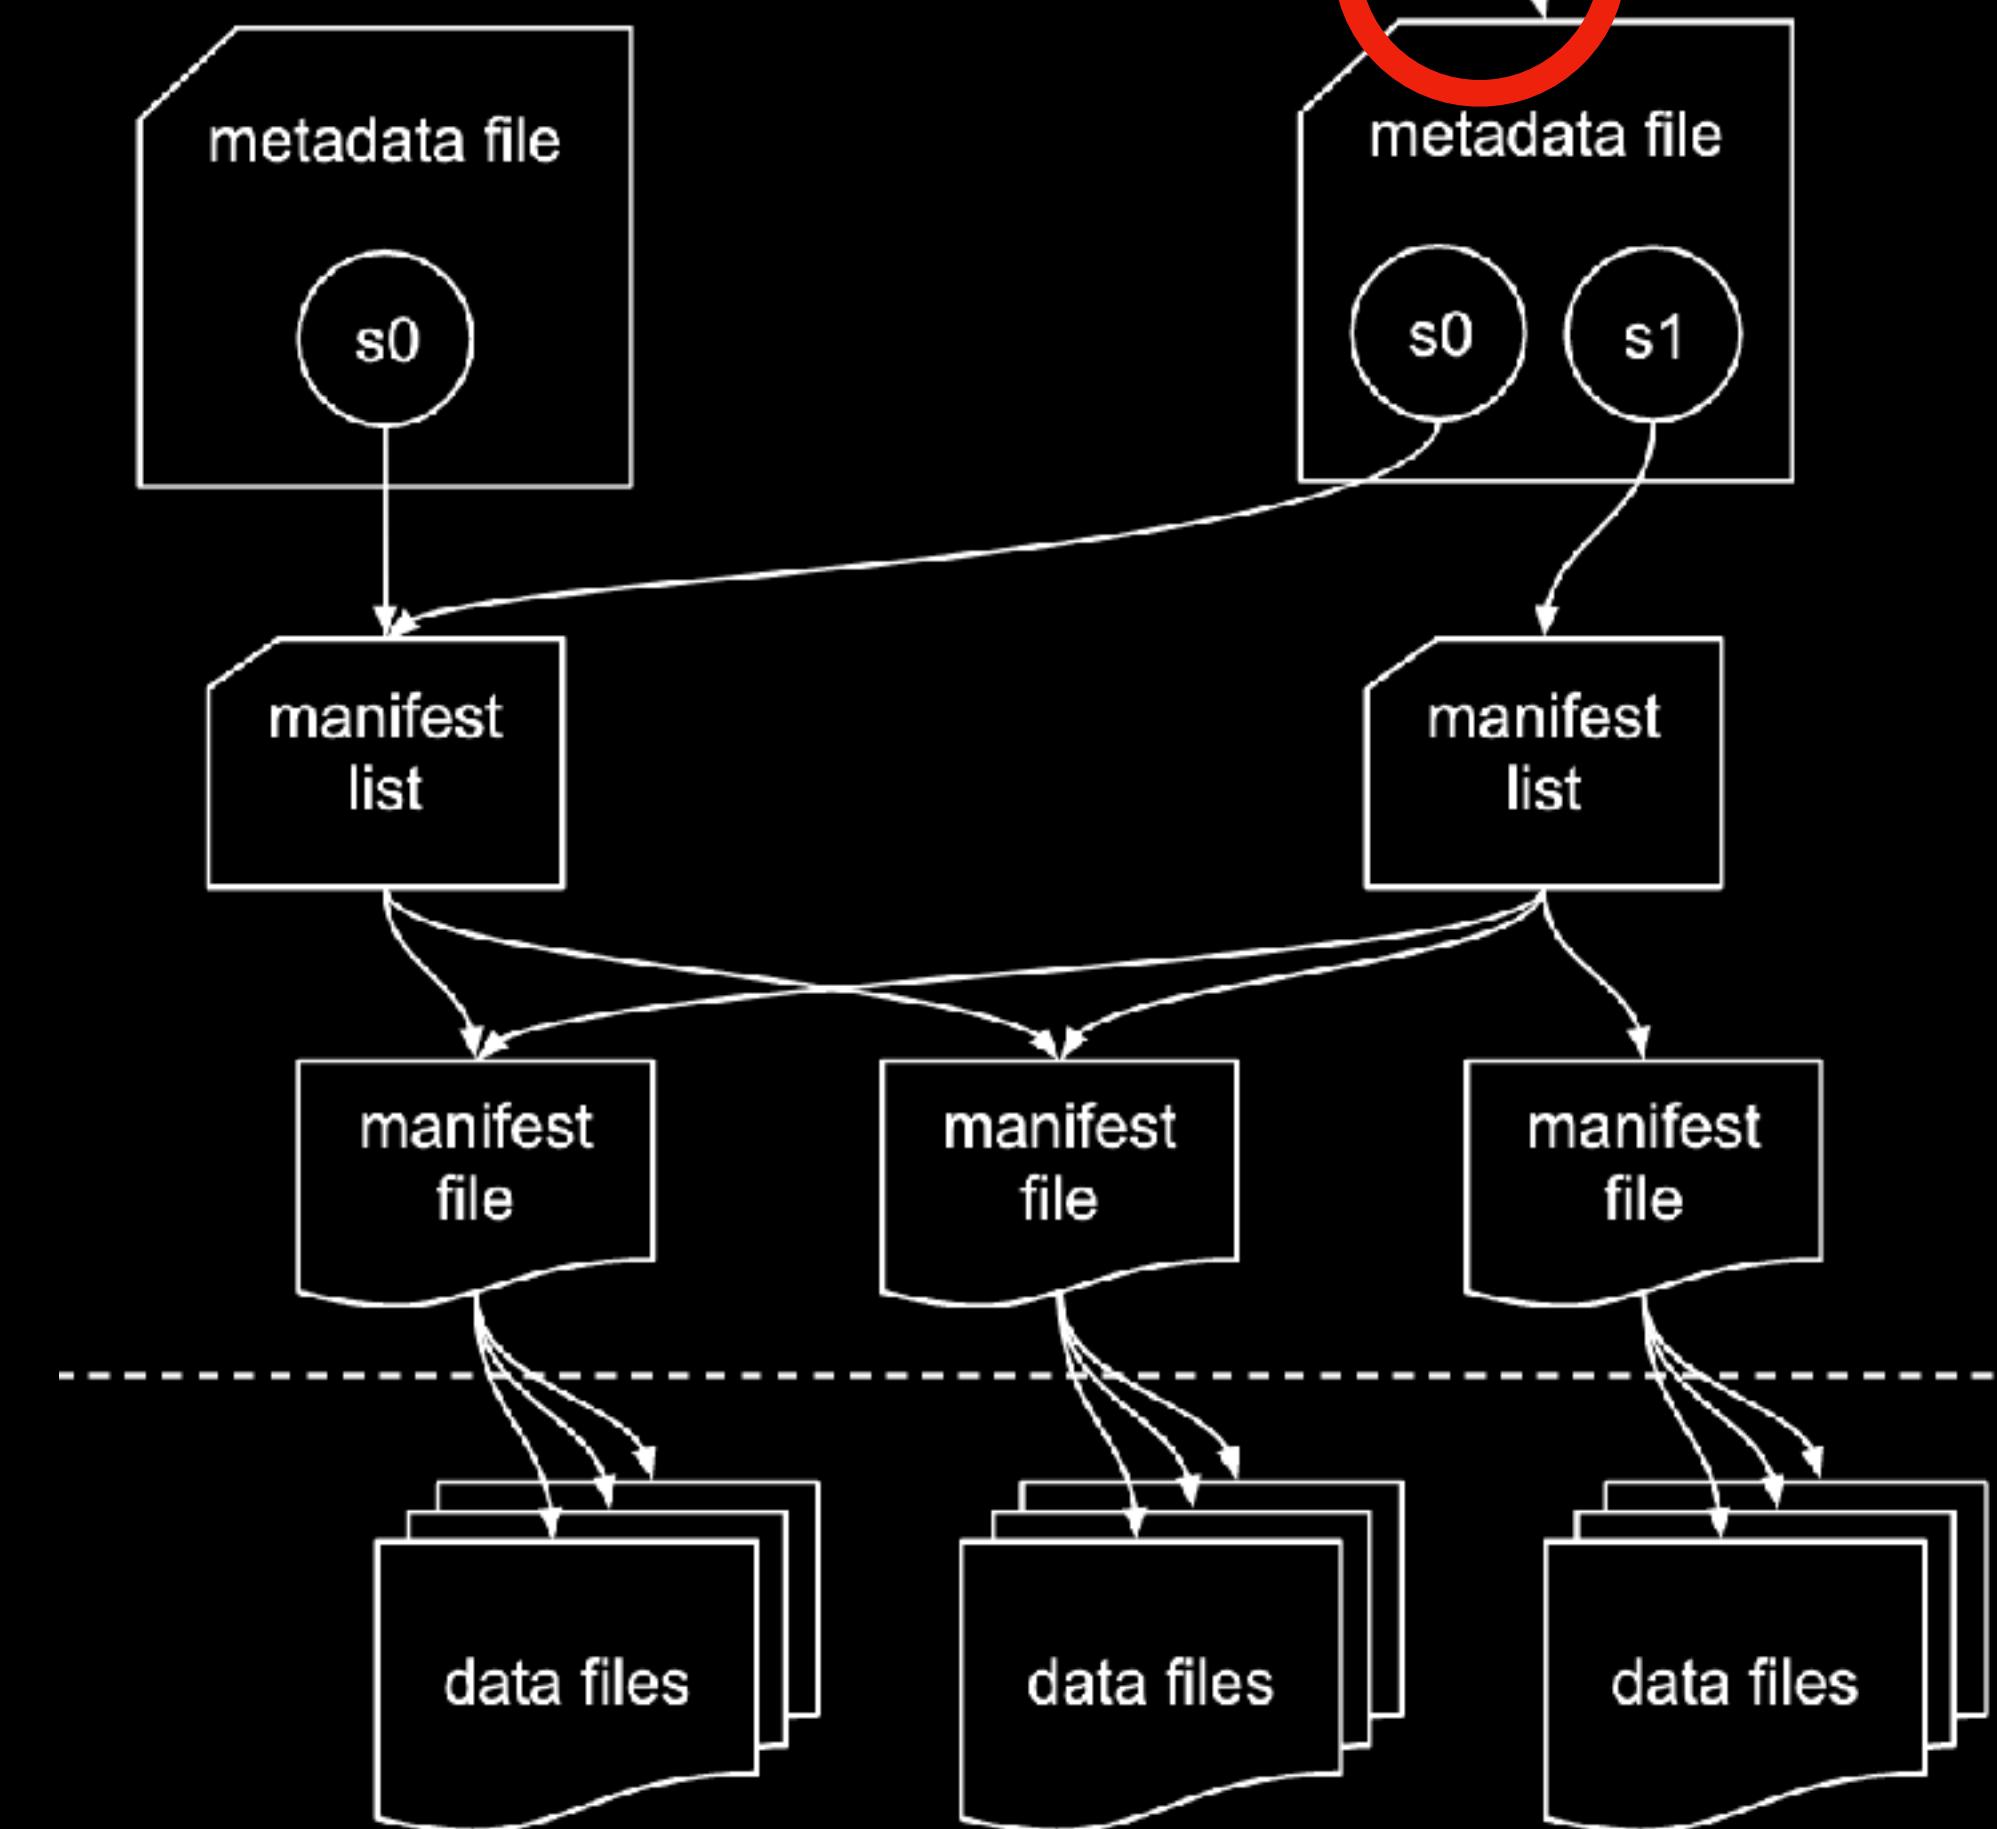
v3

Name	Speed	Tonnage

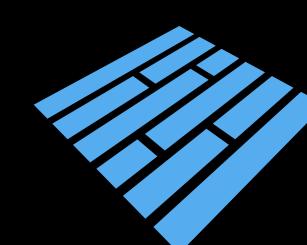
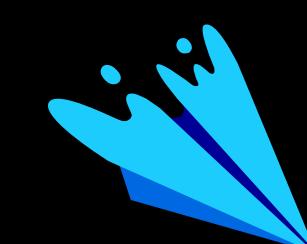
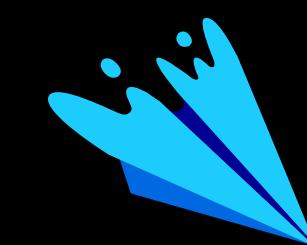
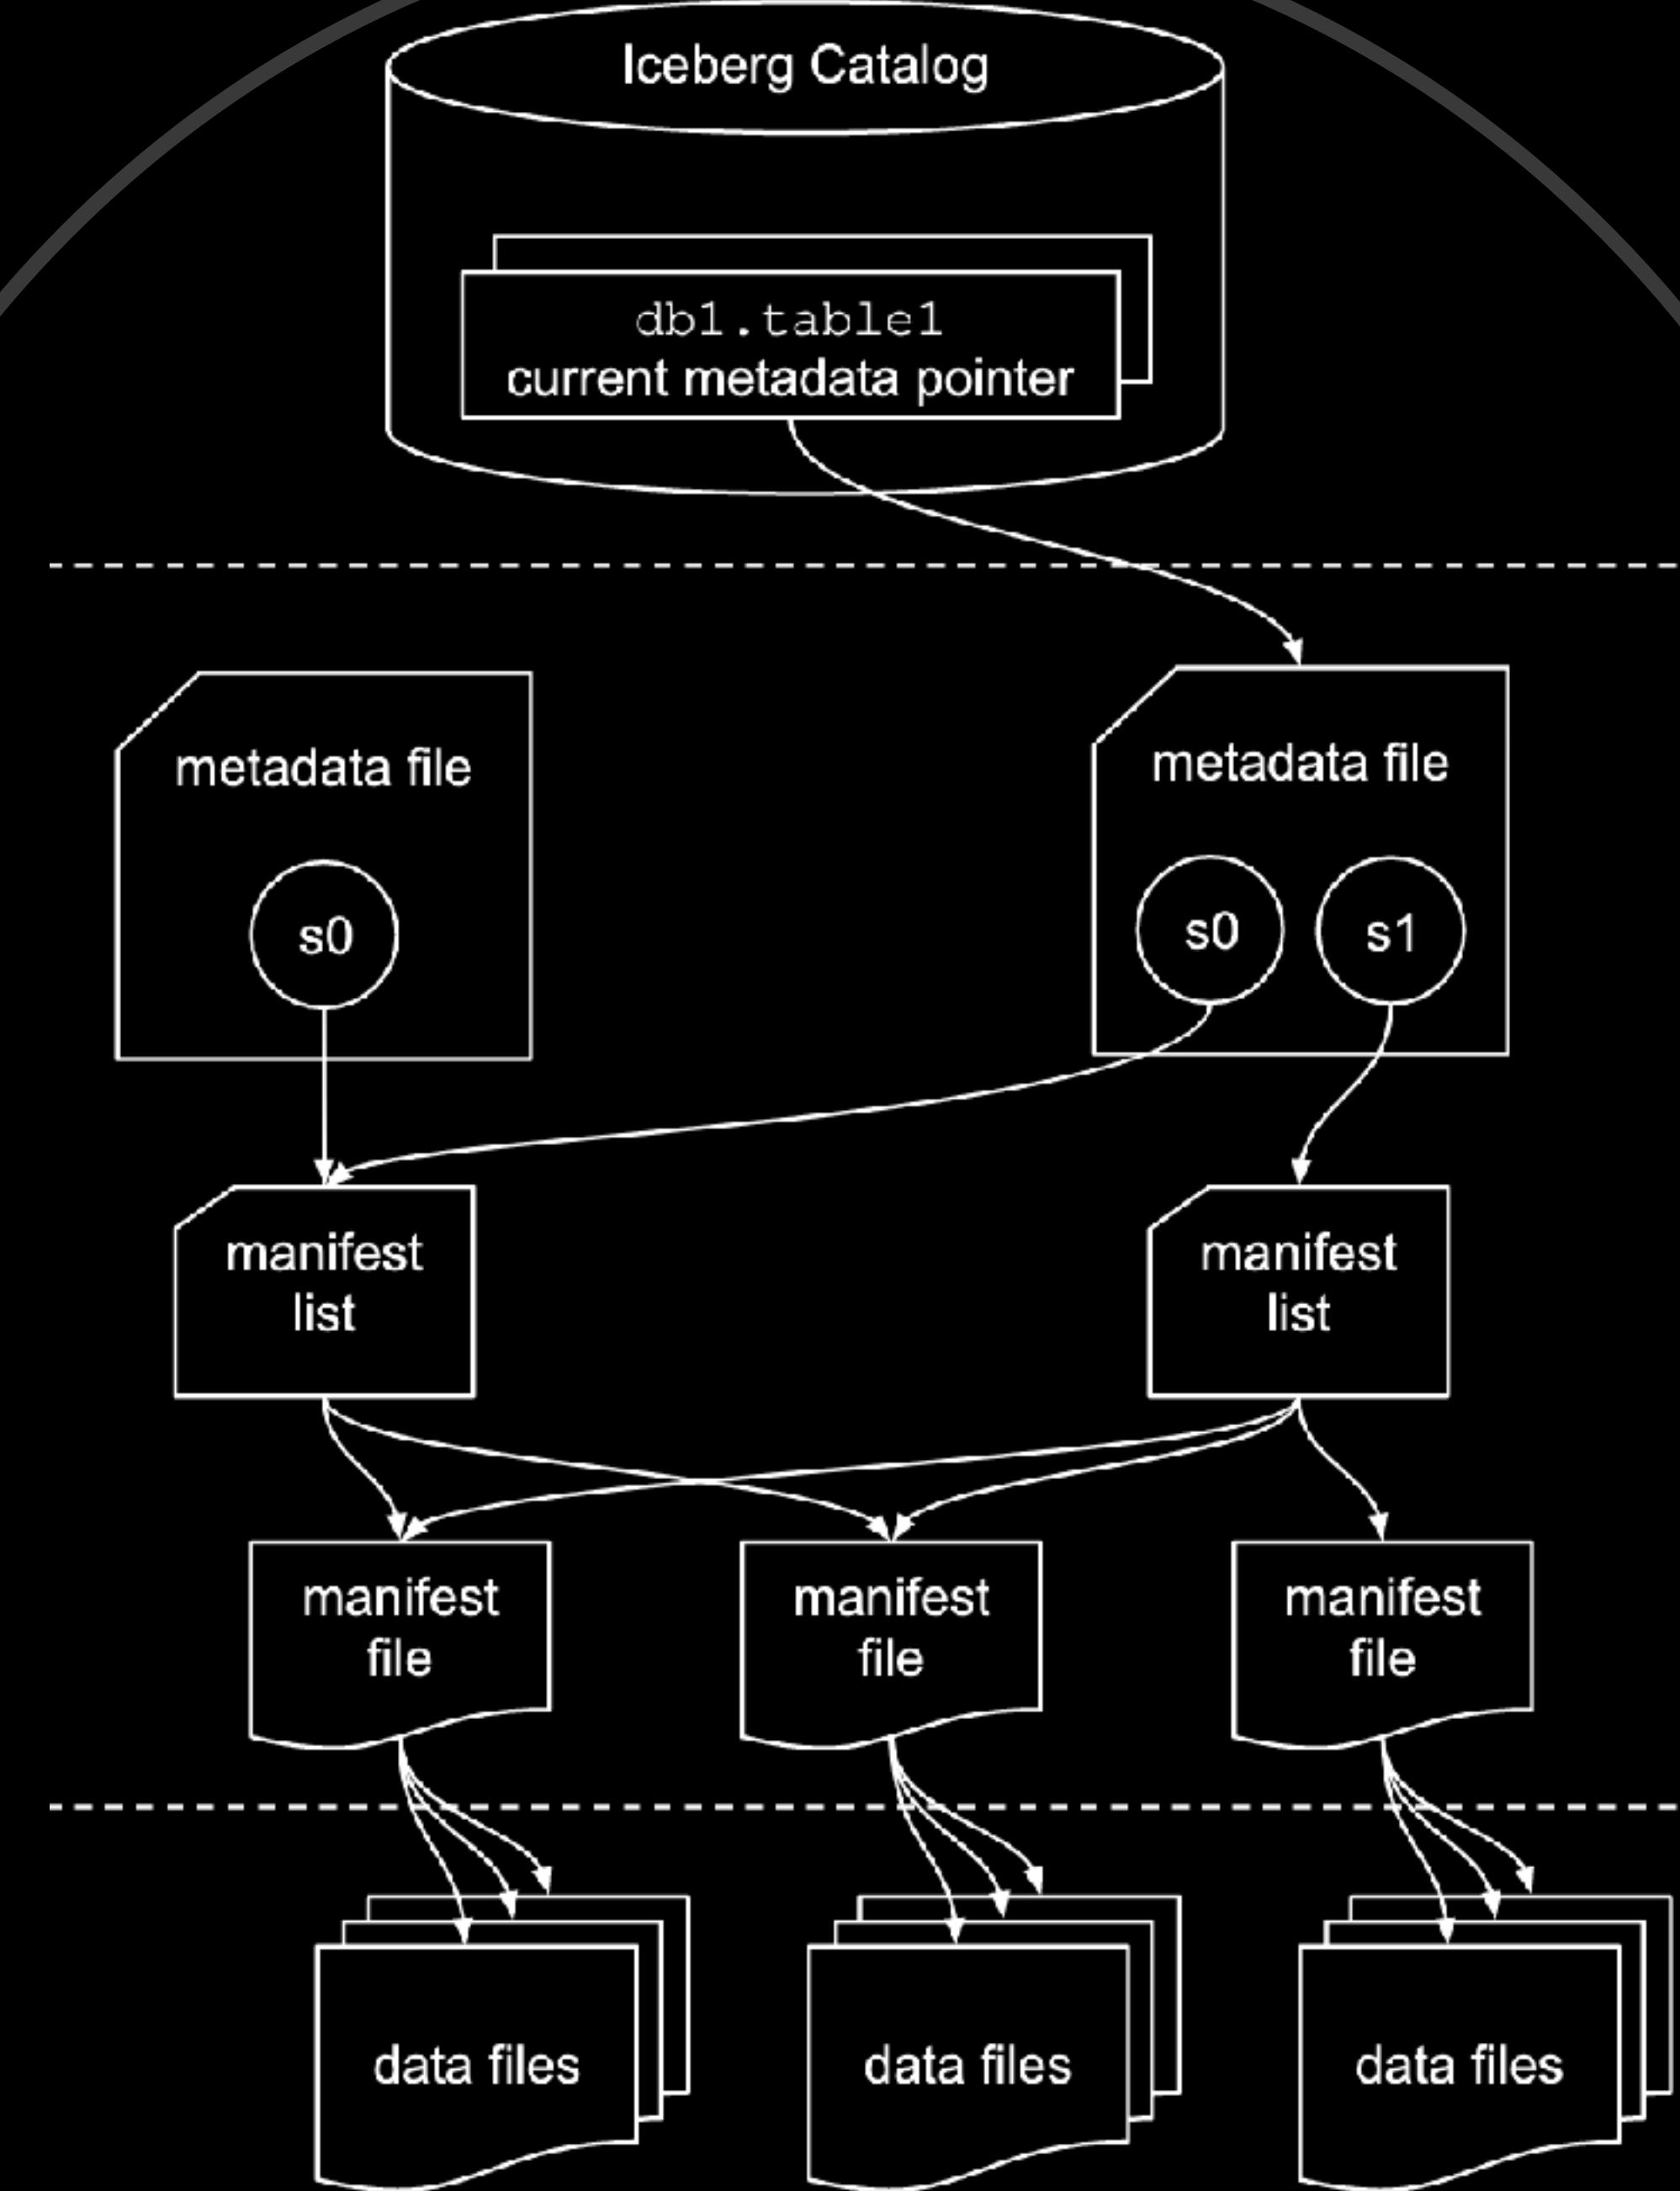




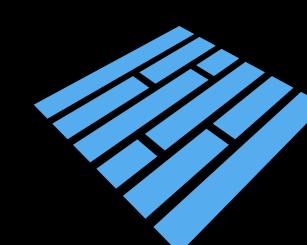
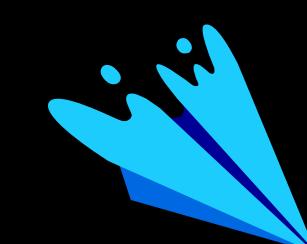
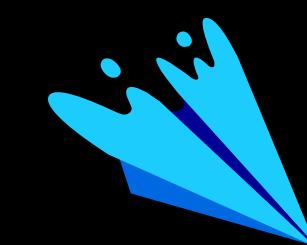
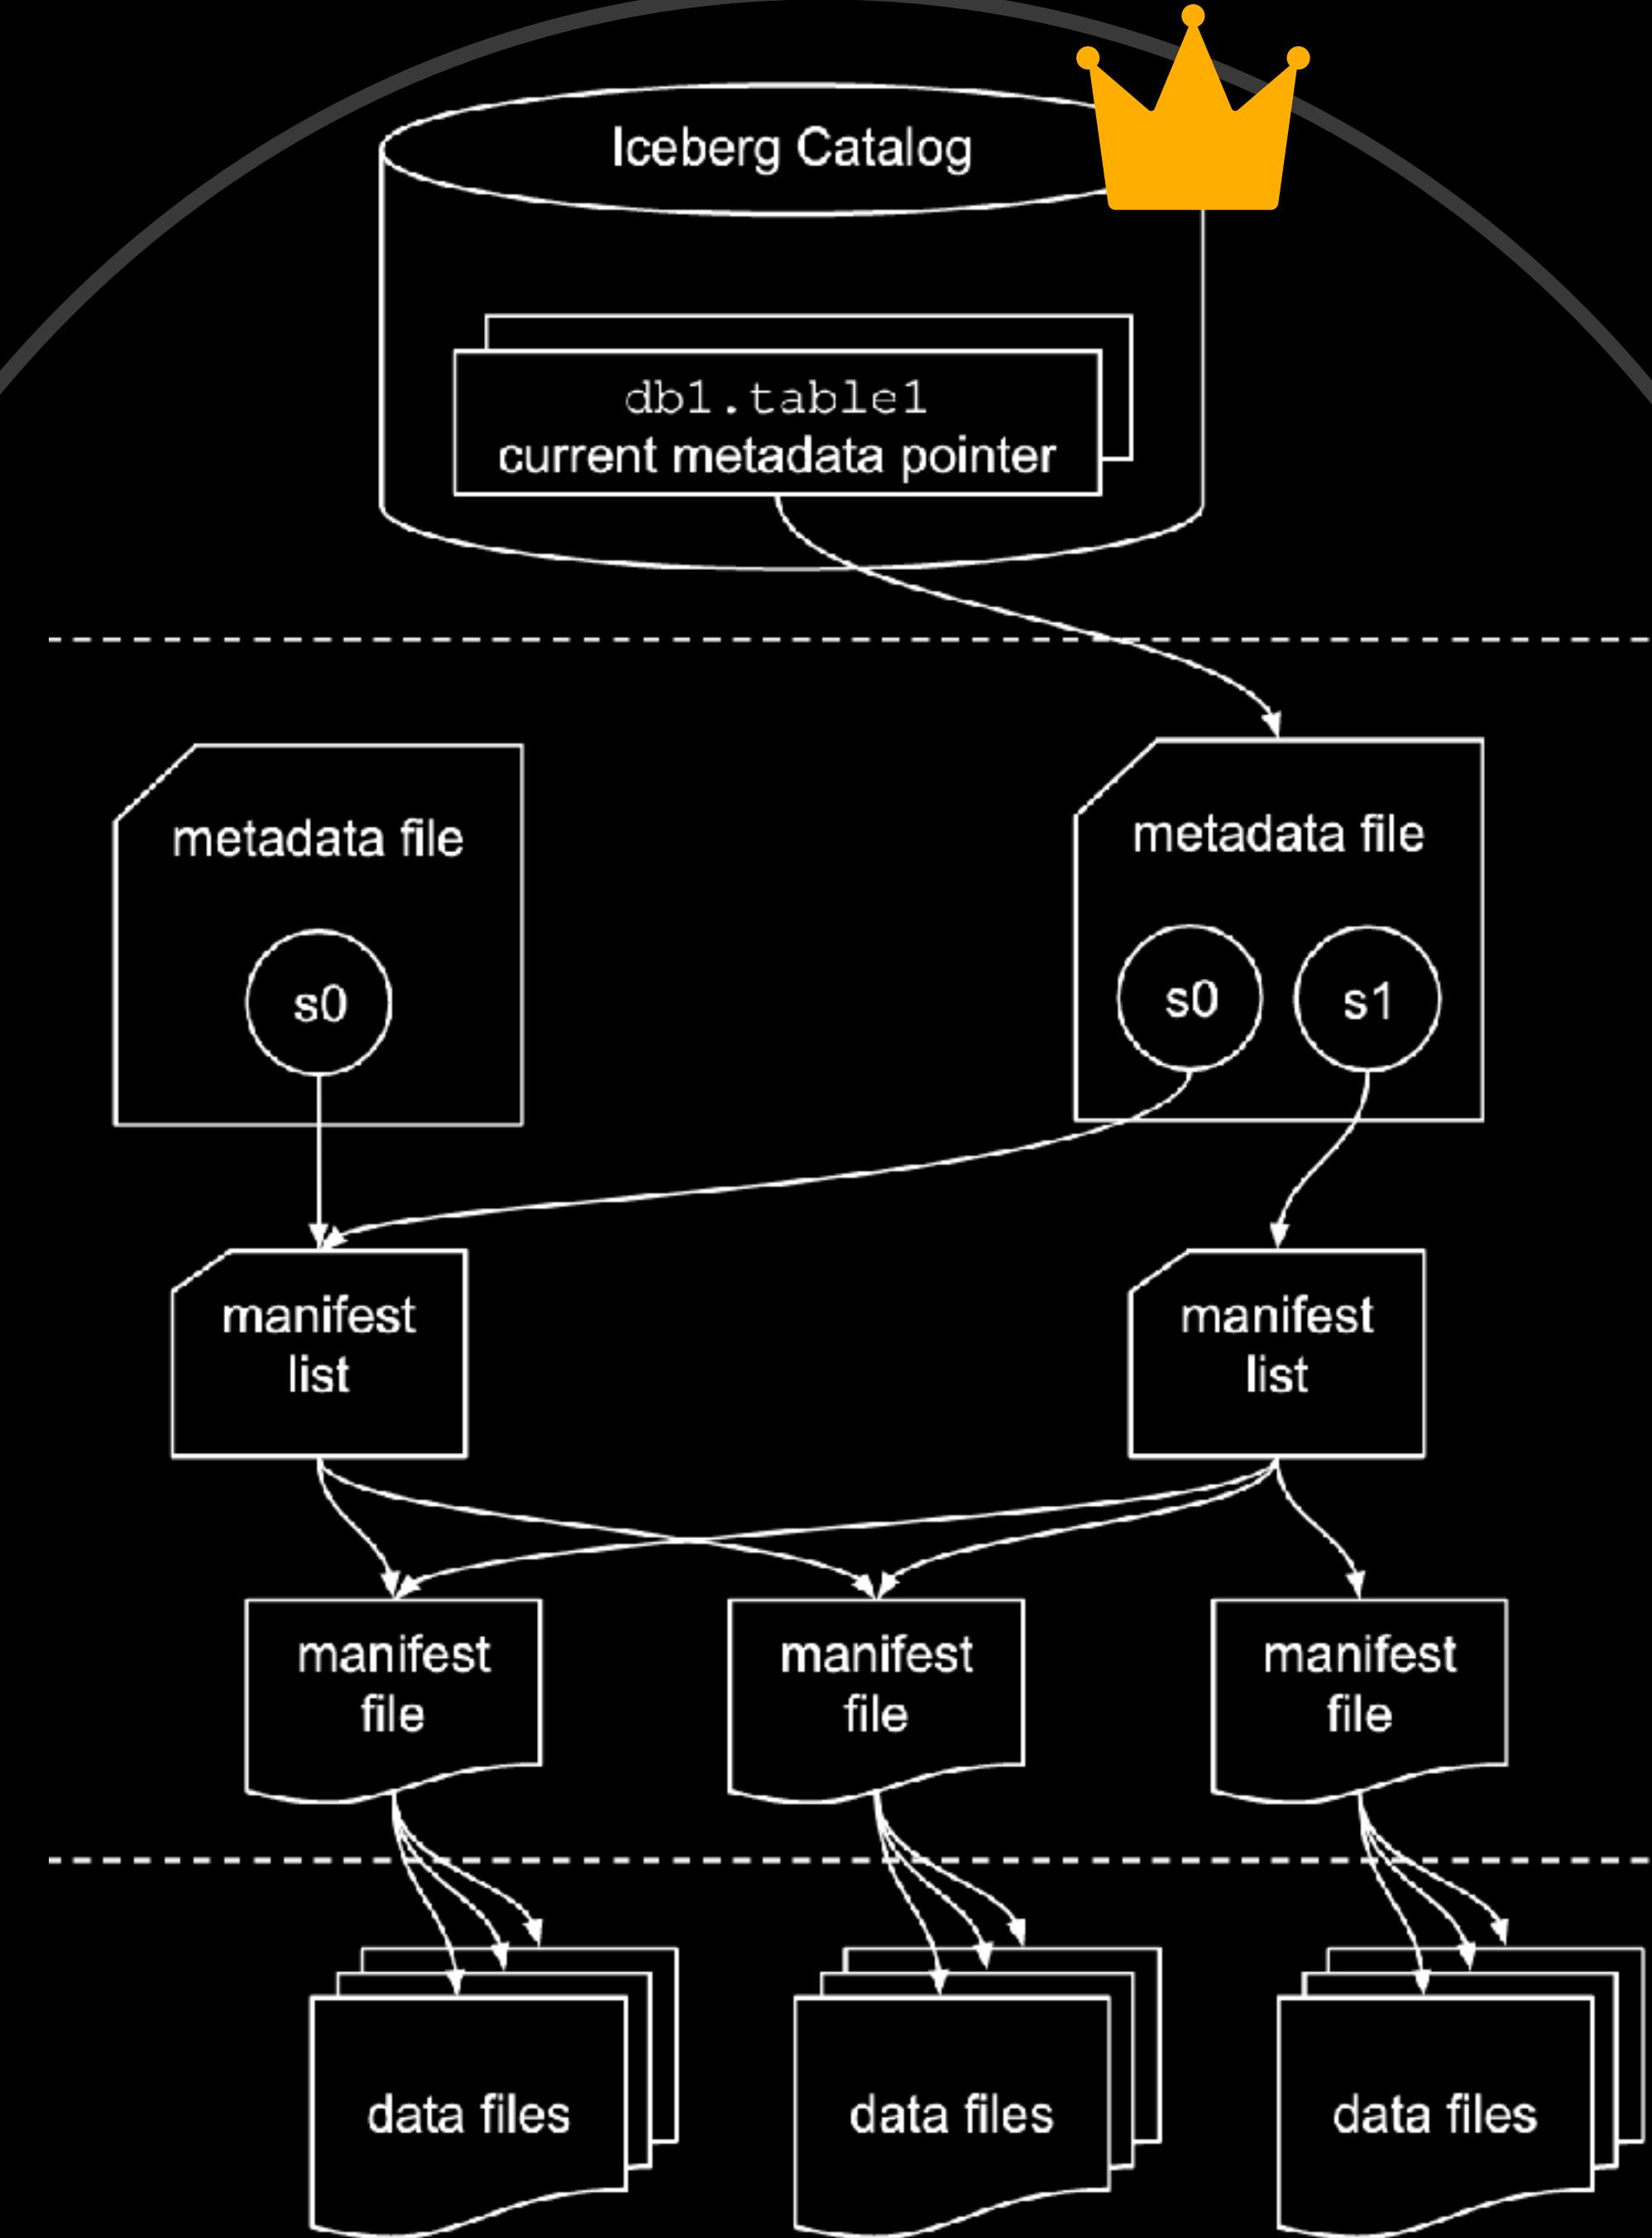


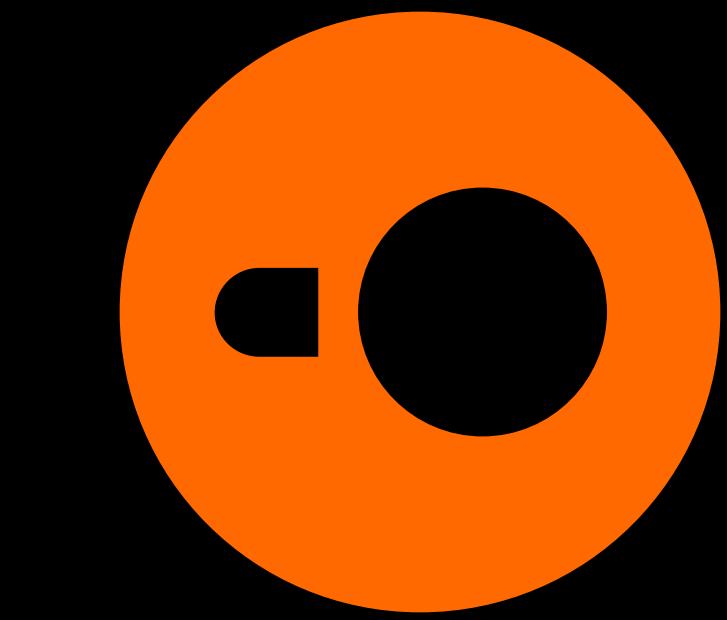


Java  
REST  
Postgres



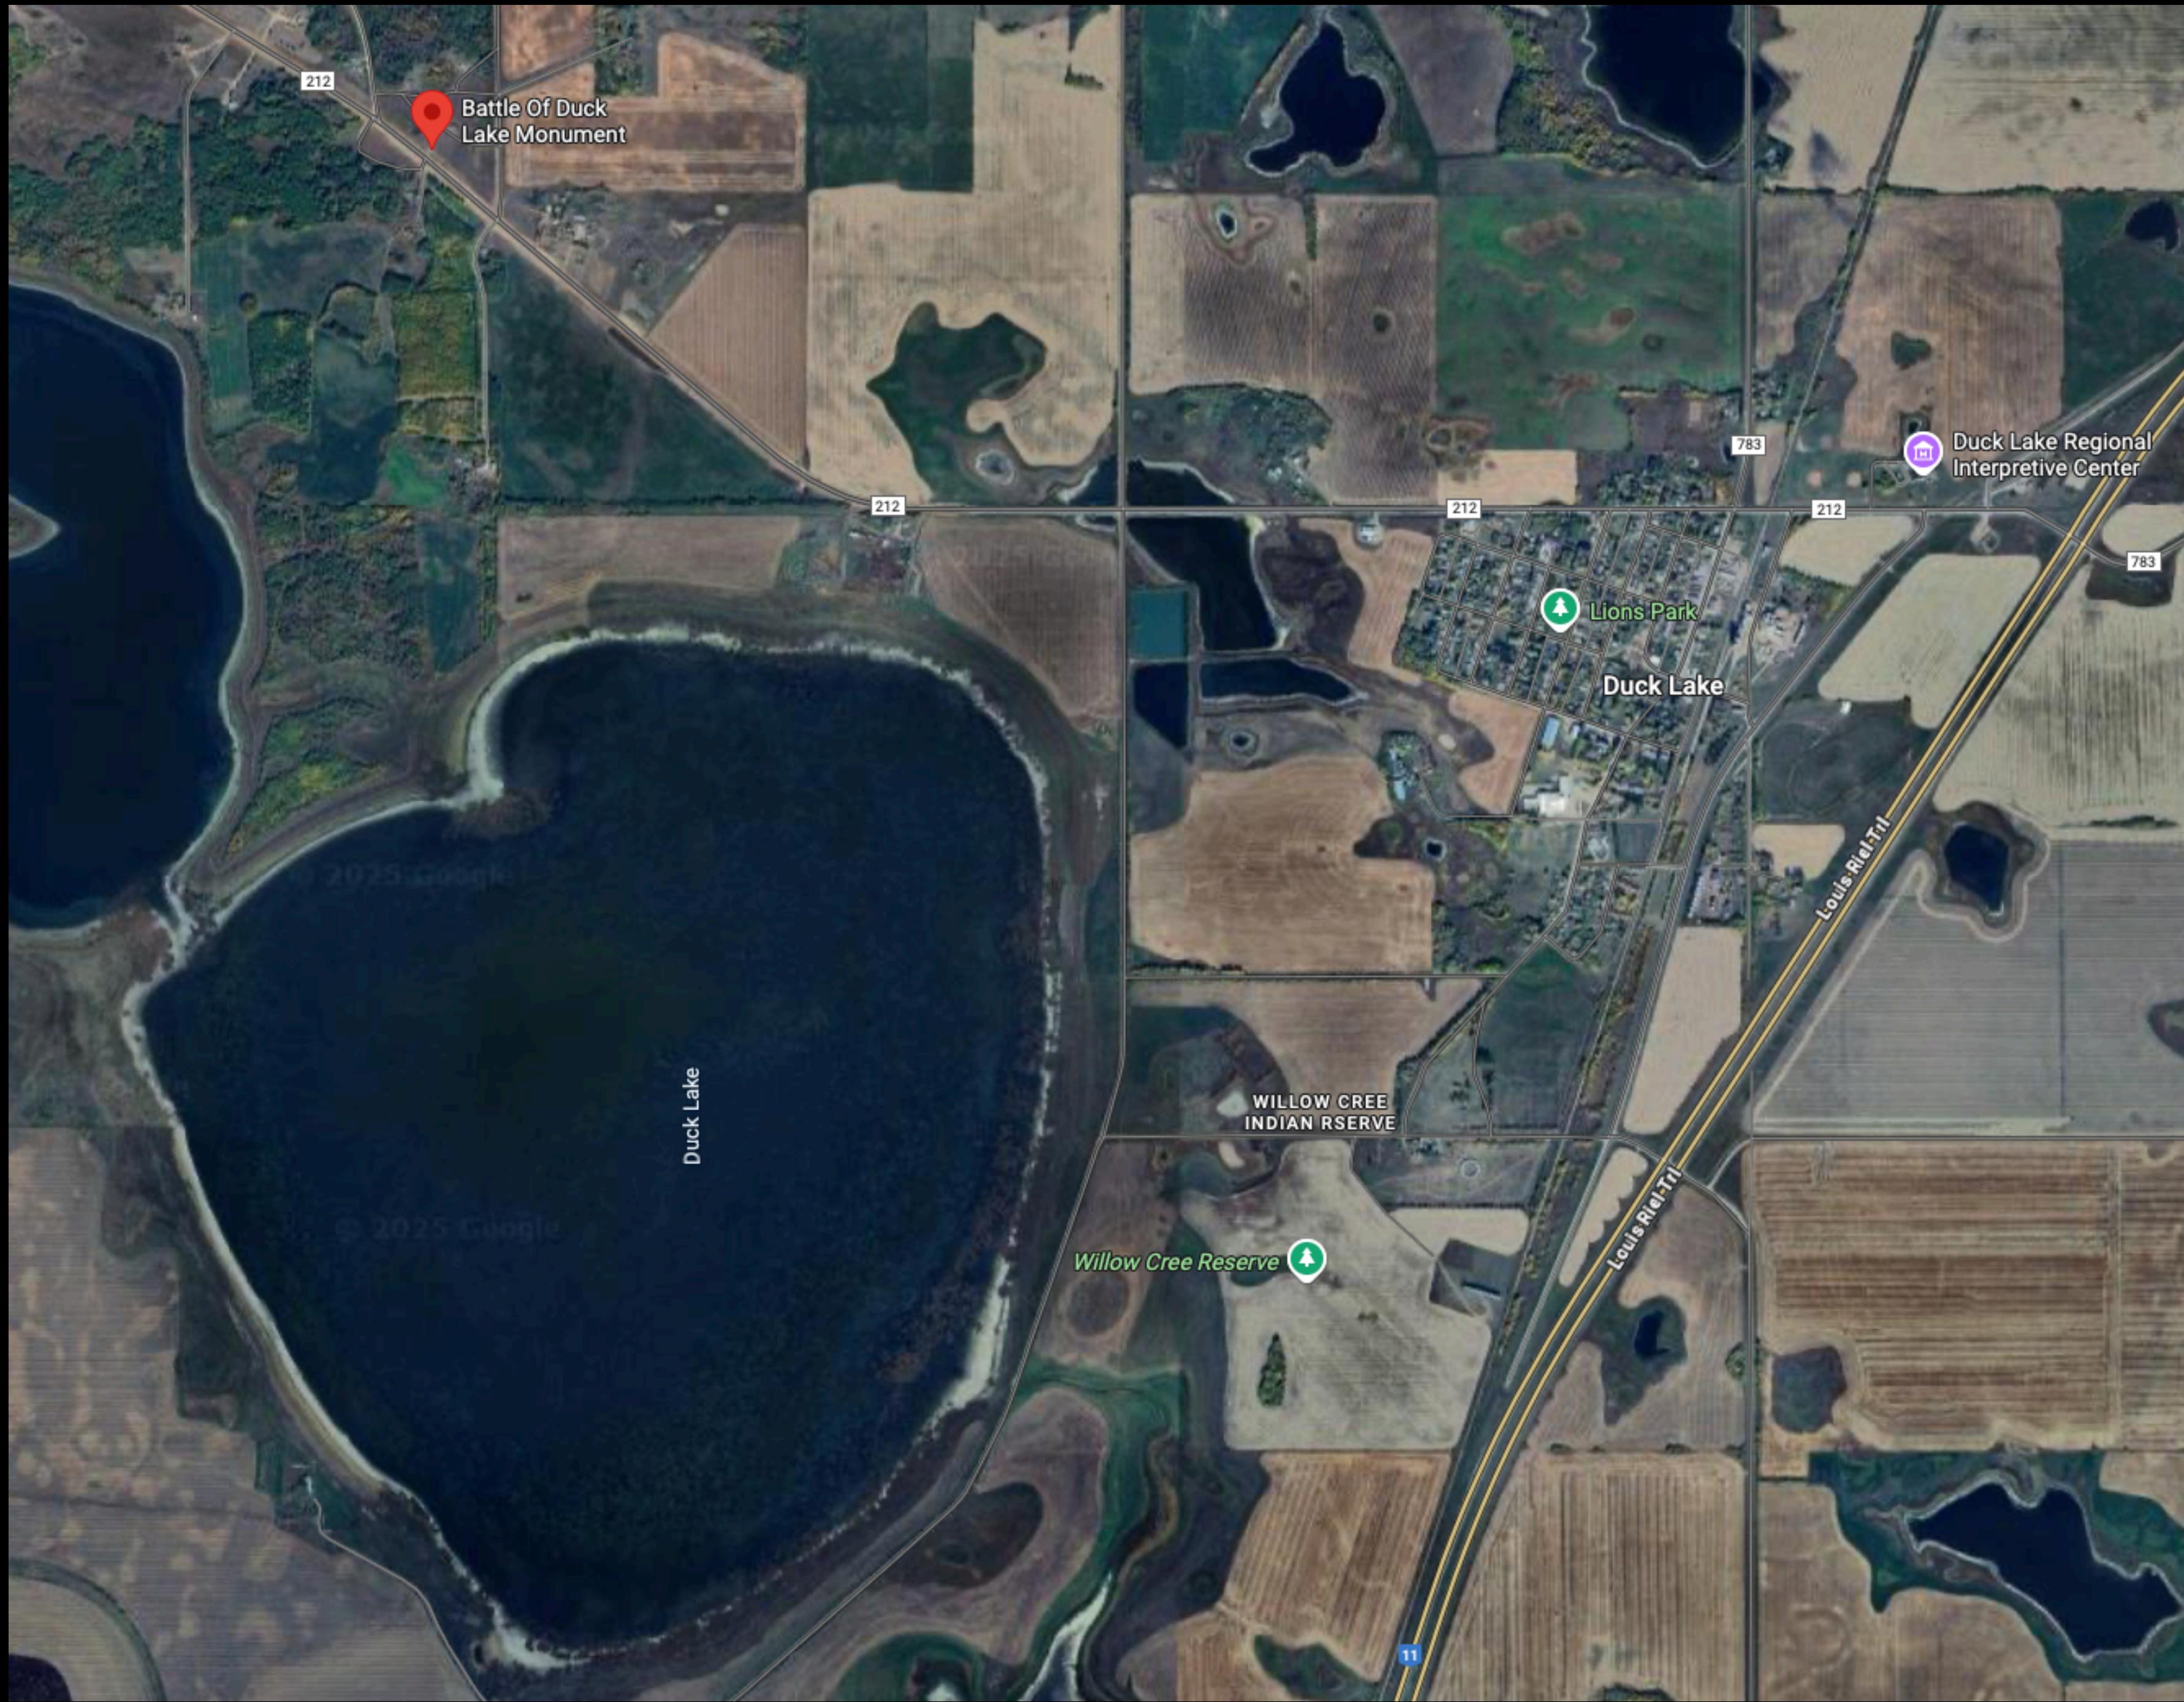
# Java REST Postgres

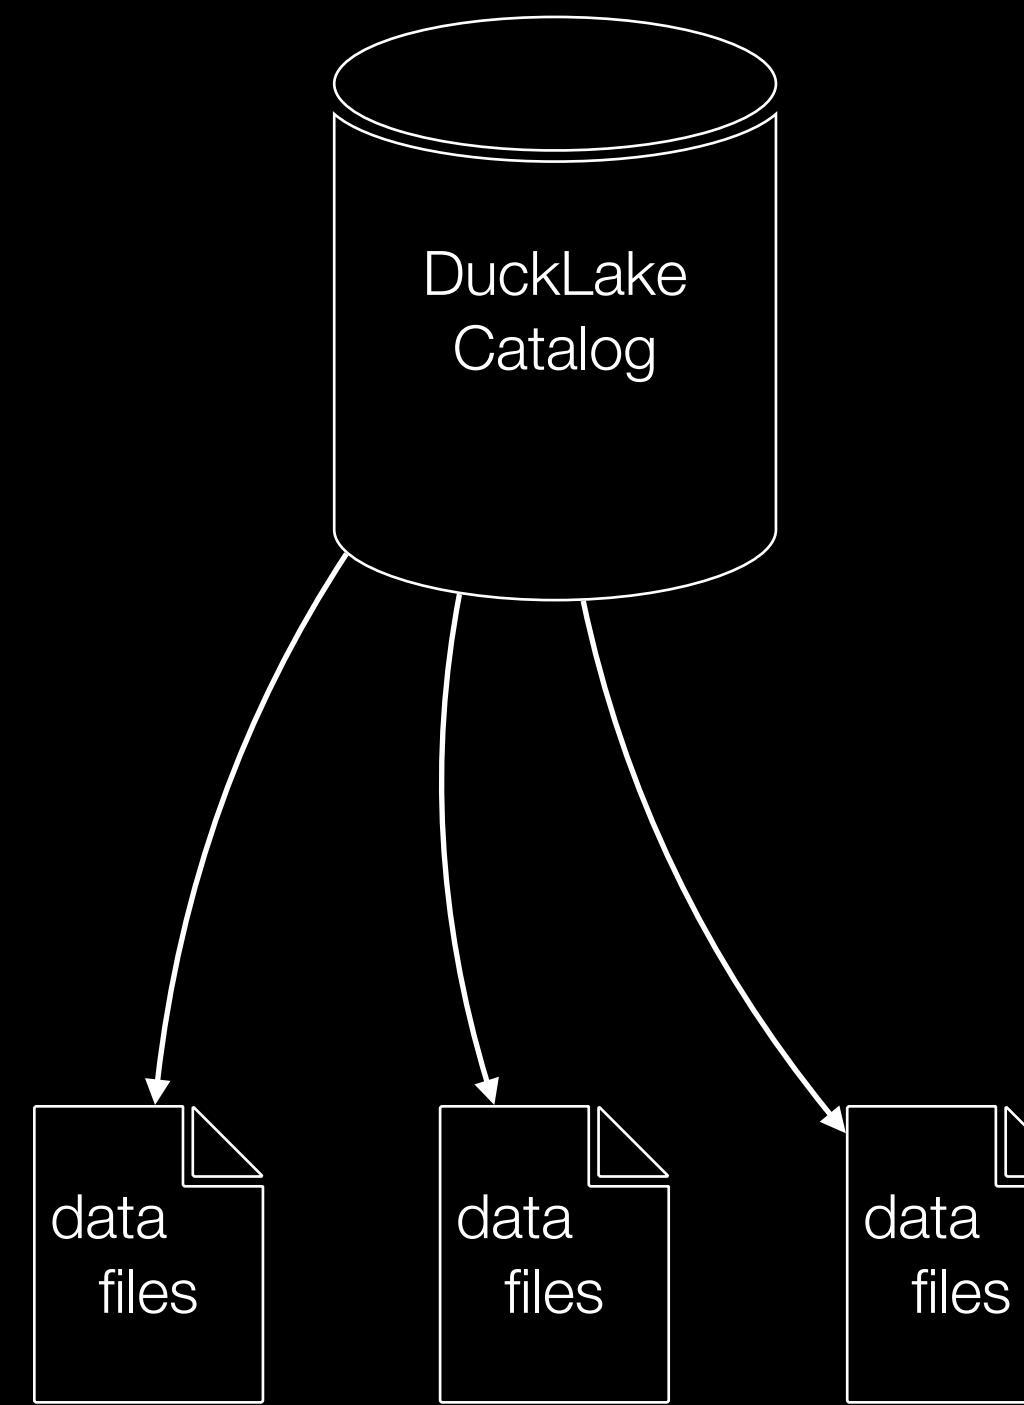
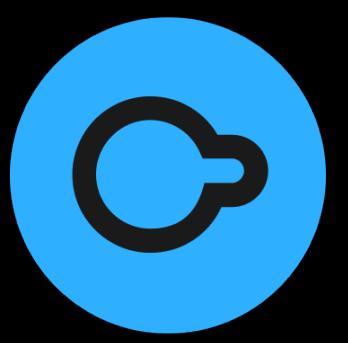
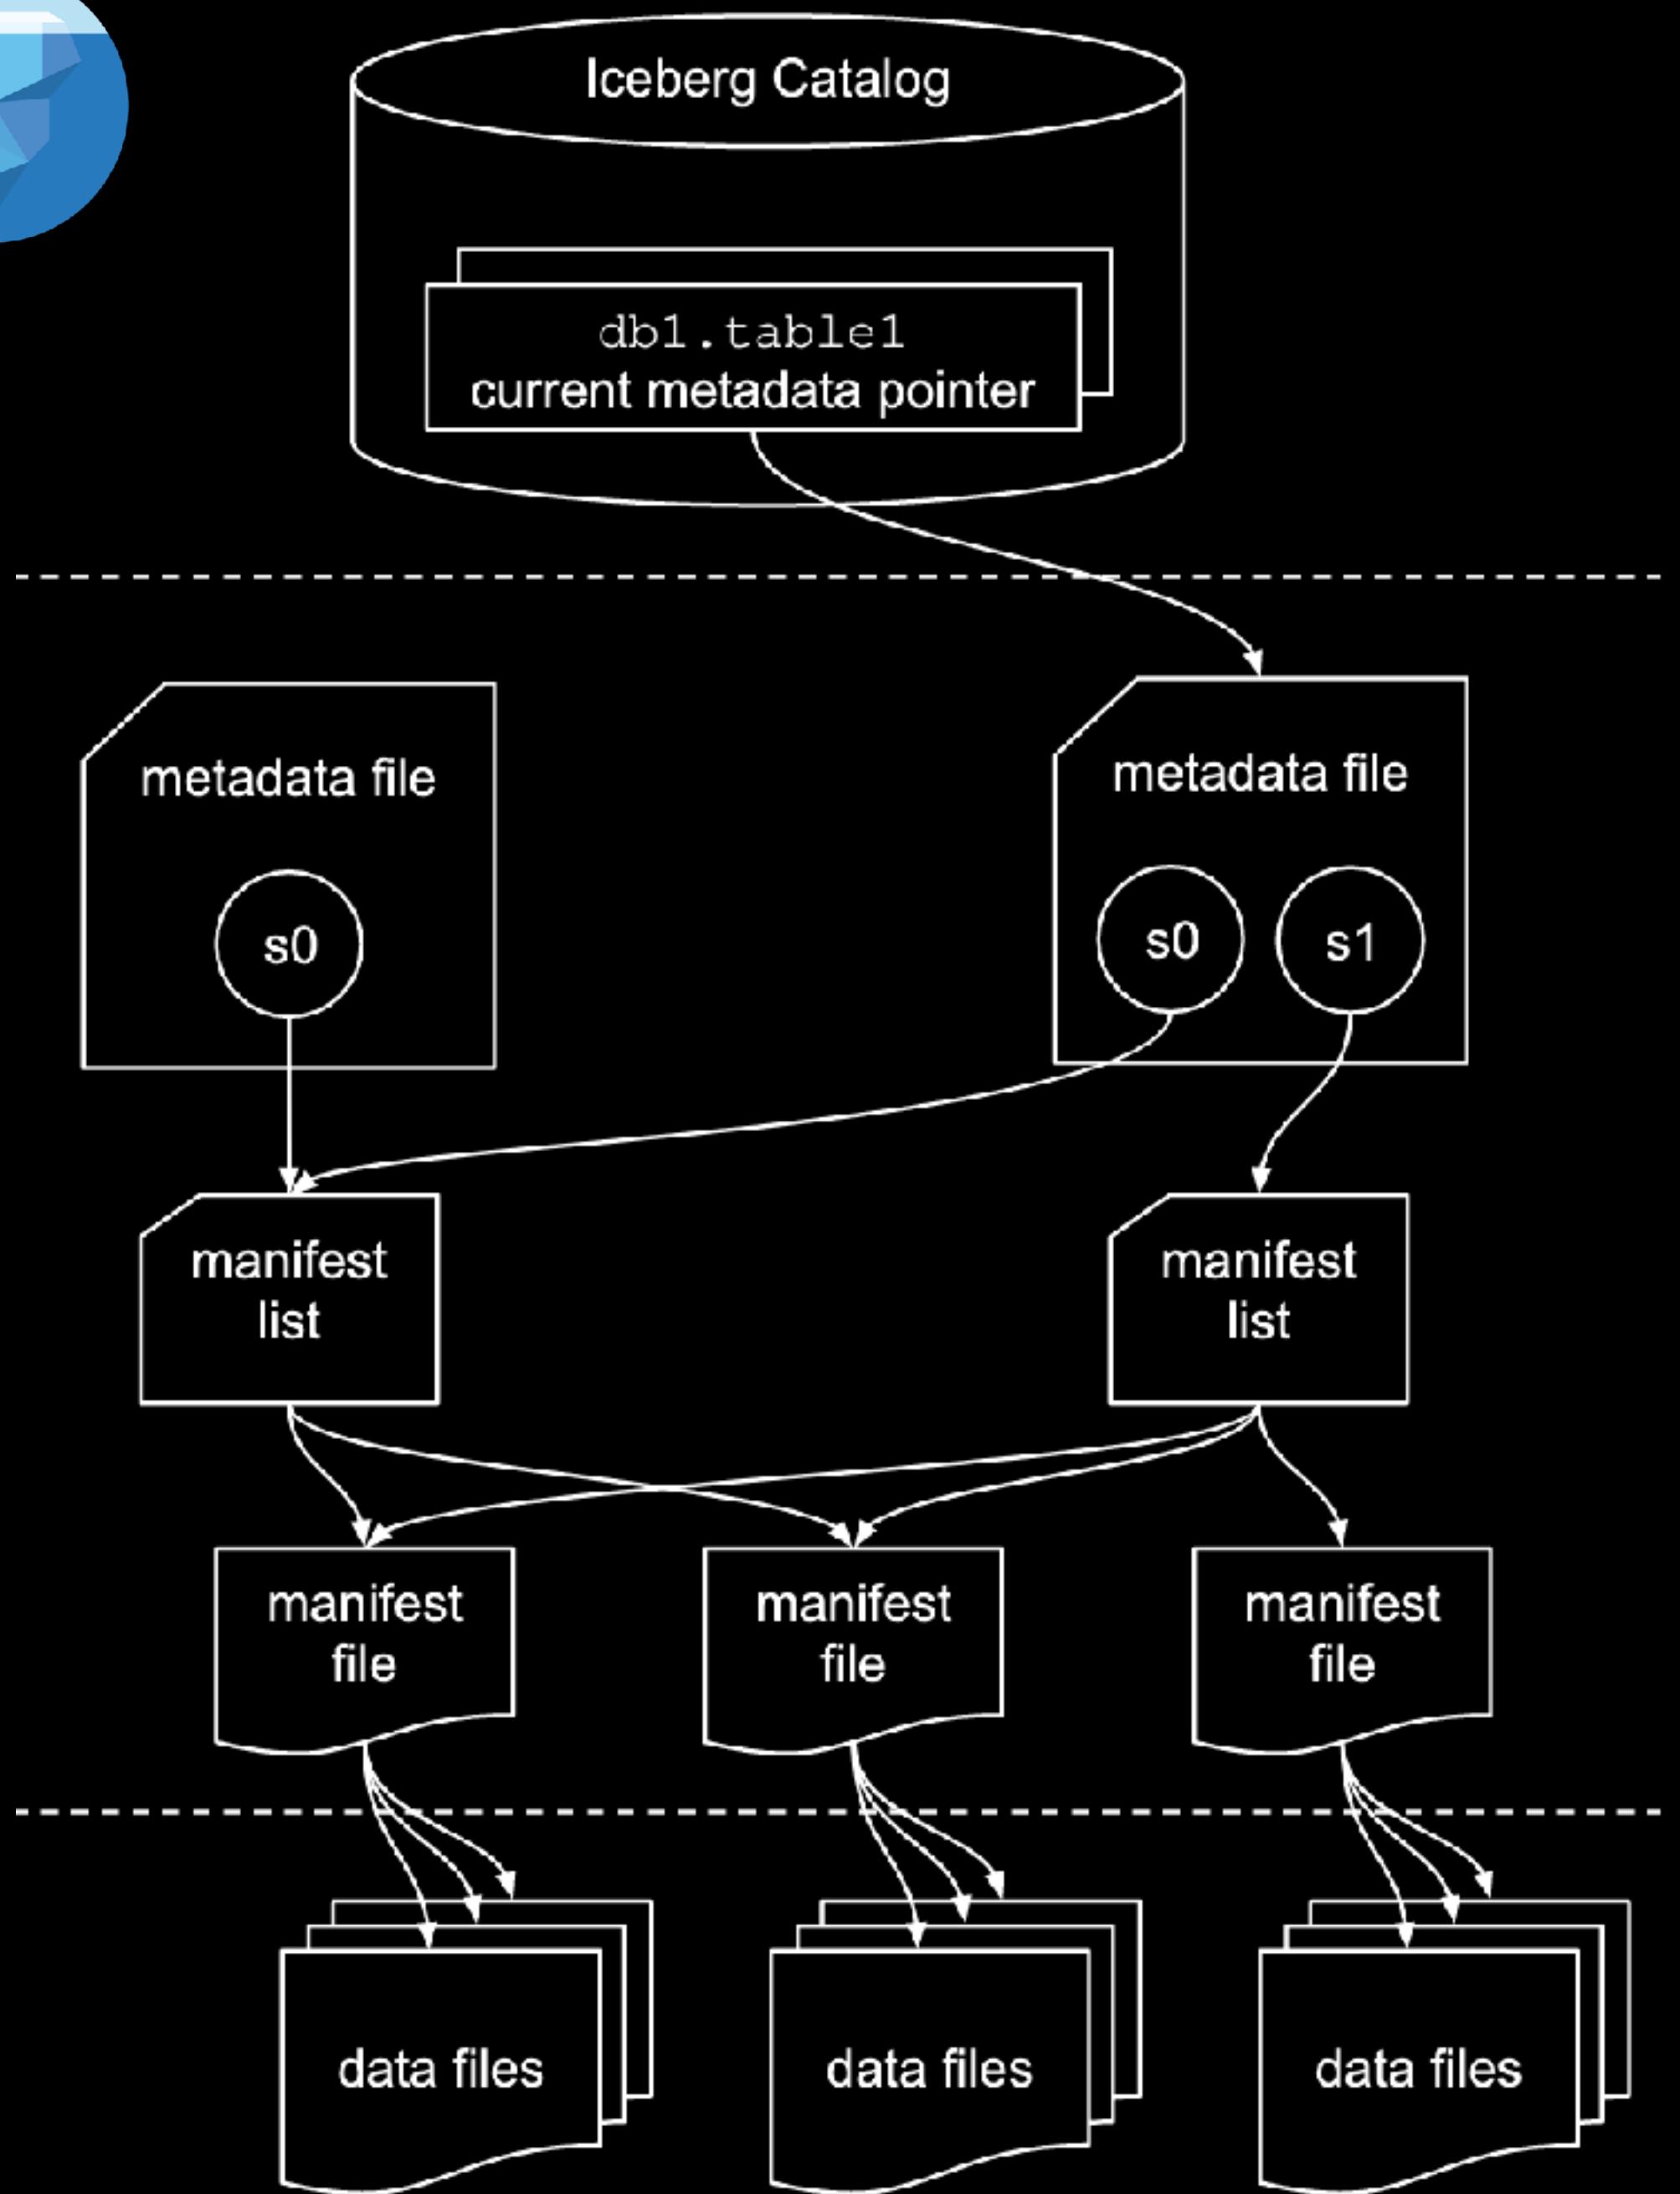






DuckLake





Simplicity  
Scalability  
Speed

# Simplicity

```
$ curl install.duckdb.org | bash
```

```
$ curl install.duckdb.org | bash  
$ ~/.duckdb/cli/latest/duckdb
```

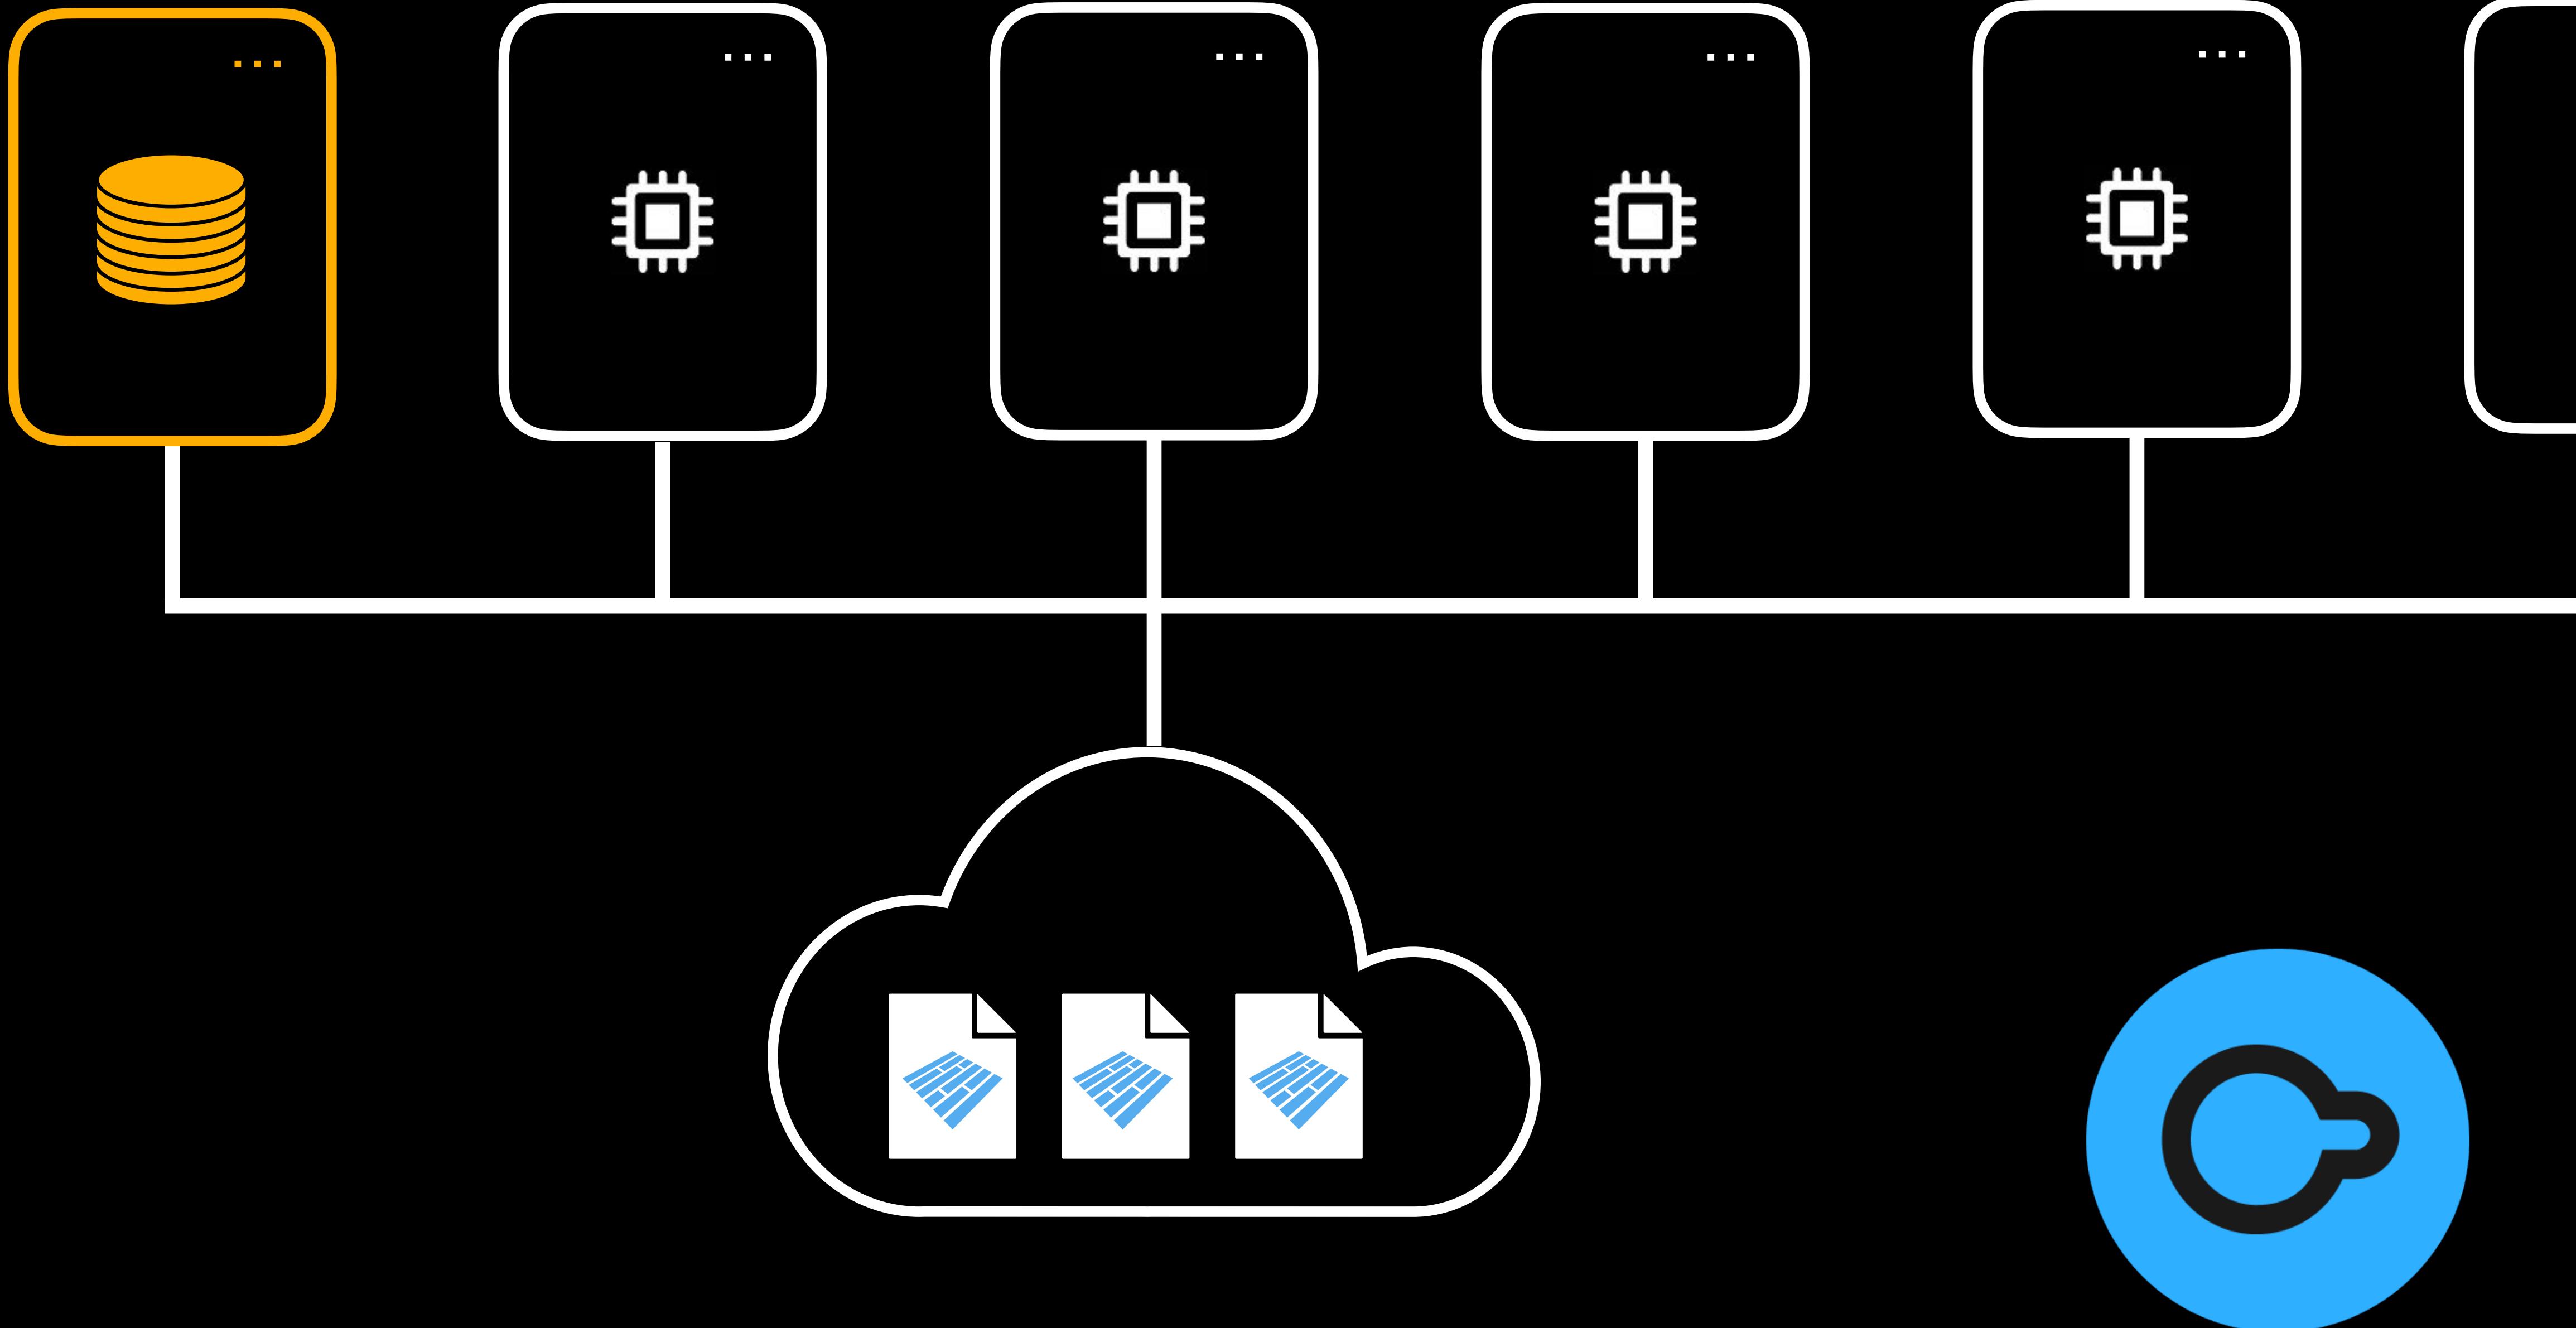
```
$ curl install.duckdb.org | bash
$ ~/.duckdb/cli/latest/duckdb
▶ ATTACH 'ducklake:metadata.ducklake' AS dl1;
```

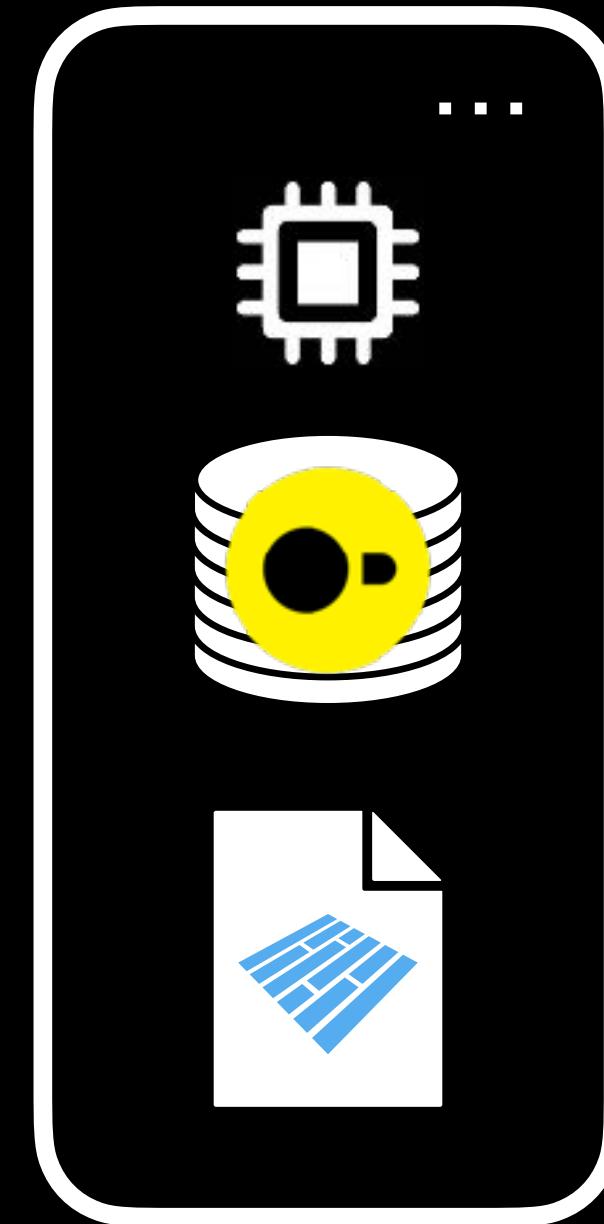
## Database



# Scalability

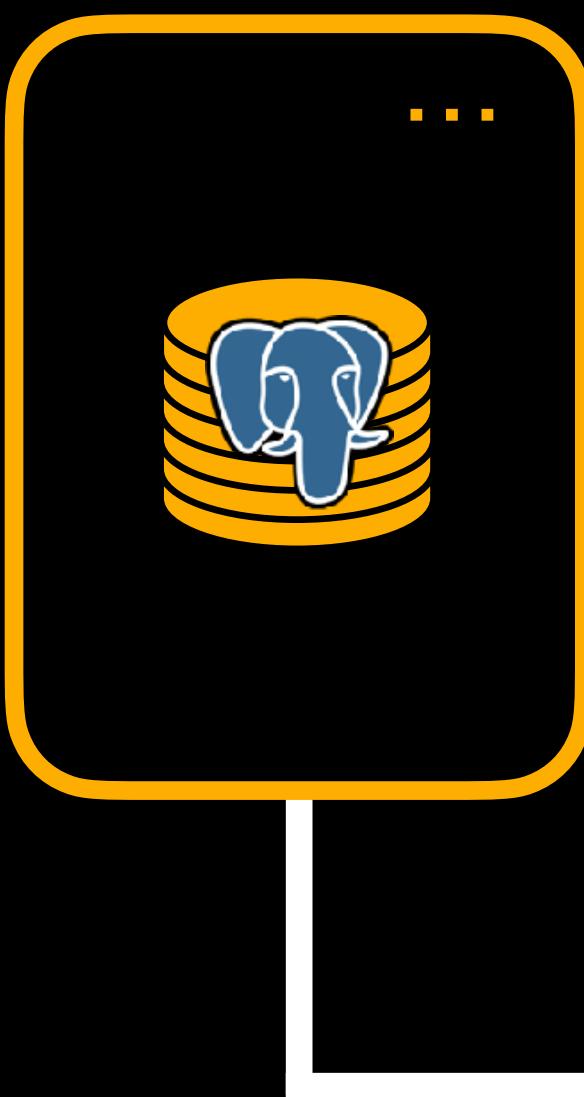




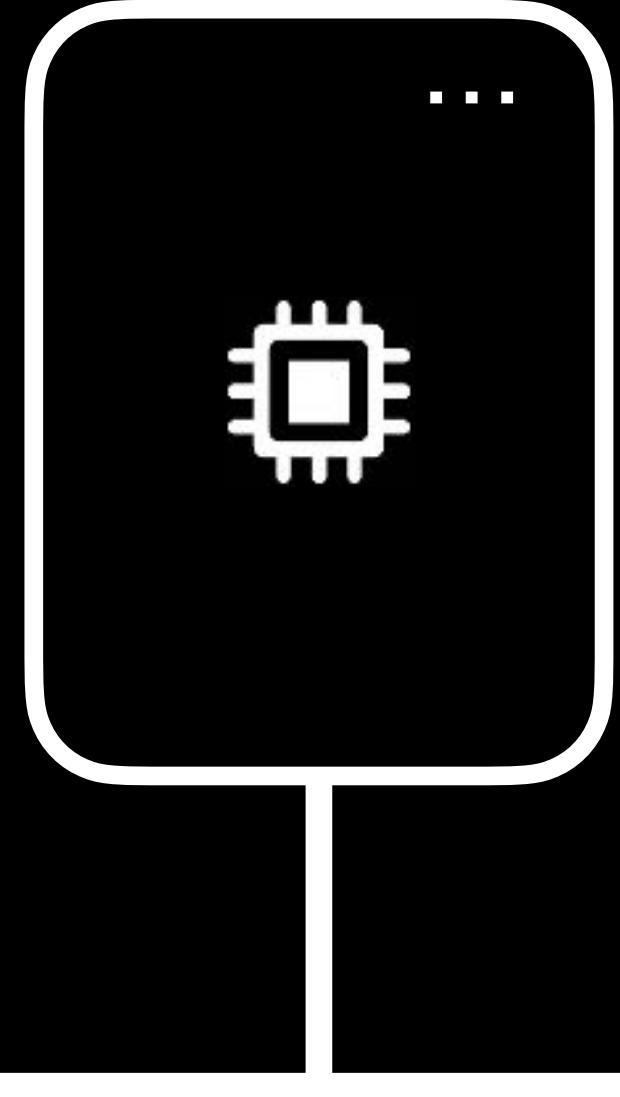
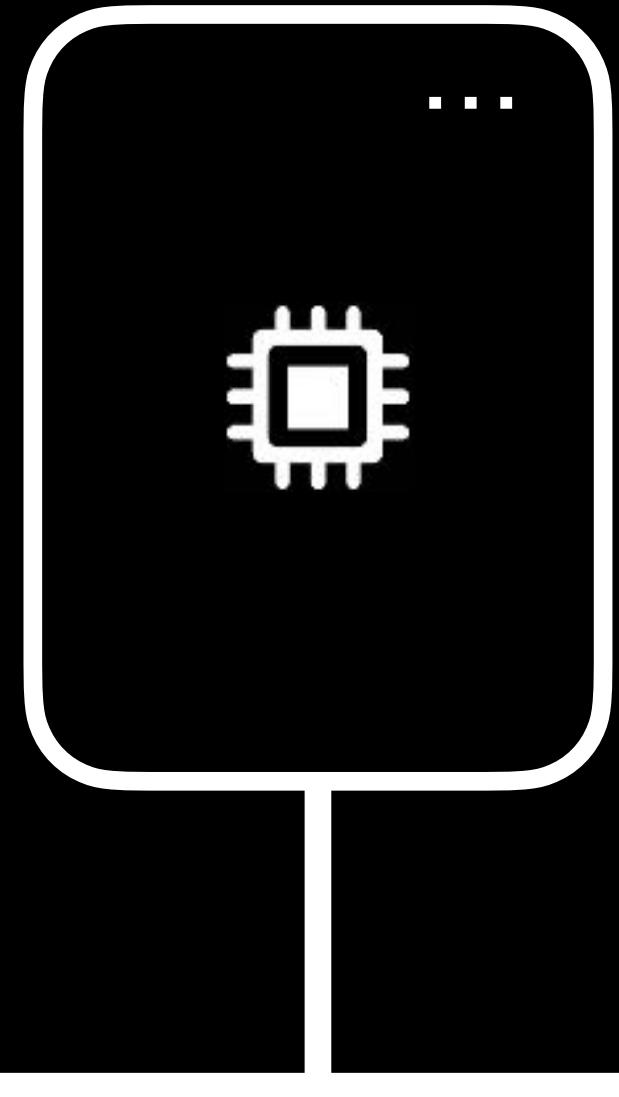
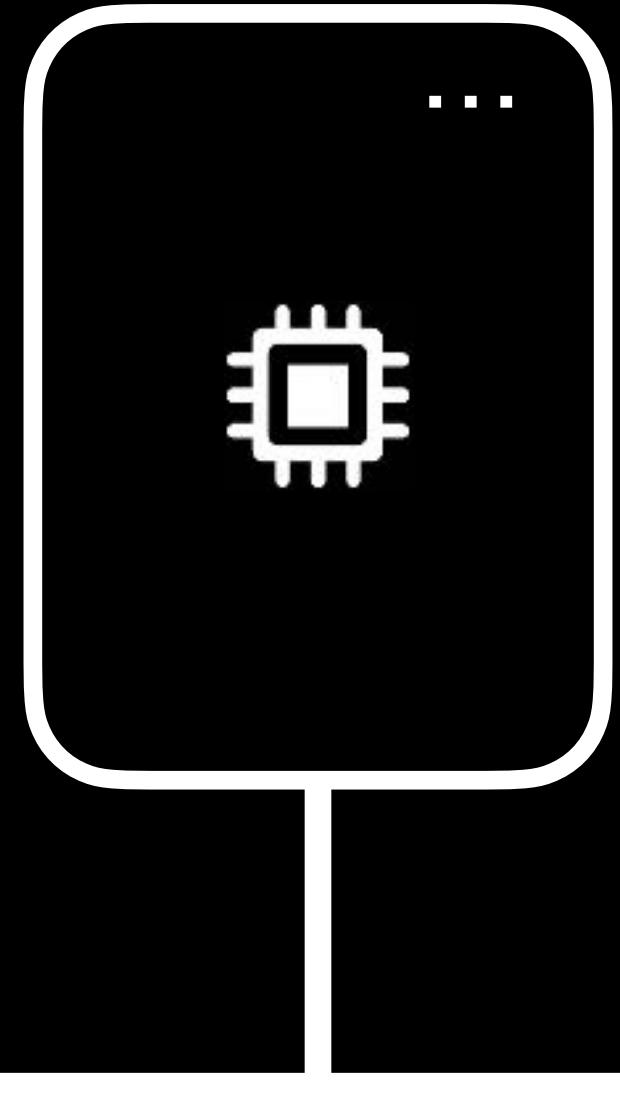
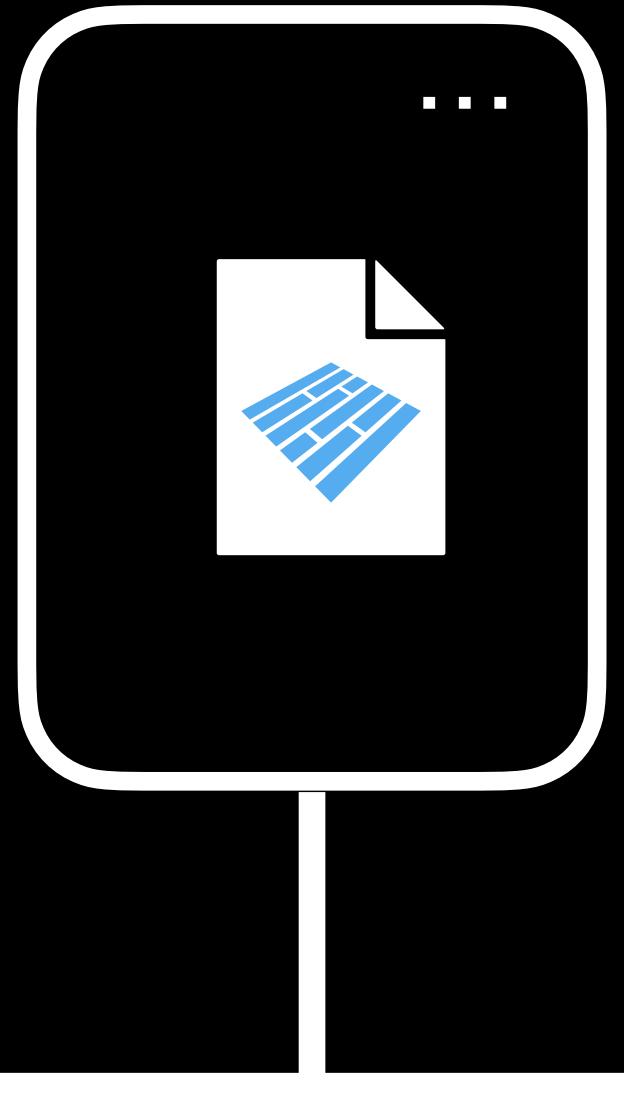


**ATTACH 'ducklake:metadata1.ducklake' AS dl1;**

metadata

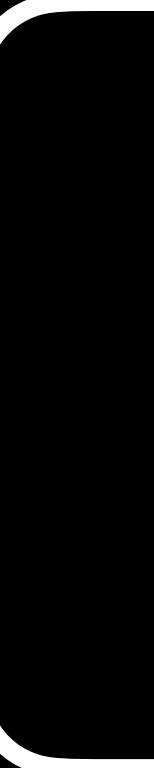
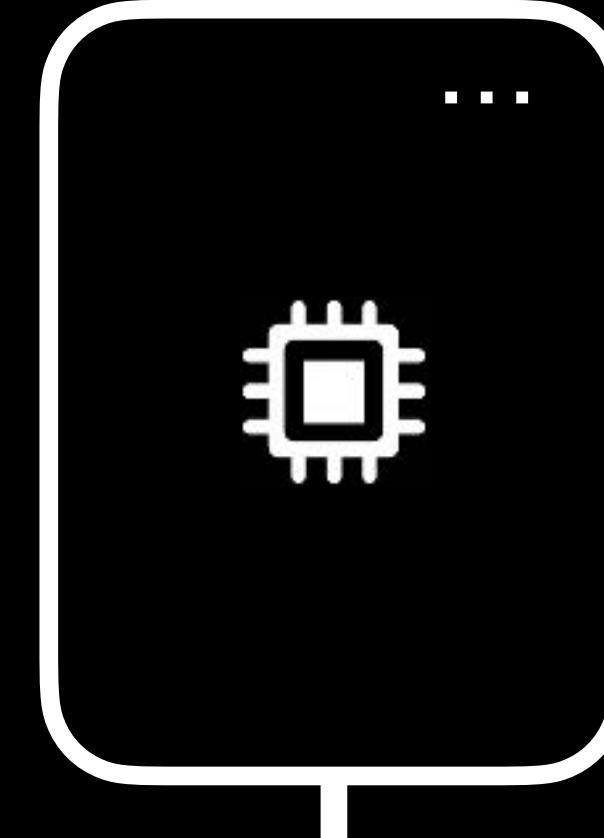
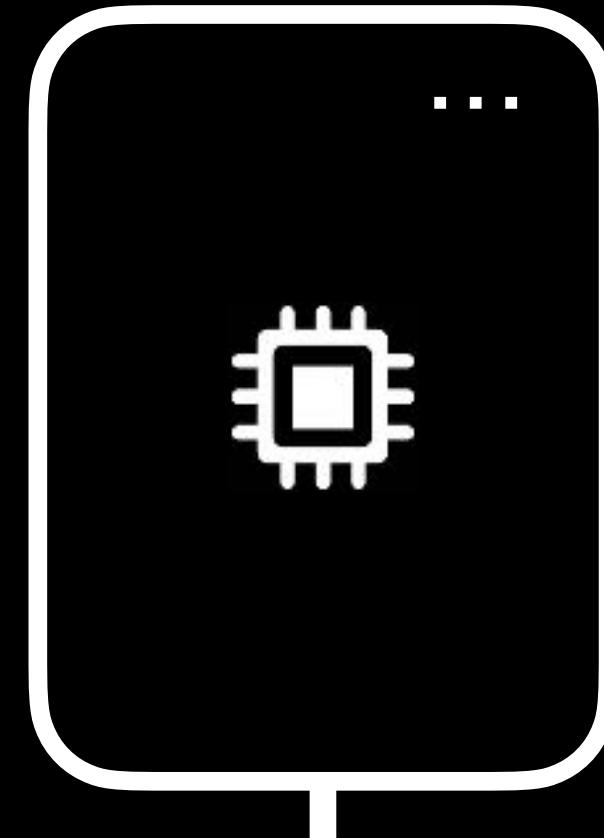
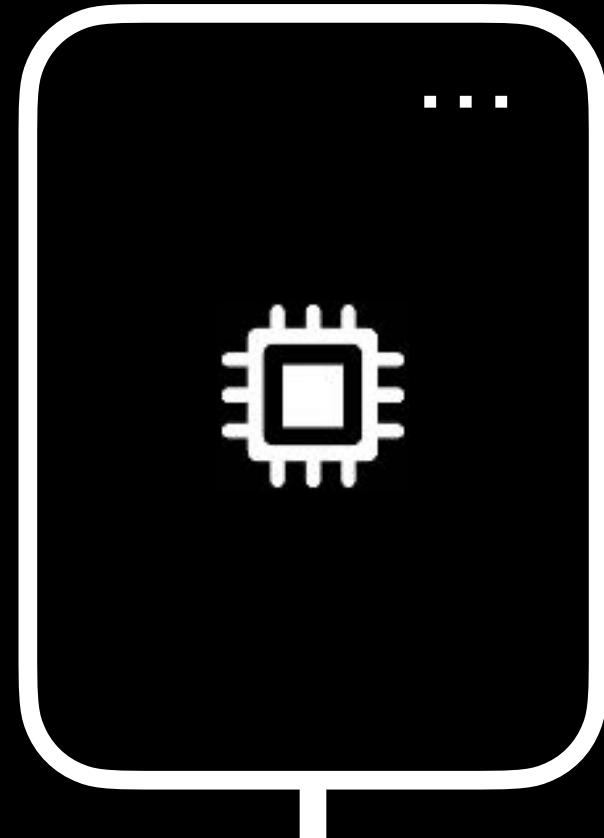
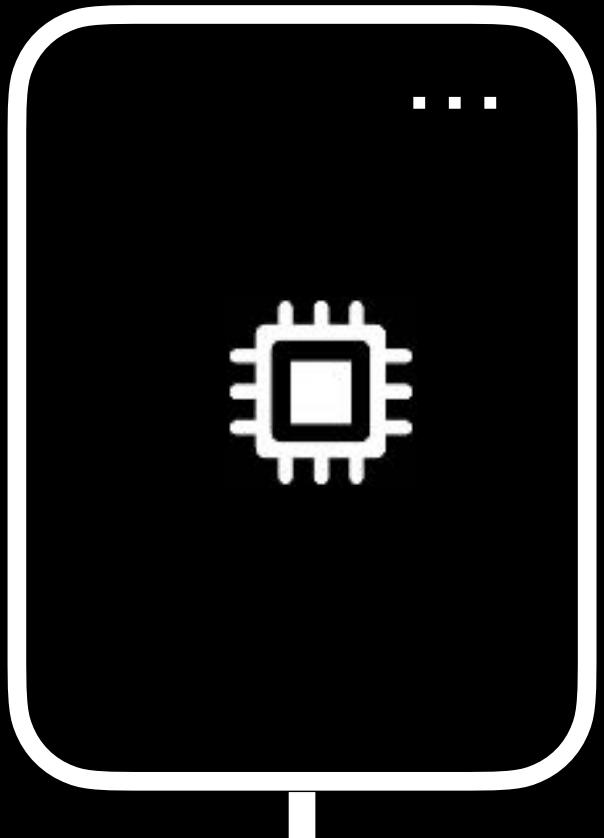
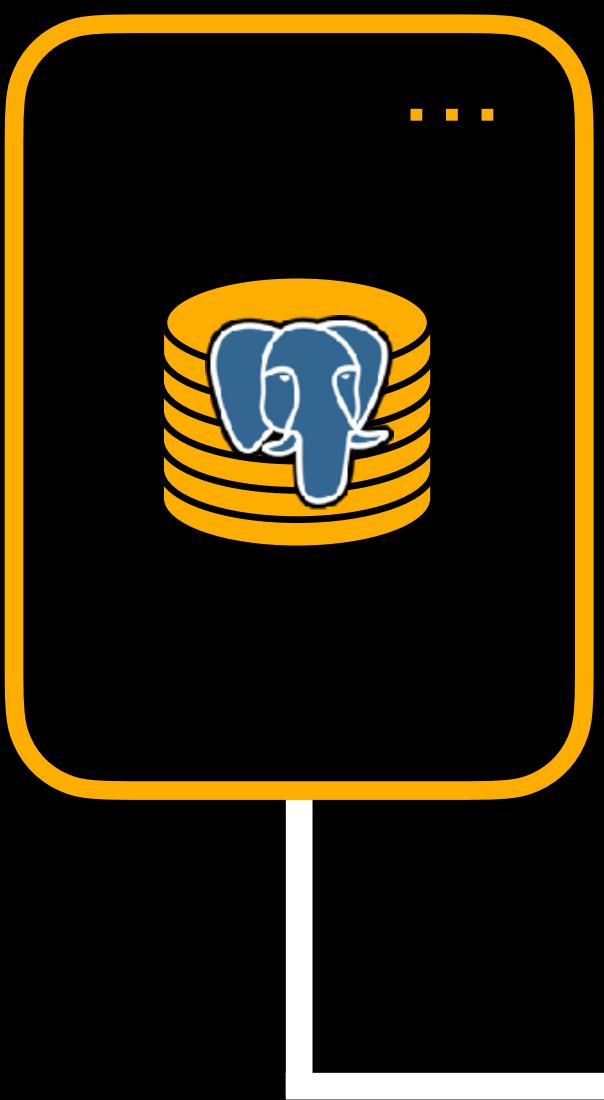


nfs

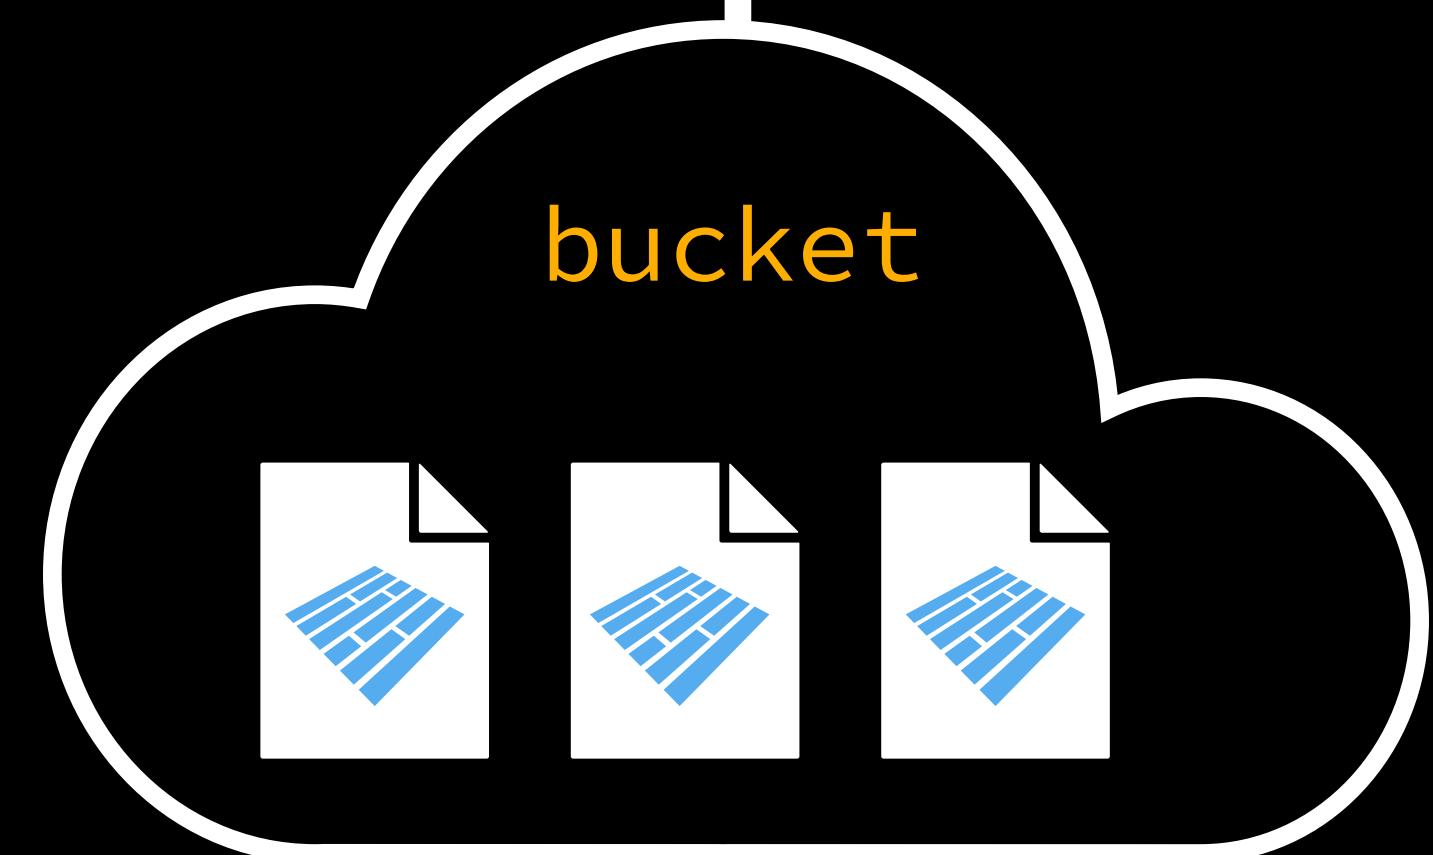


```
ATTACH 'ducklake:postgres:host=metadata' AS dl  
(DATA_PATH '/nfs/data_files/');
```

metadata

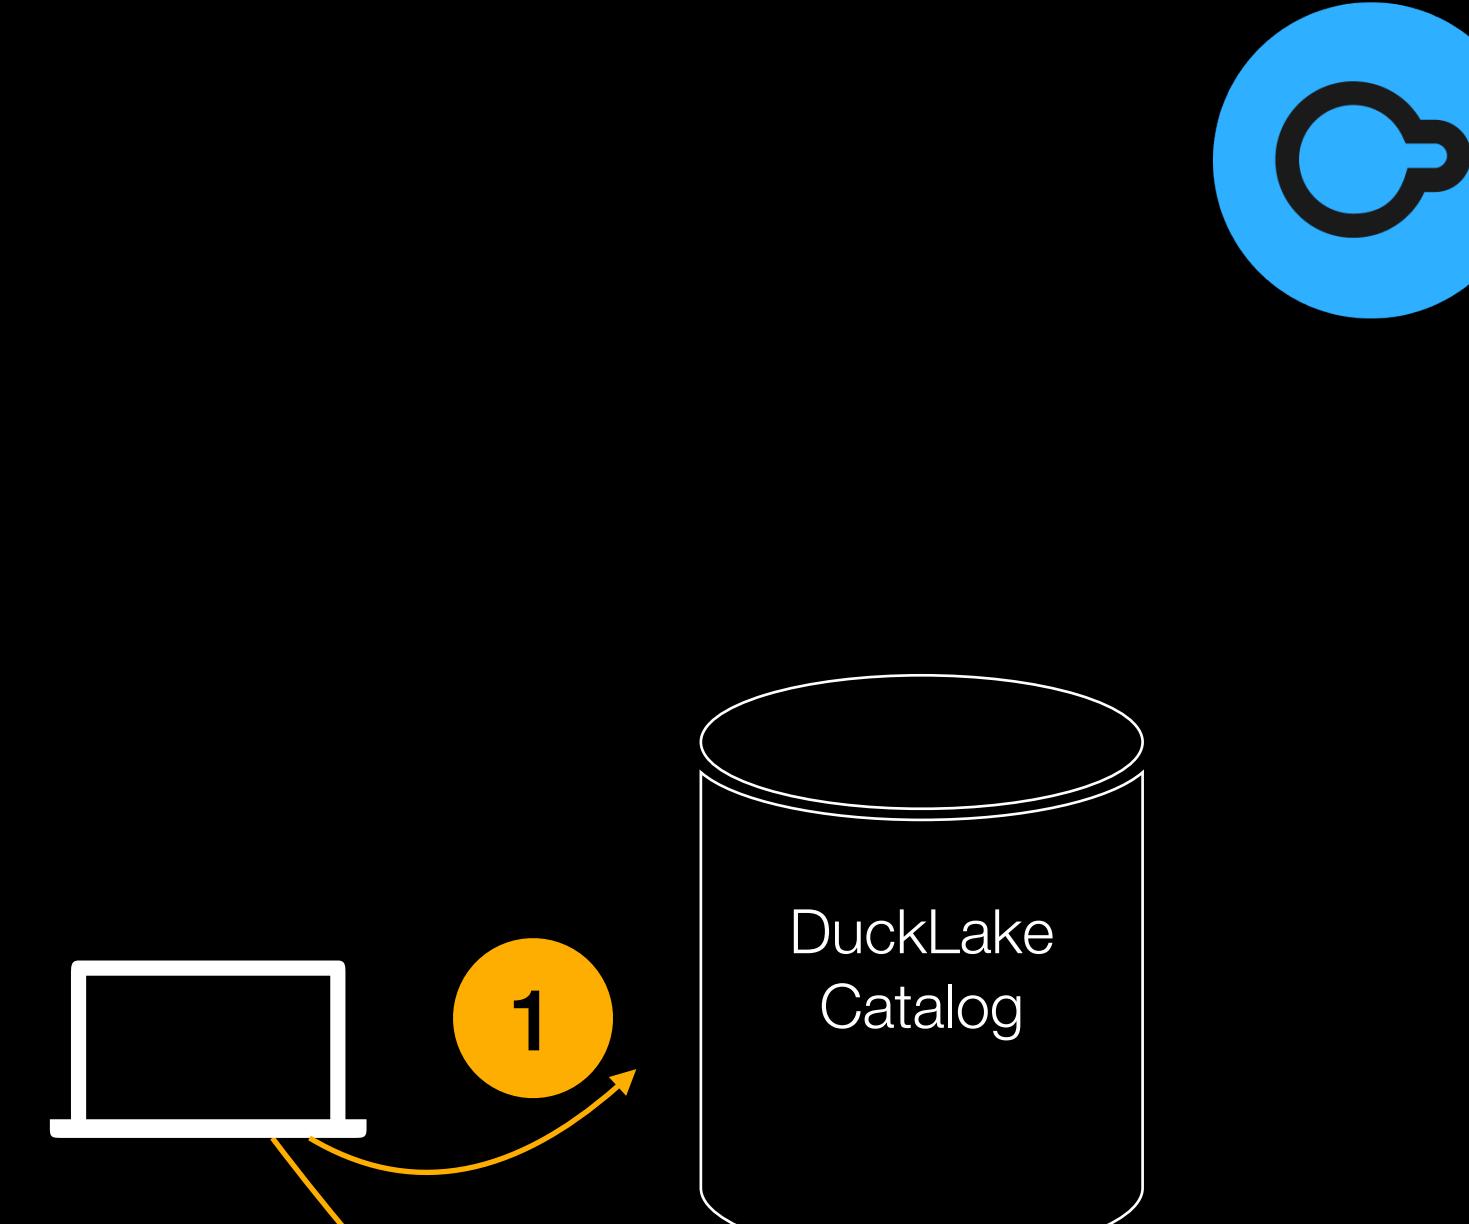
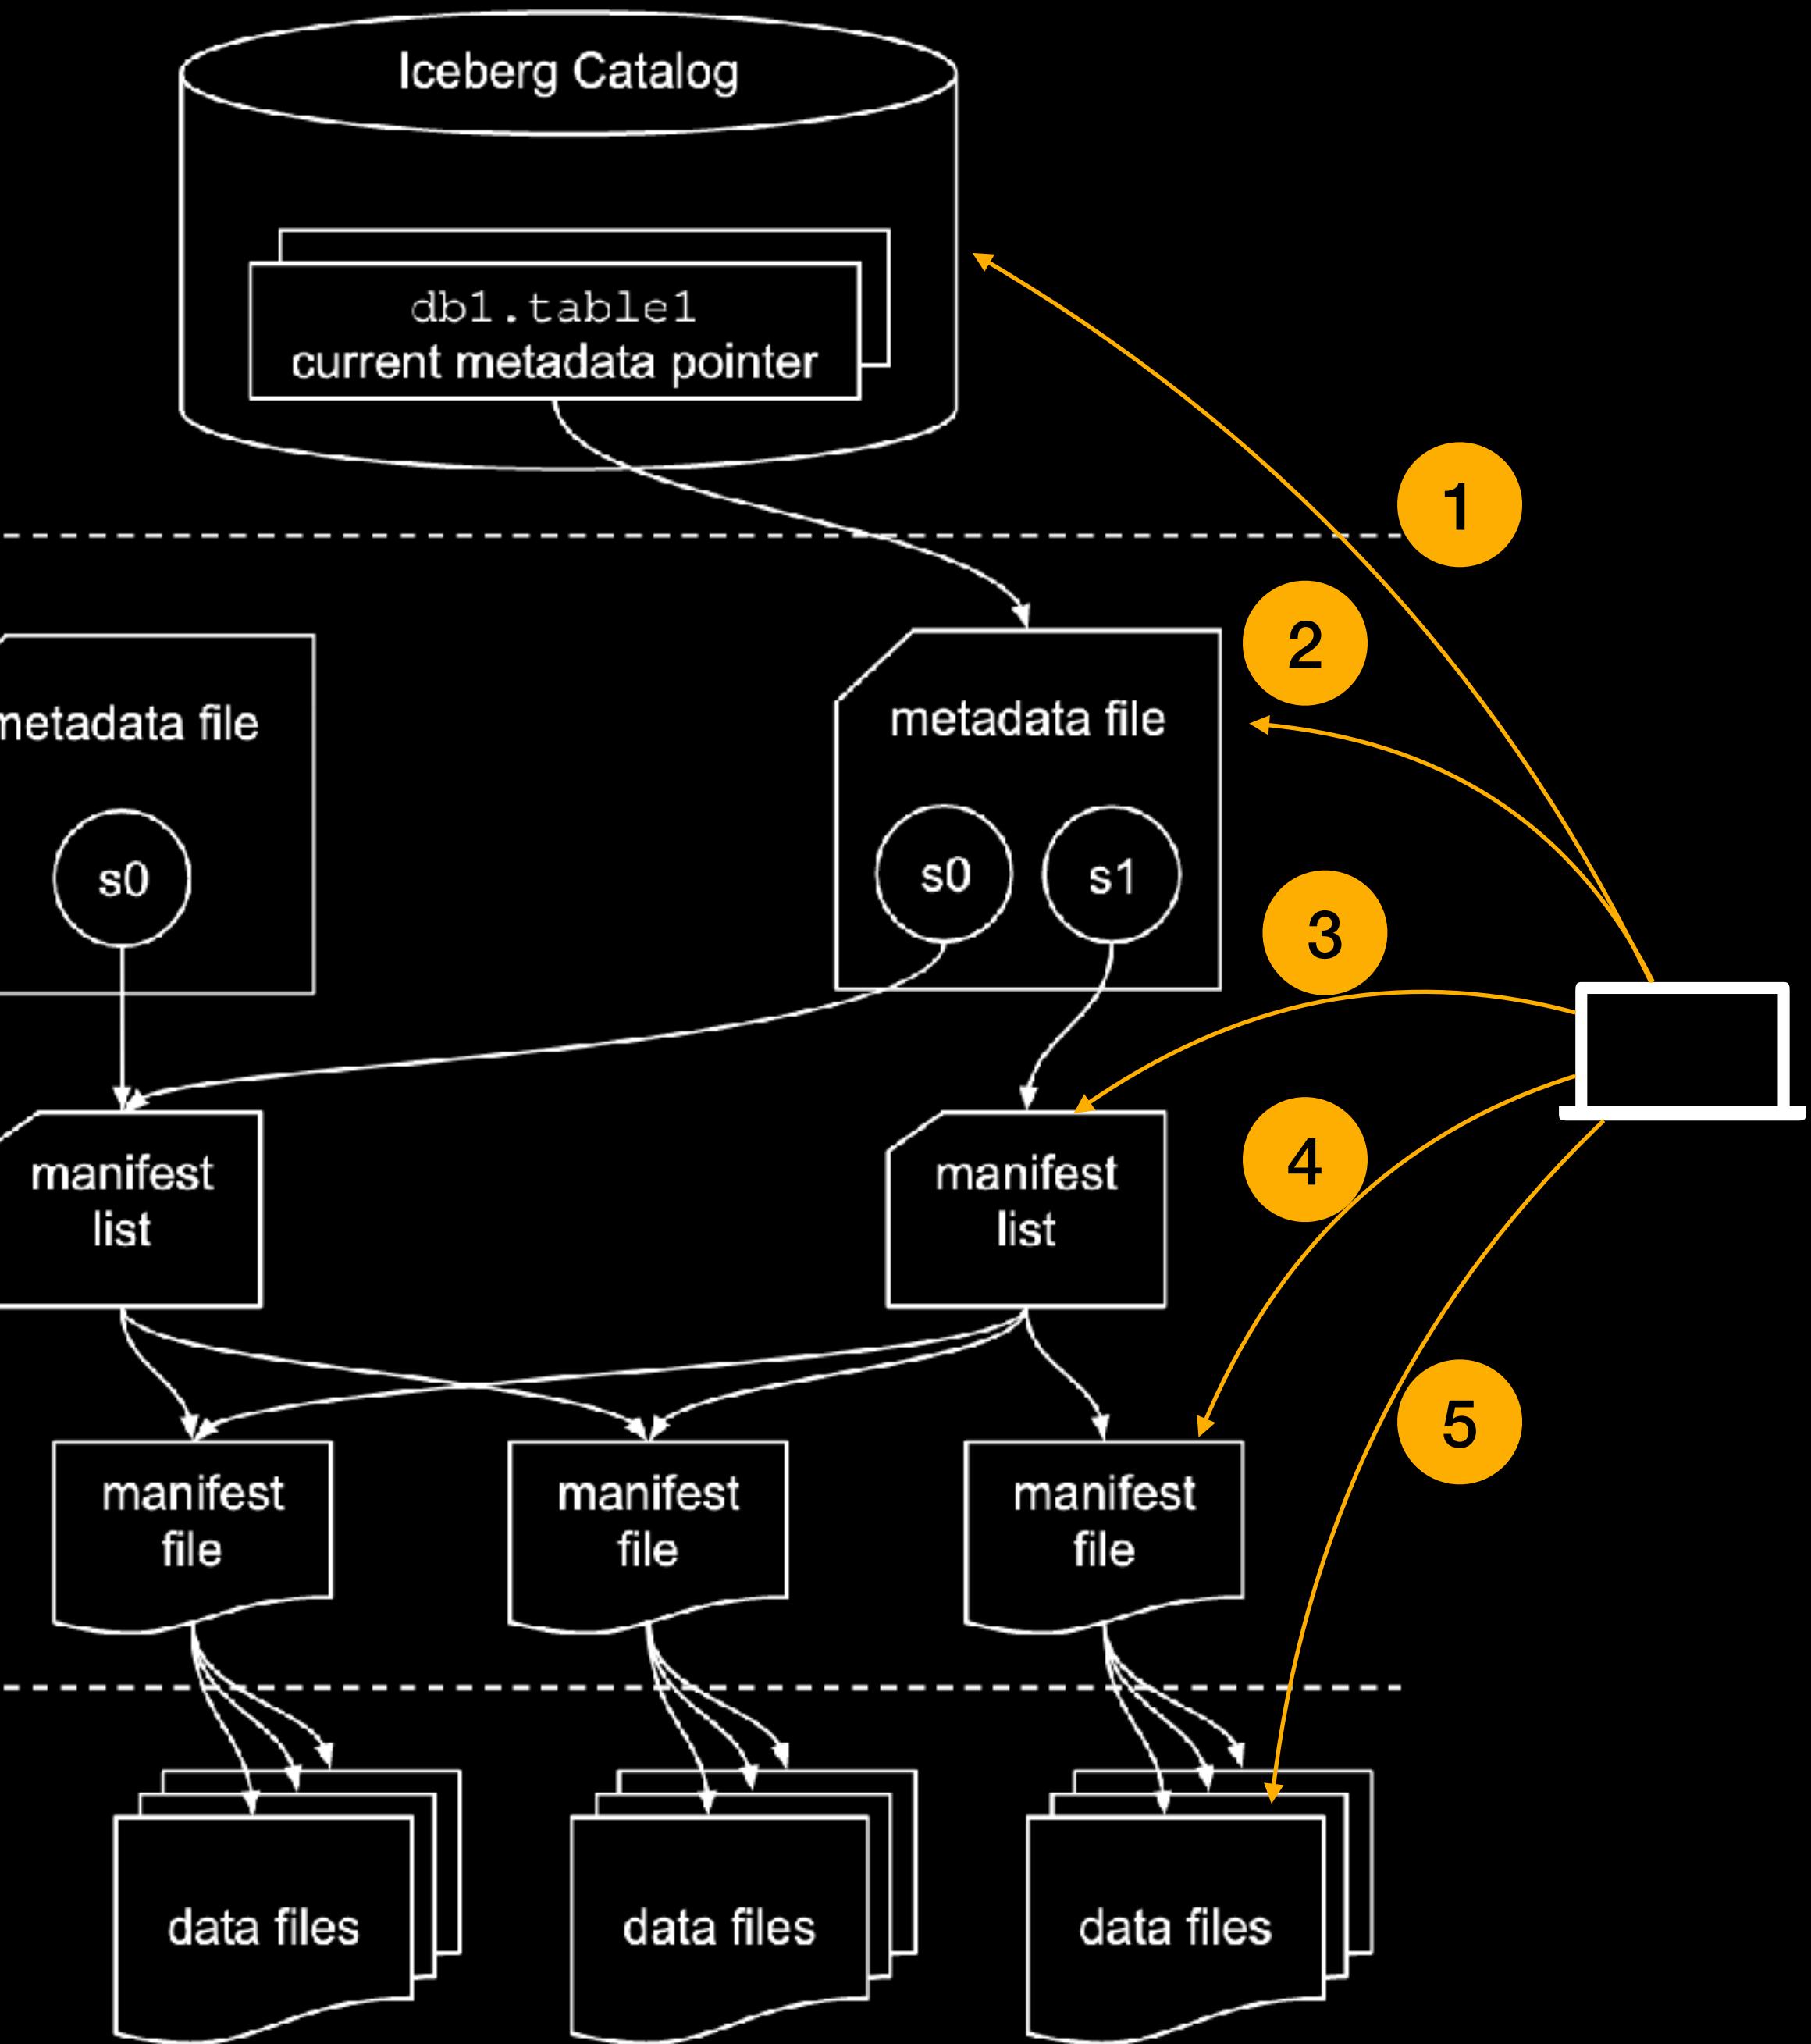


bucket



```
ATTACH 'ducklake:postgres:host=metadata'  
AS dl (DATA_PATH 's3://bucket/...');
```

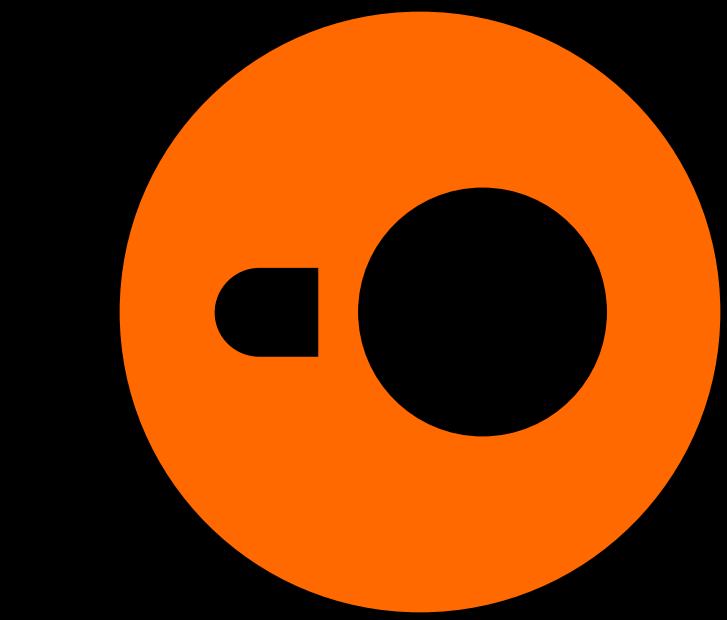
# Speed

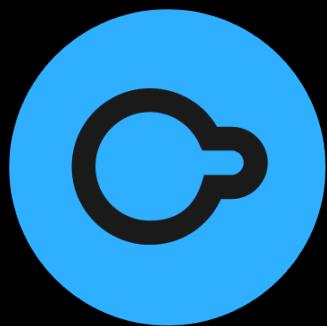




- 1 Petabyte of data
- 10 Columns
- 100 M snapshots, 10 years
  - 270 GB/day

Planning Time	DuckDB	Postgres
Insert	~5ms	~3ms
Select All Files	~0.8s	~200s
Select 1 Month	~0.25s	~0.25s
Select 1 Day	~0.25s	~0.25s





# Multi-Table

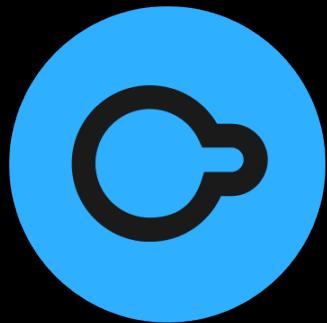
```
ATTACH 'ducklake:metadata.ducklake' AS dl;
```

```
CREATE SCHEMA dl.s1;
```

```
CREATE TABLE dl.s1.t1 ...;
```

```
CREATE TABLE dl.s1.t2 ...;
```

```
...
```



# Holistic Transactions

**BEGIN**

CREATE TABLE ...

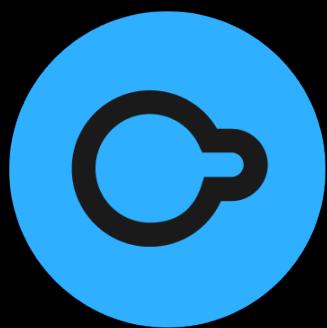
ALTER TABLE ...

INSERT ... INTO t1

UPDATE t2 ...

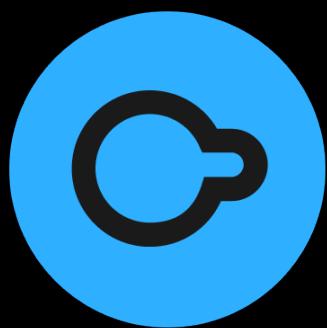
DELETE FROM t3

**COMMIT**



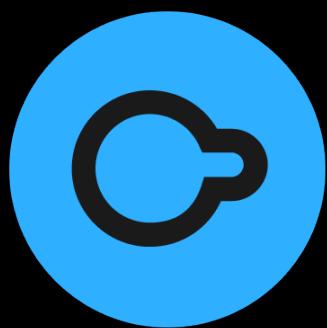
# Inlining

```
ATTACH 'ducklake:metadata.ducklake' AS dl  
  (DATA_INLINING_ROW_LIMIT 100);
```



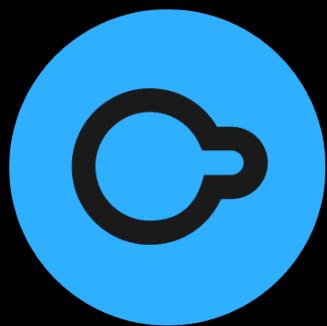
# Time Travel

```
FROM tbl AT (VERSION => 3);  
FROM tbl AT (TIMESTAMP => ...);
```



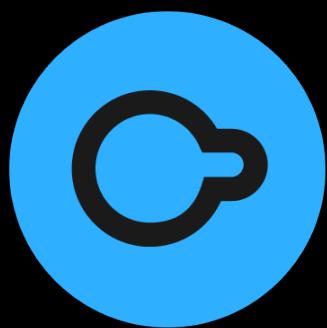
# Change Feed

```
ATTACH 'ducklake:metadata.ducklake' AS dl;  
FROM dl.table_changes('tbl', 3, 4);
```

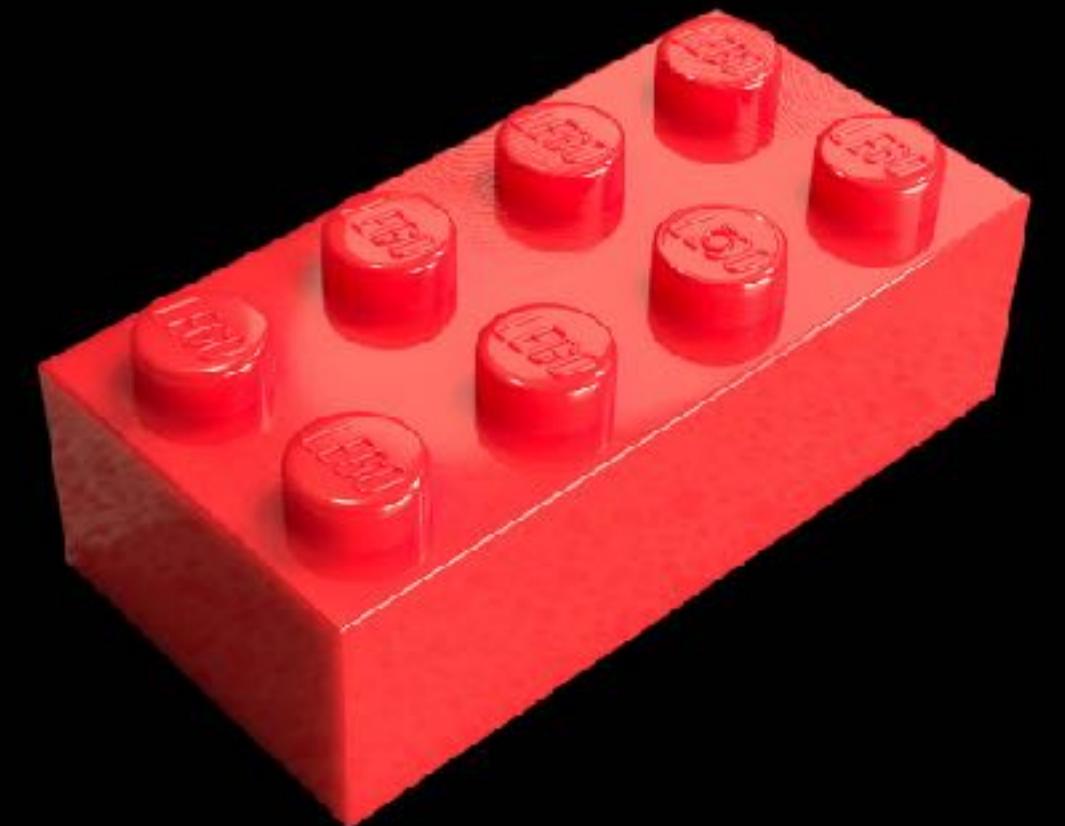


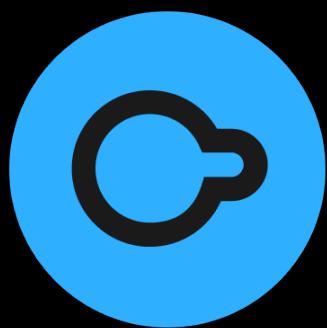
# Encryption

```
ATTACH 'ducklake:metadata.ducklake' AS dl  
(DATA_PATH '/some/untrusted/location' ENCRYPTED);
```



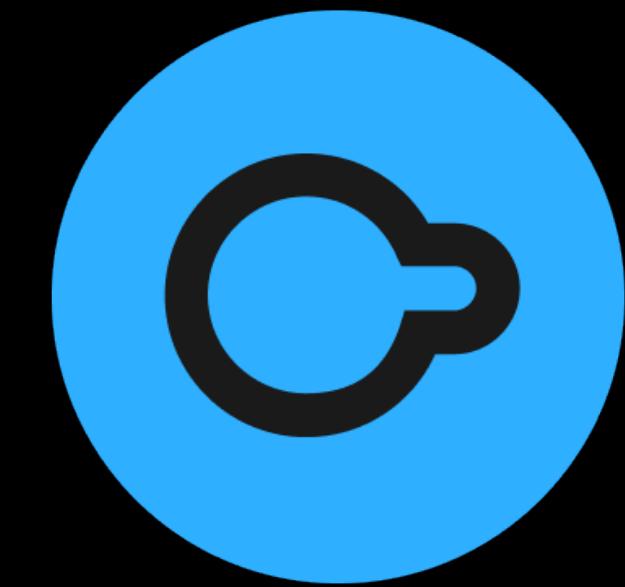
# Interoperability



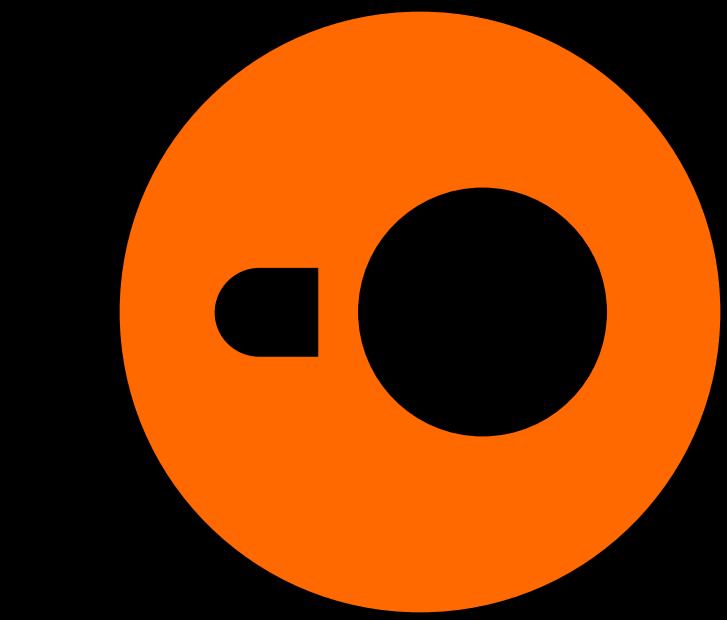


Free





Standard <>  
Implementation



```
$ curl install.duckdb.org | bash  
$ ~/.duckdb/cli/latest/duckdb  
▶ ATTACH 'ducklake:metadata.ducklake' AS dl1;
```



# DuckLake

ducklake.select  
hannes.muehleisen.org

