



Delighting users with RESTful APIs and DuckDB

2024-08-15, DuckCon #5, Seattle, USA

Miguel Filipe, Principal Engineer @ Dune Analytics

Miguel Mascarenhas Filipe



- Love DBs & Dist Systems since 2001
- At AWS, part of the DynamoDB launch team
- At Skype worked on Distributed Timeseries DB
- Working in Startups since 2015
- Building systems & Databases @ Dune
- github.com/msf || miguel@dune.com



Intro to Dune

**We are on a mission to
make crypto data accessible**

Dune, a community data platform



[Dune.com](https://dune.com) is a platform for querying public blockchain data & building beautiful dashboards

The screenshot shows the Dune.com web interface. On the left, there's a sidebar with navigation options like 'Essentials', 'Raw blockchain data', 'Decoded projects', 'Spells', and 'Community'. The main area features a SQL query editor with the following code:

```
1 WITH
2 users_weekly AS (
3 SELECT
4 DATE_TRUNC('week', block_time) AS block_date,
5 COUNT(DISTINCT "from") AS active_users_count
6 FROM
7 arbitrum_transactions
8 GROUP BY
9 1
10 ORDER BY
11 1 NULLS FIRST
12 )
13 SELECT
14 block_date,
15 active_users_count,
16 SUM(active_users_count) OVER (
17   ORDER BY
18   block_date NULLS FIRST
19 )
20 FROM users_weekly
```

Below the editor is a dashboard titled 'Arbitrum Active Users Weekly' showing a bar chart of active users from May 2021 to April 2023. The chart shows a steady increase in users over time, with a significant spike in early 2023.



This dashboard provides a comprehensive overview of Ethereum staking. It features several key metrics:

- ETH Staked:** 22,803,007 ETH Staked
- Validators:** 712,594 Validators
- Net Flow - Since Shanghai:** 3,499,782 ETH
- Percentage of Staked ETH:** 19.14% Staked Share of ETH Supply
- Lido Staking Market Share:** 31.54% Staked Through Lido
- Net Flow (Excl. Rewards) - Since Shanghai:** 4,636,856 ETH

The 'ETH Stakers' table is ordered by amount staked:

#	Entity	Category	ETH Staked	Validators	1MΔ	1MΔ%	6MΔ	6MΔ%	Marketshare	ETH Earned	Last Dep
1	Lido	Liquid Staking	7,191,558	225k	1%	15%	56%	31.5%	296k	2023-06	
3	Coinbase	CEX	2,247,840	70k	-4%	-2%	4%	9.9%	169k	2023-06	
4	Binance	CEX	1,230,400	38k	3%	10%	25%	5.4%	96k	2023-06	
5	Kraken	CEX	777,217	24k	0%	1%	-3%	3.4%	117k	2023-06	
6	Figment	Staking Pools	759,296	24k	3%	51%	88%	3.3%	16k	2023-06	
7	Rocket Pool	Liquid Staking	722,729	23k	2%	26%	121%	3.2%	16k	2023-06	
9	Staked.us	Staking Pools	636,071	70k	6%	1%	1%	2.9%	20k	2023-06	

A pie chart titled 'ETH Stakers Ordered by Amount Staked' shows the market share distribution, with Lido being the largest stakeholder at 31.5%.



The Challenge

Dune API & Developer use cases

Feature Request

***"Please support larger
Query Results"***





What it really means..

1. Pagination on Query-Results
2. Remove 1GB Query-Result limit
3. Filters (on rows & columns)
4. Sampling, sorting, ...
5. ?



The Challenge

Systems Architecture

- Cached Query-Results
- Max size 1GB
- Powered by *Trino*

Requirements

- *Query* Query-Results
- Latency < 100 milliseconds
- *Dirt cheap cost/query*



Dune API

+

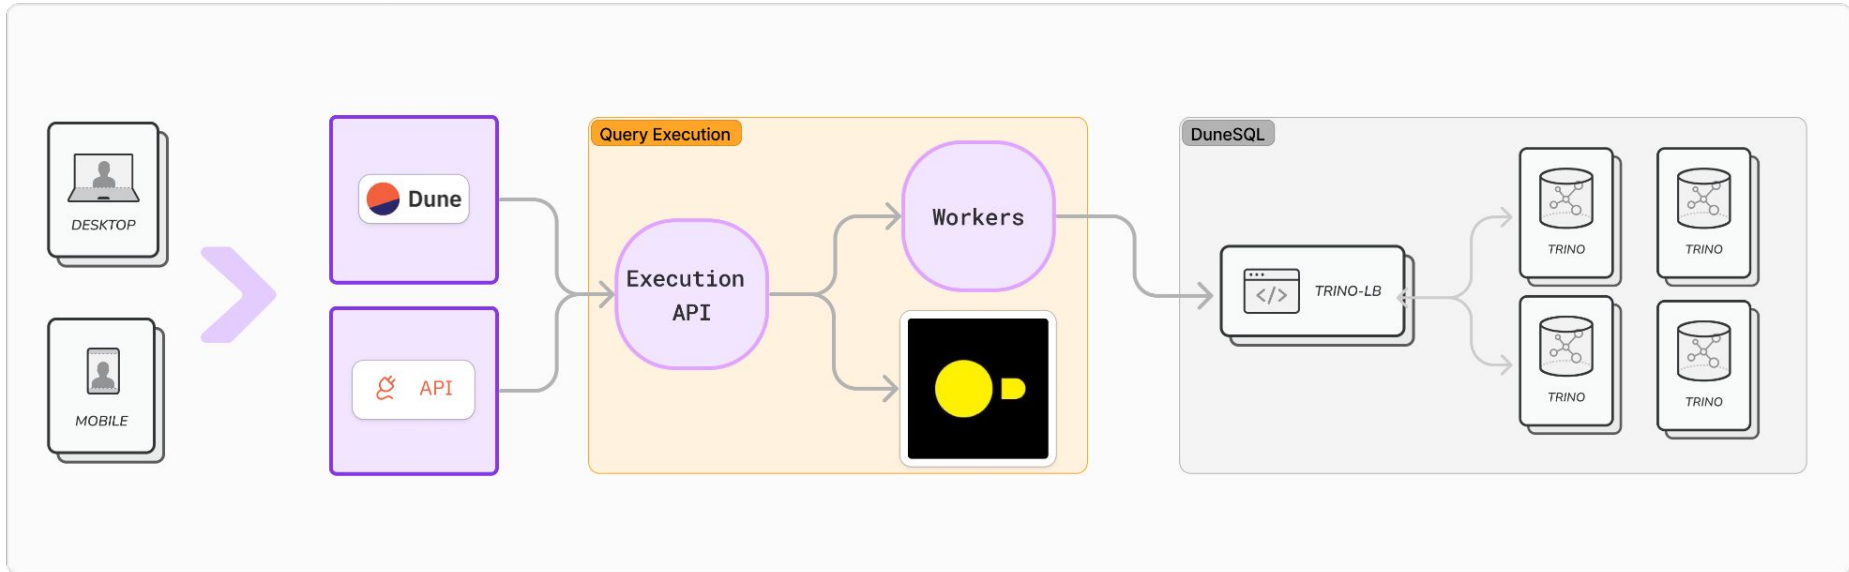


DuckDB

Solution: DuckDB

1. DuckDB on Query API Servers
2. Load Query-Result on DuckDB
3. Query DuckDB

Dune Query Execution Systems' Architecture





Technical Details

Bets Taken

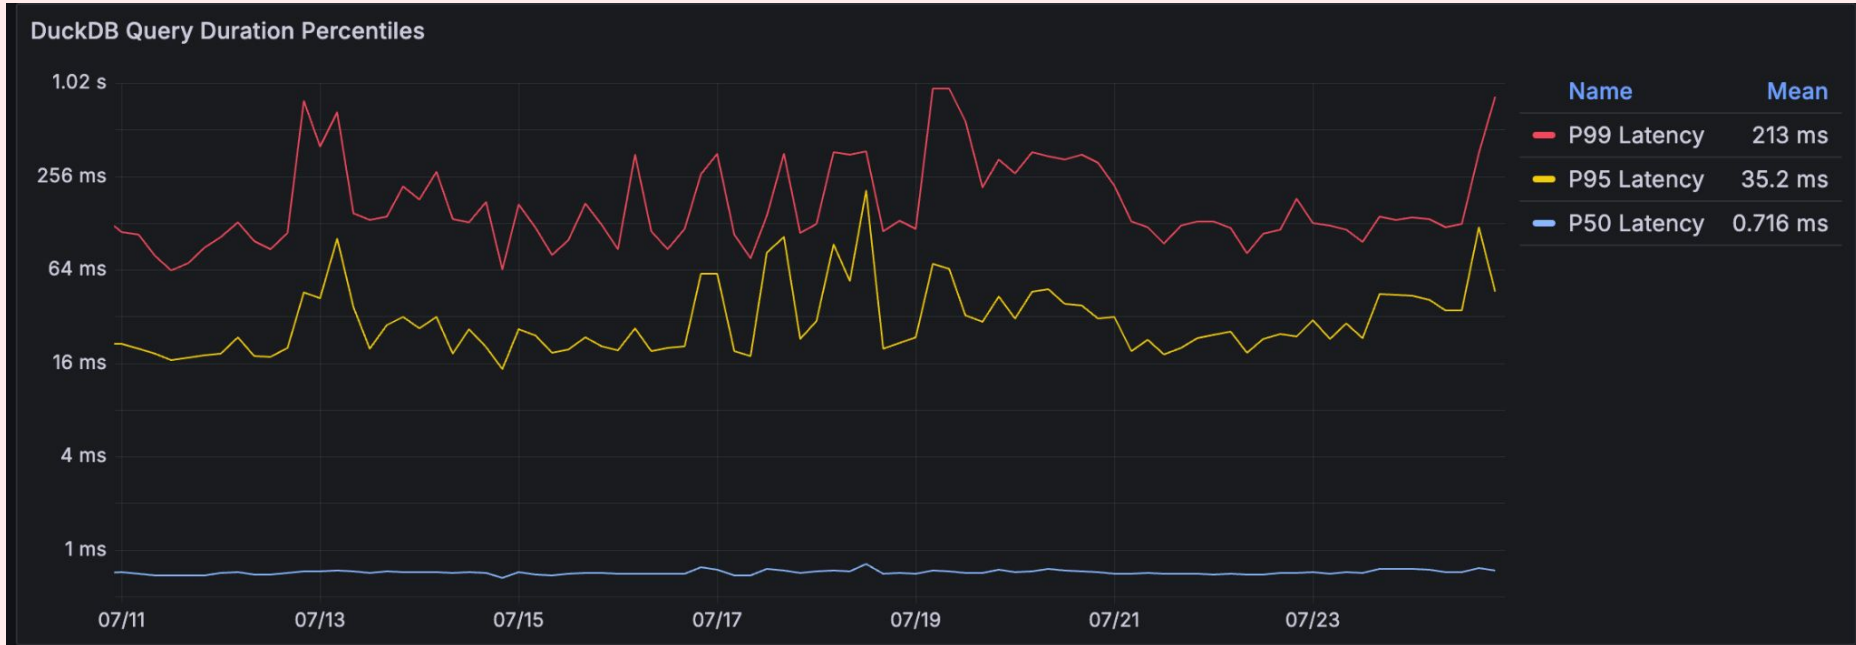
- Fast loading of query results (JSON.zstd)
- Handle \pm million “tables” on DuckDB
- Hit Rate is >90%

1. REST API is *not* SQL
 - a. Control the query complexity
 - b. Translate to SQL
2. DuckDB embedded in Servers
 - a. go-duckdb
3. Local storage as LRU cache
4. Support JSON & Parquet Query Results
5. Migrate Query Results to Parquet

Performance Metrics

Cache Query-Result on DuckDB

- Hit Rate: $\pm 95\%$
- Load on miss P99 ± 230 millis





Impact on Users

New APIs & functionalities

- Pagination, sampling, filters
- Improved dune.com
- Room to grow..

Every Dune Query

- *Is now a Data API ..*
- .. that can be queried
- .. with low latency!



Technical *Gotchas*

1. DuckDB WAL checkpointing -> global Lock ([#9150](#))
 - a. Max Wal Size: 2TB
 - b. Restarts take +2 hours
2. `SELECT * from X LIMIT 100 OFFSET 2 MILLION` ([#11261](#))
 - a. Non-constant times, +2secs latency hit ?
3. AWS + S3 + k8s + ServiceAccount Role
 - a. expired credentials? AWS plugin bugs?
4. Type mapping between Trino, JSON, Parquet, DuckDB



Conclusion

Impressed w/ DuckDB

- **Great Foundational Tech**
- Ease of use & integrate
- Fast & lightweight
- Lots of potential

Wishes & Learnings

- No global lock?
- Limit CPU secs per query
- Use Arrow APIs?
- Alternative to DataFusion?



Questions ?

dune.com



Thank you!

dune.com