

Scalable Cloud-Native

How to Build a Streaming Database (in Three Challenging Steps)

Frank McSherry, Materialize

What is a **Streaming Database?**

Streaming Database

The user is able to trade off *when* their work is done, at data ingestion or query time, through their use of

CREATE VIEW and **SELECT**

Same SQL, same scale-out dataflow infrastructure.

Demo (Streaming)



```
-- zsh
mcsberry@gallustrate ~ % psql "postgres://mcsberry%40materialize.com@51q8veosv9gow5ly1nwc488db.us-east-1.aws.staging.materialize.cloud:6875/materialize?sslmode=require"
```

```
datacouncil.sql X
1  -- Auction data as `auctions` and `bids`.
2  CREATE SOURCE auction_house
3  FROM LOAD GENERATOR AUCTION (TICK INTERVAL '1ms')
4  FOR ALL TABLES
5  WITH (SIZE = '3xsmall');
6
7  -- Brief data exploration.
8  SELECT COUNT(*) FROM auctions;
9  SELECT COUNT(*) FROM auctions WHERE end_time > mz_now();
10
11 -- Collect bids that are for currently active auctions
12 CREATE VIEW active_bids AS
13 SELECT
14     bids.id,
15     bids.buyer,
16     bids.auction_id,
17     bids.amount,
18     bids.bid_time
19 FROM bids
20 JOIN auctions ON bids.auction_id = auctions.id
21 WHERE auctions.end_time > mz_now()
22     AND bids.bid_time + INTERVAL '10 seconds' > mz_now();
23
24 -- Bids that are beating an active bid.
25 CREATE VIEW out_bids AS
26 SELECT
27     a1.id,
28     a1.buyer,
29     a1.auction_id,
30     a1.amount,
31     a1.bid_time,
32     a2.buyer AS other_buyer,
33     a2.amount AS other_amount
34 FROM active_bids a1, active_bids a2
```



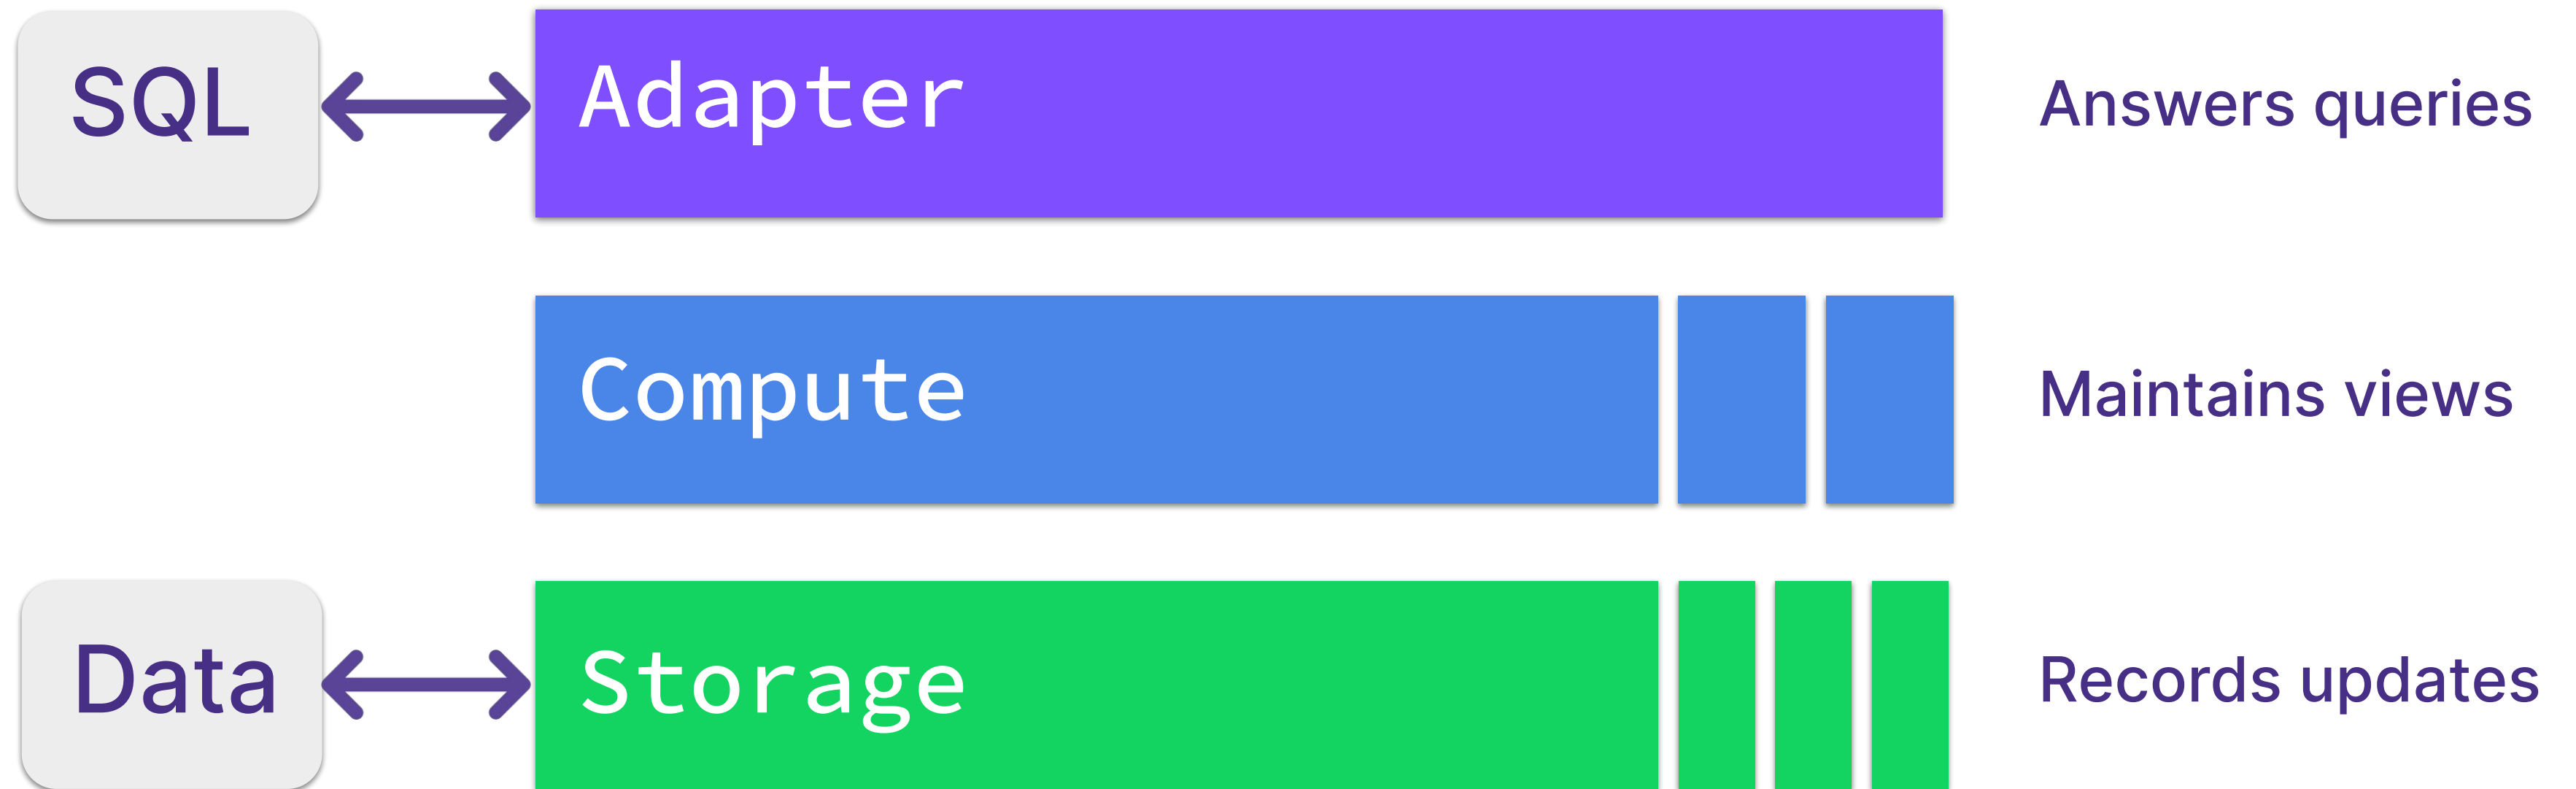
What is a

Scalable Cloud-Native


Streaming Database?



A Scalable Cloud-Native Streaming Database



Storage Layer



Records updates


Challenge:
durability

Data arrive as Change Data Capture (CDC) streams:
PostgreSQL, Debezium, Kafka Upsert, ...

... and are then written in CRDB/S3 as
“time varying collections”



Compute Layer




Maintains views

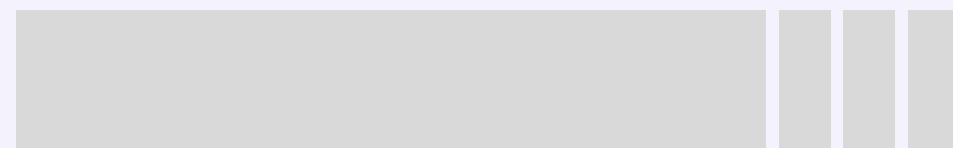
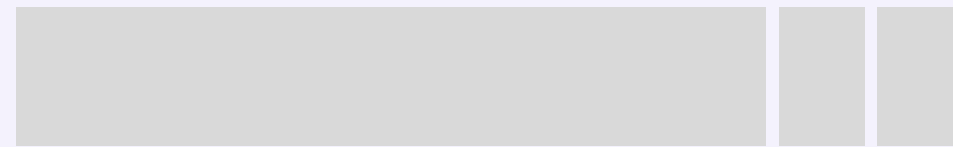
Challenge:
performance

SQL queries translate to data-parallel dataflow:
from CDC input to CDC output.

Output updates correspond exactly to inputs.



Adapter Layer




Answers queries

Challenge:
consistency

SQL commands need timestamps.

For strict serializability:
timestamps must increase.



A Scalable Cloud-Native Streaming Database



Demo (Scaling)



...e.com@51q8veosv9gow5ly1nwc488db.us-east-1.aws.staging.materialize.cloud:6875/materialize?sslmode=require					
1679843574001	t	\N	\N	\N	\N
1679843575001	t	\N	\N	\N	\N
1679843575395	f	-1	548298	3612	88
1679843575395	f	-1	548298	2089	95
1679843575395	f	-1	548298	2468	65
1679843575395	f	-1	548298	3831	49
1679843576001	t	\N	\N	\N	\N
1679843577001	t	\N	\N	\N	\N
1679843577399	f	-1	548623	382	35
1679843577399	f	-1	548623	134	54
1679843577399	f	-1	548623	668	63
1679843577399	f	-1	548623	724	55
1679843577399	f	-1	548623	3563	56
1679843577399	f	-1	548623	1012	85
1679843578001	t	\N	\N	\N	\N
1679843579001	t	\N	\N	\N	\N
1679843580001	t	\N	\N	\N	\N
1679843581001	t	\N	\N	\N	\N
1679843582001	t	\N	\N	\N	\N
1679843583001	t	\N	\N	\N	\N
1679843584000	f	1	550953	1083	86
1679843584001	t	\N	\N	\N	\N
1679843585001	t	\N	\N	\N	\N
1679843586001	t	\N	\N	\N	\N
1679843587001	t	\N	\N	\N	\N
1679843588001	t	\N	\N	\N	\N
1679843589001	t	\N	\N	\N	\N
1679843590001	t	\N	\N	\N	\N
1679843591001	t	\N	\N	\N	\N
1679843592001	t	\N	\N	\N	\N
1679843592432	f	-1	550953	1083	86
1679843593001	t	\N	\N	\N	\N
1679843594000	f	1	552514	794	37
1679843594001	t	\N	\N	\N	\N
1679843595001	t	\N	\N	\N	\N

...e.com@51q8veosv9gow5ly1nwc488db.us-east-1.aws.staging.materialize.cloud:6875/materialize?sslmode=require

materialize=>



A Scalable Cloud-Native Streaming Database



Thank you!

Register for access at
materialize.com

Visit us at booth #17

