



QuasarDB

**Drinking from the firehose:
streaming ingest benchmarks**

Preliminary report

The Problem



The Problem



How fast can my system ingest data?

How long does it take for an update to be visible?



Market Data



Repository

Queries



What is the storage efficiency?

How does ingest affect queries?

One challenge: Latency vs Throughput



The lower the latency for a specific update, the more negative impact on overall performance

In other words, low latency updates are not cooperative

- Writing 10 bytes is as fast as writing 1,000 bytes
- You can't be faster than the network
- A database has a maximum parallelization capability

In practice you want to batch your event updates as much as you can, depending on the latency you can afford



Another challenge: Performance vs Durability



The more durability you want, the higher the impact on performance

Disk is always slower than RAM

If you need data to be committed to disk, you need to wait on the storage

If data is only present in memory, a power loss will incur a data loss

Replication can mitigate the risk

In practice you decide which amount of data you are willing to lose in case of power failure

New storage technology (such as Intel Optane) can help



We think STAC Benchmarks will be helpful



Numbers are easy to throw around

But the devil is in the detail

We are investing in bringing it to life:

- **Contributions to test harness prototype**
- **Prototype QuasarDB implementation**
- **Preliminary results**



QuasarDB contributions to the test harness

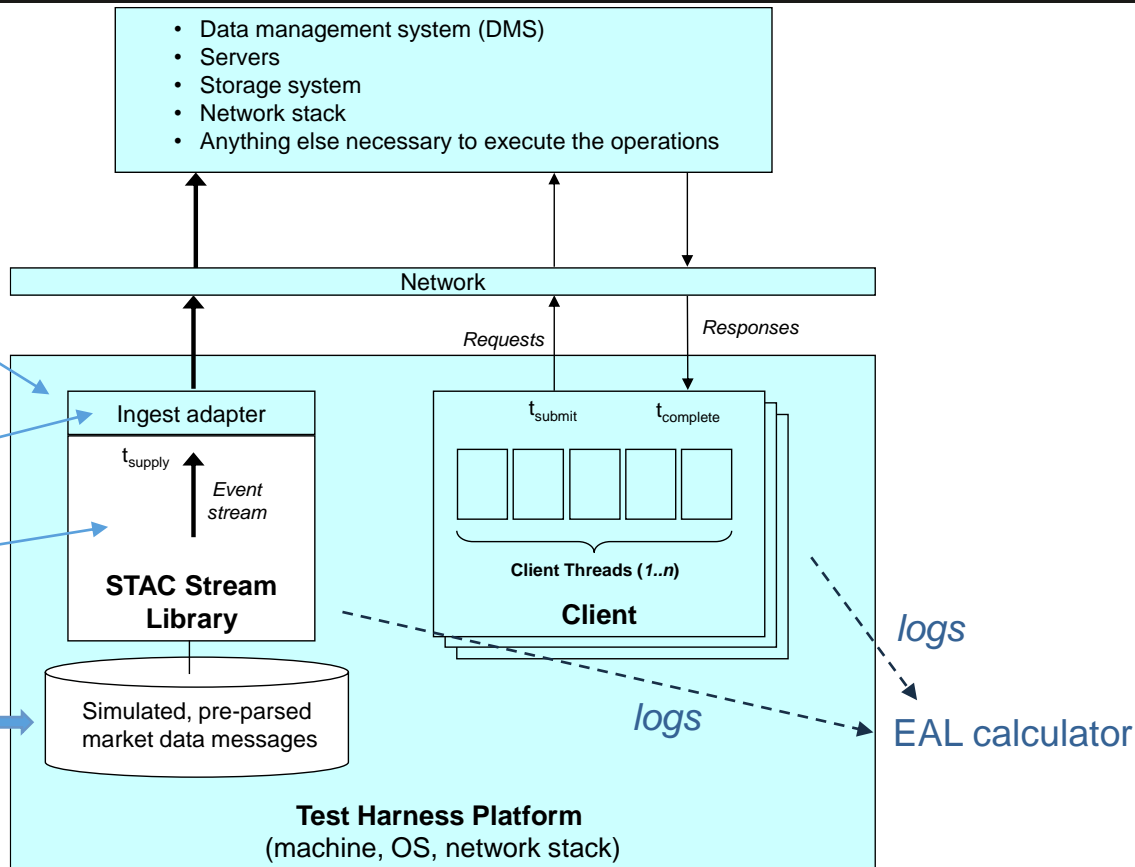


BONUS: A batch message supplier (alternative to event-driven) for troubleshooting, comparison

- Data management system (DMS)
- Servers
- Storage system
- Network stack
- Anything else necessary to execute the operations

Stub adapter
STAC Message Access Library (parses message structs)
Mods to STAC Stream Library

C++ message file generator with realistic field values



Key benchmark enabler: Data generator



No message content had been created yet for this use case

We followed the STAC-M3 record and field structures

However, STAC-M3 does not specify behavior for the data values

Values matter (e.g., for compression).

So we created naive simulation of stock behavior

- Variation is random, but values are not random
- A trade matches an actual bid
- Events are time ordered

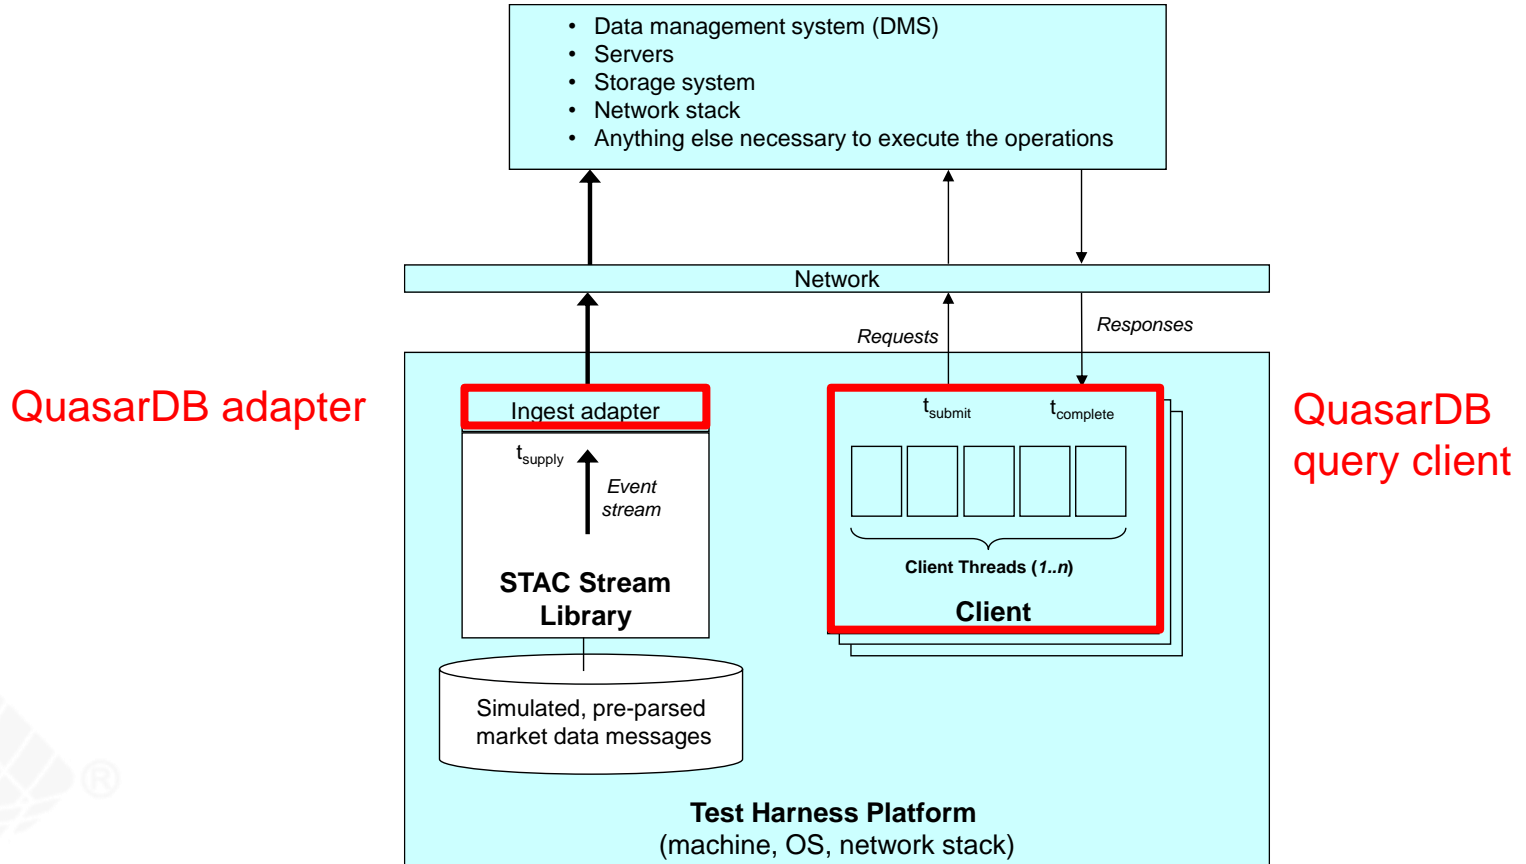
Followed STAC-M3 pattern for quotes/trades, high/low-volume symbols, etc.

~ 490K records per day

Had to write the generator in C++ for performance

Currently: 2 minute for 1 day of data (around 22 GiB)

QuasarDB implementation



Goals and non goals



Goals

Prove that it works

Get the community involved

Get some first reference numbers

Non-goals

Have the “best” possible numbers

Have auditable results

Test all possible configurations



Preliminary results (NOT STAC BENCHMARKS)



One day of ingest data

“High latency” configuration – AWS

QuasarDB configured for queries performance and ingestion efficiency

Server – m5ad.2xlarge, Client – r5d.8xlarge

Maximum ingest speed from 1 thread: 1,630,157 events per second

Fixed rate ingestion – 1,000,000 events per second

Average update latency

- 1,268 ms

MKTSNAP – 10 threads

- During ingest: 823 ms
- After ingest: 671 ms



Improvements?

STAC Stream library

Load more than one file

Test harness

Test with more queries

Test with more data

QuasarDB

Optimize setup

Run in “low latency” mode





Call to action

Get involved in Working Group

Download and test the harness

Write an adapter for a database you'd like to test

QuasarDB Community Edition is available free of charge at <https://www.quasardb.net/>

Suggest improvements

What would you like to see?

What do you think we could do better?





QuasarDB

THANK YOU

www.quasardb.net



USA

222 BROADWAY – 19TH FLOOR
NEW YORK
NY 10038



UK

40 BANK STREET
CANARY WHARF
LONDON E14 5NR



FRANCE

24, RUE FEYDEAU
75002 PARIS