



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

Peter Lankford  
Founder and Director, STAC

[peter.lankford@STACresearch.com](mailto:peter.lankford@STACresearch.com)

# Background

- Ingest and analysis of streaming data keep coming up in benchmark discussions
- Time to jump in

# Two projects to discuss

- Benchmarks of 3Forge AMI data visualization platform
- Prototype of database ingest tests

(Some tests of the two look similar but are actually quite different)

# Why did we benchmark a display platform?

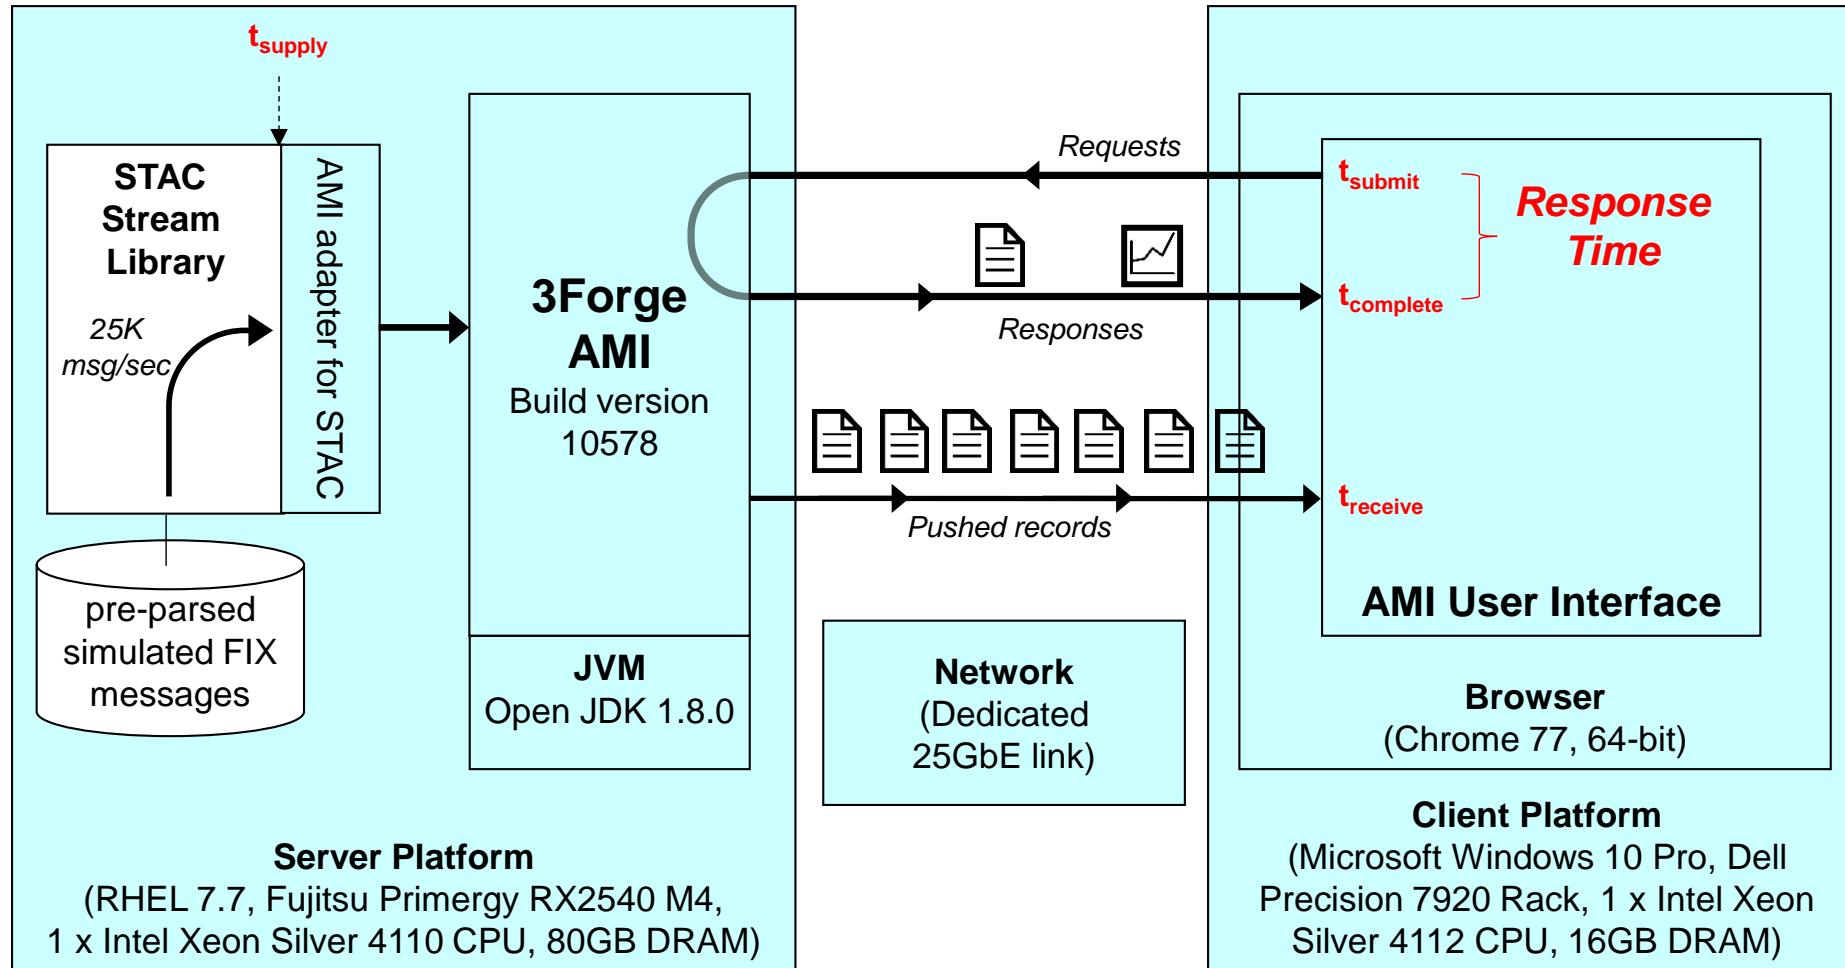
- Despite the automation trend in capital markets, human analysis is still crucial
  - In fact, automation increases the need for it (e.g., supervision of algorithms)
- Human analysis requires visual tools
  - Charts, tables, etc.
  - Dynamic exploration
- Some customers tell us:
  - Front-office tools not well suited to broader enterprise use
  - Big-name visualization tools can't cope with streaming data
- Many firms build custom solutions
- 3Forge offers a product: AMI

# Testing project

- Basic questions:
  - Query response time
  - Data freshness
- Finding the max ingress capacity was not a project objective
  - Studied the above at a reasonable fixed ingress rate
  - 25,000 messages per second = 700,000 database fields per second
- Basic setup
  - 1 server (low spec Linux box)
  - 1 client (Chrome on Windows)
  - Direct network connection

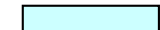
[www.STACresearch.com/3FG190902](http://www.STACresearch.com/3FG190902)

# Test setup



-50 ms < error of means < + 50ms

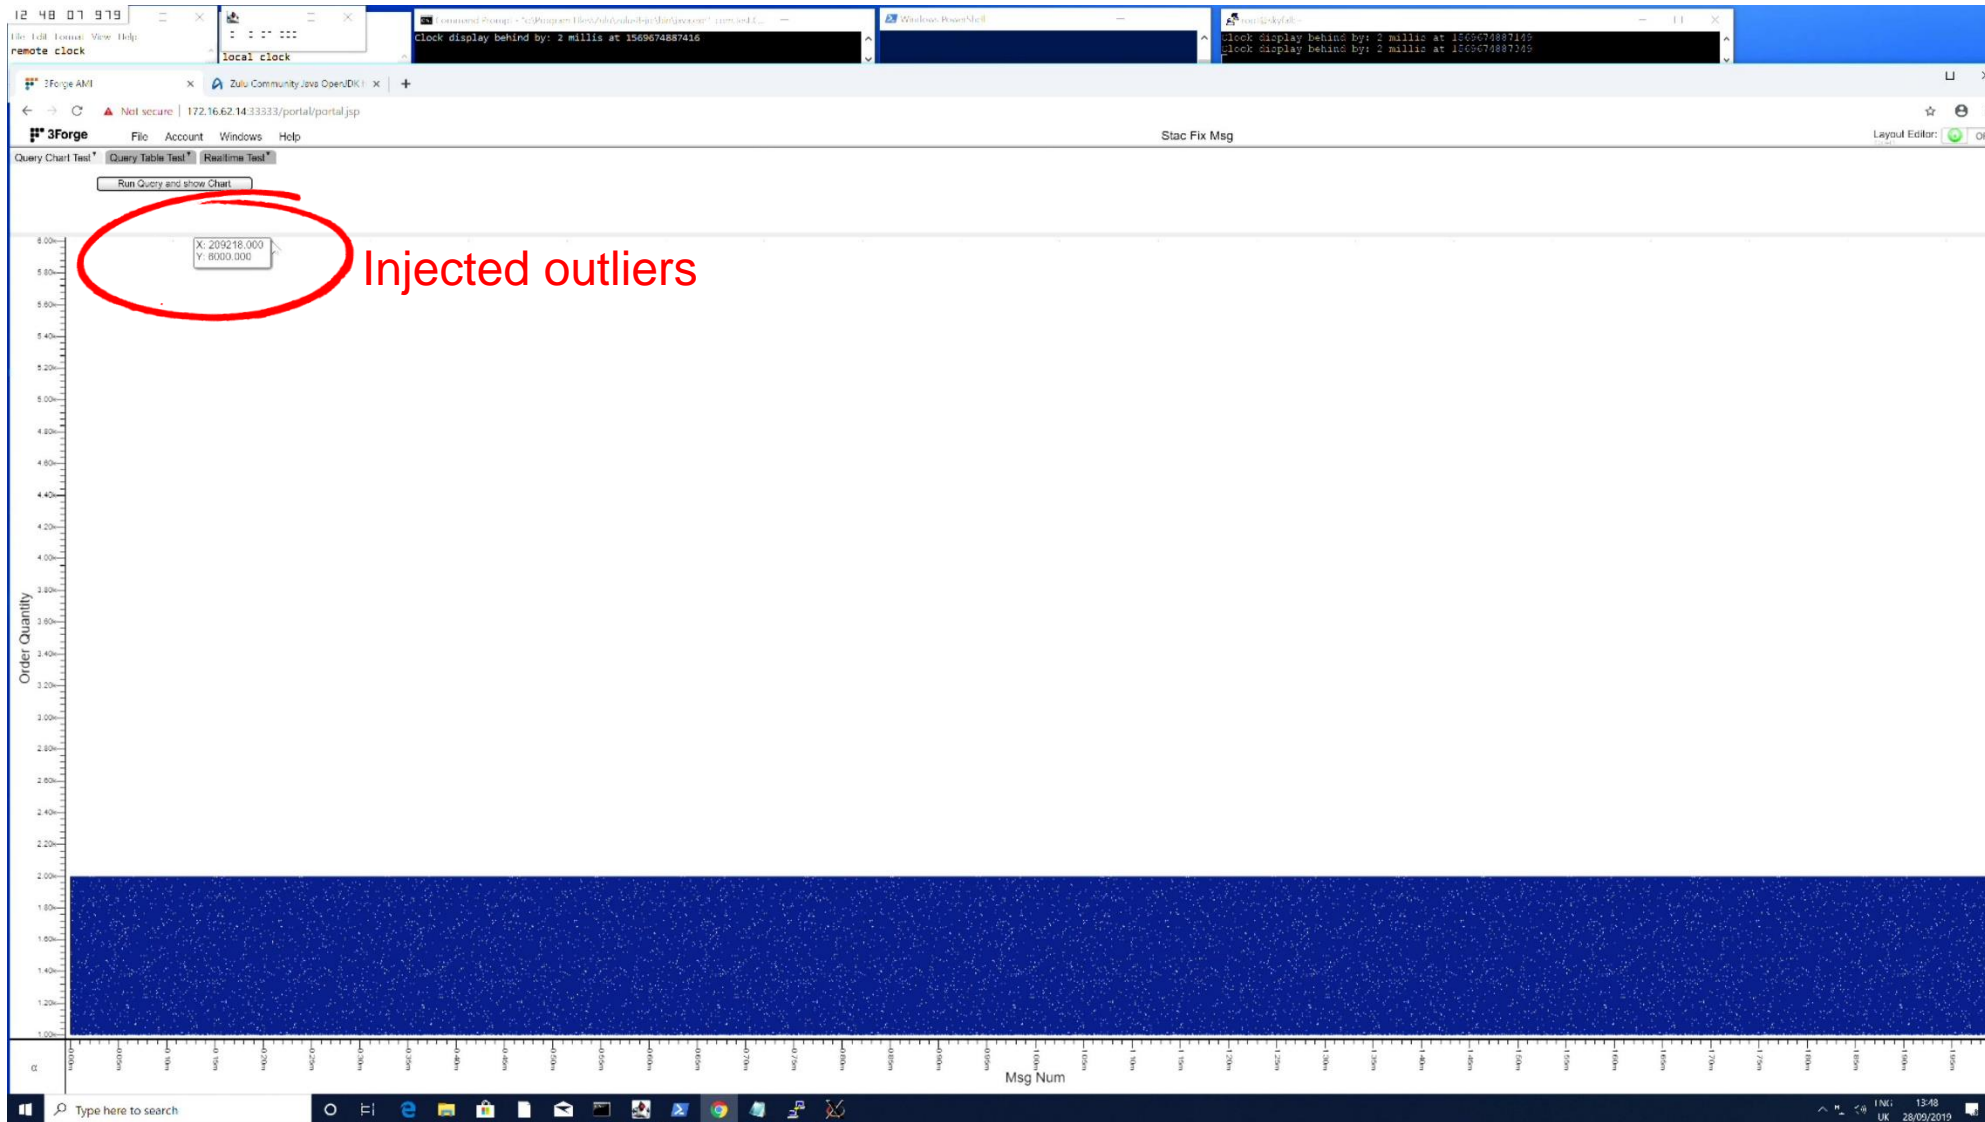
**Availability Latency**  
(data "freshness")

 = part of the stack under test (SUT)





# An example of a chart in this project





# Key findings

- ***Response times:***
  - Average response time for 40 table rows, during ingest: 317 milliseconds
  - Average response time for 40 table rows, not during ingest: 273 milliseconds
  - Average response time for a chart with 2 million data points, not during ingest: 3.67 seconds
- ***Data freshness:***
  - Average age of most recent data returned in request/response queries during ingest: 149 milliseconds
  - Average age of most recent data displayed in an auto-updating table during ingest: 247 milliseconds

# Two projects to discuss

- Benchmarks of 3Forge AMI data visualization platform
- Prototype of database ingest tests

(Some tests of the two look similar but are actually quite different)

# Motivation

- Use case categories
  - Collecting data for tick histories, like the historical data represented by STAC-M3
  - Realtime analysis of streaming data (e.g., surveillance, regime-change detection, ...)
- We hear interest in many timeseries databases. There are a lot to choose from:
  - At least 10 open source
  - At least 15 proprietary software
  - At least 6 DBaaS
  - Not to mention stuff developed in-house!
- The current STAC-M3 assesses the speed of analysis on historical data
- But firms confront two questions even before they get to analytics performance...

# Two big questions

1. Can we consume the realtime data to build up a history in the first place?
    - Ingest capacity
  2. If we want to do analysis in real time, how realtime are the data in my queries?
    - Availability latency
- And a third question (perhaps): What impact does realtime ingest have on historical query response time?
  - We have heard an interest from several user firms in having benchmarks to answer these questions

# Traditional benchmark specification approach (waterfall)

1. WG meets to decide requirements (user firms have the vote).
2. Repeat #1 until requirements set.
3. Vendors do implementations.
4. A motivated vendor engages STAC for an audit.
5. Vendor and STAC discover specification questions/problems.
6. Wild scramble to address them with the WG.

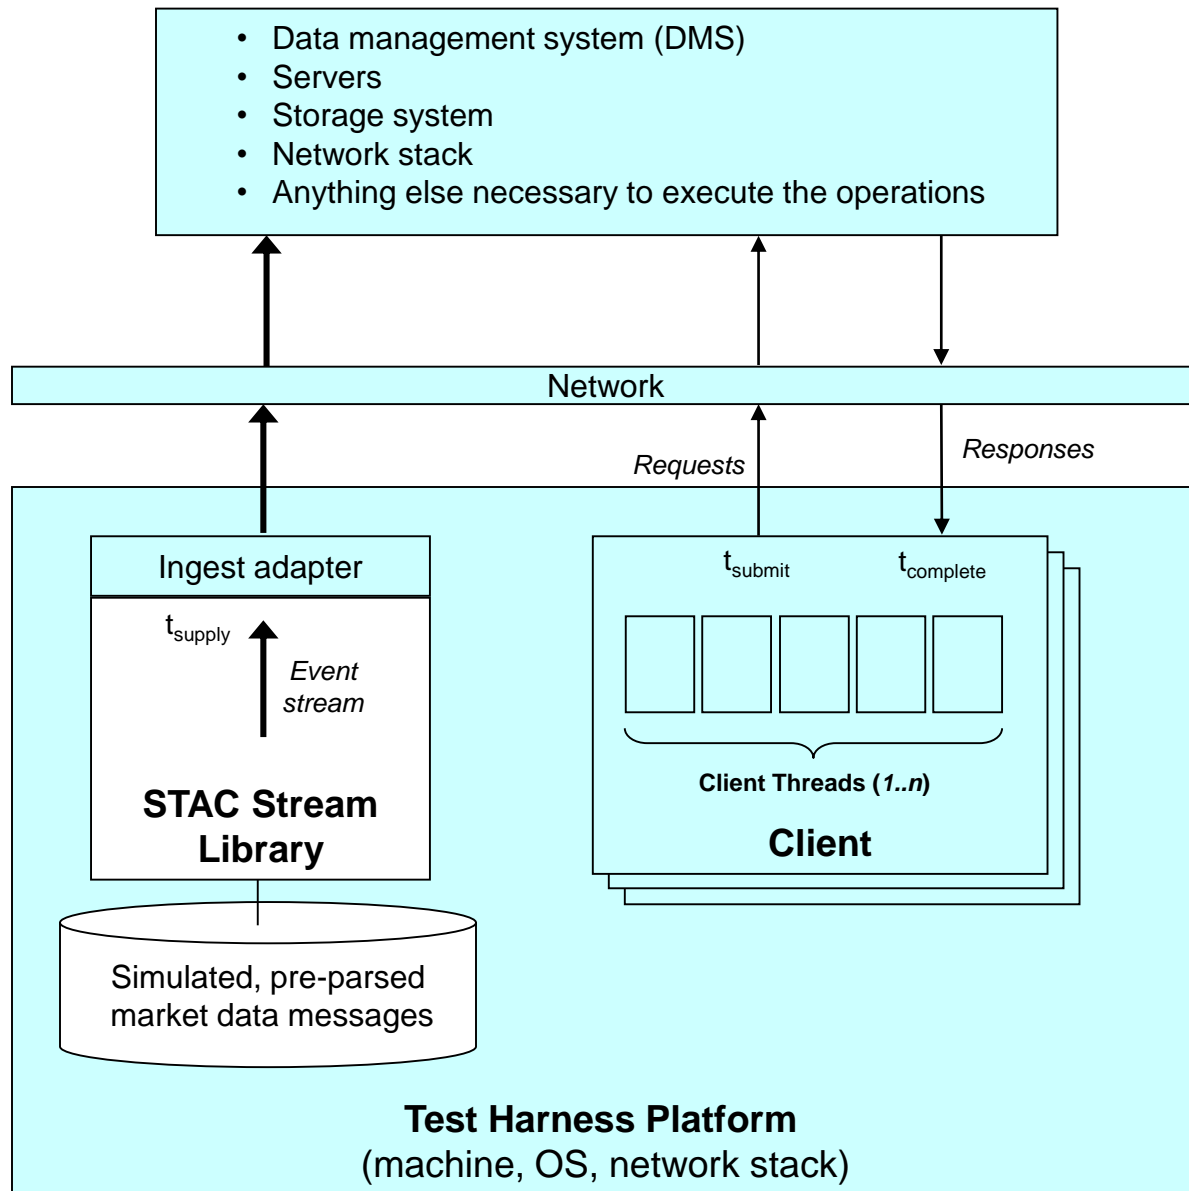
Not the fastest process in the world. (And no shortage of stress.)

# Trying a new “agile” approach

1. WG members express interest in tests, broadly defined.
2. Motivated vendor and STAC work on prototype.
3. Vendor and STAC present results to WG as a starting point.
4. WG iterates requirements (user firms have the vote)
5. Initial vendor and other vendors iterate on implementations.
6. By the time specs are finalized, one or more vendors have implementations ready.

This should yield good benchmark specs faster

# Test setup for ingest benchmarks

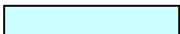



## Prototype benchmark specifications:

- Ingest capacity
- Estimated availability latency (EAL)
- As-of query performance (market snapshot)
- No persistence checks

## Key questions for the Council:

- Feedback on the details?
- Treat historical and realtime as separate?
- How to treat persistence?
- Complex queries on realtime?
- Continuous queries/filters?
- Minimum viable benchmark specs?

 = part of SUT

 = STAC component



# Motivated vendor in this case: QuasarDB

- You heard about QuasarDB in the Innovation Roundup earlier
- QuasarDB have contributed:
  - Tools to support the benchmarks
  - An implementation of proposed benchmarks
- Edouard Alligand, Founder & CEO presented those in NY
- He could not be here today, so I will give his presentation
- Just to be clear: **Results are not STAC Benchmark results**
  - The specifications have not been vetted by a STAC Working Group
  - The results are unaudited by STAC