Benchmarking streaming time series

Introducing the first step: STAC-M3 Jalua Suite

Peter Nabicht
Head of Strategy, STAC

peter.nabicht@STACresearch.com

Motivations

- Growing need to analyze streaming data
  - Market data (continues to expand)
  - IT telemetry (performance monitoring, anomaly detection)
  - Industrial IOT (manufacturing, telecom, smart grids)
- Proliferation of time series solutions
  - Number has ballooned to over 50 in the last few years
- This situation calls for benchmarks
- We've created an interest group to:
  - Discuss common concerns across use cases
  - Spin off use case specific benchmarks

www.STACresearch.com/streaming

Context for financial firms

- Firms care about streaming data from:
  - The markets
  - Their own transactions
  - Their systems
- Many firms want to use the same technology to satisfy all three needs
- This calls for a common benchmark framework

Where we are starting

- The first deliverable is for a market data use case
- Three existing STAC-M3 benchmark suites use historical market data
  - These suites continue providing useful insights after more than a decade
  - See Michel's update for the latest public results
- Now the STAC-M3 Working Group is adding a fourth suite, involving streaming market data:
  - STAC-M3 Jalua

STAC-M3 Suites

Suite	Dataset	Impediments to caching/pre-loading	Storage I/O	Network I/O*	Compute burden	Concurrent users
Baseline (Antuco)	Historical (~4TB)	Yes	Mostly high intensity reads	Neglible	Low to moderate	Varies
Small in-memory (Shasta)	Historical (~4TB)	No	Mostly high intensity reads	Neglible	Low to moderate	Varies
Scale (Kanaga)	Historical (theoretically unlimited TB)	No	Mostly high intensity reads	Neglible	Low to moderate	Theoretically unlimited
Streaming (Jalua)	Streaming ingest & historical (~400GB)	*tbd*	Potentially high intensity writes, Low intensity random reads	Potentially high	Low to moderate	Theoretically unlimited

"STAC-M3 Historical"

* Between SUT and test harness. Not necessarily within the SUT (e.g., a storage network)

STAC® SECURITIES TECHNOLOGY ANALYSIS CENTER

Status of the benchmark suite

- WG is still providing feedback on v1
- Some remaining actions to implement changes
- Four database implementations underway so far
- Expect version 1 signoff by end of next month
- You can still join the WG to weigh in
  - And help define additional benchmarks for the suite

www.STACresearch.com/m3

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

First questions for a market data solution

- What kind of ingest rates can it handle without loss?
- How well does it blend live and historical data?
- How fast is data available for querying?
- How do ingest and query load impact each other?

STAC-M3 Jalua test setup

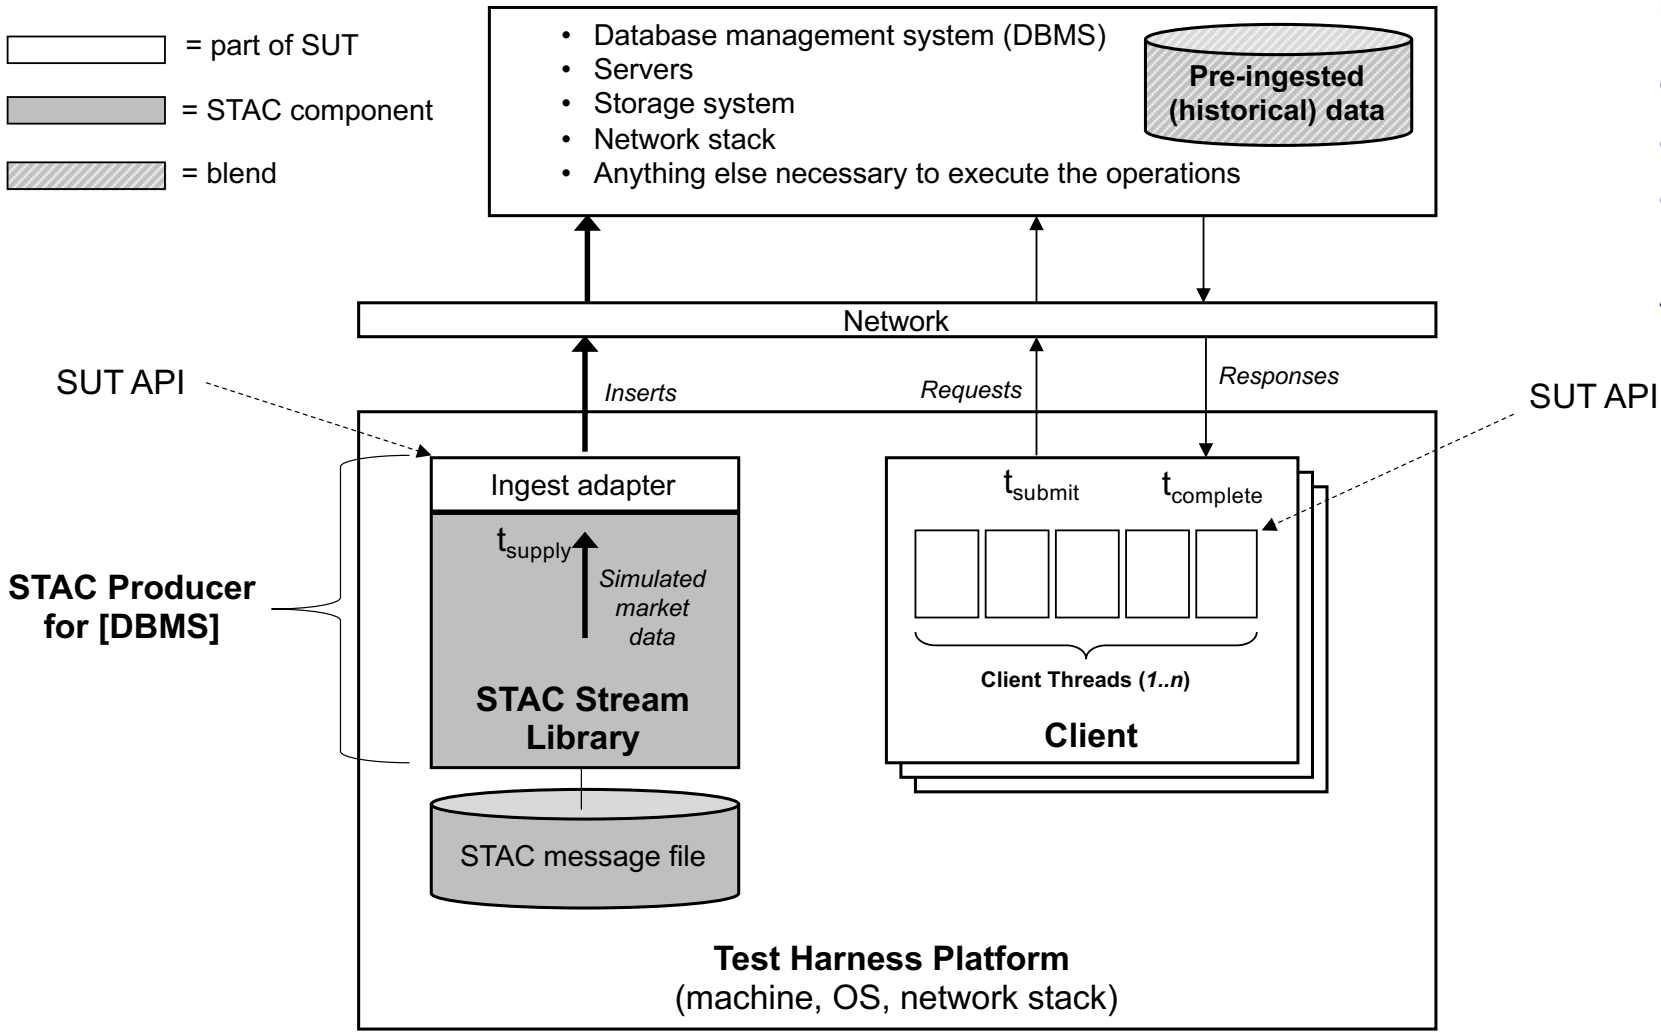

Key metrics:

- Ingest capacity*
- Estimated availability latency
- Query response time

* Relies on testing for message loss

Workload dimensions:

- Supply rate
- Supply duration
- Operation (roughly the queries)
- Number of simulated concurrent users

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

The rest of this talk: Tests and results

- Best way to understand the benchmark specs is through examples
- STAC wrote implementations for two time series databases
  - Open source
  - Commonly used
- Have interleaved preliminary results in this talk

Results are normalized to protect the innocent

Why are they innocent?

- Benchmark specs aren't finalized
- STAC configured the SUTs
- Naïve user approach.
  - STAC knew nothing about the databases.
  - No vendors were involved.
  - Hence, they probably don't represent fully optimized results.

How are they normalized?

- Message rates divided by the lowest reported message rate
- Latencies and response times divided by the lowest reported latency or response time

Same SUT platform for both databases

- Database Host:
  - Google Cloud Platform e2-standard-16 (16 vCPUs, 64 GB memory)
  - Data: 1000 GB SSD Google Persistent Disk
  - Debian Linux
- Input/output Host
  - Same as Database Host except no access to Google Persistent Disk
- Network
  - GCP Premium Tier network

NO_QUERIES

- The workload is the ingest.
- No queries running.
- Demonstrates capture capacity.
- Informative for use cases that just need capturing, not live querying.
- Also useful for any SUT as a baseline for comparison to other Operations.

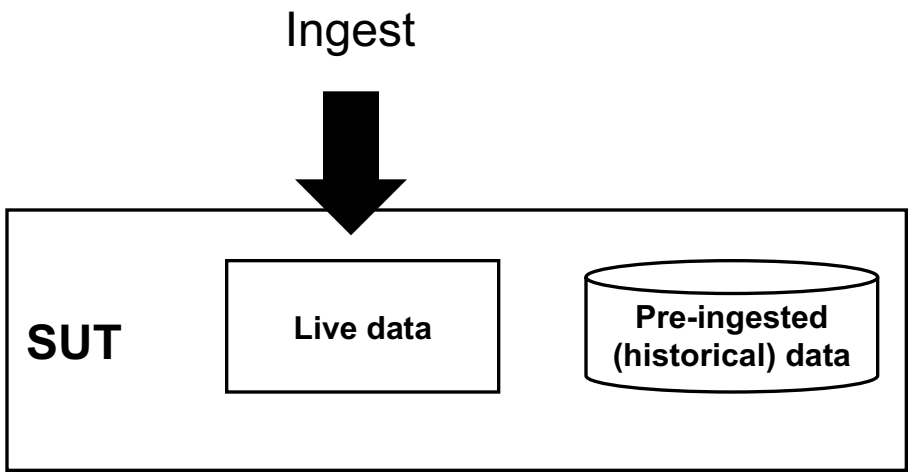

NO_QUERIES comparisons

	SUT A Rate Multiple	SUT B Rate Multiple
NO_QUERIES	2.6	2.6

- We tested to find the max ingest rate
- Coincidentally, both SUTs achieved the same rate

1MO_SNAP

- Simple as-of join on historical data. Trade & quote snapshot for 1% of symbols. Similar to MKTSNAP in STAC-M3 Historical but can't be fairly compared.
- Running with and without ingest allows impact assessments.

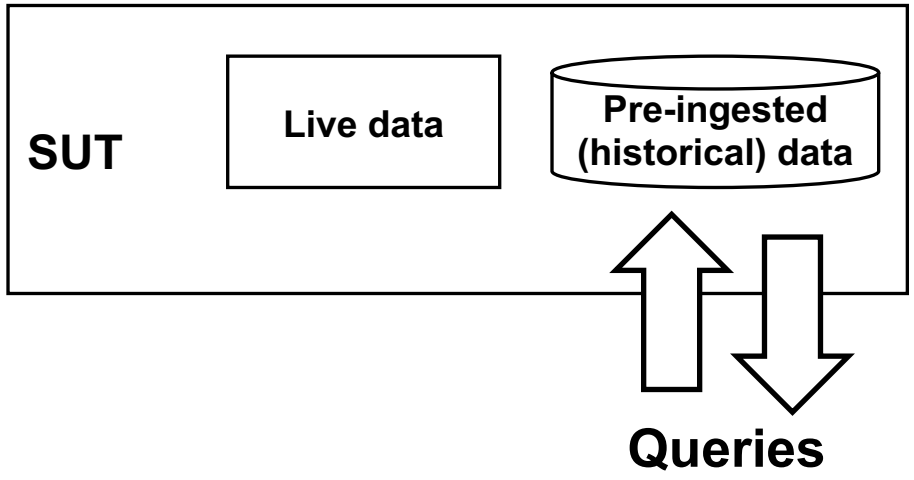

STATIC.1MO_SNAP

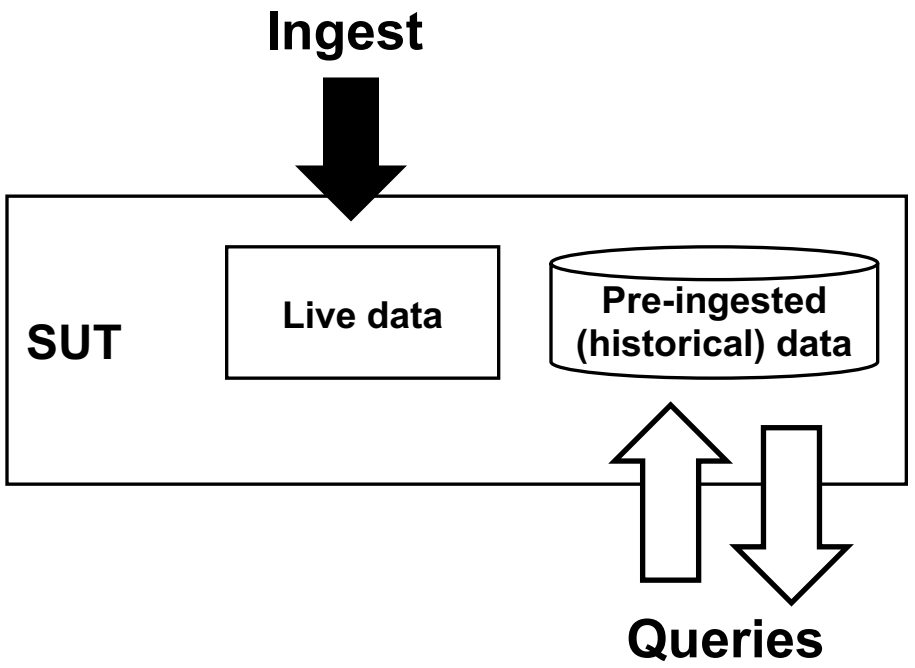

<supply_rate>.<duration>.1MO_SNAP

STAC® SECURITIES TECHNOLOGY ANALYSIS CENTER

1MO_SNAP comparisons

- Comparing response times with and without ingest gives insight into the impact of ingest on historical query performance

			SUT A response times (normalized)				SUT B response times (normalized)			
Rate	Duration	Users	Min	Median	90th	Max	Min	Median	90th	Max
--	--	10	10.40	72.93	157.49	242.42	8.85	12.57	14.39	17.60
1MULT	600S	10	9.28	82.06	215.28	295.75	7.77	10.18	11.47	11.82

- SUT A historical queries are noticeably slower during ingest
- SUT B historical queries improve slightly during ingest (may be an artifact of the test sequence)
- SUT B has both better performance than SUT A and more consistent performance, with much smaller outliers

10SEC_SNAP

- Same as 1MO_SNAP except it queries data ingested during the test run, rather than prior to test run
- One or more simulated concurrent users submit queries
- Queries are submitted repeatedly over the test run

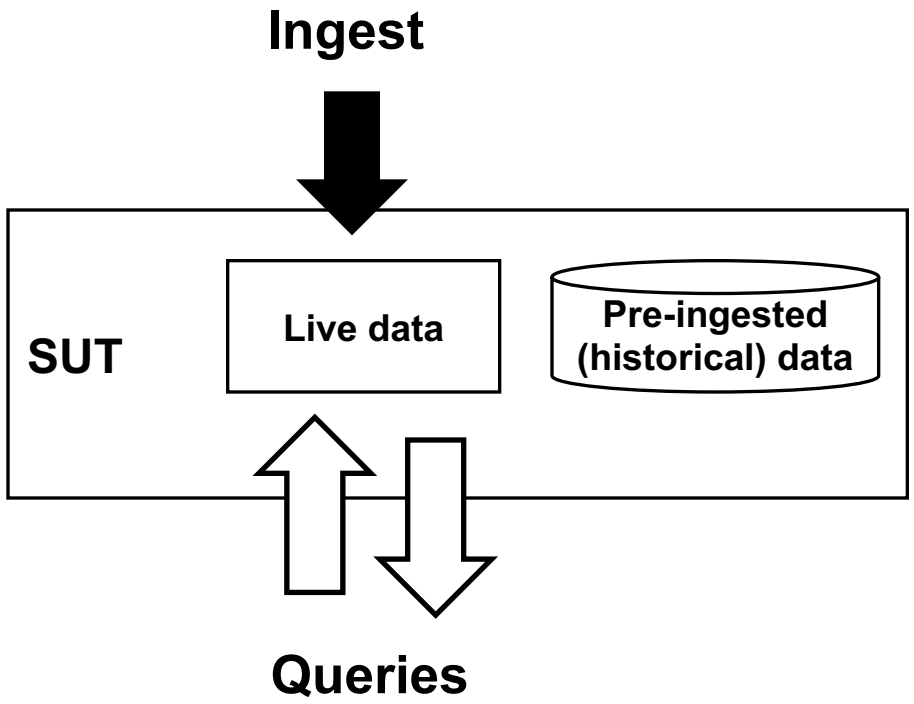

<supply_rate>.<duration>.10SEC_SNAP

10SEC_SNAP comparisons

			SUT A response times (normalized)					SUT B response times (normalized)				
Rate	Duration	Users	Min	Median	90th	99th	Max	Min	Median	90th	99th	Max
1MULT	600S	5	1	2.86	4.03	5.26	6.86	--	--	--	--	--

- SUT A was able to sustain 5 users at 1 multiplier
  - Increase the ingest rate or the number of simulated users and it failed.
- SUT B was not able to satisfy the simulated user workload at any rate
  - Due to clustered batches of outliers with the market snapshot query that prevented it from getting enough queries completed during the ingest.

Estimated availability latency (EAL)

- Measured for all tests involving ingest, except NO_QUERIES
- Measured with and without SNAP queries running
- A low overhead query to obtain the most recently ingested record
- Allows us to see the impact on latency of:
  - Ingest rate
  - Number of users (query load)

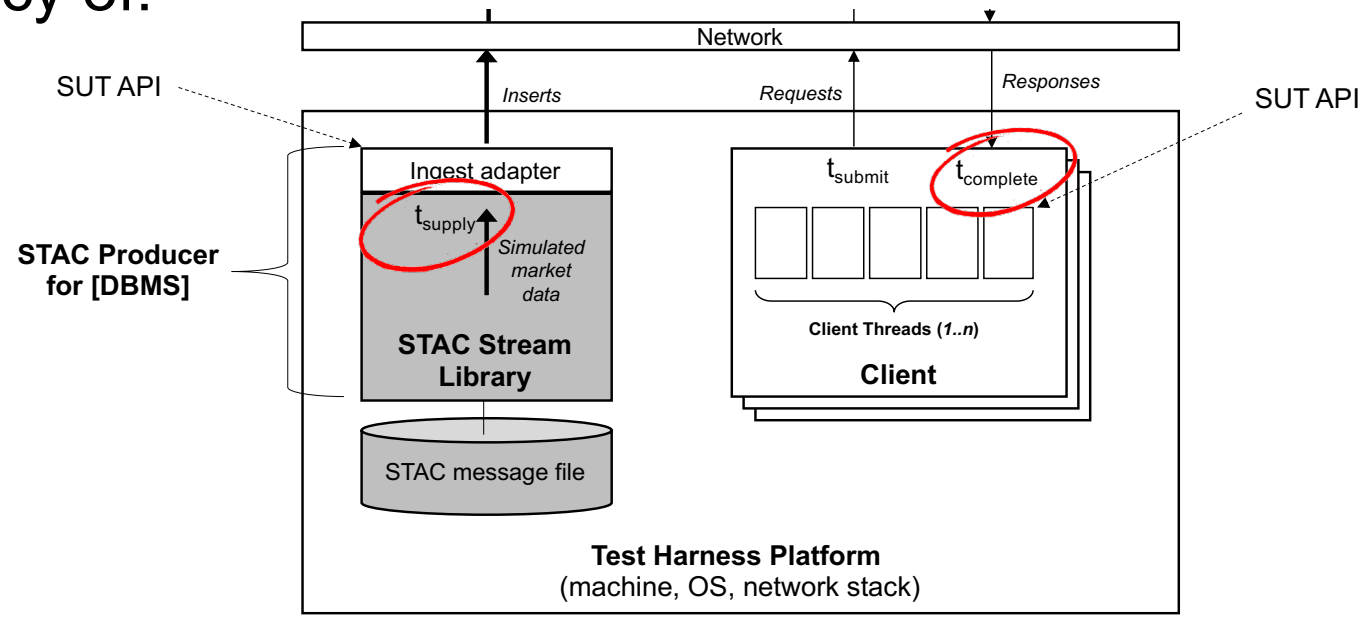

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

EAL comparisons

- EAL is run alone and along with 10SEC_SNAP to allow for insight into the impact of user workloads on EAL

			SUT A estimated availability latency (normalized)					SUT B estimated availability latency (normalized)				
Rate	Duration	Users	Min	Median	90th	99th	Max	Min	Median	90th	99th	Max
1MULT	600S	0	18.38	89.40	92.99	94.44	94.82	2.43	4.89	6.55	7.39	284.59
1MULT	600S	5	21.11	90.33	96.55	138.52	213.87	--	--	--	--	--

- SUT A's min and median stay fairly consistent with user workload, but the outliers increase greatly
- SUT B has no EAL with user workload due to the inability to complete an ingest test with user workload
- SUT B's EAL without user workload has much better performance than SUT A in most cases

Next steps

- Once benchmark specs are finalized, reports for the databases above will go into the STAC Vault.
- Let us know what you'd like to see tested
- Again, get involved:
  - Market data use cases: www.STACresearch.com/m3
  - Other use cases: www.STACresearch.com/streaming

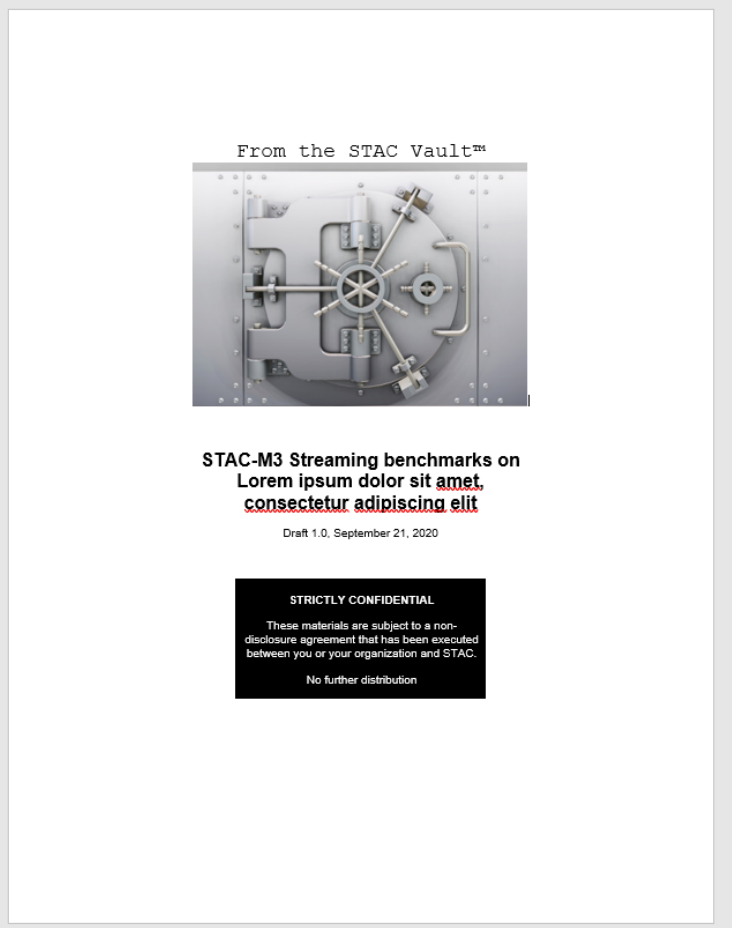

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER