6x Supermicro Petascale Servers each with 2x Intel® Xeon® 6767P, 24x Micron™ 9550 NVMe™ SSDs and 16x DDR5 RDIMMs, and KDB+ 4.1

SUT ID: KDB250929

STAC-M3™ Benchmarks

Antuco and Kanaga Suites

Test date: 29 September 2025

Release v1, 22 October 2025

STAC REPORT

These tests followed STAC benchmark specifications proposed or approved by the STAC Benchmark Community (see www.stacresearch.com). Be sure to check the version of any specification used in a report. Different versions may not yield results that can be compared to one another.

References ... 3
Summary ... 4
Vendor Commentary ... 4
Product background ... 5
Project participants and responsibilities ... 7
Contacts ... 7
Results status ... 7
Overview of the STAC-M3 Benchmark specifications ... 8
Understanding the STAC-M3 Benchmark Suites ... 9
Datasets ... 10
Metrics ... 10
Test cases ... 11
Benchmark identifiers ... 11
STAC-M3 Benchmarks in the Antuco Suite ... 13
STAC-M3 Benchmarks in the Kanaga Suite ... 16
Specification particulars ... 18
Version ... 18
Limitations ... 18
APPENDIX A: STAC-M3 ANTUCO RESULTS ... 19
Storage Efficiency ... 19
Light-Compute Benchmarks ... 19
Post-Trade Analytics Benchmarks ... 20
Research Analytics Benchmarks ... 21
NBBO Benchmark ... 22
Multi-day/Multi-User VWAB Benchmark ... 22
Chart view ... 22
APPENDIX K: STAC-M3 KANAGA RESULTS ... 32
Tabulated resuls ... 32

References

[1] *Specifications used for this benchmark (accessible by qualified members of the STAC Benchmark Community):*
STAC-M3 Benchmark Specifications, Antuco Suite, Rev Q - www.STACresearch.com/m3/antuco/revq;
STAC-M3 Benchmark Specifications, Kanaga Suite, Rev I – www.STACresearch.com/m3/kanaga/revi.

[2] *STAC Configuration Disclosure for this SUT: http://www.stacresearch.com/KDB250929.*

Summary

STAC recently performed the baseline STAC-M3™ Benchmarks on a stack involving KX's kdb+4.1 DBMS system sharded across 6x Supermicro Storage SuperServer SSG-222B-NE3X24R servers, with each system containing 24 Micron 9550 Max SSD (total of 144) and 2x Intel Xeon 6767P 64-Core CPUs. This SUT used local storage, compared with many previous that shared a share disk array.

STAC-M3 is the set of industry standard enterprise tick-analytics benchmarks for database software/hardware stacks that manage large time series of market data ("tick data"). This report highlights results from the baseline benchmark suite (code named Antuco) and the optional scaling suite (code named Kanaga).

In all, the STAC-M3 specifications deliver dozens of test results, which are presented through a variety of tables and visualizations in this report. Intel chose to highlight the following:

- This solution **outperformed** all publicly disclosed results on **19 of 24 Kanaga** mean-time response benchmarks.
  STAC-M3.β1.100T.YR{1,2,3,4,5}VWAB-12D-HO.TIME
  STAC-M3.β1.1T.{2,3,4,5}YRHIBID.TIME
  STAC-M3.β1.100T.YR{1,2,3,4,5}-12D-HO.TIME
  STAC-M3.β1.50T.YR{1,2,3,4,5}VWAB-12D-HO.TIME

- This solution **outperformed** all publicly disclosed results on **3 of 3 Antuco 50-user and 100-user** benchmarks.
  1.36x faster than a DDN based SUT (KDB221014) on STAC-M3.β1.100T.STATS-UI.TIME
  2.12x faster than an Intel Optane based SUT (KDB210428) on STAC-M3.β1.100T.VWAB-12D-NO.TIME
  1.29x faster than an Intel Optane based SUT (KDB210428) on STAC-M3.β1.50T.STATS-UI.TIME

- This solution **outperformed** all publicly disclosed results on **3 of 5 Kanaga throughput** benchmarks**.**
  1.77x faster than a Pure Storage based SUT (KDB231122) on *STAC-M3.β1.1T.3YRHIBID.MBPS*
  2.70x faster than a Pure Storage based SUT (KDB231122) on *STAC-M3.β1.1T.4YRHIBID.MBPS*
  2.51x faster than a Pure Storage based SUT (KDB231122) on *STAC-M3.β1.1T.5YRHIBID.MBPS*

Vendor Commentary

Intel provided the following comments:

> *Intel deeply values its Financial Services Industry customers for pushing the boundaries of data and performance—especially in areas like high-frequency trading, where speed, precision, and real-time analysis are essential. Their relentless drive for innovation challenges Intel to build better products. Financial firms must analyze billions of tick data points per day to respond to market changes, optimize profits, and manage risk effectively. This STAC-M3 result shows Intel Xeon 6 delivering record-setting performance and efficiency—helping trading teams act faster, capture alpha, and assess risk with confidence.*
>
> *This solution marks over a decade of STAC-M3 submissions from Intel, demonstrating continuous innovation across generations of Xeon processors and showing improvements in key metrics, including greater throughput, faster speeds, improved energy efficiency and a denser solution footprint. We extend a special thank you to Micron and Supermicro for collaborating with Intel to push the frontier of computing forward.*

Micron provided the following comments:

> *The STAC-M3 benchmark is the gold standard for independent, audited validation, showing that Micron memory and storage products deliver the superior speed, capacity, and low-latency for the financial services and high-frequency trading (HFT) industries. The combination of memory and storage tested in this report delivers high performance and low latency for accelerated market data analytics from billions of data points,*

enabling faster time to insights. Micron appreciates the ongoing collaboration with Intel and Supermicro to produce a world record solution, again raising the bar for STAC performance.

Supermicro provided the following comments:

> *Supermicro brings a 31-year legacy of innovation in enterprise, data center and cloud computing based on its building block approach which focuses on design modularity and easy customization to suit specific customer needs. Supermicro's products incorporate technologies such as the versatile Super Cloud Composer software for large-scale distributed fleet management, rack-scale and data center scale integration and installation services and end-to-end liquid cooling solutions. All of this enables Supermicro's customers to maximize power, space and density within their data center environment and achieve the fastest time-to-money possible. The Supermicro team thanks the teams at Intel and Micron for enabling this record-breaking M3 result.*

Product background

This section provides a high-level overview of the SUT in this report. A detailed STAC Configuration Disclosure for the SUT in this report is available to premium members of the STAC Benchmark Community at the same web page as this report [2]. That document provides the exact product version numbers, detailed tuning options, and other important information. Additional configuration details such as a SOS report may also be available, depending on the SUT platform.

The top of the stack under test was the benchmark implementation code (i.e., the STAC-M3 Clients and supporting scripts):

• For STAC-M3 Antuco, this was the STAC-M3 Pack for kdb+ Rev 4.1 Antuco, Compatibility Rev (I).

• For STAC-M3 Kanaga, this was the STAC-M3 Pack for kdb+ Rev 4.1 Kanaga, Compatibility Rev (Q)

Key products in the SUT included:

- STAC Pack
  - STAC-M3 Pack for kdb+, Rev I[1]
- Database
  - KDB 4.1 (2025.02.18) running in shard mode
- Database nodes
  - 6x Supermicro Storage SuperServer SSG-222B-NE3X24R
  - 2x Intel® Xeon® 6767P @ 2.40GHz
  - 16x 128GB DDR5-6400 @ 6400MT/s [2 TB total]
  - NVIDIA® ConnectX-7 400Gb IPoIB
  - OS RHEL9.6 (Plow)
- Storage system (direct attached / local to each node)
  - 24x 12.8 TB Micron™ 9550 MAX NVMe™ SSD per server
  - PCIe® Gen5
  - 144 SSDs in total, 23x6 were used for data (1.84PB)
- Network
  - MQM9700-NS2F, NVIDIA® 64-Port NDR 400G InfiniBand Data Center Switch

* In addition to hosting kdb+ server processes, one of the Database Server Nodes hosted all of the STAC-M3 Clients.

Servers in the SUT were configured to mitigate the full range of Spectre/Meltdown threats checked by the Spectre/Meltdown checker tool. Details are available in the STAC Configuration Disclosure [2], along with the detailed Spectre/Meltdown tool output.

Intel submitted the following information and claims about its products:

Intel Xeon 6 Processors
Intel Xeon 6 processors with P-cores are optimized for high performance per core. Compared to the previous generation, these processors offer higher socket performance over the same thermal design power (TDP) range. They also provide increased memory bandwidth and higher input/output (I/O), making them ideal for a wide range of workloads. The innovative microarchitecture of Intel Xeon 6700-series processors delivers many advanced features and benefits including:

- *Up to 86 cores in a single socket, enabling ultra-high density compute performance and scalability.*
- *Support for up to eight sockets to support in-memory databases and mission-critical workloads.*
- *Intel AVX-512 encompasses unique instructions and two 512-bit fused-multiply add (FMA) units per core, boosting the speed of vector mathematics common to AI, HPC, and database workloads.*
- *Rich one-socket systems support storage systems by providing high I/O lanes and the ability to attach a greater number of drives for a lower TCO.*
- *Up to 88 lanes of PCIe® 5.0 for two-socket servers, with options of up to 136 lanes for one-socket server designs, to allow for significant I/O add-in components including accelerators, network adapters, storage controllers, and storage.*
- *Up to 336 MB L3 cache per processor and exceptionally low latency at large L3 access sizes.*
- *Intel AMX provides up to 16x more multiply accumulate (MAC) operations than Intel AVX-512 for BF16- and FP16-based models to enhance AI performance.*
- *Built-in security meets FIPS 140-3 Cryptographic Algorithm Validation Program (CAVP) certification.*

Micron submitted the following information and claims about its products:

Micron 9550 NVMe SSD
The Micron 9550 SSD delivers up to 14 GB/s sequential reads and 3.3M random read IOPS, enabling lightning-fast access to time-series data for trading and risk analytics. With up to 30.72 TB capacity per drive, high throughput, and ultra-low latency, the Micron 9550 SSD supports massive datasets while accelerating real-time analytics and decision-making. This combination ensures speed, responsiveness, and scalability for financial services workloads.

Micron 128GB DDR5-6400 RDIMM
The Micron 128GB DDR5-6400 RDIMM delivers 6400 MT/s speeds to accelerate compute-heavy workloads like portfolio optimization and fraud detection. It ensures faster decision-making and improved customer experience in capital market workloads. With 128GB capacity per module, Micron DDR5 RDIMMs allow financial institutions to run large time-series databases and simulations without compromise. High-capacity memory is essential for real-time analytics and overnight risk calculations.

Supermicro submitted the following information and claims about its products:

Supermicro Petascale SSG-222B-NE3X24R
Supermicro's Petascale line of all-flash storage servers combines the latest industry standards including E3.S form-factor SSDs, PCIe® Gen 5 I/O and DDR5 DRAM with Supermicro's balanced I/O design which distributes equal amounts of I/O bandwidth to SSDs and networking for optimal performance. The SSG-222B-NE3X24R is Supermicro's latest X14 generation of Intel-based storage servers featuring two Intel Xeon 6 6500 or 6700 series

processors. Other features include up to 32 front hot-swap E3.S NVMe drive bays, up to 32 DDR5-6400 DIMMs and up to five PCIe® add-in and AIOM cards.

This 2U system is designed for high-performance storage workloads, often used as a scale-out storage node for block, file or object storage. Supermicro has validated the Petascale line with many of the leading software-defined storage ISVs.

KX submitted the following information and claims about its products:

KX is the high-performance vector and time series analytics platform powering mission-critical AI and data infrastructure for the real-time economy. Built on kdb+, the world's fastest time series database, KX enables enterprises to unify streaming, historical, and unstructured data for low-latency analytics and decision-making at speed and at scale.

Trusted by the world's leading investment banks, hedge funds, and engineering-led firms across defense, manufacturing, and capital markets, KX delivers the performance, flexibility, and reliability required to support everything from algorithmic trading and operational automation to real-time simulation and vertical agentic AI. In the race to real-time, speed wins—and KX is how our customers stay ahead.

Project participants and responsibilities

The following firms participated in the project, with the associated responsibilities:

- KX implemented the STAC-M3 STAC Pack using the STAC-M3 Benchmark specifications.
- Intel, Micron, Supermicro and KX collaborated to develop this full-stack solution.
- STAC conducted the STAC-M3 Benchmark Audit, which included validating the database; inspecting any source-code revisions to the STAC Pack; validating the operation results; executing the tests and documenting the results.

Contacts

- Intel: Kevin Gildea, FSI Solutions Architect, kevin.gildea@intel.com; Evgueny Khartchenko, Performance Software Engineer, evgueny.khartchenko@intel.com; Kevin Bleckmann, Global Lab Management, kevin.bleckmann@intel.com.
- Micron: Ryan Meredith, Director of Data Center Workload Engineering, rmeredith@micron.com; Jay Walstrum, Principal Data Center Solutions Architect, jwalstrum@micron.com
- Supermicro: Wendell Wenjen, Director of Storage Marketing Development, wendellw@supermicro.com; Sufyan Shahzada, Senior Director of Solution Management, shahzadas@supermicro.com.
- KX: Holly Spiers, VP of Product Marketing, hspiers@kx.com.

Results status

- These benchmark specifications were developed by the STAC-M3 Working Group of the STAC Benchmark Community.
- These test results were audited by STAC or a STAC-certified third party, as indicated in the Responsibilities section above. As such, they are official results. For details, see www.STACresearch.com/reporting.

- The vendors attest that they did not modify the SUT during the Audit.

Overview of the STAC-M3 Benchmark specifications

Analyzing time-series data such as tick-by-tick quote and trade histories is crucial to many trading functions, from algorithm development to risk management. But the domination of liquid markets by automated trading—especially high-frequency trading—has made such analysis both more urgent and more challenging. As trading robots try to outwit each other on a microsecond or sub-microsecond scale, they dish out quotes and trades in ever more impressive volumes. This places a premium on technology that can store and analyze that activity efficiently. For example, the faster an algorithm developer can backtest and discard a haystack of unprofitable ideas, the faster he will find the needle of a winning algorithm, leaving more time to exploit it in the market.

The STAC Benchmark Community has developed the STAC-M3 Benchmarks in order to provide a common basis for quantifying the extent to which emerging software, cloud, and hardware innovations improve the performance of the storage, retrieval, and analysis of market data time series (“tick”) data.

STAC-M3 tests the ability of a complete solution stack of database software and infrastructure to perform a variety of operations on a large store of market data. The STAC-M3 Working Group designed these test specs to enable useful comparisons of entire solution stacks (i.e., to gauge the state of the art) as well as comparisons of specific stack layers while holding other layers constant. Comparisons can include (but are not limited to tick-database software products (typically columnar), storage architectures (including media, interconnects, and file systems), server products (including processors, chipsets, and memory), and cloud infrastructure (IaaS, DBaaS, etc.).

As shown below, the test setup for STAC-M3 consists of the “stack under test” (SUT) and client applications. No restrictions are placed on the architecture of the SUT or clients (though members of the STAC-M3 Working Group frequently provide input on architectures they would like to see tested). Threads within the clients take in Randomized Reference Data (RRD) such as dates and symbols, submit requests for the required operations, receive responses, and store the timings and results from these queries. Vendor-supplied code for the operations and response-time calculations are subjected to a combination of source-code inspection and empirical validation.

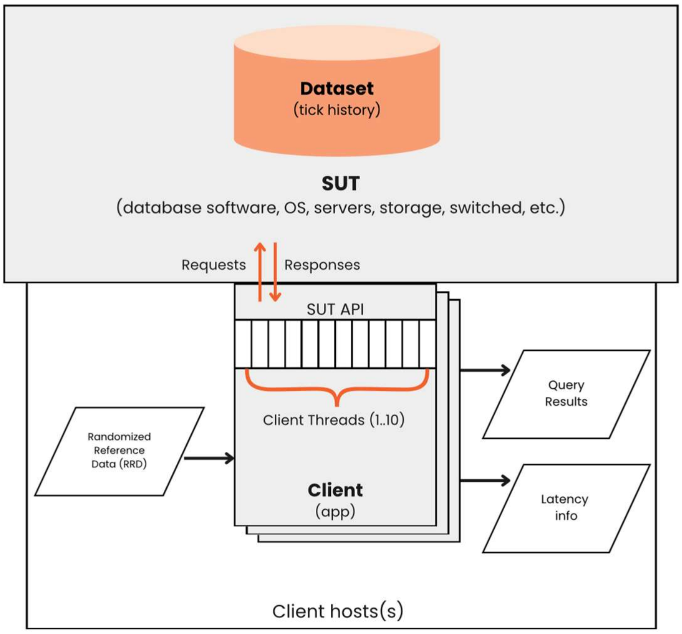

Understanding the STAC-M3 Benchmark Suites

The STAC-M3 Working Group has developed three benchmark suites that address different testing needs, as explained in the table below.

Suite	Purpose	Dataset size*	Concurrent requests	Operations	Constraints related to memory and storage
Antuco	Using a limited dataset size for convenience, simulate performance that would be obtained with a larger real-world dataset residing mostly on non-volatile media. Study a broad range of read and write operations.	4.5 TB	1 to 100	Range of compute-bound and storage-bound analytics. A few operations involving writes.	• No pre-loading into memory • File system cache cleared at several points in test run
Shasta	Study performance across a broad range of operations for datasets that are relatively small in the real world. (While the dataset tested is the same size as in Antuco, there is no attempt to simulate the storage-access pattern of a larger dataset.)	4.5 TB	1 to 100	Same as Antuco except operations involving writes are optional.	• Pre-loading into memory is allowed (most recent data first) • Caches not cleared during test run

Kanaga	Study performance on large datasets with large numbers of concurrent requests.	33 TB to 897 TB	1 to 450	A few storage-intensive queries.	• Pre-loading into memory is allowed (most recent data first) • Caches not cleared during test run • Storing certain data into faster storage tiers is allowed

* Reference size is based on a “standard” representation for each data type, making no allowance for optimizations or compression, nor for any overhead such as file headers, delimiters, indices, etc. Actual space requirements will vary by implementation and in practice tend to be smaller.

** Benchmark IDs that are identical except that one ends in ".TIME" and the other ends in ".LAT2" can be fairly compared. Prior to 2014, benchmarks in the STAC-M3 Antuco suite had two metrics: LAT1 (time to receive first result) and LAT2 (time to receive all results). Given that LAT1 and LAT2 results were identical for all systems reported from 2011 to 2013, LAT1 was eliminated in 2014. In addition, LAT2 was redesignated TIME in order to clarify that the measurement represents a response time at the application level and to avoid confusion with micro-level storage latency.

Datasets

STAC-M3 draws from client experience with equities and FX use cases. The database is synthetic, modeled on NYSE TAQ data (US equities). While testing with real data is also desirable, synthetic data has three advantages that make it compelling for STAC-M3:

- Synthetic data allows us to control the database properties exactly, which in turn allows us to randomize elements of queries from project to project while keeping the resulting workload exactly the same (for example, we control how much volume is associated with each symbol).
- Synthetic data does not incur fee liability from a third party such as an exchange.
- Synthesizing the data makes it easy to scale the database to an arbitrarily large size.

The dataset consists of high-volume symbols and low-volume symbols in proportions based on observed NYSE data. The data volume per symbol in the baseline dataset was based on doubling the typical volume in NYSE TAQ in 1Q10. The resulting database is considerably smaller than databases in use at customer sites, but the benchmarks impose policies that force the database to access storage. This approach was the STAC-M3 Working Group’s way of minimizing the cost of running baseline benchmarks while still yielding results indicative of those that would occur with large databases. Benchmarks that scale the database much larger are contained in the Kanaga suite of STAC-M3 Benchmark specifications. These are described in the Kanaga section below.

Metrics

The key metric in STAC-M3 is query response time. This measurement is performed in the client. A client thread gets a local timestamp (t_{submit}) just before submitting a query. When it receives the complete results of the query (sorted appropriately), the client immediately gets a second timestamp (t_{last}). Query response time is (t_{last}) - (t_{submit}).

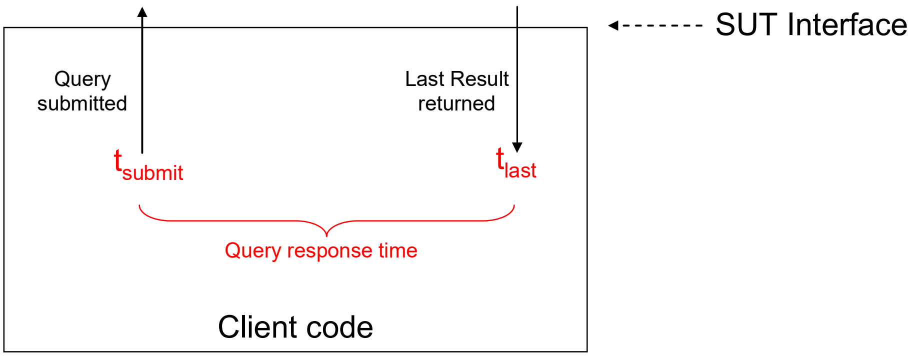

Timestamps and response time

Some of the I/O-focused benchmarks also measure the bytes read per second from persistent storage (i.e., excluding server cache), which is computed from the output of appropriate system utilities.

The algorithms in all benchmarks are defined so as to keep the result sets small. This ensures that network I/O between the test clients and server(s) is negligible compared to back-end processing times.

Test cases

The tests in the baseline STAC-M3 suite (Antuco) are listed in the "STAC-M3 Antuco Benchmark Operations" table below. These benchmarks operate on baskets of instruments, accessing many fields of underlying tick data for both trades and quotes across varying time windows. The table classifies each test case as relatively heavy on I/O, compute, or both.

The tests require a client application that is written to a product API and is capable of submitting requests from 10 independent threads, each of which simulates a user. As detailed in the table, some of the benchmarks call for one client instance making requests from a single thread, while others call for one client using 10 threads, and still others require 10 clients each using 10 threads (100 total requesting threads). One set of benchmarks (using the STATS-UI operation) tests multi-user scaling by running with 1, 10, 50, and 100 client threads. In all cases, benchmark results refer to per-request response times. For example, the mean of 10T.MKTSNAP.TIME is the mean time to satisfy a market-snapshot request from one of the threads, not the total time to satisfy requests from all 10 client threads. (Note, however, that a single request typically requires access to multiple instruments, fields, dates, and/or times.)

The range of dates eligible for querying depends on the benchmark. For example, some algorithms operate on dates randomly chosen throughout the year, some stick to a recent date range, and some always run on the most recent date (see the "Input Date Range" column of the table). The purpose of this differentiation is to provide a "recency bias" for those workloads where such bias is observed in the real world, while preventing such bias for those workloads that do not exhibit it in the real world.

Benchmark identifiers

The STAC-M3 Report Card and accompanying charts identify each benchmark unambiguously, as follows:

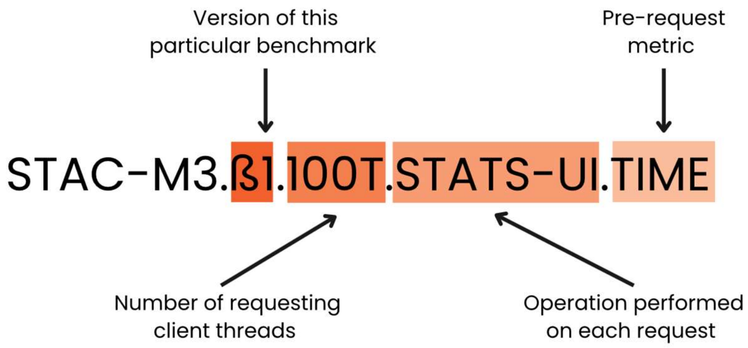

In charts, the ID is sometimes decomposed, with part of it in the chart title or labels. Each individual STAC Benchmark™ specification has its own version number. The same version of a given spec may appear in multiple benchmark suites. Thus, the code names of the suites are irrelevant when making comparisons. Versioning individual specs enables the reader to compare a discrete result from this "stack under test" (SUT) to the corresponding result from another SUT. When making comparisons, be sure that the identifiers match exactly. If they do not, the benchmark results may not be capable of fair comparison.

STAC-M3 Benchmarks in the Antuco Suite

The table below gives a brief overview of each test in this STAC-M3 suite. Version numbers of 1 or greater indicate benchmark specs that have been approved. Versions less than 1 are proposed by the STAC-M3 Working Group but not yet voted on by the full STAC Benchmark Community.

STAC-M3 Antuco Benchmark Operations

Root ID	Operation name	Ver	Number of requesting Client Threads	Algorithm performed on behalf of each requesting Client Thread	Algorithm I/O intensity	Algorithm compute intensity	Input date range*
VWAB-D	VWAB-Day	1	1	4-hour volume-weighted bid over one day for 1% of symbols (like VWAP but operating on quote data, so much higher input volume).	Heavy read	Light	Last 30 days
VWAB-12D-NO	VWAB-12DaysNoOverlap	1	100	4-hour volume-weighted bid over 12 days for 1% of symbols. No overlap in symbols among client threads.	Heavy read	Light	Full year
YRHIBID	Year High Bid	β1	1	Max bid over the year for 1% of symbols.	Heavy read	Light	Full year
YRHIBID-2	Year High Bid Re-run	β1	1	Re-run of YRHIBID (same symbols) without clearing the cache.	Heavy read†	Light	Full year
QTRHIBID	Quarter HighBid	β1	1	Max bid over the quarter for 1% of symbols.	Heavy read	Light	Most recent quarter
MOHIBID	Month High Bid	β1	1	Max bid over the month for 1% of symbols.	Heavy read	Light	Most recent month
WKHIBID	Week High Bid	β1	1	Max bid over the week for 1% of symbols.	Heavy read	Light	Most recent week

STATS-AGG	Aggregate Stats	ß1	10	One set of basic statistics over 100 minutes for all symbols on one exchange. Each 100-minute range crosses a date boundary.	Heavy read	Heavy	Full year
STATS-UI	Stats - Unpredictable Intervals	ß1	1, 10, 50, 100 (more optional)	Per-minute[‡] basic statistics over 100 minutes for all high-volume symbols on one exchange. Each 100-minute range crosses a date boundary.	Heavy read	Heavy	Full year
MKTSNAP	Market Snapshot	ß1	10	Most recent trade and quote information for 1% of symbols as of a random time.	Heavy read	Heavy	Full year
VOLCURV	Volume Curves	ß1	10	Create an average volume curve (using minute intervals aligned on minute boundaries) for 10% of symbols over 20 days selected at random.	Light read	Heavy	Full year
THEOPL	Theoretical P&L	ß1	10	For a basket of 100 trades on random dates, find the future times at which 2X, 4X, and 20X the trade size traded in each symbol. Trade sizes cause up to 5 days of forward searching. Calculate the corresponding VWAP and total volume traded over those periods.	Light read	Heavy	Full year
NBBO	NBBO	ß1	1	Create the NBBO across all 10 exchanges for all symbols on the most recent day. Write to persistent storage.	Heavy read and write	Heavy	Most recent day
WRITE	Write	1	1	Write one day's quote data to persistent storage, following the same algorithm used to generate the randomized dataset used in the other Operations.	Heavy write	Light	n/a

STORAGE.EFF	Storage efficiency	1.1	n/a	Reference Size of the Dataset divided by size of the Dataset in the SUT format used for the performance benchmarks. Expressed as as percentage.	n/a	n/a	n/a

* In some cases, one or more dates at the end of the year were excluded from eligibility to prevent an algorithm that crosses days from running out of input data.

† Typically this will be reads from DRAM cache.

‡ In this case, interval start times are offset from minute boundaries by a consistent random amount per test run, so that the SUT cannot rely on pre-calculated minute statistics.

STAC-M3 Benchmarks in the Kanaga Suite

The optional "Kanaga" suite of STAC-M3 consists of three benchmarks involving larger quantities of data than the Antuco suite in order to measure the volume-scalability of a database stack. The ability of a tick analytics stack to handle increasing volumes of historical data is important to today's trading organizations. Market data volumes continue to grow quickly, sometimes in step-function increments. Engineering a solution that delivers consistent, high performance across ever larger datasets and a large number of users can be a challenge. STAC-M3 Kanaga provides some insight into how well a given solution scales.

The STAC-M3 Kanaga dataset is an extension of the STAC-M3 Antuco dataset described above (a hypothetical year's worth of data in 2011). The Kanaga suite calls for additional years to be added to the Antuco database, using the same data structures, where each year's size is a multiple of the previous year's. Using an overly aggressive rule of thumb that assumes total market data volumes doubled roughly every 18 months, Kanaga sets the annual scaling factor at 1.6. In addition to scaling forward in time, the specs also call for a 2003 dataset (1.6^8 times smaller than the 2011 dataset) in order to test how the technology stack handles small datasets. The full STAC-M3 Kanaga dataset in a given database format is typically around 50TB when scaling through 2015.

These sizes do not represent actual NYSE TAQ volumes in the given years. Using years as the basis for increasing the volume of quotes and trades in the benchmarks is simply a convenience. Scaling up a single level-1 dataset is not necessarily the most realistic way to emulate the large tick data stores in trading institutions today. Real deployments typically also involve level 1 and level 2 data from many asset classes. However, the STAC-M3 Working Group determined that scaling the existing dataset would provide significant insight without the expense and complexity of specifying additional trade and quote record formats and designing entirely new queries to operate on them.

The STAC-M3 Kanaga operations are summarized in the table below. [n]YRHIBID is a good test of sequential read performance, while YR[n]-MKTSNAP is a good test of random read performance. As with YRHIBID, [n]YRHIBID has an additional metric: the bytes read per second from persistent storage (i.e., excluding server cache), which is computed from the output of appropriate system utilities. YR[n]VWAB-12D-HO is a good test of user scaling and how that varies with the size of the dataset.

In addition to the standard response-time metric, Kanaga also includes a volume-adjusted response time metric (VTIME), which divides the response time (TIME) by a weighted-volume factor (WVF). The (WVF) is a relative measure of the number of quotes and/or trades in the dataset subject to the given query compared to the number of quotes and trades subject to the same query in the base year (2011). For example, 2YRHIBID covers 2.6 times the quotes and trades of YRHIBID (2012 is 1.6 times the size of 2011, and 2YRHIBID covers both 2011 and 2012). Volume-adjusted response times normalize response times in order to see how the response time per quote or trade changes with the size of the dataset.

Note that some of the tables and charts in this report also display results of the corresponding tests from the Antuco suite (STAC-M3.ß1.1T.YRHIBID, STAC-M3.ß1.10T.MKTSNAP) on this SUT as the baseline for scale comparisons.

STAC-M3 Kanaga Benchmark Operations

Root ID	Operation name	Ver	Number of requesting Client Threads	Algorithm performed on behalf of each requesting Client Thread	I/O intensity	Algorithm compute intensity	Input date range
[n]YRHIBID	Multi-year high bid	ß1	1	Return the highest bid price for each of a certain 1% of symbols over a particular range of years in the dataset. The range for 2YRHIBID is from the first day of 2011 through the last day of 2012. The range for 3YRHIBID is from the first day of 2011 through the last day of 2013, and so on.	Heavy read	Light	Varies
YR[n]-MKTSNAP	Year-*n* market snapshot	ß1	10	Returns the price and size for the latest quote and trade for each of a certain 1% of symbols at a unique time on a unique date in the given year of the dataset. YR2-MKTSNAP queries dates and times in 2012, while YR3-MKTSNAP queries dates and times in 2013, and so on.	Heavy read	Heavy	Varies
YR[n]VWAB-12D-HO		ß1	Varies	Similar to the VWAB-12D-NO operation in the Antuco suite (4-hour volume-weighted average bid for 12 randomly-selected days) except for three things: 1) it varies the number of concurrent requests (client threads, the “c” in the benchmark ID); 2) it operates in multiple years of the Kanaga dataset rather than just within the Antuco year (2011); and 3) the dates and symbols are chosen so as to ensure heavy overlap among requests, since this is a common pattern in the real world. The tester chooses three scale points in terms of client threads. These scale points must top out at the maximum points to be tested for the SUT.	Heavy read	Light	Varies

Specification particulars

Version

This project followed the benchmark specifications in [1]. Qualified members of the STAC Benchmark Community can access these specifications and download the programs used in this project in order to run the same tests on systems in the privacy of their own labs.

Limitations

- As discussed in the overview of the benchmark specifications, the STAC-M3 Antuco suite was designed to test operations on a limited amount of purely historical data. The STAC-M3 Kanaga suite involves larger amounts of historical data.
- As discussed in Section 1, the dataset used in this version of STAC-M3 is synthetic. The algorithm to generate the dataset creates random values for prices and sizes that can vary widely from tick to tick. In the real world, by contrast, there is significant correlation of successive prices (i.e., large differences from tick to tick are relatively rare). Compression algorithms often take advantage of this fact, such as by focusing on deltas between successive values. Hence, the storage efficiency of a SUT may be higher when working with real data than with the synthetic dataset of this version of STAC-M3.
- MKTSNAP is a random-access operation, and because systems tend to have areas of storage that differ in access time, the MKTSNAP response times can vary considerably. STAC-M3 requires a limited number of test runs, which means that the standard deviation of response times can be quite large relative to the mean. Thus, the mean MKTSNAP.TIME in a given year of data is, in general, not the best statistic to use from these tests (i.e., two systems with the same performance could get quite different mean response times simply by chance). Median and max are probably more instructive indicators, which is why these are used in Figure 11.

Appendix A: STAC-M3 Antuco results

Below are the results from benchmarks in the Antuco suite of STAC-M3, in tabular and graphical forms.

Storage efficiency

Storage Efficiency	
The reference size of the dataset divided by the size of the dataset as stored by the SUT. The less storage space required, the higher the percentage.	
STAC-M3.v1.1.STORAGE.EFF	**149%**

Light-Compute Benchmarks

High Bid (1 Client Thread Requesting)

Return the high bid for a certain 1% of symbols over varying timeframes. Run the year-high bid a second time (YRHIBID-2) without clearing the cache.

	Response time (milliseconds)				
Spec ID	MEAN	MED	MIN	MAX	STDV
STAC-M3.β1.1T.YRHIBID.TIME	59	60	58	60	1
STAC-M3.β1.1T.YRHIBID-2.TIME	38	38	37	39	1
STAC-M3.β1.1T.QTRHIBID.TIME	56	56	52	58	2
STAC-M3.β1.1T.MOHIBID.TIME	56	56	54	57	1
STAC-M3.β1.1T.WKHIBID.TIME	51	52	44	55	4

	Megabytes read per second*				
Spec ID	MEAN			MAX	
STAC-M3.β1.1T.YRHIBID.MBPS	79,228			80,747	
STAC-M3.β1.1T.YRHIBID-2.MBPS	0			0	
STAC-M3.β1.1T.QTRHIBID.MBPS	19,597			21,005	
STAC-M3.β1.1T.MOHIBID.MBPS	6,499			6,771	
STAC-M3.β1.1T.WKHIBID.MBPS	1,804			2,086	

* Megabytes read per second from persistent media, according to iostat. That is, cache hits do not count as bytes read.

Write Test

Perform the Basic Data Generation Algorithm for 1 day's data.

	Response time* (milliseconds)				
Spec ID	MEAN	MED	MIN	MAX	STDV
STAC-M3.v1.1T.WRITE.TIME	1,440	1,503	1,246	1,606	163

* Time to write all results.

Post-Trade Analytics Benchmarks

VWAB on 1 Day's Data
(1 Client Thread Requesting)

Return ~4-hour volume-weighted bid over a single day for certain 1% of symbols

	Response time (milliseconds)				
Spec ID	MEAN	MED	MIN	MAX	STDV
STAC-M3.v1.1T.VWAB-D.TIME	16	16	15	17	1

Theoretical P&L
(10 Client Threads Requesting)

For each of 10 Client Threads querying a unique set of 100 trades, find the amount of time until 2x, 4x, and 20x the size of each trade was traded in the market, and return the VWAP and total volume over those times intervals.

	Response time (milliseconds)				
Spec ID	MEAN	MED	MIN	MAX	STDV
STAC-M3.β1.10T.THEOPL.TIME	390	390	374	402	7

Market Snapshot
(10 Client Threads Requesting)

To each of 10 Client Threads querying a unique date, time, and set of symbols (1% of the total symbols), return the price and size information for the latest quote and trade for each symbol.

	Response time (milliseconds)				
Spec ID	MEAN	MED	MIN	MAX	STDV
STAC-M3.β1.10T.MKTSNAP.TIME	149	114	48	312	72

Research Analytics Benchmarks

Volume Curves
(10 Client Threads Requesting)

To each of 10 Client Threads querying a unique set of 20 dates and set of symbols (10% of the total symbols), return the average proportion of volume traded in each minute interval for each symbol across the date set.

	Response time (milliseconds)				
Spec ID	MEAN	MED	MIN	MAX	STDV
STAC-M3.β1.10T.VOLCURV.TIME	1,374	1,338	1,125	1,852	188

Aggregated Stats
(10 Client Threads Requesting)

For each of 10 Client Threads querying a unique exchange, date, and start time, return basic statistics calculated for the entirety of the 100-minute time range following the start time. Time ranges always cross a date boundary.

	Response time (milliseconds)				
Spec ID	MEAN	MED	MIN	MAX	STDV
STAC-M3.β1.10T.STATS-AGG.TIME	1,401	1,216	829	3,996	778

Stats Over Unpredictable Intervals
(Variable Client Threads Requesting)

To each of some number of Client Threads querying a unique exchange, date, and start time, return basic statistics calculated for each minute interval in a 100-minute time range following the start time. Start times are offset from minute boundaries by a random amount. Time ranges always cross a date boundary. Tests must be run with 1, 10, 50, and 100 Client Threads. Tests with other numbers of Client Threads are optional.

	Response time (milliseconds)				
Spec ID	MEAN	MED	MIN	MAX	STDV
STAC-M3.β1.1T.STATS-UI.TIME	246	251	234	254	8
STAC-M3.β1.10T.STATS-UI.TIME	278	258	120	489	78
STAC-M3.β1.50T.STATS-UI.TIME	630	621	101	1,589	359
STAC-M3.β1.100T.STATS-UI.TIME	872	747	98	2,164	564

NBBO Benchmark

NBBO

Calculate NBBO across all exchanges for all symbols on one day.

	Response time* (milliseconds)				
Spec ID	MEAN	MED	MIN	MAX	STDV
STAC-M3.β1.1T.NBBO.TIME	23,158	23,188	22,781	23,763	362

* Time to write all results.

Multi-day/Multi-User VWAB Benchmark

VWAB for 12 Days with No Overlap in Interest (100 Client Threads Requesting)

To each of 100 Client Threads querying unique symbol sets, return 4-hour volume-weighted bid for 12 random days per thread for 1% of symbols per thread

	Response time (milliseconds)				
Spec ID	MEAN	MED	MIN	MAX	STDV
STAC-M3.v1.100T.VWAB-12D-NO.TIME	191	193	36	368	91

Chart view

The charts that follow illustrate or elaborate on the results above:

- Figure 1 through Figure 4 plot the mean response time (TIME) benchmarks for all of the operations.
- Figure 5 and Figure 6 analyze the individual response-time observations for the multi-user/multi-day VWAB benchmark (STAC-M3.v1.100T.VWAB-12D-NO.TIME), first by sorting the results by response time, then by plotting them in a histogram.
- Figure 7 provides a more explicit look at multi-user scaling by plotting the response time for the intervalized statistics benchmark (STAC-M3.s1.[n]T.STATS-UI.TIME) against the number of simultaneously requesting client threads (n).
- Figure 8 and Figure 9 take the 100-client-thread case of Figure 7 and analyze the individual responsetime observations, first by sorting the results by response time, then by plotting them in a histogram.

Refer to the overview of the benchmark specifications below or the tables above for explanations of the benchmark IDs used in the charts.

The axes in the bar charts are fixed, so that results from this SUT may be visually compared to those of other SUTs. Because the results of future SUTs are unpredictable, the axes use a log scale.

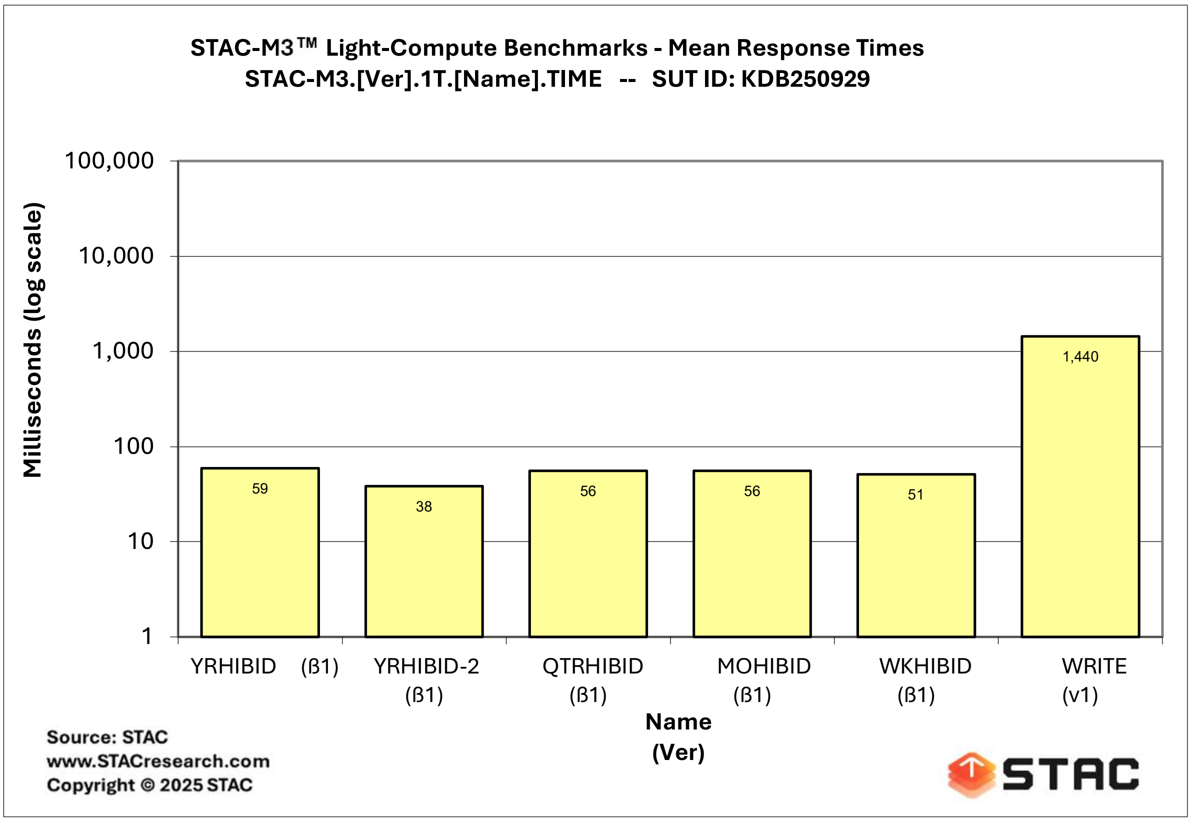
STAC-M3™ Light-Compute Benchmarks - Mean Response Times
STAC-M3.[Ver].1T.[Name].TIME -- SUT ID: KDB250929
Milliseconds (log scale)
100,000
10,000
1,000
100
10
1
59
38
56
56
51
1,440
YRHIBID (ß1)
YRHIBID-2 (ß1)
QTRHIBID (ß1)
MOHIBID (ß1)
WKHIBID (ß1)
WRITE (v1)
Name (Ver)
Source: STAC
www.STACresearch.com
Copyright © 2025 STAC

Figure 1

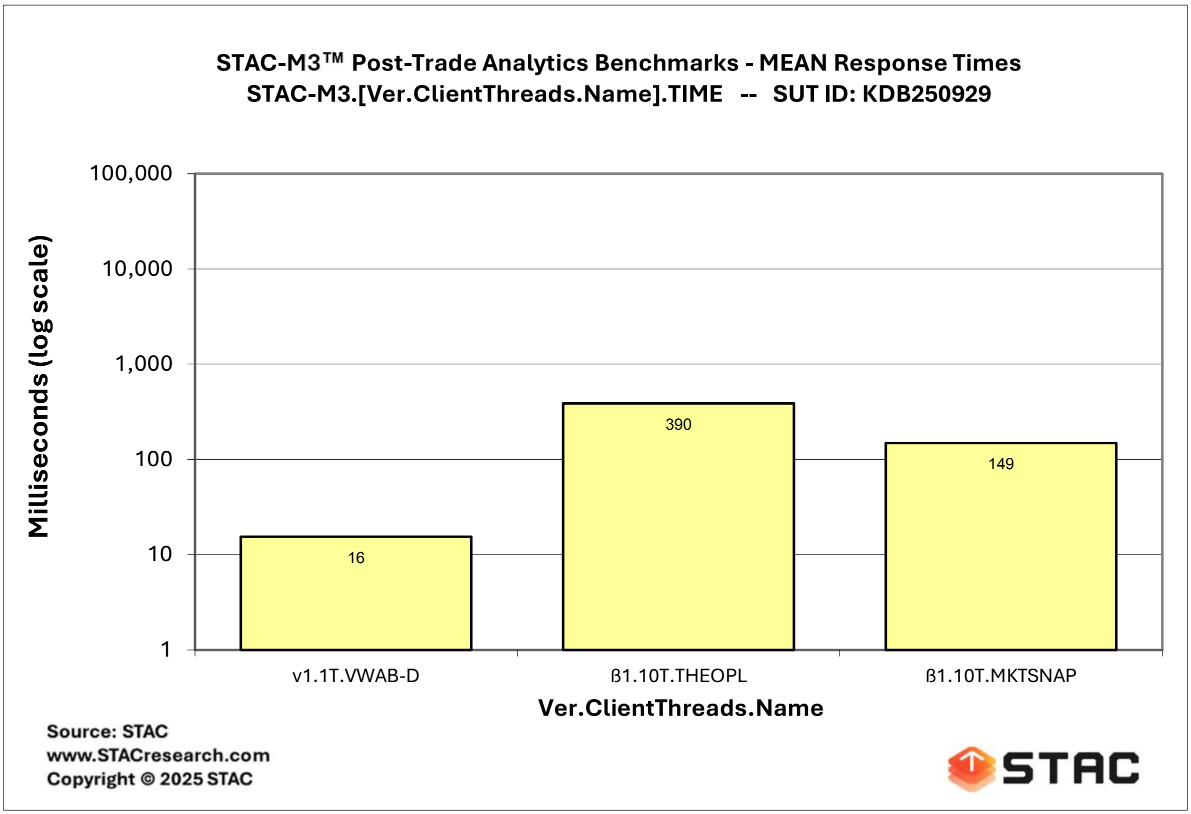
STAC-M3™ Post-Trade Analytics Benchmarks - MEAN Response Times
STAC-M3.[Ver.ClientThreads.Name].TIME -- SUT ID: KDB250929
Milliseconds (log scale)
100,000
10,000
1,000
100
10
1
16
390
149
v1.1T.VWAB-D
ß1.10T.THEOPL
ß1.10T.MKTSNAP
Ver.ClientThreads.Name
Source: STAC
www.STACresearch.com
Copyright © 2025 STAC
STAC

Figure 2

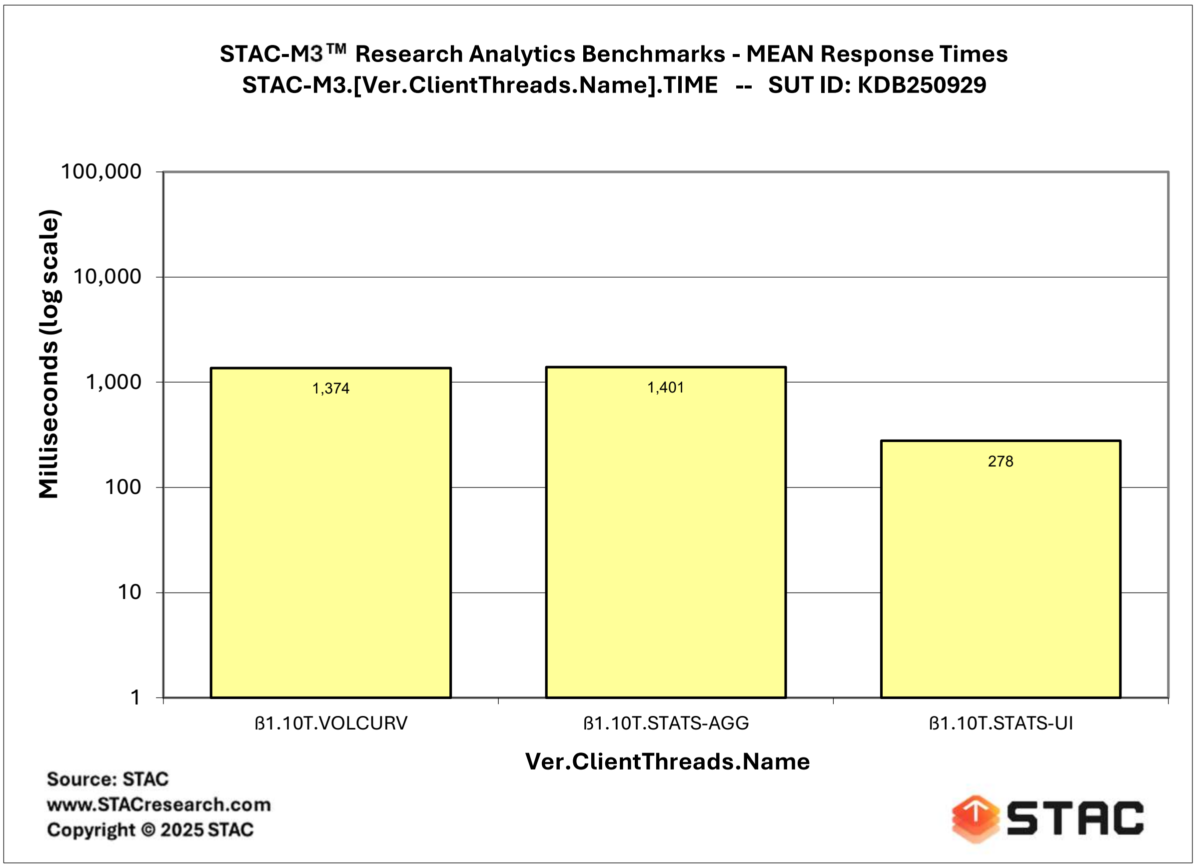
STAC-M3™ Research Analytics Benchmarks - MEAN Response Times
STAC-M3.[Ver.ClientThreads.Name].TIME -- SUT ID: KDB250929
Milliseconds (log scale)
100,000
10,000
1,000
100
10
1
1,374
1,401
278
β1.10T.VOLCURV
β1.10T.STATS-AGG
β1.10T.STATS-UI
Ver.ClientThreads.Name
Source: STAC
www.STACresearch.com
Copyright © 2025 STAC
STAC

Figure 3

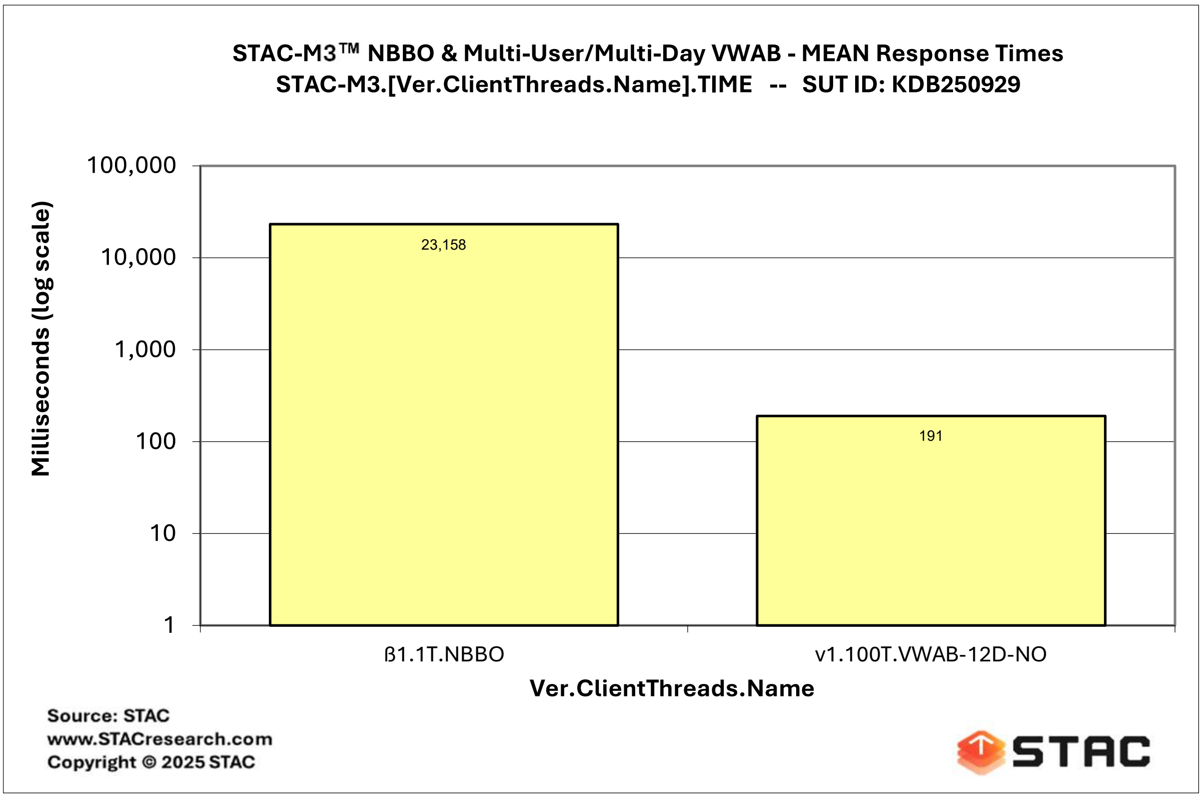
STAC-M3™ NBBO & Multi-User/Multi-Day VWAB - MEAN Response Times
STAC-M3.[Ver.ClientThreads.Name].TIME -- SUT ID: KDB250929
Milliseconds (log scale)
100,000
10,000
1,000
100
10
1
23,158
191
ß1.1T.NBBO
v1.100T.VWAB-12D-NO
Ver.ClientThreads.Name
Source: STAC
www.STACresearch.com
Copyright © 2025 STAC
STAC

Figure 4

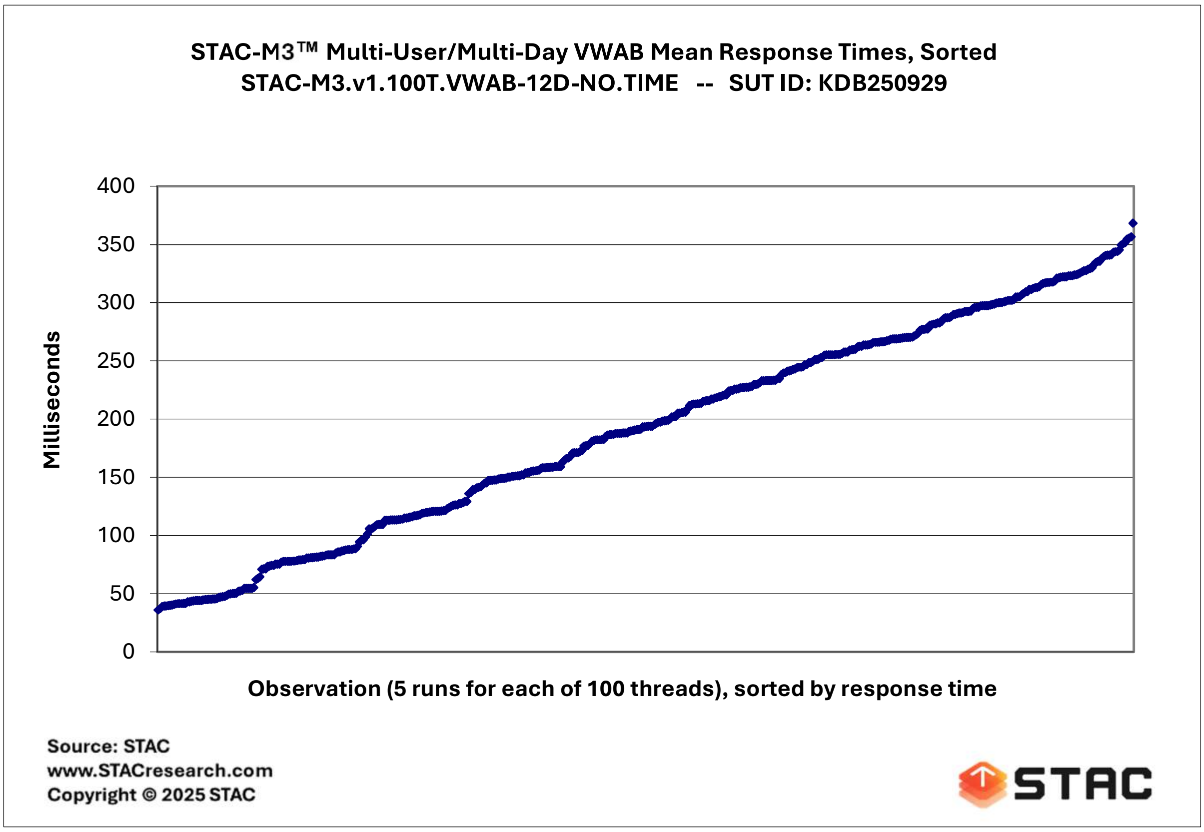
STAC-M3™ Multi-User/Multi-Day VWAB Mean Response Times, Sorted
STAC-M3.v1.100T.VWAB-12D-NO.TIME -- SUT ID: KDB250929
Milliseconds
400
350
300
250
200
150
100
50
0
Observation (5 runs for each of 100 threads), sorted by response time
Source: STAC
www.STACresearch.com
Copyright © 2025 STAC
STAC

Figure 5

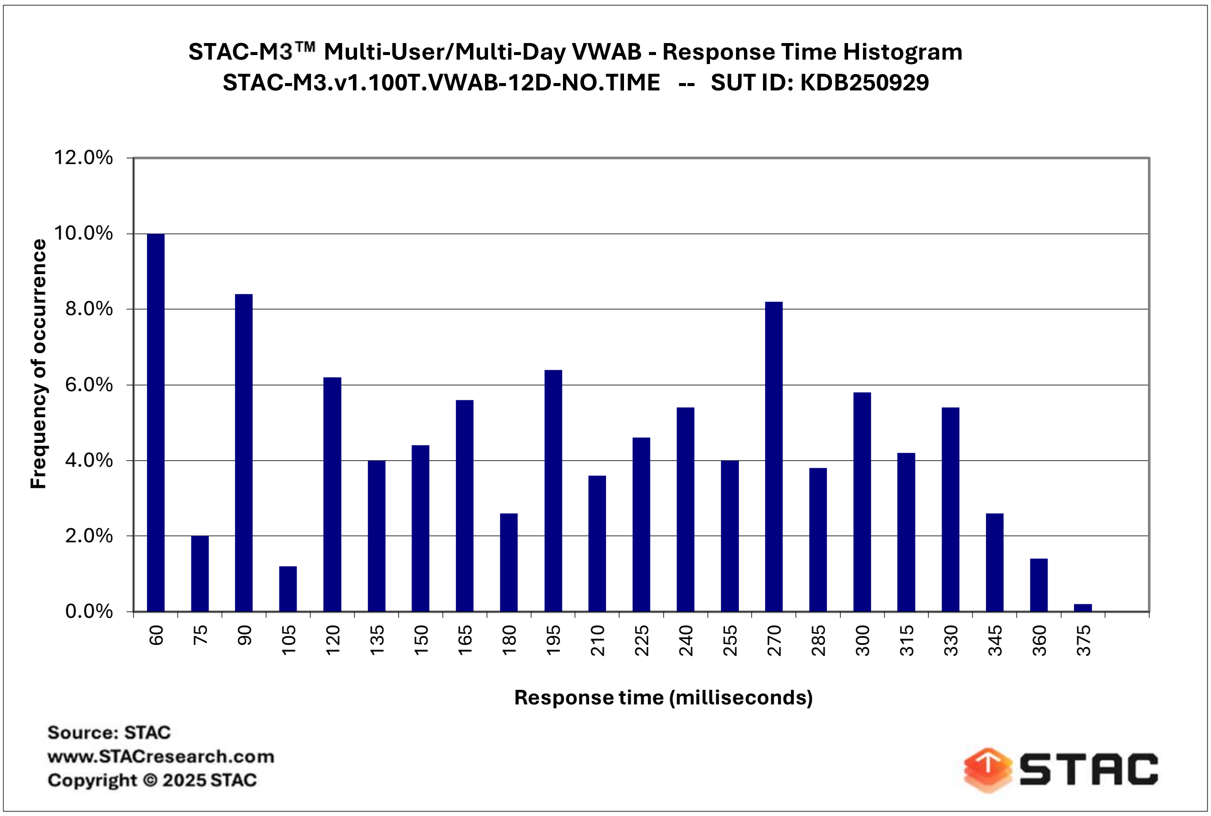
STAC-M3™ Multi-User/Multi-Day VWAB - Response Time Histogram
STAC-M3.v1.100T.VWAB-12D-NO.TIME -- SUT ID: KDB250929
Frequency of occurrence
12.0%
10.0%
8.0%
6.0%
4.0%
2.0%
0.0%
60
75
90
105
120
135
150
165
180
195
210
225
240
255
270
285
300
315
330
345
360
375
Response time (milliseconds)
Source: STAC
www.STACresearch.com
Copyright © 2025 STAC
STAC

Figure 6

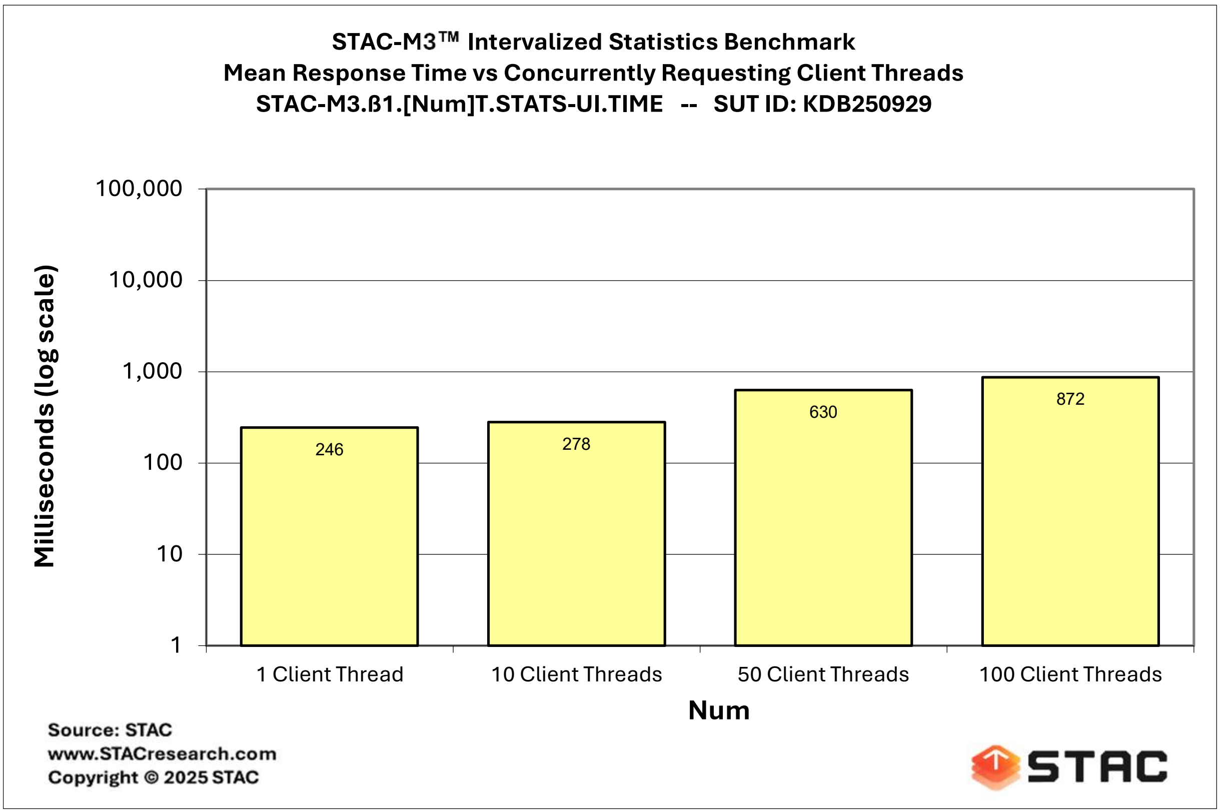
STAC-M3™ Intervalized Statistics Benchmark
Mean Response Time vs Concurrently Requesting Client Threads
STAC-M3.β1.[Num]T.STATS-UI.TIME -- SUT ID: KDB250929
Milliseconds (log scale)
100,000
10,000
1,000
100
10
1
246
278
630
872
1 Client Thread
10 Client Threads
50 Client Threads
100 Client Threads
Num
Source: STAC
www.STACresearch.com
Copyright © 2025 STAC
STAC

Figure 7

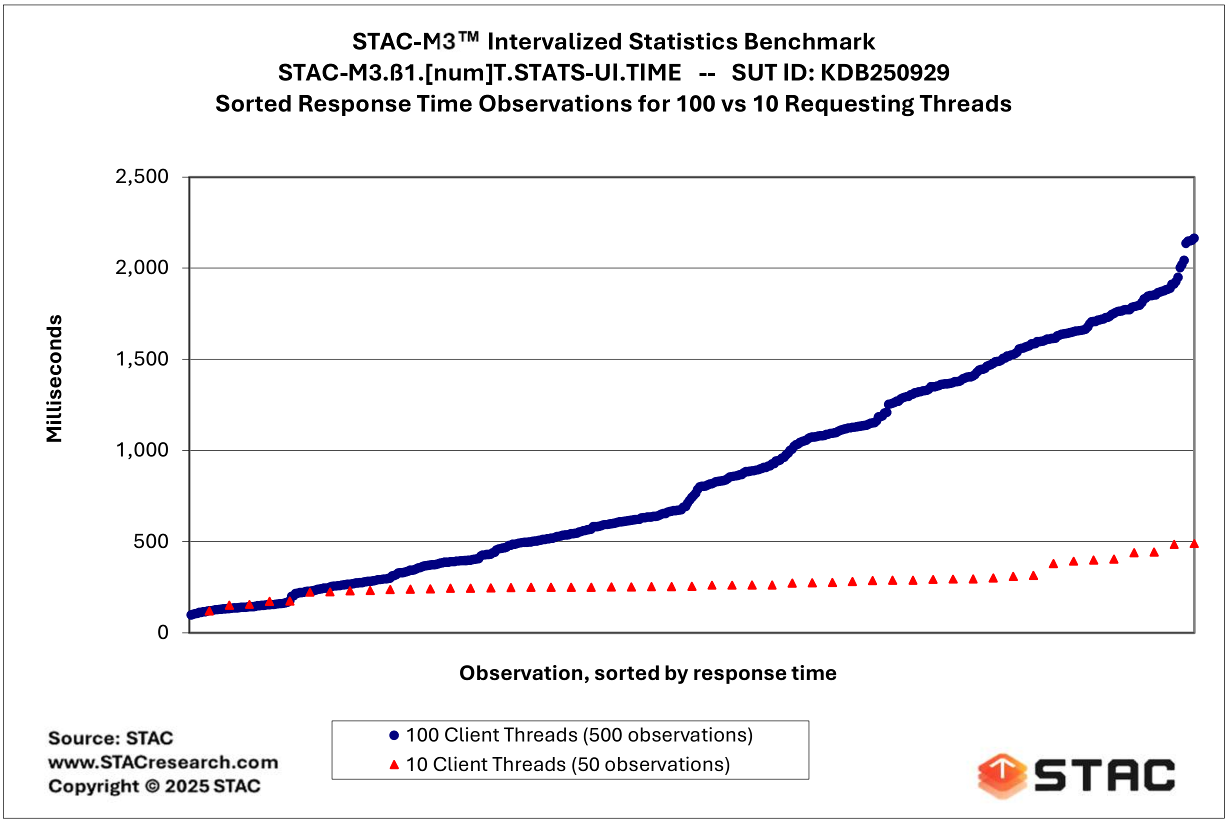
STAC-M3™ Intervalized Statistics Benchmark
STAC-M3.ß1.[num]T.STATS-UI.TIME -- SUT ID: KDB250929
Sorted Response Time Observations for 100 vs 10 Requesting Threads
2,500
2,000
1,500
1,000
500
0
Milliseconds
Observation, sorted by response time
Source: STAC
www.STACresearch.com
Copyright © 2025 STAC
100 Client Threads (500 observations)
10 Client Threads (50 observations)
STAC

Figure 8

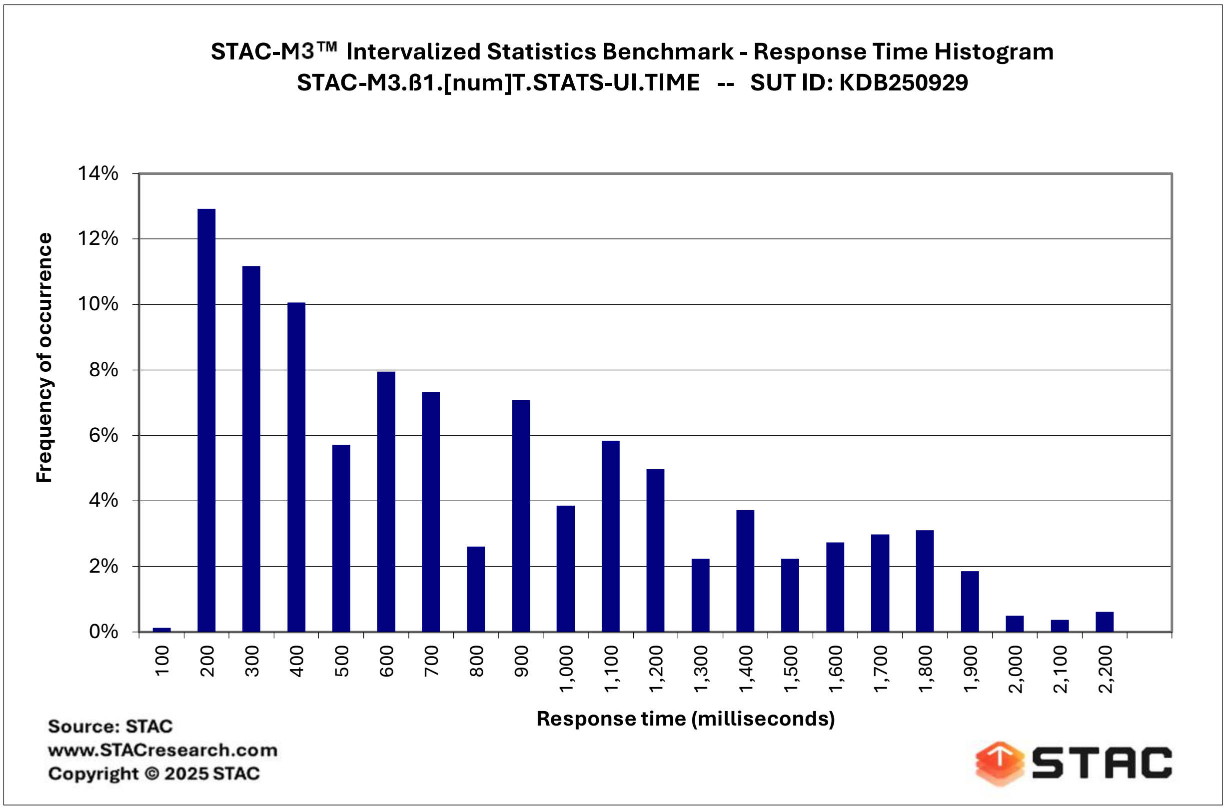
STAC-M3™ Intervalized Statistics Benchmark - Response Time Histogram
STAC-M3.ß1.[num]T.STATS-UI.TIME -- SUT ID: KDB250929
Frequency of occurrence
14%
12%
10%
8%
6%
4%
2%
0%
100
200
300
400
500
600
700
800
900
1,000
1,100
1,200
1,300
1,400
1,500
1,600
1,700
1,800
1,900
2,000
2,100
2,200
Response time (milliseconds)
Source: STAC
www.STACresearch.com
Copyright © 2025 STAC
STAC

Figure 9

Appendix K: STAC-M3 Kanaga results

Tabulated results

High Bid Over Varying Intervals (1 Client Thread Requesting)

Return the high bid for a certain 1% of symbols over a particular range of years in the dataset.

	Raw last-result latency (TIME) - Milliseconds		Volume-adjusted last-result latency (VTIME) - Milliseconds		Megabytes read per second (BPS)*	
Spec ID	MEAN	MAX	MEAN	MAX	MEAN	MAX
STAC-M3.1T.OLDYRHIBID.TIME	43	44	1,842	1,910	6,122	6,346
STAC-M3.ß1.1T.YRHIBID	59	60	23	23	79,228	80,747
STAC-M3.ß1.1T.2YRHIBID	78	80	18	19	150,459	153,575
STAC-M3.ß1.1T.3YRHIBID	103	106	15	16	218,832	225,191
STAC-M3.ß1.1T.4YRHIBID	138	140	13	13	297,934	307,394
STAC-M3.ß1.1T.5YRHIBID	203	223	12	13	360,687	373,495

* Megabytes read per second from persistent media, according to iostat. That is, cache hits do not count as bytes read.

Market Snapshot Within Varying Years (10 Client Threads Requesting)

To each of 10 Client Threads querying a unique time, and set of symbols (1% of the total symbols) on a unique date in the given year of the dataset, return the price and size information for the latest quote and trade for each symbol.

	Raw last-result latency (TIME) - Milliseconds				
Spec ID	MEAN	MED	MIN	MAX	STDV
STAC-M3.ß1.10T.MKTSNAP	149	114	48	312	72
STAC-M3.ß1.10T.YR2-MKTSNAP	143	116	81	351	71
STAC-M3.ß1.10T.YR3-MKTSNAP	149	127	86	319	66
STAC-M3.ß1.10T.YR4-MKTSNAP	159	133	89	331	68
STAC-M3.ß1.10T.YR5-MKTSNAP	173	148	99	337	69

	Volume-adjusted last-result latency (VTIME) - Milliseconds				
Spec ID	MEAN	MED	MIN	MAX	STDV
STAC-M3.ß1.10T.MKTSNAP	149	114	48	312	72
STAC-M3.ß1.10T.YR2-MKTSNAP	89	72	51	219	44
STAC-M3.ß1.10T.YR3-MKTSNAP	58	49	33	125	26
STAC-M3.ß1.10T.YR4-MKTSNAP	39	33	22	81	17
STAC-M3.ß1.10T.YR5-MKTSNAP	26	23	15	51	11

Multi-Day VWAB with Varying Concurrent Requests and within Varying Years

To each of n Client Threads querying a set of symbols on 12 random days in the given year of the dataset, return the 4-hour volume-weighted bid for each date and symbol. Date/symbol combinations are designed with heavy overlap among threads.

		Raw last-result latency (TIME) - Milliseconds				
Spec ID	**Concurrent Requests**	MEAN	MED	MIN	MAX	STDV
STAC-M3.β1.1T.YR1VWAB-12D-HO	1	23	20	18	31	5
STAC-M3.β1.50T.YR1VWAB-12D-HO	50	224	215	32	424	120
STAC-M3.β1.100T.YR1VWAB-12D-HO	100	222	197	34	475	126
STAC-M3.β1.1T.YR2VWAB-12D-HO	1	28	25	23	40	6
STAC-M3.β1.50T.YR2VWAB-12D-HO	50	325	333	24	667	183
STAC-M3.β1.100T.YR2VWAB-12D-HO	100	282	227	37	638	185
STAC-M3.β1.1T.YR3VWAB-12D-HO	1	41	34	32	53	10
STAC-M3.β1.50T.YR3VWAB-12D-HO	50	487	480	28	1,072	283
STAC-M3.β1.100T.YR3VWAB-12D-HO	100	404	313	41	951	274
STAC-M3.β1.1T.YR4VWAB-12D-HO	1	59	50	47	77	13
STAC-M3.β1.50T.YR4VWAB-12D-HO	50	787	784	38	1,677	457
STAC-M3.β1.100T.YR4VWAB-12D-HO	100	602	462	56	1,418	414
STAC-M3.β1.1T.YR5VWAB-12D-HO	1	90	86	71	113	17
STAC-M3.β1.50T.YR5VWAB-12D-HO	50	1,390	1,409	103	3,150	804
STAC-M3.β1.100T.YR5VWAB-12D-HO	100	1,028	743	96	2,752	748

		Volume-adjusted last-result latency (VTIME) - Milliseconds				
Spec ID	**Concurrent Requests**	MEAN	MED	MIN	MAX	STDV
STAC-M3.β1.1T.YR1VWAB-12D-HO	1	23	20	18	31	5
STAC-M3.β1.50T.YR1VWAB-12D-HO	50	224	215	32	424	120
STAC-M3.β1.100T.YR1VWAB-12D-HO	100	222	197	34	475	126
STAC-M3.β1.1T.YR2VWAB-12D-HO	1	11	10	9	15	2
STAC-M3.β1.50T.YR2VWAB-12D-HO	50	127	130	9	261	71
STAC-M3.β1.100T.YR2VWAB-12D-HO	100	110	89	15	249	72
STAC-M3.β1.1T.YR3VWAB-12D-HO	1	16	13	13	21	4
STAC-M3.β1.50T.YR3VWAB-12D-HO	50	190	187	11	419	110
STAC-M3.β1.100T.YR3VWAB-12D-HO	100	158	122	16	372	107
STAC-M3.β1.1T.YR4VWAB-12D-HO	1	14	12	11	19	3
STAC-M3.β1.50T.YR4VWAB-12D-HO	50	192	191	9	409	111
STAC-M3.β1.100T.YR4VWAB-12D-HO	100	147	113	14	346	101
STAC-M3.β1.1T.YR5VWAB-12D-HO	1	22	21	17	28	4
STAC-M3.β1.50T.YR5VWAB-12D-HO	50	339	344	25	769	196
STAC-M3.β1.100T.YR5VWAB-12D-HO	100	251	181	23	672	183

Chart view

The charts that follow illustrate or elaborate on the results above:

- Figure 1 plots the mean response time (TIME) and volume-adjusted response time (VTIME) for all of the HIBID operations.
- Figure 2 plots the median and maximum response times (TIME) for all of the MKTSNAP operations. To understand why median and max were chosen for this chart instead of mean, see the Limitations section.
- Figure 3 analyzes the individual response-time observations for MKTSNAP benchmarks.

Refer to Section 1 (Overview) and the tables above for explanations of the benchmark IDs used in the charts.

The axes in the bar charts are fixed, so that results from this SUT may be visually compared to those of other SUTs. Because the results of future SUTs are unpredictable, the axes use a log scale. The axes in Figure 3 are not fixed.

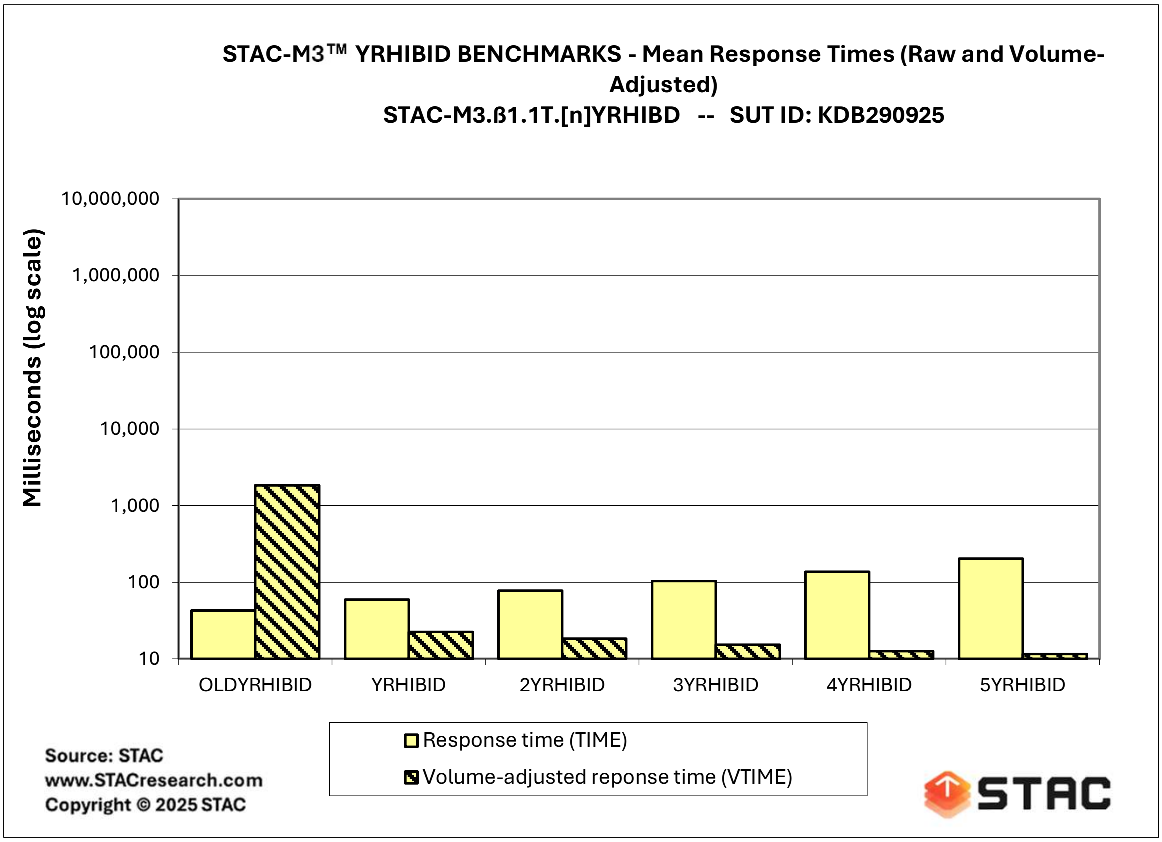

Figure 1

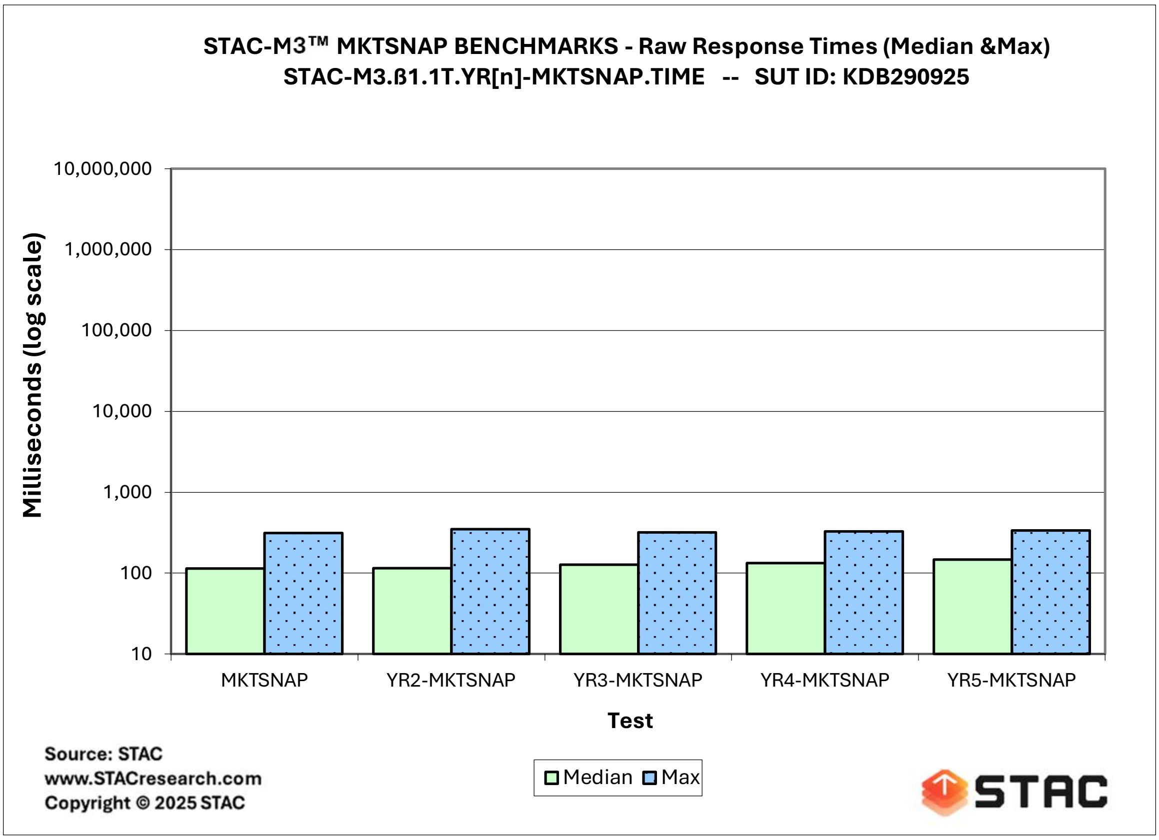
STAC-M3™ MKTSNAP BENCHMARKS - Raw Response Times (Median &Max)
STAC-M3.ß1.1T.YR[n]-MKTSNAP.TIME -- SUT ID: KDB290925
Milliseconds (log scale)
10,000,000
1,000,000
100,000
10,000
1,000
100
10
MKTSNAP
YR2-MKTSNAP
YR3-MKTSNAP
YR4-MKTSNAP
YR5-MKTSNAP
Test
Source: STAC
www.STACresearch.com
Copyright © 2025 STAC
Median
Max
STAC

Figure 2

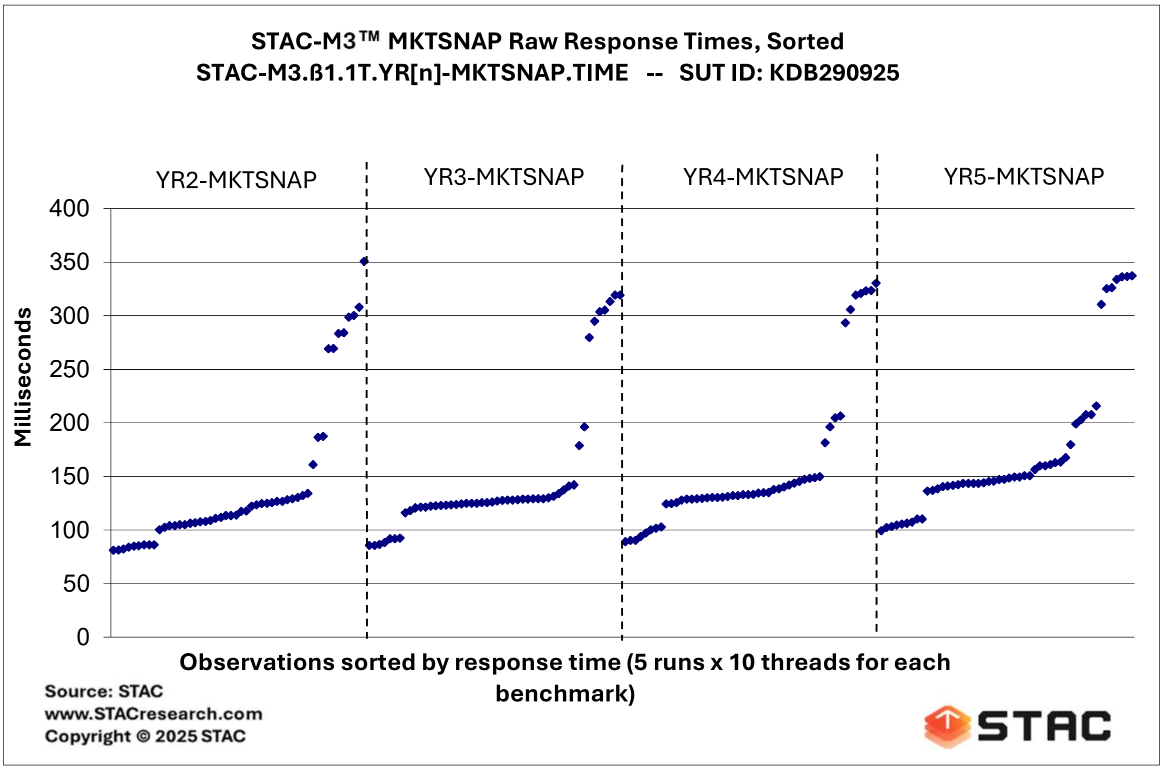

STAC-M3™ MKTSNAP Raw Response Times, Sorted
STAC-M3.β1.1T.YR[n]-MKTSNAP.TIME -- SUT ID: KDB290925
YR2-MKTSNAP
YR3-MKTSNAP
YR4-MKTSNAP
YR5-MKTSNAP
Milliseconds
400
350
300
250
200
150
100
50
0
Observations sorted by response time (5 runs x 10 threads for each benchmark)
Source: STAC
www.STACresearch.com
Copyright © 2025 STAC
STAC

Figure 3

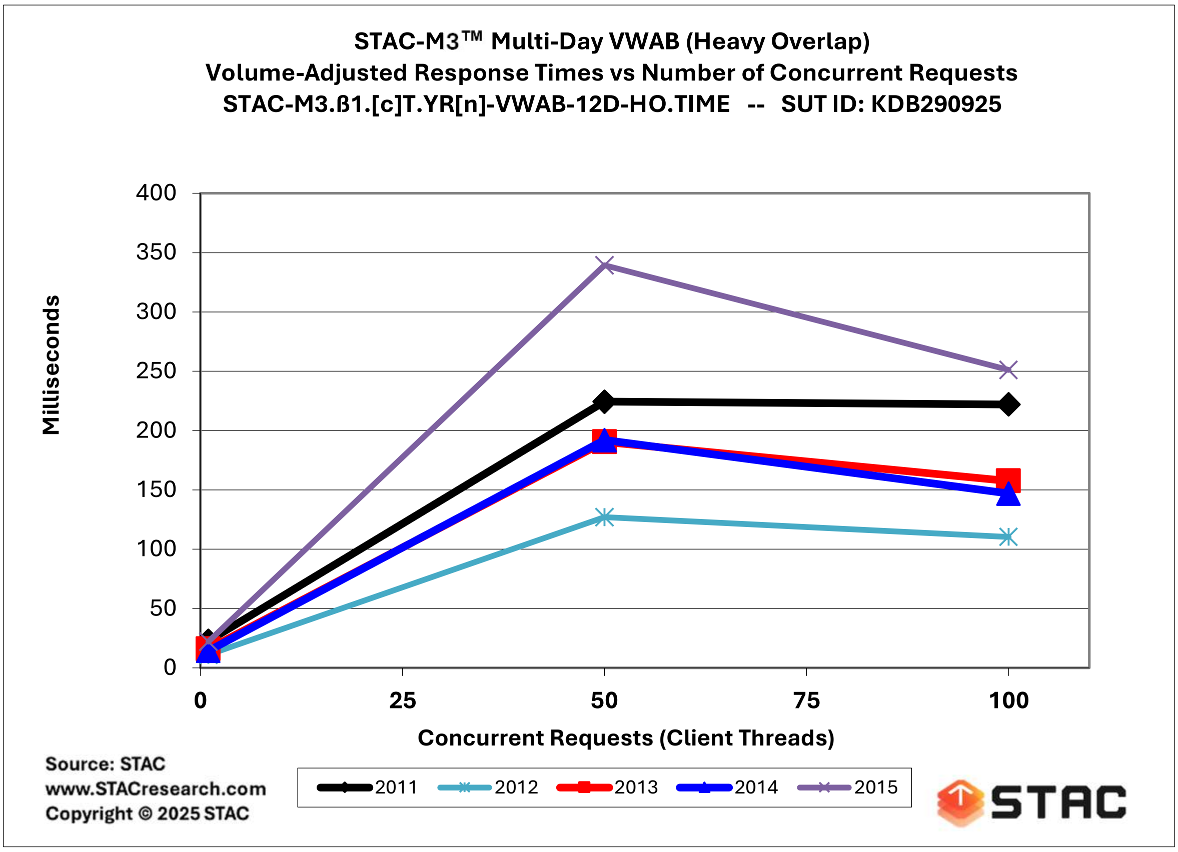
STAC-M3™ Multi-Day VWAB (Heavy Overlap)
Volume-Adjusted Response Times vs Number of Concurrent Requests
STAC-M3.ß1.[c]T.YR[n]-VWAB-12D-HO.TIME -- SUT ID: KDB290925
Milliseconds
400
350
300
250
200
150
100
50
0
0
25
50
75
100
Concurrent Requests (Client Threads)
Source: STAC
www.STACresearch.com
Copyright © 2025 STAC
2011
2012
2013
2014
2015
STAC

Figure 4