Kx Systems kdb+ 2.8 on an on IBM x3850 X5 with ION Computer Systems® STORION™ SR-71 Storage Appliance with Intel® SSD 320 Series

SUT ID KDB120507

STAC-M3™ BENCHMARKS (Antuco Suite)

Test date: 07 May 2012

Draft 1.0, 5 June 2012

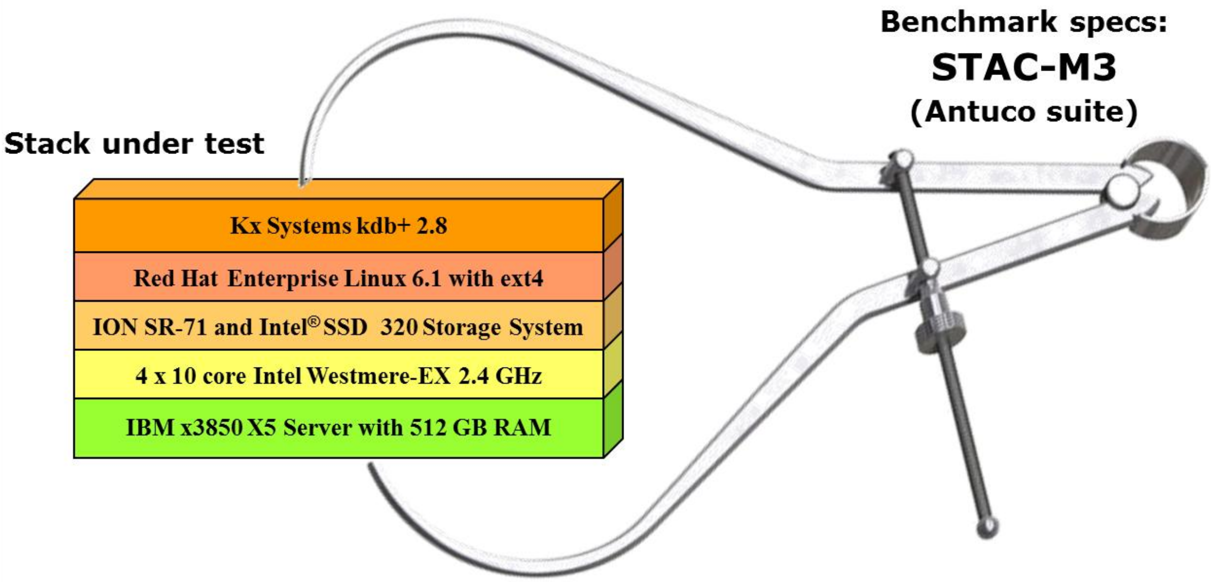

THESE TESTS FOLLOWED STAC BENCHMARK SPECIFICATIONS PROPOSED OR APPROVED BY THE STAC BENCHMARK COUNCIL (SEE WWW.STACRESEARCH.COM). BE SURE TO CHECK THE VERSION OF ANY SPECIFICATION USED IN A REPORT. DIFFERENT VERSIONS MAY NOT YIELD RESULTS THAT CAN BE COMPARED TO ONE ANOTHER.

Disclaimer

Contents

References

[1] Specifications used for this benchmark: STAC-M3 Benchmark Specifications, Antuco Suite, Rev M – http://www.stacresearch.com/node/8777. Accessible by qualified members of the STAC Benchmark Council.

Summary

STAC recently performed the baseline STAC-M3™ Benchmarks on a stack involving Kx Systems kdb+ 2.8 and the STORION SR71 iSCSI Storage Appliance, 100% SSD appliance and is based on the Intel® SSD 320 Series 600GB offering. The kdb+ software was hosted on an IBM x3850 X5 Server with four Intel Xeon E7-4870 CPUs. The server was connected to the storage using 10 Gigabit Ethernet (10 GbE). Four Intel x520 10 GbE adapters in the server communicated via a Netgear XSM7224S switch to four 10 GbE adapters in the STORION appliance.

This report documents the benchmark results. In all, the STAC-M3 specifications deliver dozens of test results, which are presented through a variety of tables and visualizations in this report. Of these, Intel chose to highlight a few, as follows:

> *Results for the MKTSNAP, THEOPL, and VWAB-D benchmarks were 2.2x to 2.7x better than the best corresponding public results to date for SUTs using spinning disks. These speeds were accomplished with a storage subsystem that takes up only 2U of rack space, consumes less than 500 Watts at full load, and costs much less than comparable spindle-based storage systems.*

A STAC-M3 Report Card appears at the top of the report as a convenience for readers who want to get straight to the results. However, we recommend that readers who are not part of the STAC-M3 Working Group first read Section 1 (Overview) to get a feel for the test cases and metrics.

Report Card

STAC-M3™ Benchmarks for SUT ID KDB120507:
Kdb+ 2.8 / IBM x3850 X5 / Intel E7-4870 2.4 GHz / ION SR-71 / Intel® SSD 320 Series / RHEL 6.1 / ext4

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	**172%**

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.

		Latency (milliseconds)	
Spec ID		MEAN	MAX
STAC-M3.ß1.1T.YRHIBID.LAT2	Last-result latency	11,555	11,654
STAC-M3.ß1.1T.YRHIBID-2.LAT2	Last-result latency	4,114	4,179
STAC-M3.ß1.1T.QTRHIBID.LAT2	Last-result latency	3,114	3,231
STAC-M3.ß1.1T.MOHIBID.LAT2	Last-result latency	1,104	1,117
STAC-M3.ß1.1T.WKHIBID.LAT2	Last-result latency	500	504
		MB/second	
Spec ID		MEAN	MAX
STAC-M3.ß1.1T.YRHIBID.BPS	Bytes-read per second*	1,319	1,338
STAC-M3.ß1.1T.YRHIBID-2.BPS	Bytes-read per second*	0	0
STAC-M3.ß1.1T.QTRHIBID.BPS	Bytes-read per second*	1,175	1,243
STAC-M3.ß1.1T.MOHIBID.BPS	Bytes-read per second*	1,102	1,123
STAC-M3.ß1.1T.WKHIBID.BPS	Bytes-read per second*	608	613

* Bytes 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.

		Latency (milliseconds)	
Spec ID		MEAN	MAX
STAC-M3.v1.1T.WRITE.LAT2	Write-completion latency	18,089	18,204

STAC Report Card (cont'd)

STAC-M3™ Benchmarks for SUT ID KDB120507:

Kdb+ 2.8 / IBM x3850 X5 / Intel E7-4870 2.4 GHz / ION SR-71 / Intel® SSD 320 Series / RHEL 6.1 / ext4

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

		Latency (milliseconds)	
Spec ID		MEAN	MAX
STAC-M3.v1.1T.VWAB-D.LAT1	First-result latency	969	1,004
STAC-M3.v1.1T.VWAB-D.LAT2	Last-result latency	969	1,004

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.

		Latency (milliseconds)				
Spec ID		MEAN	MED	MIN	MAX	STDV
STAC-M3.ß1.10T.THEOPL.LAT1	First-result latency	2,741	2,680	2,578	3,073	146
STAC-M3.ß1.10T.THEOPL.LAT2	Last-result latency	2,741	2,680	2,578	3,073	146

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.

		Latency (milliseconds)				
Spec ID		MEAN	MED	MIN	MAX	STDV
STAC-M3.ß1.10T.MKTSNAP.LAT1	First-result latency	2,094	2,054	1,760	2,847	224
STAC-M3.ß1.10T.MKTSNAP.LAT2	Last-result latency	2,094	2,054	1,760	2,847	224

STAC Report Card (cont'd)

STAC-M3™ Benchmarks for SUT ID KDB120507:

Kdb+ 2.8 / IBM x3850 X5 / Intel E7-4870 2.4 GHz / ION SR-71 / Intel® SSD 320 Series / RHEL 6.1 / ext 4

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.

		Latency (milliseconds)				
Spec ID		MEAN	MED	MIN	MAX	STDV
STAC-M3.ß1.10T.VOLCURV.LAT1	First-result latency	26,275	26,247	21,695	30,135	1,775
STAC-M3.ß1.10T.VOLCURV.LAT2	Last-result latency	26,275	26,247	21,695	30,135	1,775

Aggregated Stats (10 Client Threads Requesting)

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

		Latency (milliseconds)				
Spec ID		MEAN	MED	MIN	MAX	STDV
STAC-M3.ß1.10T.STATS-AGG.LAT1	First-result latency	73,307	71,543	62,356	99,501	8,982
STAC-M3.ß1.10T.STATS-AGG.LAT2	Last-result latency	73,307	71,543	62,356	99,501	8,982

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.

		Latency (milliseconds)				
Spec ID		MEAN	MED	MIN	MAX	STDV
STAC-M3.ß1.1T.STATS-UI.LAT1	First-result latency	26,816	26,636	26,035	28,475	870
STAC-M3.ß1.10T.STATS-UI.LAT1	First-result latency	41,806	41,313	32,681	59,074	6,122
STAC-M3.ß1.50T.STATS-UI.LAT1	First-result latency	116,998	119,816	33,087	210,298	51,606
STAC-M3.ß1.100T.STATS-UI.LAT1	First-result latency	197,659	199,622	32,877	371,318	97,168
STAC-M3.ß1.1T.STATS-UI.LAT2	Last-result latency	26,816	26,636	26,035	28,475	870
STAC-M3.ß1.10T.STATS-UI.LAT2	Last-result latency	41,806	41,313	32,681	59,074	6,122
STAC-M3.ß1.50T.STATS-UI.LAT2	Last-result latency	116,998	119,816	33,087	210,298	51,606
STAC-M3.ß1.100T.STATS-UI.LAT2	Last-result latency	197,659	199,622	32,877	371,318	97,168

STAC Report Card (cont'd)

STAC-M3™ Benchmarks for SUT ID KDB120507:

Kdb+ 2.8 / IBM x3850 X5 / Intel E7-4870 2.4 GHz / ION SR-71 / Intel® SSD 320 Series / RHEL 6.1 / ext 4

NBBO Benchmark

NBBO

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

		Latency (milliseconds)	
Spec ID		MEAN	MAX
STAC-M3.ß1.1T.NBBO.LAT2	Write-completion latency	247,730	252,878

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

		Latency (milliseconds)				
Spec ID		MEAN	MED	MIN	MAX	STDV
STAC-M3.v1.100T.VWAB-12D-NO.LAT1	First-result latency	29,569	29,573	4,390	66,096	15,049
STAC-M3.v1.100T.VWAB-12D-NO.LAT2	Last-result latency	29,569	29,573	4,390	66,096	15,049

Chart view

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

- Figures 1 through 4 plot the mean last-result latency (LAT2) benchmarks for all of the operations.
- Figures 5 and 6 analyze the individual latency observations for the multi-user/multi-day VWAB benchmark (STAC-M3.v1.100T.VWAB-12D-NO.LAT2), first by sorting the results by latency, then by plotting them in a histogram.
- Figure 7 provides a more explicit look at multi-user scaling by plotting the latency for the intervalized statistics benchmark (STAC-M3.ß1.[n]T.STATS-UI.LAT2) against the number of simultaneously requesting client threads (n).
- Figures 8 and 9 take the 100-client-thread case of Figure 7 and analyze the individual latency observations, first by sorting the results by latency, then by plotting them in a histogram.

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.

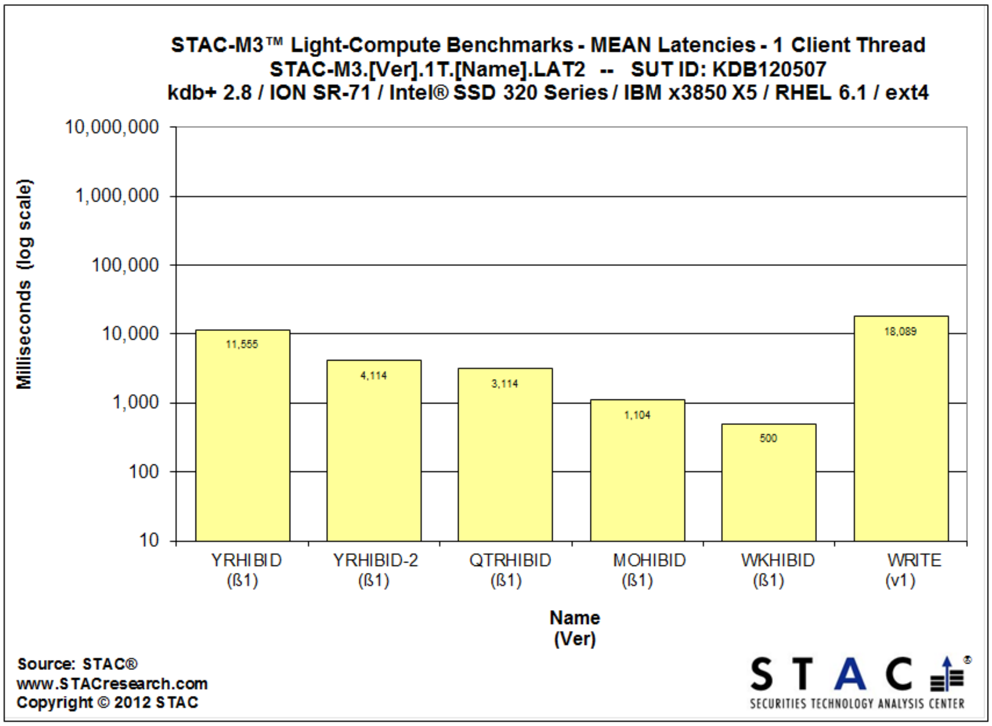

Figure 1

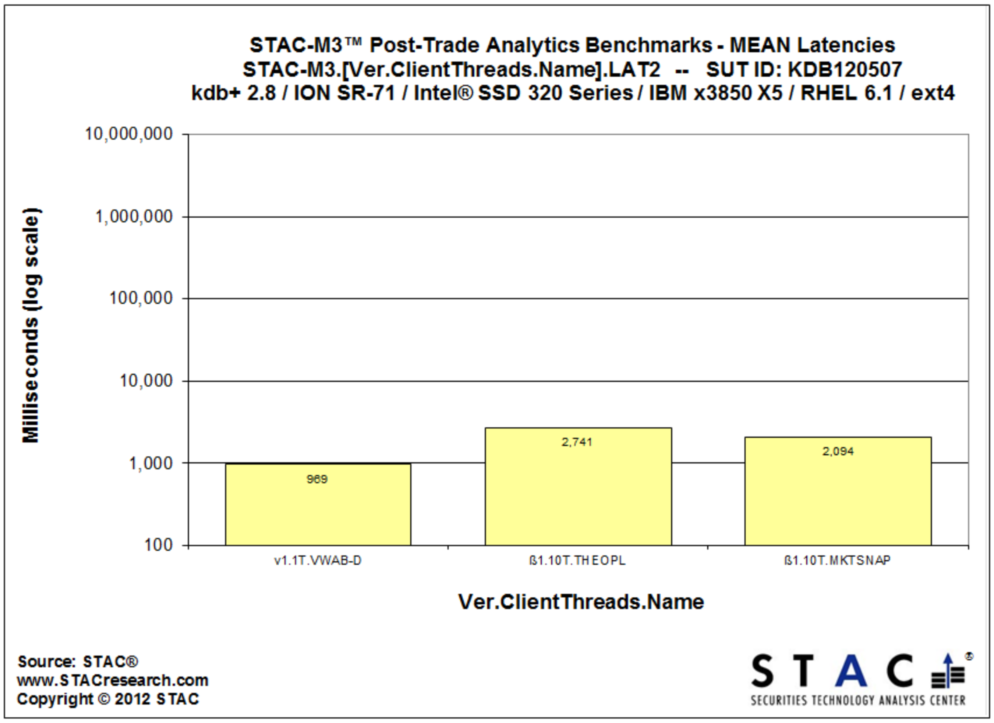

Figure 2

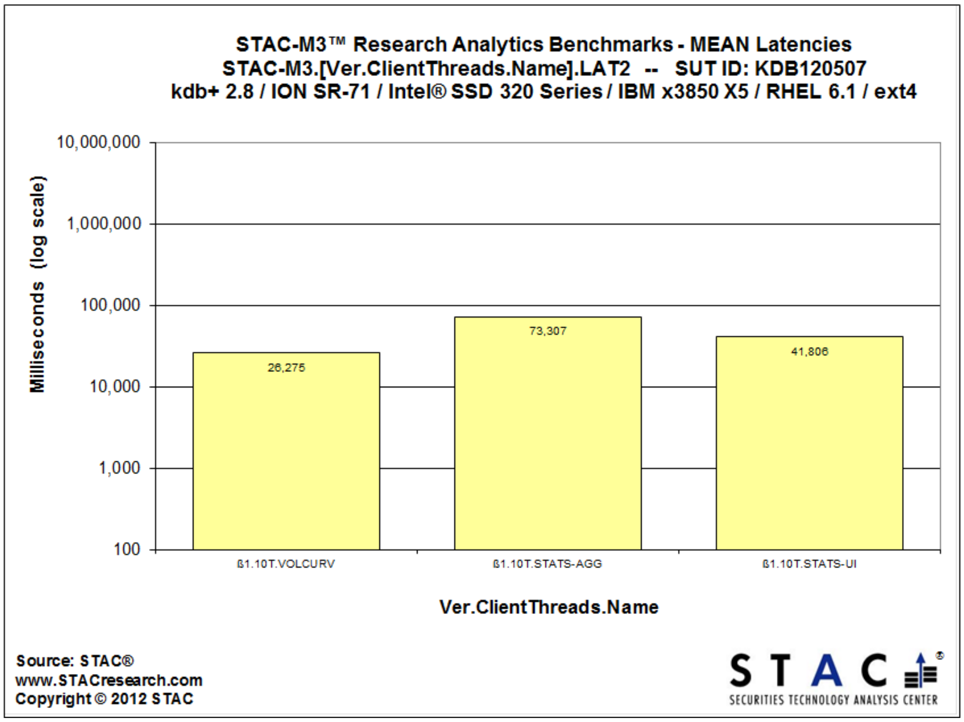

Figure 3

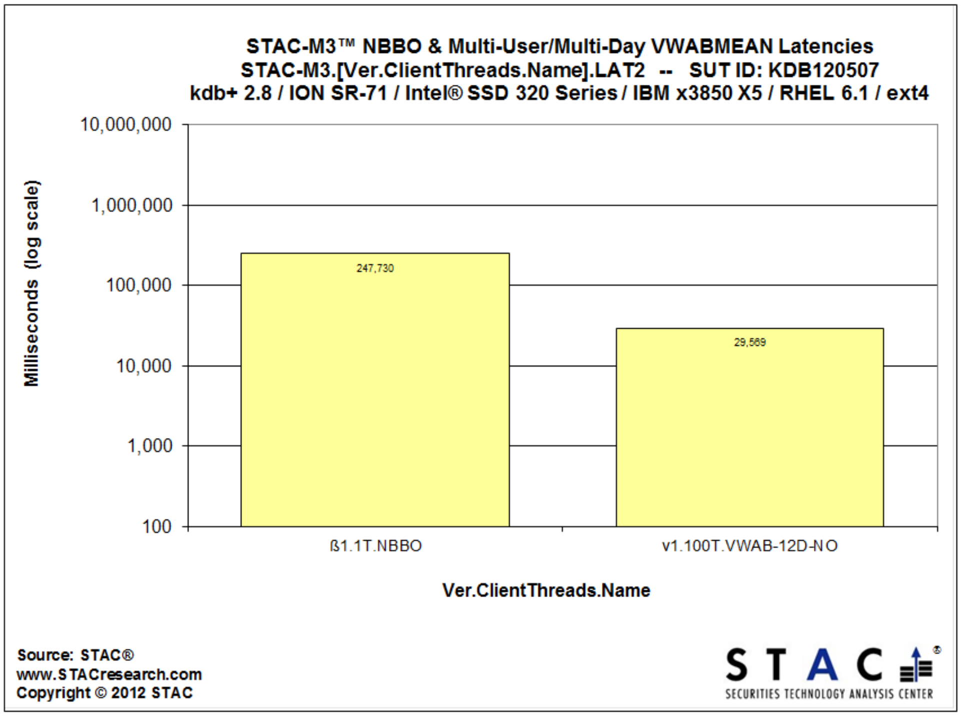

Figure 4

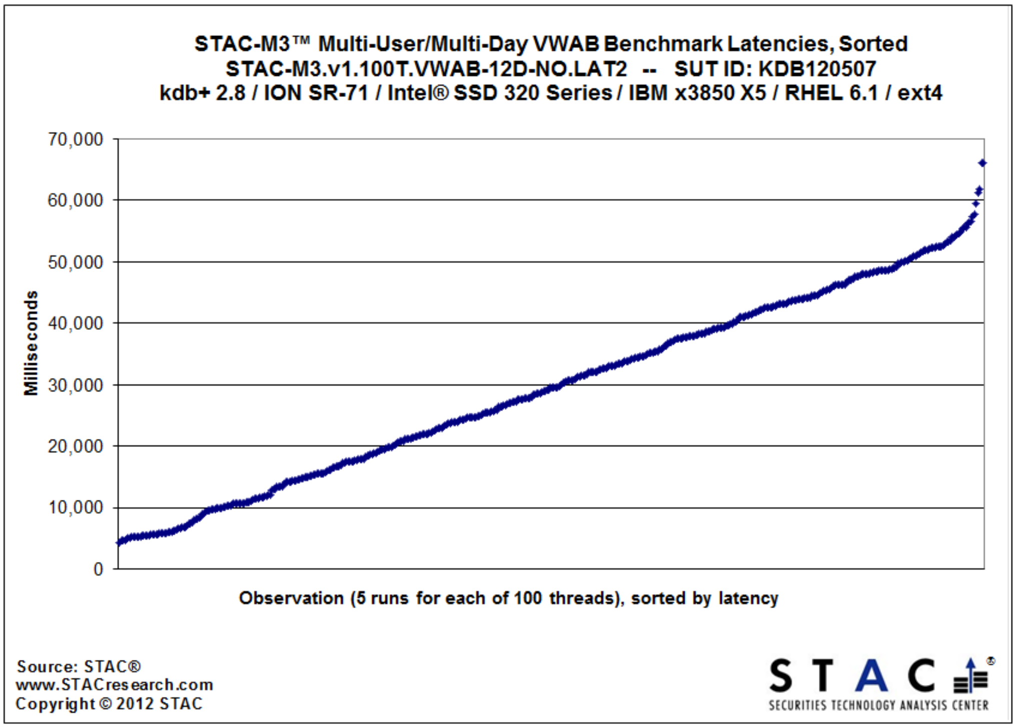

Figure 5

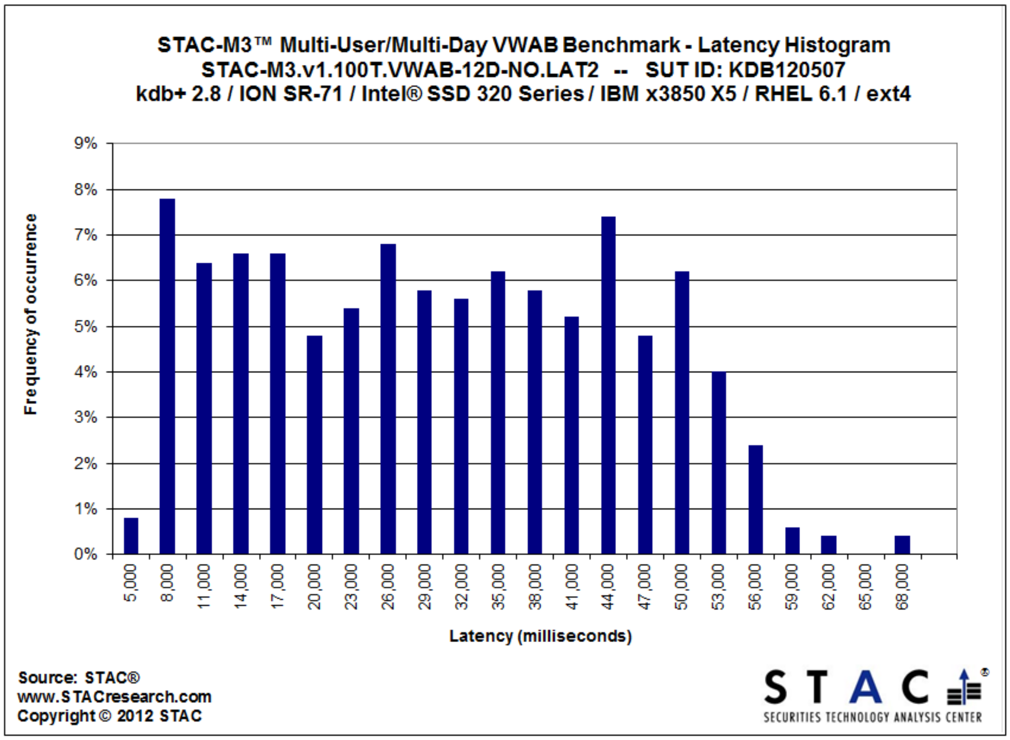

Figure 6

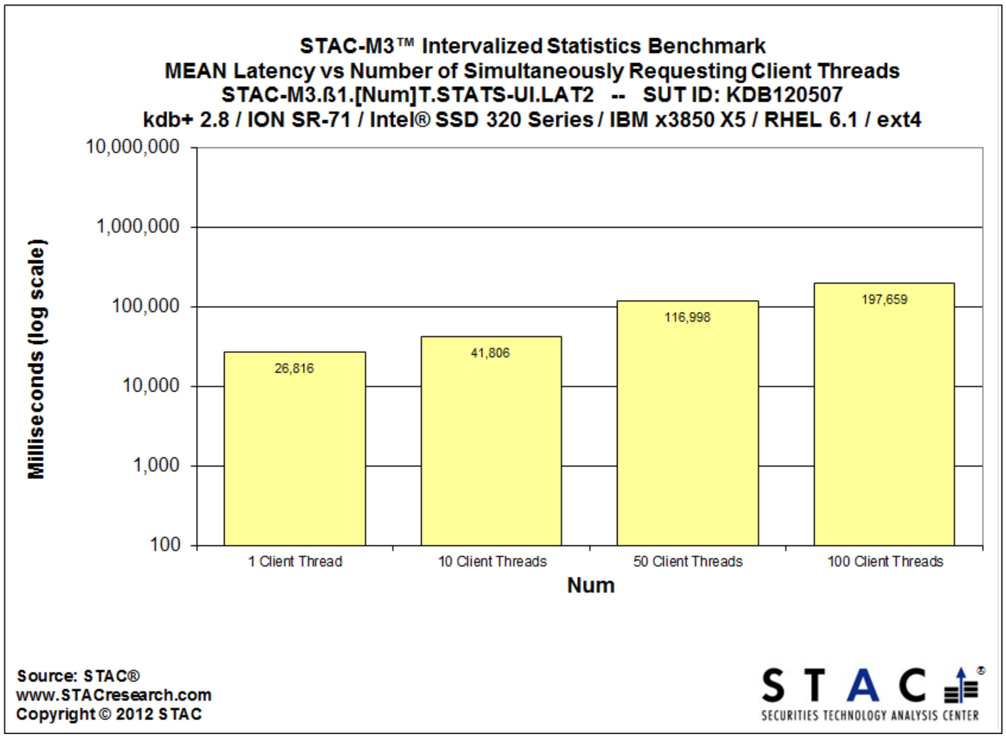

Figure 7

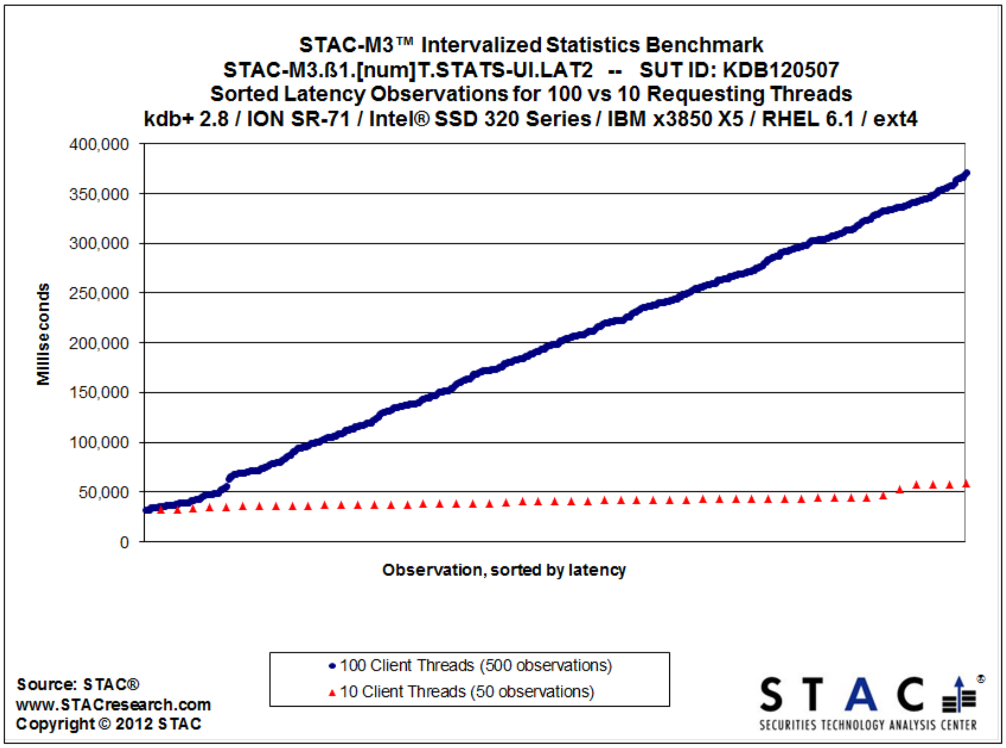

Figure 8

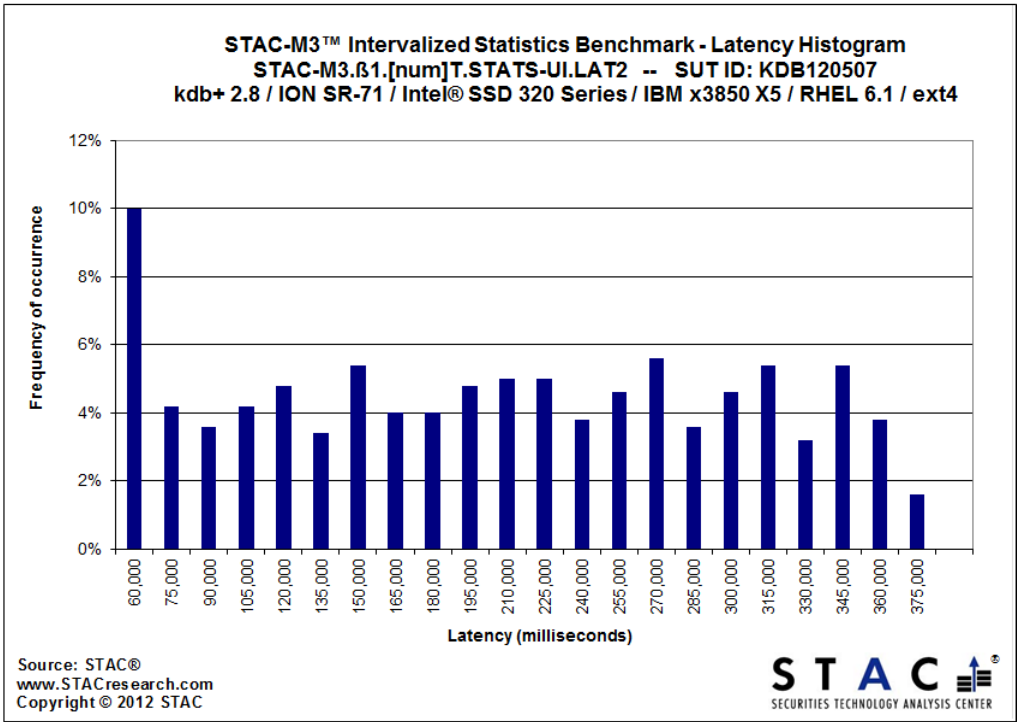
STAC-M3™ Intervalized Statistics Benchmark - Latency Histogram
STAC-M3.ß1.[num]T.STATS-UI.LAT2 -- SUT ID: KDB120507
kdb+ 2.8 / ION SR-71 / Intel® SSD 320 Series / IBM x3850 X5 / RHEL 6.1 / ext4
Frequency of occurrence
12%
10%
8%
6%
4%
2%
0%
60,000
75,000
90,000
105,000
120,000
135,000
150,000
165,000
180,000
195,000
210,000
225,000
240,000
255,000
270,000
285,000
300,000
315,000
330,000
345,000
360,000
375,000
Latency (milliseconds)
Source: STAC®
www.STACresearch.com
Copyright © 2012 STAC
STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

Figure 9

1. 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 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 back-test 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 Council has developed the STAC-M3 Benchmarks in order to provide a common basis for quantifying the extent to which emerging hardware and software innovations improve the performance of tick storage, retrieval, and analysis.

STAC-M3 tests the ability of a solution stack such as columnar database software, servers, and storage, 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):

- Different storage systems, including SSD, DRAM, interconnects, and file systems
- Different server products, processors, chipsets, and memory
- Different tick-database products

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 latency calculations are subjected to a combination of source-code inspection and empirical validation.

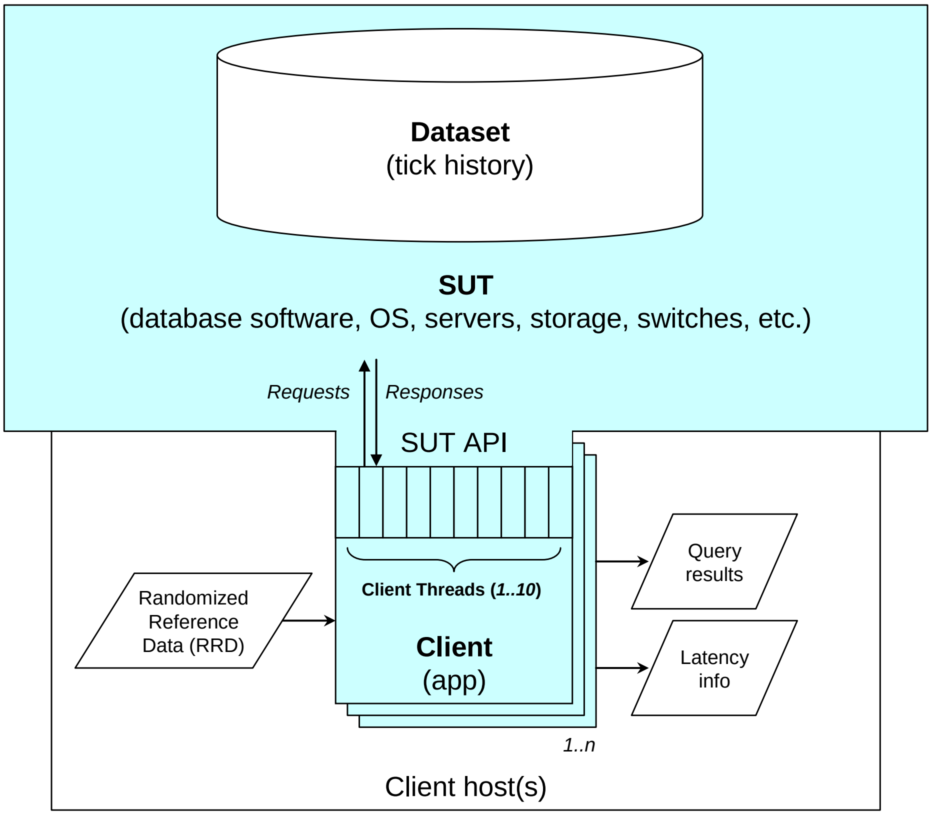

Dataset

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

- 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 and run benchmarks against projected future data volumes.

The dataset consists of high-volume symbols and low-volume symbols in proportions based on observed NYSE data. The data volume per symbol 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. This was a deliberate choice by the STAC-M3 Working Group to minimize the cost of running benchmarks while still yielding valuable results. Benchmarks that scale the database to the size of existing customer footprints and well beyond are under consideration for a future STAC-M3 suite.

Metrics

The key metric in STAC-M3 is the latency of query responses (aka response times). Latency measurements are performed in the clients. A client thread gets a local timestamp (t_{submit}) just before submitting a query. When the first results arrive, the client gets another timestamp (t_{first}). When it receives the complete results (sorted appropriately), the client immediately gets a third timestamp (t_{last}). For systems that return all results in one chunk, the first-result and last-result timestamp are identical. As the diagram below illustrates, latencies are defined as follows:

$$\text{First-result latency (LAT1)} = t_{first} - t_{submit}$$

$$\text{Last-result latency (LAT2)} = t_{last} - t_{submit}$$

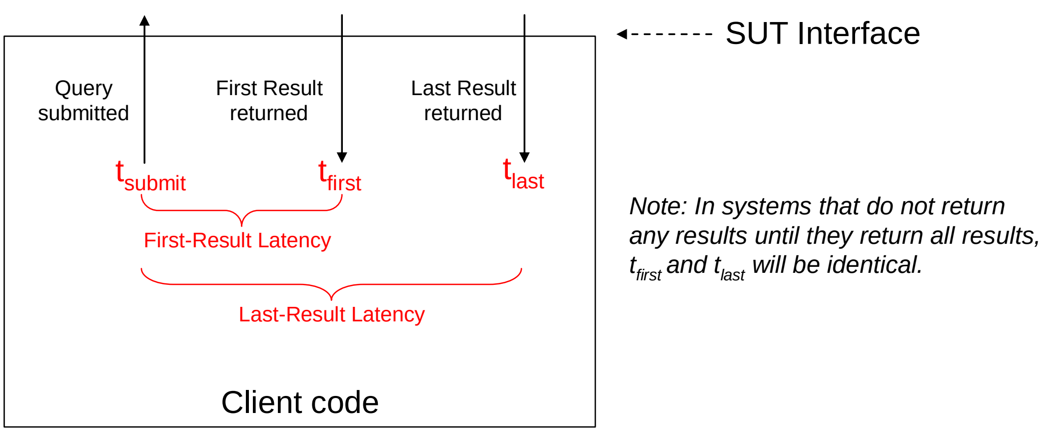

Timestamp and latency meanings

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.

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.

Test cases

The current tests in the STAC-M3 suite are listed in the Summary Table below. As the versioning illustrates, the STAC-M3 Working Group is bringing benchmark specifications to market in phases. The first set of approved specs (those marked as “v1” or, in one case, “v1.1”) focused on storage-system performance with respect to heavy historical data loads. These workloads were deliberately light on compute and heavy on I/O. The second phase, (the Antuco suite, see [1]) added benchmark specs that involve more compute-intensive analytics. These new specifications, marked as “ß1”, have not been put to a vote by the full STAC Benchmark Council and will become v1 specs if approved. These new benchmarks operate by symbol on many fields of underlying tick data for both trades and quotes across varying time windows.[1] 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. 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 and allowing the tester to scale to even higher numbers of concurrent threads. In all cases, benchmark results refer to per-request latency. For example, the mean of 10T.MKTSNAP.LAT2 is the mean time to satisfy a single market-snapshot request, not the total time to satisfy requests from all 10 client threads.

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 is to provide a realistic optimization strategy for systems with multiple storage tiers of different speed, such as solid-state disk (SSD) and spindle-based storage. For example, suppose the tester wanted to demonstrate the benefit of adding an SSD to a system that otherwise relied on spindle-based storage. Furthermore, suppose the SSD was only large enough to store 1/5 of the test database. Benchmarks that accessed dates throughout the entire database would show some performance improvement when the SSD is added, since 1/5 of the queries would enjoy acceleration. In the real world, however, customers don’t simply add faster storage and randomly allocate data to it. Rather, they typically reserve the fastest storage tier for the most frequently accessed data or data accessed by especially time-sensitive algorithms. In STAC-M3, the tester can load the most recent data into the fastest storage. In the case above, the benchmarks that operate on, say, the most recent day (1/252 of the database) or the most recent month (1/12 of the database) will show the maximum improvement possible from the SSD, since their entire queries can be satisfied from SSD. This provides the clearest indication of acceleration possibilities for datasets that can fit within the faster tier.

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

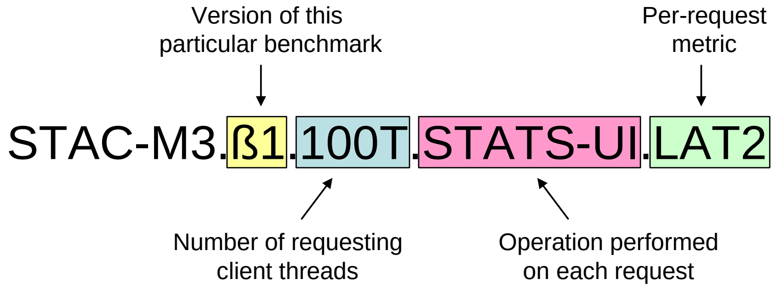

In charts, the ID is sometimes decomposed, with part of it in the chart title or labels. Each individual STAC Benchmark™ specification has its own version number. The same version of a given spec may appear in multiple benchmark suites. Thus, the code names of the suites are irrelevant when making comparisons. 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 cannot be fairly compared.

[1] Future phases will supplement these with benchmarks that provide insight into additional aspects of system performance. Contact council@STACresearch.com if you would like to be part of the process.

Summary Table – 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 Council.

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	ß1	1	Max bid over the year for 1% of symbols.	Heavy read	Light	Full year
YRHIBID-2	Year High Bid Re-run	ß1	1	Re-run of YRHIBID (same symbols) without clearing the cache.	Heavy read†	Light	Full year
QTRHIBID	Quarter HighBid	ß1	1	Max bid over the quarter for 1% of symbols.	Heavy read	Light	Most recent quarter
MOHIBID	Month High Bid	ß1	1	Max bid over the month for 1% of symbols.	Heavy read	Light	Most recent month
WKHIBID	Week High Bid	ß1	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.

2. Product background

The stack under test (SUT) included the following:

- Kx Systems kdb+ 2.8
- Red Hat Enterprise Linux 6.1 with ext4
- IBM Systems x3850 X5 Server
- ION STORION™ SR-71 iSCSI Storage Appliance with Intel SSD

Kx Systems submitted the following information and claims about its products:

Kx delivers a combination of kdb+, an ultra high-performance database with a unified format for real-time and historical data and q, an exceptionally efficient proprietary language, which enable trading operations to manage risk and implement sophisticated trading strategies in real-time. Kx simultaneously supports thousands of real-time custom queries and analyses on historical / in-memory data. kdb+ is a scalable, column-oriented database for processing massive data volumes.

The kdb+tick solution is a tick data capture application layered on the kdb+ database. kdb+ includes q, a high-level language for querying and programming designed to work with time-series data, providing interoperability with other databases and open interfaces to C, C++, Java and .Net applications. kdb+ runs on 32-bit and 64-bit versions of Linux, Mac OS X, Solaris and Microsoft Windows operating systems.

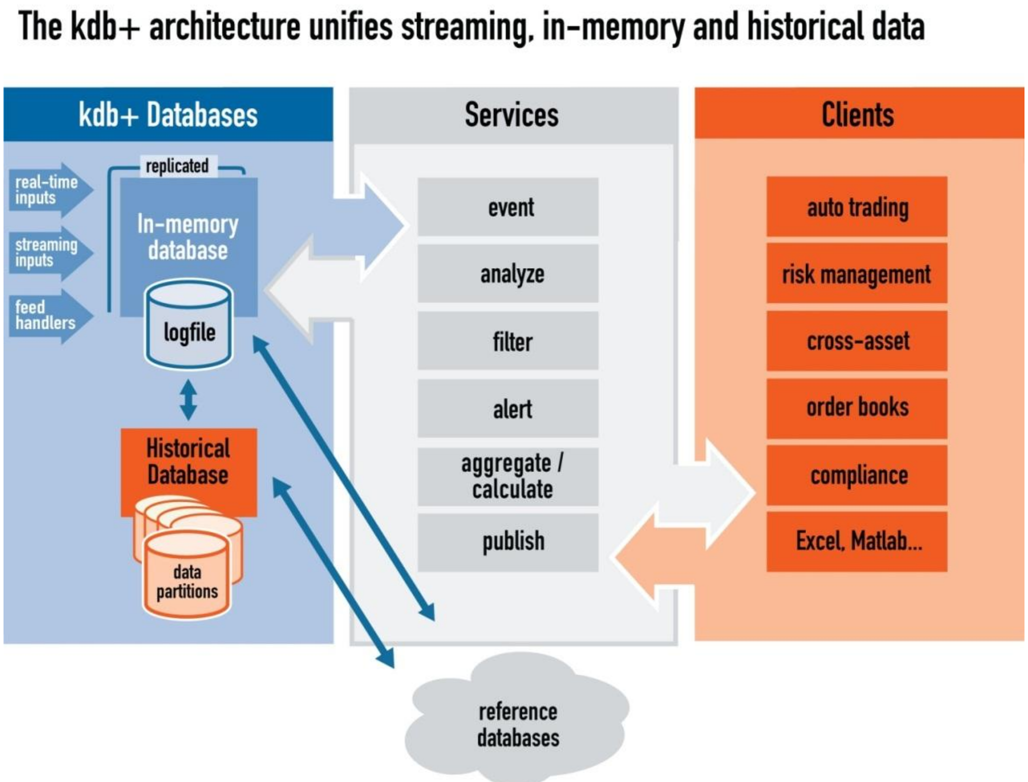

Figure 9 - View of kdb+ system architecture

ION submitted the following information and claims about its products:

The STORION™ SR-71 iSCSI Storage Appliance is unique in the market as it has been designed from the ground up as a platform to utilize SSD devices for primary storage. Other vendors have taken existing platforms that have used spinning disks in prior generations and simply changed to SSD. This type of *design drives up the price of the platform dramatically, but fails to deliver the potential performance that SSDs are capable of, resulting in a poor ROI.*

The STORION™ SR-71 iSCSI Storage Appliance utilizes only the optimal components to deliver I/O performance. These include:

- *Two Eight Core E-5 Series Intel Xeon Processors*
- *SSD Optimized Hardware RAID Controllers*
- *Industry-leading MLC and eMLC SSDs*
- *iSCSI Software Stack optimized for performance*
- *10 GB Ethernet interfaces optimized for TCP and iSCSI offload*

Additional attributes of the STORION™ SR-71 iSCSI Storage Appliance are:

- *2U Rackmount Chassis*
- *(24) Storage SSDs in front hot-swap trays*
- *(2) Boot SSDs in rear hot-swap trays*
- *3.45"(H) X 17.5"(W) X 26"(D)*
- *650W Redundant Power Supply*
- *44 lbs.*

iSCSI connectivity via 10 GB Ethernet allows the STORION™ SR-71 iSCSI Storage Appliance to deliver 9.1 TB of usable capacity in a reliable RAID 5 configuration with optimal overprovisioning of the SSD devices to your environment. In this configuration, performance can be delivered at high IOPS and throughput with very low latencies. To increase capacity to your storage system, you can easily add another STORION™ SR-71 iSCSI Storage Appliance to your network. You just need to connect it to your switch and connect your application servers.

The STORION™ SR-71 iSCSI Storage Appliance is the best option for price/performance in a 2u platform operating under 500 Watts. To view additional features and options please go to http://www.storion.net

IBM submitted the following information and claims about its products:

Built on the next generation of enterprise IBM® X-Architecture® technology and Intel® Xeon® processors, the IBM System x3850 X5 offers superior performance and unbeatable reliability within an energy - and wallet-friendly - design. The x3850 X5 server allows freedom of choice with extremely flexible configurations plus memory expansion capabilities. A modular building block design lets you customize your system for current needs while providing the ability to react to changing workloads. Expand your 4-socket, 64-DIMM x3850 X5 to 4 sockets and 96 DIMMs or up to 8 sockets and 128 DIMMs. Reallocate resources as your environment changes. The x3850 X5 meets your needs today, while providing an easy, cost-effective upgrade path to change your environment when you're ready. For more information, please visit: http://www-03.ibm.com/systems/x/hardware/enterprise/x3850x5/

Intel submitted the following information and claims about the SSD product used in these tests:

The next generation Intel SSD 320 Series offers built-in data protection features, better performance, larger capacities and more value for your money. Built with 25 nanometer (nm) compute-quality Intel® NAND Flash Memory, the Intel SSD 320 Series accelerates C performance where it matters most. With very high random read and sequential read performance, your Server will blaze through the most demanding applications and will handle intense multi-tasking needs. Couple that with very high performance on random and sequential writes to unleash your system.

3. Project participants and responsibilities

The following firms participated in the project:

- Kx Systems
- ION Computer Systems
- Intel Corporation
- IBM
- STAC

The Project Participants had the following responsibilities:

- Kx implemented the STAC-M3 Clients and Operations using the STAC-M3 Benchmark specifications.
- ION Computer supplied the storage hardware for the test.
- IBM provided the SUT host (IBM x3850 X5) for the test.
- Intel configured and optimized the stack under test (SUT).
- Intel sponsored the Audit.
- 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.

4. Contacts

- Kx Systems: Fintan Quill (info@kx.com, tel: +1 212 792 4230)
- ION Computer Systems Cameron Campbell, Director of Business Development, 631-630-3237, cameron.campbell@ioncomputer.com
- Intel: David O'Shea, Financial Services Software Manager (David.Oshea@intel.com) or Rick Carlin, Senior Performance Engineer (Rick.Carlin@intel.com). See also information posted at http://ipip.intel.com/FSI and http://software.intel.com/en-us/articles/financial-services-industry-community/
- IBM: http://www-03.ibm.com/systems/services/benchmarkcenter/

5. Results status

- These benchmark specifications were developed by the STAC-M3 Working Group of the STAC Benchmark Council. Benchmarks with a "v1" or higher have been approved by the full Council. Those with a "ß" designation have been proposed by the STAC-M3 Working Group but have not yet been approved.
- 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.

6. Specifications

This project followed the Antuco suite of STAC-M3 Benchmark specifications. Full members of the STAC Benchmark Council can access these specifications at www.STACresearch.com/m3. Premium subscribers to the Market Data or Trade Flow STAC Tracks can download the programs used in these benchmarks in order to run the same tests on systems in the privacy of their own labs.

7. Limitations

- As discussed in Section 1, this suite of STAC-M3 Benchmarks was designed to test operations on a limited amount of purely historical data. Tests involving much larger amounts of historical data and tests involving real-time data are under consideration for future STAC-M3 suites.
- 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.

8. Stack under test

As described above, the stack under test (SUT) included the following:

- Kx Systems kdb+ 2.8
- Red Hat Enterprise Linux 6.1 with ext4
- IBM Systems x3850 X5 Server
- ION STORION™ SR-71 iSCSI Storage Appliance with Intel SSD

Detailed configuration information (e.g., software load numbers, kdb+ par.txt, storage/file system/volume/partition configuration, affinity settings, OS configuration, etc.) is available to premium members of the STAC Benchmark Council. If this document is not in your myVault account (www.STACresearch.com/myVault), and you would like to request access, use the request form at:

www.STACresearch.com/node/11940

9. Vendor Commentary

Kx Systems provided the following comments:

- *Because kdb+ is an array database the complete results of queries are returned from the calls, so the concept of "first result" (LAT1) has no meaning in the context of the Kx product.*