STAC Update:
Big Workloads

Bishop Brock
Head of Research, STAC
bishop.brock@STACresearch.com

Peter Nabicht
President, STAC
peter.nabicht@STACresearch.com

SECURITIES TECHNOLOGY ANALYSIS CENTER

Overview

- STAC-A3 (backtesting)
- STAC-A2 (derivatives risk computation)
- STAC-M3 (tick history / timeseries analysis)

STAC-A3

- Overview:
  - Workloads that emulate real-world backtesting jobs
  - Measures speed, throughput, scalability, efficiency of any architecture
- Architectures: scale-up and scale-out, cloud and bare metal
- Languages: Python, C++, Scala
- Hardware: CPU, GPU
- Different approaches to parallelization and optimizations are very informative

www.stacresearch.com/A3

STAC-A3: Backtesting Benchmark – Current State

- STAC-A3 Mean Revert suite
  - parameterized technical strategy for intraday trading
  - portfolio of instruments
  - Simulated market data provided at the order level
    - Both high- and low-volume instruments are represented
  - SUT must build the order book
  - Strategy evaluated every second
  - Sweeps over EMA parameters
  - Low to moderate I/O
  - Branch-heavy computation

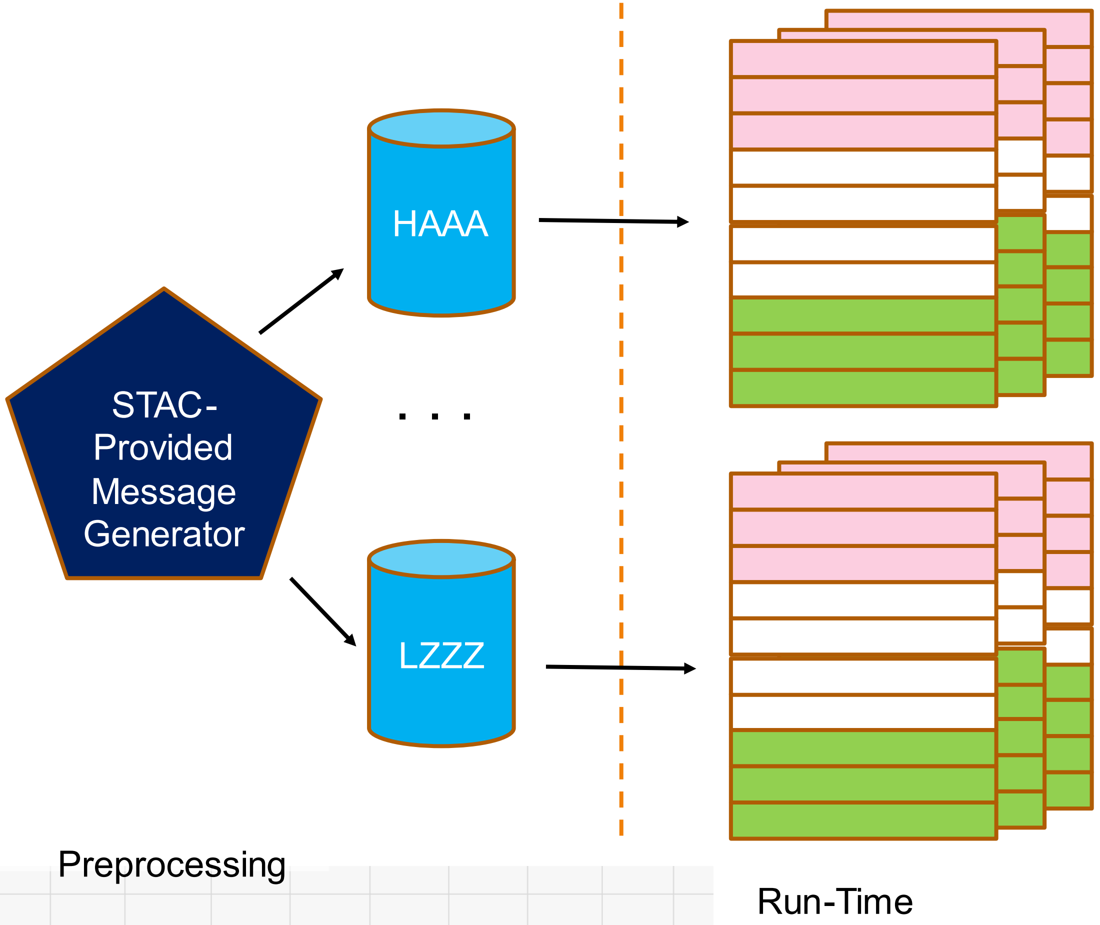

Preprocessing

STAC-A3 Mean Revert: Scaling

- Scaling dimensions
  - Number of instruments traded each day
  - Number of parameter combinations per instrument
- Values chosen by testing vendor
- Looking to standardize scaling points
  - What values are impactful to you?

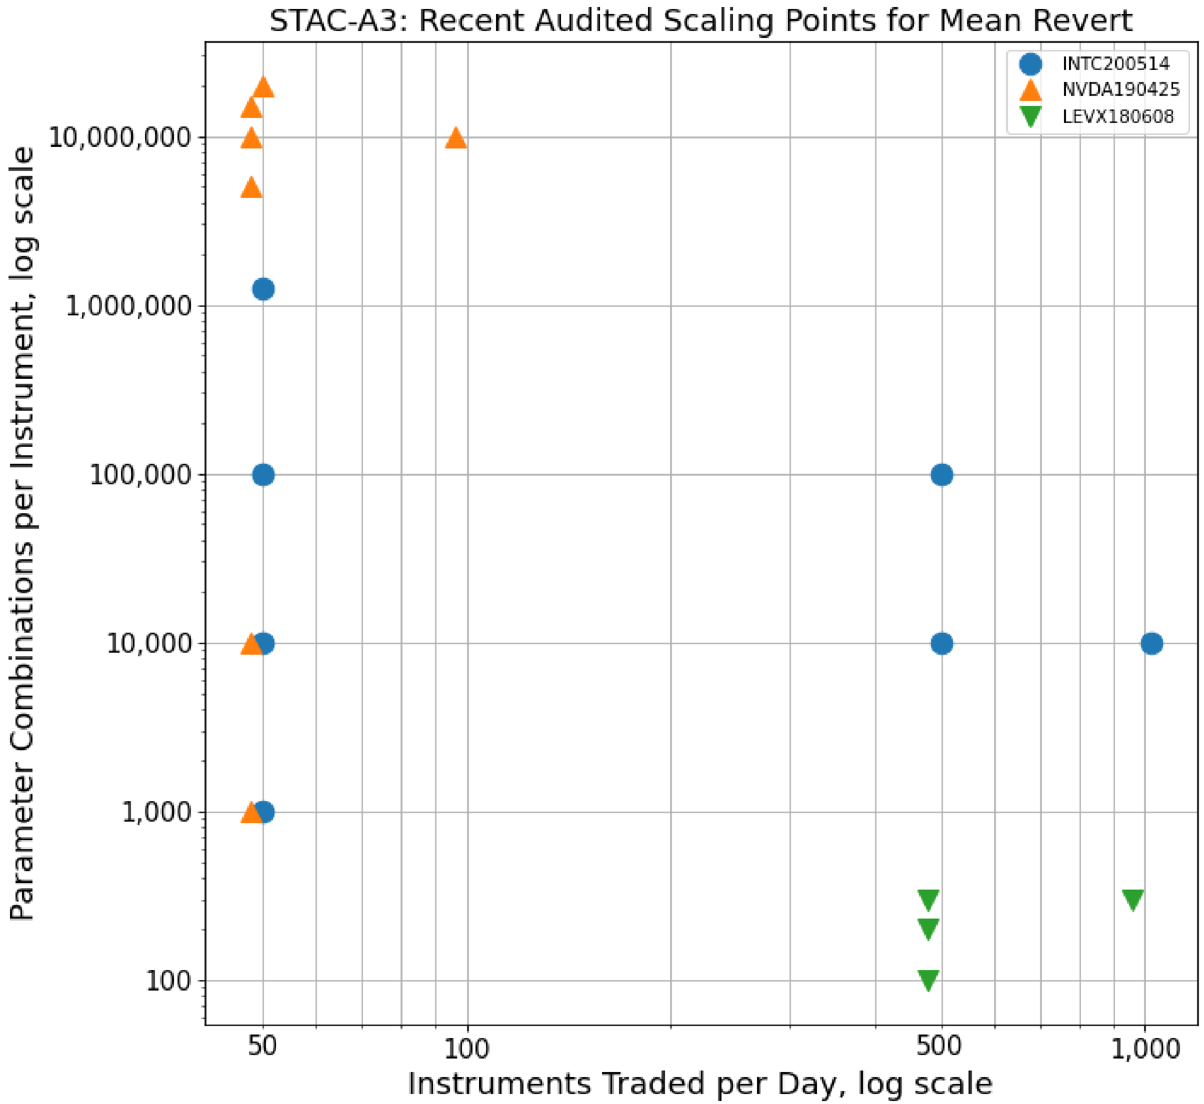

STAC SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-A3: Next Steps

- Mean revert representative of simple strategies
- Working group requested a second, more complex suite

Mean revert suite	Portfolio trading suite proposal
“Embarrassingly parallel”	Fine grained interlock
Low # of instruments = low IO	Portfolio evaluation requires high IO for all instrument scales
Strategy / parameterization has low complexity (low compute)	Strategy requires analysis of orderbook (high compute)
No position sizing	Portfolio optimization

Working Group discussions underway - **Join us!**

www.STACresearch.com/A3

STAC-A2: Risk computation

- Non-trivial Monte Carlo calculations
  - Heston-based Greeks for multi-asset, path-dependent options with early exercise
  - Metrics: Speed, capacity, quality, efficiency
- Numerous reports
  - Some public, some in the STAC Vault
- Premium STAC members get:
  - Reports in STAC Vault
  - Detailed config info on public and private reports
  - Code from vendor implementations of the benchmarks

www.STACresearch.com/a2

A few points on STAC-A2 for the uninitiated

- Some tests measure **response time** for a single option of given problem size
- **Throughput** measures time to handle a portfolio of options
- **Efficiency** relates throughput to power and space
- Each response-time workload is tested 5 times, back-to-back:
  - First run is the **COLD** run
  - Subsequent 4 are **WARM** runs
- COLD relates to real-world systems that must respond to heterogeneous problem classes
  - COLD time includes building memory structures, loading kernels, etc.
- WARM relates to real-world systems configured to handle numerous requests for the same problem class

STAC-A2 / Dell PowerEdge XE8545 / 4 x A100 SXM4 40GB

- First STAC-A2 from Dell Technologies
- STAC-A2 Pack for CUDA (Rev G)
- Stack:
  - NVIDIA CUDA 11.6
  - Dell PowerEdge XE8545
    - 2 x AMD EPYC 7713 64-core processor @ 2.0GHz
    - 4 x NVIDIA A100 SXM4 40GiB GPU
    - 32 x 16GiB Dual Rank ECC DDR4 @ 2933Mhz
  - Red Hat Enterprise Linux 8.3

www.STACresearch.com/NVDA221007

STAC SECURITIES TECHNOLOGY ANALYSIS CENTER

Compared to all publicly reported solutions to date

- Set 3 records:
  - The highest space efficiency
    - STAC-A2.β2.HPORTFOLIO.SPACE_EFF
  - The fastest cold times in the baseline Greeks benchmark
    - STAC-A2.β2.GREEKS.TIME.COLD
  - The fastest cold times in the large Greek benchmark
    - STAC-A2.β2.GREEKS.10-100k-1260TIME.COLD

www.STACresearch.com/NVDA221007

STAC SECURITIES TECHNOLOGY ANALYSIS CENTER

Vs. a solution with 8 x GPUs*

- Had 1.2x the space efficiency
  - STAC-A2.β2.HPORTFOLIO.SPACE_EFF
- Was 2.9x the speed in the cold runs of the baseline Greeks benchmark
  - STAC-A2.β2.GREEKS.TIME.COLD
- Was 1.1x the speed in the cold runs of the Greeks benchmark
  - STAC-A2.β2.GREEKS.10-100k-1260.TIME.COLD
- Was 67% of the speed in the warm runs of the baseline Greeks benchmark
  - STAC-A2.β2.GREEKS.TIME.WARM

www.STACresearch.com/NVDA221007

* SUT ID NVDA210914

STAC® SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-A2 / AWS c5.metal / 2 x Intel® Xeon® Platinum 8275CL

- First AWS-based solution with publicly released STAC-A2 results
- Cloud server with no hypervisor
- STAC-A2 Pack for Intel® oneAPI (Rev N)
- Stack:
  - Intel® oneAPI Base Toolkit 2022.3
  - Intel® oneAPI HPC Toolkit 2022.3
  - c5.metal Amazon Web Services Instance
    - 2 x Intel® Xeon® Platinum 8275CL (Cascade Lake) CPU @ 3.00GHz
    - 96 logical cores
    - 192 GiB of DRAM
  - Amazon Linux release 2 (Karoo)

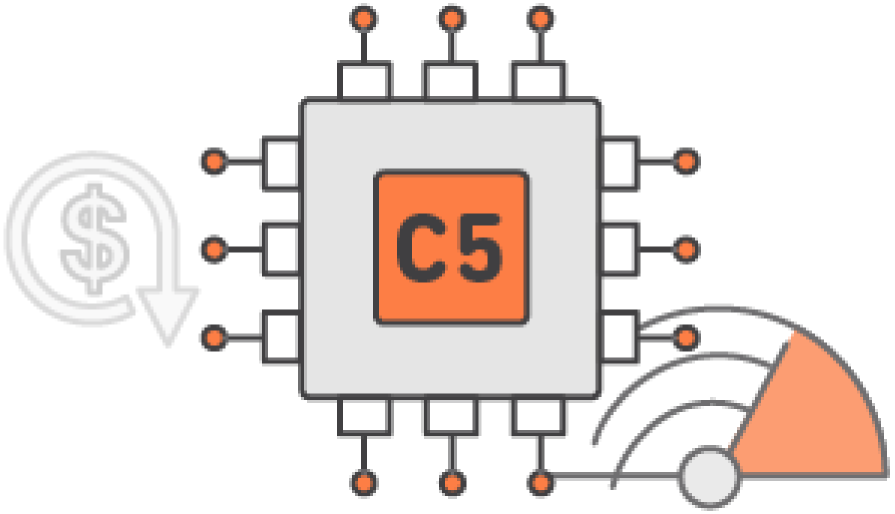

www.STACresearch.com/INTC221006a

Vs. cloud-based SUT with 10 instances and Cascade Lake CPUs*

- Completed 1.1x the options per dollar over a 1 hour burst and a 3-day period
  - STAC-A2.β2.HPORTFOLIO.PRICE_PERF .BURST and .PERIODIC
- Completed 1.2x the options per dollar (reflecting reserve instance pricing discounts) over a 1-year period
  - STAC-A2.β2.HPORTFOLIO.PRICE_PERF .CONTINUOUS

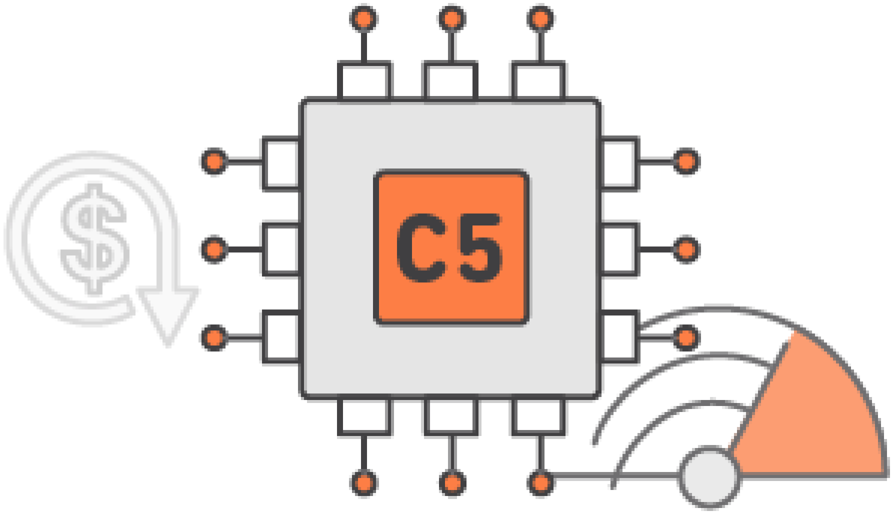

www.STACresearch.com/INTC221006a

* SUT ID INTC210331

Vs. on-prem solution with 2 x Sky Lake CPUs*

- Had 2x the throughput
  - STAC-A2.β2.HPORTFOLIO.SPEED
- In the baseline problem size
  - Was 2.4x the speed in cold runs
    STAC-A2.β2.GREEKS.TIME.COLD
  - Was 1.74x faster in warm runs
    STAC-A2.β2.GREEKS.TIME. WARM
- Was 2.3x the speed in the cold and warm runs of the large problem size
  - STAC-A2.β2.GREEKS.10-100k-1260TIME.COLD
  - STAC-A2.β2.GREEKS.10-100k-1260TIME.WARM

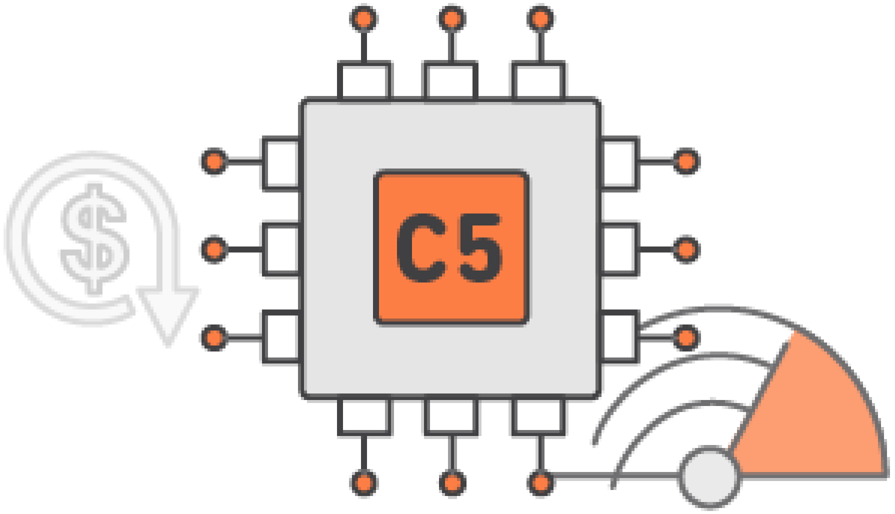

www.STACresearch.com/INTC221006a

* SUT ID INTC190401

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-A2 / AWS c6i.metal / 2 x Intel® Xeon® Platinum 8375C

- Also, cloud server with no hypervisor but with new generation CPU (Ice Lake)
- STAC-A2 Pack for Intel® oneAPI (Rev N)
- Stack:
  - Intel® oneAPI Base Toolkit 2022.3
  - Intel® oneAPI HPC Toolkit 2022.3
  - c6i.metal Amazon Web Services Instance
    - 2 x Intel® Xeon® Platinum 8375C CPU @ 2.90GHz
    - 128 logical cores
    - 256 GiB of DRAM
  - Amazon Linux release 2 (Karoo)

www.STACresearch.com/INTC221006b

STAC® SECURITIES TECHNOLOGY ANALYSIS CENTER

Vs. all publicly reported cloud-based solutions

- The highest options per dollar over a 1 hour burst and a 3-day period
  - STAC-A2.β2.HPORTFOLIO.PRICE_PERF.BURST
  - STAC-A2.β2.HPORTFOLIO.PRICE_PERF.PERIODIC
- The highest options per dollar (reflecting reserve instance pricing discounts) over a 1-year period
  - STAC-A2.β2.HPORTFOLIO.PRICE_PERF.CONTINUOUS

www.STACresearch.com/INTC221006b

Vs. on-prem solution with 2 x Cascade Lake CPUs*

- Had 1.5x the throughput
  - STAC-A2.β2.HPORTFOLIO.SPEED
- In the baseline problem size
  - Was 1.3x the speed in cold runs STAC-A2.β2.GREEKS.TIME.COLD
  - Was 1.4x the speed in warm runs STAC-A2.β2.GREEKS.TIME.WARM
- In the large problem size
  - Was 1.8x the speed in cold runs STAC-A2.β2.GREEKS.10-100k-1260TIME.COLD
  - Was 1.4x the speed in warm runs STAC-A2.β2.GREEKS.10-100k-1260TIME.WARM
- Handled 2 x the paths in the max paths test
  - (STAC-A2.β2.GREEKS.MAX_ASSETS)

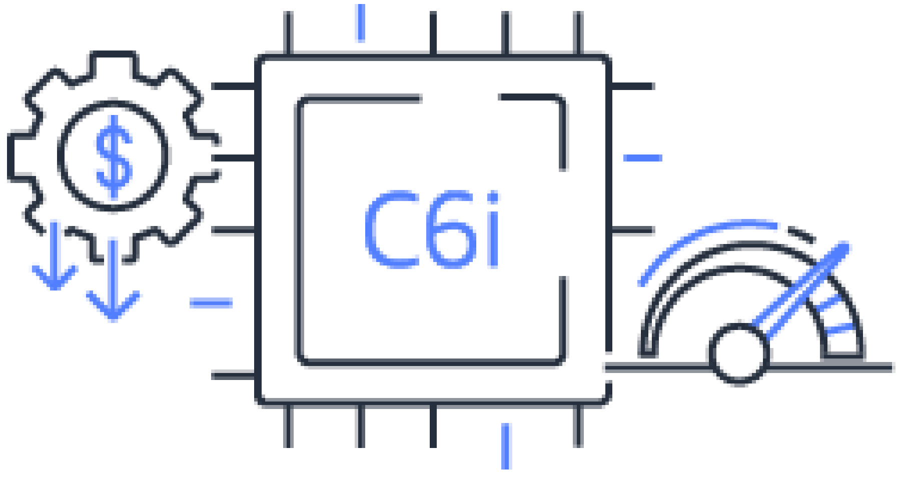

www.STACresearch.com/INTC221006b

* SUT ID INTC190402

STAC-M3

- Performance benchmarks for enterprise tick analytics
  - Language/DBMS neutral
  - Developed by banks and hedge funds
- Workload:
  - Synthetic data modeled on NYSE TAQ
  - Simulates concurrent access with varying number of users
  - Mix of I/O- and compute-intensive operations
- Many years of comparison points on diverse architectures

www.STACresearch.com/m3

STAC® SECURITIES TECHNOLOGY ANALYSIS CENTER

Multiple Suites

Suite	Dataset (reference size*)	Purpose	Impediments to caching/ pre-loading	Storage I/O	Network I/O**	Compute burden	Concurrent users
Baseline (Antuco)	Historical (~4TB)	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.	Yes	Mostly high intensity reads	Neglible	Low to moderate	Varies
Small db in-memory (Shasta)	Historical (~4TB)	Study 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.)	No	Mostly high intensity reads	Neglible	Low to moderate	Varies
Scale (Kanaga)	Historical (theoretically unlimited TB)	Study a few operations on large datasets with large numbers of concurrent requests.	No	Mostly high intensity reads	Neglible	Low to moderate	Theoretically unlimited
Streaming (Jalua)	Streaming ingest & historical (~400GB)	Study ingest capacity, how long it takes ingested data to be available for querying, and query response times on both live and historical data.	*tbd*	Potentially high intensity writes, Low intensity random reads	Potentially high	Low to moderate	Theoretically unlimited

* 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.

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

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC Packs

- Wide range of implementations
  - Databases: kdb+, shakti, eXtremeDB
  - Clustered file systems, parallel file systems, NFS, flash arrays, NVME over Fabric, direct-attached SSD, NAND and post-NAND Flash (e.g. Optane)
  - Single database server, database cluster (bare metal and cloud)
- Analytics STAC Track subscribers can access STAC Pack source code
  - Understand how to develop for a given database
  - Run tests: Mark your own stacks to market
  - Discover code optimizations

council@*STACresearch*.com

STAC-M3 / kdb+ / 2 x DDN AI400X2 All-Flash / 16x DB servers

- Ran baseline (Antuco) and scale (Kanaga) benchmarks
- Demonstrates scaling with DDN AI400X2 and EXAScaler software
- STAC-M3 Pack for kdb+: Compatibility Rev H

www.STACresearch.com/KDB221014

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3 / kdb+ / 2 x DDN AI400X2 All-Flash / 16x DB servers

- Stack:
  - kdb+ 4.0 Cloud Edition in distributed mode
  - 2 x DDN AI400X2 All-Flash appliances, each with:
    - DDN EXAScaler Parallel Filesystem version 6.1.0
    - 24 x 3.8TB NVMe SSD
  - 16 x GIGABYTE H262-Z62, each with:
    - DDN EXAScaler 6.1.0 software
    - Centos 8.3.2011
    - 2 x AMD EPYC 7763 64-core CPUs @ 2.45 GHz
    - 512GiB memory

www.STACresearch.com/KDB221014

STAC SECURITIES TECHNOLOGY ANALYSIS CENTER

Compared to all publicly disclosed mean-response time results

- Outperformed in 4 of 7 100-user benchmarks:
  - 100-user intervalized statistics (STAC-M3.β1.100T.STATS-UI.TIME)
  - 1-, 2-, and 3-year 100-user 12-day VWAB (STAC-M3.β1.100T.YR{1,2,3}VWAB-12D-HO.TIME)

www.STACresearch.com/KDB221014

* SUT ID KDB190430

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

Vs. stack w/ network-attached flash storage and 9 DB servers*

- Faster in 13 of 17 Antuco mean-response time benchmarks, including:
  - 6x speedup in 50-user intervalized stats (STAC-M3.β1.50T.STATS-UI.TIME)
  - 5x speedup in 10-user aggregate stats (STAC-M3.β1.10T.STATS-AGG.TIME)
  - 4.9x speedup in single-user intervalized stats (STAC-M3.β1.1T.STATS-UI.TIME)
- Faster in 21 of 24 Kanaga mean-response time benchmarks, including:
  - 2.1 – 4.4x speedup in single-user high-bid (STAC-M3.β1.1T.{2,3,4,5}YRHIBID.TIME)

www.STACresearch.com/KDB221014

* SUT ID KDB220506

Vs. public cloud stack with 15 DB & 40 storage servers*

- Faster in 13 of 17 Antuco mean-response time benchmarks, including:
  - 6.5x speedup in 10-user volume curve (STAC-M3.β1.10T.VOLCURV.TIME)
- Faster n 16 of 24 Kanaga mean-response time benchmarks, including:
  - 15.5 – 19.3x speedup in 10-user market snapshots (STAC-M3.β1.10T.YR{2,3,4,5}-MKTSNAP.TIME)
- Comparison SUT used previous version of the STAC Packs.

www.STACresearch.com/KDB221014

* SUT ID KDB210507