STAC Update:
Time series stacks

Peter Nabicht
President, STAC

peter.nabicht@STACresearch.com

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

www.STACresearch.com/m3

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3

- Wide range of implementations
  - Databases: kdb+, 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)

www.STACresearch.com/**m3**

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3/kdb+/Lenovo SR860 V2/Intel 8380HL & Optane Pmem v2

- First SUT with second-gen PMEM and Cooper Lake
- Also wanted to compare two different namespaces with Storage Over App Direct
  - Public SUT used Sector namespaces
  - Report in STAC Vault used Filesystem-DAX
- Ran baseline (Antuco) benchmarks and 2-year scaling (Kanaga) benchmarks

www.STACresearch.com/KDB201109

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3/kdb+/Lenovo SR860 V2/Intel 8380HL & Optane Pmem v2

- Stack for public report:
  - Kx Systems kdb+ v4.0
  - SUSE 15.2 with XFS
  - Lenovo ThinkSystem SR860 V2 server
  - 4 x 28-core Intel Xeon Platinum 8380HL (Cooper Lake) CPU @ 2.90GHz
  - 12 TiB (24 x 512GiB) Intel Persistent Memory 200 Series
  - Storage Over App Direct Mode with Sector mode namespaces
  - 6 x 8TB Intel DC P4600 SSD

www.STACresearch.com/KDB201109

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Versus all publicly disclosed mean response-time results

- Fastest in 4 of 17 baseline (Antuco) benchmarks
- Fastest single-server solution in 8 of 17 Antuco benchmarks

www.STACresearch.com/KDB201109

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Versus previous generation CPU, Optane Memory, and kdb+ 3.6

Compared to SUT ID KDB190322b:

- Was faster in 14 of 17 Antuco mean-response time benchmarks, including:
  - 6.7x speed-up in 10-user aggregate stats (STAC-M3.β1.10T.STATS-AGG.TIME)
  - 4.4x speed-up in 10-user market snapshot (STAC-M3.β1.10T.MKTSNAP.TIME)
- Was faster in 8 of 10 Kanaga (2-year) mean-response time benchmarks, including:
  - 5.0x speed-up for year-1 market snapshot (STAC-M3.β1.TRUNC.10T.YR2-MKTSNAP.TIME)
  - 4.9x speed-up for year-2 market snapshot (STAC-M3.β1.TRUNC.10T.YR1-MKTSNAP.TIME)

www.STACresearch.com/KDB201109

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Versus same SUT with Filesystem-DAX Namespaces

- Companion report in the confidential STAC Vault
- Stack:
  - kdb+ v4.0
  - SUSE 15.2 with XFS
  - Lenovo ThinkSystem SR860 V2 server
  - 4 x 28-core Intel Xeon Platinum 8380HL CPU @ 2.90GHz
  - 24 x 512GiB Intel Persistent Memory 200 Series (Storage Over App Direct Mode with Filesystem-DAX namespaces)
  - 1 x 8TB Intel DC P4600 SSD

www.STACresearch.com/KDB201110

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3 / kdb+ 4.0 / Lenovo SR650 V2 / Intel Ice Lake

- First STAC-M3 SUT with Ice Lake processors
- Tested with STAC-M3 baseline (Antuco) suite
- Stack:
  - kdb+ 4.0
  - Lenovo ThinkSystem SR650 V2 server
  - 2 x Intel® Xeon® Platinum 8360Y (Ice Lake) @ 2.40GHz
  - 1 TiB DRAM
  - 4 TiB Intel PMem 200 (16 x 256GiB DDR4-3200 DIMMs)
  - Storage Over AppDirect mode

STAC Report coming soon

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Compared to other public mean response-time Antuco results

- Outperformed all publicly disclosed results in 4 of 17 benchmarks
- Outperformed all publicly disclosed results for solutions involving kdb+ on a single, 2-socket server in 14 of 17 benchmarks
- Compared to a similar system with the previous generations of CPU, Optane memory and kdb+ 3.6 (SUT ID KDB190320b):
  - Faster in 14 out of 17 benchmarks, including:
  - 6.1x speed-up in 10-user market snapshot (STAC-M3.β1.10T.MKTSNAP.TIME)

STAC Report coming soon

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3 / kdb+ 4.0 / 2 x Lenovo SR860 / Intel Cooper Lake

- Stack:
  - kdb+ 4.0
  - New STAC Pack (benchmark implementation code): Compatibility Rev F
  - 2 x Lenovo ThinkSystem SR860 V2 each with:
    - 4 x Intel® Xeon® Platinum 8380H @ 2.90GHz
    - 3 TiB DRAM
    - SUSE EL 15 SP2
    - 24TiB Intel PMem 200 (24 x 512GiB DIMMs)
    - Storage Over AppDirect with Filesystem-DAX namespaces

STAC Report coming soon

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Results comparisons

- Outperformed all publicly disclosed results in 10 of 17 Antuco benchmarks
- Outperformed all publicly disclosed results in 9 of 10 Kanaga (2 year) benchmarks

STAC Report coming soon

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3 / kdb+ 4.0 / WekaFS / AWS

- First STAC-M3 results with Weka on a public cloud
- Stack:
  - kdb+ 4.0 (distributed mode)
  - WekaIO WekaFS Storage Software Release v3.10.1
  - Database Servers: 15 x AWS EC2 C5n.9xlarge
  - Storage Servers: 40 x AWS EC2 i3en.6xlarge (426 TiB usable storage)
  - AWS Ethernet ENI – Elastic Network Interface

STAC Report coming soon

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3 / kdb+ 4.0 / WekaFS / AWS

- Outperformed all publicly disclosed results in 3 of 24 Kanaga benchmarks
- Versus a kdb+ 4.0 solution on a 10-node cluster with 60TB of PMEM (KDB200603):
  - Faster in 16 of 24 Kanaga and 9 of 17 Antuco
- Versus a kdb+ 3.6 solution on 15 database servers with a parallel file system and all-flash appliances (KDB200915):
  - Faster in 20 of 24 Kanaga and 4 of 17 Antuco

STAC Report coming soon

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3 / kdb+ 4.0 / WekaFS / AWS

- Versus a kdb+ 3.6 solution involving a previous version of Weka with 14 on-prem database nodes and 18 storage nodes (KDB200401):
  - Faster in 3 of 24 Kanaga and 4 of 17 Antuco

STAC Report coming soon

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Streaming Update

STAC-M3 Jalua

- Streaming data ingest benchmarks for the market data use case
- Questions it answers
  - What kind of ingest rates can it handle without loss?
  - How well does it blend live and historical data?
  - How fast is data available for querying?
  - How do ingest and query load impact each other?

www.STACresearch.com/m3

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3 Jalua status

- Working group signed off on the specifications in Q4 of 2021
  - STAC-M3 Benchmark Specs - Jalua Rev B (streaming)
- STAC-M3 Jalua Test Harness is available
- Full reference implementation available
  - Built by STAC on InfluxDB
  - Naïve implementation
- STAC is currently working on vault reports as a “naïve user”
  - InfluxDB
  - TimescaleDB
- Not too late to get involved

www.STACresearch.com/m3

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Streaming IT Telemetry

- Some user firms on the STAC Benchmark Council asked for & provided input to proposed IT Telemetry streaming data benchmarks
- STAC currently meeting with users and gathering requirements for a proposal
- Working group meetings will be held this summer
- Want more voices in the discussion
- You can participate by joining the Streaming Analytics SIG

www.STACresearch.com/streaming

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER