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
- Many years of comparison points on diverse architectures

www.STACresearch.com/m3

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Advances in STAC Packs

- kdb+ STAC Pack(s): Compatibility Rev H
  - Used in audits this fall
  - Threading improvements take advantage of kdb+ 4.0 “thread-aware” features
  - Optimized some of the operations to benefit from high thread count parallelism
- Shakti STAC Pack coming soon
- Analytics STAC Track subscribers can access STAC Pack source code
  - Understand how to develop for a given database
  - Discover code optimizations
  - Run tests: Mark your own stacks to market

council@STACresearch.com

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3/kdb+/Dell EMC PowerScale F900 All-Flash NAS, 3-node cluster

- Stack involved the densest Dell EMC PowerScale storage to feature all-NVMe media
- Ran baseline (Antuco) benchmarks
- STAC-M3 Pack for kdb+: Compatibility Rev H

www.STACresearch.com/KDB210929

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3/kdb+/Dell EMC PowerScale F900 All-Flash NAS, 3-node cluster

- Stack:
  - kdb+ v4.0
  - CentOS 7.8 with NFS v3
  - 9 x Dell EMC PowerEdge R640 servers, each with:
    - 2 x Intel® Xeon® Platinum 8360Y (Ice Lake) @ 2.40GHz
    - 384GiB DRAM
  - 3 x Dell EMC PowerScale F900 All-Flash Scale-Out NAS
    - 251TiB total physical capacity
    - Dell EMC OneFS 9.2.0.0 storage cluster operating system

www.STACresearch.com/KDB210929

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

Versus previous generation of Dell EMC flash storage appliance*

- Faster in 14 of 17 mean-response time benchmarks, including:
  - 5.3x speed-up in the 10-user market snapshot (STAC-M3.β1.10T.MKTSNAP.TIME)
  - 4.9x speed-up in the 10-user volume curve (STAC-M3.β1.10T.VOLCURV.TIME)
  - 2.9x speed-up in the 100-user 12-day VWAB (STAC-M3.v1.100T.VWAB-12D-NO.TIME)
  - 2.7x speed-up in the 100-user unpredictable interval stats (STAC-M3.ß1.100T.STATS-UI.TIME)
- Comparison SUT used kdb+ 3.6 and STAC Pack Compatibility Rev E

www.STACresearch.com/KDB210929

* SUT ID KDB190430

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

Versus parallel filesystem w/ 14 DB servers & 18 storage servers*

- Faster in 3 of 17 mean-response time benchmarks:
  - 9.8x speed-up in the single-user NBBO (STAC-M3.β1.1T.NBBO.TIME)
  - 3.6x speed-up in the 10-user volume curve (STAC-M3.β1.10T.VOLCURV.TIME)
  - 17% faster in the single-user VWAB (STAC-M3.v1.1T.VWAB-D.TIME)
- Comparison SUT used kdb+ 3.6 and STAC Pack Compatibility Rev E

www.STACresearch.com/KDB210929

* SUT ID KDB200401

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3/kdb+ 4.0/10-nodes/Intel 6240L/60 TiB Optane PMem

- First STAC-M3 SUT with a 5-year data set fully in PMEM
- Tested with STAC-M3 baseline (Antuco) and 5-year scaling (Kanaga) suites
- STAC-M3 Packs for kdb+: Compatibility Rev H
- Leveraged kdb+ sharded mode to distribute data in PMem across the 10 nodes

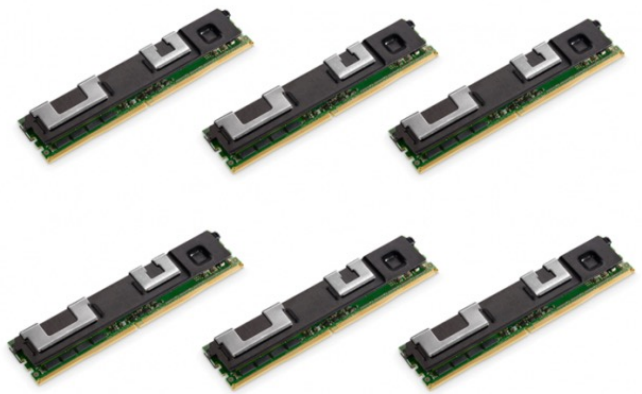

www.STACresearch.com/KDB211006

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3/kdb+ 4.0/10-nodes/Intel 6240L/60 TiB Optane PMem

- Stack:
  - kdb+ 4.0, running in sharded mode
  - 10 x Dell PowerEdge R640 servers, each with:
    - 2 x 18-core Intel® Xeon® Gold 6240L @ 2.6GHZ
    - 384GiB DRAM: 12 x 32GiB DDR4-3200 DIMMs @ 2666 MT/s
    - 6TiB of Intel® Optane™ Persistent Memory 100 Series (Storage over App Direct Mode with Filesystem-DAX namespaces)
    - Red Hat Enterprise Linux release 8.4
  - 2 x Arista 7170-32CD 100GbE switches

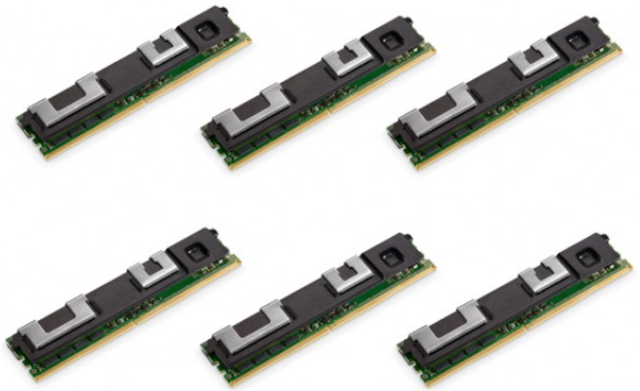

www.STACresearch.com/KDB211006

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

Compared to all publicly disclosed results

- Outperformed in 16 of 24 STAC-M3 Kanaga benchmarks
- 7 of the 16 had speed-ups exceeding 2.2x compared to the previous best result, including:
  - STAC-M3.β1.1T.YR{1,2,3,4,5}VWAB-12D-HO.TIME
  - STAC-M3.β1.50T.YR{1,2}VWAB-12D-HO.TIME

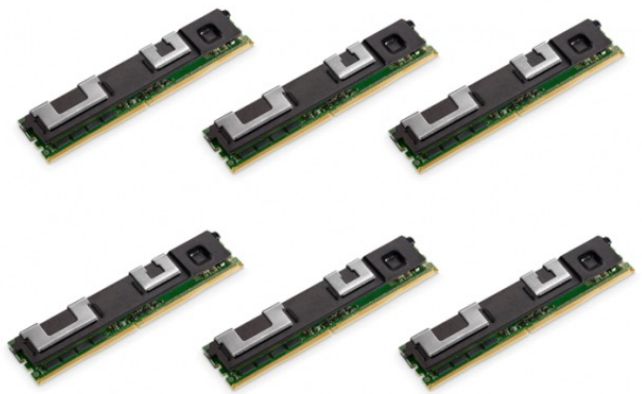

www.STACresearch.com/KDB211006

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Compared to a solution with 32 servers*

- Comparison SUT: 14 DB servers, 18 storage servers (NAND), parallel filesystem, kdb+ 3.6, STAC Pack Compatibility Rev E
- The PMEM-based 10-server solution was faster in 20 of 24 Kanaga benchmarks, including:
  - 57x to 60x speed-up in the four 10-user market snapshot benchmarks (STAC-M3.β1.10T.YR{2,3,4,5}-MKTSNAP.TIME)
  - 3.9x to 7.0x speed-up in the 5 single-user 12-day VWAB benchmarks (STAC-M3.β1.1T.YR{1,2,3,4,5}VWAB-12D-HO.TIME)
  - 4.8x speed-up in the 50-user year-one 12-day benchmark (STAC-M3.β1.50T.YR1VWAB-12D-HO.TIME)

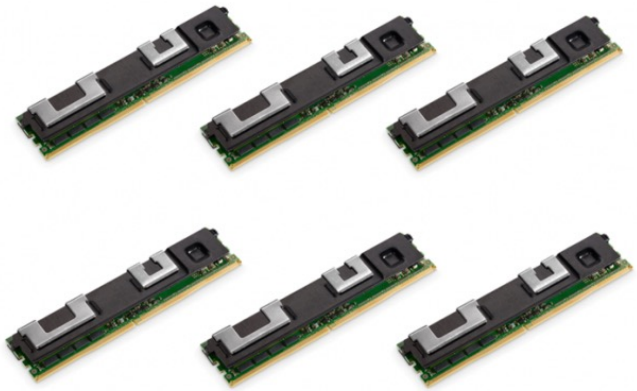

www.STACresearch.com/KDB211006

** SUT ID KDB200401*

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3 / kdb+ 4.0 / DDN AI400X2 / 15 DB server nodes

- First STAC-M3 audit of DDN's newest all-flash storage appliance: AI400X2
- Newest version of EXAScaler: 6.0.0
- Tested with STAC-M3 baseline (Antuco) and 5-year scaling (Kanaga) suites
- STAC-M3 Packs for kdb+: Compatibility Rev H

STAC Report coming soon

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3 / kdb+ 4.0 / DDN AI400X2 / 15 DB server nodes

- Stack:
  - kdb+ v4.0
  - 15 x Intel Server Board S2600BPB in, each with:
    - Intel® Xeon® Gold 6138 Single CPU
    - 96GiB DDR4 DIMMS @ 2666MT/s
    - Mellanox InfiniBand ConnectX-6 dual-port Adapter
  - Storage:
    - 1 x DDN AI400X2 All-Flash storage appliance
    - 24 x 3.8TB NVMe SSD
    - DDN EXAScaler Parallel Filesystem version 6.0.0
    - 84TiB total physical capacity (62TiB usable)
  - Mellanox QM8700 40-port Non-blocking HDR 100GbE InfiniBand Smart Switch

STAC Report coming soon

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Versus 2 previous gen DDN appliances and 15 DB servers*

- Faster in 10 of 17 Antuco benchmarks, including:
  - 3.1x speedup in 10-user volume curve (STAC-M3.β1.10T.VOLCURV.TIME)
  - 2.9x speedup in 10-user market snapshot (STAC-M3.β1.10T.MKTSNAP.TIME)
  - 2.1x speedup in 10-user theoretical P&L (STAC-M3.β1.10T.THEOPL.TIME)
- Faster in 19 of 24 Kanaga benchmarks, including:
  - 2.1x - 2.3x speedup in the four 10-user market snapshots (STAC-M3.β1.10T.YR{2,3,4,5}-MKTSNAP.TIME)
- Comparison SUT used kdb+ 3.6 and STAC Pack Compatibility Rev E

STAC Report coming soon

* SUT ID KDB200915

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

Vs. parallel FS with 15 DB & 40 storage servers on public cloud*

- Faster in 12 of 17 Antuco benchmarks, including:
  - 4.4x speedup in 10-user volume curve (STAC-M3.β1.10T.VOLCURV.TIME)
  - 4.0x speedup in 10-user market snapshot (STAC-M3.β1.10T.MKTSNAP.TIME)
  - 2.6x speedup in 10-user theoretical P&L (STAC-M3.β1.10T.THEOPL.TIME)

- Faster in 12 of 24 Kanaga benchmarks, including:
  - 16.2x to 19.9x speedup in the four 10-user market snapshots (STAC-M3.β1.10T.YR{2,3,4,5}-MKTSNAP.TIME)
- Comparison SUT used STAC Pack Compatibility Rev E

STAC Report coming soon

* SUT ID KDB210507

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

Versus network-attached flash storage & 9 DB servers*

- Faster in 11 of 17 Antuco benchmarks, including:
  - 11.4x speedup in single-user intervalized stats (STAC-M3.β1.1T.STATS-UI.TIME)
  - 6.6x speedup in 10-user market snapshot (STAC-M3.β1.10T.MKTSNAP.TIME)
  - 3.5x to 5.9x speedup in the three multi-user intervalized stats (STAC-M3.β1.{10,50,100}T.STATS-UI.TIME)
  - 4.2x speedup in 10-user aggregate stats (STAC-M3.β1.10T.STATS-AGG.TIME)
- Comparison SUT used same kdb+ version and STAC Pack Compatibility Rev

STAC Report coming soon

* SUT ID KDB210929

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER