STAC Update:
Big data benchmarks

Peter Lankford
Founder and Director, STAC

peter.lankford@STACresearch.com

STAC-M3

- Performance benchmarks for enterprise tick analytics
  - Language/DBMS neutral
  - Developed by banks
- Workload:
  - Synthetic data modeled on NYSE TAQ
  - Mix of I/O- and compute-intensive operations (read-heavy)
  - Scalable volume and number of users

www.STACresearch.com/m3

Copyright © 2017 Securities Technology Analysis Center LLC

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3 / kdb+ / Lenovo SR950 / 4 x Skylake

- SUT ID: KDB170629
- Stack:
  - Software: kdb+ 3.5 / RHEL 7.3 / ext4
  - Processor: 4 x Intel Platinum 8180 (Skylake)
  - Server: 1 x Lenovo ThinkSystem SR950
  - Storage: 6 x 1.6TB Intel SSD U.2 P4600
- Shasta suite (baseline database in-memory)

www.STACresearch.com/lenovo

Copyright © 2017 Securities Technology Analysis Center LLC

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Highlights

- Compared to all previous publicly released STAC-M3 Shasta results, this solution using four Intel “Skylake” processors:
  - Set records in 9 of the 15 required mean response-time benchmarks
  - Set records in both of the optional mean response-time benchmarks (NBBO and WRITE)
  - Over 2x the next highest score in the year-high bid benchmark (STAC-M3.β1.1T.YRHIBID/s.TIME)

Copyright © 2017 Securities Technology Analysis Center LLC

STAC-M3 / kdb+ / Lenovo SR650 / 2 x Skylake

- SUT ID: KDB170703
- Stack:
  - Software: kdb+ 3.5 / RHEL 7.3 / ext4
  - Processor: 2 x Intel Platinum 8180 (Skylake)
  - Server: 1 x Lenovo ThinkSystem SR650
  - Storage: 4 x Intel SSD DC P4600 + 4 x Intel SSD DC P4500
- Baseline tests (Antuco)

www.STACresearch.com/lenovo

Copyright © 2017 Securities Technology Analysis Center LLC

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Highlights

- Compared to previous publicly reported results for single-node 2-socket servers running kdb+
  - This solution set records in 15 of the 17 mean response-time benchmarks.

Copyright © 2017 Securities Technology Analysis Center LLC

STAC-M3 / kdb+ / Lenovo SR650 / 2 x Skylake / Intel Optane

- SUT ID: KDB171010
- Stack:
  - Software: kdb+ 3.5 / RHEL 7.3 / ext4
  - Processor: Intel Platinum 8180 (Skylake)
  - Server: 1 x Lenovo ThinkSystem SR650
  - Storage: 8 x Intel SSD DC P4800X (375 GB ea)
- Baseline tests (Antuco)
- Report will be published soon

www.STACresearch.com/intel

Copyright © 2017 Securities Technology Analysis Center LLC

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Highlights

- Compared to previous publicly reported results for single-node 2-socket servers running kdb+:
  - This solution set records in 11 of the 17 mean response-time benchmarks
- Compared to the nearly identical 2-socket system above with previous-gen (NAND flash) drives:
  - 7.5x in the market snap benchmark (STAC-M3.β1.10T.MKTSNAP.TIME)
  - 3.0x in the week-high bid benchmark (STAC-M3.β1.1T.WKHIBID.TIME)
  - 2.6x in the theoretical P&L benchmark (STAC-M3.β1.10T.THEOPL.TIME)
  - 2.5x in the multi-user VWAB benchmark (STAC-M3.β1.100T.VWAB-12D-NO.TIME)

Copyright © 2017 Securities Technology Analysis Center LLC

STAC-M3 / X-IO Axellio

- X-IO Technologies submitted results to the STAC Vault
  - www.STACresearch.com/KDB170916
- SUT ID: KDB170916
- Stack:
  - Software: kdb+ 3.5 / CentOS 7.3 / ext4
  - 1 x X-IO Axellio Edge Computing System containing 2 x nodes
  - Each node: 2 x 8-core Intel Xeon E5-2699 CPUs
- Both baseline and scale tests

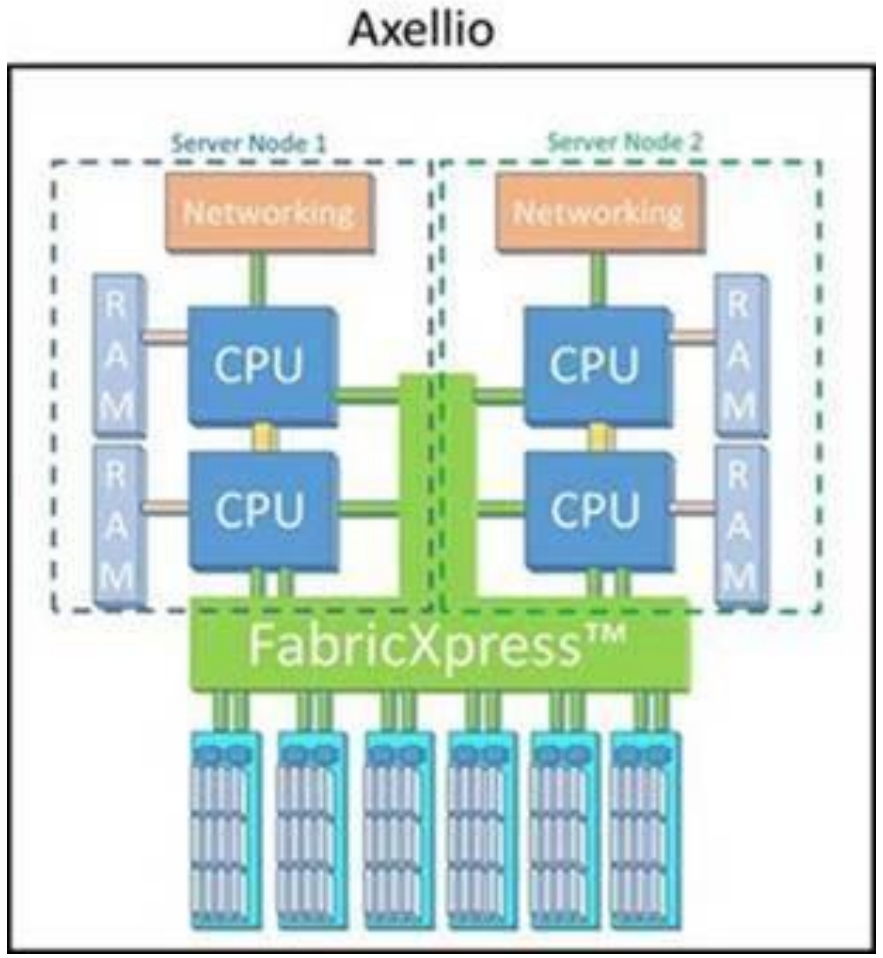

Copyright © 2017 Securities Technology Analysis Center LLC

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3 / X-IO Axellio

- STAC audited an enhanced configuration
- SUT ID: KDB171020
- Stack:
  - Software: kdb+ 3.5 / CentOS 7.3 / ext4
  - 1 x X-IO Axellio Edge Computing System containing 2 x nodes
  - Each node: 2 x 22-core Intel Xeon E5-2699 CPUs
- Both baseline and scale tests

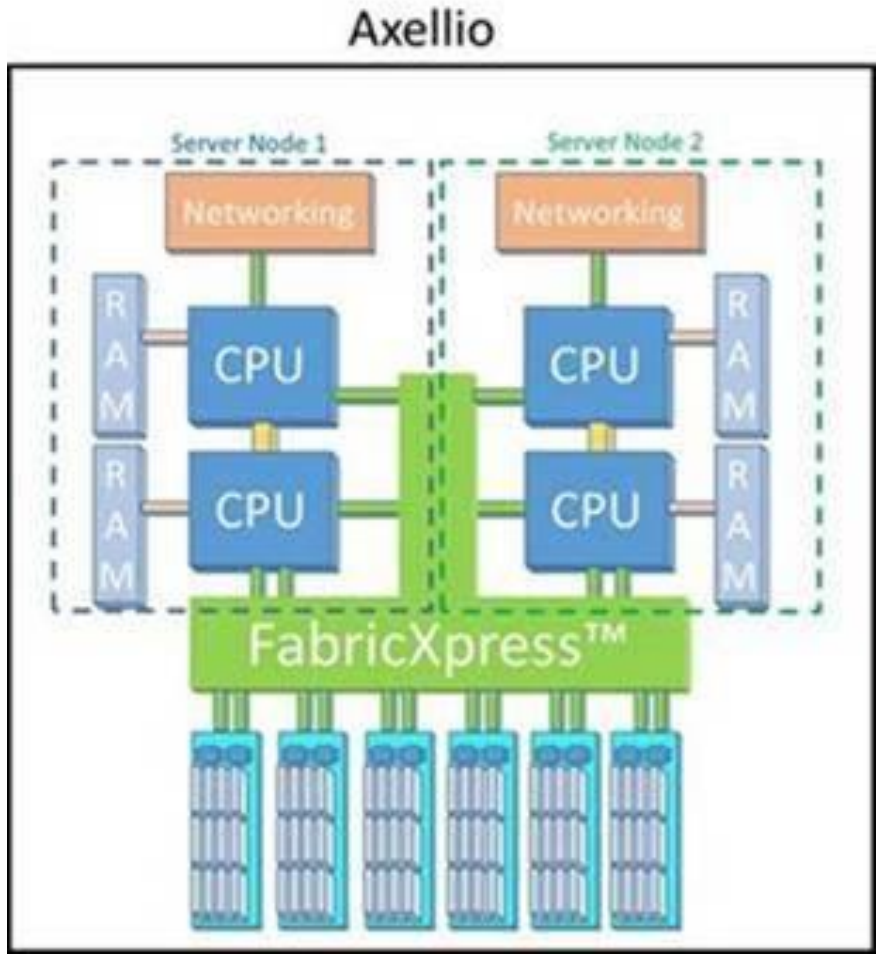

www.STACresearch.com/XIO

Copyright © 2017 Securities Technology Analysis Center LLC

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

Highlights

The solution based on this Axellio Edge Computing System:

- Set 5 new mean-response time records in the baseline benchmarks and 5 in the scaling benchmarks
- Set mean response-time records in 11 of 17 baseline benchmarks when compared to all publicly reported submissions to date that have elected to undergo both baseline and scale tests (Kanaga)

Copyright © 2017 Securities Technology Analysis Center LLC

Highlights

- But X-IO pointed something out
- The STAC-M3 benchmark suites are an all-around competition
- Thus, it makes sense to compare all benchmarks of a new submission 1:1 with those of previous submissions
- When you do that…




Copyright © 2017 Securities Technology Analysis Center LLC

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Highlights

The solution based on this Axellio Edge Computing System:

- Beat each of the other solutions in more than 50% of the benchmarks
- Over 70% of the benchmarks, if you exclude the Optane based solution mentioned earlier (SUT ID KDB171010)

Copyright © 2017 Securities Technology Analysis Center LLC

STAC-A3

- Recap:
  - Workloads that emulate real-world backtesting jobs
    - Range of parallelism and IO/compute intensity
  - Measure speed, scalability, efficiency of any architecture
- One of the two algorithms is currently auditable
  - Parameter sweep of a mean-reversion algorithm
  - Test harness hands the implementation jobs to execute
  - Measures the throughput and efficiency of the SUT

Copyright © 2017 Securities Technology Analysis Center LLC

STAC-A3 in Python using HPAT on Google Cloud Platform

- HPAT:
  - High Performance Analytic Toolkit
  - Compiles Python for clusters and cloud
  - Open source, contributed by Intel Labs
  - Ehsan will describe later today
- Stacks under test:
  - STAC-A3 Pack for Python
  - GCP Skylake instances with 64 vCPU and 270GB memory
  - 5 nodes (HPAT171028) and 20 nodes (HPAT171029)
  - Google Persistent SSD
  - HDF5 format
- Report coming soon

Copyright © 2017 Securities Technology Analysis Center LLC

Highlights

- Compared to previously reported system with I/O-accelerated Spark (Scala) on 5 Broadwell-based GCP nodes, the 5-node GCP/HPAT (Python) solution had:
  - 2.4x the throughput at 50 instruments/50 simulations
  - 33% higher throughput in STAC-A3.β1.SWEEP.SPEED1
- Moving from 5 to 20 nodes did not require any data rebalancing
  - Google Persistent SSD is independent of compute nodes
- Throughput of 20-node system in the workload scaling test was 3.2x to 3.9x that of the 5-node system

Copyright © 2017 Securities Technology Analysis Center LLC