STAC Update:
Big data benchmarks

Peter Lankford
Founder and Director, STAC

peter.lankford@STACresearch.com

STAC-M3

- Performance benchmarks for 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

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

eXtremeDB 7 / IBM Power8 / IBM Flashsystem

- SUT:
  - Software: eXtremeDB Financial Edition 7.0 with RHEL 7.4, ext4
  - Server: IBM S824L with 2 x POWER8 CPU @ 3.52 GHz
  - Storage: IBM Flashsystem 900
- Baseline tests (Antuco)
- Scale tests (Kanaga)

www.STACresearch.com/mcobject

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Highlights – baseline suite (Antuco)

- Compared to all publicly benchmarked solutions:
- Fastest mean response times in 4 of the 17 response-time benchmarks
- Most consistent results (lowest stdv) in 4 of the 17 benchmarks
- 1.93x previous best 10T.THEOPL.TIME
- 1.30x previous best 1T.STATS-UI.TIME
- Compared to all publicly benchmarked **2-socket** solutions:
  - Fastest results in 4 of the 17 benchmarks
  - 12 of the remaining records still held by previous eXtremeDB/Power8 systems

Highlights – volume & user scaling (Kanaga)

- For all combinations of query type, data volume, and concurrent users:
  - Fastest mean response times ever reported
  - Most consistent response times (lowest standard deviation) ever reported
- Each mean response time was 5.5x to 212x the previous best result, including:
  - 21x to 212x previous best 10T.YR[n]-MKTSNAP.TIME
  - 21x previous best 1T.OLDYRHIBID.TIME
  - 8-10x previous best 100T.YR[n]VWAB-12D-HO.TIME
  - 5-8x previous best 1T.[n]YRHIBID.TIME

kdb+ 3.3 / Broadwell EX / Lenovo X6 / 6TB RAM

- SUT:
  - Software:
    - Kx Systems kdb+ 3.3
    - RHEL 7.2
  - Server:
    - Lenovo System x3850 X6
    - 4 x 24-core Intel Xeon E7-8890 v4 (“Broadwell EX”) @ 2.20 GHz
    - 6TB SK hynix DDR4 SDRAM
  - Storage:
    - 3 x Intel 2TB SSD DC P3700
    - Data pre-loaded into RAM
- Shasta suite
  - Real-world performance for smaller datasets (few TB)
  - No restrictions on pre-loading/caching of data in memory

www.STACresearch.com/kx

STAC® SECURITIES TECHNOLOGY ANALYSIS CENTER

Highlights

- Compared to all publicly benchmarked solutions using the Shasta suite:
  - Fastest mean response times in 13 of 15 required benchmarks
- Compared to a 4-socket Ivy Bridge server with 6TB DRAM (SUT KDB140116):
  - Faster in all 15 required benchmarks
  - 6 of which > 2x faster
  - 10 of which >45% faster
  - 2.96x previous best 10T.THEOPL/s.TIME
  - 2.67x previous best 10T.VOLCURV/s.TIME
  - 64% faster in 10T.STATS-AGG/s.TIME
  - 53% faster in 1T.YRHIBID/s.TIME

STAC-A3

- Recap:
  - Workloads that emulate real-world backtesting jobs
    - Range of parallelism and IO/compute intensity
  - Measure speed, scalability, efficiency of any architecture
- Research agenda:
  - Shared storage (e.g., parallel FS) and shared-nothing (e.g., HDFS)
  - Drive insight & improvements at both software and hardware level
- Status:
  - Test harness almost complete
  - First vendor implementation almost complete
  - First audited results very soon

www.STACresearch.com/**backtesting**

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

First STAC-A3 results

- Stack:
  - Authored by Think Big Analytics
  - Python
  - Hadoop Streaming
  - Cloudera Enterprise 5.4
  - Dell PowerEdge servers
  - Intel Xeon E5-2680 v3 (Haswell) processors
- Validation implementation
- Also a performance baseline
- Not highly optimized
- Still demonstrated good scaling
  - With respect to both simulations and symbols per simulation.

Spark resource managers

- Spark supports many resource managers
  - Plugin architecture
- We tested three:
  - YARN
  - Mesos
  - IBM Platform Conductor for Spark
- Two phases. Completed Phase 1.
- Used a benchmark proposed by IBM
  - Exercises the resource manager
- Consulted YARN/Mesos expert on benchmark and on config of YARN and Mesos

Job-submission pattern

- Same job (2 GB sort)
- Synchronous submission
- Ramp up and down

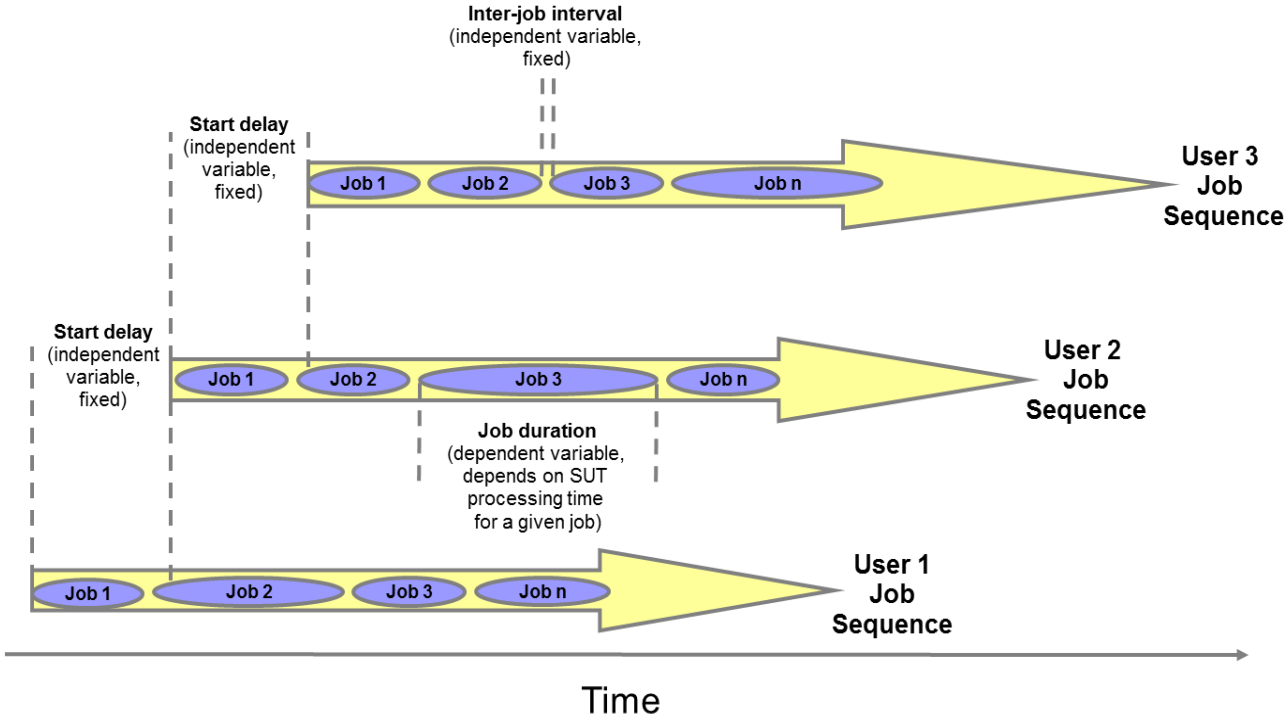

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Results

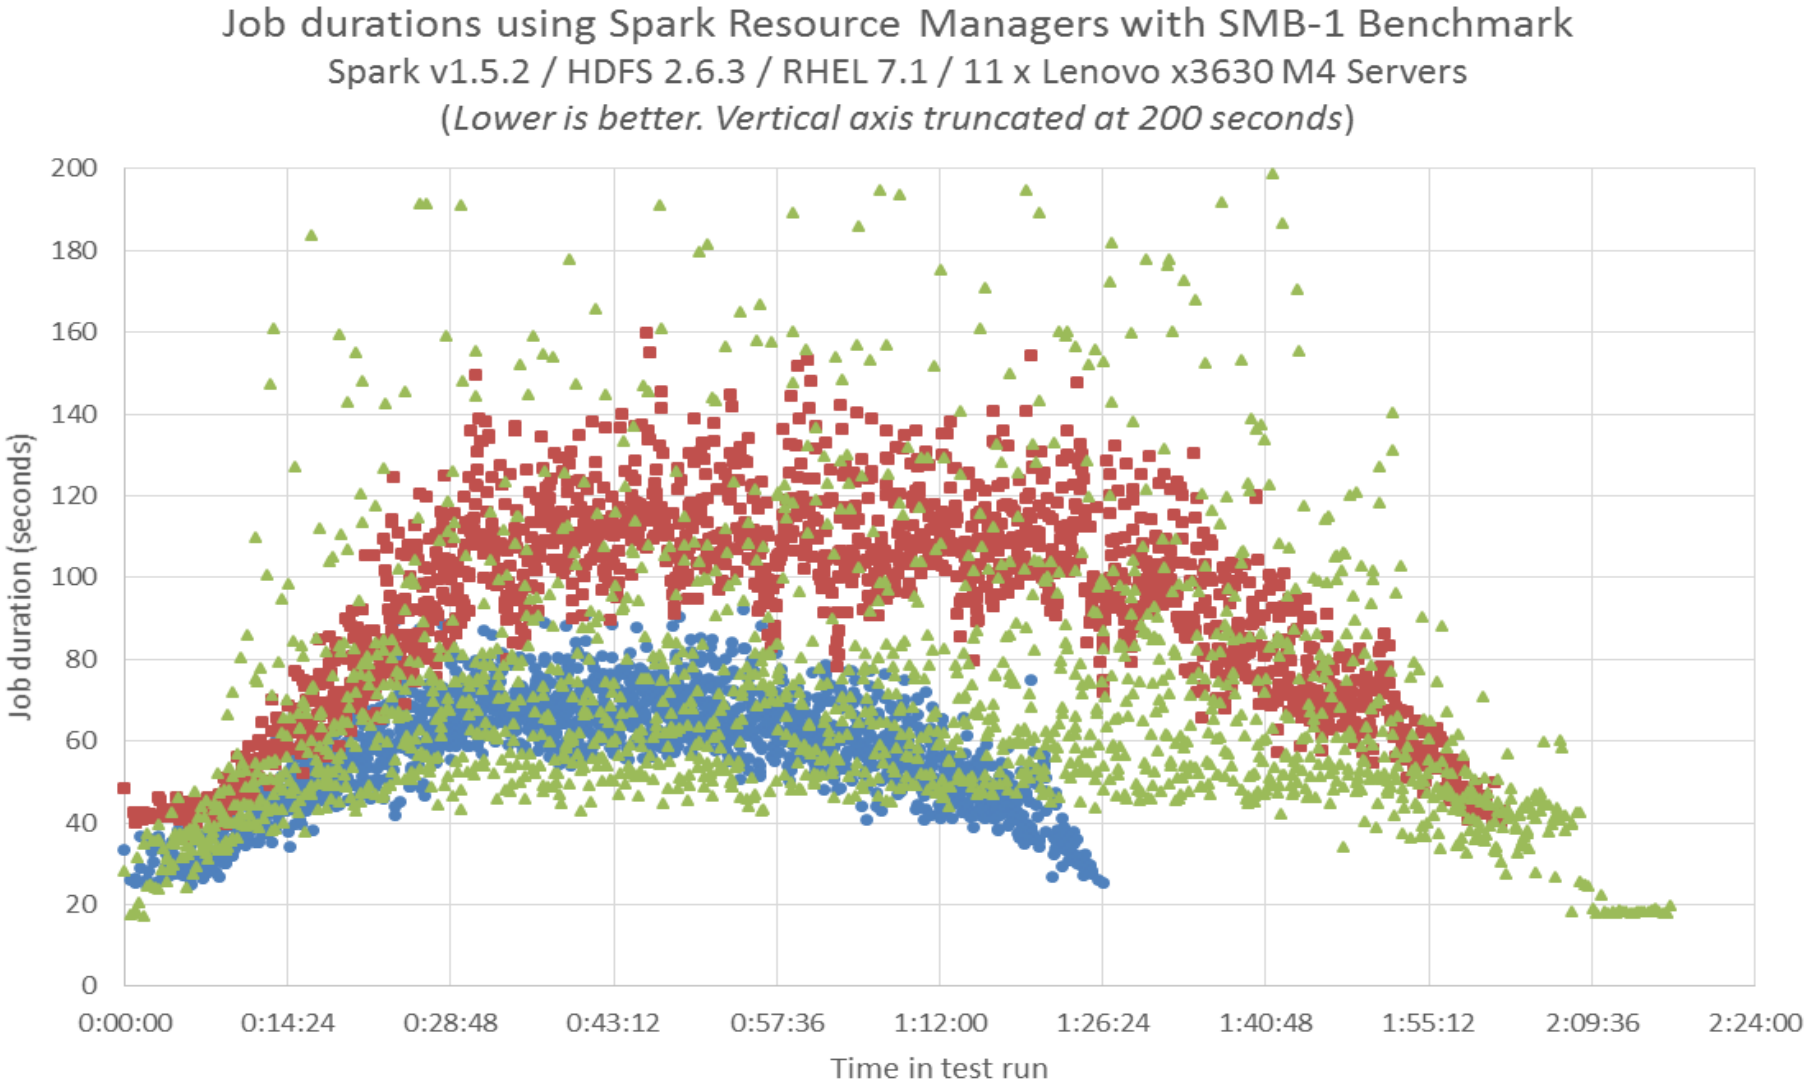

Source: STAC®
www.STACresearch.com

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

SECURITIES TECHNOLOGY ANALYSIS CENTER

Results highlights

- YARN and Mesos had similar throughput
  - YARN was about 12% higher than Mesos
- IPCS throughput was substantially higher:
  - 58% higher than Mesos
  - 41% higher than YARN.
- YARN and Mesos: same average run time per job (94 sec)
- IPCS was 59% faster at 59 seconds
- IPCS and YARN were fairer than Mesos
  - Fairness defined as per-job stdev as percent of mean
  - IPCS: 22%
  - YARN: 27%
  - Mesos: 235%.