



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

Advances in 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 / shakti 2.0 / AMD EPYC 7742 / DDN ES200NVX All-Flash

- First stack involving Shakti Software's shakti 2.0 database management system
- First STAC-M3 benchmark results using AMD EPYC CPUs
- Ran baseline (Antuco) benchmarks
- STAC-M3 Pack for shakti, Compatibility Rev A



www.STACresearch.com/SHK211203

STAC-M3 / shakti 2.0 / AMD EPYC 7742 / DDN ES200NVX All-Flash

- Stack:
 - shakti 2.0
 - GIGABYTE™ R282-Z90 server
 - 2 x AMD EPYC 7742 64-core CPU
 - 512GiB DRAM
 - CentOS 8.3.2011
 - EXAScaler Parallel Client 5.2.2
 - DDN ES200NVX appliance
 - DDN EXAScaler Parallel Filesystem 5.2.2
 - 84TiB total physical capacity (61TiB usable)
 - Mellanox SB7700 InfiniBand EDR 100Gbps switch



www.STACresearch.com/SHK211203

Versus all previously published results


- Had the highest storage efficiency at 196% (STAC-M3.v1.1.STORAGE.EFF)
- Had 1.7x the speed in STAC-M3.β1.1T.NBBO-Q.TIME versus the best previously published results for the less-demanding STAC-M3.β1.1T.NBBO.TIME benchmark



www.STACresearch.com/SHK211203

Vs. all-flash storage appliance w/ kdb+ 4.0 & 15 x 20-core servers*

- Outperformed in 3 of 17 mean response-time benchmarks, including:
 - 3.7x the speed in the version of Year-High Bid that allows caching (STAC-M3.β1.1T.YRHIBID-2.TIME)
 - 3.3x the speed in vs. the less-demanding STAC-M3.β1.1T.NBBO.TIME benchmark (STAC-M3.β1.1T.NBBO-Q.TIME)




www.STACresearch.com/SHK211203

* SUT ID KDB211014

STAC-M3/kdb+ 4.0/12x GCE N2 inst./Cascade Lake CPUs/105TiB NVMe


- Generational & architectural update to previous Google Cloud SUT (KDB181001)
- 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 across the 12 nodes



www.STACresearch.com/KDB211210

STAC-M3/kdb+ 4.0/12x GCE N2 inst./Cascade Lake CPUs/105TiB NVMe


- Stack:
 - kdb+ 4.0, running in sharded mode
 - 12 x Google Compute Engine N2 instances, each with:
 - 32 x Intel® Xeon® Cascade lake vCPUs @ 2.8GHz (reported base frequency)
 - 160GiB of memory
 - 24 x 375GiB local NVMe (105TiB in cluster)
 - 1 x 32Gbps network interface
 - Google Virtual Private Cloud software-defined network
 - Premium tier: 32Gbps



www.STACresearch.com/KDB211210

Compared to all publicly disclosed results


- Fastest Year-High Bid that allows caching (STAC-M3.β1.1T.YRHIBID-2.TIME)



www.STACresearch.com/KDB211210

Compared to a previous Google Cloud SUT*

- Comparison SUT: 13 Google Cloud Skylake instances, Google Persistent SSD, kdb+ 3.6, STAC Pack Compatibility Rev E
- Was faster in 13 of 17 Antuco mean response-time benchmarks, including:
 - 18x the speed of in version of Year-High Bid that allows caching (STAC-M3.β1.1T.YRHIBID-2.TIME)
- Was faster in 22 of 24 Kanaga mean response-time benchmarks, including:
 - over 10x the speed in market snapshot (STAC-M3.β1.10T.YR[2,3,4,5]-MKTSNAP.TIME)




www.STACresearch.com/KDB211210

* SUT ID KDB181001

Vs. parallel file system and 55 instances in public cloud*

- Comparison SUT: parallel filesystem across 40 instances, kdb+ 4.0 on 15 instances, STAC Pack Compatibility Rev E
- Was faster in 11 of 17 Antuco mean response-time benchmarks, including:
 - 8x the speed of in version of Year-High Bid that allows caching (STAC-M3.β1.1T.YRHIBID-2.TIME)
- Was faster in 13 of 24 Kanaga mean response-time benchmarks, including:
 - over 14x the speed in market snapshot (STAC-M3.β1.10T.YR[2,3,4,5]-MKTSNAP.TIME)



www.STACresearch.com/KDB211210

* SUT ID KDB210507

STAC-M3 / kdb+ 4.0 / 8x AWS EC2 x2iedn.32xlarge / io2 Block Express


- Amazon's first engineered STAC-M3 solution
- Tested with STAC-M3 baseline (Antuco)
- STAC-M3 Packs for kdb+: Compatibility Rev H



www.STACresearch.com/KDB220410

STAC-M3 / kdb+ 4.0 / 8x AWS EC2 x2iedn.32xlarge / io2 Block Express

- Stack:
 - kdb+ 4.0, running in distributed mode
 - 8 x EC2 x2iedn.32xlarge instances, each with:
 - 128 x Intel® Xeon® (Ice Lake) vCPUs @ 2.9GHz (reported base frequency)
 - 4096GiB of memory
 - 1 x 100Gbps Amazon Elastic Network Interface
 - Amazon Linux 2
 - Only mod: limiting CPU sleep states
 - Amazon Elastic Block Store io2 Block Express
 - 4 x 4000GiB multi-attach volumes
 - 1 x 5000GiB volume
 - Network
 - Amazon Virtual Private Cloud



www.STACresearch.com/KDB220410

Compared to all publicly disclosed results

- Fastest 10-user volume curve mean-response time (STAC-M3.β1.10T.VOLCURV.TIME)



www.STACresearch.com/KDB220410

Vs. parallel file system and 55 EC2 instances in AWS*

- Comparison SUT: parallel filesystem across 40 Amazon EC2 instances, kdb+ 4.0 on 15 Amazon EC2 instances, STAC Pack Compatibility Rev E
- Was faster in 15 of 17 Antuco mean response-time benchmarks, including:
 - 21x speed-up in 10-user volume curve (STAC-M3.β1.10T.VOLCURV.TIME)
 - 5.1x speed-up in 10-user theoretical P&L (STAC-M3.β1.10T.THEOPL.TIME)



www.STACresearch.com/KDB220410

* SUT ID KDB210507

Vs. network-attached flash storage & 9 DB servers on-prem*

- Comparison SUT: network-attached flash storage, kdb+ 4.0 on 9 database servers, STAC Pack Compatibility Rev H
- Was faster in 17 of 17 Antuco mean response-time benchmarks, including:
 - 11x speed-up in single-user stats with unpredictable intervals (STAC-M3.β1.1T.STATS-UI.TIME)
 - 5.0x speed-up in 10-user stats with unpredictable intervals (STAC-M3.β1.10T.STATS-UI.TIME)



www.STACresearch.com/KDB220410

* SUT ID KDB210929