STAC Update:
Big workloads

Peter Lankford
Founder and Director, STAC

peter.lankford@STACresearch.com

STAC-A2

- Non-trivial Monte Carlo
  - Heston-based Greeks for multi-asset, path-dependent options with early exercise
  - Metrics: Speed, capacity, quality, efficiency
- Numerous reports
  - Some public, some in the STAC Vault
- Premium STAC members get:
  - Reports in STAC Vault
  - Detailed config info on public and private reports
  - Code from vendor implementations of the benchmarks

www.STACresearch.com/**a2**

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Dell cluster using Intel Omni-Path Architecture and Intel MPI

- SUT ID: INTC181012
- Intel Parallel Studio XE 2018 Update 3, including Intel TBB and Intel MPI
- 2 x Dell EMC PowerEdge C6400 chassis with [4 x PowerEdge C6420 server sleds with (2 x 28-core Intel Xeon Platinum 8176F CPU)]
- Dell Networking H1048-OPF switch
  - Intel Omni-Path Architecture

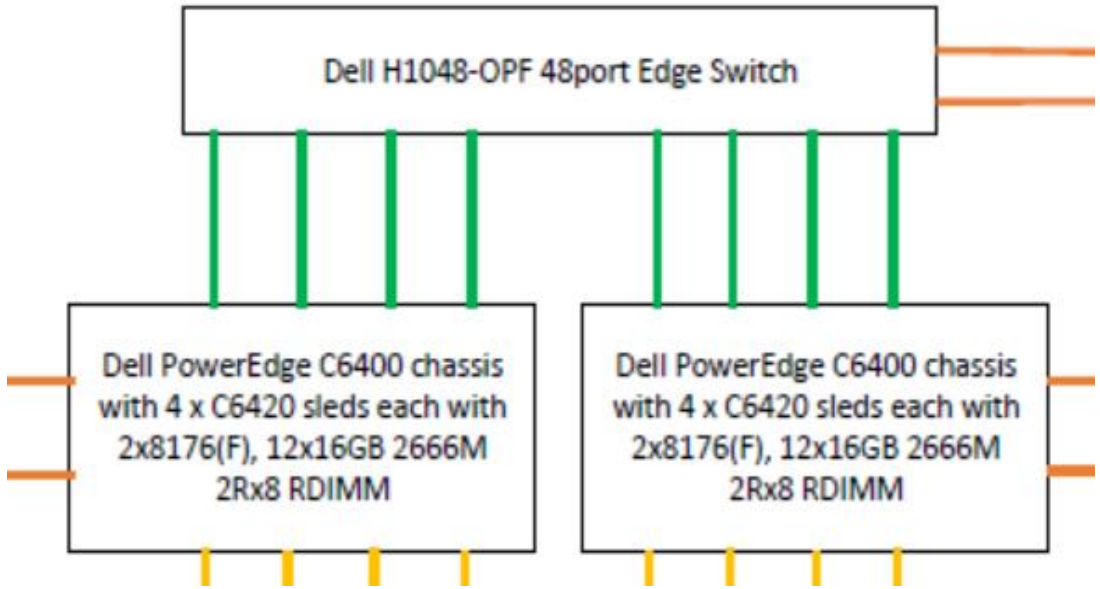

Report coming soon

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Dell cluster using Intel Omni-Path Architecture and Intel MPI

- Total SUT rack units: 4.2U
  - 4U for servers
  - 1/5 of 1U OPA switch
- 448 total cores, 1.5 TB total DRAM
- Operated as single unit of compute
  - Cluster cooperated on computation of single problems
- Intel's objectives with this project:
  - Maximize throughput and efficiency
  - Keep base response times low

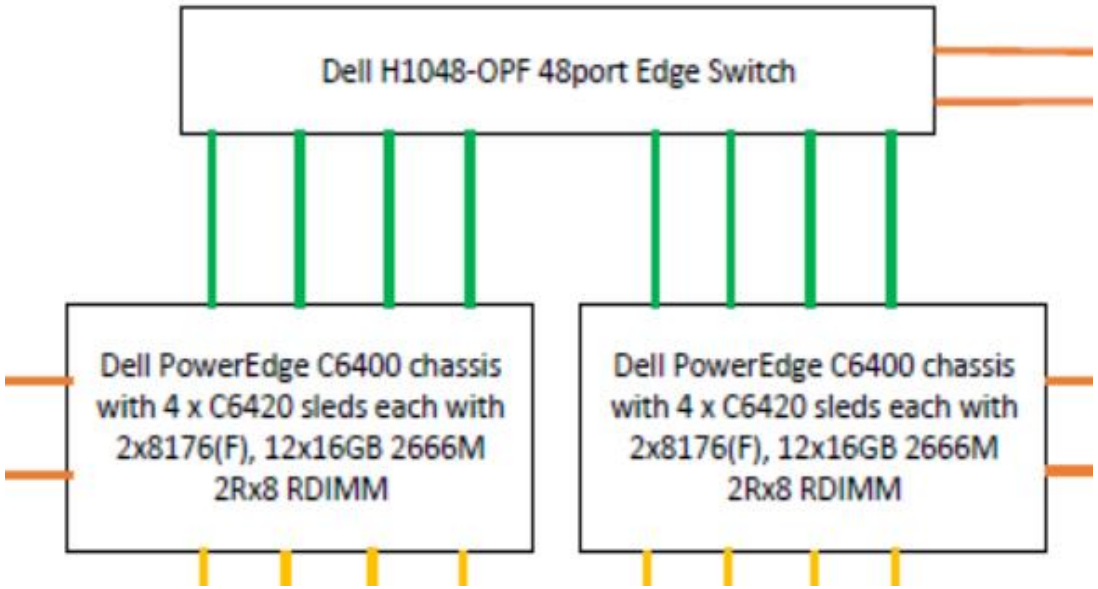

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Throughput and efficiency

- Highest throughput of any solution (STAC-A2.β2.HPORTFOLIO.SPEED)
  - 56% higher than the best throughput from a solution using GPUs (SUT ID NVDA171020)
- Highest space efficiency of any solution (STAC-A2.β2.HPORTFOLIO.SPACE_EFF)
  - 36% higher than SUT ID NVDA171020
- 91% higher energy efficiency than the previous best Intel-only solution (SUT ID INTC170503)
  - Only 30% less energy efficiency than SUT ID NVDA171020

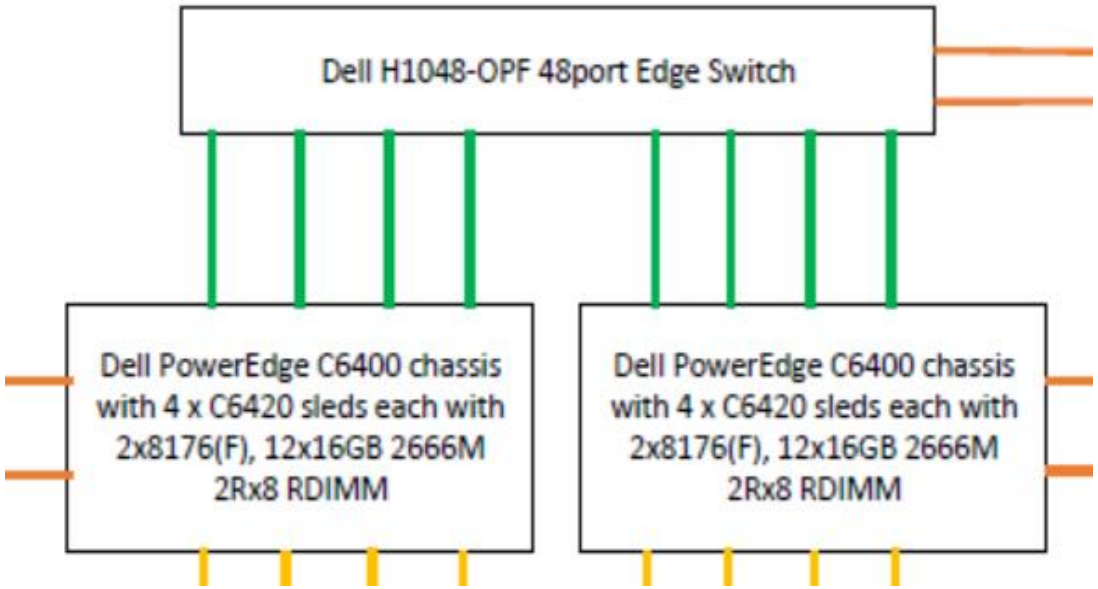

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Response times

- Fastest WARM time in the large problem size (STAC-A2.β2.10-100K-1260.TIME.WARM)
  - 2.5x the speed of SUT ID NVDA171020
- Fastest COLD time in the large problem size (STAC-A2.β2.10-100K-1260.TIME.COLD)
  - 63% faster than SUT ID NVDA171020
- Largest basket size achieved (STAC-A2.β2.GREEKS.MAX_ASSETS)
- WARM runs of baseline GREEKS (STAC-A2.β2.GREEKS.TIME.WARM): same order of magnitude as SUT ID NVDA171020
  - 34 milliseconds for this SUT,
    21 milliseconds for NVDA171020

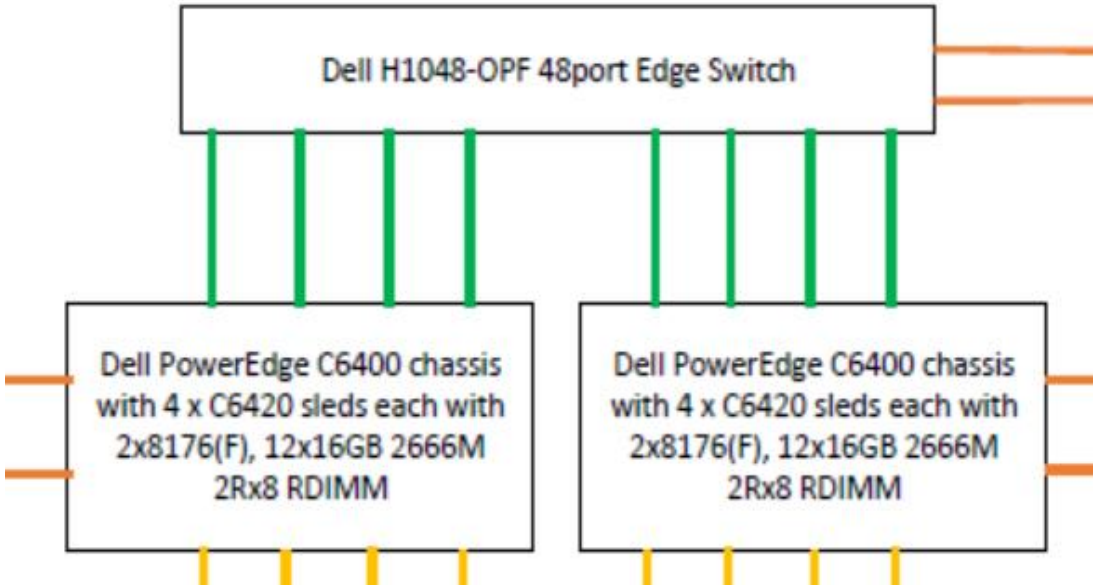

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

8 x NVIDIA Volta (V100) in HPE ProLiant XL270d Gen10 Server

- SUT ID: NVDA181105
- Update of SUT ID NVDA171020
- SUT:
  - STAC-A2 Pack for CUDA (Rev D)
  - NVIDIA CUDA 9.2
  - 8 x NVIDIA Volta V100 GPUs
  - 2 x Intel Xeon Gold 6150 @ 2.7GHz
  - HPE ProLiant XL270d Gen10
  - 768 GB DRAM
  - RHEL 7.5

STAC Report coming soon

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

Compared to all publicly reported solutions

- *Fastest warm time in baseline Greeks*
  - *STAC-A2.β2.GREEKS.TIME.WARM*
  - *80% faster than the best result from a non-NVIDIA-based solution (SUT ID INTC181012)*
- *Highest maximum paths*
  - *STAC-A2.β2.GREEKS.MAX_PATHS*
  - *48% higher than the best result from a solution based solely on standard CPUs (INTC150811)*
- *43% higher energy efficiency than the best non-NVIDIA-based solution (SUT ID INTC181012)*
  - *STAC-A2.β2.HPORTFOLIO.ENERG_EFF*

STAC Report coming soon

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Compared to all publicly reported single server solutions

- Highest throughput
  - STAC-A2.β2.HPORTFOLIO.SPEED
- Fastest warm and cold times in the large problem size
  - STAC-A2.β2.10-100K-1260.TIME.WARM
  - STAC-A2.β2.10-100K-1260.TIME.COLD
- Largest basket size
  - STAC-A2.β2.GREEKS.MAX_ASSETS

STAC Report coming soon

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3

- Performance benchmarks for enterprise tick analytics
  - Language/DBMS neutral
  - Developed by banks and hedge funds
- 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

STAC-M3 Shasta / kdb+ / Google Cloud n1-Ultramem-160

- SUT ID: KDB180713
- Stack:
  - Software: kdb+ 3.5 / CentOS 7.5
  - Instance: GCP n1-Ultramem-160 (160 vCPU, 3.97 TB DRAM)
  - Storage: Google Persistent SSD (but data pre-loaded into memory)
- Point of STAC-M3 Shasta
  - Assess real-world performance when using a relatively small database

www.STACresearch.com/KDB180713

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3 Shasta / kdb+ / GCP n1-Ultramem-160

- Results highlights:
  - Outperformed bare metal solution based on Broadwell EX and 6TB DRAM (SUT ID KDB160425) in 8 of the 15 required benchmarks.
  - Outperformed bare metal solution based on Ivy Bridge EX and 6TB DRAM (SUT ID KDB140116) in 14 of the 15 required benchmarks.
    - In 5 of these, the GCP solution was more than 2x the speed

www.STACresearch.com/KDB180713

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3 / kdb+ / GCP Cluster with Persistent SSD

- SUT ID: KDB181001
- Stack:
  - Software: kdb+ 3.6 / CentOS 7.5 / xfs
  - 13 instances, each with 32vCPU, 128GB DRAM
  - Storage: Google Persistent SSD
- STAC-M3 Antuco and Kanaga
  - Baseline and scale tests

Reports coming soon

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Compared to a Lustre-based on-prem cluster (KDB150528)

- Baseline (Antuco) results:
  - GCP-based solution outperformed in 14 of the 17 required benchmarks
  - From 1.3x to 7.8x speedup
- Scale (Kanaga) results:
  - GCP-based solution outperformed in 16 of 16 benchmarks reported for KDB150528*
  - From 1.6x to 12.6x speedup

* KDB150528 operated on only 4 years of data. For that dataset size, the Kanaga suite has 16 benchmarks. The GCP solution operated on 5 years of data, which results in 24 benchmarks.

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3 / kdb+ / Lenovo SR950 / 3D+NAND + a little Optane

- SUT ID: KDB181009
- Stack:
  - kdb+ 3.5
  - Lenovo ThinkSystem SR950
  - 4 x Intel Xeon Platinum 8180 CPUs
  - 3TB DRAM
  - 6TB Intel DC P4800X (Optane) SSDs
  - 56TB Intel DC P4510 (3D NAND) SSDs
  - Patches for Spectre & Meltdown
- STAC-M3 Antuco and Kanaga
  - Baseline tests and scale tests

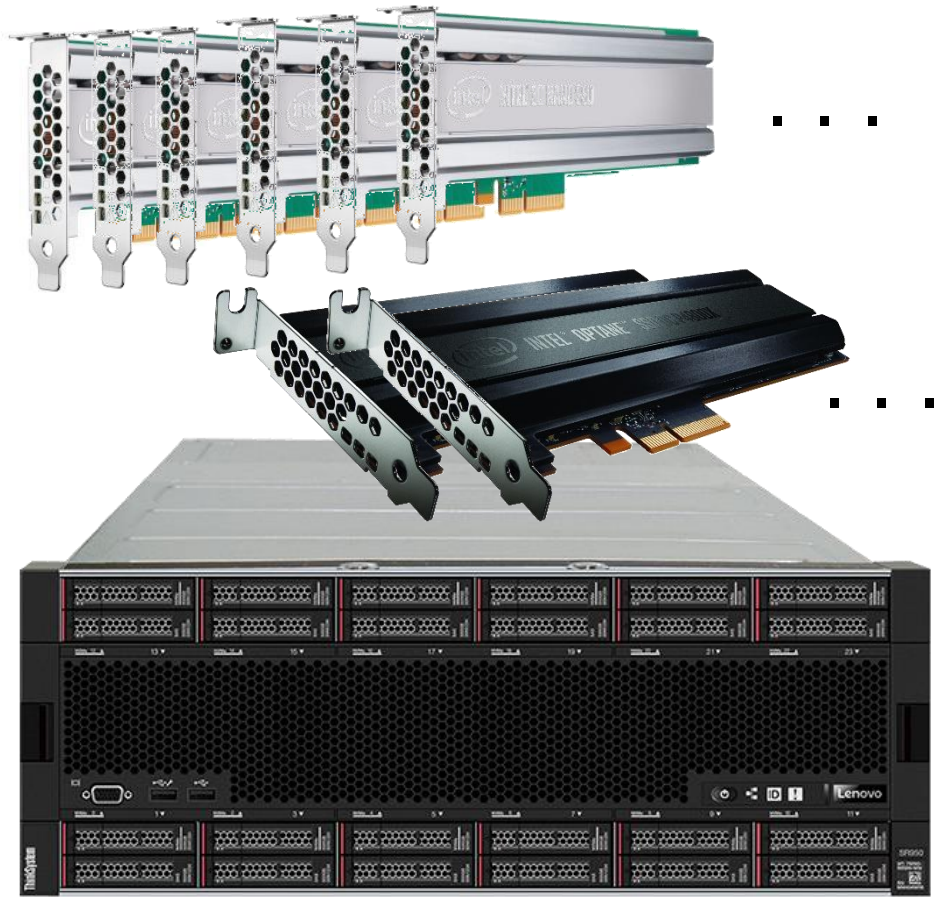

www.STACresearch.com/KDB181009

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

Results

- Set records in all scales of the market snapshot benchmark
  - STAC-M3.β1.10T.YR[n]-MKTSNAP.TIME
  - Year 5: 58% faster than the next best result
  - Year 4: 46% faster than the next best result
  - Year 3: 20% faster than the next best result
  - Year 2: 13% faster than the next best result
- Set a new record in the 50-user 12-day VWAB in Year 1
  - STAC-M3.β1.50T.YR1VWAB-12D-HO.TIME
  - 40% faster than the next best result

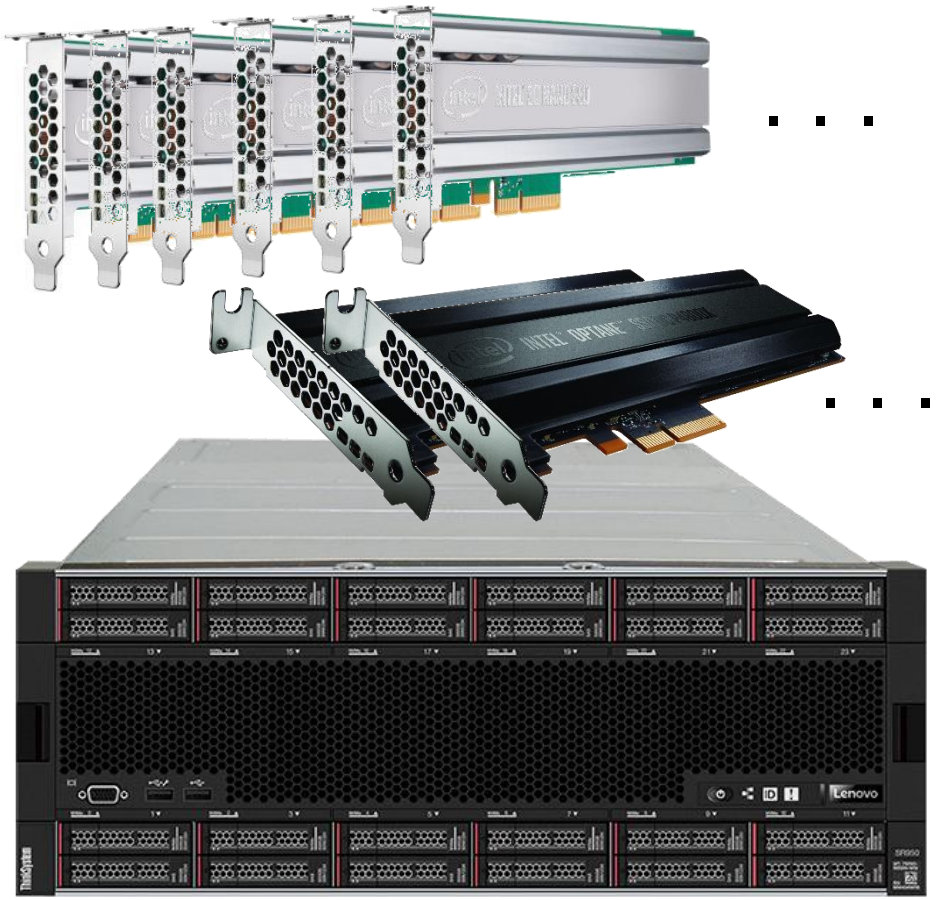

www.STACresearch.com/KDB181009

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

STAC-M3 Working Group – Important meeting(s) coming up

- How to assess price performance with deployed infra & IaaS?
- How to assess price performance with DBaaS & FaaS (“serverless”)?
- Is the set of STAC-M3 operations still representative?
- Should we make the scale tests part of the baseline?
- Should we have a STAC-M3 “teaser suite” for quick-and-dirty evaluation of emerging databases?

STAC-A3

- Recap:
  - Workloads that emulate real-world backtesting jobs
  - Measure speed, scalability, efficiency of any architecture
- Test harness hands the implementation jobs to execute
- Measures the throughput and efficiency of the SUT
- Currently defined algos are almost always bottlenecked on I/O

STAC-A3 Working Group – Important meeting(s) coming up

- Clarifying and streamlining the benchmark results set for SWEEP
  - Making it easier to do apples-to-apples comparisons
- Confirming benchmarks for SWEEP using options
- Confirming benchmarks for BLASH algorithm
- Defining portfolio optimization algorithm

Refresh: STAC Cloud SIG

- Subgroup to guide dialog and research on common issues
- Public, private, hybrid cloud
- Several financial firms
  - Banks, hedge funds, prop shops (soon asset managers)
  - Can have more
- Several vendors
  - Major players in private/public/hybrid cloud ecosystems are involved
  - Look for announcements soon on additional vendors

STAC Cloud SIG way forward

- Cloud-specific issues will be in Cloud SIG meetings
  - For example, will be setting up Q&As with each of the big 3 public cloud providers on security
- For workload-specific issues, Cloud SIG needs representation on domain working groups to provide input/ensure consistency. E.g.:
  - How to apply pricing to use cases
  - Dealing with conditions (regions, time of day, day of week) and performance variability
- For example, STAC-M3 WG will drive STAC-M3 price-performance but needs insights from cloud experts

To get involved

- Click the “Enable me” button at

www.STACresearch.com/**cloudsig**

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER