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

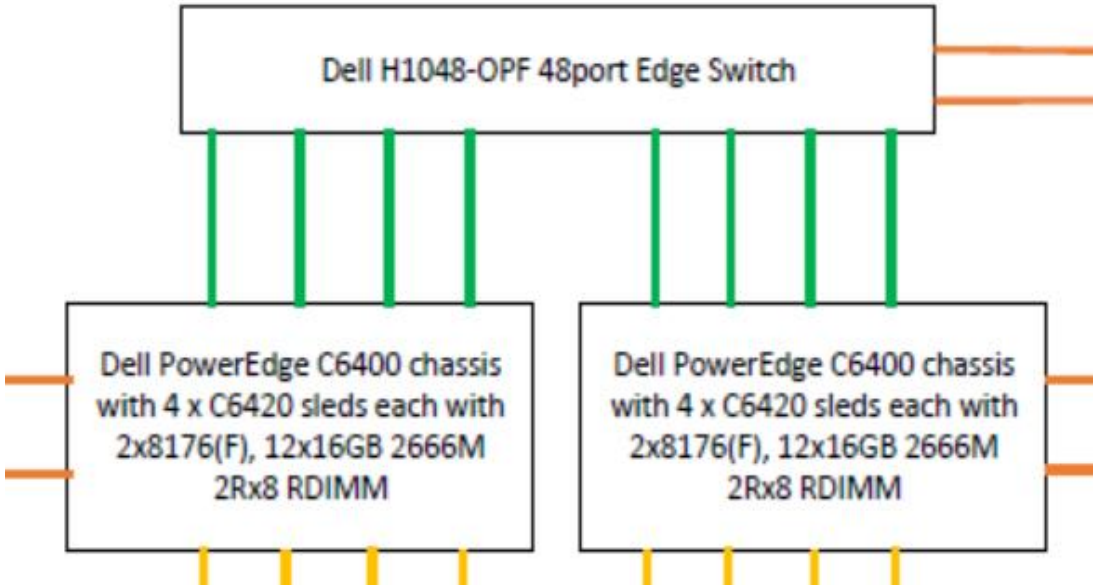

www.STACresearch.com/INTC181012

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

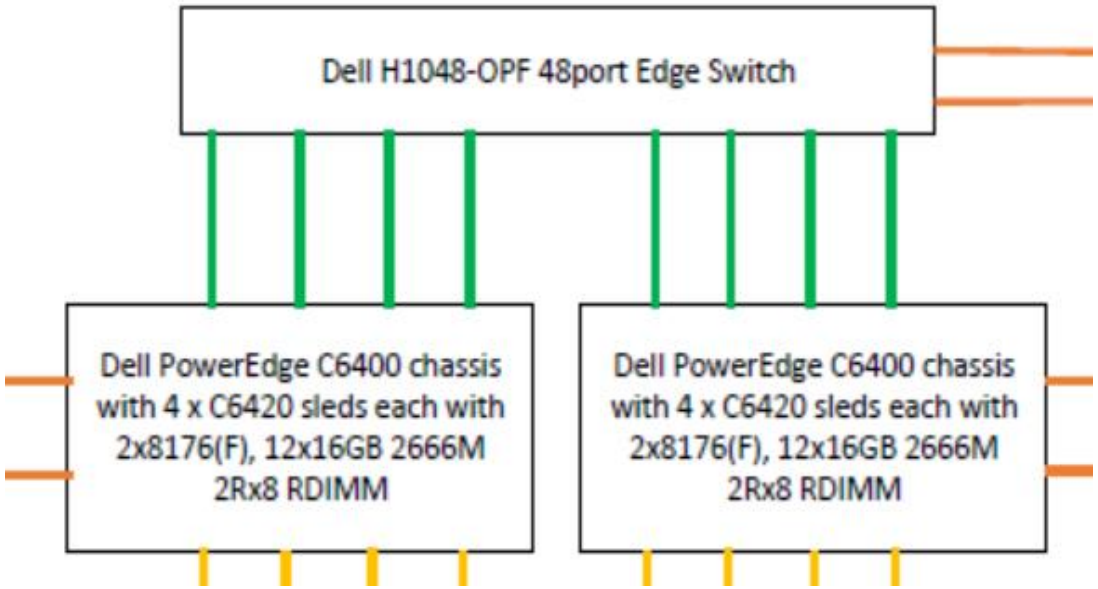

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

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
- 86% higher energy efficiency than the previous best Intel-only solution (SUT ID INTC170503)
  - Only 32% less energy efficiency than SUT ID NVDA171020

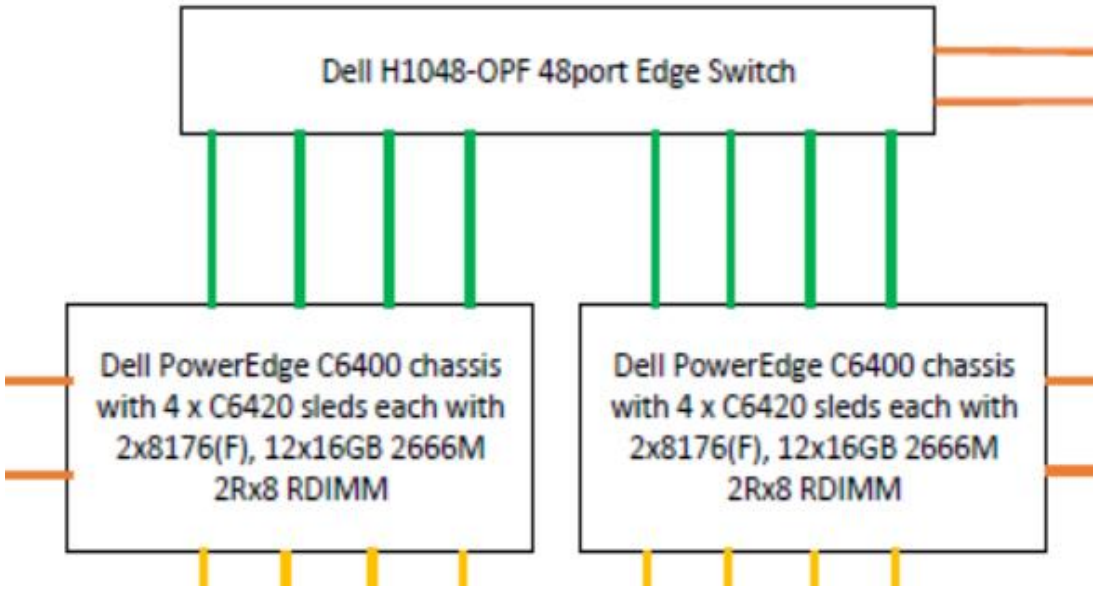

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

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

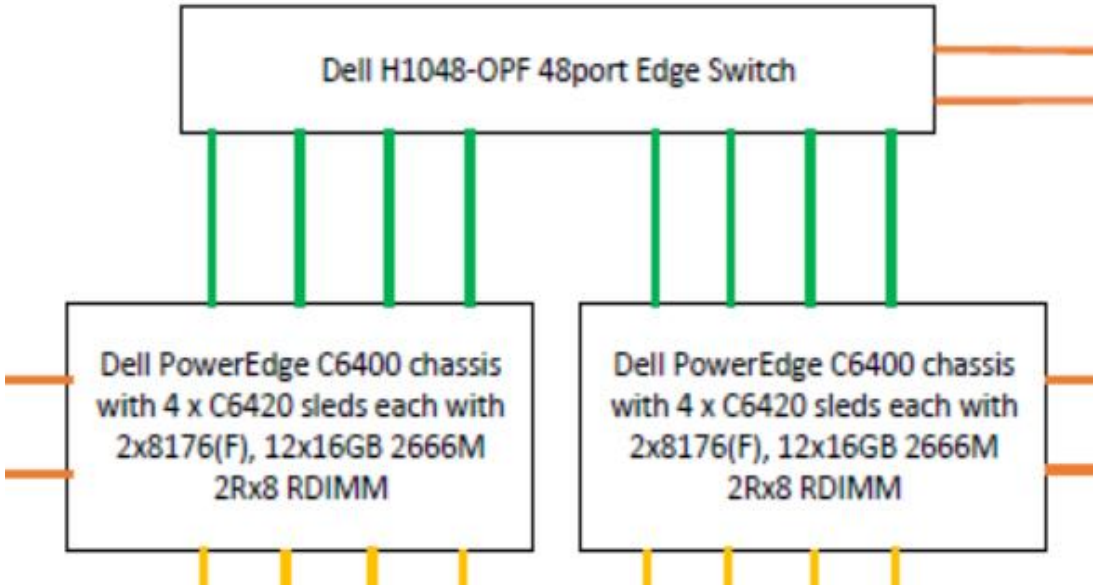

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

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
- Google has now disclosed results for all three STAC-M3 suites

www.STACresearch.com/KDB181001

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.

www.STACresearch.com/KDB181001

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

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

- 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?
- How to assess price performance with deployed infra & IaaS?
- How to assess price performance with DBaaS & FaaS (“serverless”)?

www.STACresearch.com/m3

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

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

- STAC-A3 refresh:
  - Workloads that emulate real-world backtesting jobs
  - Measure speed, scalability, efficiency of any architecture
- 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

www.STACresearch.com/a3

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER