Adventures on AWS with High-Performance Workloads

Peter Lankford
Founder and Director, STAC

peter.lankford@STACresearch.com

Why research public cloud?

- Some of you are using it today
- Some of you are evaluating it
- Most of you find the comparisons useful
  - Comparisons within a cloud service (e.g., price-performance impact of a new processor)
  - Comparisons to your internal price-performance

Why research AWS?

- Market share leader
- Many configurations to choose from and compare
  - CPU and GPU
  - Different processor types, speeds
  - Different memory configurations
  - Etc.
- This means:
  - Could be daunting for a user to figure out instance with best price-performance
  - Can make many useful comparisons even for deployed systems

Instance Types Matrix

Instance Type	vCPU	Memory (GiB)	Storage (GB)	Networking Performance	Physical Processor	Clock Speed (GHz)	Int AV
t2.micro	1	1	EBS Only	Low to Moderate	Intel Xeon family	2.5	Ye
t2.small	1	2	EBS Only	Low to Moderate	Intel Xeon family	2.5	Ye
t2.medium	2	4	EBS Only	Low to Moderate	Intel Xeon family	2.5	Ye
m3.medium	1	3.75	1 x 4 SSD	Moderate	Intel Xeon E5-2670 v2*	2.5	Ye
m3.large	2	7.5	1 x 32 SSD	Moderate	Intel Xeon E5-2670 v2*	2.5	Ye
m3.xlarge	4	15	2 x 40 SSD	High	Intel Xeon E5-2670 v2*	2.5	Ye
m3.2xlarge	8	30	2 x 80 SSD	High	Intel Xeon E5-2670 v2*	2.5	Ye
c4.large	2	3.75	EBS Only	Moderate	Intel Xeon E5-2666 v3	2.9	Ye
c4.xlarge	4	7.5	EBS Only	High	Intel Xeon E5-2666 v3	2.9	Ye
c4.2xlarge	8	15	EBS Only	High	Intel Xeon E5-2666 v3	2.9	Ye
c4.4xlarge	16	30	EBS Only	High	Intel Xeon E5-2666 v3	2.9	Ye
c4.8xlarge	36	60	EBS Only	10 Gigabit	Intel Xeon E5-2666 v3	2.9	Ye
c3.large	2	3.75	2 x 16 SSD	Moderate	Intel Xeon E5-2680 v2	2.8	Ye
c3.xlarge	4	7.5	2 x 40 SSD	Moderate	Intel Xeon E5-2680 v2	2.8	Ye
c3.2xlarge	8	15	2 x 80 SSD	High	Intel Xeon E5-2680 v2	2.8	Ye
c3.4xlarge	16	30	2 x 160 SSD	High	Intel Xeon E5-2680 v2	2.8	Ye
c3.8xlarge	32	60	2 x 320 SSD	10 Gigabit	Intel Xeon E5-2680 v2	2.8	Ye
g2.2xlarge	8	15	1 x 60 SSD	High	Intel Xeon E5-2670	2.6	Ye
g2.8xlarge	32	60	2 x 120 SSD	10 Gigabit	Intel Xeon E5-2670	2.6	Ye
r3.large	2	15.25	1 x 32 SSD	Moderate	Intel Xeon E5-2670 v2	2.5	Ye
r3.xlarge	4	30.5	1 x 80 SSD	Moderate	Intel Xeon E5-2670 v2	2.5	Ye
r3.2xlarge	8	61	1 x 160 SSD	High	Intel Xeon E5-2670 v2	2.5	Ye
r3.4xlarge	16	122	1 x 320 SSD	High	Intel Xeon E5-2670 v2	2.5	Ye
r3.8xlarge	32	244	2 x 320 SSD	10 Gigabit	Intel Xeon E5-2670 v2	2.5	Ye
i2.xlarge	4	30.5	1 x 800 SSD	Moderate	Intel Xeon E5-2670 v2	2.5	Ye
i2.2xlarge	8	61	2 x 800 SSD	High	Intel Xeon E5-2670 v2	2.5	Ye
i2.4xlarge	16	122	4 x 800 SSD	High	Intel Xeon E5-2670 v2	2.5	Ye
i2.8xlarge	32	244	8 x 800 SSD	10 Gigabit	Intel Xeon E5-2670 v2	2.5	Ye
d2.xlarge	4	30.5	3 x 2000	Moderate	Intel Xeon E5-2676 v3	2.4	Ye
d2.2xlarge	8	61	6 x 2000	High	Intel Xeon E5-2676 v3	2.4	Ye
d2.4xlarge	16	122	12 x 2000	High	Intel Xeon E5-2676 v3	2.4	Ye
d2.8xlarge	36	244	24 x	10 Gigabit	Intel Xeon E5-2676	2.4	Ye

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

What have we been doing?

- Use AWS as a customer
- Paid for an AWS Business Support plan
- Document our experiences for the benefit of STAC subscribers
- Test a bunch of instance types

STAC®
SECURITIES TECHNOLOGY ANALYSIS CENTER

What instance types did we choose for STAC-A2?

- What AWS calls “compute-optimized” and “memory-optimized” types
- Sadly, GPU instance types were not compatible with the STAC-A2 Pack for CUDA 5.5 or STAC-A2 Pack for CUDA 6.5
  - Future?
- Used latest STAC-A2 Pack for Intel Composer XE
- OS: chose RHEL 6.5 because it’s common
- Virtualiztion: chose HVM rather than PV
- Chose (mostly) Dedicated instance types because Multi-tenant introduces another variable
  - Did do a couple multi-tenant tests. But testing the impact of multi-tenancy is tricky.
- Chose On-Demand instance types because use case was cloud bursting

What instance types did we choose for STAC-A2?

The list:

- c3.4xlarge, dedicated, on-demand, RHEL 6.5
- c3.4xlarge, multi-tenant, on-demand, RHEL 6.5
- c3.8xlarge, dedicated, on-demand, RHEL 6.5
- c3.8xlarge, dedicated, on-demand, RHEL 6.6 kernel*
- c4.4xlarge, dedicated, on-demand, RHEL 6.5*
- c4.8xlarge, dedicated, on-demand, RHEL 6.6 kernel*
- c4.8xlarge, multi-tenant, on-demand, RHEL 6.6 kernel*
- r3.2xlarge, dedicated, on-demand, RHEL 6.5
- r3.4xlarge, dedicated, on-demand, RHEL 6.5
- r3.8xlarge, dedicated, on-demand, RHEL 6.5

* See the caveats in the Tech Note. There were configuration conflicts related to Xen/RHEL that limited what could be done with some instance types. RH created bug reports, and AWS corrected their instance type descriptions in response to our findings.

STAC® SECURITIES TECHNOLOGY ANALYSIS CENTER

Methodology

- Tests were standard STAC-A2
  - Including new 10-100k-1260 benchmark
- Price-performance extrapolated from WARM times to infer the cost to run:
  - 1 million jobs of the baseline workload (GREEKS) in one hour
  - 1 million jobs of the large workload (GREEKS.100-100k-1260) in one hour
- A customer can plug in its internal costs to compare

Selected results – baseline price performance

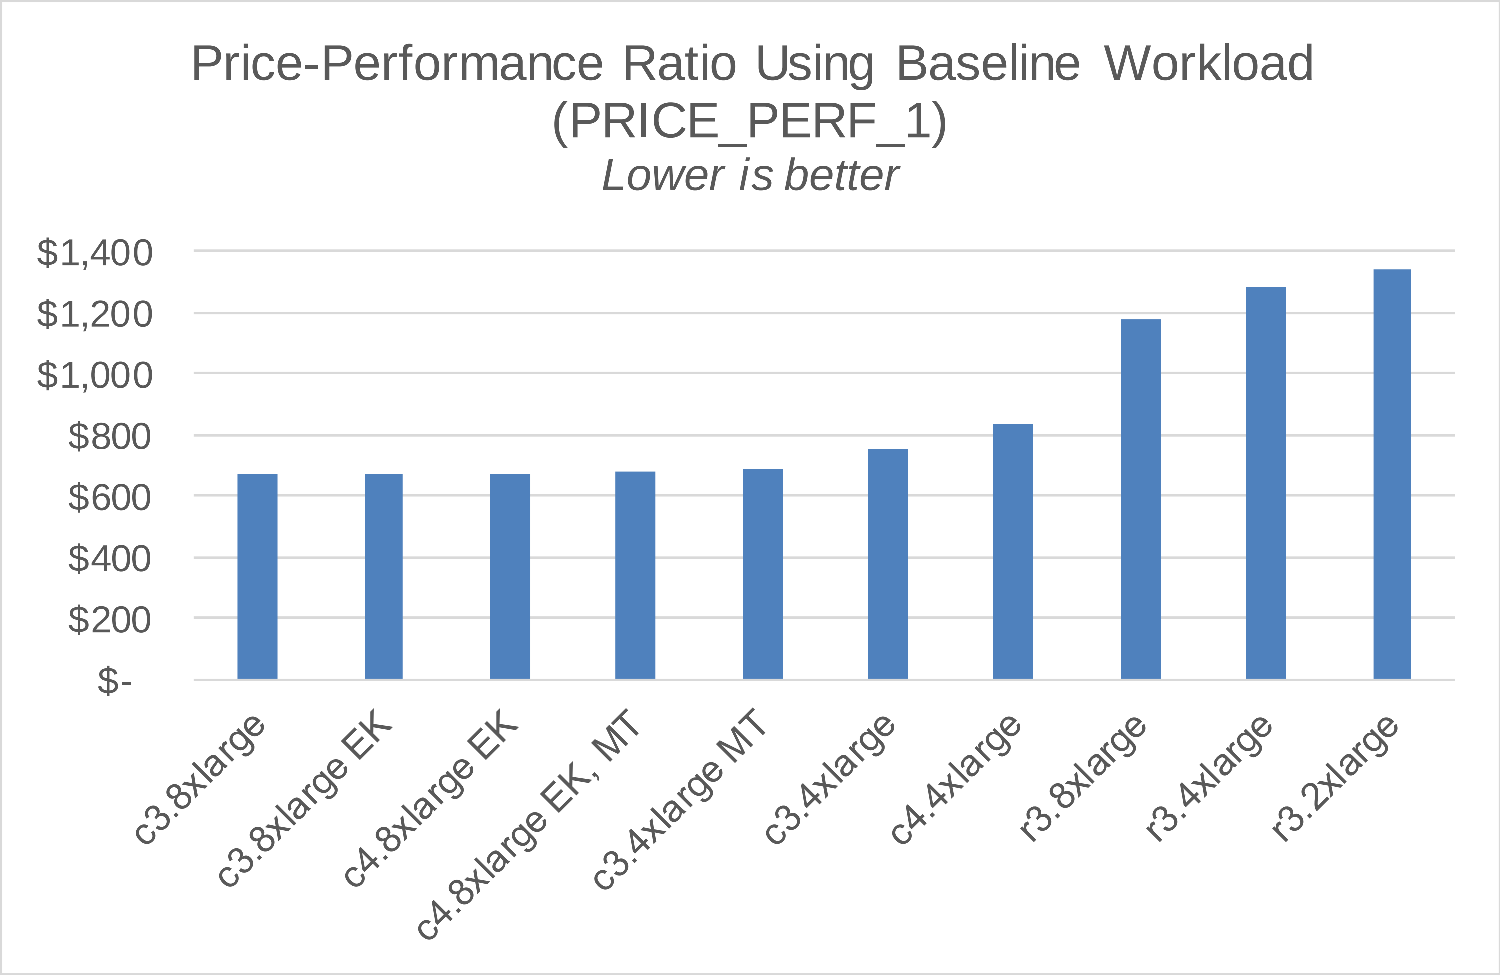

STAC SECURITIES TECHNOLOGY ANALYSIS CENTER

Selected results – absolute performance vs standalone

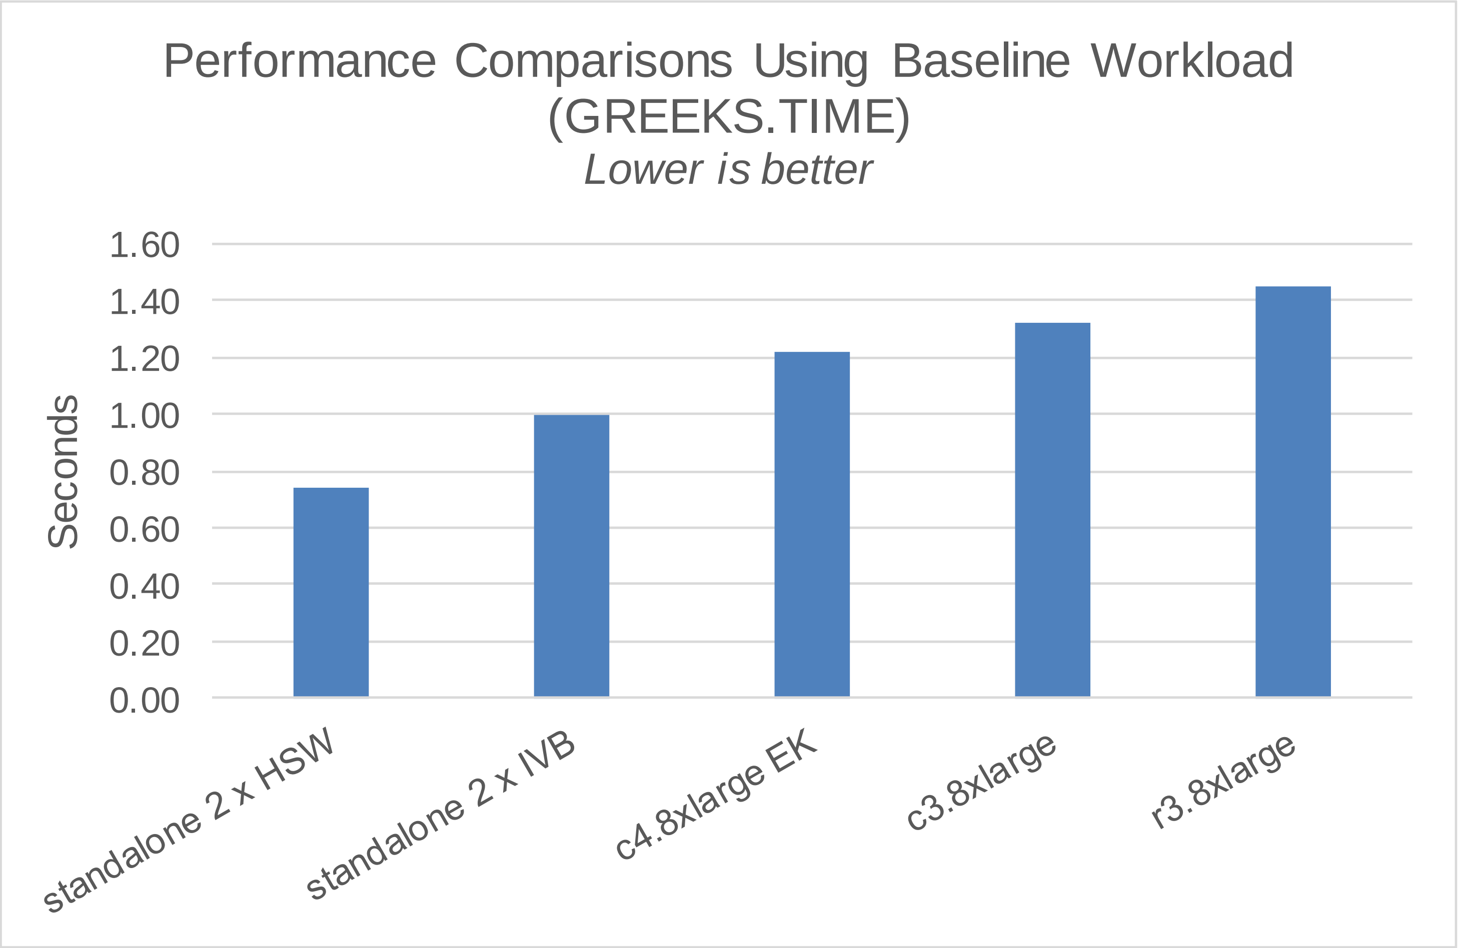

STAC SECURITIES TECHNOLOGY ANALYSIS CENTER

What instance types did we choose for STAC-M3?

- What AWS calls “storage-optimized” types
  - “very fast SSD-backed instance storage optimized for very high random I/O performance, and provide high IOPS at a low cost.”
- What AWS calls “dense storage” types
  - “lowest price per disk throughput performance on Amazon EC2”
- Only had scope to test two instance types so far
- There are many other configuration possibilities
  - Elastic Block Storage (EBS) General Purpose SSD
  - EBS Provisioned IOPS SSD
  - EBS Magnetic Volumes
  - Simple Storage Service (S3)
  - Combinations of these with various EC2 instance types
- Used latest STAC-M3 Pack for kdb+ 3.2
  - Based on interest expressed by customers
- Chose RHEL 6.5, Dedicated, On-Demand for same reasons as in STAC-A2 research

What instance types did we choose for STAC-M3?

The two:

- d2.8xlarge, dedicated, on-demand, RHEL 6.5
- i2.8xlarge, dedicated, on-demand, RHEL 6.5

* See configuration notes in the Tech Note.

Methodology

- Tests were standard STAC-M3 Antuco
  - Based on results, held off Kanaga until more feedback
- Price-performance methodology not settled
  - Two cases: batch & interactive
  - Proposal for batch: extrapolate to resources required to complete a large number of batch jobs
  - Proposal for interactive: extrapolate to resources required to maintain a response-time SLA for a large volume of queries
- Goal: Let a customer plug in its internal costs to compare
- Still working this up. Here's a sneak peek.

Selected results – absolute performance of i2 vs d2

- i2.8xlarge vs d2.8xlarge

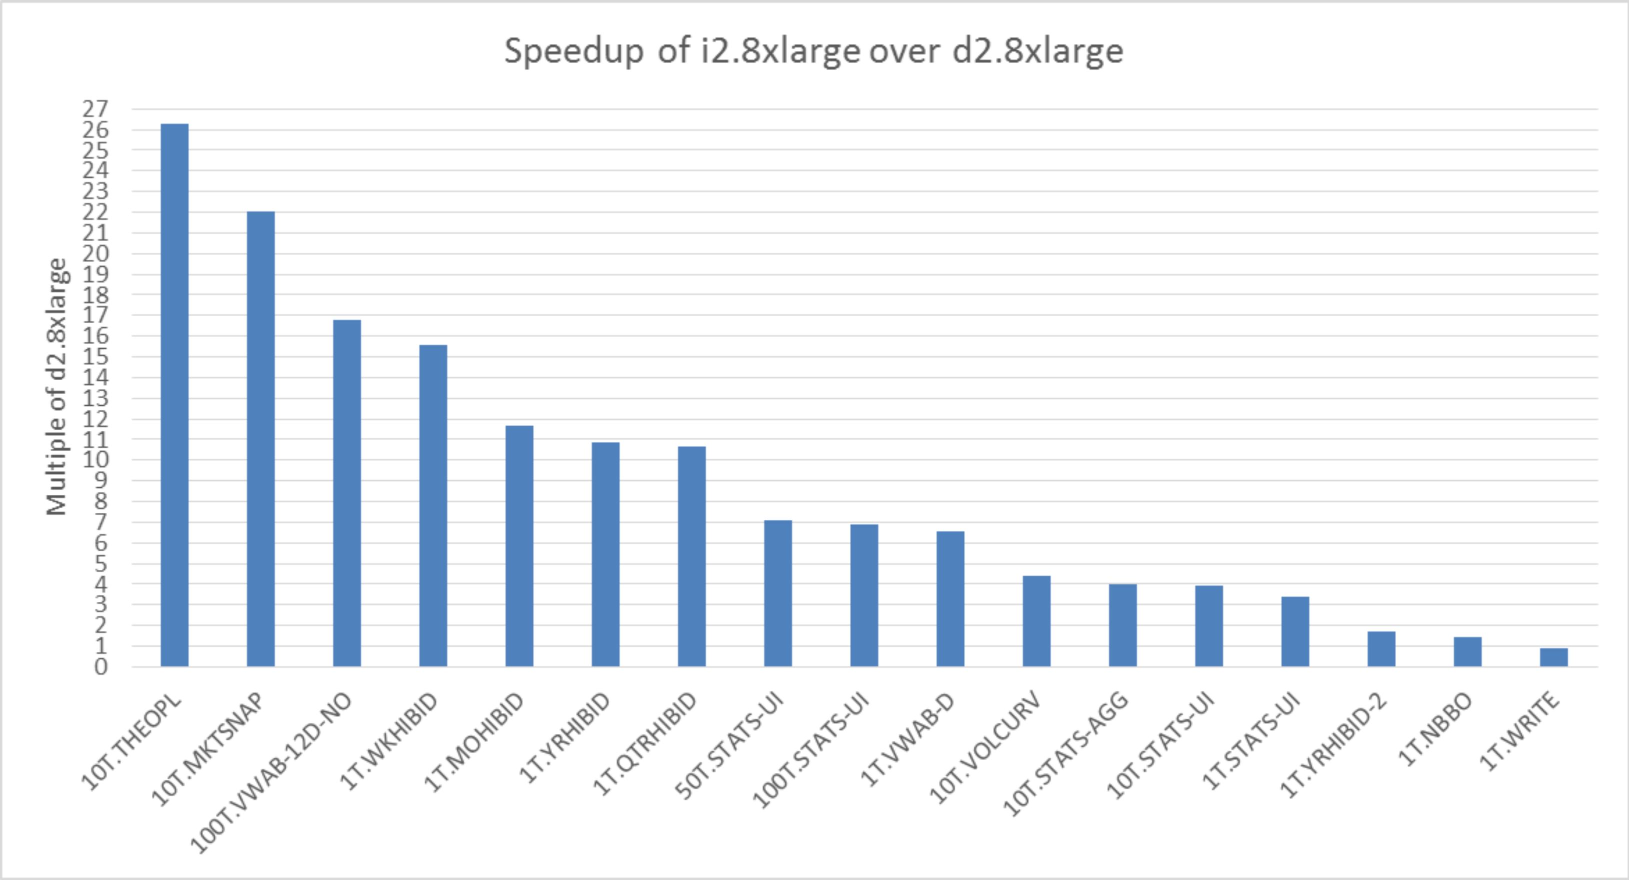

Absolute performance vs standalone systems

- i2.8xlarge vs Scalable Informatics 2-socket Ivy Bridge server with Optimus SSD

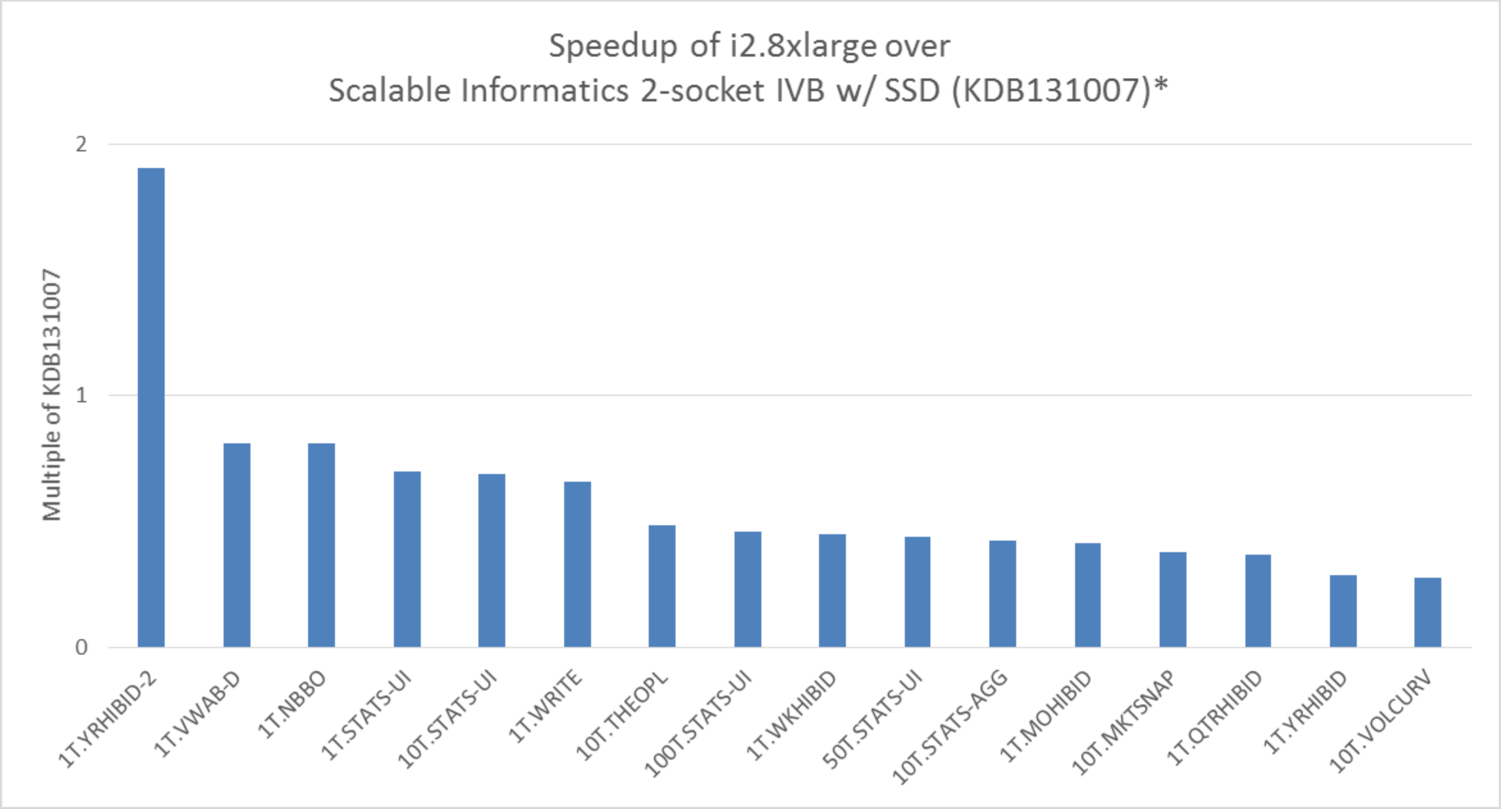

* 10T.MKTSNAP omitted because Kx radically improved performance of that benchmark after the Scalable tests.

Selected results: batch-based price performance

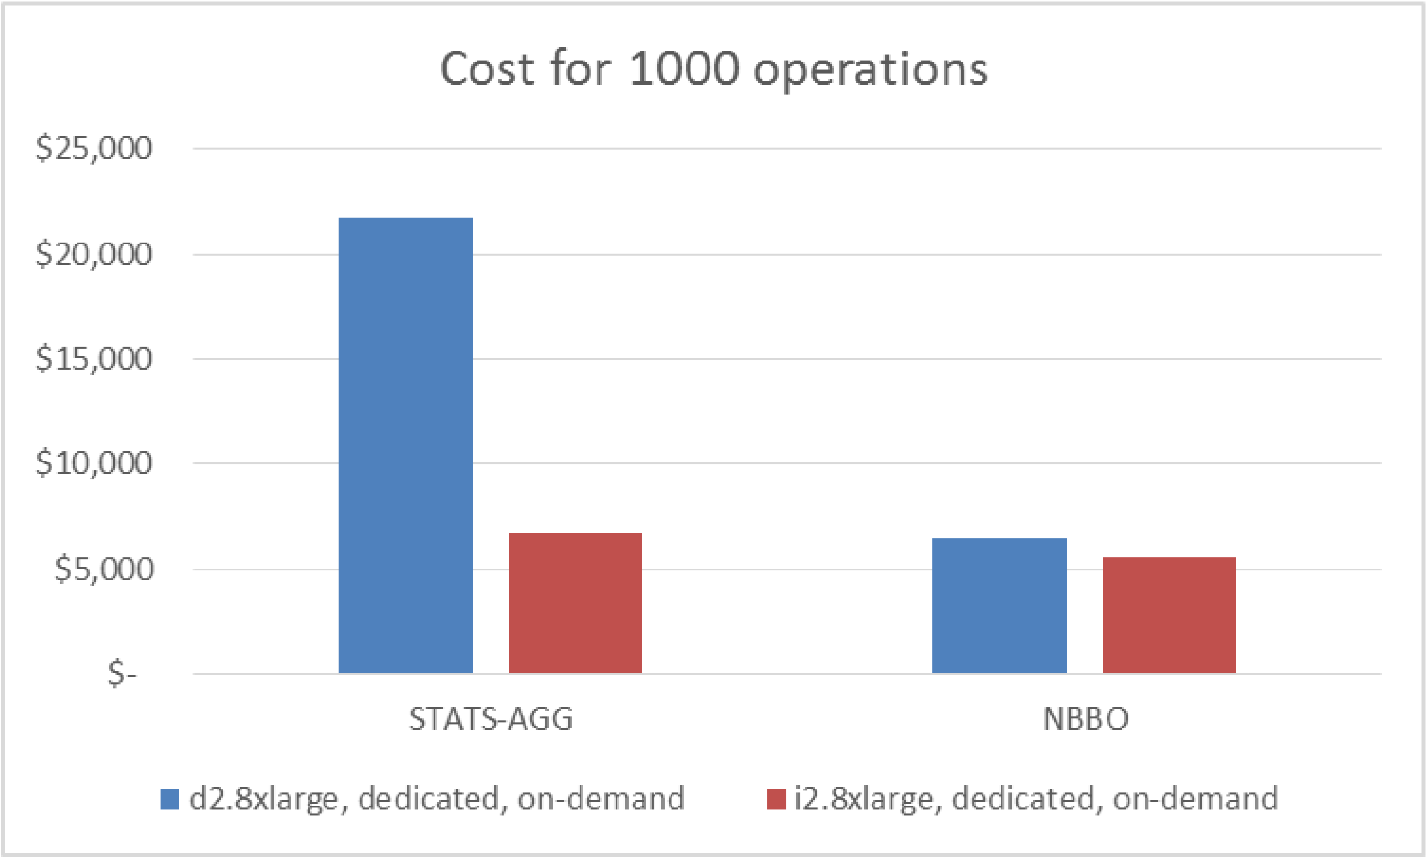

STAC SECURITIES TECHNOLOGY ANALYSIS CENTER

Next steps (need your help prioritizing)

- Other instance types
  - Under STAC-A2
  - Under STAC-M3
- Other clouds
  - IBM Softlayer, Microsoft Azure, Google Cloud
  - Specialty high-performance providers
- Other workloads?
  - STAC-A3 (backtesting)
  - STAC-M2 (messaging)
  - Other streaming benchmarks
- Getting cloud providers involved

Question

- Do we create a cloud interest group?
- Or is cloud simply one of the things to study within each workload domain?