Real-Time Persistent Computing for Big Data

Bernie Wu

Levyx

Time to Value & Cost of Latency are Increasing for Batch and "Real-time" Operations

Examples:

Amazon – 100 msec E-commerce latency increase = -1% sales

Google – 500 msec page search delay = -20% Traffic

Hi-Freq Trading House - -1 msec = $-100M annual profits

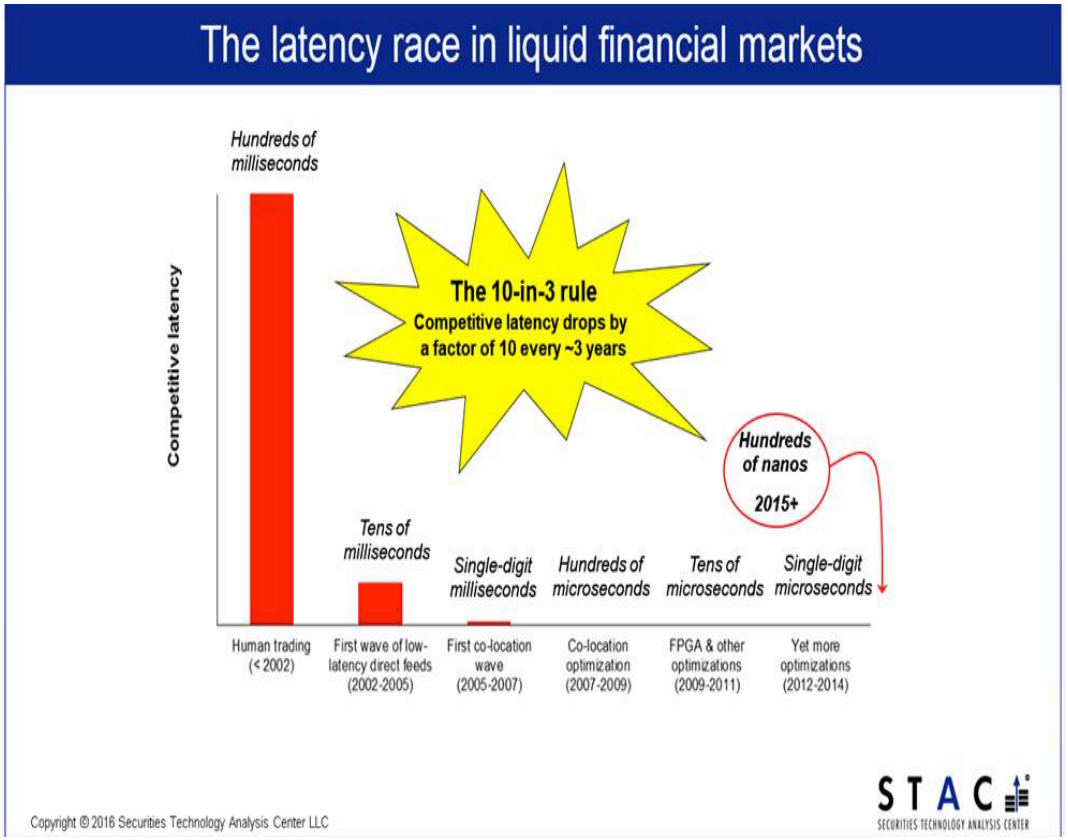

Software is now becoming the Biggest Component of Overall Application Latency

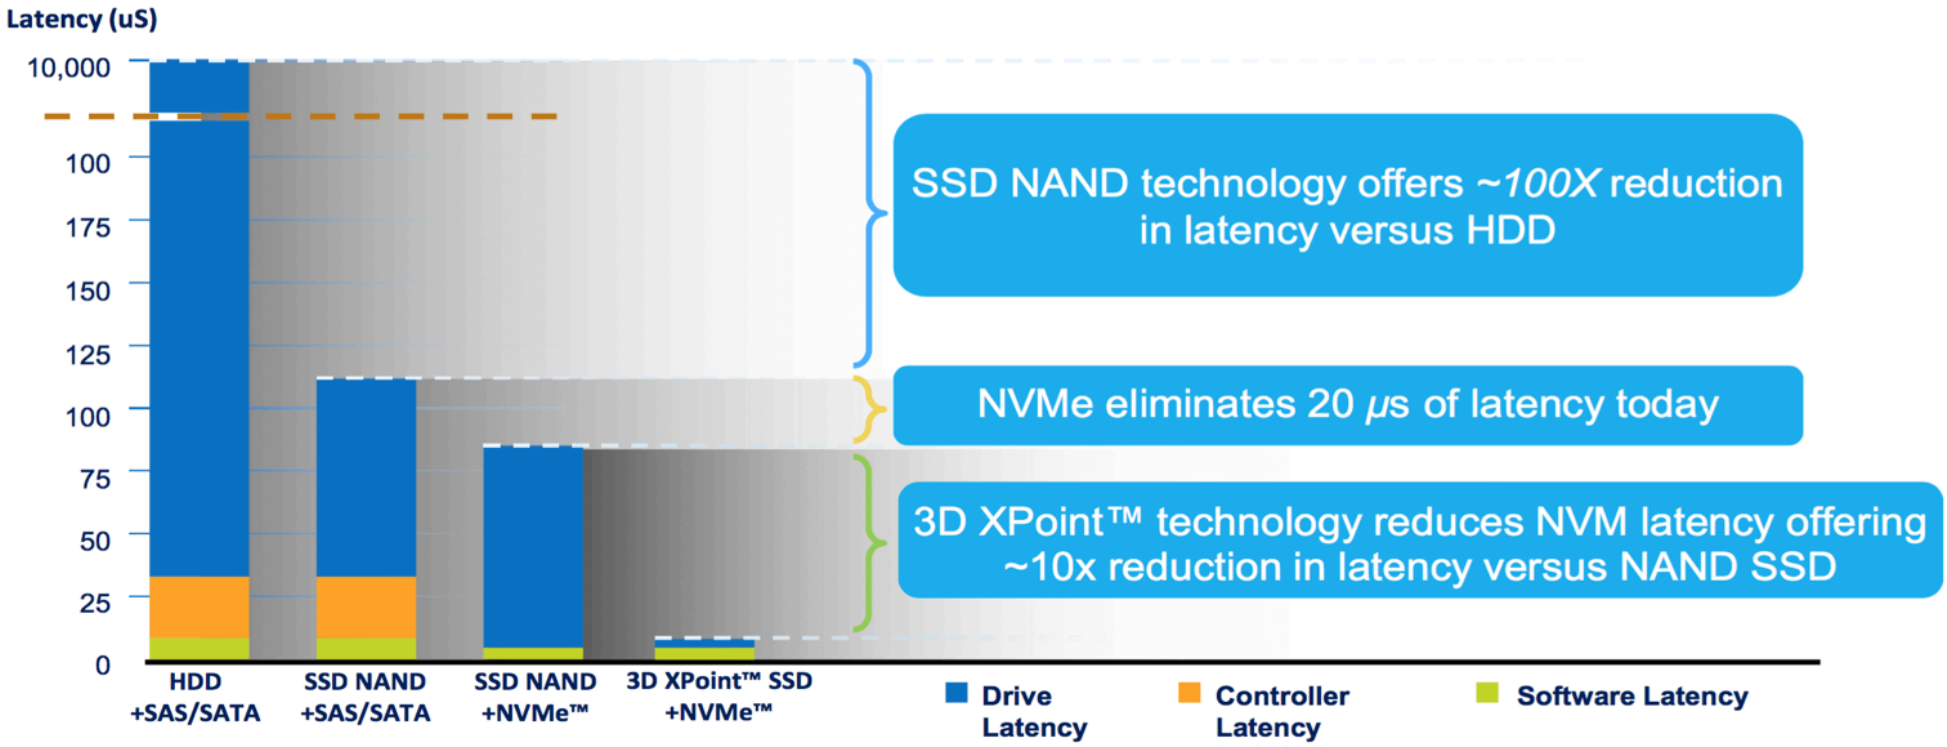

Time to Value & Cost of Latency are Increasing for Batch and "Real-time" Operations

Google Big Table (Circa 2006)

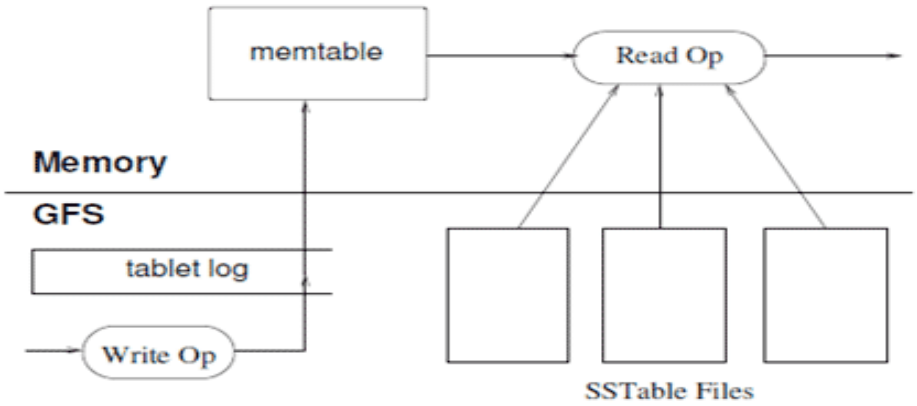

Mid-range Server - 2006

Processor	2 GHz
Cores	2
Storage	400GB HDD
RAM	8GB

Log-Structure Merge (LSM) from 2006 still widely used for NoSQL DB's

Reducing SW Latency and leveraging the Shift towards Persistent Memory Computing

Step 1:

Create a next-generation Big Data Storage Engine that leverages latest HW advances in multi-core, multi-channel flash with a scale-in first architecture which uses parallelism and compact indexing

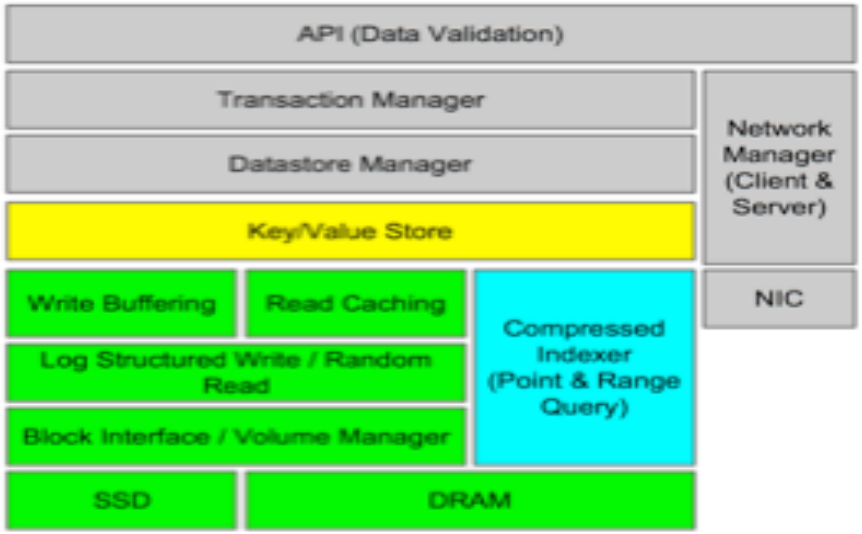

Result: KVS Performance Scales with HW Resources/More Deterministic Latency

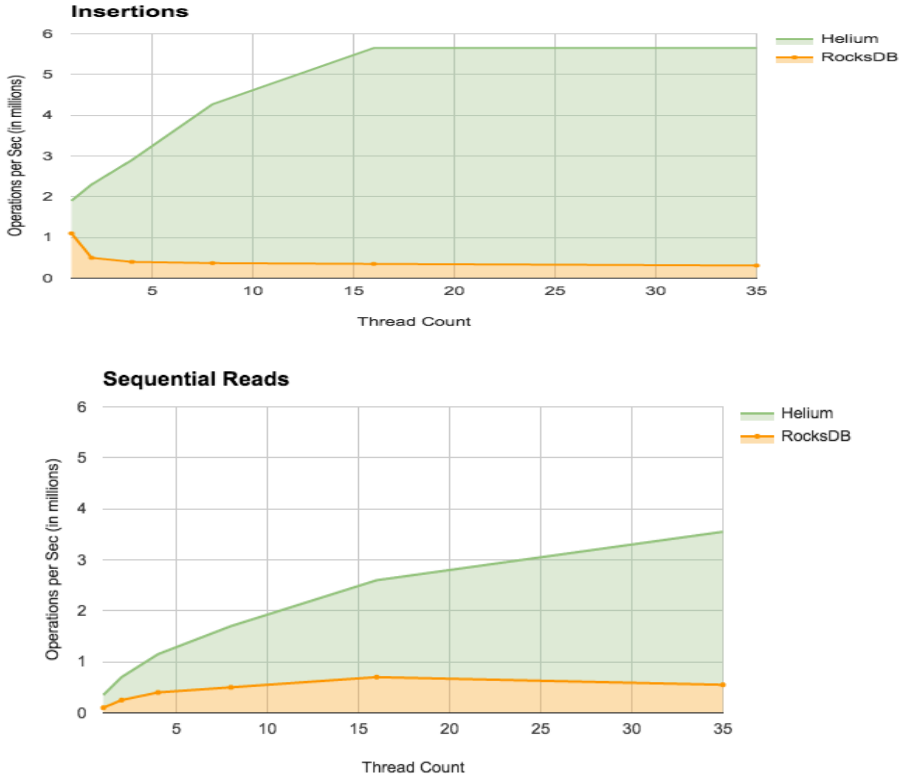

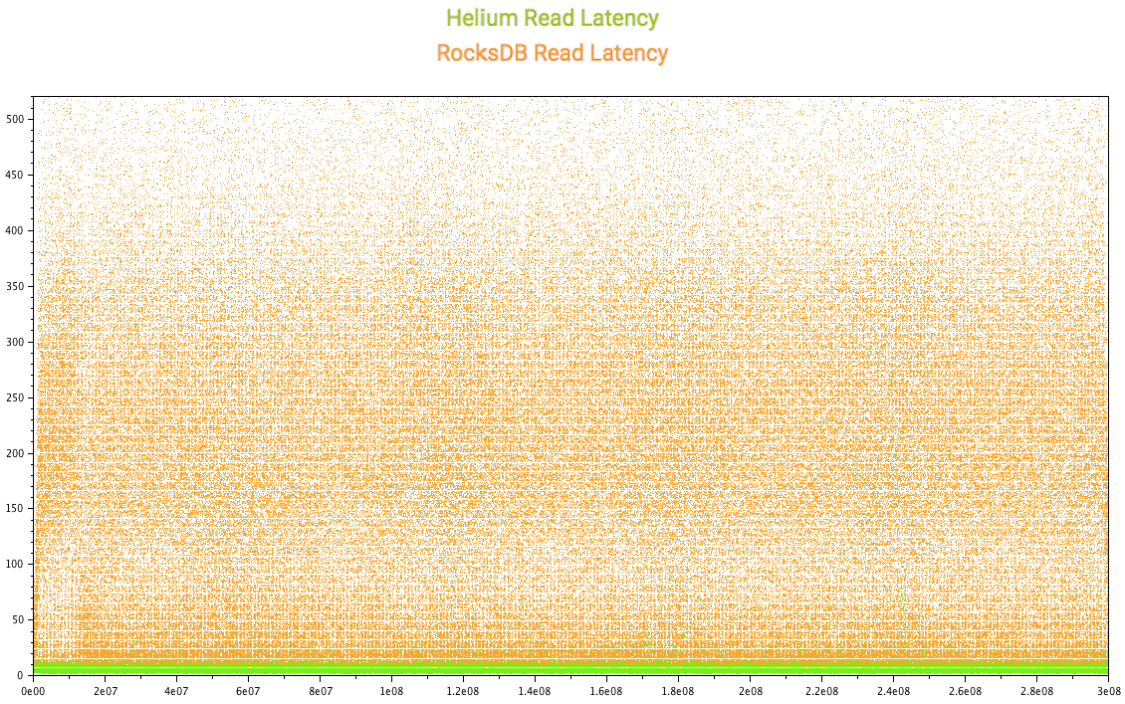

Cumulative Update Operations- Optane

Introduce Persistent Computing to Spark….

- Step 2:
  Leverage advance KVS and add Query/JITC techniques to accelerate queries/analytics.

Scale out and integrate with Apache Spark

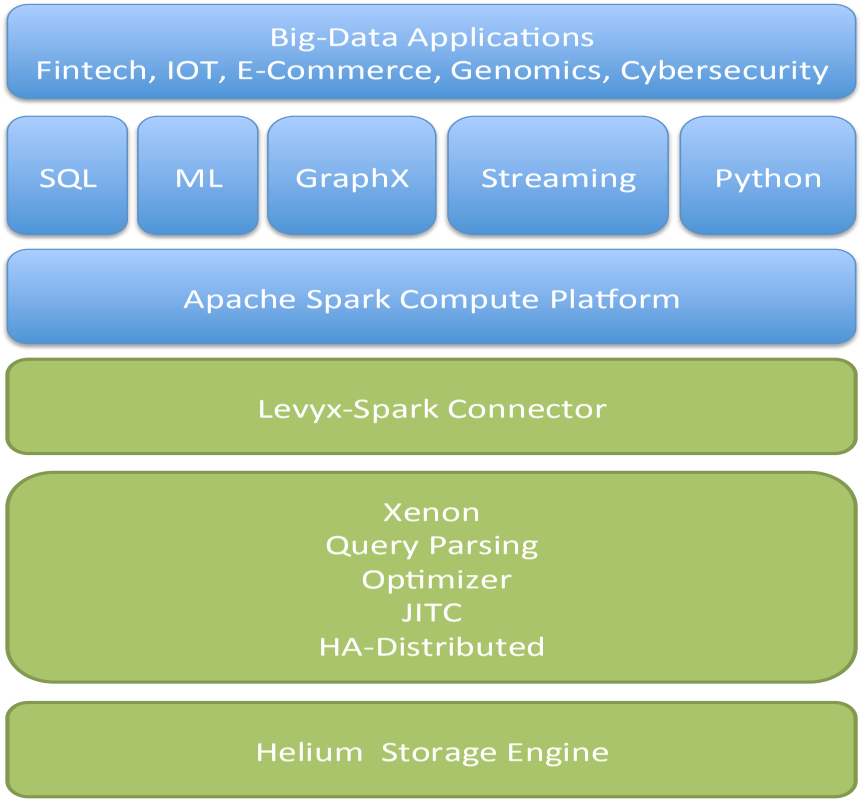

Xenon Use Case:
Financial Algorithm Backtesting Workflow

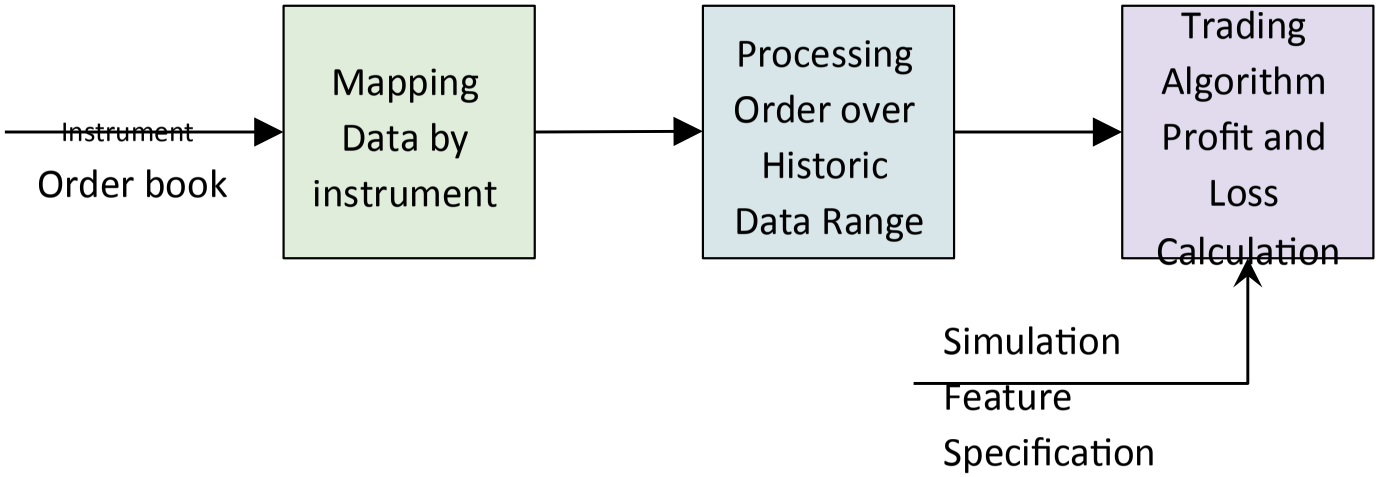

Transforming to Distributed-Parallelized Workflow

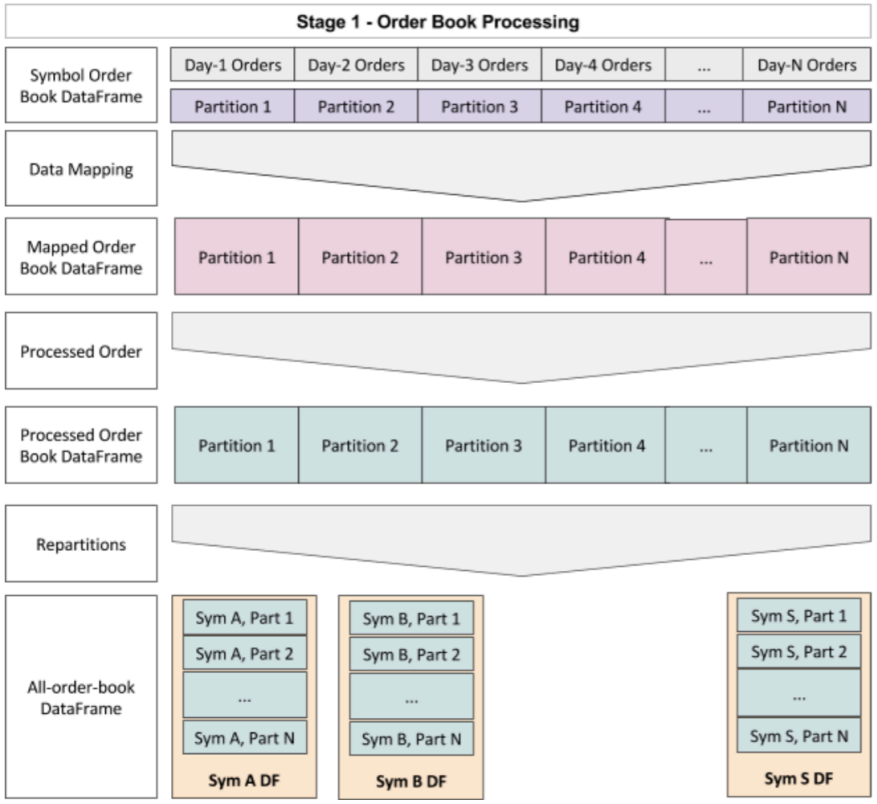

Results- 32X Shortening of Time to Value with less HW resources

- Compared to SUT ID INTC141220-VI
  - –Hadoop Streaming/Python (CDH 5.4.8)
  - –14 x Dell PowerEdge servers
  - –Direct-attached spindle-based storage

	INTC141220-VI	SPRK170603	Improvement
STAC-A3.ß1.SWEEP.MAX60	25	200	8x
STAC-A3.β1.SWEEP.SPEED	0.43	14.08	32.7x
Max instrument simulations/second	2.5	12.6	5.0x

Only 4 (Spark-Xenon) vs 14 nodes (Hadoop) required

Full Benchmark Report Available at : https://stacresearch.com/levyx

Financial Backtesting Using Optane

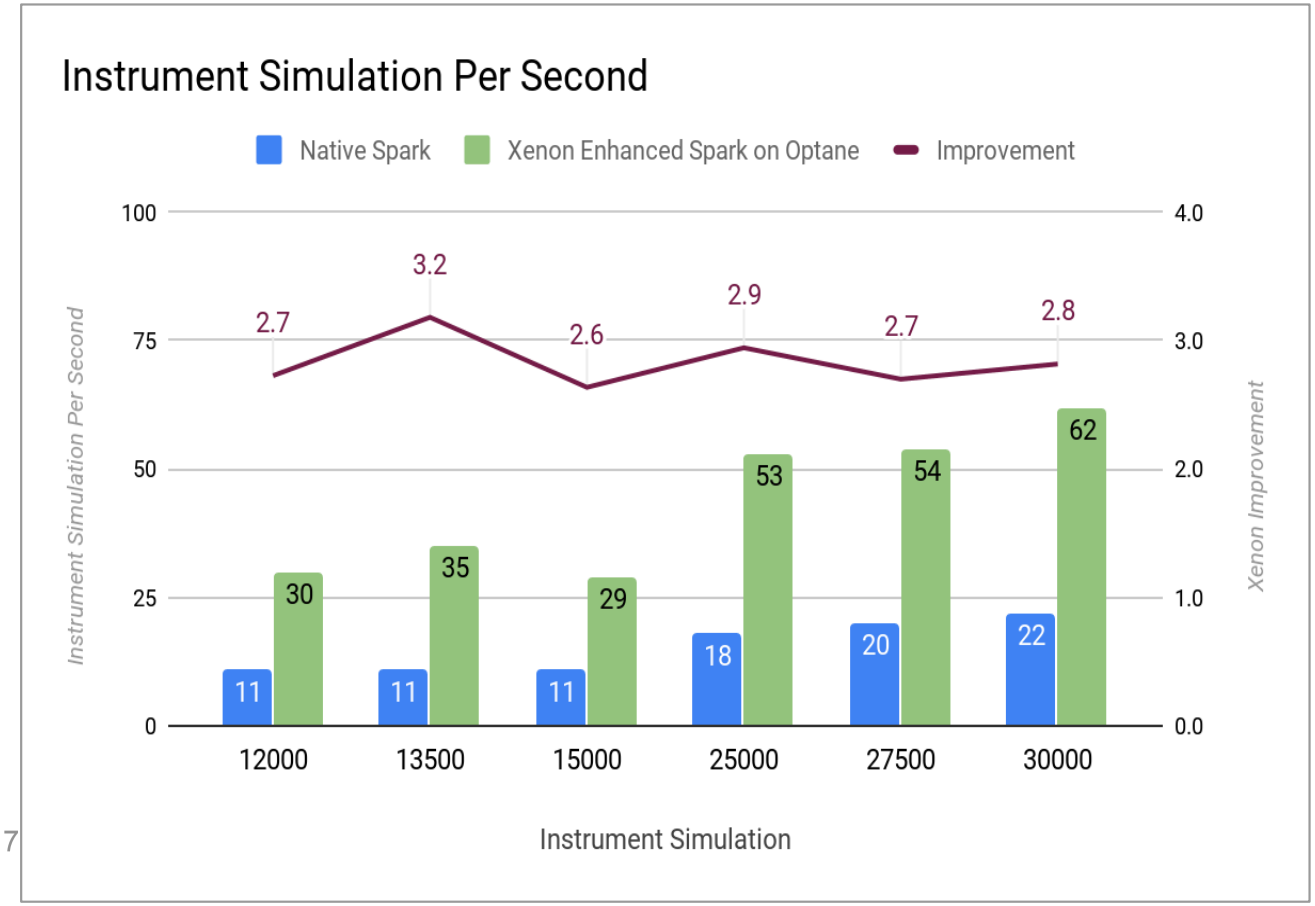

Summary

- Workloads such as Financial Algorithm Back-testing can be accelerated and cost reduced simultaneously by:
  - Re-factoring the storage engine layer to take full advantage of latest multi-core and multi-channel Flash and Optane technologies resulting in increased performance/lower latencies.
  - Build and Leveraging Persistent DataFrames and Computing on top of this new engine and integrate into Big Data's popular in-memory compute platforms, such as Apache Spark.

Questions?

Thank You!

Bernie Wu, CBDO, Levyx
bernie@levyx.com