



# Using Flash to attain Ultra-low latency, Hi-performance for Hybrid Transactional/Analytical Processing

**Faster...Simpler...Cheaper..**

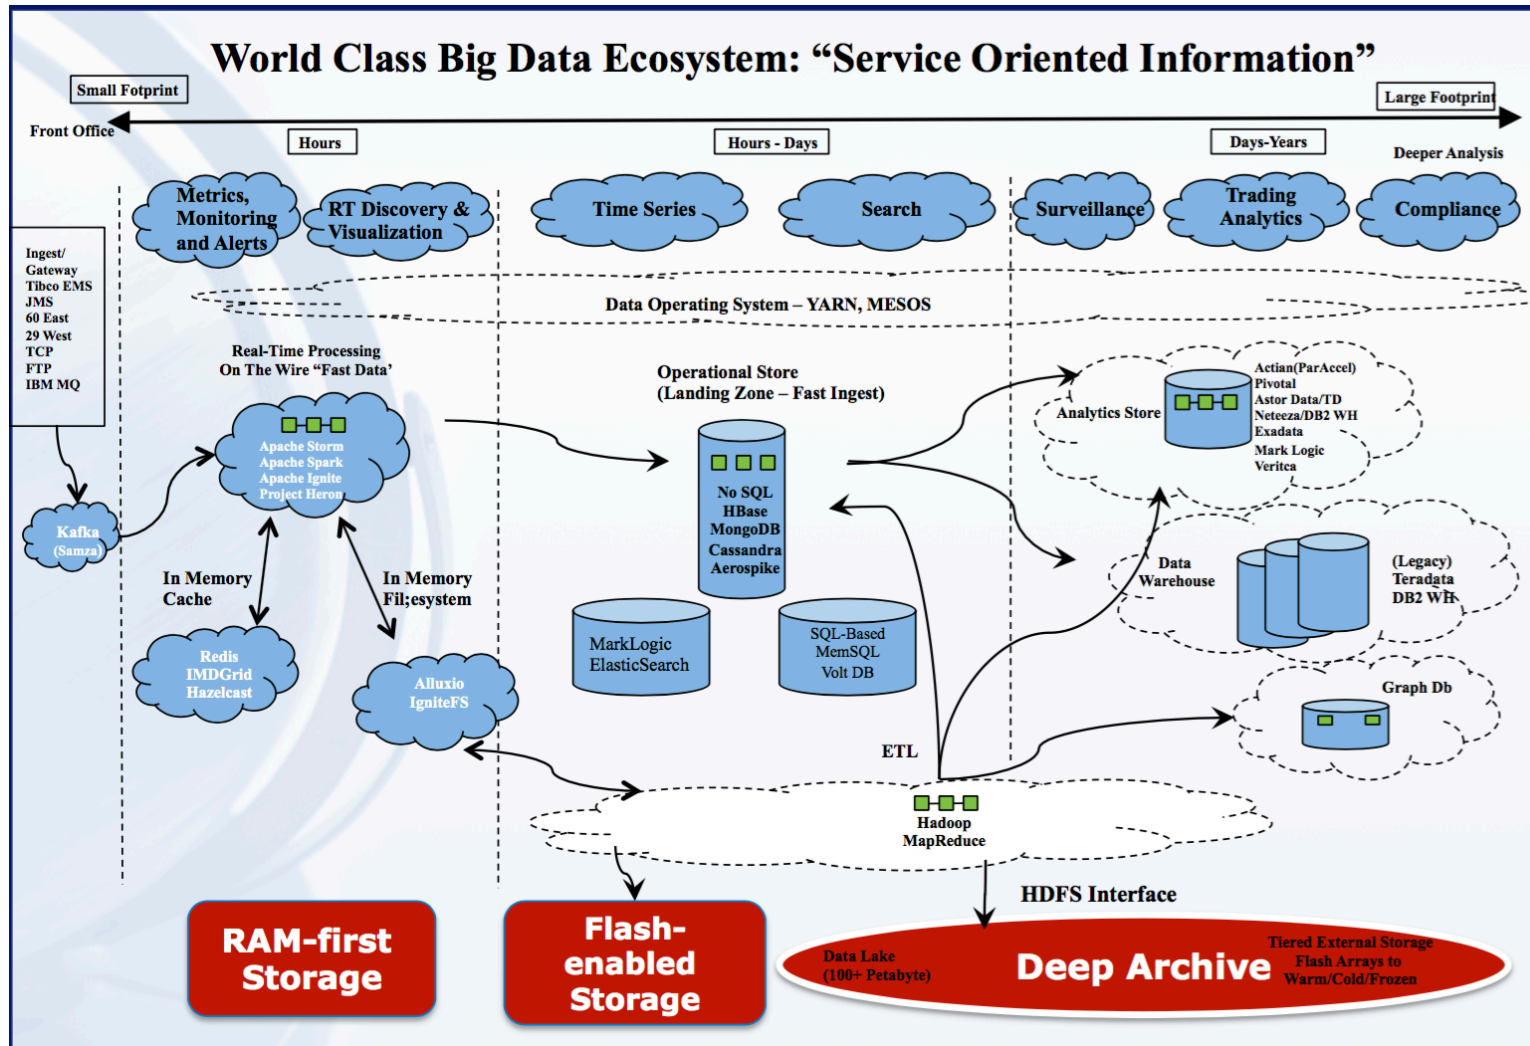
October 2016

# Levyx Company Overview

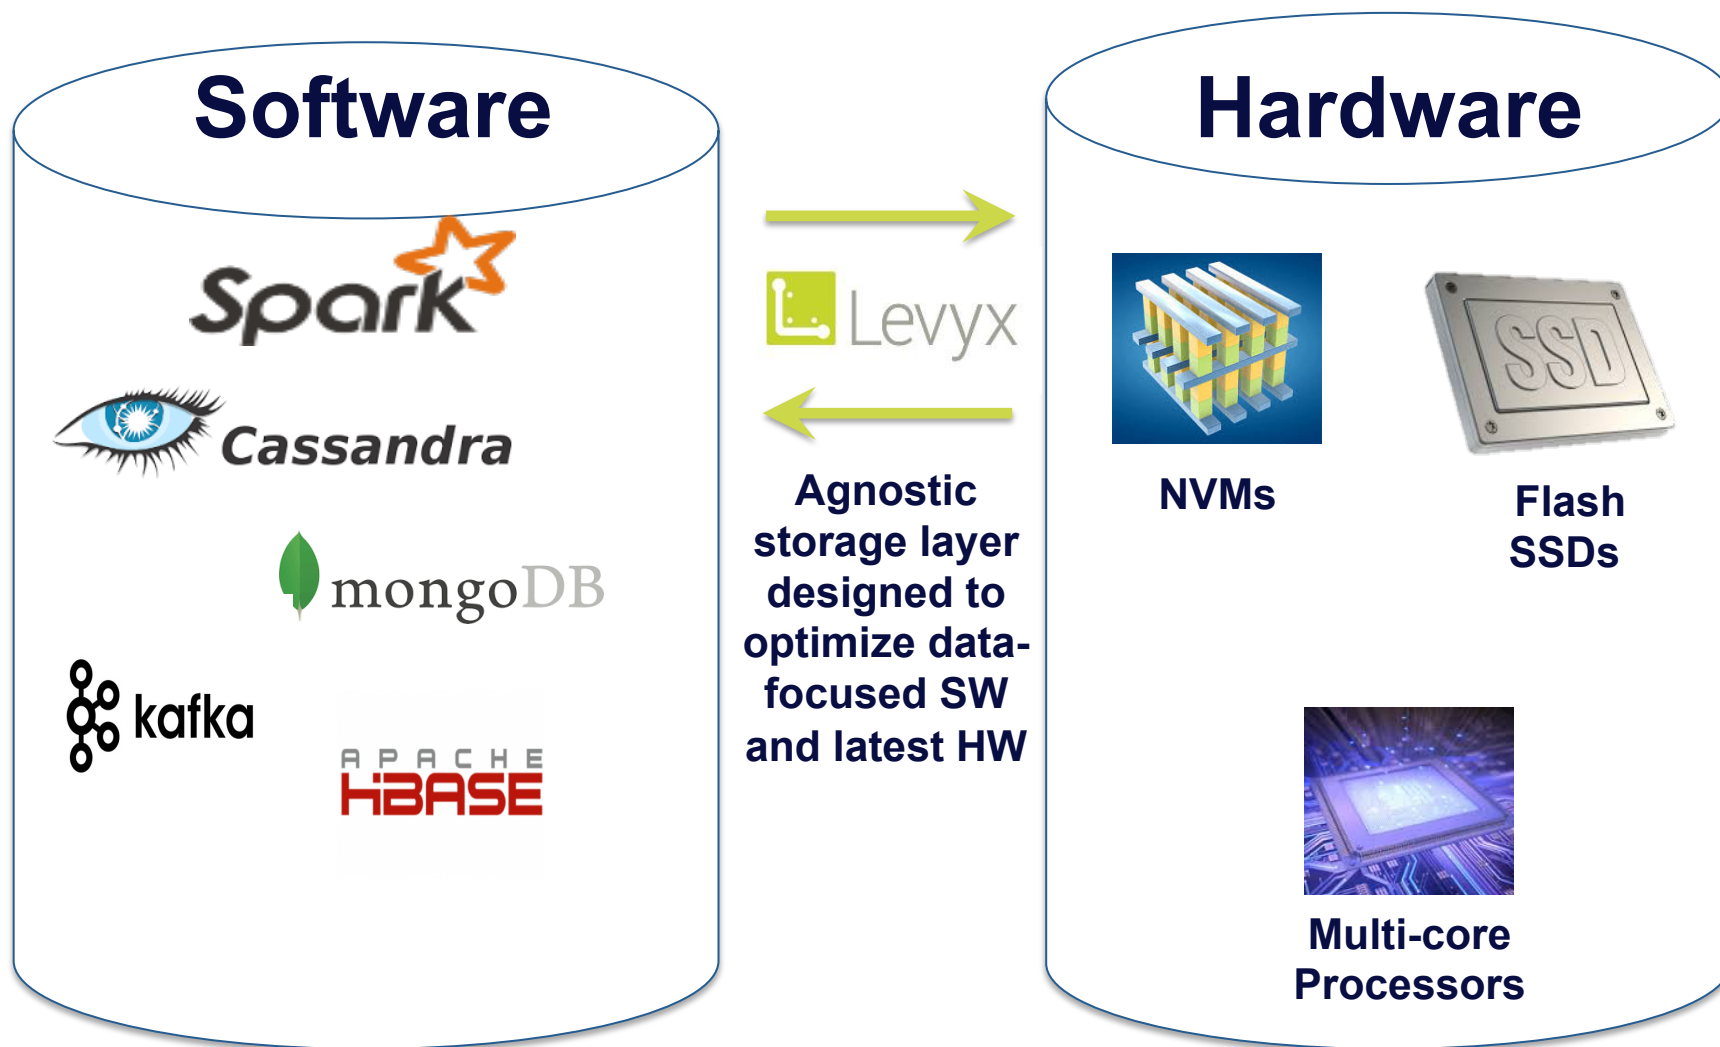
- Mission:
  - Provide Middleware Software that maximizes performance, minimizes latency, and reduces cost/complexity for Big Data and other database scale out server Platforms
- Founded in 2013 , Headquartered in Irvine, CA
  - Reza Sadri, CEO
    - Entrepreneur, PhD UCLA CS. Database specialization
  - Tony Givargis, CTO
    - UC Irvine Professor, PhD CS, Embedded Systems
- Series “A” funding led by OCA Ventures
- One Major OEM (Data Grid use case)
- Patent-Pending Indexing technology



# Trading Analytics Architecture




# Levyx Key-Value Storage Layer Bridges Gap



# Helium:

Leverage Multi-core/Multi-Channel Flash Parallelism to boost performance/  
reduce latency. Reduce layers of abstraction/overhead



# Levyx Products

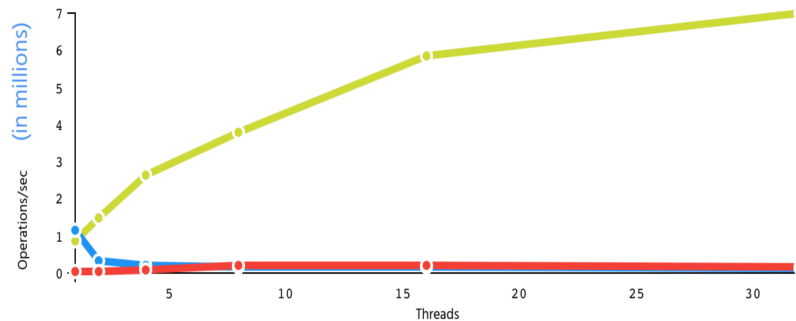
- **Helium: Pluggable DB Storage Engine**
  - World's Fastest Key Value store for Big Data Analytics and Operational Databases
  - Greater than In-Memory Speeds with Scale-in SSD Persistence
  - Native APIs plus wrappers for RocksDB and Memcached
- **Xenon: Distributed Storage and Query Engines**
  - Scales out Helium via distributed architecture
  - Query and Aggregation offload/acceleration using JITC
- **LevyxSpark Connector: Apache Spark + Xenon**
  - Storage Optimized and Accelerated Open Source Spark for real-time/hi IO performance applications
  - Integrated with Spark Catalyst to provide SQL query pushdown (join, groupby, filter, etc) and acceleration to machine code speeds
  - Node consolidation with combined memory-flash storage layer

# Helium vs leading Open Source KV Stores


<http://www.levyx.com/content/helium-demo>

Lexyx Performance scales with # of Cores

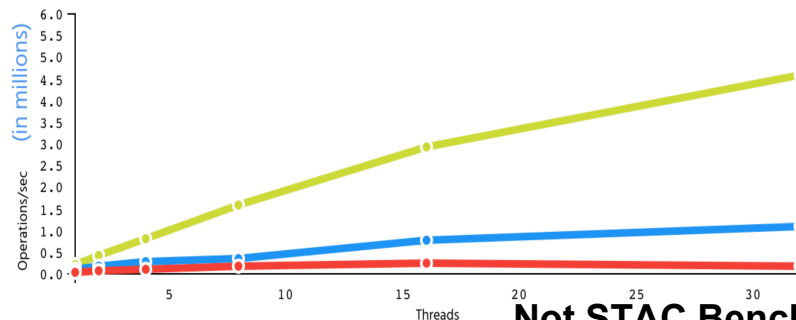
Inserts



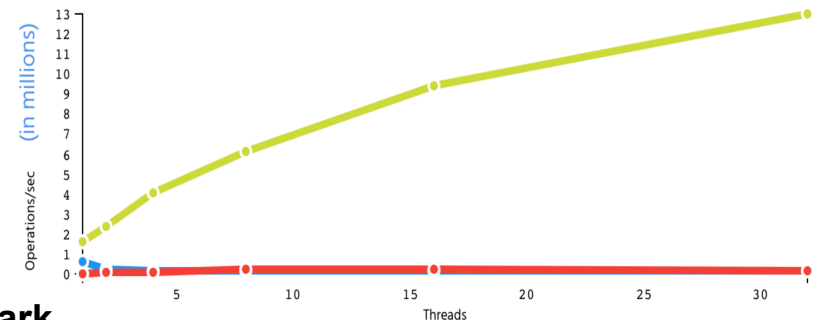
Sequential Lookups



Random Lookups



Deletes



Not STAC Benchmark

# Helium Accelerated Memcached (Open Source)

**Faster** : Standard 90:10 (get:set) Helium-Memcached is at least 10x better in TPS on cloud and on-prem.

**Cheaper** : Single Helium-Memcached scales with cores/SSD vs. stock memcached (needs multiple nodes, large amounts of RAM)


**Simpler**: Plug and Play with existing Memcached applications. Rapid Automatic recovery from persisted SSD simplifies operations.

Small: 100 byte x 100 Million total objects

Configuration	TPS stock memcached	TPS Levix memcached
100% set	359 K	1.5 Million
90% get, 10% set	322 K	7.36 Million

Large: 1K x 10 Million total objects

Configuration	TPS stock memcached	TPS Levix memcached
100% set	276 K	1.32 Million
90% get, 10% set	192 K	2.5 Million




Not STAC Benchmark



# Levyx Spark Connector Benefits on IBM Power 8

- [https://www-304.ibm.com/events/tools/edge/2016ems/REST/presentations/PT/PDF/Edge2016\\_2442.pdf](https://www-304.ibm.com/events/tools/edge/2016ems/REST/presentations/PT/PDF/Edge2016_2442.pdf)

## Iterative Graph Test Bench



Edge 2016



Outthink status quo.

#ibmedge

## Not STAC Benchmark

# Levyx-Spark Connector Benefit Summary

- **Faster:**

- Higher Performance on IO intensive and/or larger data sets and workloads
- Persistent RDD/DataFrame option (reduce disk operations-spillovers, ser-des)

- **Cheaper:**

- Fewer nodes (Freed up CPU Utilization, 10:1 cost advantage of Flash vs RAM)
- Storage-Class Memory architecture

- **Simpler:**

- Persistence Option for RDD/DataFrames – simplify multi-tenant/multi-user operations
- Simpler networks/less shuffling, reduce Spark context switching/scaling/GC/crash issues

# Thank You

Further Inquiries:  
[bernie@levyx.com](mailto:bernie@levyx.com)

Bernie Wu  
Chief Business Development Officer

