



Realtime Persistent Computing on Google Cloud and Intel Optane

June 2017

Matt Meinel

mmeinel@levyx.com

615-972-1259

Levyx's Differentiated Position

Software Delivering on our Mission of **Faster, Simpler and Cheaper**

HELIUM™ ULTRA HIGH PERFORMANCE KEY/VALUE STORE

- Designed to use Optane/Flash and index billions of objects on a single node
- High-level functionality: range-query, search and transactions
- Integrated into major NoSQL and caching solutions and can be easily embedded into new SQL engines, graph databases, search platforms, and streaming platforms

XENON™ SQL / ANALYTICS OFF-LOAD ENGINE

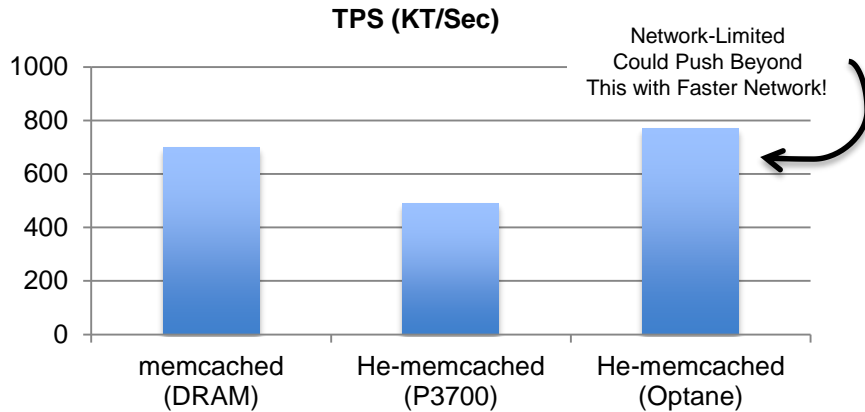
- Performs on-the-fly translation of SQL queries to create a custom data-path that is compiled and executed on bare metal
- Provides large dataset persistence and multi-tenancy
- Connectors that seamlessly integrate the engine with Apache Spark, allowing immediate and automatic offload of Spark operations to SSD or Optane
- Can connect to any major analytics or Machine Learning platform

KEY FEATURES OF LEVYX SOFTWARE ARE IDEAL FOR INTEGRATION INTO EXISTING INFRASTRUCTURE

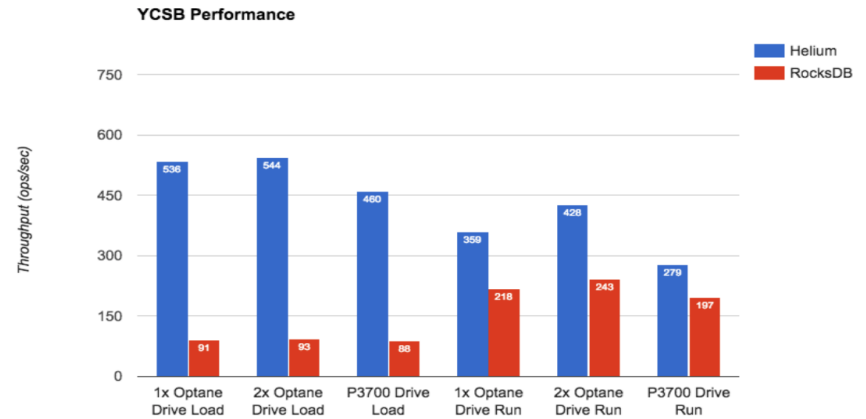
- 1 Easily embeddable
- 2 Provides a distributed storage class memory abstraction for SSD or Optane.
- 3 Immediately optimizes major Big Data and ML applications for SSD or Optane



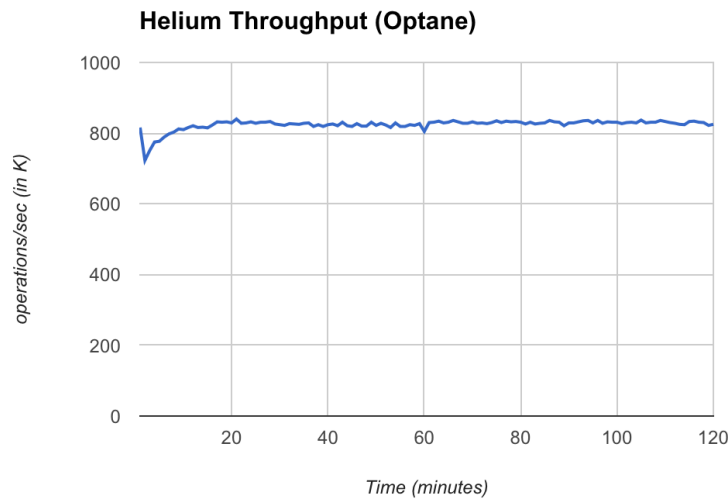
Levyx' Early Intel Optane Test Results Very Strong



Intel Optane + He-memcached: Outperforming DRAM!




Intel Optane + Helium™ vs. RocksDB: Gap Grows!




Helium + Optane = Outstanding Steady-State Performance

Execution Time (Seconds) – Lower is Faster



Optane + Xenon™ Outperforming Apache Spark

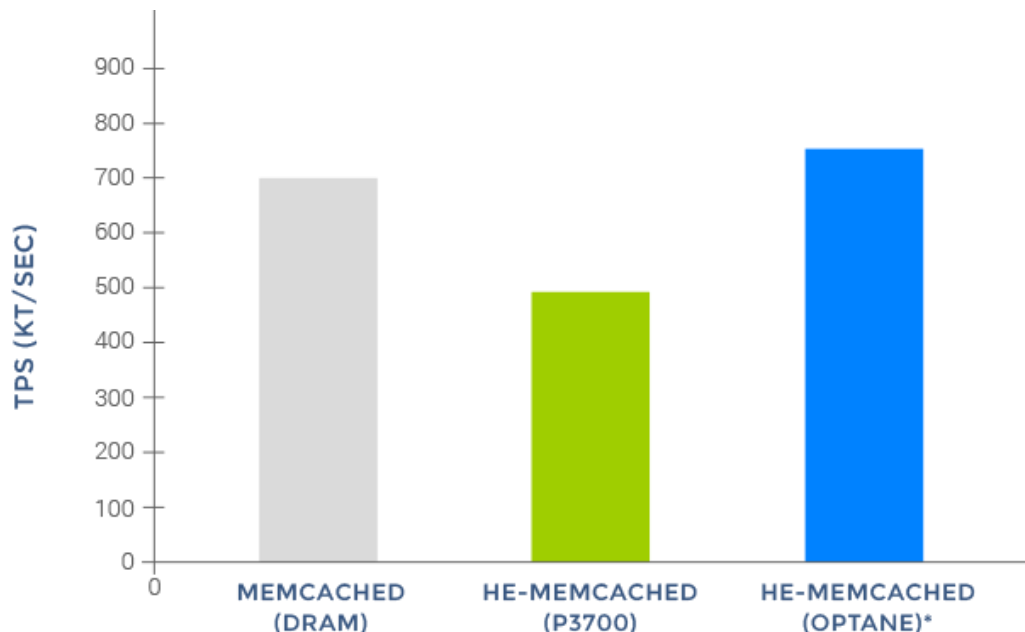
Getting the Most Out of Open-Source SW: *Levyx's Helium Vs Facebook's RocksDB on SSD and Optane*



- Data here shows Helium has 2.5X better throughput than RocksDB and the gap widens to almost 3.0X using Intel Optane!

- The average latency of Levyx running on Intel Optane is over 3.5X lower than running conventional RocksDB

Value Proposition of Helium™ in Key Applications: Vast Improvements Over Standard Memcached



**Network-Limited, Could Push Beyond This with Faster Network!*

Result: Better performance at a lower cost and with all of these additional benefits!

What is Helium™?

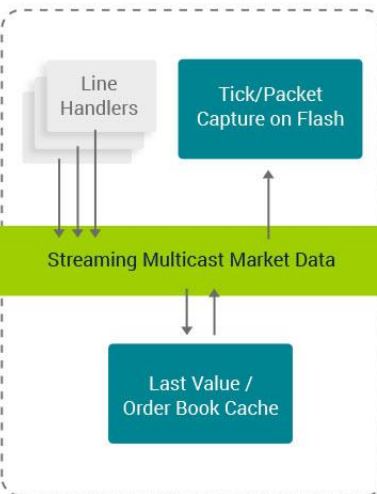
- Helium makes processing huge workloads in real-time faster, cheaper and simpler
- Highly embeddable, high-performance core data processing engine (i.e. KVS)
- Optimized for NVM and multi-core; supercharges open-source caching software
- Example here illustrates how an Intel Optane-based solutions outperforms a conventional DRAM-based solution

✓ Persistence ✓ Lower TCO ✓ Instant Recovery ✓ Sharable ✓ Low Power

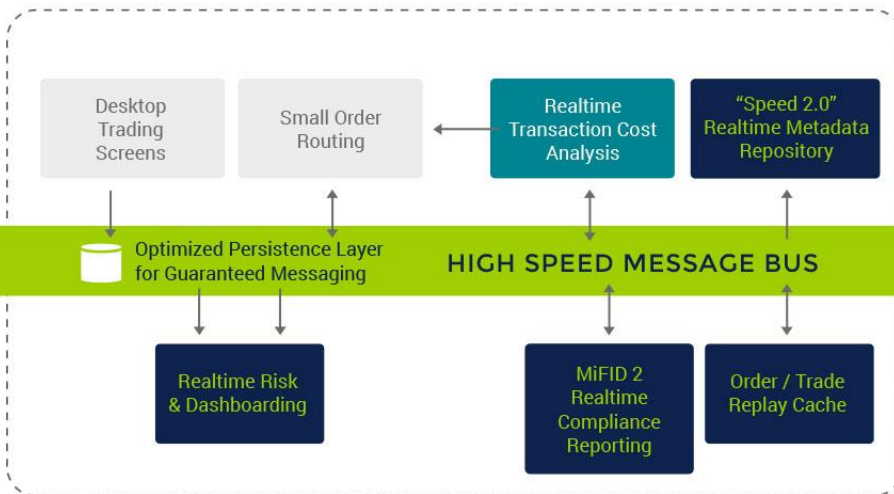
The Financial Sector Use Cases for Levyx

TRADING USE CASES

REALTIME MARKET DATA DISTRIBUTION & CACHING



"SMART" TRADING, GLOBAL ORDER ROUTING, AND TRADE EXECUTION



CLEARING AND SETTLEMENT



■ Helium Opportunities

■ Xenon / Levyx Spark Opportunities

STAC-A3 Backtesting: Spark DataFrames on Xenon

- Levyx Contributed A3 code for Spark with Dataframes to STAC
- Platforms Tested with Spark/Xenon
 - Google Cloud 5 node cluster (4 workers) (STAC Audited)
 - X-IO Axellio single 2U 2 node cluster
 - Dell with Intel P3700's
- Same backtesting technique also used in analyzing:
 - Insurance risk
 - Alternative Credit scores
 - Consumer / Sentiment Analysis
- Although A3 Focuses on 1 year of Tick data with 50 symbols, our lab tests showed excellent scaling for larger numbers of symbols and simulations.
- Having DataFrames persisted in Xenon allows realtime updates to Spark Dataframe in place as well as DataFrames shared between Spark jobs
- Solution: By moving the computation to a cluster including SSDs running Levyx Xenon software, time to value/alpha is reduced by a factor of 10.



- Recap:
 - Workloads that emulate real-world backtesting jobs
 - Range of parallelism and IO/compute intensity
 - Measure speed, scalability, efficiency of any architecture
- Research agenda:
 - Shared storage (e.g., parallel FS) and shared-nothing (e.g., HDFS)
 - Drive insight & improvements at both software and hardware level
- Spark implementation now available:
 - STAC-A3 Pack for Spark rev B
 - Scala implementation originally by Cloudera and Intel
 - Enhanced by Levyx (moved to data frames, changed data layout)

www.STACresearch.com/backtesting

STAC-A3 / Spark with Levyx Xenon on GCP

- SUT ID: SPRK170603
- Stack:
 - STAC-A3 Pack for Spark rev B
 - Apache Spark 1.6.1
 - Levyx Xenon 3.1.0 software
 - Google Cloud Platform 5-node cluster
 - n1-standard-64, 3TB local SSD per node
- Results (STAC Report forthcoming shortly):
 - **STAC-A3.β1.SWEEP.MAX60 = 200 simulations**
 - Max simulations using 50 instruments that could be completed in 60 minutes
 - **STAC-A3.β1.SWEEP.SPEED = 14.08 simulations per minute**
 - SWEEP.MAX60 divided by the elapsed time of the SWEEP.MAX60 sequence
 - **Max instrument simulations/second in scaling tests = 12.6**
 - 100 simulations on 200 instruments in ~26.6 minutes



STAC-A3 / Spark with Levyx Xenon on GCP

- 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.β1.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

Value Proposition of Xenon™ in Key Applications: Large-Scale Real-Time Analytics Using Apache Spark

Welcome to the New Age: Intel can now target major DRAM applications with Optane-based Solutions!

CHARACTERISTIC	DRAM	OPTANE™ DRIVES + XENON™
CAPACITY	• 750 GB	• 750 GB
PRICE OF CAPACITY	• >\$11,000	• \$3,000
PERFORMANCE (TPCH, TPCDS)	• 650 MB/sec Analytics Processing	• 1.3 GB/sec Analytics Processing
TIME TO LOAD	• 89 Seconds	• 1.3 seconds
EXECUTION TIME	• 330 seconds	• 150 seconds
PERSISTENCE		✓
INSTANT RECOVERY		✓
SHARABLE		✓
LOW POWER		✓


Equivalent to Status Quo

Advantage Levyx/Intel

WHAT IS XENON™?

- Low latency, scalable data analytics solution designed to manage the retrieving, processing, and indexing of very large datasets
- Delivers bare-metal, in-memory-like performance using Intel Optane-based solutions
- Example here illustrates how an Intel Optane-based solution compares with DRAM solution, both running Apache Spark

EXECUTION TIME (TPC-DS/H)



About Levyx

MISSION	Deliver software that is hardware optimized to enable Big Data processing in a way that is faster, simpler, and cheaper
ORIGIN	<ul style="list-style-type: none">• Year founded: 2013• HQ: Irvine, CA• Founders:<ul style="list-style-type: none">✦ Reza Sadri, CEO - Entrepreneur, PhD C.S. Database specialization✦ Tony Givargis, CTO, UC-Irvine Professor of C.S., PhD C.S.• 25 total headcount; 15 engineers; 5 PhDs; 5 Eng. Masters
FUNDING	<ul style="list-style-type: none">• Seed (2014)• Series A (April 2016) - \$5.4 M
CUSTOMER MILESTONES	<ul style="list-style-type: none">• First Major OEM (data grid use case) signed (2015)• Currently under eval by OEMs and large customers for IoT, FinTech, and Cloud Infrastructure
IP PORTFOLIO	<ul style="list-style-type: none">• Patent-pending indexing technology• Multiple fundamental patents filed