

AEROSPIKE  
**NEXTGEN**  
**NOW**  
SUMMIT '20

# No data left behind: Ultra-Reliable Low-Latency Solutions for Industry 4.0



**Hewlett Packard**  
Enterprise



**Theresa Melvin**

GLOBAL CHIEF ARCHITECT, BD/AI  
SOLUTIONS

Hewlett Packard Enterprise

# Agenda

- A year in review
- Industrial Revolution: Generation 4
- Use Cases
- Solution Benchmarks
- Calling All Data Scientists ...

# A Year in Review





# 2019 Joint Solution Success Story

## A Leading Digital Payments Provider

- Software (Aerospike Database)
  - Replaced Apache Hbase
  - Beat out multiple competing NoSQL products in PoC
- Hardware (for Aerospike Database)
  - HPE ProLiant DL360 Gen10 NVMe Server
  - HPE Persistent Memory featuring Intel® Optane™ DC Persistent Memory
- Stats
  - Reduced database cluster node count by 66%
  - Tripled dataset size
  - Meeting SLA of 99.99% uptime
  - Meeting 10-20ms payment processing window



# Gen 4 Industrial Revolution

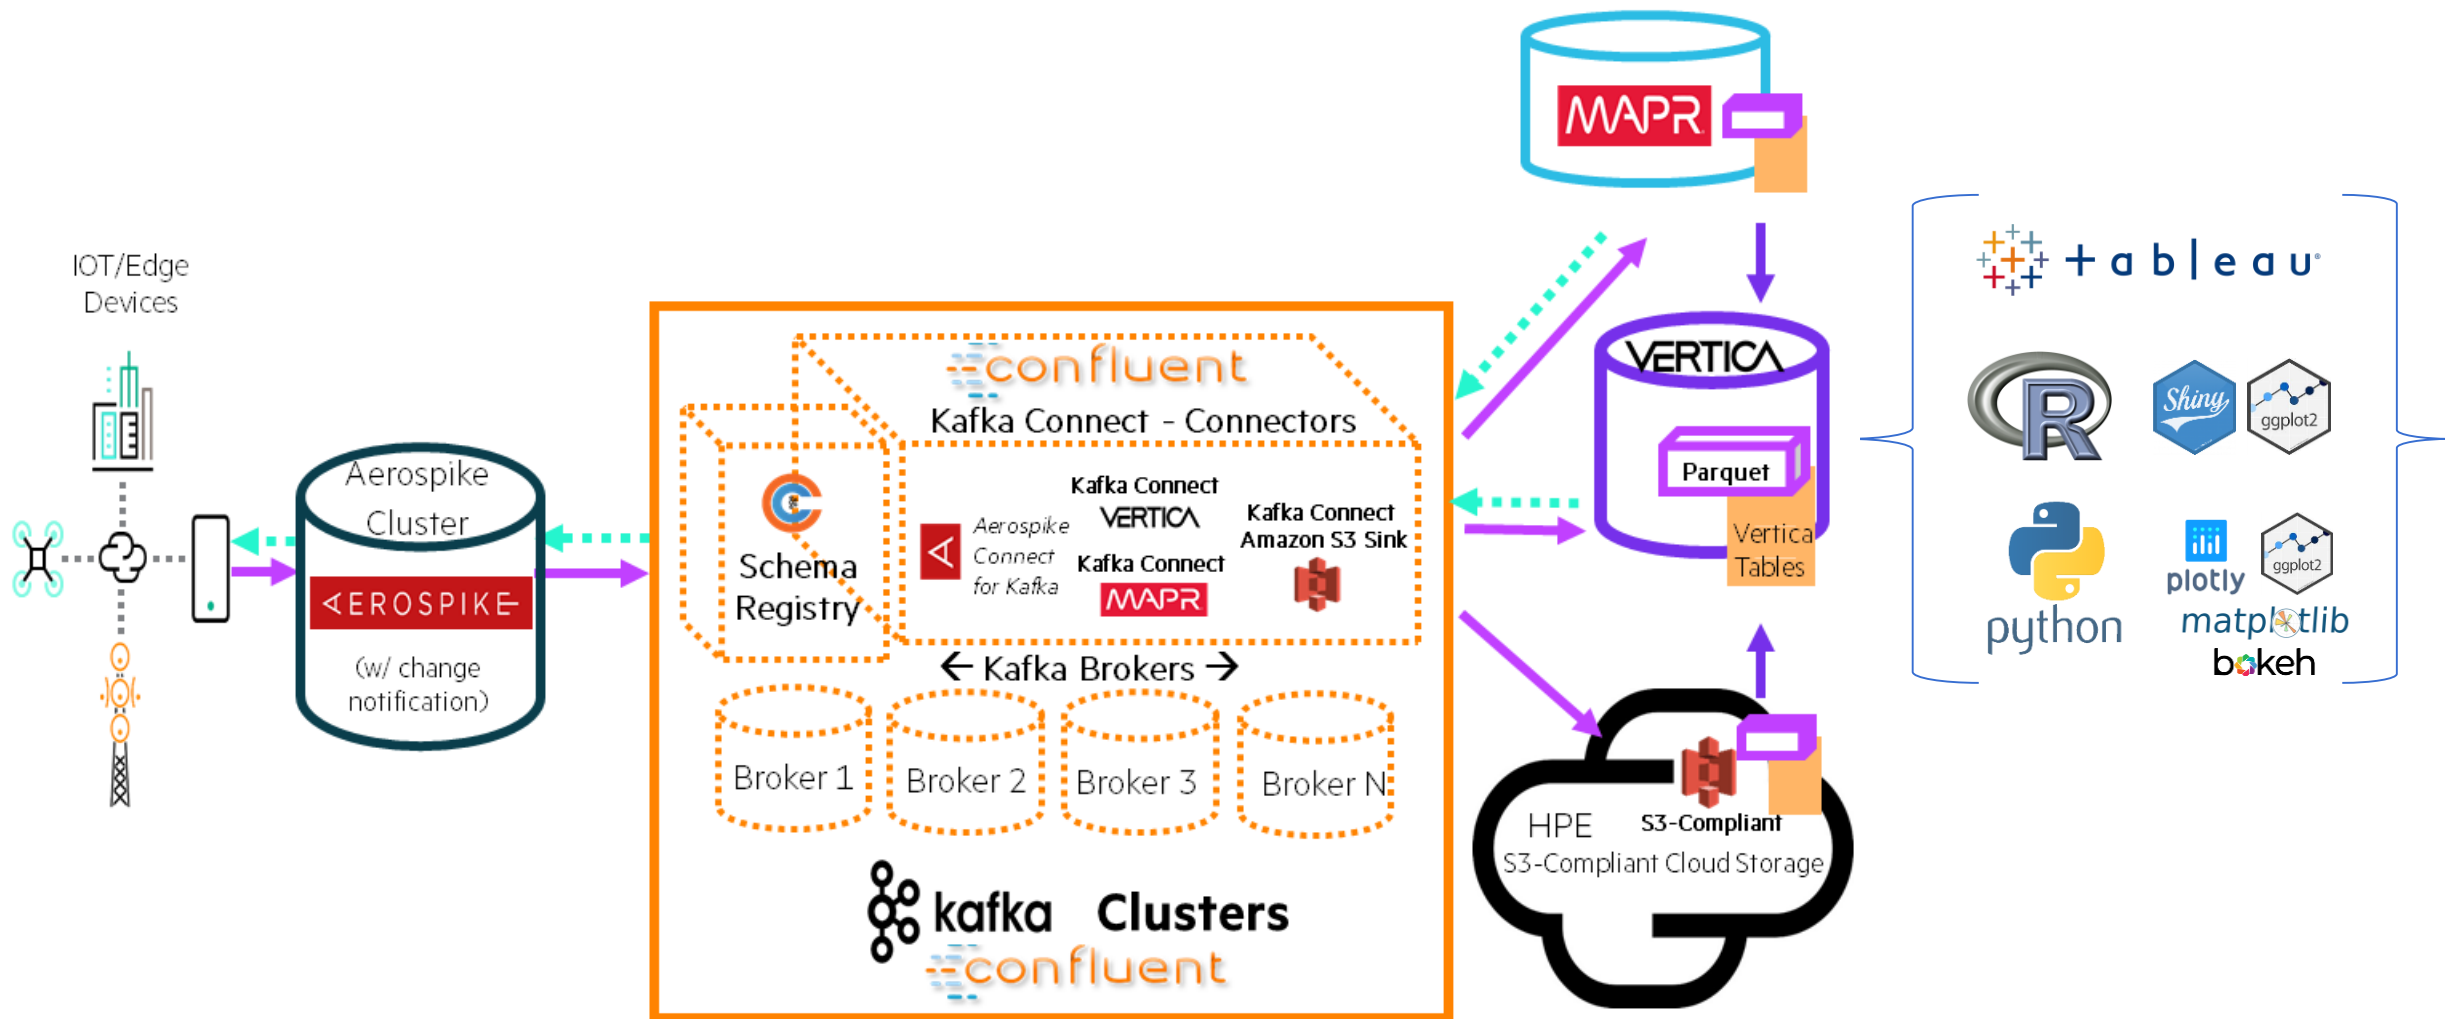


## Industry V4 Use Cases





# Industrial V4 Real-Time Lookups Using Aerospike



# A Desperate Need for Accurate Real-Time Weather Forecasting

- Many weather sources still provide their forecast updates using email updates and PDFs, very few sources offer real-time or near real-time updates
- Comprehensive weather forecasts often lack granularity, making it difficult to predict potential local weather events with any degree of accuracy, that could significantly impact a town's power grid



National Oceanic and Atmospheric Administration (NOAA) National Centers for Environmental Information (NCEI) New Data Volume Reporting Method reveals a historical weather archive of **32.6 petabytes** in 2019. NOAA's new Environmental Real-time Observation Network (NERON) will also establish a nationwide network of high quality **near real-time weather monitoring** stations across the United States. *Source NOAA/NCEI*



## Time is Money in the Energy Market

- As global markets become more interconnected, the need for a comprehensive global view of the latest forecasts, in real-time, is critical to reducing the time needed to make trading decisions and making smarter, more profitable trades
- Renewable energy sources can serve a great deal of demand; therefore, being able to accurately predict weather events is critical to taking advantage of real-time market prices



Renewable energy sources, like wind and solar, are much more **variable** than traditional energy sources, and being able to accurately predict weather events is key to taking advantage of real-time energy market prices

## Employee Assistance using Commercial Drones

- There are numerous benefits from augmenting the work force with commercial drone technology, and routine or basic inspections top this list
- Unmanned Ariel System (UAS) Project cost savings are capable of generating a quantifiable Return On Investment (ROI) within 2-3 years




Drones performing daily flights each week can average 500 images per flight, which equates to 10s of gigabytes each week in just still camera imagery. Photogrammetry operations, in turn, generate many-millions of 3D data points.



## Routine Bridge Inspections

- A task that takes a human as long as four hours to complete, a drone can perform in one hour
- Safety protocols require at least two people for every bridge inspection, these costs could be recouped using drones generating a considerable annual cost savings




Approximately 83% of bridges are inspected every two years, 12% are inspected annually, and 5% are inspected every 4 years  
*Source: US Department of Transportation*



## Basic Rail Road Track Inspections

- Safety is one of the most important benefits realized by using drones for basic track inspections
- Placing human inspectors on tracks used by regularly scheduled trains can pose a considerable danger to human life
- Drones mitigate this risk by flying above the train when inspecting the track



**GE's Evolution™ Series Tier 4 locomotive features more than 200 on-board sensors that monitor performance and health and process more than 1 billion instructions per second**

*Source: GE*

# Wind Turbine Blade Inspections

- Onsite inspections currently required, which rely on rope access technicians and ground-based cameras to check for cracks, tip erosion, delamination, and other issues.
- Drones could considerably reduce routine inspection climb rates, which in turn would significantly cut human inspection costs.



1 Turbine	50 Turbines
>100 Sensors	>5,000 Sensors
33.33 TPS	1,666 TPS
143 GB/day	7.2 TB/day
52.1 TB/year	372 TB/year

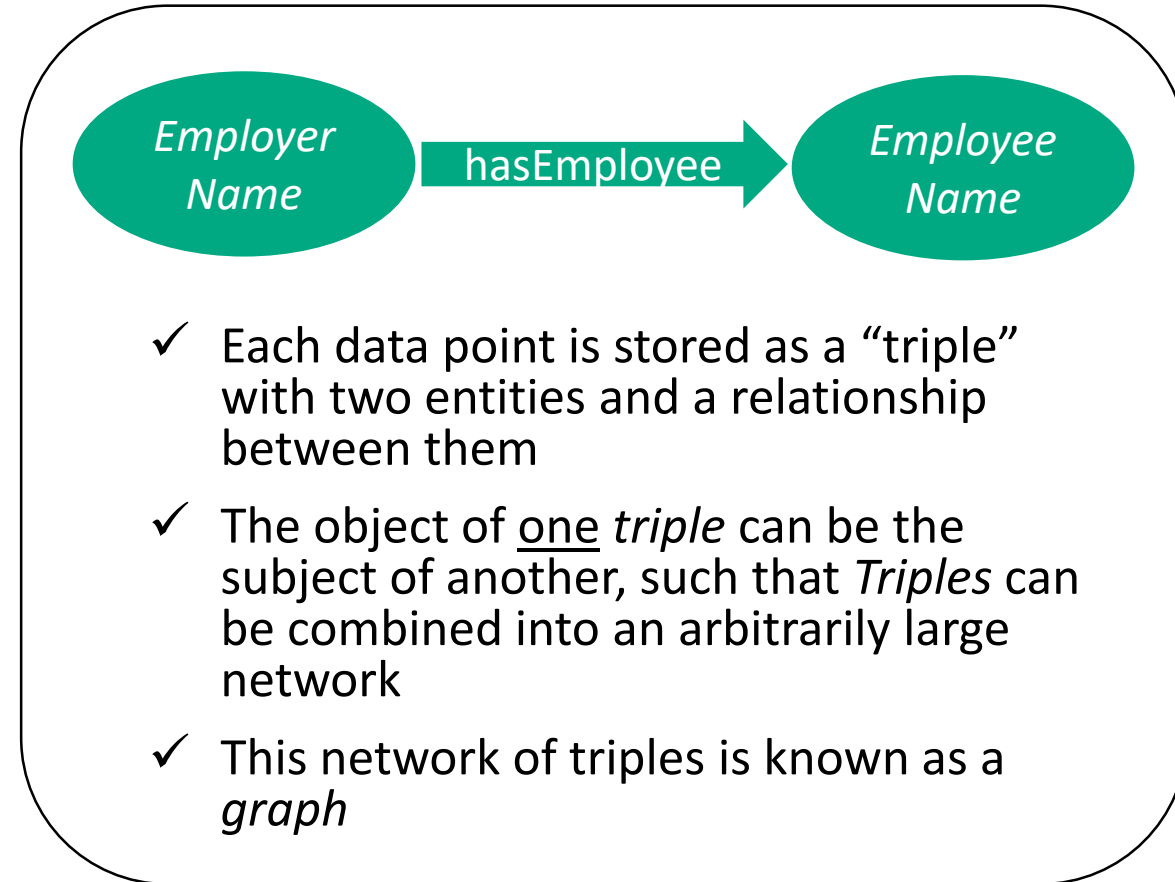
## Industry V4 and the Semantic Web





# The “Semantic Web” Concept

- An alternative method for linking and storing data
  - Semantic web technology organizes interconnected data as a network of meaningful relationships between entities (people, places, businesses, etc.) rather than as a table with variables and metadata
  - This format allows computers and people to quickly interpret the meaning of a relationship



## Benefits of Semantic Web Technology & Why Industry 4.0 will Leverage it

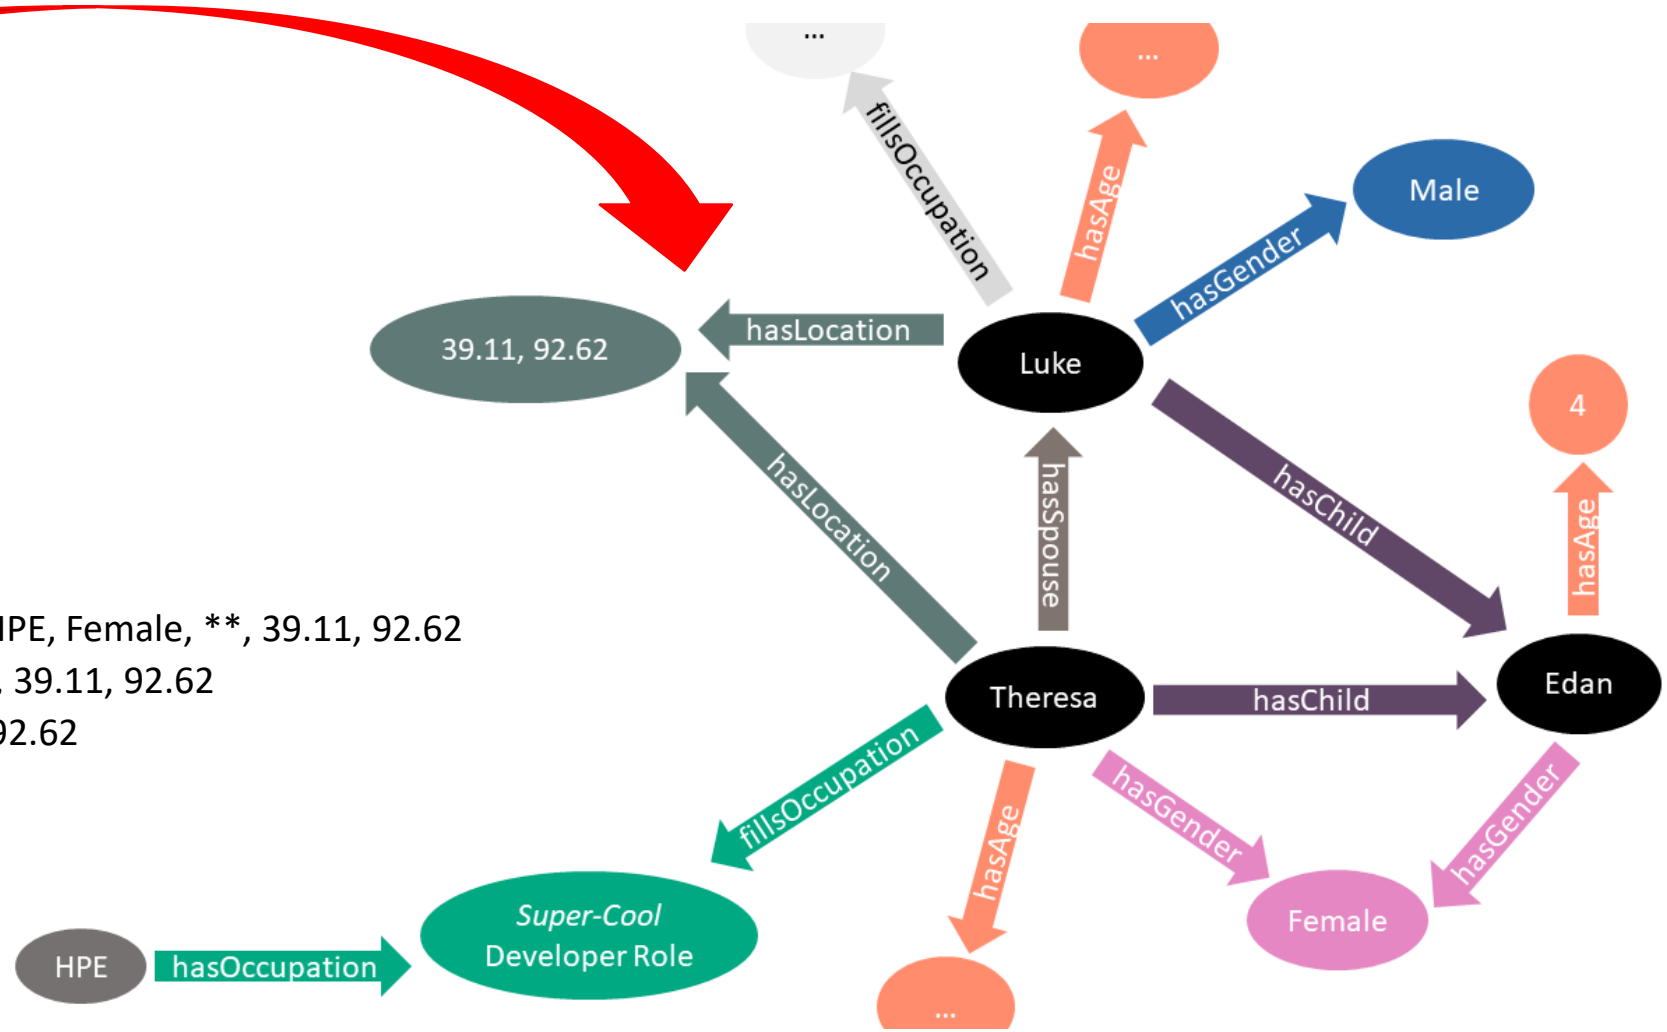
- Fast-flexible data structure
- Machine reasoning and inference against the dataset
- Shared understanding of the data's meaning
- Reusable classifications and standards
- Easy exploration of many dimensions
- Network analysis

# Aerospike's Role in Next-Gen Network Graphs for Industrial 4.0 Use Cases



AEROSPIKE

- People (Key)
  - (Values)
    - Theresa, Super Cool Developer, HPE, Female, \*\*, 39.11, 92.62
    - Luke, Analyst, US Army, Male, \*\*, 39.11, 92.62
    - Edan, NA, NA, Female, 4, 39.11, 92.62
- Company (Key)
  - (Values)
    - HPE, Super Cool Developer
    - US Army, Analyst







# New Benchmarks



# A Collaborative Effort by Three Companies

- The Benchmark Team



# Massive Transactional Workload Scale with Aerospike 4.9 and HPE Persistent Memory featuring Intel® Optane™ DC Persistent Memory DIMM Technology





# Aerospike with HPE Persistent Memory featuring Intel® Optane™ DC Persistent Memory

## The Tested HPE ProLiant Server Hardware – PMEM Servers (x3)

### ■ CPU

- 2 x 28-core **Cascade Lake** Procs
- CLX SP **28c 2.5GHz** 205W
- L1 = 1792 KB, L2 = 28672 KB, L3 = 39424 KB

### ■ Memory

- 2.6 GHz
- Total = 1.75 TB (2x12 Slots)
  - RDIMM = 16.00 GB x 6
  - 96 GB \* 2 = 192 GB Total
  - DCPMM = 126.38 GB x 6
  - 758.16 \* 2 = **1516.32 GB**

### – Network

- Adapter 1 / LOM
- HPE Eth 10/25Gb 2p 640FLR-SFP28
- Adapter 2 / 1GbE
- HPE Ethernet 1Gb 4-port 331i Adapter
- Adapter 3 / **100GbE**
- HPE InfiniBand EDR/Ethernet 100Gb 2-port
- Adapter 4 / 10GbE
- HPE Eth 10/25Gb 2p 621SFP28

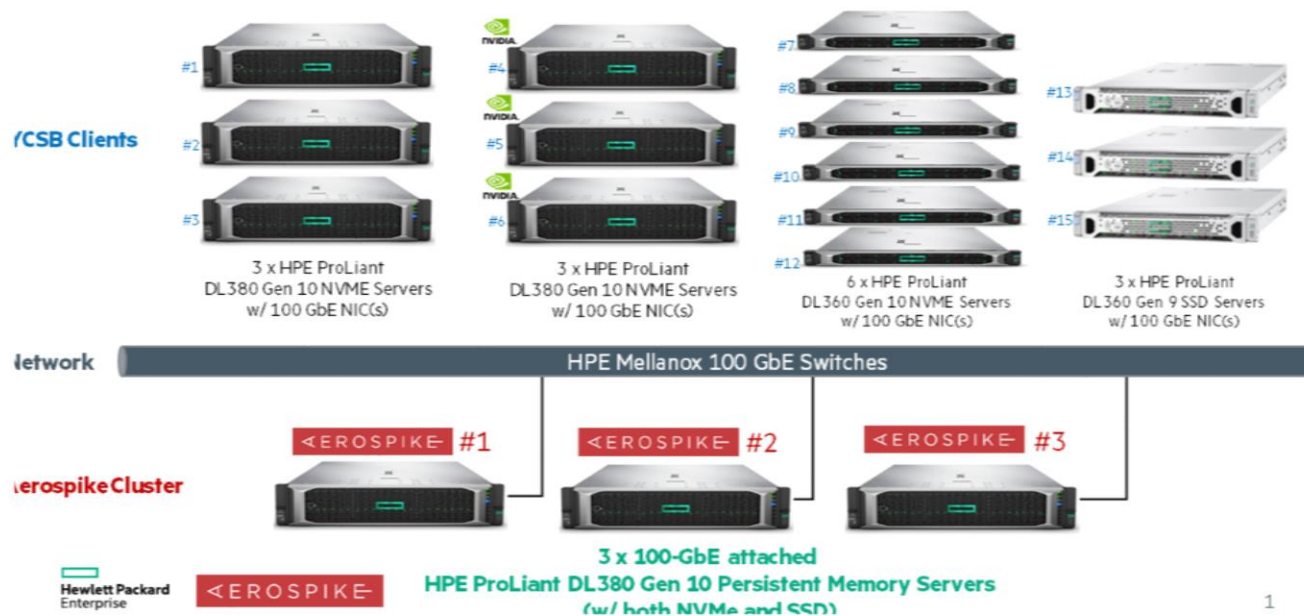
### ■ Storage

- NVMe Controller
  - 3 x 1600GB SSD NVMe
- HPE Smart Array P408i-a SR Gen10
  - 2 x 400 GB, 1 x 480 GB SSD



1

## Aerospike's YCSB Benchmark to HPE PM Servers Reference Architecture



1


## Aerospike PMEM Benchmarks

- **2019** PMEM Benchmark  
*Unsupported Test with Aerospike 4.5*
  - Indexes and Data all stored in PMEM
  - 8 Million TPS
- **2020** PMEM Benchmark  
**Fully Supported Test with Aerospike 4.9**
  - Indexes and Data all stored in PMEM
  - 14.87 Million TPS
    - Tibor: *We possibly could have hit 18 Million TPS with more client servers*

Nice Performance! But seriously, who really needs performance like this?

## *Building Next-Gen [Quantum] Machine Learning Pipelines*

- HPE Cray will provide [NOAA](#) with Two AMD-Powered Supercomputers
  - EACH Supercomputer will have a theoretical peak performance of 12 petaflops
  - NOAA's total computational power for prediction and research will be ~40 PFLOPS



*More and more supercompute-customers are looking to next-gen pipeline technologies to feed this massive Supercomputers data in real-time in order to generate even more timely and accurate predictions*

# Significant Throughput Achievements for Aerospike 4.9 using ScaleFlux Computational Storage Drive





# Aerospike ACT Benchmark Comparison Against Traditional NVMe SSD and ScaleFlux Computational Storage Drives

## The Tested HPE ProLiant Server Hardware – ScaleFlux Server

- **CPU**
  - 2 x 4-core **Intel® Xeon® Gold 5122** Procs
    - CLX SP 4c 3.6GHz 105W
    - L1 = 256 KB, L2 = 4096 KB, L3 = 16896 KB
- **Memory**
  - 2.67 GHz
  - Total = 96 GB
    - RDIMM = 16.00 GB x 6
- **Storage**
  - NVMe Controller
    - 1 x 4TB **ScaleFlux CSD2000**
    - 1 x 1600GB SSD NVMe
  - HPE Smart Array P408i-a SR Gen10
    - 6 x 600 GB HDD
- Network
  - HPE InfiniBand EDR/Ethernet 841QSFP28 100Gb 2-port



HPE ProLiant DL360 Gen10 Server



2

## Evaluating Scaleflux RA

AEROSPIKE

Aerospike Server



1x HPE ProLiant DL360 Gen10 NVMe Servers w/ 100 GbE NIC(s)



Drives



1x 1.6TB NVMe PCI-E X4 Mixed



1x Scaleflux CSD 2000




AEROSPIKE

2

# Turnkey Computational Storage Drive Solutions by ScaleFlux


- Complete, validated solution
  - Pre-Programmed FPGA
  - Hardware
  - Software
  - Firmware
- No FPGA knowledge or coding required
- Field upgradeable
- Enterprise updates
- Standard U.2 & AIC form factors

## ScaleFlux CSD 2000



ScaleFlux has a single FPGA which combines Compute and SSD functions

## Traditional SSD

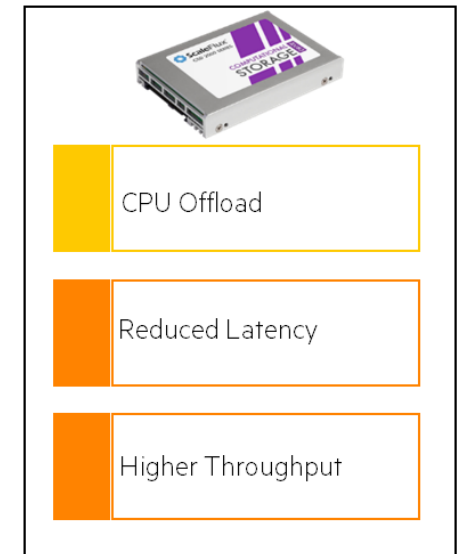
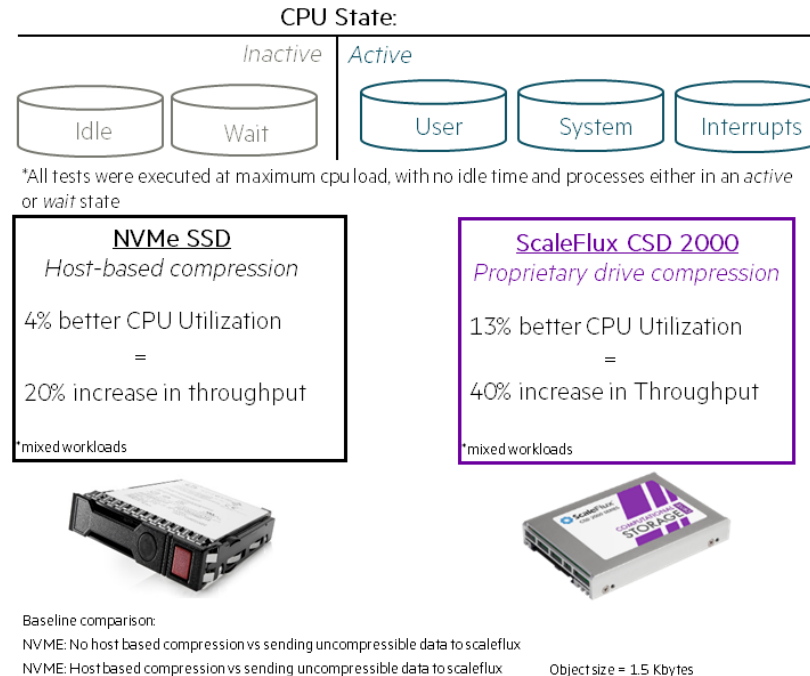


Many different components required in order to support both Compute & SSD functions

# NVMe vs. ScaleFlux Compression Performance using Aerospike's ACT Benchmark Tool

## Test Results

- ScaleFlux observed a **3X** improvement in CPU Utilization over NVMe SSD
- ScaleFlux afforded **>2X** Throughput Increase over NVMe SSD
- An **80%** reduction in latency was observed using ScaleFlux Compression over NVMe SSD host-based compression





## Cools Stats from Aerospike - ScaleFlux Benchmark

- NVMe drives
  - Not enormously parallel
  - >50 threads creates latency
  - Queue size peaked at ~15
- ScaleFlux CSD 2000
  - Highly efficient parallel operations
  - >160 threads before latency incurred
  - Performed great with queue sizes 40<>50
- 2xNVMe SSD drives equaled the performance of 1 ScaleFlux (SF) drive
- SF Read Only (RO) Workload averaged ~900,000 Operations Per Second (OPS)
- SF Write Only (WO) Workload averaged ~400,000 OPS
- RO workload = ~+10% with SF
- Mixed workload = ~+20% with SF
- WO workload = ~+25% with SF
- At 40 threads
  - NVMe latency was in the 95th percentile
  - SF latency was sub-millisecond

# The Global Data Science Community is Called Upon to Help Fight COVID-19



## COVID-19 White House Briefing on March 16, 2020

- The [White House Office of Science and Technology Policy](#) issued a “Call to Action to the Tech Community on New Machine Readable COVID-19 Dataset”
- The [COVID-19 Open Research Dataset \(CORD-19\)](#), has been compiled by the Allen Institute for AI’s Semantic Scholar team and partner research groups
- **CORD-19 is a free resource** consisting of a continually growing corpus of scientific papers on COVID-19 and related historical coronavirus research.



## Aerospike Annual Summit's Contribution to the COVID-19 Cause

- Aerospike, HPE and Intel are giving back to those on the COVID-19 front lines by donating \$30 per conference attendee to two COVID-19 Relief Efforts:

[World Central Kitchen](#)

[International Medical Corps](#)



*Stop by the HPE Virtual Booth to say "Hello" or to ask me any questions from this presentation!*

AEROSPIKE  
**NEXTGEN**  
**NOW**  
SUMMIT '20

# Thank you!

Theresa Melvin, HPE

*Global Chief Architect, BD/AI Solutions*

*Full Stack Data Science Lab, Lead*

*Open Source Profession, BD/AI Lead*

**HPE AI**

  
**Hewlett Packard**  
Enterprise

