WEKA.IO
World's Fastest File System

PENGUIN COMPUTING
A subsidiary of SMART Global Holdings, Inc.

kx

A Storage Architecture to Accelerate kdb+

Penguin Computing & WekaIO

Penguin Computing

US HQ in Fremont, CA
Offices in Japan and UK
18,000 installations, 40 countries

What We Do:
Specialize in innovative on-premise high-performance computing (HPC), bare metal HPC in the cloud, AI, and storage technologies coupled with leading-edge design, implementation, hosting, and managed services

FrostByte with WekaIO Integrated Solution

Jointly Designed, Engineered and Vetted

WEKA.IO

Founded in 2013
US HQ in San Jose, CA
Eng. in Tel Aviv, Israel
10 patents issued, 54 submitted

What We Do: Accelerate artificial intelligence and data intensive workloads, on premises and in the public cloud

High Performance Storage Requirements

Compute & KX kdb+

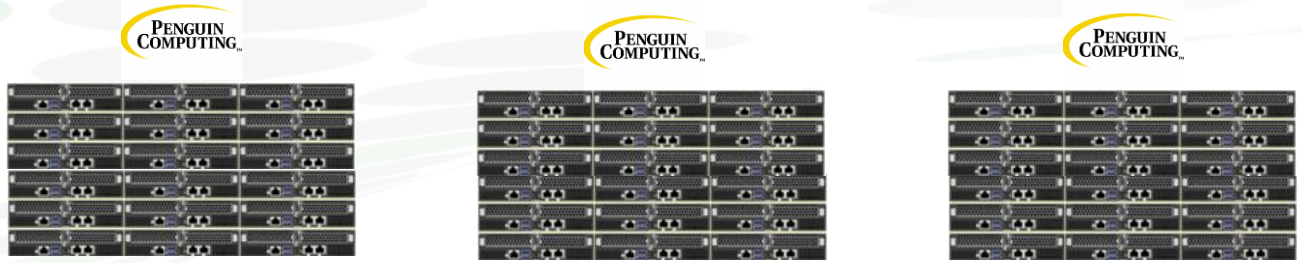

Algorithmic trading | Regulatory & Compliance | Research

Different workloads ;

- Real-time
- Streaming
- Historical

High Performance Storage

With Varied IO Requirements

- Support POSIX filesystem interfaces
- Provide suitable performance for streaming and random I/O
- Offer acceptable performance for both large and small files
- Provide good metadata performance
- Low latency

Accelerating kdb+ Time-Series Data

Workload: STAC-M3 KANAGA

- Test suite run against five years of stock market data

Test Solution:

- KX Systems kdb+, Penguin Computing FrostByte Integrated Solutions combining NVMe servers & WekaIO Matrix Parallel File System

Technical Value:

- Performance density: Up to 87.5GB/sec bandwidth & 2.5M 4K IOPS in 4u
- Fully distributed data, metadata and file system services
- Super low file system latency: avoids kernel, leverages high-speed networks and DPDK + SPDK

Test Results (www.STACresearch.com/KDB190501):

- 8 STAC-M3 mean response-time records
  - 4 STAC-M3 throughput records

Customer Value:

- Faster time to insight for market data
- Lower TCO with industry-standard hardware & object tiering

TEST - M3 KANAGA	MEAN	IMPROVEMENT
100T.YR1VWAB-12D-HO.TIME	1380.2	⟵ 23.4%
100T.YR2VWAB-12D-HO.TIME	1151.22	⟵ 23.1%
100T.YR3VWAB-12D-HO.TIME	1571.63	⟵ 21.5%
100T.YR4VWAB-12D-HO.TIME	2344.54	⟵ 41.5%
100T.YR5VWAB-12D-HO.TIME	9561.65	
10T.YR2-MKTSNAP.TIME	1115.26	
10T.YR3-MKTSNAP.TIME	1204.13	
10T.YR4-MKTSNAP.TIME	1276.46	
10T.YR5-MKTSNAP.TIME	1338.42	
1T.2YRHIBID.TIME	353.326	⟵ 5.8%
1T.3YRHIBID.TIME	578.384	⟵ 9.0%
1T.4YRHIBID.TIME	1020.01	⟵ 6.0%
1T.5YRHIBID.TIME	1673.07	⟵ 7.8%
1T.OLDYRHIBID.TIME	66.4426	
1T.YR1VWAB-12D-HO.TIME	479.166	
1T.YR2VWAB-12D-HO.TIME	589.29	
1T.YR3VWAB-12D-HO.TIME	666.148	
1T.YR4VWAB-12D-HO.TIME	801.183	
1T.YR5VWAB-12D-HO.TIME	1022.46	
50T.YR1VWAB-12D-HO.TIME	2746.58	
50T.YR2VWAB-12D-HO.TIME	1976.66	
50T.YR3VWAB-12D-HO.TIME	2567.28	
50T.YR4VWAB-12D-HO.TIME	3671.56	
50T.YR5VWAB-12D-HO.TIME	5910.83	

Performance Comparison

Compared to a solution using Lustre (KDB150528), this WekaIO/Penguin solution was:

- Better in 15 of 17 baseline benchmarks (Antuco suite), including…
- 11x the speed in STAC-M3.β1.1T.YRHIBID.TIME (purest serial read benchmark)
- 2.1x the speed in STAC-M3.β1.10T.MKTSNAP.TIME (purest random read benchmark)

Compared to solution using a network-attached flash array from a popular vendor (KDB140415), this Weka/Penguin solution was:

- Better in 12 of 17 baseline benchmarks (Antuco suite), including…
- 9.3x the speed in STAC-M3.β1.1T.YRHIBID.TIME
- 5.7x the speed in STAC-M3.β1.10T.MKTSNAP.TIME

WekaIO Matrix

One Solution, Many Use-Cases, Industry-Standard Hardware

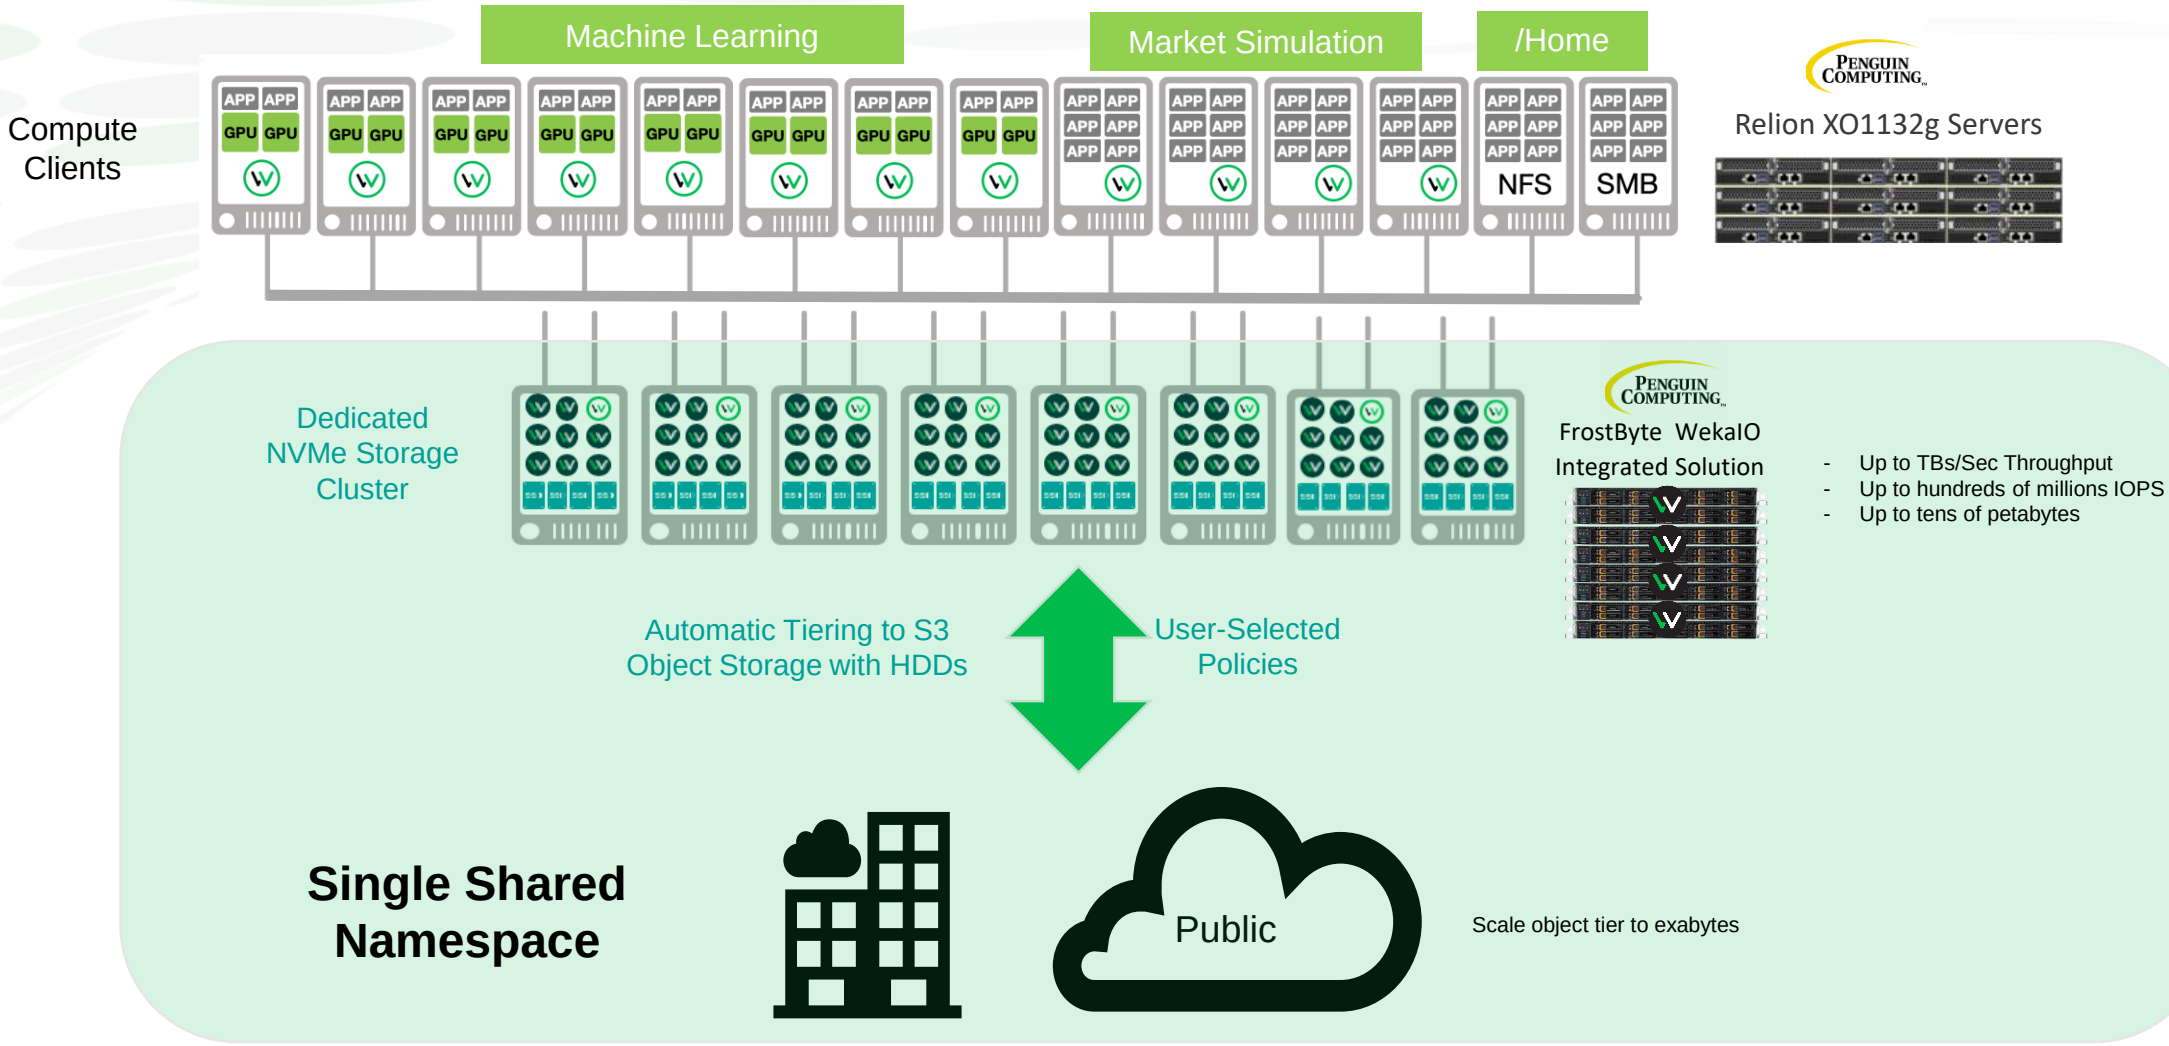

Thank You