High Speed Networking with WEKA Cuts Costs & Complexity

The "Built-For-Purpose" Series

The power of the AI-Native WEKA architecture is that we can deliver incredible IOPs, bandwidth, and latency in a highly efficient footprint.

You know us for our performance.
But we're pretty awesome when it comes to networking efficiency too.

With WEKA

- High per-port performance
- Fewer switch ports to achieve the desired performance levels required for checkpointing and inference at scale
- Significantly less cabling and fewer nodes required to reach the desired performance
- Reduced costs and complexity

A Shared-Everything architecture merely shifts the bottleneck by creating inefficiencies in the networking stack.

- Expensive and sprawling backend connections
- More points of potential failure
- Increased physical infrastructure management overhead
- Increased cost for additional switch ports and cabling
- Increased complexity
- Difficult to troubleshoot
- Challenging basic maintenance

Let's look at how a DASE architecture compares to the Networking Efficiency of WEKA. First let's consider a 1 PB Configuration.

DASE Competitor

- 14 Rack Units
- DASE: 4 Storage Nodes
- DASE: 4 Data Nodes
- DASE: 96 x 100Gb Network connection
- DASE: 32 x 1Gb Management Network
- Roughly 200GB/s read bandwidth
- 40GB/s write band
- Roughly 1 million IOPS

- 4 x storage nodes connected to 4 x compute nodes through a pair of fabric switches
- 4 x 100Gb connections per node (8 connections per box)
- Compute nodes must connect to the backend NVMe fabric using another 8 x 100Gb network cables to communicate with the storage nodes.
- Each storage node is also connected to the NVMe fabric with up to 8 x 100Gb network connections per box.

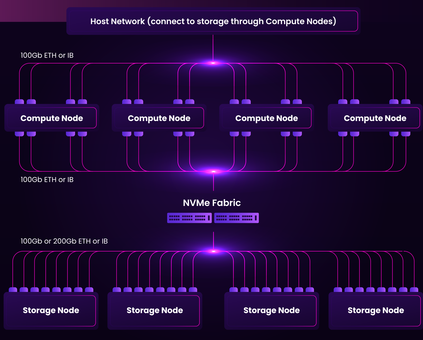

WEKA

- 8 Rack Units
- WEKA: 8 Storage Nodes
- WEKA: 0 Data Nodes
- WEKA: 16 x 400Gb Network connection
- WEKA: 8 x 1Gb Management Network
- Over 720GB/s read bandwidth
- 186GB/s write bandwidth
- 18 million+ IOPS

Look at all that bandwidth and those IOPS

1/6th the number of network cables and ports &

3.5x-18x the performance of shared-everything NAS

Now let's consider a larger configuration designed to deliver "Best" level requirements based on NVIDIA's guidance around recommended storage performance for NVIDIA Cloud Partners (NCP).

1,000 GPU infrastructure,
which consists of 4x NVIDIA Scalable Units (SU)

Shared Everything Scale-Out NAS solution

- 101 rack units
- 25 x storage nodes
- 30 x compute nodes
- 20 NVMe fabric switches since you'd be forced into a spine-leaf NVMe fabric topology
- 60 x 400Gb-to-2x200Gb splitter cables connecting compute nodes to the host network and 440 x 100Gb network cables to connect compute and storage nodes together, 220 x 1Gb management connections.

Deliver over 1,200GB/s read bandwidth
250GB/s write bandwidth

Power: ~90kW
Capacity: 7.5PB

WEKA architecture

- 12 rack units
- 12 homogenous storage nodes, each node connecting to the host network using redundant
- 400Gb networking cables
- 24 x 400Gb network connections
- 12 x 1Gb management connections.

Capable of nearly 800GB/s read bandwidth 264GB/s write bandwidth

Power: 14kW
Capacity: 0.9PB

To achieve the recommended 250GB/s, WEKA requires

20x fewer network cables and ports

1/12 the rack units

1/3 the nodes

1/8 the capacity

1/6 the power

Reduce Complexity. Reduce Costs.
Improve Performance.

Dive Into the Details

WEKA

910 E Hamilton Avenue, Suite 430, Campbell, CA 95008 | 408.335.0085 | info@weka.io | www.weka.io