



FEATURES



PRECONFIGURED AND FULLY TESTED
Best-in-class components engineered to work together



EASY TO MANAGE
Simple, intuitive management via GUI, API, or CLI



INTEGRATED STORAGE TIERING
Tier to multiple S3 targets for infinite capacity scaling and lowest cost



INDUSTRY LEADING PERFORMANCE
World's fastest file system with native NVMe and NVIDIA GPUDirect Storage



EXASCALE CAPACITY
Scale to petabytes of NVMe and exabytes of Object storage



UNIFIED ACCESS
POSIX compliant with support for NFS, SMB, S3 gateway, and NVIDIA GPUDirect Storage



HYBRID CLOUD INTEGRATION
Backup or burst to the cloud for near unlimited resources



TRANSPARENT SNAPSHOT CAPABILITY
Integrated backup and disaster recovery

SUPERMICRO

Supermicro is a global leader in high performance, green computing server technology and innovation. We provide our global customers with application-optimized servers and workstations customized with blade, storage, and GPU solutions. Our products offer proven reliability, superior design, and one of the industry's broadest array of product configurations, to fit all computational need.

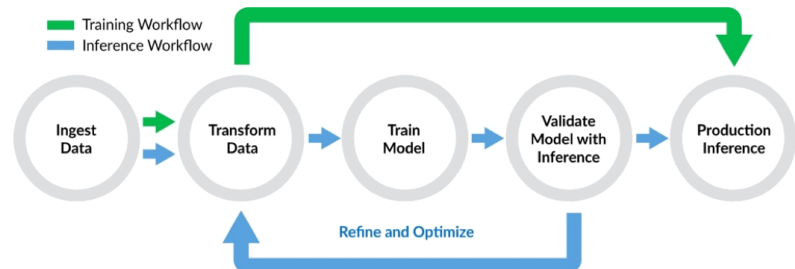
Supermicro AI / HPC Solution Bundle

Featuring WekaFSTM

Accelerate High-Performance Workloads

Designing and implementing AI and HPC infrastructure can be time consuming and daunting as workload demands resemble high performance computing compared with traditional data center applications. These workloads benefit from large amounts of data from gigabytes to petabytes including a more demanding metadata intensive workload than traditional HPC systems were designed to serve.

WekaTM software is an NVMe-native, POSIX compliant, shared filesystem that runs on Supermicro[®] servers, delivering the highest bandwidth, lowest latency performance via InfiniBand or Ethernet enabled CPU/GPU compute cluster. Having an integrated solution is critical to accelerate the I/O pipeline to minimize both time to value and time to insight.



The Solution: a pre-configured stack that delivers the right price performance solution

The combined SuperMicro[®] and WekaIO[™] solution integrates the latest NVIDIA GPUs in the latest AMD[™] PCIe Gen4 servers, and the latest high performance networking to create a tightly integrated and validated solution for the most challenging HPC and AI workloads. This solution provides several choices of NVIDIA GPUs including NVIDIA A100 Tensor Core GPUs, whether PCIe Gen4 or SXM4 with NVIDIA[®] NVLink[™] technology, several options on networking, and support for [NVIDIA[®] NGC](#); enabling access to frameworks, AI SDKs, HPC, pre-trained models and other tools.

For more information, contact us at WEKA-PM@supermicro.com

GPU Servers



H12: 4U-8GPU AS -4124GO-TNAR

Integrated Performance, NVIDIA HGX™ A100 8-GPU



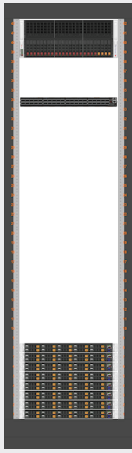
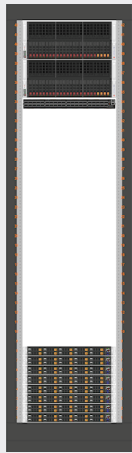
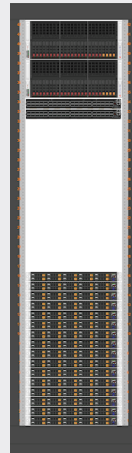
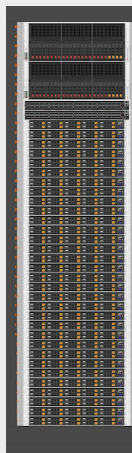
H12: 4U-8GPU AS -4124GS-TNR

Direct Attach & Low Latency, PCIe GPU



H12: 2U-4GPU AS -2124GS-NART

Scalable Performance, NVIDIA HGX™ A100 4-GPU

WekaIO ENGINEERED SOLUTIONS	8 SERVERS	8 SERVERS	16 SERVERS	32 SERVERS
<ul style="list-style-type: none"> HIGHEST PERFORMANCE LOWEST LATENCY MASSIVELY SCALABLE SPACE AND ENERGY EFFICIENT FULLY PRECONFIGURED 				
GPU Server AS-4124GS-TNR	1*	2	2	2
Switch SSE-MQM8700-HS2R	1	1	2	2
Total Rack Space	13U	17U	26U	42U
Storage Nodes	8*	8*	16	32
SSD Count	80	80	160	320
Raw Capacity** (TB)	1228.8	1228.8	2457.6	4915
Usable Capacity** (TB)	645	645	1777	3809
IOPS Performance (millions) ***	6.8M	6.8M	13.6M	27.2M
Bandwidth (GB/sec) ***	320	320	640	1280

* - Minimum configuration ** - Rounded to nearest number with N+2 data protection
 *** - Estimated

