



## 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

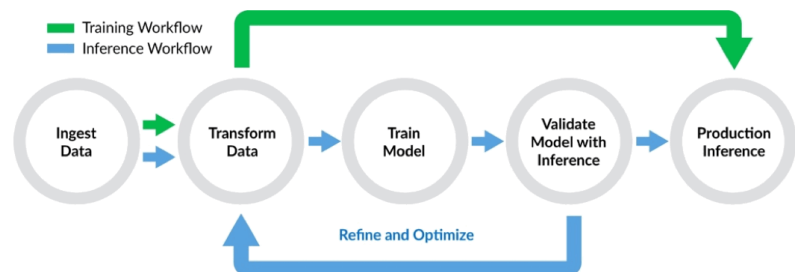
# Supermicro® AI / HPC Solution Bundle

Featuring WekaFS™

## Accelerate High-Performance Workloads

Designing and implementing Artificial Intelligence (AI) and High Performance Computing (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 workloads than traditional HPC systems were designed to serve.

Weka™ 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.



## 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 the industry's broadest array of product configurations, to fit all computational needs.

## 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; 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 the 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](mailto:WEKA-PM@supermicro.com)



## H12: 4U-8GPU

Integrated Performance, NVIDIA HGX™ A100 8-GPU



## H12: 4U-8GPU

Direct Attach & Low Latency, PCIe GPU



## H12: 2U-4GPU

Scalable Performance, NVIDIA HGX™ A100 4-GPU

AMD Bundles	8 SERVERS	10 SERVERS	12 SERVERS	18 SERVERS	18 SERVERS
<ul style="list-style-type: none"> <li>• HIGHEST PERFORMANCE</li> <li>• LOWEST LATENCY</li> <li>• MASSIVELY SCALABLE</li> <li>• SPACE AND ENERGY EFFICIENT</li> <li>• FULLY PRECONFIGURED</li> </ul>					
<b>Useable Capacity</b>	<b>250TB</b>	<b>500TB</b>	<b>1,014TB</b>	<b>1,462TB</b>	<b>2,089TB</b>
Orderable number	WEKAAMD250	WEKAAMD500	WEKAAMD1014	WEKAAMD1462	WEKAAMD2089
Network Cards/Server	200Gbps Dual Port	200Gbps Dual Port	200Gbps Dual Port	200Gbps Dual Port	200Gbps Dual Port
Weka SW License Term	3 Years, 24x7	3 Years, 24x7	3 Years, 24x7	3 Years, 24x7	3 Years, 24x7
HW Support Included	3 Years, 24x7	3 Years, 24x7	3 Years, 24x7	3 Years, 24x7	3 Years, 24x7
Cluster Read IOPS	9.5 Million	11.9 Million	14.9 Million	24 Million	24 Million
Cluster Write IOPS	1.6 Million	2 Million	2.5 Million	4.2 Million	4.2 Million
Cluster Read BW	190GB/s	240GB/s	300GB/s	415GB/s	450GB/s
Cluster Write BW	67GB/s	84GB/s	105GB/s	149GB/s	165GB/s

