



SUPERMICRO RACK SCALE SOLUTIONS: LARGE SCALE AI TRAINING WITH LIQUID COOLING

Embrace an Order-of-Magnitude Leap In Performance With Supermicro Rack Scale AI Solutions

TABLE OF CONTENTS

Contents

Executive Summary	1
Solution Highlights	1
Use Cases.....	3
Representative Performance Benchmarks	5
Supermicro Advantages with Scale AI Solutions Plug and Play.....	7
Further Information	7
Appendix	7



Executive Summary

Turbocharge AI infrastructures with Supermicro’s rack-scale complete plug-and-play AI solutions powered by Supermicro [SYS-821GE-TNHR](#) or the [AS -8125GS-TNHR](#) GPU Servers. These application-optimized servers are ideal for small and medium AI training and contain dual CPUs and eight high-performance NVIDIA GPUs. This purpose-built highest density turnkey rack scale solution is extremely scalable and customizable to meet any scale of Deep Learning workload demands. Leveraging NVIDIA’s cutting-edge NVIDIA H100 SXM GPUs and harnessing the power of Supermicro’s supreme building blocks, these rack-scale AI solutions deliver unprecedented Deep Learning performance.

Solution Highlights

Supreme AI Cluster for Exascale Computing

The Supermicro Rack Scale AI Solution is powered by Supermicro GPU servers - the highest density and compact computation powerhouse. The cluster utilizes the latest NVIDIA HGX™ H100 GPUs to deliver incomparable performance. The design features 32 GPUs in the Base Package (Scalable Unit-SU), scaling up to 128 GPUs per POD and 256 GPUs per SuperPOD.

Scalable Design achieving unprecedented peak performance

Supermicro Rack Scale AI Solution is designed to provide outstanding scalability for faster and easier future deployments. Starting with the Base Package (Scalable Unit-SU) delivering 1.1 PetaFlops, the cluster is seamlessly scalable to reach close to 4.5 Peta Flops (FP64) per POD and up to 8.7 Peta Flops (FP64) per SuperPOD. In addition, Supermicro offers to deploy Rack Scale AI Solutions with 1 to 4 nodes as a proof of concept (POC). It provides flexibility to quickly scale to hundreds of servers via SuperPODs to meet workload demands.

Most Advanced Processors & Networking

The clusters feature the latest and state-of-the-art CPUs (both [4th Gen Intel® Xeon® Platinum Processors](#) and AMD EPYC™ 9004 Processors), achieving an unprecedented number of cores. Furthermore, each cluster is powered by NDR InfiniBand, allowing virtually unlimited scalability for large data aggregation through the network. Along with three terabytes per second (TB/s) of memory bandwidth per GPU and scalability with NVLink and NVSwitch—further enhancing system performance under heavy Deep Learning workloads.

Customizable Highest Quality Storage Options

Depending on the scale and end application requirements, Supermicro's AI Solution offers additional [Storage Solutions](#) that are seamlessly integrated and completely tested with the compute cluster. Utilizing Supermicro's application-optimized, high performance storage blocks, along with storage software integration (parallel filing systems like WekaIO, BeeGFS, etc.), Supermicro's AI solutions are complete offerings – capable of meeting any scale of DL workloads.

Flexible and Superior Cooling Options

With the rising number of TDPs for both CPUs and GPUs, large-scale AI clusters will soon demand superior cooling technologies compared to air cooling. Supermicro Rack Scale AI Solution offers air and liquid cooling options, which include Direct To Chip, Rear Door Heat Exchangers, and Immersion Cooling. Powered by Supermicro's high-quality liquid cooling components, Supermicro's AI solution provides dramatic savings in PUE and OPEX. In addition, the building blocks for this solution can be either air cooled or liquid cooled.




Figure 1 - Liquid Cooled or Air Cooled Supermicro Servers with H100 GPUs

Use Cases

Supermicro's rack-scale AI solutions are designed to remove AI infrastructure obstacles and bottlenecks, accelerating Deep Learning (DL) performance to the max.

Primary Use Case – Large Scale Distributed DL Training

Deep Learning Training requires high-efficiency parallelism and extreme node-to-node bandwidth to deliver faster training times. Supermicro's rack-scale design facilitates training massive neural network models with millions of training instances and billions of parameters - in the most optimized and cost-efficient manner. While the Base Package (SU) is a great starting point for training DL models, the PODs and SuperPODs (and beyond) ensure to shorten enterprise level DL training times to a minimum.

Secondary Use Cases – DL Inference & Hyperparameter Search

In addition to parallel computing, production Inferencing requires deploying models at scale and high availability. Supermicro's GPU servers include (N+N) power redundancy, capable of delivering seamless inferencing performance. Regarding Hyperparameter search, where parallel computation matters but node-to-node bandwidth is not critical, Supermicro's customizable networking options make these solutions an excellent fit for this application.

Deep Learning Performance, Redefined



Image Recognition



Speech Processing



Data Analytics



Forecasting



Chat Bots

... & BEYOND

Supermicro's versatile array of end-application and delivery focused total AI solutions offer great flexibility to choose from the latest and greatest compute platforms. Whether the application is DL Training focused or both DL Training and Inferencing focused (ChatGPT, GPT-4, etc.), our wide range of powerful building blocks and seamless solution integration has it covered.

The Supermicro Scalable Rack Scale AI Solutions are based on a single rack as a building block, referred to as a Scalable Unit, which can then be expanded to four or eight racks.



Scaleable Unit
Base
Package

POD

SuperPOD

	SU/Base Package SRS-42UGPU-AI-SU1	POD SRS-42UGPU-AI-SU2	SuperPOD SRS-42UGPU-AI-SU3
<i>GPU Server (8U 8GPU)</i>	4x SYS-821GE-TNHR / 4x AS -8125GS-TNHR	16x SYS-821GE-TNHR / 16x AS -8125GS-TNHR	32x SYS-821GE-TNHR / 32x AS -8125GS-TNHR
<i>Total CPUs</i>	8x Intel® Xeon® Platinum 8480+ Processor or 8x AMD EPYC™ 9004 Processor	32x Intel® Xeon® Platinum 8480+ Processor or 16x AMD EPYC™ 9004 Processor	64x Intel® Xeon® Platinum 8480+ Processor or 64x AMD EPYC™ 9004 Processor
<i>Total GPUs</i>	32x NVIDIA HGX H100 SXM5	128x NVIDIA HGX H100 SXM5	256x NVIDIA HGX H100 SXM5
<i>Rack</i>	1x 42U (Optional 48U)	4x 42U (Optional 48U)	8x 42U (Optional 48U)
<i>Memory</i>	32TB DDR5 (X13) / 24TB DDR5 (H13)	128TB DDR5 / 96TB DDR5	256TB DDR5 / 192TB DDR5
<i>Estimated Total Power Per Rack</i>	Max 45 kW	Max 180 kW	Max 360 kW
<i>Networking</i>	1x 400G 64-port NDR IB: SSE- MQM9700-NS2F	1x 400G 64-port NDR IB: SSE- MQM9700-NS2F	3x 400G 64-port NDR IB: SSE- MQM9700-NS2F
	1x SMC 100G Eth Switch (Storage)	1x SMC 100G Eth Switch (Storage)	1x SMC 100G Eth Switch (Storage)
	1x SMC 1G/25G MGT Switch	1x SMC 1G/25G MGT Switch	2x SMC 1G/25G MGT Switch

Representative Performance Benchmarks

X13-H100 GPU Super Server: Training Performance - ResNet-50 v1.5 for MXNet (Config 1)

Performance Unit: Images/ sec (Higher is better) -

# of GPUs	SYS-821GE-TNHR	Ref NVIDIA DGX A100
1	6385	3411
4	25022	13443
8	49433	26674

H13-H100 GPU Super Server: Training Performance - ResNet-50 v1.5 for MXNet (Config 2)

Performance Unit: Images/ sec (Higher is better)

# of GPUs	AS-8125GS-TNHR	Ref NVIDIA DGX A100
1	6393	3411
4	24705	13443
8	49057	26674

H13-H100 GPU Super Server: TensorRT BERT Large Inference Performance (Higher is Better)


Data Type	Batch Size	Sequence Length	AS-8125GS-TNHR	Ref A100 SXM4 80GB	Ref H100 SXM5 80GB
INT8	128	128	9743	4887	9622
INT8	8	128	5056	2679	5024
INT8	128	384	2826	1412	2819
INT8	8	384	2047	1071	2016

Building Block Server: GPU Super Server SYS-821GE-TNHR



Overview	8U Dual Socket (4 th Gen Intel® Xeon® Scalable Processors), up to 8 SXM5 GPUs
CPU	2x 4 th Gen Intel Xeon Scalable Processors
Memory (additional memory available)	32 DIMM slots Up to 8TB: 32x 256 GB DRAM
Graphics	8x HGX H100 SXM5 GPUs (80GB, 700W TDP)
Storage (additional storage available)	8x 2.5" SATA 8x 2.5" NVMe U.2 Via PCIe Switches Additional 8x 2.5" NVMe U.2 Via PCIe Switches (option) 2x NVMe M.2
Power	3+3 Redundant 6x 3000W Titanium Level Efficiency Power Supplies *These are max Turbo frequencies

Building Block Server: GPU Super Server AS -8125GS-TNHR



A+ Server 8125GS-TNHR

Overview	8U Dual Socket (4 th Gen AMD EPYC™), up to 8 SXM5 GPUs
CPU	2x 4 th Gen AMD EPYC™ Processors
Memory (additional memory available)	24 DIMM slots Up to 6TB ECC DDR5-4800 RDIMM
Graphics	8x HGX H100 SXM5 GPUs (80GB, 700W TDP)
Storage (additional storage available)	8x 2.5" SATA 8x 2.5" NVMe U.2 Via PCIe Switches Additional 8x 2.5" NVMe U.2 Via PCIe Switches (option) 2x NVMe M.2
Power	3+3 Redundant 6x 3000W Titanium Level Efficiency Power Supplies

Building Block Switch: NDR 400G IB SSE-MQM9700-NS2F



Overview	64-ports NDR, 32 x NDR 400Gb/s OSFP ports, managed, power-to-connector (P2C) airflow (forward)
Max Throughput	51.2 Tb/s
Power	Typical power with passive cables (ATIS): 747W Max power with active cables: 1,703W

Supernano Advantages with Scale AI Solutions Plug and Play

One-Stop-Shop: From initial cluster design (extreme optimization for end user DL applications), assembly and configuration, testing and validation, delivery and deployment, all the way up to support and service – Supernano is the ultimate one-stop-shop for AI infrastructure building.

Supernano’s comprehensive AI packages are entirely tested and validated at rack scale. Extensive testing includes L10 (system level tests), L11 (cluster level tests), and L12 (application level optimization and benchmarking).

Further Information

To learn more about Supernano’s Rack Scale AI Solutions, please visit:

- <https://www.supernano.com/en/products/gpu>
- https://www.supernano.com/en/products/gpu?filter_gen=H13&filter_group=4u_gpu_lines#models
- <https://www.supernano.com/en/support/resources/cpu-4th-gen-intel-xeon-scalable>

Appendix

Config 1 -

The AS-8125GS-TNHR system has the following configuration

Hardware	Description	Quantity
Motherboard	H13DSG-O-CPU	1
Firmware	BIOS-Version=F.0, BMC-Version=09.02.01 BETA, CPLD-Version=F5.12.DE	N/A
CPU	AMD EPYC 9634 84C168T	2
Memory	Samsung 128GB DDR5 4800MHz ECC REG DIMM	24
GPU	NVIDIA H100 5XM5 80GB (Delta-Next)	8
Drive	Micron 7450 PRO 480GB NVMe PCIe 4.0 x4 3D TLC M.2 22x80mm (1 DWPD)	1
Drive	Micron 7400 PRO 3.84TB NVMe PCIe 4.0 x4 3D TLC U.3 2.5" 15mm (1 DWPD)	1
NIC	NVIDIA ConnectX-6 VPI Dual Port HDR InfiniBand and 200Gb Ethernet	6

Config 2 -

The SYS-821GE-TNHR system has the following configuration

Hardware	Description	Quantity
Motherboard	X13DEG-OAD	1
Firmware	BIOS-Version=1.0, BMC-Version=01.00.02, CPLD-Version=FS.11.C9	N/A
CPU	Intel Xeon Gold 8462Y+ 32C64T	2
Memory	Samsung 64GB DDR5 4800MHz ECC REG DIMM	32
GPU	NVIDIA H100 SXM5 80GB (Delta-Next)	8
Drive	Samsung PM9A3 7.68TB NVMe PCIe 4.0 x4 TLC U.2 2.5" 7mm (1 DWPD)	1
NIC	NVIDIA ConnectX-7 Single OSFP InfiniBand NDR and 400Gb Ethernet	8

Software specifications for benchmarks

Software	Version
Operating System	Ubuntu 22.04 LTS
NVIDIA CUDA	12.0
NVIDIA Driver	525.85.12
NVIDIA Fabric Manager	525.85.12
DOCKER CE	20.10.23
NVIDIA Docker	2.11.0