


X14 GPU-Optimized

Maximum acceleration for AI Training, LLMs, and Generative AI



Engineered to Perform

Designed for the specific requirements of AI data centers, Supermicro X14 GPU-optimized systems provide maximum acceleration for large-scale AI training, large language models, and generative AI applications. In addition to support for next-generation GPUs, these systems have been completely re-engineered to take advantage of the latest interconnect, memory, storage, and cooling technologies to ensure significant performance increases over prior generations. With X14, Supermicro builds upon its proven AI-optimized system architecture with thermal-design improvements to handle the most powerful AI GPUs

Maximum Acceleration with NVIDIA GPUs

In close partnership with NVIDIA, Supermicro delivers one of the broadest selections of GPU-accelerated systems providing the maximum performance and efficiency for a range of deployments from small enterprises to massive, unified AI training clusters. New Supermicro X14 GPU-optimized systems support NVIDIA HGX H100/H200 8-GPU as well as the next generation HGX B200 8-GPU, allowing organizations to take advantage of the industry's most powerful GPU configurations using a common server architecture. The modular design includes a dedicated GPU tray which houses an SXM5 (H100/H200) or SXM6 (B200) 8-GPU baseboard and is easily accessible from the cold aisle, simplifying installation and servicing.

Next-generation architectures for the most intensive AI workloads

- Dual Intel® Xeon® 6900 series processors with P-cores
- Support for the latest GPUs including NVIDIA HGX™ H100/H200 8-GPU, as well as upcoming HGX B200 8-GPU baseboards
- Up to 10 PCIe 5.0 slots
- Support for DDR5-6400 and 8800MT/s MRDIMMs
- Optimized thermal designs to support 8 GPUs with free-air cooling
- High-density 4U liquid cooled configuration with direct-to-chip CPU and GPU cold plates
- Modular CPU and GPU trays enhance serviceability and maintenance

Ready to Scale

As the fundamental building blocks of AI superclusters, X14 GPU-optimized systems can easily scale up or out as application requirements increase. These systems feature up to 10 PCIe 5.0 x16 slots for a 1:1 GPU-to-NIC ratio which can support high speed NICs and DPUs including NVIDIA ConnectX®-7 and BlueField®-3 to for networking up to 400Gb/s and enabling every GPU in the cluster to communicate with one another directly. As part of Supermicro's SuperCluster fully validated reference architecture, 32 system nodes paired with a non-blocking spine/leaf networking topology via Ethernet or InfiniBand switches form a functional unit of compute that can easily be scaled to thousands of nodes depending on application requirements.


Air or Liquid Cooling

Several system configurations and form factors are available depending on the specific cooling requirements and infrastructure of the data center. An all-new 10U form factor enables support for top-bin GPUs in air-cooled environments, with a modular GPU tray capable of accommodating enlarged heatsinks for enhanced thermal performance. For maximum GPU density, a liquid-cooled 4U architecture can be integrated using Supermicro's complete direct-to-chip liquid cooling solution, allowing up to 8 systems in a standard 48U rack for a total of 64 GPUs.

Powered by Intel Xeon 6 Processors

Supermicro X14 takes performance to the next level, with support for the new generation of Intel Xeon 6900 series processors with P-cores that deliver the highest performance per-core of any Intel Xeon processor ever. Designed for maximum performance and ideal for the most demanding AI, HPC, and cloud environments, Intel Xeon 6900 series processors with P-cores feature up to 128 cores per socket, include new FP16 instructions on the

built-in Intel AMX accelerator to further enhance AI workload performance, and bring new support for MRDIMMs up to 8800MT/s for up to 37% faster memory bandwidth than standard RDIMMs. P-cores are optimized for high performance per core and excel at the widest range of workloads, including better AI performance than any other general-purpose CPU. X14 GPU-Optimized systems will also support Intel Xeon 6900 series processors with E-cores in 1Q'25.



GPU-Optimized	SYS-A22GA-NBRT	SYS-422GA-NBRT-LCC
Processor Support	Dual Intel® Xeon® 6900 series processors with P-cores Up to 500W TDP (air cooled) [†]	Dual Intel® Xeon® 6900 series processors with P-cores Up to 500W TDP (liquid cooled) [†]
GPU Support	NVIDIA SXM: HGX B200 8-GPU (180GB)	NVIDIA SXM: HGX B200 8-GPU (180GB)
Memory Slots & Capacity	24 DIMM slots up to 6TB DDR5-6400MT/s up to 6TB MRDIMM 8800 MT/s	24 DIMM slots up to 6TB DDR5-6400MT/s up to 6TB MRDIMM 8800 MT/s
I/O Ports	2 RJ45 10GbE with Intel® X710 1 VGA port 1 RJ45 dedicated BMC LAN port 2 USB 3.0 ports (rear)	2 RJ45 10GbE with Intel® X710-AT2 1 VGA port
Motherboard	X14DBG-DAP	X14DBG-DAP
Form Factor	10U Rackmount 843mm/33.2" depth	4U Rackmount 879.4mm/34.62" depth
Expansion Slots	8 PCIe 5.0 x16 LP slots 2 PCIe 5.0 x16 (in x16) FHHL slots	8 PCIe 5.0 x16 LP slots 2 PCIe 5.0 x16 FHHL slots
Drive Bays	10 front hot-swap 2.5" PCIe 5.0 x4 NVMe drive bays	8 front hot-swap 2.5" NVMe drive bays
Cooling	19 counter-rotating 80x80x38mm fans	4 Fan 8cm fans Direct to chip cold plates
Power	6x 5250W Redundant (3 + 3) Titanium (certification pending) Level (96%) power supplies	4x 5250W Redundant Titanium (certification pending) Level (96%) power supplies

[†] CPUs with high TDP supported under specific conditions. Contact Technical Support for details.