



SUPERMICRO X14 SERVERS OPTIMIZED FOR DATA CENTER WORKLOADS - INCORPORATING INTEL® XEON® 6 SERIES PROCESSORS

Supermicro X14 Servers Give Users the Flexibility to Match Workloads to Systems



Executive Summary

Supermicro, a Total IT Solution Provider for AI, HPC, Cloud, Storage, and 5G/Edge, delivers a wide range of servers that are workload-optimized. The Supermicro product portfolio includes multi-node servers that share power supplies and cooling technology, the highest performing rack mount servers, GPU optimized servers, and specifically designed systems for use at the edge. Within each product family, there are a number of options, such as the number of storage devices, memory, CPUs, and network capacities. Supermicro designs and manufactures a wide range of servers that are built on the latest technologies. Typically, Supermicro is first to market with CPU and GPU servers, storage servers, and workstations built with the newest technologies.

Table of Contents


Executive Summary	1
What's New with Supermicro X14 Products	3
Supermicro System Advantages.....	8
Intel Xeon 6 Details.....	8
Generation to Generation Comparison.....	9
Summary and Further Information	10



Supermicro Plug-and-Play Deployment at Scale

Supermicro, with worldwide manufacturing facilities, is a leader in supplying enterprises and cloud service providers with the latest server technology. Racks and clusters are delivered ready to operate with a Plug-and-Play design, requiring the user to simply uncrate and plug into power and networking infrastructures. This concept consists of a number of integration services:

- **Consultation** – once a project has been identified, a dedicated project manager is assigned to work with the customer through delivery and beyond.
- **Design** – expert product managers work with the customer to identify the most optimized systems, storage, software, and networking to ensure a reduced TCO once installed in a data center.
- **Assembly** – the entire cluster is assembled in one of Supermicro's worldwide facilities, including liquid-cooling if needed.
- **Testing** – the entire cluster is tested at the L12 level, according to industry standards, with additional testing as specified by the customer.
- **Deployment** – dedicated and experienced Supermicro and partner employees can assist with the deployment of the clusters at the customer site, including networking and complete application testing.
- **Support** – Ongoing maintenance and support are offered to ensure the entire data center operates smoothly and as specified, with agreed upon SLAs.



With the new Intel Xeon 6 Processors, further optimization is now possible, incorporating specially architected CPUs for different workloads.

There are two main product lines for the Intel Xeon 6 processors. Briefly, this includes:

- **Intel Xeon 6 processors with E-cores –Efficient Cores:** These CPUs are designed with significant performance per watt improvement over previous generations of Intel Xeon processors. The primary workloads for this processor can be categorized as cloud-native and scale-out workloads, where the high core count and low energy usage are critical. Performance per watt is the primary goal of these CPUs.
- **Intel Xeon 6 processors with P-cores – Performance Cores:** The P-core products refers to CPUs with performance per core as the primary goal of the CPU. P-cores are designed for the highest performance in a core and excel when an application is required to perform very well. HPC, AI, and Analytics will excel with the Intel Xeon 6 processors with P-Cores.

What's New with Supermicro X14 Products?



New Intel® Xeon® 6 Processors
Workload-optimized & pin compatible between E-core and P-core CPUs (up to 500W TDP*)

*Available Q3, 2024 in future Supermicro products



Up to 288 cores per node today (dual-socket) and up to 576 cores per node in the future. Increased computing density.



Up to DDR5-6400, MCR-DIMMs (8800 MT/s*), and CXL 2.0

Faster memory bandwidth. New capabilities to extend capacity.

*Available Q3, 2024 in future Supermicro products



EDSFF E1.S and E3.S NVMe drives supported on more families
High throughput, higher density



Data Center Modular Hardware System
DC-MHS Reduces complexity and simplifies maintenance

Supermicro X14 systems are designed for a range of workloads and take advantage of the latest system level technologies.

Supermicro's X14 portfolio with Intel Xeon 6 processors with E-cores are applications optimized for different workloads. Using a building block architecture, Supermicro is able to offer a wide range of servers designed for various applications. Included in the Supermicro X14 server portfolio with Intel Xeon 6 processors are:

Multi-Node Servers


Multi-Node servers from Supermicro give customers several benefits, including a very high core density per rack and up to 34,560 cores of Intel Xeon 6 processors with E-Cores. Shared power and cooling fans which can help to reduce the PUE in a data center. These product families are available with increased storage performance and density using the latest E1.S and E3.S NVMe storage devices.

SuperBlade® – 8U: Future-proof, Resource-saving Architecture

Supermicro's high-performance, density-optimized, and energy-efficient SuperBlade® can significantly reduce initial capital and operational expenses for enterprises and data centers. SuperBlade® utilizes shared, redundant components, including power supplies, cooling fans, chassis management modules (CMMs), integrated switches, or pass-thru modules to deliver the most cost-effective, green computing solutions. The X14 8U SuperBlade® architecture maximizes rack density, with up to 120 dual-processor servers in a 48U rack. The number of cables are reduced by up to 95% when compared to rackmount servers. Optional direct liquid cooling (DLC) can support the highest Intel Xeon 6 TDP CPU SKUs to achieve the lowest PUE with the best TCO in the data center.




Key Applications: AI Inferencing / Hybrid and Private Cloud / Cloud Computing / Big Data Analytics / Financial Services / HPC / CDN / vSAN



SuperBlade – 6U: Future-proof, Resource-saving Architecture

Supermicro's X14 6U high-performance, density and memory-optimized, and energy-efficient SuperBlade® can significantly reduce customers' initial capital and operational expenses. SuperBlade® utilizes shared, redundant components, including power supplies, cooling fans, chassis management modules (CMMs), switches, or pass-thru modules to deliver the most cost-effective, green computing solutions. The X14 6U SuperBlade® architecture maximizes rack density, with up to 100 servers per rack. Optional direct liquid cooling (DLC) can support servers with the highest power CPUs to achieve the lowest Power Usage Effectiveness (PUE) with the best TCO. Supermicro's X14 6U SuperBlade® architecture is optimized for performance with maximum memory capacity (32 DIMMs – DP, 16 DIMMs - UP). Up to 20 GPUs can be installed in 6U enclosures for AI/ML, acceleration, graphics, and 3D rendering. It has the ability to house up to 10 NVMe SSDs per server, which can be ideal for vSAN, EDA, big data analytics, and financial services applications.

Key Applications: AI/ML/ Hybrid and Private Cloud / Cloud Computing / Big Data Analytics / Financial Services / HPC / CDN / vSAN / EDA



BigTwin®: Highly Modular Multi-Node Systems with Tool-Less Design

Supermicro X14 BigTwin® systems provide superior performance and serviceability with dual Intel® Xeon® 6 processors per node and hot-swappable tool-less design. Optimized for density (2U4N) or storage (2U2N), BigTwin® architectures can be more cost-effective than standard 1U servers thanks to shared power and cooling while also increasing compute density and reducing overall TCO. The modular mid-plane design provides NVMe Gen 5 storage controller options with support for E3.S devices, and a new riser card design can support up to 4 M.2 drives for boot/OS or metadata/caching.



Key Applications: HCI / HPC / CDN / Hybrid Cloud Container-as-a-Service / Cloud Computing / Big Data Analytics / Back-up and Recovery / Scale-Out Storage

GrandTwin®: Highly Configurable Single Processor Systems with Front or Rear I/O

The GrandTwin® architecture is purpose-built for single-processor performance. The design maximizes compute, memory and efficiency to deliver maximum density. Powered by Intel Xeon 6 processors, GrandTwin's flexible modular design can be easily adapted for a wide range of



applications, with the ability to add or remove components as required, reducing cost. For front configurations, all I/O and node trays are fully accessible from the cold aisle, simplifying installation and servicing in space-constrained environments. Flexible storage with up to eight E1.S storage devices, a significant density increase over Supermicro X13 systems, and networking options are available via front AIOM modules, allowing countless custom configurations.

Key Applications: MEC (Multi-Access Edge Computing) / HPC / Cloud Gaming / Multi-Purpose CDN / High-Availability Cache Cluster / Telco Edge Cloud / EDA (Electronic Design Automation) / Mission-Critical Web Applications

Rackmount Optimized Servers

Supermicro's X14 Rackmount servers offer customers several choices, including single or dual processor configurations, a significant amount of memory capacity, and PCIe expansion slots. In addition, the new DC-MHS (Data Center Modular Hardware System) servers are designed to be installed when designing and implementing a multi-vendor rack scale facility.

	X14 Hyper	X14 CloudDC	X14 WIO
Positioning	Flagship Enterprise	CSP/Scale-out	Entry-level, Efficiency
Segments	HPC, Enterprise, Cloud Service Providers, Storage, Virtualization, Networking	Cloud Service Providers, Hyperscale Data Centers, Enterprise	Virtualization, Cloud Computing, Data Center
Sockets	2/1	2/1	1
DIMMs	32/16	32/16	8
PCIe	Up to 3 slots (1U) Up to 8 slots (2U)	Up to 3 slots (1U) Up to 6 slots (2U)	Up to 3 slots (1U) Up to 5 slots (2U)
Drives	Up to 24 hot-swap 2.5"	Up to 24 hot-swap 2.5"	10 hot-swap 2.5"/8 hot-swap 3.5"
Key Feature	Maximum I/O flexibility	DC-MHS	UP platform

Hyper: Flagship Performance and Flexibility for Enterprise Data Centers



The new X14 Hyper series brings next-generation performance to Supermicro's range of rackmount servers, built to take on the most demanding workloads in the most proven 1U and 2U form factors. Supermicro's modular design allows customization of storage, expansion slots, network, and power supplies to meet the application requirements. The Supermicro Hyper server balances compute, storage, and expansion in a tool-less rackmount design for optimization, flexibility, and serviceability. The X14 Hyper lineup includes the best-selling dual-socket configurations designed for maximum power and compute density and new single-socket architectures to provide balanced performance with only one processor.

Key Applications: Enterprise Server / Cloud Computing / Big Data Analytics / Hyperconverged Storage / AI Inference and Machine Learning / Network Function / Virtualization

CloudDC with DC-MHS: High-density, Tool-less Mechanical Design for Rapid Cloud Deployment and Easy Maintenance

The new Supermicro X14 CloudDC with DC-MHS delivers ultimate flexibility on I/O and storage to support a range of cloud and data center workloads. The systems are designed to meet OCP DC-MHS specifications, improving modularity and flexibility for large-scale enterprises and cloud service providers to simplify data center management with DC-SCM modules. X14 CloudDC also features tool-less brackets, hot-swap drive trays, and redundant power supplies that ensure rapid deployment and more efficient maintenance in data centers. High-efficiency Titanium Level, redundant power supplies provide resiliency and a low carbon footprint. Rich security features include Intel® SGX, TPM 2.0, signed firmware, Silicon Root of Trust, Secure Boot, System Erase, Runtime FW protection, FIPS Compliance, and Trusted Execution Environment.



Key Applications: Private, Public, Hybrid Cloud / Big Data Analytics / AI Inferencing / Machine Learning / Network Appliance / Virtualization / Open BMC

WIO (UP): Wide-Ranging Flexibility for any Enterprise Workload



Supermicro WIO systems offer a wide range of I/O options to deliver truly optimized systems for specific requirements. Users can optimize the storage and networking alternatives to accelerate performance, increase efficiency, and find the perfect fit for their applications. In addition to enabling customizable configurations and optimization for multiple application requirements, Supermicro WIO SuperServers also provide attractive cost advantages and investment protection.

Appliance / General Purpose Computing

Key Applications: Enterprise / Firewall / Security Appliances / Cloud / Network

Storage Servers

Supermicro's Petascale storage systems are ideal for deployments where storage throughput and latency are paramount, including generative AI, mission-critical databases, virtualization, next-gen big data, HPC, media & entertainment, and hot-tier caching. Supermicro's open architectures are designed to work with the broadest range of software partners to create a solution to drive every application. The symmetrical dual-CPU architecture balances resources and reduces latency by minimizing the length of data paths and maximizing airflow over critical components for optimal thermal performance.



With the latest industry-standard EDSFF E3.S form factors explicitly designed for high-performance solid-state media, these storage drives facilitate maximum performance from PCIe 5.0 interconnects, which ensures compatibility with future iterations of the PCIe protocol. These storage systems support the new Gen 5 drives from all major vendors, allowing customers to choose the best components for their specific applications. Supermicro Petascale systems also support the industry's first CXL expansion modules, adding up to 1TB of DDR memory to the already powerful 32-DIMM solution. This emerging CXL technology is now available to add capacity and bandwidth for memory-bound applications

Key Applications: Data Intensive HPC and AI / Private & Hybrid Cloud / Software-Defined Storage / NVMe Over Fabrics Solution / In-Memory Computing / Composable Infrastructure Platform

Edge/Telco

These systems are designed to be installed outside of a traditional data center. These compact systems can accommodate up to 3 GPUs and, with the new Intel Xeon 6 processors with E-Cores, have 2.2X more cores and improved performance per watt, ideal for edge locations where power may be limited. Additionally, these systems feature DC power NEBS compliance and are available in rack mount and wall mounting options.

Hyper-E: Data Center Performance at the Edge

Hyper-E delivers the performance and flexibility of Supermicro's flagship rackmount server family in a compact form factor optimized for telco and micro data center deployments. A mid-depth chassis and front I/O make incorporating Hyper-E into existing edge and telco infrastructure easier, while carrier grade (NEBS Level 3) design and optional DC power options further enhance flexibility in non-traditional data center environments. Storage and expansion configurations can be adjusted depending on the application, while maintenance-friendly design innovations eliminate the need for tools when servicing the system to simplify rollout and installation.



Key Applications: 5G Core and Edge / Telco Micro Data Center

Telco/Edge: Optimized Designs for 5G, Edge Computing and Emerging IoT Systems

Supermicro provides innovative and first-to-market technologies that are the building blocks for today's computing platforms. Rapid growth in embedded markets and open standards are driving the need for higher levels of product integration and optimization through virtualization, AI inferencing, network connectivity, remote management, mobile communication, expanded I/O, and device-to-device communications using space and power-efficient configurations. Supermicro's family of high-performance embedded products is optimized for a wide range of applications and solutions. Supermicro offers many flexible and customized solutions for critical OEM projects, as well as advanced designs for stringent environments, firmware customization, BOM enhancements, and a wide range of legacy IO support.



Key Applications: Multi-Access Edge Computing / Flex-RAN / Open RAN / Edge AI Outdoor 5G

Supermicro System Advantages

Supermicro works closely with Intel to bring the latest Intel CPUs and GPUs to market. Supermicro delivers rack and cluster-level solutions to many industries with a range of server products and configurations. In addition, Supermicro offers a range of services and solutions to address the requirements of CSPs, Enterprise, and SMB customers.

Supermicro Total IT Solutions

- Industry's broadest portfolio of systems based on Intel® Xeon® 6 processors
- Rack Scale plug-and-play service to deliver complete, validated solutions within weeks, not months
- Production capacity of up to 5,000 racks per month worldwide
- Made in the USA program with manufacturing in San Jose headquarters
- Industry standard compliance for hardware and silicon Root of Trust (RoT) and cryptographical attestation of components throughout the entire supply chain
- Supermicro liquid cooling, including CPU/GPU cold plate, Cooling Distribution Unit, and Cooling Distribution Manifolds for a complete integrated solution

Optimized, Open Architectures

- More than 15 families of systems optimized for AI, Cloud, 5G Edge, and more
- Modular Building Block architecture enables customization for specific workloads and configurations.
- Resource saving architecture to reduce materials and energy usage
- Enhanced thermal capacity to support next-gen CPUs, GPUs, and other components
- Flexible networking with Advanced I/O Modules (AIOM) up to 400G per card
- High ambient temperature operation up to 40°C with liquid cooling options
- Support for open and industry standards, including OCP 3.0, DC-MHS, OAM, ORV2, OSF, Open BMC, and EDSFF

Intel Xeon 6 Processor Details

The latest Intel Xeon family of CPUs, the Intel Xeon 6 processors with E-Cores, is designed to bring maximum performance per watt to the market for cloud-native and scale-out applications. The high core counts enable large numbers of applications or existing multi-core applications to run on a single CPU and even more in dual processor systems. Intel Xeon 6 processors are designed to be socket compatible within the same platform.

Intel Xeon 6 processors with E-cores: Technical Details

- Up to 8 channels of DDR5 support per CPU
- Up to 24 GT/s and 6 links of UPI (1.8x gen-over-gen increase in future products)
- Support for secure UPI/PCIe/CXL link encryption
- Support for CXL 2.0 for all device types with Intel® Flat Memory Mode (FMM)
- Support for Intel® Accelerator Engines including Intel® AMX (future products), Intel® DLB, Intel® QAT, Intel® DSA and Intel® IAA
- Maximum performance-per-watt and core density for cloud, networking, analytics, and scale-out workloads
- Future CPUs with up to 288 cores
- Up to 2.7x better 5G core performance per rack vs 2nd Gen Intel® Xeon® processors (Intel internal measurements)
- 1S-2S scalability
- TDP starting at 205W

Intel® Xeon® 6 processors with P-cores: Technical Details (available Q3, CY2024)

- Maximum performance-per-core for AI, HPC, storage, and Edge workloads
- Up to 2.9x better throughput on AI workloads vs 4th Gen Intel® Xeon® processors (Intel internal measurements)
- Up to 2-3x better performance for AI workloads vs 4th Gen Intel® Xeon® processors (Intel internal measurements)
- Up to 2.8x better memory bandwidth 4th Gen Intel® Xeon® processors (Intel internal measurements)
- 1S-8S scalability
- New FP16 support added to Intel® AMX accelerator
- Support for MCR DIMMs

Generation to Generation Comparisons

Each generation of CPUs from Intel improves a number of essential features. The improvement over multiple generations of CPUs is quite impressive and can be viewed below.

	X12	X14/Future Products	Benefit
CPU	3 rd Gen Intel Xeon	Intel Xeon 6	
Memory	DDR4-3200	DDR5-6400 MCR DIMM 8800 (Q3 2024)	2x memory bandwidth 2.7x memory bandwidth
PCIe	PCIe 4.0	PCIe 5.0	2x throughput increase
Storage	U.2	E1.S, E3.S	Increased density, throughput and better thermals
Cores/socket	Up to 40	Up to 288 (Q1 2025)	7.2x increase
CXL	-	CXL 2.0 (all device types)	Increased shared memory pool

Xeon 6700 series					Xeon 6900 series			
	Max DIMM Count/Socket	Memory Channels	Memory Speed (1 DPC)	Memory Speed (2 DPC)	Max DIMM Count/Socket	Memory Channels	Memory Speed (1 DPC)	Memory Speed (2 DPC)
E-core	16	8	6400	5200	12	12	6400	-
P-core	16	8	6400 8000 (MCR)*	5200	12	12	6400 8800 (MCR)*	-

* Available Q3, 2024

Over many generations, progress has been made when looking at the amount of work per watt. The significant improvement over several generations of Intel processors demonstrates how leading-edge technology can significantly improve server performance and enable new types of computing at lower costs.




Image Courtesy of Intel

Intel Xeon 6 CPUs with P-cores

Looking forward, Intel will enhance its family of Intel Xeon 6 CPUs with the addition of P-cores or performance cores. These CPUs will have up to 128 cores per CPU, with a top TDP of 500W and up to 12 channels of memory. These CPUs are designed for one or two socket servers and will have 12 memory channels, with up to DDR5-6400 and MCR DIMMs at up to 8800 MT/s. In addition, each CPU will be able to support up to 96 PCIe 5.0 lanes.

Summary

The new Supermicro X14 product lines are designed for a range of workloads. With the new Intel Xeon 6 processors, specific workloads will show increased performance per core and a better performance/watt compared to previous generations. The Supermicro X14 products are positioned to take advantage of the latest CPU and GPU technologies.

For More Information:

Supermicro X14 Summary: www.supermicro.com/x14

SuperBlade: https://www.supermicro.com/en/products/superblade?pro=generation_new%3DX14

Hyper: https://www.supermicro.com/en/products/hyper?pro=generation_new%3DX14

Hyper-E: <https://www.supermicro.com/en/products/embedded/5g-telecom-systems>

BigTwin: https://www.supermicro.com/en/products/bigtwin?pro=generation_new%3DX14

GrandTwin: v https://www.supermicro.com/en/products/grandtwin?pro=generation_new%3DX14

WIO: https://www.supermicro.com/en/products/wio?pro=generation_new%3DX14

CloudDC with DC-MHS: https://www.supermicro.com/en/products/clouddc?pro=generation_new%3DX14

Telco/Edge: <https://www.supermicro.com/en/products/embedded/5g-telecom-systems>