

H14 FlexTwin System

Liquid-Cooled Multi-Node Server for Maximum Density and Efficiency



A+ Server AS -2126FT-HE-LCC “FlexTwin”



(rear view)

Purpose built for high-performance computing (HPC) applications, the H14 FlexTwin™ system is designed to help overcome challenges including delivering high performance, optimizing costs, working within limited space, power, and cooling constraints, and easing the challenges of cluster management and maintenance.

Designed for High-Performance Computing

This multi-node system is designed for high density, with four 2-socket nodes per 2U chassis. Equipped with the latest AMD EPYC™ 9005 Series CPUs with up to 192 cores, each chassis can host up to 1536 CPU cores for up to 36,864 cores per 48U rack. The capability to use the most dense server processors available is made possible through direct liquid cooling for each CPU combined with midplane fans to cool memory and I/O components.

Built for Increased Density and Energy Efficiency

Deployed in large clusters, the H14 FlexTwin system combines AMD’s tradition of developing highly energy-efficient CPUs with the precise nature of Supermicro’s liquid cooling. These two factors help to lower your data center’s Power Usage Effectiveness (PUE) and aid in reaching your sustainability goals. Combining direct liquid cooling with rear-door heat exchangers can make your FlexTwin deployments neutral with respect to data center cooling.

4-Node-Per-Chassis 2U High-Density High-Performance Compute System

Maximize performance through liquid cooling

- Liquid cooled, multi-node system for CPUs up to 500W
- Dual AMD EPYC™ 9005 Series CPUs per node for up to 8 CPUs and 1536 cores per chassis
- 24 DIMM slots per node for up to 9 TB of DDR5-6000 memory
- Flexible options for I/O and storage
- N or N+N redundant Titanium level power supplies

Cost-Optimized Deployment

The system is designed to integrate through an optimized 200-Gbps HDR InfiniBand fabric architecture using in-row switching. This, along with improved serviceability, helps to reduce both capital and operating expenses.




Improved RAS

When you build a cluster of high-density servers, the last thing you need is for a failure to interrupt your workload. Each pair of servers is supported by up to two power supplies for N+N power. Redundant power combined with liquid cooling contribute to reliability, especially when using Supermicro cooling systems. In the event of a I/O interface, drive, or node failure, all components are hot pluggable in both the front and rear of the chassis. Serviceability is enhanced with L10 through L12 rack-level validation, and easily accessible front- and rear-panel I/O components. Internal M.2 drives can be configured with hardware RAID to minimize the impact of data errors on overall availability.

Power Your HPC Workloads

You can configure nodes to meet your HPC needs. For example with an InfiniBand card in the front-panel-accessible PCIe 5.0 slot and 100 Gigabit Ethernet in the rear AIOM/OCP 3.0 slot. The system can be equipped with up to two front-panel E1.S drives and two internal M.2 drives. This flexibility enables you to power HPC workloads including:

- **Enterprise data center applications** including private cloud, virtualization clusters, and network infrastructure applications
- **High-performance computing** including EDA simulation, computational fluid dynamics, oil and gas exploration, and weather modeling
- **Financial services** including high-speed trading
- **Manufacturing**
- **Scientific research** ranging from simulations to drug discovery



H14 Generation

Dual-socket AS-2126FT-HE-LCC "FlexTwin"

Processor Support	<ul style="list-style-type: none"> • Dual SP5 socket for up to two AMD EPYC™ 9005 Series processors • Up to 192 cores, up to 500W TDP per socket¹
Memory Slots & Capacity	<ul style="list-style-type: none"> • 24 DIMM slots for DDR5-6000 MHz RDIMM; up to 9 TB registered ECC
On-Board Devices	<ul style="list-style-type: none"> • System on Chip • IMPI 2.0 with virtual-media-over-LAN and KVM-over-LAN support • ASPEED AST2600 BMC graphics with RoT
I/O Ports	<ul style="list-style-type: none"> • 1 RJ45 Dedicated IPMI LAN port • 2 USB 3.0 ports (front) • 1 VGA (front) • 1 TPM 2.0 header • Hardware root of trust
Storage	<ul style="list-style-type: none"> • Up to 2 E1.S drives per node (optional) • Up to 2 M2 drives with optional hardware RAID per node
Expansion Slots	<ul style="list-style-type: none"> • 1 full-height, half-length, PCIe Gen 5 x16 slot (front, optional) • 1 low-profile PCIe Gen 5 x16 (rear) • 1 PCIe Gen 5 x16 AIOM/OCP 3.0 slot (rear)
BIOS	<ul style="list-style-type: none"> • AMI 256 Mb (32 MB) SPI Flash ROM
System Management	<ul style="list-style-type: none"> • Integrated IPMI 2.0 plus KVM with dedicated LAN • Supermicro Server Manager (SSM) and Supermicro Update Manager (SUM) • Supermicro SuperDoctor® 5 and Watch Dog
Cooling	<ul style="list-style-type: none"> • Cold-plate liquid-cooled liquid-cooled CPUs only, or liquid cooled CPU, VRM, and DIMMs • Four 4 cm PWM fans
Chassis	
Form Factor	<ul style="list-style-type: none"> • 2U rackmount
Front Panel	<ul style="list-style-type: none"> • On/off and Universal Information (UID) buttons • Power status, HDD activity, network activity, and UID LEDs
Shared Power	<ul style="list-style-type: none"> • 2-4 Titanium Level power supplies up to 3200W with N or N+N availability configurations

¹Certain CPUs with high TDP may be supported only under specific conditions. Please contact Supermicro Technical Support for additional information about specialized system optimization

Open Architecture

Our open-architecture approach for memory and storage gives you the convenience of pre-installed, Supermicro-qualified devices or the freedom to populate the chassis on your own. Whatever your choice, our approach gives you flexibility, easy maintenance, and operating cost reduction.

Open Management

Our approach to management enables you to deliver the scale your organization requires. SuperCloud Composer software helps you configure, maintain, and monitor all of your systems using single-pane-of-glass management. If your DevOps teams prefer to use their own tools, our accessible RedFish-compliant API provides access to higher-level tools and scripting languages. Regardless of your datacenter's management approach, our open management APIs and tools are ready to support you.