

H15 FlexTwin System

Flexible, Dense, Multinode Design with Liquid Cooling



AS-1127FT7-DEO-ALC



(rear view)

Purpose built for high-performance computing (HPC) and AI rendering applications, the all new H15 FlexTwin™ system is designed to Open Compute Project (OCP) specifications including its 10U form factor and 48-54 VDC rack-level power distribution. The system helps 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 two 2-socket nodes per 10U chassis. Equipped with the latest AMD EPYC™ 9006 Series CPUs with up to 256 cores, each chassis can host up to 1024 CPU cores for up to 49,152 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. 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.

Built for Increased Density and Energy Efficiency

Deployed in large clusters, the H15 FlexTwin system combines AMD's tradition of developing highly energy-efficient CPUs with the precise nature of Supermicro's liquid cooling. These two

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

Maximize performance through liquid cooling

- Built for Open Compute Project (OCP) 3.0 racks
- Liquid cooled, multi-node system for CPUs up to 600W
- Dual AMD EPYC™ 9006 Series CPUs per node
- Up to 4 CPUs and 1024 cores per chassis
- Maximum memory capacity (1DPC): Up to 8TB ECC DDR5-8000 MT/s or up to 8TB DDR5-MRDIMM 12800 MT/s
- Flexible options for I/O and storage

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.



Cost-Optimized Deployment

The system is designed for OCP 3.0 racks that feature centralized power and cooling which is economically shared between up to 44 servers per 48OU rack.

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 a 48-54 VDC bus bar fed by two power supplies providing 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 L11 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 and AI Workloads

You can configure nodes to meet your HPC needs. For example with up to 3 InfiniBand cards in the front-panel-accessible PCIe 6.0 slots, or up to 2 InfiniBand cards plus AIOM slot (OCP 3.0 compatible) with 10, 25, or 200 Gigabit Ethernet connectivity. 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 and AI 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

- **Scientific research** ranging from simulations to drug discovery
- **AI rendering** powered by up to 512 cores of AMD EPYC processor performance

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.

H15 Generation	Dual-socket AS-1127F77-DE0-ALC "FlexTwin"
Processor Support (per node)	• Dual SP7 socket for up to two AMD EPYC™ 9006 Series processors • Up to 256 cores/512 threads, up to 600 TDP per socket[†]
Memory Slots & Capacity (per node)	• Slot Count: 32 DIMM slots • Max Memory (1DPC): Up to 8TB ECC DDR5-8000 MT/s or up to 8TB DDR5-MRDIMM 12800 MT/s
On-Board Devices (per node)	• Network Connectivity: Via AIOM • IMPI 2.0 with virtual-media-over-LAN and KVM-over-LAN support
I/O Ports (per node)	• LAN: 1 RJ45 1 GbE Dedicated BMC LAN port • USB: 1 USB 3.0 Type-A Port • Video: 1 Mini-Display Port • TPM: 1 TPM onboard/port 80
Storage Drive Bays (per node)	• 2 E1.S NVMe drive bays (Default) • M.2: 2 M.2 PCIe 5.0 x4 NVMe slots (M-key 22110 or 2280m)
Expansion Slots (per node)	• PCI-Express (PCIe) Configuration: Optional • 1 PCIe 6.0 x16 (in x16) FHHL slot or 2 PCIe 6.0 x16 (in x16) FHHL slot • 1 PCIe 6.0 x16 (in x16) LP slot or 1 AIOM (OCP 3.0)
System BIOS	• BIOS Type: AMI 64MB SPI Flash EEPROM
System Management	• Integrated IPMI 2.0 plus KVM with dedicated LAN • ASPEED AST2700 BMC graphics with RoT • Supermicro Server Manager (SSM); Super Diagnostics Offline (SDO); KVM with dedicated LAN; IPMI 2.0; • Supermicro Thin-Agent Service (TAS); SuperServer Automation Assistant (SAA)
Cooling	• Advanced liquid cooling solution support up to 100% Heat Capture. • Liquid cooling solution with 8 non-hot-swappable 40x56 mm fan support up to 96% Heat Capture. (The heat transfer may be affected by the ambient temperature)
Power Source	• Via Power bus bar: • 48-54V DC power busbar • 1-OU 33kW or 72kW power shelves
Chassis	
Form Factor	• 1-OU Rackmount
Front Panel	• On/off and Universal Information (UID) buttons • Power status, HDD activity, network activity, and UID
Shared Power	• Power shelf for OCP rack
Dimensions (H x W x D)	• 1.89" (48mm) x 21.14" (537mm) x 33.66" (855mm)

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