



# CONSTRUCT OPTIMIZED IT SOLUTIONS WITH THE SUPERMICRO X12 TWIN FAMILY OF MULTI-NODE SERVERS

*Select From the Twin Family of Multi-Node Servers to Exactly Match Workload Requirements with Superior Functionality and Efficiency*



TwinPro®

BigTwin®

FatTwin®

## TABLE OF CONTENTS

Executive Summary .....	1
Supermicro Twin Family.....	1
Supermicro TwinPro.....	3
Supermicro BigTwin .....	3
Supermicro FatTwin .....	4
Example Use Cases .....	7
Summary & Additional Information.....	8

## Executive Summary

The Supermicro Twin family of products has been designed for the most demanding applications while also reducing OPEX through innovative design that reduces electricity usage and E-waste. The Twin family consists of three product families that all offer the benefits of shared power and cooling systems and are based on the 3<sup>rd</sup> Gen Intel® Xeon® Scalable processors.

## Supermicro Twin Family

Multi-node systems are designed so that all servers (nodes) are located within a single chassis with shared power supplies and fans. Supermicro has developed a product family that takes advantage of the latest available computing, storage, and networking technologies in the Twin product family. These systems are available in the TwinPro®, BigTwin®, and FatTwin® product lines. These product lines offer the latest CPU technology in various form factors and offer a significantly increased density

## SUPERMICRO

As a global leader in high performance, high efficiency server technology and innovation, we develop and provide end-to-end green computing solutions to the data center, cloud computing, enterprise IT, big data, HPC, and embedded markets. Our Building Block Solutions® approach allows us to provide a broad range of SKUs, and enables us to build and deliver application-optimized solutions based upon your requirements.



of compute power and memory capacity compared to traditional rackmount servers while reducing power consumption. For example, a Supermicro BigTwin with four nodes and dual processors in each node results in 8 CPUs in just 2U. This increased density compares to a typical 1U rackmount server with 2 CPUs in 1U, a 100% increase in computing density. In addition, a single chassis containing multiple compute nodes reduces the amount of sheet metal used and allows for upgrading the nodes as needed. This also lowers possible E-waste as only one node needs to be updated or upgraded at a time.

The airflow over a hot surface (CPU) needs to be removed so that the CPU does not overheat. Since shared fans can be larger in size (diameter), they can run at lower speeds to produce the same airflow CFM (cubic feet per minute) needed to cool the processors. Larger fans running at lower speeds use less electricity than smaller fans do. The CFM required can be calculated and fan speed adjusted accordingly.

Building blocks enable IT administrators to create environments from multiple components that match their specific workloads. The Twin product lines with flexible computing, storage, and networking capabilities allow an entire IT environment to be constructed and deployed to match specific workloads and the physical design of a data center.

The Twin family consists of three product lines, the TwinPro®, BigTwin®, and FatTwin®. All are configured with the 3<sup>rd</sup> Gen Intel Xeon Scalable processors, with a maximum thermal design profile (TDP) of 270 Watts supported in specific configurations. The nodes in the BigTwin series can also be liquid-cooled for environments where maximum cooling is needed, and data center cooling systems cannot maintain proper operating temperatures.

## Supermicro TwinPro

The Supermicro TwinPro is the entry point into the Twin family. Four nodes, each containing two CPUs, can be configured in the TwinPro. Each node has 16 DIMMs, which give each node up to 4TB of high speed memory. On the motherboard are dual 10GbE network interfaces, ideal for OS access and cluster management. TwinPro systems can also be configured with two redundant 2200W Titanium Level (96%) power supplies. Up to 6 hot-swappable SAS or SATA drives and two internal M.2 NVMe SSDs per node can be used for direct attached storage. Expansion slots for each node include two low-profile PCI-E 4.0 slots x16, ideal for hand-bandwidth network adapters or FPGAs. Both a 1U 2-Node and 2U 4-Node (pictured here) are available.



Figure 1 - TwinPro Servers

## Supermicro BigTwin

The Supermicro BigTwin continues to lead the industry with systems designed for low OPEX and high performance. The BigTwin is available in a 2U form factor, with either two or four nodes available. Each node can accommodate up to two 3<sup>rd</sup> Gen Intel Xeon Scalable processors. The 2-Node system contains CPUs of up to 270W TDP, the maximum available in the Intel product lineup. 4-Node systems can accommodate CPUs up to 205W TDP, which is a major thermal design feature optimized for hyperconverged and high-performance storage solutions. Various storage options are available for the nodes, including up to twelve 2.5" hot-swap NVMe/SAS/SATA or six 3.5" drives, including Helium HDDs to maximize storage capacities for scale-out storage solutions. Additionally, two M.2 SATA or NVMe drives are supported with optional HW RAID 1 for boot redundancy, based on Marvell® 88NR2241 intelligent NVMe switch. Each compute-optimized node contains two PCI-E 4.0 slots and SAS support with Broadcom SAS3808/SAS3908 controllers. Each storage-optimized node has three PCI-E 4.0 slots and SAS support with Broadcom SAS3808/SAS3816 and an optional SAS3916, supporting up to 12 SAS/SATA drives. In addition, the BigTwin family includes up to 16 DIMM slots per server, accommodating up to 6TB per node when using Intel® Optane™ Persistent

Memory 200 series. Onboard networking options for BigTwin nodes are highly flexible with AIOM PCI-E 4.0 add-on cards, compliant with OCP (Open Compute Project) 3.0 SFF standards.




Figure 2 - BigTwin Servers




Figure 3 - BigTwin Server with Liquid Cooling

## Supermicro FatTwin

Supermicro's FatTwin product family is a 4U multi-node system containing eight or four nodes. A true workhorse, the FatTwin is a battle tested platform deployed in some of the largest cloud datacenters. In addition, its flexibility allows customers to standardize on one platform for both their storage and compute needs in today's infrastructure environments.

Each node can contain up to two 3<sup>rd</sup> Gen Intel Xeon Scalable processors, and each node is designed to run with only a single CPU installed on the dual-CPU motherboard. These systems are designed with front facing nodes for easy serviceability. In a hot-aisle, cold-aisle data center, the ability to service these systems from the front (cold aisle) makes servicing more accessible and faster than working in a hot-aisle containment setup. The hot-swappable drive bays can accommodate NVMe, SAS, or SATA storage devices. A single GPU can also be installed for HPC and AI/ML applications. Additionally, 16 DIMM slots are available on

each node, accommodating up to 4TB of DRAM or up to 6TB when using Intel Optane™ Persistent Memory. Network connectivity is through AIOM support.



*Figure 4 – FatTwin Servers*

	TwinPro (X12)		BigTwin (X12)		FatTwin(X12)			
<b>Height</b>	1U	2U	2U		4U			
<b>Number of Nodes</b>	2	4	2	4	4	8		
<b>Max CPUs Per Node</b>	2	2	2	2	2	2		
<b>Max TDP Watts Per Node</b>	185	185	270	205	Single	Dual	Single	Dual
					250	205	250	185
<b>Max DIMMS/Socket (Including PMem)</b>	16	16	20	20	16	16		
<b>NVMe Support</b>	No		Yes		Yes			
<b>Intel® Optane™ PMem Support</b>	No		Yes		Yes			
<b>Max GPUs/Node</b>	0	0	2	2	0	1		
<b>Max. Number of Drives/Node</b>	4	6	12	6	6	8		
<b>Networking</b>	2x10G		AIOM*		AIOM*			

\*AIOM: The Supermicro® Advanced I/O Module (AIOM) is the latest form factor designed to provide a wide range of networking options and other I/O technologies for server, storage, and IoT systems. The Supermicro AIOM card contains many options for networking. See more at:

<https://www.supermicro.com/en/support/resources/aoc/aiom>

*Table 1 - High-Level Comparison of TwinPro, BigTwin, FatTwin*

## Example Use Cases

### TwinPro:

- **Data Center Cloud Computing:** The Supermicro TwinPro is an economical choice when a lower cost multi-node solution is required. With lower cost CPUs than other systems, the Supermicro TwinPro can be used for various data center and cloud computing workloads, where high-density servers are required. The TwinPro reduces E-waste with up to four nodes in a compact 2U chassis, as each node can be replaced or upgraded without replacing an entire chassis. Raw materials are also reduced compared to standard 1U rackmount servers.
- **Virtualization:** For workloads that require advanced virtualization implementation for standard workloads, the Supermicro TwinPro is an ideal solution. With up to 4 nodes in a 2U chassis, these systems can easily satisfy the SLAs of many organizations that require fast responses with economic CPU choices.

### BigTwin:

- **Hyperconverged Infrastructure:** The Supermicro BigTwin is an excellent server for Hyperconverged Infrastructure (HCI) deployments. Many workloads in various industries benefit from the BigTwin, including Big Data Analytics, DevOps & IT Operations, Cloud-Native Applications, High-Performance Databases, and Mission-Critical Applications. With a range of configurations of CPUs and very flexible I/O options, the Supermicro BigTwin allows mission critical applications to be run with the highest performance available. Many Nutanix customers have chosen to use the Supermicro BigTwin for their most demanding enterprise applications. See more at: (<https://www.supermicro.com/en/pressreleases/supermicro-expands-total-solution-portfolio-offering-nutanix-nx-platforms-hybrid>). These environments require flexible storage options addressed with NVMe/SAS/SATA support and multiple network connectivity options that are easily satisfied using AIOM options and high-bandwidth low-profile cards.
- **Hyperscale Cloud Solutions:** With various configurations supporting the latest, highest performing CPUs, the BigTwin is an ideal building block for organizations creating and building their own clouds, whether on-prem or a publicly available service. With the density available in a 2U 4-Node setup with dual sockets per node, the BigTwin can also address a wide range of applications, such as virtualization, software-defined storage, and general hosting requirements.
- **HPC:** The Supermicro BigTwin, with its ability to house the most powerful CPUs and GPUs available, is an ideal choice for High Performance Computing environments. With a high number of CPUs available in a dense package and a range of I/O performance options, the BigTwin accelerates time to solution. Whether running more solutions or increasing the data and physics used in simulations, organizations can bring products to market faster and reduce costs in a wide range of industries.
- **Other Applications:** Other applications that the Supermicro BigTwin is ideal for include: **Back-up & Recovery, All-Flash Object Storage** ([https://www.supermicro.com/white\\_paper/white\\_paper\\_Ceph\\_Cluster\\_BigTwin.pdf](https://www.supermicro.com/white_paper/white_paper_Ceph_Cluster_BigTwin.pdf)), **High-Density RAID storage**, and **Big Data Analytics**.

## FatTwin:

- **Cloud and Web Applications:** Organizations that need to provide services to millions to billions of users, the Supermicro FatTwin is an ideal solution, built with the highest density available. The front-facing (cold aisle) serviceability and ease of maintenance reduce OPEX and allow faster replacement than a server with rear-facing (hot aisle) access. With the design capability of running the nodes with only a single CPU, infrastructure executives can optimize exactly the best amount of computing cores for extremely dense data center environments
- **EDA and Semiconductor Simulation:** With the increasing demand for accurate simulations of multi-billion transistor CPUs and other microprocessors, a system with high speed CPUs and large memory access is required. The Supermicro FatTwin, optimized for high density computing, is an ideal server for simulations that require access to significant amounts of memory. In addition, the Supermicro FatTwin includes Titanium level redundant power supplies, ensuring that simulations do not abort when there is a slight power loss to the critical computing components.
- **Other Applications:** Other applications that the Supermicro FatTwin excels at include: **Cloud Optimized Servers, Data Center Enterprise Applications, and Virtualization Servers.**

Read more about how a public company uses the Supermicro FatTwin to handle trillions of messages daily. [https://www.supermicro.com/CaseStudies/CaseStudy\\_Fortune100FatTwin.pdf](https://www.supermicro.com/CaseStudies/CaseStudy_Fortune100FatTwin.pdf)

## Summary of Supermicro Twin Product Family

The Supermicro Twin Product Family consists of three product lines designed to use the 3<sup>rd</sup> Gen Intel Scalable processors. The Twin family starts with the TwinPro systems. These flexible 1U and 2U systems offer a range of I/O expansion options and are designed to reduce CAPEX and OPEX costs. Next up is the 2U 2-Node and 2U 4-Node BigTwin®, a breakthrough multi-node server system with many innovations and industry firsts, designed to support the highest performance CPUs and GPUs, maximum memory capacity, flexible networking, and all-flash NVMe support. These systems are optimized for answering challenging business, finance, or scientific research requirements. Finally, the 4U FatTwin® series is configurable up to 8 nodes and is built for advanced deployment scenarios with numerous memory capacities, storage technologies, PCI alternatives, networking capabilities, and GPU support options for even more demanding applications.

## Additional Information

<https://www.supermicro.com/en/products/twin>

<https://www.supermicro.com/en/products/twinpro>

<https://www.supermicro.com/en/products/bigtwin>

<https://www.supermicro.com/en/products/fattwin>