

Supermicro AI Networking Switch Portfolio

Pioneering Green Solutions for a Sustainable AI Infrastructure

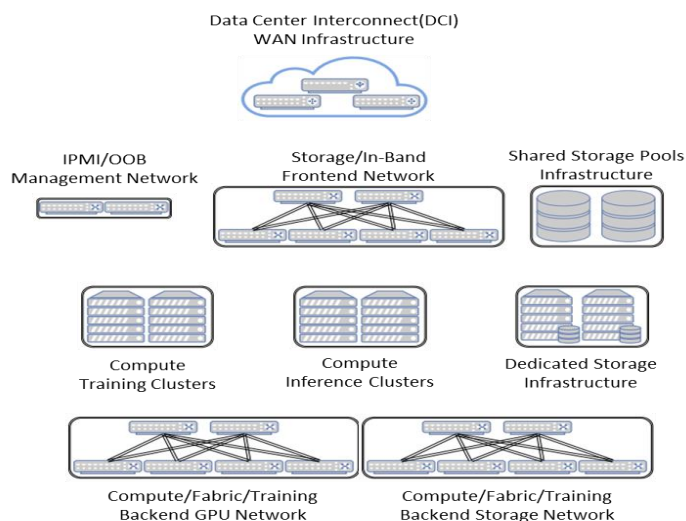


Figure 1: Typical AI/ML Training & Inference DC Infrastructure
The Supermicro AI Networking Solution

AI workloads demand networking that is fast, open, and sustainable. Supermicro delivers a complete switch portfolio, purpose-built for AI and high-performance data centers. Running Enterprise Advanced SONiC or Community SONiC, our solutions provide enterprise-grade features, open networking, and future-ready flexibility.

Unparalleled Performance

Supermicro switches combine high-throughput ASICs with ultra-low latency to accelerate GPU training and inference. The SSE-T8xxx series

delivers massive bandwidth with RoCEv2 optimization and advanced congestion control for faster training cycles and efficient inference.

Scalable by Design

From small pilots to hyperscale AI clusters, Supermicro switches scale seamlessly using advanced layer 3 capabilities, including various hashing algorithms and VXLAN-EVPN. Customers can tailor airflow, cooling, and power options while deploying consistently across 1G, 10/25G, 100G, 400G, and 800G platforms.

Optimized for AI Fabrics

Features such as QoS with end-to-end congestion management, streaming telemetry, and programmable APIs ensure critical AI traffic is prioritized, monitored, and automated. Open integrations through REST and gNMI APIs enable full lifecycle automation and visibility.

Sustainability First

Supermicro's eco-conscious design emphasizes energy efficiency. High-efficiency PSUs, hot-swappable modules, and support for air or liquid cooling reduce power consumption and e-waste, aligning high-performance networking with sustainability goals.

Optimal TCO for AI Networking

Supermicro delivers pre-validated rack-scale designs qualified with leading GPU, NIC, and storage vendors, enabling fast turn-key deployments with shorter lead times. Our Ethernet switches further reduce costs by supporting standard cables and optics without vendor lock-in, all backed by responsive on-site support and services.

Designed for AI/ML pipelines, Supermicro switches combine high speed, high density, and broad compatibility to form the backbone of scalable, high-performance infrastructure. With our solutions, organizations can maximize the value of their AI investments and accelerate innovation with confidence.

Speed/Use Case	Product SKUs	Hardware Details	Software
64x800G Leaf/Spine	SSE-T8164S/SR/D/LCC	Ideal for hyperscale AI leaf and spine fabrics <ul style="list-style-type: none"> 2RU 51.2T Capacity with 64x800GbE QSFP Ports 2x25GbE SFP28 & 1x1GbE RJ45 OOB Port AC & DC models with redundant hot-pluggable 1+1 power supplies Regular & Reverse airflow models with hot swappable 3+1 fans Air-cooled & Liquid-cooled models 	<ul style="list-style-type: none"> Supermicro Enterprise Advanced SONiC with Industry-Standard CLI Rich AI/ML Fabric Optimizations Customizable Congestion Controls Advanced Telemetry Capabilities
64x400G Leaf	SSE-T8196S/SR/D/LCC	Ideal for large AI leaf with 400G to GPUs and 800G uplinks <ul style="list-style-type: none"> 3RU 51.2T Switch with 64x400GbE QSFP112 & 32x800GbE QSFP Ports 2x25GbE SFP28 Port & 1x1GbE RJ45 OOB Management Port AC & DC models with redundant hot-pluggable 1+1 power supplies Regular & Reverse airflow models with hot swappable 3+1 fans Air-cooled & Liquid-cooled models 	<ul style="list-style-type: none"> Supermicro Enterprise Advanced SONiC with Industry-Standard CLI Rich AI/ML Fabric Optimizations Customizable Congestion Controls Versatile connectivity Options Advanced Telemetry Capabilities
64x400G Leaf	SSE-T8032S	Ideal for mid-scale AI and HPC 400G leaf deployments <ul style="list-style-type: none"> 1RU 25.6T Total Capacity with 32x2x400GbE QSFP Twin Ports 2x10GbE SFP+ Port & 1x1GbE RJ45 OOB Management Port AC model with redundant hot-pluggable 1+1 power supplies Regular airflow model with hot swappable 6+1 fans 	<ul style="list-style-type: none"> Supermicro Enterprise Advanced SONiC with Industry-Standard CLI Rich AI/ML Fabric Optimizations Customizable Congestion Controls Advanced Telemetry Capabilities
32x400G Leaf/Spine	SSE-T7132S/SR/D/DR	Ideal for 100/200G aggregation or 400G leaf/spine roles <ul style="list-style-type: none"> 1RU 25.6T Total Capacity with 32x400GbE QSFP-DD Ports 2x10GbE SFP+ Port & 1x1GbE RJ45 OOB Management Port AC & DC models with redundant hot-pluggable 1+1 power supplies Regular & Reverse airflow models with hot swappable 6+1 fans 	<ul style="list-style-type: none"> Supermicro Community SONiC OS Customizable Congestion Controls L2 and L3 VXLAN Support
32x100G Leaf/Agg	SSE-C4632SB/SRB	Ideal for high-density 100G aggregation and leaf networks <ul style="list-style-type: none"> 1RU 3.2T Total Capacity with 32x100GbE QSFP28 Ports 1x10GbE SFP+ Port & 1x1GbE RJ45 OOB Management Port AC model with redundant hot-pluggable 1+1 power supplies Regular model with hot swappable 3+1 fans 	<ul style="list-style-type: none"> Supermicro Enterprise Advanced SONiC with Industry-Standard CLI Customizable Congestion Controls L2 and L3 VXLAN Support Advanced Telemetry Capabilities
48x10/25G Access	SSE-F3648S/SR	Ideal for 10/25G server access with 100G uplinks <ul style="list-style-type: none"> 1RU 2.0T Total Capacity with 48x10/25GbE SFP+/SFP28 & 8x100GbE QSFP28 Ports 1x1GbE RJ45 OOB Management Port AC & DC models with redundant hot-pluggable 1+1 power supplies Regular model with hot swappable 3+1 fans 	<ul style="list-style-type: none"> Supermicro Enterprise Advanced SONiC with Industry-Standard CLI Customizable Congestion Controls L2 and L3 VXLAN Support Advanced Telemetry Capabilities
48x1G Management	SSE-G3748R/D/DR-SONiC	Ideal for secure 1G out-of-band management <ul style="list-style-type: none"> 1RU 396G Total Capacity with 48x1GbE RJ45 & 6x10G/25G SFP+/SFP28 Ports 1x1GbE RJ45 OOB Management Port AC & DC models with redundant hot-pluggable 1+1 power supplies Regular & Reverse airflow models with hot swappable 6+1 fans 	<ul style="list-style-type: none"> Supermicro Community SONiC OS L3 Routing Protocol Support Basic Congestion Controls

Figure 2: Supermicro AI Networking Switch Portfolio

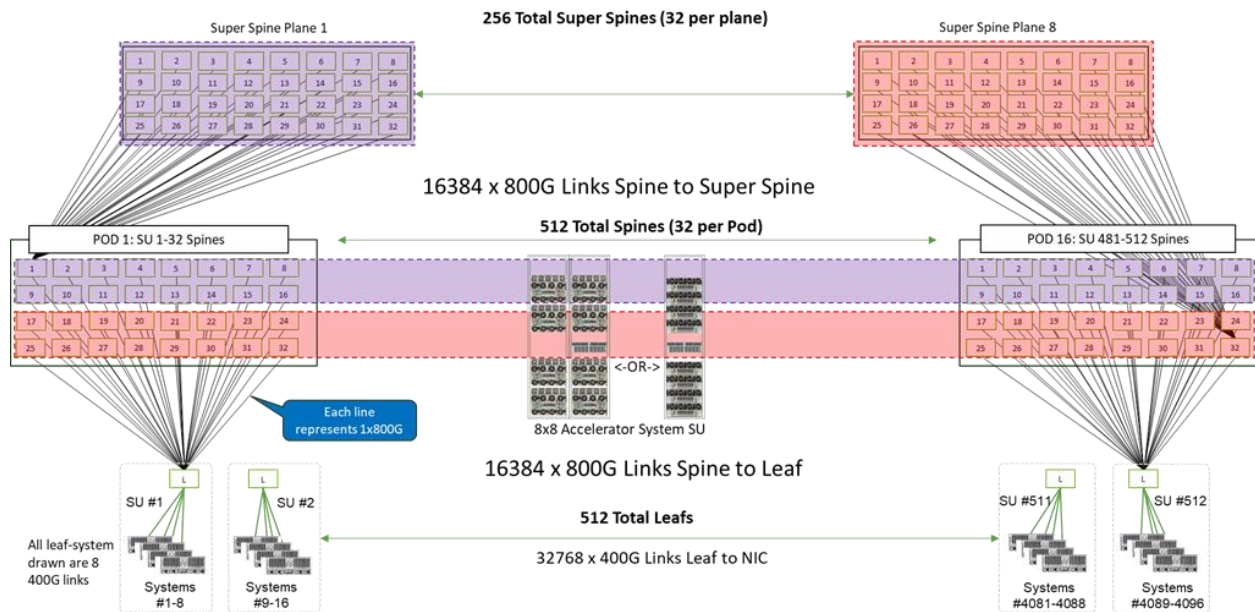


Figure 3: Rail-optimized cluster for 32,768 GPUs using SSE-T8164 and SSE-T8196 switches (Can scale much higher)

Highlights of Rail-optimized cluster for 32,768 GPUs using SSE-T8164 and SSE-T8196 switches:

- Scales from ≤64 to 32,768+ GPUs, with higher scales supported by adding more super spine planes
- Comprises of 3-tier Clos fabric with 512 Leaf, 512 Spine, and 256 Super Spine nodes
- Total number of interconnects/cables between leaf-spine-superspine nodes scales to 32,768 in the above design
- Total number of 8-GPU computes with 8-NICs part of this cluster scales to 4,096

Ideal Applications:

- **SSE-T8164** excels as the ultra-high-density leaf, spine, or super spine switch for hyperscale AI clusters, delivering 800G leaf-and-spine fabric performance akin to industry-leading 800G AI fabric platforms.
- **SSE-T8196** is ideal for high-density AI leaf nodes, balancing 400G downlinks to GPU RNICs with 800G uplinks to spine, ensuring congestion controls are less likely to be needed, perfect for scalable GPU deployments.
- **SSE-T8032** fits midsize AI or HPC leaf deployments, providing 400G bandwidth scalable for compute clusters without over-provisioning.
- **SSE-T7132** serves as an efficient dense aggregation or spine option for 100/200/400G fabrics, streamlining connectivity in modern AI-optimized data centers.
- **SSE-C4632** offers high-density 100G aggregation or leaf capability, enabling seamless migration from legacy 10/25/50G environments with modern fabric readiness.
- **SSE-F3648** delivers 10/25G server access with 100G uplinks, making it an excellent fit for access-layer connectivity in data center and edge deployments.
- **SSE-G3748** offers secure and reliable 1G out-of-band management, ensuring consistent monitoring and control across the infrastructure.

Why Choose Supermicro?

- **Building Block Solutions:** Part of Supermicro's Building Block architecture, our switches integrate seamlessly into pre-validated AI clusters with optimized topologies and templated designs.
- **Performance at Scale:** Deliver industry-leading throughput and ultra-low latency across fabrics, with flexible 1G to 800G platforms that grow from small to hyperscale.
- **Optimized for AI:** Advanced congestion control, telemetry, and automation APIs ensure reliable, efficient, and visible AI/ML pipelines.
- **Simplified Operations:** Industry-standard CLI and SONiC-native tooling simplify Day-0 to Day-2+ configuration, monitoring, and troubleshooting, with support for templated designs to speed deployment.

Next Steps:

- Visit the SMC website for detailed product information: <https://www.supermicro.com/en/products/networking/switches>
- Contact your SMC sales representative to discuss your specific AI networking needs.

Discover how SMC network switches can power your next-generation AI and data center infrastructure.

Disclaimer: Specifications and features may vary by model. Please refer to the SMC website for detailed product information.

