Supermicro SSE-T8196 800G/400G High Density Port Ethernet Switch

Supermicro introduces the SSE-T8196, a high-density Ethernet switch supporting both 800G and 400G in one SKU to be the suitable companion with other product offerings such as SSE-T8164 and SSE-8032S switches. At the heart of ultra high-performance applications like hyper-scale cloud computing and AI/ML clusters lies the pivotal role of 800G and 400G Ethernet. With advancements in today's datacenter deployments with high capacity and large GPU clusters become common, the market has a need for high bandwidth, lossless, and low latency transport, which the SSE-T8196 is designed to excel at.

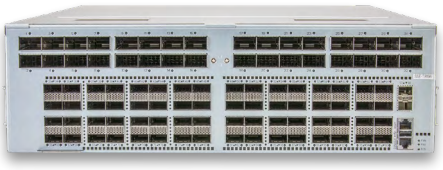

(Front View)

(Rear View)

With a throughput of up to 51.2Tbps, Supermicro SSE-T8196 stands out for offering the highest density of 400G and 800G in a 3RU form factor. With proven layer 2 and layer 3 capabilities, advanced adaptive routing, dynamic load balancing, and support for end-to-end congestion control, these switches provide the ideal foundation for Large-scale enterprise backbone, AI/ML clusters, and Cloud computing.

The SSE-T8196 presents a diverse range of port speeds and densities, supporting 50GbE, 100GbE, 200GbE, 400GbE, and 800GbE, facilitating consistent network architectures that seamlessly scale from small, dedicated clusters to the requirements of expansive multi-tier networks.

SSE-T8196 comes with 64 port 400G QSFP112 downlinks and 32 port 800G OSFP uplinks in a 3RU chassis and the included rail kit facilitates rack-mounting installation.

With emerging Supermicro Advanced Enterprise SONiC OS support, the SSE-T8196 switch provides greater flexibility to applications such as large-scale enterprises, AI/ML and high-performance cloud computing. Supermicro provides user friendly enhancements to the OS that suit the various data center applications and deployment.

Benefits & Advantages

Target Use Cases

- Data Center AI/ML, Core, Edge, ToR and DCI applications

Key Advantages

- Low-latency (<1 usec) and high-throughput (51.2Tbps) Ethernet switch
- High radix of 64 400G and 32 800G ports in a 3 RU form-factor
- Multi-rate 800G/400G/200G/100G ports in direct/breakout mode
- Operating System : Supermicro Advanced Enterprise SONiC OS
- Advanced AI/ML features for optimized job completion times.
- Option for Direct Liquid Cooling system

1K Systems (8192 GPUs) 2 Tier Topology

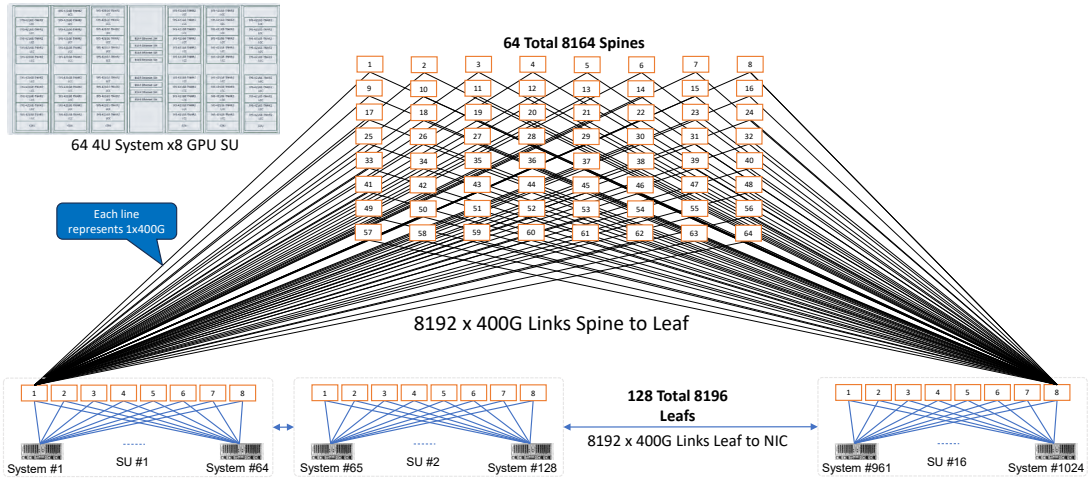

SSE-T8196 is a high radix Ethernet switch with 32 800G OSFP ports and 64 QSFP112 400G ports that can deliver 51.2 Tbps of switching capacity in a compact 3RU form factor. The combination of 800G OSFP and 400G QSFP112 ports make SSE-T8196 more suitable for leaf nodes in large scale non-blocking AI clusters with the 64 QSFP112 ports used to connect to the GPU systems and the remaining 32 OSFP to connect to the spines. The topology above is a representation of how a non-blocking 2-tier scale-out fabric comprising 128 leaf nodes of 8196 and 64 spine nodes of 8164 switches can be built to connect 1k GPU systems (8192 GPUs) into a single cluster.

Hardware Specifications

Key Switch Specs

- Switch Fabric Capacity: 51.2Tbps
- Non-blocking, wire-speed data forwarding
- 165MB memory packet buffer
- BMC AST2600

Physical Ports

- 32x800G OSFP Ports
- 64x400G QSFP112 Ports
- 2x25G SFP28 ports
- 1xRJ45 console and 1xOOB Ethernet management ports
- 1x USB 2.0 Type A interface

CPU Engine

- Intel ICX-D LCC x86 CPU, D1736, 8C
- 32GB DDR4
- 128GB SSD

Physical and Environmental

- Hot Swappable 3+1 load balancing Fan Modules
- Capability for future Direct Liquid Cooling
- Front to rear/Rear to Front airflow
- 3RU, Mountable in 19" or 21" racks
- Dimensions: (WxDxH) 438.4 x 650 x 131.1mm
- Weight: 30kg
- Temperature: Operating 0C to 40C Storage -40C to 70C
- Operating Humidity : 5 to 95%
- Storage Humidity : 8 to 90%

Power

- 3200W AC PSU, support 1+1 Hot-pluggable, load balancing, and redundant, Input Voltage: 200-240VAC, 200-277VAC, 50/60Hz
- 1600W DC PSU, Shared, Input Voltage: 48-72 VDC
- Support Optics/Cables: OSFP ports upto 18W, QSFP112 ports upto 10W

IEEE compliance

- 802.1p QOS/COS
- 802.1Q VLAN Tagging
- 802.1AB Link Layer Discovery Protocol
- 802.3ad Link Aggregation with LACP
- 802.3ae 10 Gigabit Ethernet
- 802.3ba 100 Gigabit Ethernet
- 802.3bs 400 and 200 Gigabit Ethernet
- 802.3cm 400 Gigabit over multimode fiber
- 802.1Qbb PFC 802.1Qaz ETS
- 802.1X Network Access Control
- 802.3x Flow Control
- Cables and Transceivers
  OSFP DAC options from 2-4m (4m on certain NICs)
  OSFP Active Electrical Copper
- OSFP Active Optical to 10m SR4/SR8/VR4/VR8 MMF transceivers to 50m
  DR4/DR8 SMF transceivers to 500m
  FR8 SMF transceivers to 2km

Safety and Compliance	Electromagnetic Compatibility	Environmental
802.1p QOS/COS **EU:** EC/EN 62368-1 **North America:** UL62368-1 CAN/CSA No. 62368-1 **Taiwan (BSMI):** NS 15598-1	**EU:** EN300 389, EN55032, EN IEC 61000-3-2, EN 61000-3-3, EN 55035, B5 EN 55032, BS EN IEC 61000-3-2, BS EN 61000-3-3, BS EN 55035 **North America:** FCC 47 CFR Part 15 Subpart B, ICES-003 issue 7 **Australia/New Zealand:** AS/NZS CISPR 32 **Japan:** VCCI CISPR 32 **Taiwan (BSMI):** CNS 15936 **Korea (KCC):** KS 9832, KS C 9835	**EU:** RoHS, REACH, WEEE **Taiwan (BSMI):** CNS 15663

Supermicro Enterprise SONiC Features

System and Platform Infrastructure Features

- Dynamic Port Breakout
- DOM Information Display
- System Locator LED Support (Beacon)
- CMIS 4.0 Optics Support
- Hardware Watchdog
- CoPP (Control Plane Policing)
- Transceiver Parameter Tuning
- Third-Party Container Management
- PDDF and PDK Framework
- Interface Aliasing (IS-Standard and IS-Standard-Extended Interface Naming)
- Kdump Support
- Maintenance Mode
  - LACP Graceful Shut
  - BGP Graceful Shut
  - OSPFv2 Maximum Metrics
- Multi-Instance Redis DB
- Hardware Resource Allocation and Reservation
- Zero Touch Provisioning (ZTP)
- Auto Negotiation and Link Training
- Link Statistics Enhancements
- Link-Down Reason Codes
- Link Flap Error-Disable
- Forwarding Plane Drop Counters
- Time Zone Command Support
- 2 × 50G Speed Support
- Broadcom Debug Tool
- Memory Histogram
- System Ready for Services and Applications
- Secure Boot Process and Reference Implementation
- Syslog High Threshold notifications and clear for CPU/Temperature
- Per Platform CoPP
- Interface Beacon LED
- ZTP Provisioning using a USB Drive
- Flexible DPB
- Support Option to Bind the Third Party Container to the Management VRF
- Limiting CPU/Memory/Disk Usage for Third Party Containers
- Patching Support in SONiC (Patch Host/Containers)
- Option to Send Audit Log Messages to Syslog Server
- Ability to Filter Logs based on Facility and Severity
- Media AutoFEC for FEC Type automation

Layer 2 Features

- VLAN Auto-state
- Interface Hold-Down
- LACP Graceful Shutdown
- link Tracking
- LACP Fast Rate and LACP Fallback
- Static LAG
- LLDP
- UDLD
- PVST and RPVST+
- MSTP
- DHCP snooping
- Port Channel Min-Links configuration enhancement

Layer 3 Features

- DHCP Relay Enhancements
  - Overlay Interfaces
  - DHCP Relay Source Interface Selection (e.g. loopback)
  - DHCP Relay over IPv6 Link-Local Interfaces with RFC5549 Routes
  - DHCP Relay Hop Count Configuration
  - DHCP Relay Over IPv4 Unnumbered Interfaces
  - DHCP Relay Option 82, Sub Option 151 VRF Name/ID Option
  - DHCP Relay Option 82, Sub Option 5 Link-Selection Option RFC3527
  - Support for Circuit-Id Formats
  - DHCP Relay Circuit-Id Option
  - DHCP Relay and DHCP snooping support on the same VLAN
- Dynamic BGP Neighbor
- IP Fabric over IPv6 underlay RFC5549
- IP Helper
- Route Leaking across VRFs including Management VRF
- BGP Docker Warm Restart
- Avoid Netlink for Handling IPv6 Link-Local Address
- BFD Optimizations to Support 5x100msec Aggressive Timers in SW
- IP SLA (ICMP and TCP tracker)
- IPv4 Unnumbered Interfaces
- RPVST+ over MC-LAG
- RPVST+ Scaling to 3500 VLAN Ports
- Symmetric Hashing
- BFD IS-CLIs
- BFD with VRF
- VRF support for syslog
- VRF support for SSH.in
- VRRPv3, VRRP/VRRPv3 over VRF
- Management VRF Hardening
- OSPFv2
- Multi Site Data Center Interconnect (DCI)
- RIB/FIB Consistency Checker
- Next Hop Group (NHG) Support
- RIF Counters for L3 Interfaces
- BFD Profile
- 1 Million Route Scale
- Route Updates Performance Improvements
- Drop Neighbor Entry to Protect CPU from Unknown IP Packets Hitting the CPU
- CPU Offload for Neighbor Suppression
- CPU Offload for Slowpath ARP Flooding
- OSFPv2 GR
- Router Advertisement
- KLISH/REST/gNMI Support
  - CLI Commands for RA Retransmission Interval, RADv Disable
  - RFE-8106

ACL and Flow-Based Services

- PBR Enhancements for Service Chaining
- ACL-based CoPP
- ACL DSCP Map/Remarking
- ACL Rate Limiting
- Control Plane ACL
- Policy-based Routing (IPv4 and IPv6)
- ACL-based Packet Replication
- ACL Consistency Checker

Security Features

- RADIUS and TACACS
- RADIUS/TACACS Password Obfuscation
- NTP Server and NTP Authentication
- NTP Prefer Option
- LDAP Integration
- AAA Authorization support with TACACS+

Manageability Features

- Industry Standard CLI (IS-CLI)
- REST and gNMI Interfaces via OpenConfig YANG (OC-YANG)
- Role-Based Access Control (RBAC)
- RBAC and HAMd Enhancements
- SNMP Configuration Traps and OIDs
- gNMI Subscription Support for Limited YANG Paths (OnChange, Interval, Once, Poll, Target defined)
- Bulking support in both REST(YANG patch) and gNMI
- Query parameter for REST and filtering support for gNMI
- Scalar encoding support for gNMI
- Support to Read Service Tag via SNMP
- SNMP Trap Enablement on Interface Instead of Global

Multicast Features / Enhancements

- L3 Multicast with PIM operates on L3 interfaces only
- IGMP
- IGMP Snooping (v1, v2, v3)
- IPv4 PIM-SSM Support

Debuggability / Serviceability features

- In-memory Debug Logging
- Audit Logging and Syslogs
- Command to Return Interfaces to the Default Configuration
- Port Mirroring on Port Channel and VLAN

Scalability improvements

- L3 VLAN Scale to 4K
- Host Table Resource Reservation for Local Hosts

QoS

- DSCP Marking Preservation for
- QoS Map Support for Remarking and SVI
- Port and Priority Shaping

Telemetry and Instrumentation

- sFlow on Management VRF
- BST – Watermarks,
- Port MAC Security
- LLDP-MED

Warmboot

Key Specifications	
800GbE (8x100Gb/s)	Yes
400GbE (4x100Gb/s)	Yes
400GbE (8x50Gb/s)	Yes (OSFP)
200GbE (2x100Gb/s)	Yes
200GbE (4x50Gb/s)	Yes
100GbE (1x100Gb/s)	Yes
100GbE (2x50Gb/s)	Yes
100GbE (4x25Gb/s)	Yes
50GbE (1x50Gb/s)	Yes
50GbE (2x25Gb/s)	Yes
1GbE (1x1Gb/s)	OOB only
USB2.0	yes
Console Ports	yes
OOB Management Port	yes
Throughput	51.2Tbps
Latency	<1µs
CPU	Xeon D1736
Memory	32GB
Packet Buffer	165.2MB
Flash Storage	128G
ECMP Groups	4K
MAC Addresses	128K
IPv4 Routes	500K
IPv6 Routes	600K
Max. ACLs	2K
Max. Egress ACLs	512
Hot-Swappable PSUs	1+1 (AC ONLY)
Hot-Swappable FANs	3+1 redundant
Power Budget Typical/Max (F2R SKU)	540W/2330W (18W OSFP, 10W QSFP112)
Size (WxHxD)	438.4 x 650 x 131.8mm
Weight	30Kg
Input Voltage	200-240VAC 3200W 200-277VAC 3200W 48-72V DC 1600W
Output Power	3200W

Switch Hardware SKUs

Product SKU	Description
SSE-T8196S	64 port 400G QSFP112 & 32 port 800G OSFP -200-240VAC PSU, Front-to-rear airflow - Air Cooling
SSE-T8196SR	64 port 400G QSFP112 & 32 port 800G OSFP - 200-240VAC PSU, Rear-to-front airflow - Air Cooling
SSE-T8196D	64 port 400G QSFP112 & 32 port 800G OSFP – 48-72VDC PSU, Front-to-rear airflow - Air Cooling

Switch Software SKUs

Product SKU	Description
SFT-BCM800G-3YR	3 year software support
SFT-BCM800G-5YR	5 year software support

**A 3 year or 5 year software SKU is mandatory per switch*

Cable/Transceiver Parts

Part Number	Length	Description
CBL-NTWK-1107-20M-G	2m	800G OSFP to OSFP, DAC, 2m, 30 AWG
CBL-NTWK-1107-20M-E	2m	800G OSFP to OSFP(close top), DAC, 2m, 30 AWG
CBL-NTWK-1226-35M-H	3.5m	800G OSFP to OSFP AEC, 3.5m, 30 AWG
CBL-NTWK-0976-20M-G	2m	800G OSFP to 2x400G QSFP112, DAC, 2m, 30 AWG
CBL-NTWK-0976-20M-E	2m	800G OSFP to 2x400G QSFP112, DAC, 2m, 30 AWG
CBL-NTWK-1110-20M-G	2m	800G OSFP to 2x400G OSFP RHS, DAC, 2m, 30 AWG
CBL-NTWK-1115-20M-G	2m	800G OSFP to 4x200G QSFP112, DAC, 2m, 30 AWG
Part Number	**Reach**	**Description**
TRX-1105-VR4-R	upto 50m	400G QSFP112 VR4 Transceiver, MPO12, 50m (host side)
TRX-0964-VR4-R	upto 50m	400G OSFP VR4 Transceiver, MPO12, 50m (host side)
TRX-1107-VR8-R	upto 50m	800G OSFP VR8 Transceiver, Dual MPO12, 50m
TRX-1107-DR8-R	upto 500m	800G OSFP DR8 Transceiver, Dual MPO12, 500m
TRX-1107-FR8-R	upto 2km	800G OSFP FR8 Transceiver, LC/LC, 2km
Part Number	**Length**	**Description**
CBL-NTWK-0973-50M-P	5m	MPO12/MPO12, OM4, MMF, 50/125, 5m APC
CBL-NTWK-0973-100M-P	10m	MPO12/MPO12, OM4, MMF, 50/125, 10m APC
CBL-NTWK-0973-200M-P	20m	MPO12/MPO12, OM4, MMF, 50/125, 20m APC
CBL-NTWK-0973-500M-P	50m	MPO12/MPO12, OM4, MMF, 50/125, 50m APC
CBL-NTWK-1217-30M-P	3m	MPO12/2x MPO12, OM4, MMF, 50/125, 3m APC
CBL-NTWK-1217-50M-P	5m	MPO12/2x MPO12, OM4, MMF, 50/125, 5m APC
CBL-NTWK-1217-70M-P	7m	MPO12/2x MPO12, OM4, MMF, 50/125, 7m APC
CBL-NTWK-1217-100M-P	10m	MPO12/2x MPO12, OM4, MMF, 50/125, 10m APC
CBL-NTWK-0983-150M-P	15m	MPO12/MPO12, OS2, SMF, 15m APC
CBL-NTWRK-0982-150M-P	15m	LC to LC, SMF, 15m

Warranty:

SSE-T8196 comes with a standard (3-1-1) warranty which covers 3 years of labor, 1 year of parts and 1 year of cross-shipment warranty. The warranty can be extended up to total 5 years. For more information, please visit the warranty page.

On-site Service:

Supermicro Hardware Maintenance provides flexible and customizable Service Level Agreements for remote help desk and rapid onsite support to cover Supermicro hardware solutions. Our Onsite Service Programs offer a 4-hour Onsite Response time option for mission-critical uptime or any tailored solution that will meet your specific business requirements. For more details, please visit the Onsite Services page.

MKT-0002-05/2025-R05