

Supermicro SSE-T8164 800G Ethernet Switch

The network bandwidth requirement has been on a constant rise over the years. The recent advancement in the Artificial Intelligence (AI) and Machine Learning (ML) applications has further accelerated the need for higher bandwidth. These new classes of AI workloads have led to development of a wide range of innovations in the AI data center infrastructure ranging from high-speed CPUs, GPUs, accelerator cards, fast storage arrays etc. all of which need fast and efficient interconnects to work cohesively. 800G Ethernet technology provides the higher speed and increased throughput that is needed to address the scaling network connectivity requirements in the large-scale next generation data centers.

Supermicro SSE-T8164 switch is a high through-put, low-latency Ethernet switch that comes with 64 800G OSFP ports and 2 25G SFP28 ports in a compact 2 RU form factor. The switch provides the highest density of 800G Ethernet ports with 51.2 Tbps of switching capacity making it highly suitable for building large flat scale-out network



(Front View)



(Rear View)

fabrics. The SSE-T8164 high-performance switch has been designed to meet the complex requirements of AI data center networks. The cutting-edge switch silicon in the SSE-T8164 and the software comes with advanced features that are needed for ensuring faster job completion time in large AI training networks. The switch can inter-operate with all the latest AI accelerators and industry standard fiber optic and copper cables making it a great choice for all the standard AI cluster designs architectures. The switch supports Direct Attached Copper (DAC) cables up to 4 meters in length which facilitate building dense low-power cost-effective AI training clusters.

The 800G ports in the SSE-T8164 are flexible and can be broken down to 400G, 200G, 100G using breakout cables. The choice of multiple port speeds from 25G to 800G and the advanced software features enables the switch to operate at all the tiers, as a leaf, spine or a border leaf in a multi-tier CLOS network fabric. SSE-T8164 is a versatile switch that is suitable for building networks of all sizes for various use cases, from smaller enterprise data centers to large scale AI/ML training and hyper scale cloud networks.

SSE-T8164 switches run the Supermicro Advanced Enterprise SONiC Operating system which provides a range of features needed for different types of applications such as enterprise, AI/ML, and high-performance computing and cloud computing. In addition to the standard layer 2, layer 3 and tunneling features, it supports advanced features like RoCEv2, packet spraying, Dynamic load balancing coupled with advanced traffic engineering and QoS features making SSE-T8164 is an optimal solution for building large-scale high-performance low-latency AI networks. The Supermicro Enterprise SONiC Operating system can be configured using industry standard CLI, REST/gNMI etc. and can work with all industry standard management applications.

The SSE-T8164 provides both front to back and back to front airflow options. The switch is air cooled and has been designed with Direct Liquid Cooling capability as well which can be used to significantly reduce the power consumption and building energy efficient networks. The switch comes with a choice of AC and DC power supply options, and the included rail kit facilitates simple and easy installation of the switch.

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 800G ports in a 2 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-T8164 is a high radix Ethernet switch with 64 800G ports that can deliver 51.2 Tbps of switching capacity in a compact 2RU form factor that is optimized for building AI clusters of all sizes. The topology above is a representation of how a non-blocking 2-tier scale-out fabric comprising 128 leaf nodes and 64 spine nodes using the 8164 switches can be built to connect 1k GPU systems (8192 GPUs) into a single cluster.

Safety and Compliance	Electromagnetic Compatibility	Environmental
EU: IEC/EN 62368-1 North America: UL62368-1 CAN/CSA No. 62368-1 Taiwan (BSMI): CNS 15598-1	EU: EN 300 386, EN 55032, EN IEC 61000-3-2, EN 61000-3-3, EN 55035, BS 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 C 9832, KS C 9835	EU: RoHS, REACH, WEEE Taiwan (BSMI): CNS 15663

Hardware Specifications

Key Switch Specs

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

Physical Ports

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

CPU Engine

- Intel Xeon-D x86 CPU, 8C at 2.3GHz
- 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
- 2RU, Mountable in 19" or 21" racks
- Dimensions: (WxDxH) 438.4 x 650 x 87.4mm
- Weight: 26kg
- 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: Up to 18W transceivers.

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
- 802.3-2022 25 Gigabit Ethernet
- 802.3cd 50 Gigabit Ethernet
- 802.3ba 100 Gigabit Ethernet
- 802.3bs 400 and 200 Gigabit Ethernet
- 802.3cm 400 Gigabit over multimode fiber
- 802.3df 800 Gigabit Ethernet
- 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

Supermicro Enterprise SONiC Features

System and Platform Infrastructure Features

- Dynamic Port Breakout
- DOM Information Display
- Port Locator Support (use port LED to identify
 NTP Prefer Option the port)
- CMIS 4.0 Optics Support
- Hardware Watchdog
- CoPP (Control Plane Policing)
- 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 •
- 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 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
- 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, OSPFv2 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) •

© Copyright Super Micro Computer, Inc. All rights reserved

- ACL-based Packet Replication
- ACL Consistency Checker

Security Features

- RADIUS and TACACS
- RADIUS/TACACS Password Obfuscation
- NTP Server and NTP Authentication •
- 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 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

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 VLA

Scalability improvements

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

OoS

- 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

AI/ML Focused Features

- RoCEv2
- PFC
- FCN
- **Enhanced DCQCN**
- Dynamic Load Balancing (DLB) •
- Packet Spraying •
- Enhanced Hashing and Load Balancing with 100mSec granularity
- **Rich Congestion and Load Balancing** Telemetrv

Layer 2 Features

 VLAN Auto-state Interface Hold-Down

PVST and RPVST+

DHCP snooping

enhancement

Layer 3 Features

loopback)

Interfaces

Name/ID Option

the same VLAN

support for gNMI

Dynamic BGP Neighbor

Management VRF

BGP Docker Warm Restart

Aggressive Timers in SW

IP SLA (ICMP and TCP tracker)

RPVST+ Scaling to 3500 VLAN Ports

VRRPv3, VRRP/VRRPv3 over VRF

Multi-Site Data Center Interconnect (DCI)

SUPERMICRO

Management VRF Hardening

RIB/FIB Consistency Checker Next Hop Group (NHG) Support

RIF Counters for L3 Interfaces

IPv4 Unnumbered Interfaces

RPVST+ over MC-LAG

Symmetric Hashing

VRF support for syslog

VRF support for SSH.in

BFD IS-CLIs

OSPFv2

BFD Profile

BFD with VRF

and gNMI

of Global

Warmboot

IP Helper

Address

•

•

•

•

.

Overlay Interfaces

with RFC5549 Routes

DHCP Relay Enhancements

link Tracking

Static LAG

LLDP

MSTP

•

UDLD

LACP Graceful Shutdown

LACP Fast Rate and LACP Fallback

Port Channel Min-Links configuration

DHCP Relay Source Interface Selection (e.g.

DHCP Relay over IPv6 Link-Local Interfaces

DHCP Relay Option 82, Sub Option 151 VRF

DHCP Relay Option 82, Sub Option 5 Link-Selection Option RFC3527

DHCP Relay and DHCP snooping support on

Limited YANG Paths (OnChange, Interval,

Bulking support in both REST(YANG patch)

Query parameter for REST and filtering

Support to Read Service Tag via SNMP

SNMP Trap Enablement on Interface Instead

Scalar encoding support for gNMI

IP Fabric over IPv6 underlay RFC5549

Route Leaking across VRFs including

Avoid Netlink for Handling IPv6 Link-Local

BFD Optimizations to Support 5x100msec

DHCP Relay Hop Count Configuration

DHCP Relay Over IPv4 Unnumbered

Support for Circuit-Id Formats

DHCP Relay Circuit-Id Option

Once, Poll, Target defined)

Switch Hardware SKUs

Product SKU	Description
SSE-T8164S	64 port 800G OSFP/800, 200-240 VAC PSU, Front-to-rear airflow – Air Cooling
SSE-T8164SR	64 port 800G OSFP/800, 200-240 VAC PSU, Rear-to-front airflow – Air Cooling
SSE-T8164D	64 port 800G OSFP/800, 48-72 VDC 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-T8164S 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 <u>warranty</u> 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 <u>Onsite Services</u> page.

© Copyright Super Micro Computer, Inc. Specifications subject to change without notice. All other brands and names are the property of their respective owners. All logos, brand names, campaign statements and product images contained herein are copyrighted and may not be reprinted and/or reproduced, in whole or in part, without express written permission by Supermicro Corporate Marketing.

SUPERMICRO