


Network Switching Modules for Large Model Training Clusters - Official Technical Overview & Hardware Datasheet

EXECUTIVE SUMMARY

Network Switching Modules for Large Model Training Clusters represent a purpose-built, high-density switching platform engineered to meet the extreme bandwidth, ultra-low latency, and deterministic synchronization demands of AI/ML training fabrics. As large language models and multimodal architectures scale beyond trillions of parameters, the interconnection network becomes the critical determinant of cluster efficiency. These modules deliver non-blocking 400G/800G switching capacity with hardware-accelerated collective communication primitives, enabling GPU clusters to achieve near-linear scaling efficiency across thousands of nodes. The platform integrates advanced congestion management, adaptive routing, and precision time synchronization into a compact modular chassis, providing a resilient and future-proof foundation for next-generation AI infrastructure.



ARCHITECTURE & CHASSIS DESIGN

The switching module chassis employs a modular, orthogonal midplane architecture that separates the data plane, control plane, and power delivery into independent subsystems. This design ensures signal integrity at extreme speeds while enabling field-replaceable line cards, switch fabric modules, and management controllers without disrupting adjacent components. The chassis supports front-to-back or back-to-front airflow configurations, accommodating diverse datacenter thermal profiles. A dedicated out-of-band management plane provides always-on visibility and control, independent of the data plane operating state.

HARDWARE FEATURES

The platform supports up to 64 ports of 800G QSFP-DD or 128 ports of 400G QSFP112 in a 4RU form factor, with a total switching capacity of 51.2 Tbps. Each line card integrates a programmable forwarding ASIC optimized for AI workloads, featuring hardware-accelerated RoCEv2, GPUDirect Async, and collective offload engines for AllReduce, AllGather, and Broadcast operations. The system includes 32GB of packet buffer per line card, dynamic shared buffer allocation, and per-port configurable latency profiles. Redundant 1+1 hot-swappable power supplies and N+1 fan trays ensure continuous operation. Precision Time Protocol (PTP) IEEE 1588v2 hardware timestamping is supported on every port with sub-10ns accuracy.

COMPLIANCE & STANDARDS

The Network Switching Modules for Large Model Training Clusters comply with IEEE 802.3cd, 802.3ck, and 802.3db Ethernet standards, RoCEv2 (RDMA over Converged Ethernet), and PCIe Gen5/Gen6 host interfaces. The platform is certified for NEBS Level 3, UL 60950-1, IEC 62368-1, and FCC Part 15 Class A. It supports SNMPv3, NETCONF/YANG, gNMI, and Redfish management interfaces, ensuring seamless integration into existing orchestration frameworks.

TECHNICAL SPECIFICATIONS

The following table summarizes the key physical, performance, and environmental parameters of the switching module platform. All specifications are subject to change without notice and represent typical values at nominal operating conditions.

Parameter	Specification
Form Factor	4RU Modular Chassis
Switching Capacity	51.2 Tbps
Port Configuration	Up to 64x 800G QSFP-DD or 128x 400G QSFP112
Forwarding ASIC	Programmable, AI-optimized, collective offload capable
Packet Buffer	32GB per line card, dynamic shared allocation
Latency	Sub-600ns port-to-port (cut-through mode)
Power Supply	1+1 Redundant, hot-swappable AC/DC
Cooling	N+1 Redundant fan trays, front-to-back or back-to-front airflow
Management	SNMPv3, NETCONF/YANG, gNMI,

	Redfish
PTP Accuracy	IEEE 1588v2, sub-10ns
Operating Temperature	0°C to 40°C (32°F to 104°F)
Compliance	NEBS Level 3, UL 60950-1, IEC 62368-1, FCC Part 15 Class A

ORDERING OPTIONS

Base chassis configurations are available with or without integrated management modules. Line card options include 800G QSFP-DD, 400G QSFP112, and 200G QSFP56 variants. Optional accessory kits provide redundant power supplies, airflow kits, and rack mounting hardware. Software licenses are available for advanced telemetry, collective offload, and adaptive routing feature sets. Please consult the official ordering guide for specific SKU combinations and availability.

