OEM Bare-Metal SDN Switch Hardware Interface and Compatibility Reference - Official Technical Overview & Hardware Datasheet

PRODUCT IDENTIFICATION

This document serves as the authoritative technical reference for the OEM Bare-Metal SDN Switch Hardware Interface and Compatibility Reference, a family of high-performance, open-networking switching platforms engineered for software-defined network (SDN) deployments. These bare-metal switches are designed to provide maximum flexibility, allowing network operators to deploy their choice of network operating system (NOS) while leveraging a robust, carrier-grade hardware foundation. This datasheet details the physical interfaces, hardware compatibility, and technical specifications essential for system integration and procurement.



CORE INTERNAL TOPOLOGY

The architectural foundation of the OEM Bare-Metal SDN Switch is built upon a high-capacity, merchant-silicon switching ASIC, delivering non-blocking, line-rate performance across all ports. The internal topology is structured to minimize latency and maximize data throughput.

- Switching Fabric: A centralized, shared-memory switching fabric provides full connectivity between all front-panel interfaces, ensuring that any port can communicate with any other port at wire speed without oversubscription.
- CPU Subsystem: An integrated, powerful x86 or ARM-based CPU complex is dedicated to system management, control plane operations, and protocol processing. This subsystem runs the user-selected NOS and is isolated from the data plane to ensure deterministic forwarding performance.
- Memory Architecture: Packet buffers are strategically allocated across ingress and egress stages to absorb micro-bursts and ensure smooth traffic flow. Management memory includes high-speed DDR4 RAM and persistent storage (e.g., eMMC or SSD) for OS and configuration files.
- Physical Interfaces: All data plane interfaces are routed from the ASIC directly to the front-panel connectors, utilizing controlled-impedance PCB traces to maintain signal integrity for high-speed SERDES (Serializer/Deserializer) links.

SWITCHING SUBSYSTEM CAPABILITIES

The switching subsystem is the heart of the platform, providing a comprehensive set of hardware-accelerated features that offload the CPU and enable high-performance networking.

- Forwarding Engine: The ASIC employs a ternary content-addressable memory (TCAM) for high-speed lookups, supporting advanced packet classification for Access Control Lists (ACLs), Quality of Service (QoS), and virtualization. The forwarding database (FDB) is capable of storing up to 128,000 MAC addresses, while the routing table supports up to 32,000 IPv4 and 16,000 IPv6 routes.
- Programmable Pipeline: The switch pipeline is highly programmable, allowing for the implementation of custom matching and action logic. This is critical for SDN use cases, enabling operators to define novel forwarding behaviors via OpenFlow or the switch's native API.
- Queueing and Scheduling: Hierarchical QoS (HQoS) capabilities allow for fine-grained traffic management. Each physical port supports up to 8 hardware queues, enabling sophisticated traffic shaping, policing, and scheduling algorithms (e.g., Strict Priority, Weighted Round Robin, Deficit Weighted Round Robin) to ensure service-level agreements (SLAs) are met.

OPERATIONAL ADVANTAGES

- Software Flexibility: Freedom to choose and deploy the most suitable NOS, including but not limited to SONiC, Pica8, IP Infusion, and traditional open-source Linux-based stacks, preventing vendor lock-in.

- Cost Efficiency: Competitive total cost of ownership (TCO) by decoupling hardware from software licensing. High port density reduces the number of devices required in a rack, optimizing space and power consumption.

- Scalability: Designed for large-scale leaf-spine architectures. The hardware supports advanced networking features like Virtual Extensible LAN (VXLAN), Segment Routing, and MPLS, facilitating the creation of highly scalable and resilient networks.

- Reliability and Redundancy: Hardware-level failover redundancy is achieved through features such as Link Aggregation Control Protocol (LACP) for port-level redundancy, a redundant hot-swappable power supply, and multiple fan modules for thermal resilience.

SPECIFICATION MATRIX

Parameter	**Specification**
Form Factor	1RU Standard Rackmount
Port Configuration	48 x 25G SFP28, 8 x 100G QSFP28
Switching Capacity	Up to 2.0 Tbps (Full-Duplex)

Throughput	Up to 2.97 Bpps (Packets per second)
CPU Complex	Intel Xeon D-1527 (4-Core, 2.2GHz) or equivalent
System Memory	16GB DDR4
Storage	64GB eMMC / 240GB SSD (optional)
Power Supply	Dual, Hot-Swappable Redundant AC/DC (Platinum-rated)
Power Consumption	Typical: 150W; Max: 250W
Cooling	4+1 Hot-Swappable Fan Modules (Front-to-Back Airflow)
Operating Temperature	0°C to 45°C
Operating Humidity	10% to 90% (non-condensing)
MTBF	> 250,000 hours
Management Interfaces	1 x RJ45 Serial Console, 1 x 1G Management Ethernet, USB Type-A

REGULATORY COMPLIANCE

The OEM Bare-Metal SDN Switch is designed and certified to meet a comprehensive suite of global regulatory and environmental standards, ensuring safe and compliant operation in any datacenter or enterprise

environment.

- Safety Standards: UL 60950-1, CSA C22.2 No. 60950-1, EN 60950-1, IEC 60950-1 (CB Scheme).

- EMC Standards: FCC Part 15 Subpart B (Class A), ICES-003 (Class A), EN 55032 (Class A), EN 55035, EN 61000-3-2, EN 61000-3-3.

- Environmental: RoHS (Restriction of Hazardous Substances) and WEEE (Waste Electrical and Electronic Equipment) compliant.

- NEBS: Optional Level 3 compliance available upon request for carrier-grade installations.

