240G Capacity Hardware Processing - Official Technical Overview & Hardware Datasheet

PRODUCT OVERVIEW & MARKET POSITIONING

The modern network edge and campus aggregation layer demand deterministic, high-throughput processing at a accessible power envelope. The 240G CAPACITY HARDWARE PROCESSING platform delivers wire-speed, non-blocking Layer 2 switching and Layer 3 routing with a total switching fabric capacity of 240 Gigabits per second. Designed for small-to-medium enterprise cores, regional carrier aggregation, and high-performance data center top-of-rack (ToR) deployments, this system bridges the gap between legacy 1G/10G infrastructures and next-generation 25G/100G uplinks.



HARDWARE ARCHITECTURE & FORWARDING ENGINE SILICON

The platform is built around a purpose-designed network processor (NPU) and packet forwarding engine (PFE) capable of line-rate processing on all active interfaces. The internal switching fabric employs a shared-memory architecture with output queuing, ensuring predictable latency under full 240G bidirectional load. The hardware supports cut-through and store-and-forward switching modes with a configurable maximum transmission unit (MTU) of up to 12,000 bytes for jumbo frame efficiency in storage or backup environments.

KEY FEATURES & ROUTING CAPABILITIES

- Aggregate switching capacity: 240 Gbps full-duplex
- Forwarding rate: Up to 180 Mpps (million packets per second) for 64-byte frames
- Hardware-based IPv4/IPv6 routing: Static, RIP, OSPF, BGP, and IS-IS offload
- Flow-based hashing for ECMP and Link Aggregation Groups (LAG) up to 8 ports per group
- Multi-stage priority queuing: 8 hardware queues per port with strict priority or weighted fair queuing (WFQ)
- Data center bridging features: PFC, ETS, and DCBX support for converged Ethernet (FCoE readiness)
- Hardware telemetry: sFlow, NetFlow/IPFIX line-rate export without CPU intervention

REDUNDANCY, COOLING, AND PHYSICAL INTERFACES

Active-active redundant hot-swappable AC/DC power supplies (1+1) and N+1 variable-speed fan trays maintain carrier-grade uptime. Front-to-back airflow aligns with standard hot-aisle/cold-aisle data center designs. The base configuration includes:

- 24 x SFP+ ports (1G/10G) supporting SR, LR, and copper SFP transceivers
- 4 x QSFP28 ports configurable as 40G or 100G uplinks (breakout support 4x25G)
- 1 x RJ45 console port; 1 x USB-C for file transfer; 1 x 10/100/1000BASE-T management Ethernet

Parameter	**Specification**
Form Factor	1RU Standard Rack-mountable (19-inch)
Switching Capacity	240 Gbps (full-duplex non-blocking)
Power Supply	Dual Hot-Swappable Redundant AC (100-240V) or DC (-48V)
Port Configuration	24x SFP+ (1G/10G) + 4x QSFP28 (40G/100G or 4x25G breakout)
Forwarding Rate	180 Mpps (64-byte packets)

Latency (10G)	1.5 µs (cut-through), 4 µs (store-and-forward)
Routing Table (IPv4)	64,000 prefixes (hardware FIB)
Operating Temperature	0°C to 45°C (short-term -5°C to 55°C)
MTBF	310,000 hours (Telcordia SR-332)

COMPLIANCE & STANDARDS CERTIFICATIONS

The 240G platform complies with NEBS Level 3 (Telcordia GR-63-CORE, GR-1089-CORE), ETSI EN 300 386, and FCC Part 15 Class A. Safety certifications include IEC 62368-1, UL/CSA, and CE marking. The hardware is MEF 3.0 compliant for Carrier Ethernet services including E-Line, E-LAN, and E-Tree. All optical interfaces conform to SFF-8431 (SFP+) and SFF-8665 (QSFP28) specifications.

TECHNICAL SPECIFICATIONS TABLE (Detailed)

- System backplane: 480 Gbps (full-duplex fabric oversubscription ratio 2:1 at chassis level, 1:1 at line-rate ports)

- CPU subsystem: Quad-core ARM Cortex-A72 at 1.5 GHz, 4 GB DDR4 RAM, 8 GB eMMC flash

- Latency (10G SFP+): < 1.5 microseconds (cut-through), < 4 microseconds

(store-and-forward)

- MAC address table: 128,000 entries

- IPv4 routing table (FIB): 64,000 prefixes (hardware); IPv6 FIB: 32,000 prefixes

- ACL / QoS entries: 4,000 ingress + 2,000 egress (TCAM-based)

- Operating temperature: 0°C to 45°C (short-term -5°C to 55°C)

- Acoustic noise: < 55 dBA at 25°C ambient

- MTBF: 310,000 hours (Telcordia SR-332, Issue 4)

ORDERING SKU INFORMATION (Extract)

- ACC-240G-BASE: Chassis, backplane, fan tray, no PSU, no transceivers

- ACC-240G-PSU-AC: 300W AC power supply (hot-swappable)

- ACC-240G-PSU-DC: 300W DC power supply (-48V)

- ACC-240G-24S: 24-port SFP+ line card (optional modular variants)

- ACC-240G-4Q28: 4-port QSFP28 uplink module

- ACC-240G-FAN-01: Spare fan tray unit



DEPLOYMENT ARCHITECTURE SUMMARY

In a typical campus backbone or data center leaf-spine topology, two 240G platforms can form a stacked virtual chassis (up to 4 units) providing 960 Gbps of logical switching capacity. MLAG (multi-chassis link aggregation) ensures server or switch dual-homing without spanning tree blocking. For service provider environments, the hardware supports MPLS/VPLS at the edge with hardware-based label switching up to 240 Gbps. The small 1RU form factor and low sub-80W typical power consumption make this platform optimal for distributed aggregation points, remote PoP sites, or greenfield enterprise networks requiring predictable 240G processing without thermal or budgetary escalation.