Border Router Quality of Service: Hierarchical Queuing (HQoS) Implementation - Official Technical Overview & Hardware Datasheet

EXECUTIVE SUMMARY

This document defines the architectural implementation, hardware-level quality of service (QoS) capabilities, and deployment guidelines for the Hierarchical Queuing (HQoS) forwarding engine within carrier-grade border router platforms. As service provider and enterprise WAN edge traffic profiles converge toward multi-service, low-latency, and high-density aggregation, the requirement for deterministic, multi-level traffic scheduling has become a fundamental architectural criterion. The HQoS implementation described herein provides a two-tier and three-tier scheduling hierarchy that enables service providers to enforce per-subscriber, per-service, and per-class bandwidth and priority guarantees at line rate, without compromising the non-blocking forwarding plane. This datasheet details the silicon-level queuing engine, the hardware buffer hierarchy, the operational advantages for MPLS and Carrier Ethernet deployments, and the exact technical parameters required for evaluation and procurement.



HARDWARE ARCHITECTURE & SILICON

The HQoS engine is integrated directly into the forwarding ASIC, leveraging a dedicated on-chip Traffic Manager (TM) that operates independently from the packet parsing and route lookup pipelines. This separation ensures that hierarchical scheduling, shaping, and marking operations do not introduce head-of-line blocking or reduce the line-rate forwarding capacity of the border router.

At the hardware level, the HQoS implementation utilizes a three-tier memory hierarchy:

1. PACKET BUFFER MEMORY: A high-bandwidth, ultra-low-latency on-die SRAM block provides per-queue buffering for burst absorption and

latency-sensitive traffic. The buffer is dynamically shared across ports and queues based on programmable thresholds, preventing a single congested egress port from monopolizing the global buffer pool.

2. QUEUE DESCRIPTOR ENGINE: The descriptor engine maintains the state of each scheduler node, including queue occupancy, weighted fair queuing (WFQ) weights, strict priority status, and shaping token buckets. It supports a hierarchical structure of up to three levels: Port (Level 1), Subscriber/Group (Level 2), and Class of Service (Level 3).

3. SCHEDULING FABRIC: A dedicated scheduling crossbar connects the queue descriptor engine to the egress pipeline. It executes the Hierarchical Token Bucket (HTB) and Deficit Round Robin (DRR) algorithms in parallel, enabling simultaneous shaping and priority scheduling across thousands of queues.

The ASIC supports a physical queue count of 8,192 queues per forwarding complex, with 8 queues per port and up to 8,192 hierarchical scheduler nodes. The scheduling granularity is 1 kbps for shaping and 1 byte for DRR quantum adjustment, ensuring precise bandwidth allocation for voice, video, and business-critical data services.

KEY FEATURES & ROUTING LOGIC

- Multi-Level Hierarchy: Supports 3-level scheduling (Port, Group, Class) with 8

strict priority queues and 8 WFQ/DRR queues per level, enabling full HQoS separation for wholesale, retail, and enterprise VPN services.

- Per-Subscriber Shaping: Enforces Committed Information Rate (CIR) and Peak Information Rate (PIR) at the subscriber aggregation level, eliminating the need for external BRAS or edge policers.

- MPLS-Aware Queuing: The HQoS engine parses MPLS EXP bits and maps them directly to internal queue priorities, ensuring that transport labels and service labels receive consistent scheduling treatment across the core.

- Congestion Avoidance: Integrated Weighted Random Early Detection (WRED) and Explicit Congestion Notification (ECN) operate per-queue, with configurable thresholds that trigger before the queue reaches its maximum depth.

- Latency Guarantee: Strict priority scheduling for voice and real-time control traffic ensures a maximum queuing latency of less than 10 microseconds under full line-rate load, validated at 100 Gbps and 400 Gbps interface speeds.

- Hardware-Based Policing: Dual-rate, three-color marker (trTCM) policers are implemented in hardware at the ingress pipeline, ensuring that non-conforming traffic is marked or dropped before entering the switching fabric.

- Buffer Optimization: The global packet buffer is configurable from 16 MB to 64 MB per line card, with per-queue reserved thresholds and dynamic shared thresholds that adapt to traffic patterns in real time.

COMPLIANCE & STANDARDS

The HQoS implementation complies with the following industry standards and carrier specifications:

- IETF RFC 2474: Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers
- IETF RFC 2597: Assured Forwarding PHB Group
- IETF RFC 3246: An Expedited Forwarding PHB
- IETF RFC 4115: A Differentiated Service Two-Rate, Three-Color Marker with Efficient Handling of In-Profile Traffic
- MEF 10.3: Ethernet Services Attributes Phase 3
- MEF 23.1: Carrier Ethernet Class of Service Phase 2
- ITU-T Y.1564: Ethernet Service Activation Test Methodology
- ITU-T G.8013/Y.1731: OAM Functions and Mechanisms for Ethernet-Based Networks
- IEEE 802.1Q: Virtual Bridged Local Area Networks
- IEEE 802.1ad: Provider Bridges (QinQ)

TECHNICAL SPECIFICATIONS

The following table summarizes the exact hardware and operational parameters of the HQoS-enabled border router platform. These values are critical for capacity planning, service level agreement (SLA) validation, and interoperability testing.

Parameter	**Specification**
Form Factor	2RU / 4RU / 8RU Modular Chassis
Switching Capacity	Up to 25.6 Tbps per chassis (full-duplex)
Forwarding Rate	Up to 9.5 Bpps (billion packets per second)
HQoS Hierarchy	3 levels: Port, Subscriber/Group, Class of Service
Queues per Forwarding Complex	8,192 physical queues, 8 queues per port
Scheduler Nodes	Up to 8,192 hierarchical nodes per ASIC
Strict Priority Queues	8 per port, per hierarchy level
WFQ/DRR Queues	8 per port, per hierarchy level
Shaping Granularity	1 kbps per queue and per scheduler node
Packet Buffer	16 MB to 64 MB per line card,

	dynamically shared
Buffer Memory Type	On-die SRAM with ECC protection
Congestion Avoidance	WRED, ECN, tail-drop with programmable thresholds
Policing	Dual-rate, three-color marker (trTCM) at ingress
MPLS Support	MPLS EXP bit mapping to internal queue priorities
Latency (Strict Priority)	< 10 microseconds under full line-rate load
Interface Options	100GE QSFP28, 400GE QSFP-DD, 10GE SFP+
Power Supply	Dual hot-swappable redundant AC/DC
Redundancy	1+1 power, N+1 fan, redundant route processor (modular)
Operating Temperature	0°C to 45°C (32°F to 113°F)
Storage Temperature	-40°C to 70°C (-40°F to 158°F)
Humidity	5% to 95% non-condensing
MTBF	> 250,000 hours (Telcordia SR-332)
Regulatory Compliance	FCC Part 15 Class A, CE, VCCI, RoHS, WEEE
Telecom Standards	MEF 10.3, MEF 23.1, ITU-T Y.1564,

	ITU-T G.8013/Y.1731
Management	CLI, NETCONF/YANG, SNMPv3, gRPC/gNMI
Warranty	Limited lifetime hardware warranty with advanced replacement

ORDERING SKU INFORMATION

Base System Configurations:

- BR-HQOS-2RU-AC: 2RU chassis, dual AC power, 16x 100GE QSFP28, 8x 400GE QSFP-DD, HQoS engine enabled, 32 MB packet buffer per line card.
- BR-HQOS-4RU-DC: 4RU chassis, dual DC power, 32x 100GE QSFP28, 16x 400GE QSFP-DD, HQoS engine enabled, 64 MB packet buffer per line card.
- BR-HQOS-8RU-MOD: 8RU modular chassis, 4x line card slots, dual AC/DC power, HQoS engine enabled, 64 MB packet buffer per line card, redundant route processor.

Software Licenses:

- LIC-HQOS-BASE: Enables 3-level hierarchical queuing and 8,192 queues per

forwarding complex.

- LIC-HQOS-ADV: Enables MPLS-aware queuing, per-subscriber shaping, and WRED/ECN advanced congestion management.

- LIC-HQOS-SEC: Enables hardware-based trTCM policing and MACsec integration for secured HQoS deployments.

Accessories:

- ACC-2RU-RACK: 2RU rack mount kit, standard 19-inch EIA.

- ACC-4RU-RACK: 4RU rack mount kit, standard 19-inch EIA.

- ACC-PWR-AC-1600: 1600W AC hot-swappable power supply.

- ACC-PWR-DC-2000: 2000W DC hot-swappable power supply.

- ACC-FAN-HQOS: High-efficiency cooling fan module for HQoS line cards.

