


ZXR10 5950 Series Edge Switch QoS Queue Scheduling Policies - Official Technical Overview & Hardware Datasheet

EXECUTIVE SUMMARY

The ZXR10 5950 Series Edge Switch represents ZTE's carrier-grade answer to the escalating demands of multi-service convergence at the network edge. As enterprise and carrier networks converge voice, video, and mission-critical data onto a unified IP fabric, the ability to deterministically manage congestion and enforce service-level guarantees becomes the defining characteristic of a credible edge platform. The ZXR10 5950 Series addresses this challenge through a comprehensive Quality of Service (QoS) architecture that combines hardware-based traffic classification, flexible queue scheduling disciplines, and granular congestion management mechanisms. This document provides a technical deep-dive into the queue scheduling policies implemented on the ZXR10 5950 Series, detailing the operational characteristics of Strict Priority (SP), Weighted Round Robin (WRR), Deficit Weighted Round Robin (DWRR), and Weighted Fair Queuing (WFQ), alongside their practical application in real-world service provider and enterprise edge deployments. The platform supports up to eight hardware queues per port, enabling precise differentiation of up to eight distinct traffic classes, ensuring that latency-sensitive applications such as VoIP and real-time video receive prioritized treatment while bulk data flows are served with weighted fairness .



ARCHITECTURE & CHASSIS DESIGN

The ZXR10 5950 Series is engineered as a compact 1RU form-factor switch, optimized for space-constrained edge locations while delivering the port densities and forwarding capacities expected of a carrier-class aggregation device. The chassis supports a range of fixed-port configurations, including models with 24 or 48 Gigabit Ethernet access ports, complemented by 4 or more 10 Gigabit Ethernet uplink interfaces, with select models offering 40GE expansion capability through dedicated expansion slots . This physical architecture is purposefully aligned with the QoS processing pipeline: the switch's ASIC-based forwarding engine performs line-rate traffic classification and queue assignment at ingress, ensuring that QoS policies are enforced without introducing forwarding latency or throughput degradation. The

separation of the data plane from the control and management planes further ensures that queue scheduling operations remain unaffected by control-plane processing loads, a critical attribute for maintaining deterministic service quality under sustained traffic pressure . Power and thermal design considerations, including redundant AC or DC power supply options on higher-end models, support continuous operation in edge environments where service interruptions are unacceptable .

HARDWARE FEATURES

QUEUE SCHEDULING DISCIPLINES

The ZXR10 5950 Series implements a rich set of queue scheduling algorithms designed to address the diverse traffic profiles encountered at the network edge. Each physical port supports eight hardware queues, providing the granularity necessary to isolate up to eight classes of service . The scheduling policies available include:

STRICT PRIORITY (SP) QUEUING

SP queuing represents the most deterministic scheduling discipline available on the platform. Under SP operation, packets are classified into priority tiers

(typically four or eight levels), and the scheduler services queues in strict descending order of priority. The highest-priority queue is drained completely before any lower-priority queue receives service. This mechanism is essential for latency-critical traffic such as VoIP bearer streams and real-time control protocols, where any queuing delay beyond the strict priority queue introduces unacceptable jitter. However, SP queuing inherently risks starvation of lower-priority queues during sustained congestion, which is why the ZXR10 5950 Series supports hybrid SP+WRR configurations that allow operators to designate specific queues for strict priority treatment while allocating the remaining queues to weighted scheduling .

WEIGHTED ROUND ROBIN (WRR) QUEUING

WRR addresses the fairness limitations of pure SP scheduling by allocating bandwidth to each queue according to a configured weight value. The ZXR10 5950 Series supports eight WRR queues, each of which is assigned a weight that determines its relative share of the egress interface bandwidth. During each scheduling cycle, the WRR algorithm services queues in proportion to their weights, ensuring that no single queue can monopolize the link while still providing differentiated throughput guarantees. This discipline is particularly well-suited to environments where multiple classes of elastic data traffic—such as business applications, bulk transfers, and best-effort Internet access—must

share available bandwidth equitably according to policy-defined ratios .

DEFICIT WEIGHTED ROUND ROBIN (DWRR) QUEUING

DWRR extends the WRR concept by incorporating packet length into the scheduling decision, thereby improving fairness when traffic classes exhibit heterogeneous packet size distributions. The algorithm maintains a deficit counter for each queue, which is incremented by a quantum value proportional to the queue's configured weight during each scheduling round. The scheduler transmits packets from a queue only while its deficit counter remains non-negative, decrementing the counter by the actual packet length. This byte-oriented approach prevents queues carrying large packets from receiving disproportionate bandwidth relative to their configured weights, a refinement that is particularly relevant in mixed-traffic environments where voice (small packets) and data (large packets) classes coexist .

WEIGHTED FAIR QUEUING (WFQ)

WFQ provides flow-based scheduling that goes beyond class-based differentiation. The ZXR10 5950 Series implements WFQ with support for up to eight flow queues, where packets are classified into flows based on source and destination IP addresses, source and destination port numbers, protocol

number, and IP priority. Each flow receives a queue, and the scheduler allocates egress bandwidth in proportion to the IP priority value of the flow, with lower priority values receiving less bandwidth and higher values receiving more. This mechanism ensures fairness among flows of the same priority while maintaining weighted differentiation across priority levels, making it suitable for applications where per-flow fairness is as important as per-class service differentiation .

COMPLIANCE & STANDARDS

The ZXR10 5950 Series QoS implementation aligns with established industry standards for traffic classification and priority marking. The platform supports 802.1p Class of Service (CoS) mapping at Layer 2 and Differentiated Services Code Point (DSCP) marking at Layer 3, enabling seamless interoperability with upstream and downstream devices that honor these standard priority fields. Priority remarking and mapping functions allow the switch to rewrite CoS and DSCP values at ingress or egress, facilitating policy enforcement across administrative boundaries where trust levels differ. Additionally, the platform supports MPLS EXP bit mapping for environments where MPLS transport is employed at the edge . The switch's QoS capabilities are further complemented by congestion avoidance mechanisms including Tail Drop and Weighted Random Early Detection (WRED), which operate in conjunction with the queue

scheduler to manage buffer occupancy proactively rather than reactively .

TECHNICAL SPECIFICATIONS

The following table summarizes the core QoS and hardware parameters relevant to the ZXR10 5950 Series queue scheduling implementation. Values are representative of the series and may vary by specific model.

Parameter	Specification
Form Factor	1RU fixed-configuration chassis
Hardware Queues per Port	8 queues
Supported Scheduling Disciplines	SP, WRR, SP+WRR, DWRR, WFQ
Traffic Classification Granularity	802.1p CoS, IP DSCP, ACL-based (L2/L3/L4), VLAN ID
Priority Marking & Remarking	802.1p and DSCP remarking; CoS/DSCP/MPLS EXP mapping
Congestion Avoidance	Tail Drop, WRED
Switching Capacity (representative)	288 Gbps to 336 Gbps (model dependent)
Forwarding Rate (representative)	216 Mpps to 252 Mpps (model dependent)
Power Supply Options	AC, DC, HVDC; 1+1 redundancy on

	select models
Operating Temperature	-10°C to 50°C (standard models); -10°C to 55°C (extended range models)

ORDERING OPTIONS

The ZXR10 5950 Series is available in multiple model variants to suit different port density, power, and environmental requirements. Key ordering considerations for QoS-critical deployments include the selection of models with redundant power supplies (AC 1+1 or DC 1+1) for high-availability edge sites, and confirmation that the chosen model supports the full eight-queue-per-port hardware configuration required for advanced scheduling disciplines. PoE+ models should be specified where powered devices such as IP phones and surveillance cameras are to be connected, as these benefit directly from the platform's ability to prioritize voice and video traffic over the same physical infrastructure that delivers power .

