

Cisco Nexus 9300-EX Platform Deep Buffer Allocation Algorithms - Official Technical Overview & Hardware Datasheet

EXECUTIVE SUMMARY

The Cisco Nexus 9300-EX Platform represents a purpose-built fixed-form-factor switching solution engineered for data center environments demanding deterministic buffering behavior under bursty, incast-prone traffic patterns. At the architectural core of this platform lies the Cisco Cloud Scale Leaf Spine Engine (LSE) ASIC, which implements a sophisticated deep buffer allocation framework designed to absorb microbursts while preserving fairness across competing traffic flows. Unlike conventional shared-memory architectures that rely solely on static thresholds or random discard, the Nexus 9300-EX introduces a flow-aware buffer management paradigm anchored by Approximate Fair Drop (AFD) and Dynamic Packet Prioritization (DPP). These algorithms operate in conjunction with a segmented buffer pool architecture, wherein each LSE slice maintains an independent 18.7 MB buffer subsystem, yielding a total per-ASIC buffer capacity of 37.4 MB . The platform is offered in three primary models—the 93180YC-EX, 93108TC-EX, and 93180LC-EX—each delivering wire-rate Layer 2 and Layer 3 forwarding with a consistent 40 MB system buffer allocation . This document provides a definitive technical examination of the buffer allocation mechanisms, algorithmic controls, and operational parameters that distinguish the Nexus 9300-EX as a platform of

choice for latency-sensitive and high-density virtualized workloads.



ARCHITECTURE & CHASSIS DESIGN

The Nexus 9300-EX platform is constructed around the Cisco Cloud Scale LSE ASIC, a multi-slice system-on-chip design that partitions the switching fabric into two independent 800 Gbps slices, collectively delivering 1.6 Tbps of aggregate bandwidth at 1.25 GHz operation . Each slice functions as a self-contained switching subsystem with its own dedicated buffer memory, port scheduler, and forwarding engine. This architectural decision is fundamental to the platform's buffer behavior: ports assigned to a given slice may only access the buffer resources of that slice, ensuring predictable latency and eliminating cross-slice contention for buffer cells . The raw buffer memory within each slice is organized into 208-byte cells, with multiple cells linked together to

accommodate complete packets. This cell-based granularity enables efficient utilization of the 88,000 cells available per slice, translating to 18.7 MB per slice and 37.4 MB total per ASIC . The LSE ASIC supports 18 Class of Service (CoS) designations: 16 user-definable classes, one dedicated to Switched Port Analyzer (SPAN) traffic, and one reserved for CPU-generated control traffic. Software further partitions the buffer into up to four distinct pool groups, with drop and no-drop classes (enabled via Priority Flow Control) assigned to separate storage groups to guarantee isolation and resource availability . This pool segmentation is critical for preserving lossless semantics in lossless classes while simultaneously allowing lossy classes to utilize remaining buffer capacity without interference.

HARDWARE FEATURES

The defining buffer management capabilities of the Nexus 9300-EX are implemented through a suite of intelligent algorithms operating at the ASIC level. Approximate Fair Drop (AFD) represents the primary advancement over legacy Weighted Random Early Detection (WRED) mechanisms. Whereas WRED applies a single early-drop threshold per class-based weighted queue without regard to individual flow behavior, AFD introduces per-flow awareness by tracking active flows through a hash table and measuring byte counts at ingress . The Elephant Trap (ETRAP) mechanism enables AFD to distinguish

short-lived, loss-sensitive mouse flows from long-lived, bandwidth-consuming elephant flows. Flows transmitting fewer bytes than the user-configurable ETRAP threshold are classified as mice and are excluded from early-drop decisions, while flows exceeding the threshold are transitioned to an elephant flow table and subjected to AFD's fair-rate calculations . When an elephant flow packet arrives at an egress queue, AFD compares the flow's measured arrival rate against a dynamically computed fair share derived from queue occupancy feedback. Packets from flows exceeding their fair share are dropped proportionally to the excess rate, ensuring that competing elephant flows converge toward equitable bandwidth distribution . Dynamic Packet Prioritization (DPP) complements AFD by enabling the ASIC to prioritize short flows over long flows within the same traffic class, further reducing latency for latency-sensitive applications without requiring explicit QoS configuration. The platform also supports real-time buffer monitoring per port and per queue, providing operators with granular visibility into burst behavior and application traffic patterns .

COMPLIANCE & STANDARDS

The Cisco Nexus 9300-EX platform adheres to a comprehensive suite of regulatory and safety standards. Safety certifications include UL 60950-1 Second Edition, CAN/CSA-C22.2 No. 60950-1 Second Edition, EN 60950-1

Second Edition, IEC 60950-1 Second Edition, AS/NZS 60950-1, and GB4943 .

Electromagnetic compatibility emissions compliance encompasses 47CFR Part 15 Class A, AS/NZS CISPR22 Class A, CISPR22 Class A, EN55022 Class A, ICES003 Class A, VCCI Class A, EN61000-3-2, EN61000-3-3, KN22 Class A, and CNS13438 Class A. Immunity standards include EN55024, CISPR24, EN300386, and the KN 61000-4 series . The product is RoHS-6 compliant with exceptions for leaded-ball grid-array balls and lead press-fit connectors. NEBS compliance is maintained for deployment in carrier-grade environments. Mean Time Between Failures (MTBF) is rated at 265,570 hours for the 93180YC-EX model, reflecting the platform's suitability for continuous-operation data center deployments .

TECHNICAL SPECIFICATIONS

The following table summarizes the principal hardware and buffer-related specifications for the Cisco Nexus 9300-EX platform.

Parameter	Specification
Form Factor	1RU fixed chassis
ASIC	Cisco Cloud Scale LSE (2 slices, 1.6 Tbps aggregate)
System Buffer (Total)	40 MB

Buffer per LSE Slice	18.7 MB (88,000 cells of 208 bytes)
Buffer Cell Size	208 bytes
CoS Support	18 classes (16 user-defined, 1 SPAN, 1 CPU)
Buffer Pool Groups	Up to 4 (drop, no-drop/PFC, CPU, SPAN)
Congestion Management	AFD (Approximate Fair Drop), DPP (Dynamic Packet Prioritization), WRED
Flow Awareness Mechanism	ETRAP (Elephant Trap) with hash table flow tracking
PFC No-Drop Classes	Up to 3 classes supported via ingress buffer
System Memory	24 GB
SSD	64 GB
CPU	4 cores
Management Ports	1 RJ-45, 1 SFP
Power Supplies	500W AC, 650W AC, 930W DC, or 1200W HVAC/HVDC (1+1 redundant)
Typical Power	210W–290W (model dependent)
MTBF (93180YC-EX)	265,570 hours
Operating Temperature	0°C to 40°C (32°F to 104°F)

ORDERING OPTIONS

The Cisco Nexus 9300-EX platform is available in three distinct hardware configurations, each optimized for specific port-density and speed requirements. The Cisco Nexus 93180YC-EX provides 48 native 1/10/25-Gbps SFP28 ports and 6 QSFP28 ports supporting 40/100-Gbps connectivity, making it suitable for fiber-based server access and aggregation deployments . The Cisco Nexus 93108TC-EX offers 48 100-Mbps/1/10GBASE-T copper ports alongside 6 QSFP28 uplinks, addressing environments where copper cabling infrastructure predominates. The Cisco Nexus 93180LC-EX delivers the industry's first 50-Gbps-capable 1RU switching capability, supporting up to 32 40/50-Gbps QSFP+ ports or 18 100-Gbps QSFP28 ports, with breakout cable support enabling up to 72 10-Gbps connections . All three models share identical CPU, memory, and buffer resources — 4-core CPU, 24 GB system memory, 64 GB SSD, and 40 MB system buffer—ensuring consistent buffer management behavior across the portfolio . Power supply options include 500W AC, 650W AC, 930W DC, and 1200W HVAC/HVDC configurations, with typical power consumption ranging from 210W to 290W depending on model. Airflow is configurable for either port-side intake or port-side exhaust, accommodating diverse rack thermal architectures.

