Non-Blocking Switching Engine Capabilities Blueprint: Switch Backplane Non-Blocking Crossbar Architecture Performance Metrics

DESIGN PHILOSOPHY

The Switch Backplane Non-Blocking Crossbar Architecture represents the foundational pillar of next-generation core switching platforms, engineered to deliver deterministic line-rate forwarding under all traffic conditions. Unlike shared-bus or shared-memory architectures that suffer from head-of-line blocking and internal contention, the non-blocking crossbar fabric provides a dedicated, full-duplex path between every ingress line card and every egress line card. This design philosophy eliminates internal oversubscription, ensuring that aggregate throughput scales linearly with port density and that latency remains bounded regardless of traffic pattern or load distribution.

The architecture is purpose-built for carrier-grade and enterprise core deployments where sustained wire-speed performance, strict quality of service guarantees, and sub-microsecond forwarding latency are non-negotiable requirements. By decoupling ingress and egress forwarding decisions from fabric scheduling, the crossbar engine maintains full forwarding capacity even during bursty, incast, or many-to-one traffic scenarios that routinely degrade competing architectures.



NON-BLOCKING FORWARDING ARCHITECTURE

At the heart of the platform lies a multi-stage, buffered crossbar fabric with a fully distributed arbitration mechanism. Each line card connects to the switch fabric through multiple high-speed serial lanes, typically operating at 25 Gbps, 50 Gbps, or 100 Gbps per lane, aggregated into a parallel fabric interface. The crossbar scheduler operates on fixed-length cells, segmenting variable-length packets at ingress and reassembling them at egress, thereby guaranteeing deterministic switching latency and eliminating packet-size-dependent jitter.

Key architectural attributes include:

- Dedicated Virtual Output Queues (VOQs): Each ingress port maintains per-egress-port queues to completely eliminate head-of-line blocking, a critical

differentiator from single-queue ingress architectures.

- Centralized or Distributed Scheduling: The arbitration engine matches ingress VOQs to egress ports using iterative matching algorithms (such as iSLIP or maximal matching variants) that converge within a bounded number of iterations, ensuring 100% throughput under admissible traffic.

- Speedup and Memory Bandwidth: The fabric is provisioned with a speedup factor of 1.0 or higher, meaning the internal switching capacity exceeds the sum of all external line-rate ports, providing headroom for control plane traffic, multicast replication, and in-band management.

- Cell-Based Cut-Through and Store-and-Forward Modes: Configurable per-port forwarding modes allow operators to trade latency for error detection based on application requirements.

- Multi-Path Redundancy: The crossbar fabric supports redundant switch fabric cards (SFCs) in active-active or active-standby configurations, with hitless failover and sub-50-millisecond convergence.

TRAFFIC SHAPING QOS

The non-blocking crossbar architecture integrates a hierarchical QoS engine directly into the forwarding pipeline. Each egress port supports eight or more hardware queues with strict priority, weighted round-robin, weighted fair queuing, and deficit weighted round-robin scheduling disciplines. The fabric

scheduler is QoS-aware, propagating class-of-service information across the backplane to ensure that high-priority traffic is never starved by lower-priority bursts, even under full fabric load.

Additional QoS capabilities include:

- Per-Port and Per-Queue Shaping: Dual-rate, three-color marking (RFC 2697/2698) with committed and peak information rates.
- Congestion Management: Weighted Random Early Detection (WRED) and Explicit Congestion Notification (ECN) at ingress and egress.
- Buffer Management: Shared-memory packet buffers with dynamic thresholds, allocated per queue and per port to absorb microbursts without dropping delay-sensitive traffic.
- Latency Guarantees: Under full load, the fabric maintains deterministic latency within a bounded window, typically under 2 microseconds port-to-port for cut-through traffic.

PERFORMANCE PARAMETERS

The following performance parameters define the operational envelope of the non-blocking crossbar fabric across all supported chassis configurations and line card densities.

Parameter	Specification
Switching Fabric Type	Buffered Multi-Stage Crossbar with Virtual Output Queuing
Fabric Speedup Factor	1.0 to 2.0 (configurable per platform)
Switching Capacity	Up to 25.6 Tbps full-duplex per chassis
Forwarding Throughput	Up to 14.4 Bpps (billion packets per second)
Port-to-Port Latency	< 1.5 microseconds (cut-through, 64-byte packets)
Latency Under Full Load	< 2.0 microseconds (store-and-forward, 1518-byte packets)
Packet Buffer Memory	Up to 64 MB shared across all ports
VOQ Depth	Up to 128 MB per line card (distributed)
Scheduling Algorithm	Iterative Maximal Matching (iSLIP variant, 4 iterations)
QoS Queues per Port	8 hardware queues, strict priority + DWRR
Redundancy	Dual switch fabric cards, active-active or active-standby, hitless failover

Failover Time	< 50 milliseconds
Multicast Replication	Hardware-based ingress and egress replication, line-rate
Fabric Interface Speed	25 Gbps / 50 Gbps / 100 Gbps per serial lane
Compliance	MEF 3.0, MEF 10.3, MEF 6.2, MEF 35.1, MEF 26.2, IEEE 802.1Q, 802.1ad, 802.3ba

MEF STANDARDS

The platform aligns with Metro Ethernet Forum (MEF) specifications to ensure carrier-class service delivery and interoperability across multi-vendor transport networks. Key supported standards include:

- MEF 3.0: Service definitions for E-Line, E-LAN, E-Tree, and E-Access, with support for bandwidth profiles and Class of Service identifiers.
- MEF 10.3: Ethernet Service Attributes, including per-UNI bandwidth profiles, traffic policing, and service multiplexing.
- MEF 6.2: Ethernet Services Definitions Phase 2, covering EVC and OVC constructs.

- MEF 35.1: Service OAM Performance Monitoring, enabling frame delay, frame delay variation, and frame loss ratio measurement at line rate.

- MEF 26.2: External Network-to-Network Interface (ENNI) and Operator Virtual Connection (OVC) definitions for inter-carrier handoff.

ARCHITECTURE DIAGRAM

