


Ultra-Low Latency Forwarding Engine Factsheet: Telecom Edge Computing

Node Architecture

APPLICATION CONTEXT

The Telecom Edge Computing Node Architecture is a carrier-grade hardware platform engineered to meet the stringent demands of latency-sensitive applications at the network edge. As service providers deploy 5G, industrial IoT, and real-time video analytics, the need for deterministic, ultra-low latency forwarding at the edge becomes paramount. This factsheet details the hardware design, forwarding pipeline, and performance specifications of the edge node, enabling network architects to validate its suitability for demanding edge workloads.



FORWARDING PIPELINE DESIGN

The edge node employs a purpose-built forwarding ASIC with a fully non-blocking, cut-through switching fabric. The pipeline is optimized for minimal jitter and deterministic latency, featuring hardware-accelerated packet parsing, classification, and forwarding. Key design elements include:

- Line-rate forwarding at 400 Gbps per slot with sub-10 microsecond latency.
- Integrated Traffic Manager with 64K queues for granular QoS enforcement.
- Programmable parser for custom protocol support (e.g., SRv6, MPLS, VXLAN).
- On-chip buffer of 128 MB shared across ports to absorb microbursts.

The architecture ensures that latency remains stable under load, making it ideal for edge computing applications such as autonomous vehicle coordination, remote surgery, and augmented reality.

HARDWARE SPECS SUMMARY

The platform is housed in a compact 2RU chassis with redundant power and cooling. It supports a range of interface options, including 100GE, 400GE, and 25GE, with optional MACsec encryption. The system is managed via a secure, out-of-band management interface and supports zero-touch provisioning.

Parameter	Specification
Form Factor	2RU rack-mountable chassis
Switching Capacity	3.2 Tbps full-duplex
Forwarding Latency	< 10 microseconds (cut-through)
Port Density	32 x 100GE QSFP28, 8 x 400GE QSFP-DD
Power Supply	1+1 redundant, hot-swappable AC/DC
Cooling	Front-to-back airflow, redundant fans
Management	1 x 10/100/1000BASE-T out-of-band, console, USB
Operating Temperature	-5°C to 55°C (NEBS compliant)
Dimensions (HxWxD)	88.9 x 444.5 x 600 mm
Weight	18.5 kg (fully loaded)

QUALITY OF SERVICE PROTOCOLS

The edge node implements a comprehensive suite of QoS protocols to guarantee service level agreements. These include:

- Hierarchical QoS (H-QoS) with 3-level shaping and policing.
- Strict priority and weighted fair queuing scheduling.

- WRED and tail-drop congestion avoidance.
- PFC and ECN for lossless Ethernet fabrics.
- IEEE 802.1Qbv Time-Sensitive Networking (TSN) for deterministic delivery.

These features enable precise control over latency, bandwidth, and jitter, ensuring critical traffic is prioritized across the edge network.

COMPLIANCE VERIFICATION

The hardware has been rigorously tested to meet industry standards and regulatory requirements, including:

- NEBS Level 3 and ETSI compliance for carrier deployments.
- IEC 62368-1 safety certification.
- RoHS and WEEE environmental directives.
- CE and FCC marking for global markets.
- MEF 3.0 certification for Carrier Ethernet services.

Compliance ensures seamless integration into existing telecom infrastructures and reduces deployment risks.

SCENARIO VIEW

The following diagram illustrates a typical edge computing deployment, where the node aggregates traffic from 5G small cells and IoT gateways, performs local compute offload, and forwards latency-critical traffic to the core.

