Carrier Ethernet Technical Standard Compliance Register: QinQ Technology

CARRIER ETHERNET TECHNICAL STANDARD COMPLIANCE REGISTER: QINQ (802.1AD) TECHNOLOGY OVERVIEW & HARDWARE SPECIFICATION

EXECUTIVE SUMMARY

This document serves as the official technical compliance and product overview register for QinQ (802.1ad) technology as implemented across our Carrier-Grade Ethernet Switch portfolio. QinQ, also known as Stacked VLANs or VLAN Double Tagging, addresses the critical limitation of the standard 4096 VLAN identifier space by encapsulating a customer's original 802.1Q tag within a second, provider-assigned Service VLAN (S-VLAN) tag. This mechanism enables service providers to deliver transparent, scalable Layer 2 VPN services to thousands of unique customers over a shared infrastructure while preserving each customer's internal VLAN topology.



HARDWARE ARCHITECTURE & SILICON IMPLEMENTATION

The QinQ forwarding plane is processed natively within our proprietary THRUST-ASIC ™ (Application-Specific Integrated Circuit) forwarding engine. Unlike software-based implementations that introduce microsecond-scale latency, the THRUST-ASIC performs dual tag push/pop operations in hardware at line rate. Key internal components include:

- Tag Manipulation Pipeline: Capable of pushing an outer S-VLAN tag upon ingress from a Customer Edge (CE) port, or popping the outer tag upon egress towards the CE.

- TPID Detection Engine: Programmable to recognize both standard 0x8100 and alternative TPID values (0x88A8 or 0x9100) for maximum interoperability.

- MAC Address Table Partitioning: Dedicated memory blocks separate customer MAC learning from provider MAC learning to ensure isolation.

KEY FEATURES & ROUTING LOGIC

- Feature 1: Transparent LAN Service (TLS) – Preserves customer 802.1Q tags (C-VLANs) across the service provider backbone, allowing enterprises to retain full control over their internal VLAN numbering schemes.

- Feature 2: VLAN Scalability – Expands available VLANs from 4096 to 16,777,216 (4096 S-VLANs × 4096 C-VLANs per S-VLAN).

- Feature 3: Selective QinQ – Per-port or per-VLAN configurable mode allowing mixed deployments where some ports perform QinQ tagging while others act as standard trunk ports.

- Feature 4: EtherType Flexibility – Supports custom TPID configurations for legacy equipment or multi-vendor network environments.

- Feature 5: Layer 2 Protocol Tunneling (L2PT) – Allows transparent transport of STP, CDP, VTP, and other customer-side control protocols across the provider backbone.

COMPLIANCE & STANDARDS

This implementation fully complies with IEEE 802.1ad (Provider Bridges). Additional compliance certifications include:

- MEF 6.1 (Ethernet Private Line and Ethernet Virtual Private Line)

- MEF 9 (Abstract Test Suite for Ethernet Services)

- MEF 14 (Abstract Test Suite for Traffic Management)

- RFC 4542 (Provider VLANs)

- RoHS and WEEE environmental directives

TECHNICAL SPECIFICATIONS

Parameter	**Specification**
Form Factor	1RU/2RU Rack-Mountable or Modular 5RU Chassis
Switching Capacity	Up to 1.2 Tbps (non-blocking, full-duplex)
Power Supply	Dual Hot-Swappable Redundant AC (100-240V) or DC (-48V)
Supported Outer Tag TPID	0x8100, 0x88A8, 0x9100 (user-configurable)
Max QinQ VLAN Entries	4096 S-VLANs x 4096 C-VLANs (hardware table limited to 128k entries)
Frame Size	Jumbo frame support up to 9,600 bytes (including double tags)
MTBF	> 320,000 hours at 40°C ambient



OPERATIONAL ADVANTAGES

Service providers deploying QinQ on our hardware achieve three primary operational benefits. First, the elimination of per-customer VLAN planning reduces operational expenditure (OPEX) significantly. Second, the hardware-based tag manipulation ensures zero performance degradation when QinQ is enabled. Third, the ability to offer per-customer QinQ and non-QinQ service modes on the same physical chassis provides unmatched deployment flexibility. For enterprise customers, QinQ resolves the classic "VLAN count exhaustion" problem in large multi-tenant data centers and campus backbones.