

Network Virtualization Foundation: H3C S12500X-AF Data Center Switch

VXLAN Routing Capacity

SDN/NFV HARDWARE ROLE

The H3C S12500X-AF Data Center Cloud Core Switch Series is positioned as the foundational hardware engine for large-scale software-defined data center fabrics. In the context of VXLAN routing capacity, the platform serves a dual role: it operates as a high-performance VXLAN Tunnel Endpoint (VTEP) at the spine layer for distributed Layer 3 gateways, and as a centralized VXLAN gateway for legacy or non-EVPN workloads requiring inter-subnet routing. Unlike fixed-form-factor leaf switches constrained by ASIC table sizes, the S12500X-AF modular chassis architecture decouples control-plane scalability from forwarding-plane density, enabling service providers and large enterprises to sustain routing state for tens of thousands of tenant endpoints without chassis replacement. The platform directly addresses the SDN/NFV requirement for hardware that can terminate overlay tunnels at line rate while maintaining the massive ARP, MAC, and FIB tables necessary for multi-tenant cloud infrastructure.



CENTRALIZED CONTROL INTEGRATION

The S12500X-AF supports the full spectrum of VXLAN control-plane deployment models, ensuring architectural flexibility across brownfield and greenfield data centers. For SDN-centric environments, the platform integrates with OpenFlow and NETCONF controllers, allowing a centralized controller to manage VXLAN tunnel establishment and VNI-to-VSI mappings without relying on dynamic protocol adjacencies. For standards-based distributed control, the switch implements MP-BGP EVPN as the overlay control plane, automating VTEP peer discovery, VXLAN tunnel creation, and VNI association. This EVPN implementation eliminates the need for manual tunnel provisioning and reduces flood-and-learn traffic through BGP Route Reflector hierarchies. The platform further supports IS-IS+ENDP for distributed VXLAN control, providing an alternative interior gateway-based mechanism for environments where BGP

policy complexity is undesirable. Critically, the S12500X-AF can simultaneously function as a Spine node in an EVPN-VXLAN fabric while also acting as a border gateway for Data Center Interconnect (DCI) scenarios, extending VXLAN routing domains across geographically dispersed sites.

VIRTUAL FORWARDING CAPACITIES

The VXLAN routing capacity of the S12500X-AF is defined not merely by tunnel count, but by the interplay of hardware table resources, line-card ASIC selection, and software feature licensing. Independent validation by Tolly Group confirms the platform sustains up to 9,215 IPv4 VXLAN tunnels in tested configurations, with capacity scaling according to the installed line card generation (HB or HF series). For Layer 3 VXLAN gateway functions, the system supports up to 16,000 VXLAN VSI instances and 16,000 VXLAN ACs, enabling granular tenant segmentation at scale. Hardware forwarding tables accommodate up to 1,048,576 ARP entries and 2,097,152 IPv4 FIB routes, ensuring that distributed VXLAN routing does not become table-bound even under dense multi-tenant addressing schemes. MAC address capacity reaches 750,000 entries, sufficient for large-scale virtual machine mobility and EVPN MAC advertisement. The platform also implements VXLAN Overlay security isolation, allowing distinct VLANs—or even identical VLANs—to be logically separated through VNI-based policies, a requirement for multi-tenant cloud environments where overlapping

IP addresses are common. VXLAN OAM tooling, including VXLAN Ping and Tracert, provides operational visibility into overlay path health without requiring external probes.

SYSTEM PARAMETERS

The S12500X-AF family is available in three chassis configurations — S12504X-AF (4 line-card slots), S12508X-AF (8 slots), and S12516X-AF (16 slots) — all sharing the same supervisor engine and line-card portfolio. Switching capacity ranges from 57.6 Tbps to 230.4 Tbps depending on chassis and fabric module generation, with throughput scaling from 28,800 Mpps to 115,200 Mpps. The platform employs a CLOS orthogonal midplane-free architecture, which eliminates the midplane as a signal bottleneck and improves airflow efficiency for the high-density 100GE line cards required in VXLAN spine deployments. Control, detection, and maintenance engines are physically separated: a dedicated control-plane CPU handles protocol processing and VXLAN control-plane adjacencies, while an independent detection engine runs BFD and OAM for millisecond-level failure detection, and a maintenance engine manages power sequencing and environmental monitoring. High availability is achieved through 1+1 supervisor redundancy, N+1 fabric module redundancy, and N+M power supply redundancy, with the platform supporting Hitless VXLAN routing failover through BFD for VRRP and BGP, and NSR/GR for OSPF,

BGP, and IS-IS.

Parameter	Specification
VXLAN IPv4 Tunnel Capacity	Up to 9,215 tunnels (validated configuration)
VXLAN VSI Instances	16,000
VXLAN AC (Attachment Circuit) Capacity	16,000
ARP Entries	Up to 1,048,576
IPv4 FIB Routes	Up to 2,097,152
MAC Address Table	Up to 750,000
Control Plane Options	MP-BGP EVPN, OpenFlow+NETCONF, IS-IS+ENDP
VXLAN Gateway Modes	Distributed (EVPN), Centralized, Layer 2 Bridging
Chassis Options	S12504X-AF (4 slots), S12508X-AF (8 slots), S12516X-AF (16 slots)
Switching Capacity Range	57.6 Tbps to 230.4 Tbps
High Availability	1+1 Supervisor, N+1 Fabric, N+M Power; BFD, NSR, GR
Security	MACsec (802.1ae), VXLAN Overlay Isolation, ACL-based DoS protection

COMPLIANCE MATRIX

The S12500X-AF series is designed to meet the regulatory and interoperability requirements of carrier and large-enterprise data center deployments. The platform supports IEEE 802.1Q, 802.1ad, and 802.1ae (MACsec) for Layer 2 security and encapsulation, with MACsec providing hardware-level link encryption for VXLAN underlay and overlay traffic. For EVPN-VXLAN interoperability, the implementation adheres to RFC 7432 (BGP MPLS-Based Ethernet VPN) and RFC 8365 (Network Virtualization Overlay with EVPN), ensuring multi-vendor fabric compatibility where required. The switch is RoHS compliant and designed with front-to-back airflow to satisfy data center thermal management standards. Management interfaces include CLI, NETCONF, OpenFlow, and SNMPv3, with role-based access control and TACACS+/RADIUS integration for operational security. The platform has been independently validated by Tolly Group for VXLAN Layer 2 and Layer 3 connectivity, IPv6 over IPv4 VXLAN, EVPN with distributed gateways, and centralized VXLAN Layer 3 gateway functionality, providing third-party assurance of protocol conformance and forwarding correctness.

INTEGRATION DIAGRAM

In a canonical EVPN-VXLAN spine-leaf deployment, the S12500X-AF assumes the spine role, providing high-radix 100GE connectivity to leaf switches while terminating VXLAN tunnels from server-facing VTEPs. The spine performs distributed Layer 3 gateway functions, with each leaf acting as a VTEP for its locally attached hosts. VXLAN routing occurs at the spine: packets arriving from one leaf are decapsulated, the inner destination IP is looked up in the global routing table (or VRF), and the packet is re-encapsulated toward the destination leaf. This architecture places the VXLAN routing table scale requirement squarely on the spine, where the S12500X-AF's large ARP and FIB capacities prevent table exhaustion as tenant counts grow. For DCI scenarios, the same chassis can serve as a border spine, mapping VXLAN VNIs to MPLS L3VPN VRFs or extending EVPN routes to a remote data center over a WAN underlay. The platform's support for MDC (Multitenant Device Context) allows a single physical chassis to be partitioned into up to 16 logical switches, enabling separate administrative domains for production, staging, and management VXLAN fabrics on shared hardware—a critical TCO lever for service providers operating multi-tenant cloud infrastructure.

