

## H3C S9855 Series Data Center Switches


### Performance Evaluation and Feature Validation

### Executive Summary

With the rapid development of new-generation technologies such as 5G, cloud computing, and AI, the accelerated integration of information technology and traditional industries, data centers, as the physical carrier for the operation of information systems in various industries, have become indispensable for economic and social operations.

H3C S9855 series switches are a new generation of high-performance, high-density 400GE/100GE Ethernet switches launched by H3C for data centers. Provides high-density 400GE/200GE/100GE ports; supports redundant pluggable power supplies and fans. The S9855 can be used in the core and aggregation networking of the new generation data center. It connects to the S12500 series core switches through 400GE uplinks, and connects to 200GE/100GE servers in the downlink, providing high-bandwidth and large-capacity server access. Tolly engineers evaluated H3C S9855 series data center switches' performance and validated their features.


#### H3C S9855 Series Data Center Switches




S9855-24B8D



S9855-32D



S9855-40B



S9855-48CD8D

Source: Tolly, October 2024

Figure 1

Report link: <https://www.tolly.com/publications/detail/225113>



## Test Results

Tolly engineers tested functions and performance of the H3C S9855 series data center switch (hereinafter referred to as the S9855). Unless specified, test results apply to all S9855 models. For a summary of the performance test cases, refer to Table 1. For a summary of the feature test cases, refer to Table 3 and Table 4 on Page 6 and Page 7. Test results are as follows.

## Performance

### Throughput

Tolly engineers evaluated throughput of the S9855-24B8D switch model with 24\*200GE and 8\*400GE ports using the RFC2544 test in Spirent TestCenter. It supports 100% Layer 2 line-rate for frame sizes of 360-byte and above. Please see Table 1 for detailed results.

### Energy Efficiency (TEER)

According to the ATIS Energy Efficiency for Telecommunication Equipment standard,

TEER (Telecommunications Energy Efficiency Ratio) is defined as a ratio of maximum demonstrated throughput to weighted power.

The H3C S9855-24B8D switch model demonstrated 8,000Gbps maximum throughput. The ATIS weighted power of the S9855-24B8D switch is 314.4W. So the TEER of the S9855-24B8D switch is 25.45Gbps/Watt. Please see Table 2 on Page 3 for detailed results.

## M-LAG

H3C S9855 supports multi-chassis link aggregation group (M-LAG). Two switches in an M-LAG system connects to the uplink or downlink device through links in an M-LAG. It provides high availability for traditional VLAN and new VXLAN overlay networks.

New H3C  
Technologies Co.,  
Ltd.

S9855 Series Data  
Center Switch

Performance  
Evaluation and  
Feature Validation



Tested  
October  
2024

## Link Aggregation

### LAG with Load Balancing

H3C S9855 supports static L2/L3 link aggregation group (LAG) or dynamic L2/L3 LAG. The LAG supports load balancing between member links for unicast, multicast and broadcast traffic.

**H3C S9855 Series Switch Throughput (Percentage of Line-rate)**  
**Dual-snake topology, 24\*200GE and 8\*400GE ports**  
(as reported by Spirent TestCenter)

Frame Size	360-Byte	512-Byte	1024-Byte	1280-Byte	1518-Byte	2048-Byte	4096-Byte	9416-Byte
Layer 2 Throughput	100%	100%	100%	100%	100%	100%	100%	100%

Note: The S9855-24B8D model was used in the test. Throughput was tested with bidirectional traffic using all 200GE ports in a snake topology and all 400GE ports in a snake topology. All test results had 0 frame loss. 100% line-rate means 8Tbps throughput.

Source: Tolly, October 2024

Table 1



## VXLAN

### Basic VXLAN Functions

H3C S9855 supports the VXLAN overlay network which is widely used to build modern campus networks and data center networks.

Tolly engineers verified that H3C S9855 can provide Layer 2 and Layer 3 endpoint connectivity over the VXLAN overlay network. It can provide IPv6 endpoint connectivity over the IPv4 VXLAN overlay network. It can also provide endpoint connectivity between VXLAN and traditional VLAN (non-VXLAN) networks.

### EVPN VXLAN with Distributed Gateways

VXLAN with the BGP-EVPN control plane and distributed gateways is the most popular VXLAN implementation with many advantages including high scalability and easy deployment. H3C S9855 supports EVPN VXLAN with distributed gateways.

S9855 also supports centralized L3 gateway for traditional VXLAN deployment.

### VXLAN DCI

H3C S9855 supports VXLAN Data Center Interconnect (DCI) to provide L2 connectivity between data centers.

### VXLAN Capacity

H3C S9855 supports 8,191 IPv4 VXLAN tunnels.

## Routing

### Routing Protocols

H3C S9855 supports OSPF, IS-IS and BGP IPv4 routing protocols.

H3C S9855 supports OSPFv3, IS-IS for IPv6, and BGP4+ IPv6 routing protocols.

H3C S9855 also supports IPv4 and IPv6 dual-stack network with routing protocols.

### Graceful Restart (GR)

H3C S9855 supports graceful restart for OSPF.

Tolly engineers verified that there was no packet loss when the OSPF routing protocol process was reset on the switch.

### DLB ECMP

H3C S9855 supports Equal-Cost Multi-Path Routing (ECMP) load balancing for high availability and efficient use of available bandwidth.

### Traffic-matrix

H3C S9855 supports traffic-matrix. It can forward traffic according to the configured traffic matrix on the ECMP paths.

### LBN and LBN Enhanced

H3C S9855 supports default ECMP HASH mode and ECMP enhanced mode. In the default ECMP HASH mode, when a ECMP path is down, the switch recalculates to load balance all traffic flows. This may cause sudden high latency for all traffic. In the ECMP enhanced mode, it only recalculates paths for traffic flows previously on the down link.

## BFD

H3C S9855 supports bidirectional forwarding detection (BFD) with 100ms minimum sending interval in echo or control mode.

## H3C S9855 Series Switch Power Consumption (TEER)

(as reported by one AC clamp meter)

	0% utilization power	30% utilization power	100% utilization power	ATIS weighted power	Maximum demonstrated throughput	TEER
S9855-24B8D	280.8W	322.3W	401W	314.4W	8,000Gbps	25.45Gbps/Watt

Note: TEER = (maximum demonstrated throughput) / (ATIS weighted power) = 8,000Gbps / 314.4W = 25.45Gbps/Watt. Test used the same topology as the throughput test in Table 1. The Spirent TestCenter default iMIX profile was used. The switch was fully loaded with two AC power modules and six fan modules. According to the ATIS-0600015.03.2016 standard, ATIS weighted power = 0.1\* (0% utilization power) + 0.8\* (30% utilization power) + 0.1\* (100% utilization power).

Source: Tolly, October 2024

Table 2



## MPLS

### MPLS VPN

H3C S9855 supports MPLS VPWS L2VPN (point-to-point), MPLS VPLS L2VPN (point-to-multipoint), and MPLS L3VPN.

## MCE

### MCE-VRF

Multi-VPN-Instance Customer Edge (MCE) enables a switch to function as the CE of multiple VRF VPN instances to reduce equipment investment. H3C S9855 supports MCE-VRF.

### VPN Routing Protocols

When the H3C S9855 switch acts as the MCE, it supports RIP, OSPF, IS-IS, and BGP routing protocols for the VPN instances (VRFs).

### IP Overlap in Different VPNs

When the H3C S9855 switch acts as the MCE, the IP address in one VPN can overlap with IP addresses in other VPNs to save IP resources.

## O&M

### Telemetry

H3C S9855 supports real-time and high performance gRPC-based dial-in and dial-out Telemetry for network monitoring (e.g. device hardware and traffic statistics).

### INT

H3C S9855 supports Inband Network Telemetry (INT) which collects data from the device and sends data to a collector in real time for performance and network monitoring. INT generates INT packets by mirroring original packets, inserting INT headers, and collecting data.

### Telemetry Stream

H3C S9855 supports Telemetry Stream which collects data from the device and streams data to a collector in real time for performance and network monitoring. Telemetry stream is similar to INT, but is different in that it requires configuration only on the device to collect data from.

### MoD

H3C S9855 supports Mirror on Drop (MoD) to detect packet drops during the forwarding process on the device. It can send the packet drop reason and the characteristics of the dropped packet to the collector.

### Ansible

H3C S9855 can be configured by Ansible for automation.

### NETCONF

H3C S9855 can be controlled by a controller via NETCONF.

### SNMP

H3C S9855 can be managed by a network management system via SNMP v1, v2c and v3.

### OpenFlow

H3C S9855 supports OpenFlow to forward traffic according to the flow table.

### ZTP

H3C S9855 supports Zero Touch Provisioning (ZTP) with the DHCP option and a TFTP server.

### ERSPAN

H3C S9855 supports Encapsulated Remote Switch Port Analyzer (ERSPAN) to mirror

traffic on a port to a remote port using a GRE tunnel.

### RSPAN

H3C S9855 supports Remote Switch Port Analyzer (RSPAN) to mirror traffic on a port to a remote port on another device.

### NetStream

H3C S9855 supports NetStream to monitor traffic flow statistics on the device and send it to a NetStream server.

## Capacity

### MAC Table

H3C S9855 supports up to 224,692 MAC addresses in the MAC table.

### MAC Learning Rate

H3C S9855 supports learning up to 13,500 MAC addresses per second to the MAC table.

### ARP Table

H3C S9855 supports up to 28,669 ARP entries in the ARP table with different MAC addresses.

H3C S9855 supports up to 163,000 ARP entries in the ARP table with the same MAC address.

Tolly engineers verified that the switch forwarded traffic matching all entries in its ARP table, without any packet loss.

### ND Table

H3C S9855 supports 28,000 ND entries in the ND table with different MAC addresses.

H3C S9855 supports 81,000 ND entries in the ND table with the same MAC address.



Tolly engineers verified that the switch forwarded traffic matching all entries in its ND table, without any packet loss.

## FIBv4

H3C S9855 supports up to 980,000 IPv4 routes in the FIB.

Tolly engineers verified that the switch forwarded traffic matching all routing entries in the FIBv4 table, without any packet loss.

## FIBv4 Learning Rate

H3C S9855 supports learning up to 48,600 IPv4 routes per second to the FIB.

## FIBv6

H3C S9855 supports up to 980,000 IPv6 routes in the FIB.

Tolly engineers verified that the switch forwarded traffic matching all routing entries in the FIBv6 table, without any packet loss.

## FIBv6 Learning Rate

H3C S9855 supports learning up to 44,700 IPv6 routes per second to the FIB.

# Features

## PFC Deadlock Prevention

PFC deadlock may happen in certain networks due to a loop or other causes to block forwarding on multiple devices.

H3C S9855 supports deadlock prevention. It can automatically recover from deadlock.

## Mice Elephant Flow QoS

H3C S9855 supports the EM mode for the mice elephant flow scenario. In some data centers, a few elephant flows occupies a large portion of bandwidth and may cause packet loss for lots of mice flows. H3C S9855

can automatically detects the elephant flows and forward mice flows with higher priority.

## DCBX and PFC

H3C S9855 supports Data Center Bridging Exchange (DCBX) and Priority-based Flow Control (PFC) for a lossless data center network.

## ECN

H3C S9855 supports Explicit Congestion Notification (ECN) for a lossless data center network.

## RoCE Lossless Ethernet

H3C S9855 supports RoCE lossless Ethernet with AI ECN and one-click PFC.

# Component Redundancy

## Fan Redundancy

H3C S9855 supports fan module redundancy. When one fan module was unplugged, there was no packet loss.

## Power Redundancy

H3C S9855 supports power module redundancy. When one power module was unplugged, there was no packet loss.

## H3C S9855 Series Data Center Switch

### Tolly Verified Features - Part 1 of 2

M-LAG		Routing	
✓	Multi-chassis Link Aggregation Group (M-LAG)	✓	IPv4 routing protocols: OSPF, IS-IS, BGP
<b>Link Aggregation</b>		✓	IPv6 routing protocols: OSPFv3, IS-IS for IPv6, BGP4+
✓	L2/L3 Static Link Aggregation Group (LAG)	✓	Dual-stack OSPF and OSPFv3; IS-IS and IS-IS for IPv6; BGP and BGP4+
✓	L2/L3 Dynamic Link Aggregation Group (LAG) with LACP	✓	Graceful Restart (GR) OSPF
✓	Load Balancing for unicast, multicast and broadcast traffic with either the static LAG or dynamic LAG	✓	DLB ECMP Dynamic Load-balance Equal-cost Multi-path
<b>VXLAN</b>		✓	Traffic-matrix Forward traffic according to the configured traffic matrix on the ECMP paths
✓	Layer 2 connectivity over the VXLAN overlay network with the same VXLAN Network Identifier (VNI)	✓	LBN (ECMP HASH)
✓	Layer 3 connectivity over the VXLAN overlay network with different VNIs	✓	LBN Enhanced (ECMP Enhanced mode)
✓	IPv6 connectivity over IPv4 VXLAN overlay network	<b>BFD</b>	
✓	Connectivity between VXLAN and non-VXLAN (VLAN only) networks	✓	Minimum BFD sending interval: 100ms in control mode or echo mode
✓	Distributed VXLAN L3 gateway with an IPv4 or IPv6 underlay network	<b>MPLS</b>	
✓	Centralized VXLAN L3 gateway with an IPv4 or IPv6 underlay network	✓	MPLS VPWS L2VPN Point-to-point
✓	EVPN VXLAN Widely used modern network architecture. VXLAN with the BGP EVPN control plane	✓	MPLS VPLS L2VPN Point-to-multipoint
✓	VXLAN Data Center Interconnect (DCI)	✓	MPLS L3VPN
✓	VXLAN tunnel capacity: 8,191 IPv4 VXLAN tunnels	<b>MCE</b>	
		✓	MCE-VRF Multi-VPN-Instance Customer Edge (MCE) enables a switch to function as the CE of multiple VPN instances to reduce equipment investment
		✓	VPN routing protocol support for MCE-VRF RIP/OSPF/IS-IS/BGP
		✓	IP overlap in different VPNs

Source: Tolly, October 2024

Table 3

## H3C S9855 Series Data Center Switch

### Tolly Verified Features - Part 2 of 2

O&M		Capacity	
✓	Telemetry Real-time and high performance gRPC-based dial-in and dial-out Telemetry for network monitoring (e.g. device hardware and traffic statistics)	✓	MAC table: 224,692 MAC addresses
✓	INT	✓	MAC learning rate: 13,500 MAC addresses/sec
✓	Telemetry Stream	✓	ARP table: 28,669 ARP entries (with different MAC addresses) 163,000 ARP entries (with the same MAC address)
✓	Mirror on Drop (MoD)	✓	ND table: 28,000 ND entries (with different MAC addresses) 81,000 ND entries (with the same MAC address)
✓	Ansible	✓	IPv4 FIB: 980,000 routes
✓	NETCONF	✓	IPv4 FIB learning rate: 48,600 routes/sec
✓	SNMP v1, v2c and v3	✓	IPv6 FIB: 980,000 routes
✓	OpenFlow Forwarding traffic according to the flow table	✓	IPv6 FIB learning rate: 44,700 routes/sec
✓	Zero Touch Provisioning (ZTP)	Features	
✓	ERSPAN	✓	PFC deadlock prevention
✓	RSPAN	✓	Mice Elephant Flow QoS
✓	NetStream	✓	DCBX and PFC Data Center Bridging Exchange (DCBX), Priority-based Flow Control (PFC)
		✓	ECN Explicit Congestion Notification
		✓	RoCE Lossless Ethernet (AI ECN + One-click PFC)
		Component Redundancy	
		✓	Fan Redundancy
		✓	Power Redundancy

Source: Tolly, October 2024

Table 4





## About Tolly

The Tolly Group companies have been delivering world-class ICT services for over 35 years. Tolly is a leading global provider of third-party validation services for vendors of ICT products, components and services.

You can reach the company by E-mail at [sales@tolly.com](mailto:sales@tolly.com), or by telephone at +1 561.391.5610.

Visit Tolly on the Internet at:  
<http://www.tolly.com>

## Learn More About H3C

H3C is an industry leader in the provision of Digital Solutions, and is committed to becoming the most trusted partner of its customers in their quest for business innovation and digital transformation. H3C offers a full portfolio of Digital Infrastructure products, spanning across compute, storage, networking, 5G, security and related domains, and provides a comprehensive one-stop digital platform that includes cloud computing, big data, artificial intelligence (AI), industrial internet, information security, intelligent connectivity, new safety, and edge computing, as well as end-to-end technical services. H3C is also the exclusive provider of HPE® servers, storage and associated technical services in China.

For more information, visit: <https://www.h3c.com/en/>

## Terms of Usage

This document is provided, free-of-charge, to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase a product must be based on your own assessment of suitability based on your needs. The document should never be used as a substitute for advice from a qualified IT or business professional. This evaluation was focused on illustrating specific features and/or performance of the product(s) and was conducted under controlled, laboratory conditions. Certain tests may have been tailored to reflect performance under ideal conditions; performance may vary under real-world conditions. Users should run tests based on their own real-world scenarios to validate performance for their own networks.

Reasonable efforts were made to ensure the accuracy of the data contained herein but errors and/or oversights can occur. The test/audit documented herein may also rely on various test tools the accuracy of which is beyond our control. Furthermore, the document relies on certain representations by the sponsor that are beyond our control to verify. Among these is that the software/hardware tested is production or production track and is, or will be, available in equivalent or better form to commercial customers. Accordingly, this document is provided "as is", and Tolly Enterprises, LLC (Tolly) gives no warranty, representation or undertaking, whether express or implied, and accepts no legal responsibility, whether direct or indirect, for the accuracy, completeness, usefulness or suitability of any information contained herein. By reviewing this document, you agree that your use of any information contained herein is at your own risk, and you accept all risks and responsibility for losses, damages, costs and other consequences resulting directly or indirectly from any information or material available on it. Tolly is not responsible for, and you agree to hold Tolly and its related affiliates harmless from any loss, harm, injury or damage resulting from or arising out of your use of or reliance on any of the information provided herein.

Tolly makes no claim as to whether any product or company described herein is suitable for investment. You should obtain your own independent professional advice, whether legal, accounting or otherwise, before proceeding with any investment or project related to any information, products or companies described herein. When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from Tolly.com. No part of any document may be reproduced, in whole or in part, without the specific written permission of Tolly. All trademarks used in the document are owned by their respective owners. You agree not to use any trademark in or as the whole or part of your own trademarks in connection with any activities, products or services which are not ours, or in a manner which may be confusing, misleading or deceptive or in a manner that disparages us or our information, projects or developments.