TEST REPORT

Tolly.

#225131

March 2025

Commissioned by

New H3C Technologies Co., Ltd.

H3C S6850 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.

The S6850 series data center switches are developed by H3C for high-density 25GE switching and aggregation at data centers and cloud computing networks. They can also operate as ToR access switches on an overlay or integrated network. Tolly engineers evaluated H3C S6850 series data center switches' performance and validated their features.

H3C S6850 Series Data Center Switches

S6850-56HF

S6850-2C

Source: Tolly, May 2023

Figure 1

Report Link: https://www.tolly.com/publications/detail/225131

Test Results

Tolly engineers tested functions and performance of the H3C S6850 series data center switch (hereinafter referred to as the S6850). Unless specified, test results apply to both S6850 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 5 on Page 6 to Page 8. Test results are as follows.

Performance

Throughput

Tolly engineers evaluated throughput of the S6850-56HF switch with 48*25GE + 8*100GE ports using the RFC2544 test in Spirent TestCenter. It supports 100% Layer 2 and Layer 3 line-rate for frame sizes of 105-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 S6850-56HF switch model demonstrated 2,000Gbps maximum throughput. The ATIS weighted power of the S6850-56HF switch is 205.9W. So the TEER of the S6850-56HF switch is 9.7Gbps/Watt. Please see Table 2 on Page 3 for detailed results.

IRF2

The Intelligent Resilient Framework 2 (IRF2) technology virtualizes multiple physical devices at the same layer into one virtual fabric to provide data center class availability and scalability. An IRF2 fabric appears as one node for easy management, redundancy, multi-chassis link aggregation with load balancing, scalability and resiliency. The S6850 switch supports IRF2.

ISSU

An S6850 IRF2 fabric supports In-Service Software Upgrade (ISSU) with no packet loss to upgrade IRF2 members' software. ISSU upgrades IRF2 members in sequence and migrate traffic out of the member before upgrading it.

New H3C Technologies Co., Ltd.

S6850 Series Data Center Switch

Performance Evaluation and Feature Validation

Tolly. Certified

Tested May 2023

M-LAG

H3C S6850 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.

GIR

An S6850 M-LAG system supports Graceful Insertion and Removal (GIR) with no packet loss to upgrade M-LAG members' software. GIR upgrades M-LAG members in sequence. GIR maintenance mode automatically switch traffic on one member to the other member before upgrading it.

H3C S6850 Series Switch Throughput (Percentage of Line-rate)
Dual-snake topology, 48*25GE ports and 8*100GE
(as reported by Spirent TestCenter)

Frame Size	105-Byte	128-Byte	256-Byte	512-Byte	1024-Byte	1280-Byte	1518-Byte	9416-Byte
Layer 2 Throughput	100%	100%	100%	100%	100%	100%	100%	100%

Note: The S6850-56HF model was used in the test. Throughput was tested with bidirectional traffic using all ports in a dual-snake topology. All 25GE ports were in a snake topology and all 100GE ports were in another snake topology. All test results had 0 frame loss. 100% line-rate means 2Tbps throughput.

Source: Tolly, May 2023

Table 1

Link Aggregation

LAG with Load Balancing

H3C S6850 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.

VXLAN

Basic VXLAN Functions

H3C S6850 supports the VXLAN overlay network which is widely used to build modern campus networks and data center networks.

Tolly engineers verified that H3C S6850 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 S6850 supports EVPN VXLAN with distributed gateways.

S6850 also supports centralized L3 gateway for traditional VXLAN deployment.

VXLAN Capacity

H3C S6850 supports 15,000 IPv4 VXLAN tunnels.

VXLAN OAM

H3C S6850 supports VXLAN Ping to the VXLAN tunnel destination (loopback interface) address for VXLAN operations, administration, and maintenance (OAM).

Routing

Routing Protocols

H3C S6850 supports OSPF, IS-IS and BGP IPv4 routing protocols.

H3C S6850 supports OSPFv3, IS-IS for IPv6, and BGP4+ IPv6 routing protocols.

H3C S6850 also supports IPv4 and IPv6 dual-stack network with routing protocols.

Graceful Restart (GR)

H3C S6850 supports graceful restart for IPv4 routing protocols.

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

Non-stop Routing (NSR)

H3C S6850 supports non-stop routing for IPv4 routing protocols.

In a Layer 3 network, Tolly engineers verified that there was no packet loss when a master-standby switchover was run in an IRF2 fabric of S6850 switches.

The difference between GR and NSR is that GR can only work properly when the device neighbors also support GR. NSR does not have this requirement.

BFD

H3C S6850 supports bidirectional forwarding detection (BFD) with 3ms minimum sending interval in control mode.

H3C S6850 Series Switch Power Consumption (TEER)

(as reported by Chroma Programmable AC Source)

	0% utilization power	30% utilization power	100% utilization power	ATIS weighted power	Maximum demonstrated throughput	TEER
S6850-56HF	190W	204.5W	233.3W	205.9W	2,000Gbps	9.7Gbps/Watt

Note: TEER = (maximum demonstrated throughput) / (ATIS weighted power) = 2,000Gbps / 205.93Watts = 9.712Gbps/Watt. Test used the same topology as the throughput test in Table 1. The Spirent TestCenter default iMIX profile was used. 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, May 2023 Table 2

MPLS

MPLS VPN

H3C S6850 supports MPLS VPWS L2VPN (point-to-point), MPLS VPLS L2VPN (point-to-multipoint), and MPLS L3VPN.

MPLS TE

H3C S6850 supports MPLS Traffic Engineering (MPLS TE).

MPLS MCE

MPLS Multi-VPN-Instance Customer Edge (MPLS MCE) enables a switch to function as the CE of multiple VPN instances in a BGP/MPLS VPN network to reduce equipment investment. H3C S6850 supports MPLS MCE.

VPN Routing Protocols

When the H3C S6850 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 S6850 switch acts as the MCE, the IP address in one VPN can overlap with IP addresses in other VPNs to save IP resources.

Ring Protection Protocols

RRPP

H3C S6850 supports Rapid Ring Protection Protocol (RRPP) with less than 40ms failover time.

ERPS

H3C S6850 supports Ethernet Ring Protection Switching (ERPS) with less than 40ms failover time.

O&M

Telemetry

H3C S6850 supports real-time and high performance gRPC-based Telemetry for network monitoring (e.g. device hardware and traffic statistics.)

INT

H3C S6850 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.

Telemetry Stream

H3C S6850 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 S6850 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.

PTP

H3C S6850 supports working as the Precision Time Protocol (PTP) server or client. It supports 1588v2, 802.1AS, and SMPTE 2059 standards.

Ansible

H3C S6850 can be configured by Ansible for automation.

Puppet

H3C S6850 can be configured from the Puppet server using scripts for automation.

NETCONF

H3C S6850 can be controlled by a controller via NETCONF.

SNMP

H3C S6850 can be managed by a network management system via SNMP v1, v2 and v3..

ZTP

H3C S6850 supports Zero Touch Provisioning (ZTP) with the DHCP option and a TFTP server.

ERSPAN

H3C S6850 supports Encapsulated Remote Switch Port Analyzer (ERSPAN) to mirror traffic on a port to a remote port using a GRE tunnel.

RSPAN

H3C S6850 supports Remote Switch Port Analyzer (RSPAN) to mirror traffic on a port to a remote port on another device.

iNQA

H3C S6850 supports Intelligent Network Quality Analyzer (iNQA) which measures the service packets directly to calculate packet loss results, thus reflecting the real network quality.

NetStream

H3C S6850 supports NetStream to monitor traffic flow statistics on the device.

Capacity

MAC Table

H3C S6850 supports up to 289,149 MAC addresses in the MAC table.

MAC Learning Rate

H3C S6850 supports learning up to 2,714,843 MAC addresses per second to the MAC table. (Note: The result is reported by the Spirent TestCenter RFC2889 address learning rate test. It shows the rate of learning 70,000 MAC addresses.)

ARP Table

H3C S6850 supports up to 272,000 ARP entries in the ARP table.

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

ND Table

H3C S6850 supports 136,000 ND entries in the ND table.

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

FIBv4

H3C S6850 supports up to 324,000 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 S6850 supports learning up to 6,600 IPv4 routes per second to the FIB.

FIBv6

H3C S6850 supports up to 162,000 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 S6850 supports learning up to 16,500 IPv6 routes per second to the FIB.

Features

DCI

H3C S6850 supports Data Center Interconnect (DCI) for VXLAN fabrics. Layer 2 known unicast, broadcast, unknown unicast and Layer 3 traffic was all forwarded between two data centers properly.

DLB ECMP

H3C S6850 supports Equal-Cost Multi-Path Routing (ECMP) load balancing for high availability and efficient use of available bandwidth.

DLB Link-aggregation

H3C S6850 supports link-aggregation load balancing for high availability and efficient use of available bandwidth.

TCB

H3C S6850 supports Transient Capture Buffer (TCB) to monitor packet drop events on a set of memory management unit (MMU) resources. When a packet is dropped on a queue, the system collects drop time, drop reason, packet metadata, and other information and reports them to a network management system through gRPC.

FCoE

H3C S6850 supports FCoE for the storage network.

Micro-segment

H3C S6850 supports micro-segment. Administrators can group IP addresses into a microsegment and use ACL to define access permissions between the micro-segments.

PBR Service Chain

H3C S6850 supports Policy-Based Routing (PBR) for service chain.

PFC Deadlock Detection

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

H3C S6850 supports deadlock detection. So administrators can recover the network from PFC deadlock after receiving the alarm.

Mice Elephant Flow QoS

H3C S6850 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 S6850 can automatically detects the elephant flows and forward mice flows with higher priority.

DCBX and PFC

H3C S6850 supports Data Center Bridging Exchange (DCBX) and Priority-based Flow Control (PFC) for a lossless data center network.

ECN

H3C S6850 supports Explicit Congestion Notification (ECN) for a lossless data center network.

RoCE Lossless Ethernet

H3C S6850 supports RoCE lossless Ethernet with AI ECN and one-click PFC.

OpenFlow

H3C S6850 supports OpenFlow to forward traffic according to the flow table.

MACsec

The H3C S6850-2C model provides expansion slots. With the LSWM124XGT2Q card, it supports MACsec to encrypt traffic for security.

Component Redundancy

Fan Redundancy

H3C S6850 supports fan module redundancy. When one fan module was unplugged, there was no packet loss.

Power Redundancy

H3C S6850 supports power supply redundancy. When one power supply was unplugged, there was no packet loss.

Mixed Power Supply Modules

H3C S6850 supports using the AC power supply module and the DC power supply module together on the same switch.

H3C S6850 Series Data Center Switch
Tolly Verified Features - Part 1 of 3

IRF2 Stack		VXLAN	
✔	IRF2	✔	Layer 2 connectivity over the VXLAN overlay network with the same VXLAN Network Identifier (VNI)
✔	ISSU	✔	Layer 3 connectivity over the VXLAN overlay network with different VNIs
M-LAG		✔	IPv6 connectivity over IPv4 VXLAN overlay network
✔	Multi-chassis Link Aggregation Group (M-LAG)	✔	Connectivity between VXLAN and non-VXLAN (VLAN only) networks
✔	Graceful Insertion and Removal (GIR) No packet loss to upgrade M-LAG members' firmware. GIR upgrades M-LAG members in sequence. M-LAG maintenance mode automatically switch traffic on one member to the other member before upgrading it	✔	Distributed VXLAN L3 gateway with an IPv4 or IPv6 underlay network
Link Aggregation		✔	Centralized VXLAN L3 gateway with an IPv4 or IPv6 underlay network
✔	Static Layer 2 and Layer 3 Link Aggregation Group (LAG)	✔	EVPN VXLAN Widely used modern network architecture. VXLAN with the BGP EVPN control plane
✔	Dynamic Layer 2 and Layer 3 Link Aggregation Group (LAG) with LACP	✔	VXLAN tunnel capacity: 15,000 IPv4 VXLAN tunnels
✔	Load Balancing for unicast, multicast and broadcast traffic with either the static LAG or dynamic LAG	✔	VXLAN OAM Simulate the ping from a source endpoint to a destination endpoint via the VXLAN tunnel for VXLAN operations, administration, and maintenance (OAM)

Source: Tolly, May 2023 Table 3

H3C S6850 Series Data Center Switch
Tolly Verified Features - Part 2 of 3

Routing		Ring Protection Protocols	
✔	IPv4 routing protocols: OSPF, IS-IS, BGP	✔	Rapid Ring Protection Protocol (RRPP) less than 40ms failover time
✔	IPv6 routing protocols: OSPFv3, IS-IS for IPv6, BGP4+	✔	Ethernet Ring Protection Switching (ERPS) less than 40ms failover time
✔	Dual-stack OSPF and OSPFv3; IS-IS and IS-IS for IPv6; BGP and BGP4+	**O&M**	
✔	Graceful Restart (GR)	✔	Telemetry Real-time and high performance gRPC-based Telemetry for network monitoring (e.g. device hardware and traffic statistics)
✔	Non-stop Routing (NSR)	✔	INT
BFD		✔	Telemetry Stream
✔	Minimum BFD sending interval: 3ms in control mode	✔	Mirror on Drop (MoD)
MPLS		✔	Precision Time Protocol (PTP) server and client 1588v2, 802.1AS, and SMPTE 2059 standards
✔	MPLS VPWS L2VPN Point-to-point	✔	Ansible
✔	MPLS VPLS L2VPN Point-to-multipoint	✔	Puppet
✔	MPLS L3VPN	✔	NETCONF
✔	MPLS TE (Traffic Engineering)	✔	SNMP v1, v2 and v3
✔	MPLS MCE MPLS Multi-VPN-Instance Customer Edge (MPLS MCE) enables a switch to function as the CE of multiple VPN instances in a BGP/MPLS VPN network to reduce equipment investment	✔	Zero Touch Provisioning (ZTP)
		✔	ERSPAN
		✔	RSPAN
✔	VPN routing protocol support for MPLS MCE RIP/OSPF/IS-IS/BGP	✔	iNQA
✔	IP overlap in different VPNs	✔	NetStream

Source: Tolly, May 2023

Table 4

H3C S6850 Series Data Center Switch
Tolly Verified Features - Part 3 of 3

Capacity		✔	FCoE
✔	MAC table: 289,149 MAC addresses	✔	Micro-segment
✔	MAC learning rate: 2,714,843 MAC addresses per second to the MAC table. (Note: The result is reported by the Spirent TestCenter RFC2889 address learning rate test. It shows the rate of learning 70,000 MAC addresses)	✔	Service Chain
✔	ARP table: 272,000 ARP entries	✔	PFC deadlock detection
✔	ND table: 136,000 ND entries	✔	Mice Elephant Flow QoS
✔	IPv4 FIB: 324,000 routes	✔	DCBX and PFC Data Center Bridging Exchange (DCBX), Priority-based Flow Control (PFC)
✔	IPv4 FIB learning rate: 6,600 routes/sec	✔	ECN Explicit Congestion Notification
✔	IPv6 FIB: 162,000 routes	✔	RoCE Lossless Ethernet (AI ECN + One-click PFC)
✔	IPv6 FIB learning rate: 16,500 routes/sec	✔	OpenFlow Forwarding traffic according to the flow table
Features		✔	MACsec S6850-2C with the LSWM124XGT2Q card
✔	Data Center Interconnection (DCI)	**Component Redundancy**	
✔	ECMP Dynamic Load Balancing (DLB)	✔	Fan Redundancy When one fan module is unplugged, there is no packet loss
✔	Link-aggregation DLB	✔	Power Redundancy When one power module is unplugged, there is no packet loss
✔	TCB	✔	Mixed power supply modules Using the AC power module and the DC power module together on the same switch

Note: The switch supports multiple modes. Please contact H3C for detailed capacity in different modes.

Source: Tolly, May 2023 Table 5

About Tolly

The Tolly Group companies have been delivering world-class ICT services for over 30 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, 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

225131 ivcfs15-yx-20250309-VerE