



# Huawei CloudEngine XH16800 Data Center Switch Datasheet

Huawei CloudEngine XH16800 is a data center (DC) core switch designed for AI scenarios.

## Product Overview

CloudEngine XH16800 series switches are Huawei's DC core switches built for AI scenarios. The switches build on an all-new hardware capability system and deliver the industry's lowest inter-card forwarding latency of about 3.5  $\mu$ s. They adopt the lossless Ethernet technology to enable DC storage networks to be built over Ethernet, thereby unifying technologies and O&M for computing and storage networks.


Running on Huawei's next-generation YunShan operating system, CloudEngine XH16800 series switches deliver stable, reliable, and high-performance Layer 2 and Layer 3 switching services, and provide customers with an intelligent, simple, secure, and open DC cloud network platform.

CloudEngine XH16800 series switches use an advanced hardware architecture design. They provide high-density 100GE/200GE/400GE ports and support smooth evolution to 800GE.

In addition to offering a diverse range of DC features, CloudEngine XH16800 series switches provide industrial-grade reliability by adopting an industry-leading backplane-free Clos orthogonal architecture and a strict front-to-back airflow design. As next-generation core switches, they significantly reduce power consumption by capitalizing on multiple innovative energy-saving technologies.

## Product Models and Appearances

CloudEngine XH16800 series switches fall into the following models.



CloudEngine XH16800-16    CloudEngine XH16800-8

# Product Features

## VXLAN Working with EVPN to Enable Service Deployment Within and Across DCs

- CloudEngine XH16800 series switches allow BGP-EVPN to be used as the VXLAN control plane to simplify VXLAN deployment.
- BGP-EVPN triggers automatic VXLAN tunnel establishment between virtual tunnel endpoints (VTEPs), eliminating the need to configure full-mesh tunnels. BGP-EVPN also reduces flooding of unknown traffic by advertising MAC address tables on the control plane. DCs can use BGP-EVPN to establish large Layer 2 networks.
- As a standard protocol, BGP-EVPN enables CloudEngine XH16800 series switches to interconnect with devices from other vendors, ensuring long-term network evolution.
- CloudEngine XH16800 series switches support distributed network deployment and support various VXLAN access modes, including accessing VXLAN through QinQ sub-interfaces. This enables flexible customization of heterogeneous networks.
- EVPN and VXLAN can be used to set up Layer 2 interconnections between DCs, enabling active-active VXLAN deployment across DCs and conserving DCI bandwidth.

## Automatic and Simplified Deployment of VXLAN Dual-Stack Network Services

- CloudEngine XH16800 series switches support hybrid overlay networking of IPv4 and IPv6 to meet the requirements of two IPv6 evolution modes.
- CloudEngine XH16800 series switches can carry IPv6 overlay services, that is IPv6 first for the network — an underlay IPv6 network is built and the IPv4 services of tenants are gradually reconstructed.
- CloudEngine XH16800 series switches support IPv6 VXLAN over IPv4 — IPv6 evolution of tenant services and network device reuse. That is, smooth evolution to IPv6 is performed based on the existing IPv4 network.

## Inter-Device Link Aggregation, Ensuring High Efficiency and Reliability

- CloudEngine XH16800 series switches support Multichassis Link Aggregation Group (M-LAG) to implement link aggregation among multiple devices, improving link reliability from the card level to the device level.
- Switches in an M-LAG all work in active state to share traffic and back up each other, enhancing system reliability.
- Switches in an M-LAG system can be upgraded independently. During the upgrade, other switches in the system take over traffic forwarding to ensure uninterrupted services.
- M-LAG supports dual-homing to Ethernet, VXLAN, and IP networks, allowing for flexible networking.

## Standard Interfaces, Enabling Open Interconnection

- CloudEngine XH16800 series switches support NETCONF and can interconnect with iMaster NCE-Fabric.
- CloudEngine XH16800 series switches support Ansible — an automatic management and O&M tool — to implement unified provisioning of physical and virtual networks.
- CloudEngine XH16800 series switches are integrated with mainstream cloud platforms (including commercial and open-source cloud platforms) and third-party controllers, enabling flexible service customization and automatic management.

## ZTP, Enabling Automatic O&M

- CloudEngine XH16800 series switches support Zero Touch Provisioning (ZTP). ZTP enables the switches to automatically obtain and load version files from a file server, freeing network engineers from onsite configuration and deployment. ZTP reduces labor costs and improves device deployment efficiency.

- ZTP supports embedded script languages and provides them for users through open APIs. DC users can use a familiar programming language (such as Python) to centrally configure network devices.
- ZTP decouples the configuration time of new devices from device quantity and geographical distribution, shortening the service provisioning time and improving the service provisioning efficiency.

## Simplified DCN Deployment via Collaboration with iMaster NCE-Fabric

- CloudEngine XH16800 series switches can interconnect with iMaster NCE-Fabric through standard protocols such as NETCONF and SNMP to adapt to networks and implement automatic network management. This helps to provide more efficient and intelligent operation methods, simplifying network management, and reducing the OPEX.

## Intelligent O&M Through Collaboration with iMaster NCE-FabricInsight

- CloudEngine XH16800 series switches support telemetry technology to collect device data in real time and send the collected data to iMaster NCE-FabricInsight — the DCN analysis component of Huawei iMaster NCE. Leveraging the intelligent fault identification algorithm, iMaster NCE-FabricInsight can analyze network data, accurately display the real-time network status, locate faults and identify their root causes in a timely and effective manner, and detect network problems that can affect user experience, precisely guaranteeing user experience.
- CloudEngine XH16800 series switches can insert IFIT extension headers into packets, visualize paths, and analyze interface-level packet loss, traffic, and latency to implement high-precision service-level packet loss measurement and facilitate fault demarcation.
- CloudEngine XH16800 series switches support Packet Event. When a device discards packets due to reasons such as abnormal forwarding, specified packet discarding rules, a full buffer, or ACL rule deny actions, or when the latency of packets exceeds a specified threshold, the device reports related flow entries to the iMaster NCE-FabricInsight collector.

## Intelligent Lossless Network, Meeting High Performance Requirements of RoCEv2 Applications

- CloudEngine XH16800 series switches support the iLossless algorithm to eliminate packet loss on the conventional Ethernet. This helps to build a lossless, low-latency, and high-throughput network environment for RoCEv2 traffic, meeting high performance requirements of RoCEv2 applications.
- CloudEngine XH16800 series switches support PFC deadlock prevention. They can identify service flows that may cause PFC deadlocks and change queue priorities of these flows to prevent PFC deadlocks.
- CloudEngine XH16800 series switches support Artificial Intelligence Explicit Congestion Notification (AI ECN). This future-oriented function can intelligently adjust the ECN thresholds of lossless queues based on the live-network traffic model to ensure low latency and high throughput with zero packet loss, maximizing the performance of lossless services.
- CloudEngine XH16800 series switches support Explicit Congestion Notification (ECN) Overlay. ECN Overlay applies ECN to a VXLAN network, enabling the traffic receiver to detect congestion on the overlay network in a timely manner and instruct the traffic sender to reduce its packet sending rate to relieve network congestion.

## Large Network Throughput, Fully Unleashing Computing Power in AI Scenarios

- CloudEngine XH16800 series switches support enhanced network scale load balancing (NSLB), achieving up to 98% network throughput on the entire network in AI scenarios and improving AI training efficiency.

## High-Performance, Non-Blocking Switching Architecture

- CloudEngine XH16800 series switches have a high-performance, non-blocking switching architecture that is characterized by its orthogonal switch fabric design, Clos orthogonal architecture, cell switching, and Virtual Output Queuing (VoQ).
- Backplane-free orthogonal switch fabric design: LPUs and SFUs of the CloudEngine XH16800 use an orthogonal design in which service traffic between LPUs is directly sent to the SFUs through orthogonal connectors. This design greatly improves system bandwidth and evolution capabilities, and can smoothly expand the system capacity to several thousands of Tbps.

- Clos orthogonal architecture: The three-stage Clos architecture of CloudEngine XH16800 series switches enables flexible expansion of switch fabric capacity. All types of cards (10GE, 25GE, 40GE, 100GE, and 400GE) of the switches support cell switching and dynamic traffic steering, and implement load balancing among multiple switch fabrics, preventing the switching matrix from being blocked and coping with changes in traffic models in DCs.
- The MPU engine is separated from the SFU hardware. The fault or change of the MPU does not affect the forwarding performance of the device.
- VoQ: CloudEngine XH16800 series switches support VoQ queues for fine-grained Quality of Service (QoS) on the switch fabrics. With the VoQ mechanism and inbound interface buffer, CloudEngine XH16800 series switches create independent VoQ queues on inbound interfaces to perform E2E flow control on traffic destined for different outbound interfaces. This ensures unified service scheduling and sequenced forwarding and implements non-blocking switching.

## Highly Reliable Industrial-Grade Hardware Architecture

- CloudEngine XH16800 series switches provide industrial-grade ultra-high reliability. They can run stably as core switches for a long period, ensuring service continuity.
- Hot standby of four hardware components: MPUs work in 1+1 hot standby mode. SFUs work in N+M hot standby mode. Power modules support dual inputs, work in N+1 or N+N hot standby mode, and have their own heat dissipation system. Fan modules work in N+1 hot standby mode, ensuring efficient heat dissipation.
- Redundancy of three buses: The monitoring, management, and data buses all work in 1+1 redundancy mode, ensuring reliable transmission of various signals in the system. Strict front-to-back airflow design
- The patented front-to-back airflow design strictly isolates cold air channels from hot air channels, meeting heat dissipation requirements in DC equipment rooms.
- Fan modules support intelligent area-based speed adjustment and on-demand heat dissipation, saving energy and reducing noise.

## Efficient and Intelligent Power Supply System

- The switch uses highly efficient digital power modules with a power efficiency of up to 96%.
- The system measures power consumption in real time and intelligently puts power modules into sleep mode to save power when system power demands are low.
- The switch adjusts power consumption of major components based on changes in service traffic volume, saving energy dynamically.

# Licensing

Huawei CloudEngine XH16800 series switches support the CloudFabric IDN One Software (N1) business model, which bundles iMaster NCE-Fabric, iMaster NCE-FabricInsight, and CloudEngine switches in a range of typical scenarios. This approach simplifies transactions, provides customers with more functions and value, and protects customers' software investment through Software License Portability.

Feature	N1 Package for AI Scenarios		Add-On Package	
	Foundation	Advanced	Digital Map	Security
Basic functions (including IPv6 and VXLAN)	•	•		
AI ECN	•	•		
FusionInsight basic analysis	•	•		
Telemetry	•	•		
AI ZTP	•	•		
NSLB		•		
IFIT function		•		
IFIT service assurance function		•		
MACsec				•
Digital map function			•	

# Product Specifications

Item	CloudEngine XH16800-8	CloudEngine XH16800-16
Switching capacity (Tbps)	967/2903	1935/5806
Packet forwarding rate (Mpps)	460,800	921,600
Number of service slots	8	16
SFU slot	9	
Air duct type	Standard front-to-rear airflow	
Device virtualization	M-LAG	
Interface	Jumbo frames	
Network virtualization	VXLAN routing and VXLAN bridging	
	BGP-EVPN	
SDN	iMaster NCE-Fabric	
Network convergence	PFC and AI ECN	
	RDMA and RoCE (RoCE v1 and RoCE v2)	
Programmability	OpenFlow	
	OPS programming	
Traffic analysis	NetStream	
VLAN	Access, trunk, and hybrid ports	
	Default VLAN	
MAC address table	Automatic MAC address learning and aging	
	Static, dynamic, and blackhole MAC address entries	
	Source MAC address filtering	
	MAC address learning limiting based on ports and VLANs	
IP routing	IPv4 dynamic routing protocols such as RIP, OSPF, IS-IS, and BGP	
	IPv6 dynamic routing protocols such as RIPng, OSPFv3, IS-ISv6, and BGP4+	
IPv6	VXLAN over IPv6	
	IPv6 VXLAN over IPv4	
	IPv6 neighbor discovery (ND)	
	Path MTU discovery (PMTU)	
	TCP6, IPv6 ping, IPv6 tracer, IPv6 socket, UDP6, and raw IPv6	
Multicast	Multicast routing protocols such as IGMP, PIM-SM, and MSDP	

Item	CloudEngine XH16800-8	CloudEngine XH16800-16
	IGMP snooping and IGMP proxy	
	IPv6 Layer 3 multicast and configuration of both Layer 2 and Layer 3 multicast services	
	Fast leave of multicast member interfaces	
	Multicast traffic suppression	
<b>Reliability</b>	LACP	
	STP, RSTP, VBST, and MSTP	
	BPDU protection	
	Smart link and multi-instance	
	Hardware-based Bidirectional Forwarding Detection (BFD), with a minimum packet sending interval of 3.3 ms	
	VRRP, VRRP load sharing, and BFD for VRRP	
	BFD for BGP, IS-IS, OSPF, and static routing	
	BFD for VXLAN	
<b>QoS</b>	ACL, CAR, re-marking, and scheduling	
	Queue scheduling modes such as PQ, DRR, and PQ+DRR	
	Congestion avoidance mechanisms such as WRED and tail drop	
	Traffic shaping	
	1588v2	
<b>Intelligent O&amp;M</b>	Network-wide path detection	
	Telemetry	
	Enhanced ERSPAN	
	In-situ Flow Information Telemetry (IFIT)	
	Packet Event: packet loss visualization and ultra-long latency visualization	
	VXLAN OAM: VXLAN ping and VXLAN tracet	
<b>Intelligent lossless network</b>	PFC deadlock prevention	
	AI ECN	
	ECN Overlay	
	Enhanced NSLB	
<b>Configuration and maintenance</b>	Terminal login through the console port, Telnet, and SSH	

Item	CloudEngine XH16800-8	CloudEngine XH16800-16
	Network management protocols, such as SNMPv1/v2/v3	
	File upload and download through FTP and TFTP	
	Boot Read-Only Memory (BootROM) upgrade and remote online upgrade	
	Hot patches	
	User operation logs	
	Configuration rollback	
	ZTP	
<b>Security and management</b>	MACsec	
	Command line authority control based on user levels, preventing unauthorized users from using commands	
	Defense against DoS, ARP, and ICMP attacks	
	Port isolation, port security, and sticky MAC	
	Binding of the IP address, MAC address, port ID, and VLAN ID	
	Authentication methods, including AAA, LDAP, RADIUS, and HWTACACS	
	RMON	
<b>Dimensions (H x W x D)</b>	703 mm (15.8 U) x 483 mm x 985 mm	1436 mm (32.3 U) x 483 mm x 1144 mm
<b>Weight in full configuration</b>	129.8 kg	262.8 kg
<b>Environment requirements</b>	Operating temperature: 0°C to 40°C (0 m to 1800 m) Storage temperature: -40°C to +70°C Relative humidity: 5% RH to 95% RH (noncondensing)	
<b>Operating voltage</b>	AC: 176 V to 290 V, 45 Hz to 65 Hz High-voltage DC (HVDC): 188 V to 288 V/260 V to 400 V	
<b>Maximum power supply capability of the system</b>	30,000 W	60,000 W



## Performance and Scalability

Item	Value
Maximum number of MAC address entries	640K
Maximum number of routes (FIB IPv4/IPv6)	1.5M/750K
ARP size	128K
Maximum number of VRFs	4096
IPv6 ND table size	128K
Maximum number of VRRP groups	1024
Maximum number of ECMP paths	128
Maximum number of VXLAN bridge domains	16K
Maximum number of BDIF interfaces	16K
Maximum number of virtual tunnel endpoints (VTEPs)	16K
Maximum number of LAGs	1024
Maximum number of links in a LAG	256
Maximum number of MSTIs	1000
Maximum number of VLANs where VBST can be configured	640K

Note: This specification may vary between different scenarios. Please contact Huawei for details.

# Safety and Regulatory Compliance

The following table lists the safety and regulatory compliance of CloudEngine XH16800 series switches.

Certification Category	Description
<b>Safety</b>	EN 62368-1 IEC 62368-1 UL 62368-1 CSA-C22.2 No.62368-1 AS/NZS 62368-1 GB4943
<b>Electromagnetic Compatibility (EMC)</b>	EN 300386 EN 55032 EN 55035 IEC/EN 61000-3-2 IEC/EN 61000-3-3 AS/NZS CISPR32 FCC 47CFR Part15 ICES-003 CISPR 32 CISPR 24 VCCI- CISPR32 CISPR35 GB9254
<b>Environment</b>	EN 50581 EN 50419 (EC) No.1907/2006 GB/T 26572 ETSI EN 300 019-1-1 ETSI EN 300 019-1-2 ETSI EN 300 019-1-3 ETSI EN 300 753
<p>EMC: electromagnetic compatibility</p> <p>CISPR: International Special Committee on Radio Interference</p> <p>EN: European Standard</p> <p>ETSI: European Telecommunications Standards Institute</p> <p>CFR: Code of Federal Regulations</p> <p>FCC: Federal Communication Commission</p> <p>IEC: International Electrotechnical Commission</p> <p>AS/NZS: Australian/New Zealand Standard</p> <p>VCCI: Voluntary Control Council for Interference</p> <p>UL: Underwriters Laboratories</p> <p>CSA: Canadian Standards Association</p>	

# Ordering Information

Device Mode	Description
<b>Chassis</b>	
XH16800-8-AH	XH16800-8 AC/HVDC assembly chassis
XH16800-16-AH	XH16800-16 AC/HVDC assembly chassis
<b>Main processing unit</b>	
XH-MPUE-H	XH16800 Main Processing Unit E (half-width)
XH-MPUE-F	XH16800 Main Processing Unit E (Full-width)
<b>Switch fabric unit (G-H series)</b>	
XH-SFU08G-H	XH16808 Switch Fabric Unit G-H
XH-SFU16G-H	XH16816 Switch Fabric Unit G-H
<b>Interface card (-J2 series)</b>	
XH-L36DQ2CQJ2	36-port 400GE and 2-port 100GE interface card (XHL-J2, QSFP-DD, QSFP28)
XH-L40DQJ2	40-port 400GE interface card (XHL-J2, QSFP-DD)
<b>Fan module (Hongyan 5.0)</b>	
FAN-240SD-B	Fan Tray Assembly,FAN-240SD-B
FAN-480SE-B	Fan Tray Assembly,FAN-480SE -B
FAN-960SE-B	Fan Tray Assembly,FAN-960SE -B
<b>Power module</b>	
PAH-3000WA	3000W Dual Inputs AC&HVDC Power Module

Note: The XH-L40DQJ2 interface card will be launched in the second half of 2024.

Software Model	Software Description
<b>N1-XH-M-LIC-DM</b>	N1-XH-M-LIC-DM,N1-XH Digital Map Function-Modular
<b>N1-XHMDM-SnS1Y</b>	N1-XH Digital Map Function, Per Modular device-SnS-Year
<b>N1-XH-M-LIC-MACSEC</b>	N1-XH Data Center Switch Security Function-Modular
<b>N1-XHMMACSEC-SnS1Y</b>	N1-XH Data Center Switch Security Function, Per Modular device-SnS-1 Year
<b>N1-XH168-256T-LIC-STD</b>	N1-XH16800 256T Switch Standard Package for the AI Scenarios
<b>N1-XH168-256T-STD-SnS1Y</b>	N1-XH16800 256T Switch Standard Package for the AI Scenarios-SnS-1 Year
<b>N1-XH168-128T-LIC-ENH</b>	N1-XH16800 128T Switch Enhanced Package for the AI Scenarios
<b>N1-XH168-128T-ENH-SnS1Y</b>	N1-XH16800 128T Switch Enhanced Package for the AI Scenarios-SnS-1 Year
<b>N1-XH168-256T-LIC-ENH</b>	N1-XH16800 256T Switch Enhanced Package for the AI Scenarios
<b>N1-XH168-256T-ENH-SnS1Y</b>	N1-XH16800 256T Switch Enhanced Package for the AI Scenarios-SnS-1 Year

## Networking and Applications

On typical data center networks, CloudEngine XH16800 series switches work as core switches, whereas CloudEngine XH9000/XH8000 series switches connect to the core switches through 100GE or 200GE ports to build an end-to-end and fully-connected 100GE/200GE/400GE network. The switches use VXLAN and other fabric protocols to establish a non-blocking large Layer 2 network, which allows large-scale VM migration and flexible service deployment.



Note: VXLAN can also be used on campus networks to support flexible service deployment in different service areas.

## More Information

For more information about Huawei products, visit <https://e.huawei.com/en/> or contact Huawei's local sales office.

Alternatively, you can contact us through one of the following methods:


- Global service hotline: <https://e.huawei.com/en/about/service-hotline>
- Enterprise technical support website: <https://support.huawei.com/enterprise/en/index.html>
- Service email address for enterprise users: support\_e@huawei.com

---

**Copyright © Huawei Technologies Co., Ltd. 2024. All rights reserved.**

No part of this document may be reproduced or transferred in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

### Trademarks and Permissions

 **HUAWEI** and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

### Notice

The purchased products, services, and features are stipulated by the commercial contract made between Huawei and the customer. All or partial products, services, and features described in this document may not be within the purchased scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

### Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base, Bantian,  
Longgang, Shenzhen, People's Republic of  
China

Post code: 518129

Website: <https://e.huawei.com/en/>