

# Huawei CloudEngine 16800 Switch Datasheet

Huawei CloudEngine 16800 uses innovative iLossless algorithm to learn and train network-wide traffic, implement zero packet loss and E2E  $\mu$  s-level latency, and achieve maximum throughput.



# Product Overview

Huawei CloudEngine 16800 is the data center switch built for the AI era. It uses innovative iLossless algorithm to learn and train network-wide traffic, implement zero packet loss and E2E  $\mu$ s-level latency, achieve the maximum throughput and enables more efficient NVMe and fully unleashes the value of all-flash storage. In addition, they support NOF+ technology to implement plug-and-play of the storage system and fast fault detection.

CloudEngine 16800 series switches provide stable, reliable, and secure high-performance Layer 2/Layer 3 switching capabilities to help build an elastic, virtualized, and high-quality network.

CloudEngine 16800 series switches use an advanced hardware architecture design. The CloudEngine 16800 series provide up to 576 x 400GE, 768 x 100GE, 768 x 40GE, 792 x 25GE or 792 x 10GE line-rate ports and use an industry-leading backplane-free Clos architecture and provide industrial-grade reliability. The switches support comprehensive data center service features. Their front-to-back airflow design suits data center equipment rooms, and the innovative energy conservation technologies greatly reduce power consumption.

## Product Appearance

The CloudEngine 16800 series switches are available in three models: CloudEngine 16804, CloudEngine 16808, and CloudEngine 16816.



## Product Characteristics

The CloudEngine 16800 Provides Simplified Deployment, and Various SDN Features Implement Automatic Service Deployment

### VXLAN and EVPN Enable Flexible Expansion Within and Across Data Centers

- The CloudEngine 16800 supports Border Gateway Protocol - Ethernet VPN (BGP-EVPN), which can run as the VXLAN control plane to simplify VXLAN deployment.
- BGP-EVPN triggers automatic VXLAN tunnel setup between virtual tunnel endpoints (VTEPs), removing the need for full-mesh tunnel configuration. BGP-EVPN also reduces flooding of unknown traffic by advertising MAC routes on the control plane. With BGP-EVPN, large Layer 2 networks can be established for data centers.
- Because BGP-EVPN is a standard protocol, the CloudEngine 16800 is interoperable with devices from other vendors, enabling long-term network evolution.

- The CloudEngine 16800 supports centralized and distributed VXLAN deployment and supports various VXLAN access modes, including QinQ access VXLAN. This allows for flexible customization of heterogeneous networks.
- EVPN and VXLAN can be used to set up Layer 2 interconnections between data centers, enabling active VXLAN deployment across data centers and conserving DCI link bandwidth.
- The CloudEngine 16800 supports IP packet fragmentation and reassembly, enabling oversized IP packets to be transmitted WANs without being limited by the MTU. The CloudEngine 16800 switches can also identify fragmented packets to seamlessly interconnect with routers.

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

- The CloudEngine 16800 supports IPv4 and IPv6 networking, meeting requirements of two IPv6 evolution modes.
- For new data centers, the CloudEngine 16800 supports IPv6 service deployment on the underlay network. IPv4 services of tenants are gradually reconstructed.
- The CloudEngine 16800 supports IPv6 VXLAN over IPv4. In this case, network devices are reused, and services are smoothly evolved to IPv6 based on the existing IPv4 network.

## ZTP, Implementing Automatic O&M

- CloudEngine 16800 series support Zero Touch Provisioning (ZTP). ZTP enables the CloudEngine 16800 series to automatically obtain and load version files from a USB flash drive or file server, freeing network engineers from onsite configuration and deployment. ZTP reduces labor costs and improves device deployment efficiency.
- ZTP provides built-in scripts through open APIs. Data center personnel can use a programming language they are familiar with, such as Python, to centrally configure network devices.
- ZTP decouples the configuration time of new devices from the device quantity and area distribution, which improves service provisioning efficiency.

## iMaster NCE-based Simplified Network Deployment

- CloudEngine 16800 series switches can interconnect with iMaster NCE-Fabric through standard protocols such as NetConf and SNMP to implement network automatic management and control, providing more efficient and intelligent operation methods, simplifying network management, and reducing the OPEX.

## FabricInsight-based Intelligent O&M

- CloudEngine 16800 series switches provide telemetry technology to collect device data in real time and send the data to Huawei data center network analyzer iMaster NCE-FabricInsight. The iMaster NCE-FabricInsight analyzes network data based on the intelligent fault identification algorithm, accurately displays the real-time network status, effectively demarcates and locates faults in a timely manner, and identifies network problems that affect user experience, accurately guaranteeing user experience.
- CloudEngine 16800 series switches provide intelligent traffic analysis function to monitor and analyze network traffic. It allows O&M personnel to perform in-depth analysis on a specified service flow, obtain data about high-precision performance indicators such as the packet loss rate and latency (nanosecond-level) of the service flow, and export the analysis results to the analyzer for graphical display. This makes it easier for O&M personnel to monitor the network condition and quickly locate network faults.
- CloudEngine 16800 series switches provide AnyFlow function. It creates flow tables based on fields such as 5-tuple for TCP, UDP, and VXLAN traffic entering the device, in order to collect traffic statistics and detect abnormal traffic. It can also report those flow entries to the analyzer for network quality analysis.
- CloudEngine 16800 series switches provide Packet Event function. It can create flow entries carrying information about packet loss or high-latency packets and report them to the iMaster NCE-FabricInsight collector.

## Intelligent Lossless Network

- CloudEngine 16800 series switches provide iLossless algorithm solve the packet loss problem caused by congestion on traditional Ethernet networks and provide a network environment with zero packet loss, low delay, and high throughput for RoCEv2 traffic to meet high-performance requirements of RoCEv2 applications.
- The CloudEngine 16800 series support iNOF which enables the switch to rapidly manage the newly connected hosts to obtain information about such hosts in real time, intelligently adjust the intelligent lossless network configurations, and notify the storage system of the host information. This helps the storage system manage hosts and implements plug-and-play of the storage system. In this manner, links are automatically established in seconds based on the network plan once computing and

storage nodes are brought online. In addition, faults can be quickly detected, and switchovers between the active and standby servers can be completed within seconds.

- The CloudEngine 16800 series support Network-based Proactive Congestion Control (NPCC) which intelligently identifies the congestion status, proactively sends CNPs, and accurately controls the speed at which servers send RoCEv2 packets. This ensures timely speed reduction upon congestion and prevents excessive speed reduction when congestion is relieved, guaranteeing low latency and high throughput of RoCEv2 services in long-distance DCI scenarios.
- The CloudEngine 16800 series support PFC deadlock prevention which identifies service flows that easily cause PFC deadlocks and changes queue priorities to prevent PFC deadlocks.
- CloudEngine 16800 series switches support Artificial Intelligence Explicit Congestion Notification (AI ECN) function to intelligently adjust ECN thresholds of lossless queues based on the traffic model on the live network. This function ensures low delay and high throughput with zero packet loss, achieving optimal performance for lossless services.
- CloudEngine 16800 series switches support Explicit Congestion Notification (ECN) Overlay function, apply 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.

## The CloudEngine 16800 Is Secure and Reliable, Improving Network Service Stability

### Inter-device Link Aggregation, High Efficiency and Reliability

- CloudEngine 16800 series support multi-chassis link aggregation group (M-LAG), which enables links of multiple switches to aggregate into one to implement device-level link backup.
- 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 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.

## Openness and Programmability Enable Agile Deployment and O&M

### OPS Implements Programmability at the Control Plane

- The CloudEngine 16800 uses the Open Programmability System (OPS) to provide programmability at the control plane.
- The OPS provides open APIs. APIs can be integrated with mainstream cloud platforms (including commercial and open cloud platforms). The OPS enables services to be flexibly customized and provides automatic management.
- Users or third-party developers can use open APIs to develop and deploy specialized network management policies to implement extension of fast service functions, automatic deployment, and intelligent management. The OPS also implements automatic operation and maintenance, and reduces management costs.
- The OPS provides seamless integration of data center service and network in addition to a service-oriented, software-defined networking.

## Standard Interfaces Provide Openness and Interoperability

- The CloudEngine 16800 supports Netconf. It can work with Huawei iMaster NCE-Fabric.
- The CloudEngine 16800 provides the standard Netconf interface for third-party software to invoke. This enables programming of functions and integration with third-party software, providing openness and flexibility.
- You can use CE modules for Ansible released on open-source websites and Ansible tools to automate network deployment, simplifying device management and maintenance. Through in-depth collaboration with mainstream cloud platforms, and O&M tools, the CloudEngine 16800 series switches can be integrated into SDN and cloud computing platforms flexibly and quickly.

## Next-Generation Core Engine with the High Performance

### 1548Tbps Switching Capacity

- The CloudEngine 16800 supports 10GE, 25GE, 40GE, 100GE and 400GE interfaces. This high capacity can support sustainable development of cloud-computing data centers for the next 10 years.
- The CloudEngine 16800, together with the CloudEngine 8800, CloudEngine 6800, or CloudEngine 5800 series Top of Rack (TOR) switches, can implement the large non-blocking switching network, providing access for tens of thousands of GE/10GE/25GE/50GE/100GE servers.

## Tbps-Level High-Density Line Cards

- Line card provides up to 14.4Tbps forwarding capacity.
- The CloudEngine 16800 supports up to 36 x 400GE, 48 x 100GE, 72 x 25GE + 6 x 100GE line cards, which provide line-rate forwarding

## Advanced Architecture Ensures Industry-Leading Network Quality

### High-Performance, Non-blocking Switching Architecture

- The CloudEngine 16800 uses a non-blocking switching architecture that is characterized by its orthogonal switch fabric design, Clos architecture, cell switching, and Virtual Output Queuing (VoQ).
- Backplane-free orthogonal switch fabric design: CloudEngine 16800 service line cards and SFUs use an orthogonal design in which service traffic between line cards is directly sent to the SFUs through orthogonal connectors. This design greatly improves system bandwidth and evolution capabilities, enabling the system switching capacity to scale to more than 100 Tbps.
- Clos architecture: The CloudEngine 16800's three-level Clos architecture permits flexible expansion of switch fabric capacity. The architecture uses Variable Size Cell (VSC) and provides dynamic routing. Load balancing among multiple switch fabrics prevents the switching matrix from being blocked and allows it to easily cope with complex, volatile traffic in data centers.
- VoQ: The CloudEngine 16800 supports VoQ queues that implement fine-grained Quality of Service (QoS) based on the switch fabrics. With the VoQ mechanism and large buffer on inbound interfaces, the CloudEngine 16800 creates independent VoQ queues on inbound interfaces to perform end-to-end flow control on traffic destined for different outbound interfaces. This method ensures unified service scheduling and sequenced forwarding and implements non-blocking switching.

### Highly Reliable Industry-grade Hardware Architecture

- The CloudEngine 16800 has industrial-grade reliability and stability to ensure long-term service continuity.
- Hot backup of four key components: MPUs work in 1+1 hot backup mode. SFUs work in N+M hot backup mode. Power supplies support dual inputs and N+1 or N+N backup and have their own fans. Both fan trays work in N+1 backup mode. Each fan tray has multiple fans working in 1+1 backup mode, ensuring efficient heat dissipation.
- Redundancy of three types of major buses: Monitoring, management, and data buses all work in 1+1 backup mode. Bus redundancy ensures reliable signal transmission.

## Pioneering Energy-saving Technology

### Strict Front-to-Back Airflow Design

- The CloudEngine 16800 uses a patented front-to-back airflow design that isolates cold air channels from hot air channels. This design meets heat dissipation requirements in data center equipment rooms.
- The fan speed in each area can be dynamically adjusted based on the workload of line cards in the area. This on-demand cooling design lowers power consumption and reduces noise.

### Low Power Consumption

- The CloudEngine 16800 uses innovative energy saving technologies. The port power consumption is half the industry average. This greatly reduces power consumption in the data center equipment room.

### Efficient, Intelligent Power Supply System

- The CloudEngine 16800 incorporates efficient digital power modules, which provide power efficiency as high as 96%.
- The power supply system measures power consumption in real time and puts one or more power modules into sleep mode when system power demands are low.
- The CloudEngine 16800 can save energy dynamically by adjusting the power consumption of components to adapt to changes in service traffic volume.

## Licensing

CloudEngine 16800 supports Huawei IDN One Software (N1 mode for short) licensing mode. The CloudFabric N1 business model combines the NCE controller, analyzer, and CloudEngine switch software for use in a range of common scenarios. This simplifies transactions, provides customers with more functions and value, and protects their software investment with Software License Portability.

Product	Feature	N1 Mandatory Software Packages			N1 Add-on Software Packages			
CloudEngine 16800 series switch		Foundation	Advanced	Premium	AI Fabric RDMA Application Acceleration Function Package 2 (Applicable to HG-P/L-J series line cards)	AI Fabric Value-added Package for the NOF+ Storage Scenarios	Security Function	Multi-cloud Multi-DC Value-added Package
	Lossless upgrade	√	√	√				
	Telemetry	√	√	√				
	PTP	√	√	√				
	LLETH				√	√		
	AI ECN				√	√		
	iNOF					√		
	MACsec						√	
iMaster NCE-Fabric controller	Automatic deployment function	√	√	√				
	Basic intent functions (simulation and verification, and network-wide configuration rollback)			√				
	Multi-cloud multi-DC function							√
iMaster NCE-FabricInsight analyzer	Telemetry-based network analysis	√	√	√				
	Network health evaluation		√	√				
	Value-added functions for network traffic analysis (managing 100 VMs)			√				
Version mapping		Select one from three options. The Advanced software package contains functions of the Foundation software package, and the Premium software package contains functions of the			It is used with the Foundation, Advanced, or Premium software package.			It is used with the Foundation, Advanced

Product	Feature	N1 Mandatory Software Packages	N1 Add-on Software Packages
		Advanced software package.	ed, or Premium software package

For details about product function differences, refer to the product documentation.

Note: For detailed information of Huawei CloudFabric N1 business model, visit <https://e.huawei.com/en/material/networking/dcsswitch/03a0e69bfa2c4f168323ba94a75f1f09>.

## Product Specifications

Note: This content is applicable only to regions outside Chinese mainland. Huawei reserves the right to interpret this content.

### Functions and Features

Item	CloudEngine 16800(L-J series line cards)
Fabric architecture	Clos architecture, cell switching, VoQ
Airflow design	Strict front-to-back
Network virtualization	<p>VXLAN routing and bridging</p> <p>EVPN</p> <p>QinQ access VXLAN</p>
SDN	iMaster NCE-Fabric
DCI	<p>BGP-EVPN</p> <p>VXLAN mapping, implementing interconnection between multiple DCI networks at Layer 2</p>
Network convergence	<p>PFC and ECN</p> <p>RDMA and RoCE (RoCEv1 and RoCEv2)</p>
Programmability	<p>OPS programming</p> <p>Ansible-based automatic configuration and open-source module release</p>
Traffic analysis	NetStream,sFlow(R23.1)
VLAN	<p>Adding access, trunk, and hybrid interfaces to VLANs</p> <p>Default VLAN</p> <p>QinQ</p>
MAC address	<p>Dynamic learning and aging of MAC address entries</p> <p>Static, dynamic, and blackhole MAC address entries</p> <p>Packet filtering based on source MAC addresses</p> <p>MAC address limiting based on ports and VLANs</p>
IP routing	IPv4 routing protocols, such as RIP, OSPF, IS-IS, and BGP

Item	CloudEngine 16800(L-J series line cards)
	IPv6 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 tracert, IPv6 socket, UDP6, and raw IPv6
Multicast	Multicast routing protocols such as IGMP, PIM-SM, PIM-DM and MSDP
	IGMP snooping
	IGMP proxy
	Fast leaving of multicast member interfaces
	Multicast traffic suppression
Reliability	Link Aggregation Control Protocol (LACP)
	M-LAG
	STP, RSTP, VBST, and MSTP
	BPDU protection
	Smart Link and multi-instance
	Device Link Detection Protocol (DLDP)
	Hardware-based Bidirectional Forwarding Detection (BFD)
	VRRP, VRRP load balancing, and BFD for VRRP
	BFD for BGP/IS-IS/OSPF/static route
	BFD for VXLAN
	BFD for M-LAG
QoS	Traffic classification based on Layer 2, Layer 3, Layer 4, and priority information
	Actions including ACL, CAR, and re-marking
	Queue scheduling modes such as PQ, DRR, and PQ+DRR
	Congestion avoidance mechanisms, including WRED and tail drop
	Traffic shaping
Intelligent O&M	Telemetry
	1588v2
	ERSPAN+
	Packet Event
	Network-wide path detection
	Statistics on the buffer microburst status
	VXLAN OAM: VXLAN ping and VXLAN tracert

Item	CloudEngine 16800(L-J series line cards)
Intelligent and lossless network	RDMA and RoCE (RoCEv1 and RoCEv2)
	PFC
	AI ECN
	iNOF
	ECN Overlay
Configuration and maintenance	Console, Telnet, and SSH terminals
	Network management protocols, such as SNMPv1/v2/v3
	File upload and download through FTP and TFTP
	BootROM upgrade and remote upgrade
	Hot patches
	User operation logs
Security and management	Zero Touch Provisioning (ZTP)
	MACsec
	RADIUS and HWTACACS authentication for login users
	Command line authority control based on user levels, preventing unauthorized users from using commands
	Defense against MAC address attacks, broadcast storms, and heavy-traffic attacks
	ICMP-based ping and traceroute
	Port mirroring and flow mirroring
	Remote Network Monitoring (RMON)

## Performance and Scalability

Item	CloudEngine 16800(L-J series line cards)
Maximum number of MAC address entries	480K
Maximum number of routes (FIB IPv4/IPv6)	1M/512K
ARP size	128K
Maximum number of VRFs	8000
IPv6 ND table size	128K
Maximum number of VRRP groups	4096
Maximum number of ECMP paths	128
Maximum number of VXLAN bridge domains	32K
Maximum number of BDIF interfaces	32K
Maximum number of virtual tunnel endpoints (VTEPs)	16K
Maximum number of LAGs	1024
Maximum number of links in a LAG	128

Item	CloudEngine 16800(L-J series line cards)
Maximum number of MSTIs	64
Maximum number of VLANs where VBST can be configured	2000

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

## Hardware Specifications

Item	CloudEngine 16804	CloudEngine 16808	CloudEngine 16816
Switching capacity (Tbit/s)-L-J Series Cards	45(Single Device)	89(Single Device)	178(Single Device)
Forwarding rate (Mpps)-L-J Series Cards	14,400	28,800	57,600
Service slots	4	8	16
Switching fabric module slots	4 (scalable to 9 for future expansion)		
Dimensions (W x D x H)	483 mm x 985 mm x 438 mm	483 mm x 985 mm x 703 mm	483 mm x 1144 mm x 1436
Chassis weight (empty)	94.1 kg (207.49 lb)	129.8 kg (286.21 lb)	262.8 kg (579.47 lb)
Operating voltage	AC: 176 V to 290 V DC: -40 V to -72 V HVDC: 188 V to 288 V/260 V to 400 V		
Hot swappable (power modules)	Yes		
Maximum power supply (W)	18,000	30,000	60,000
Operating temperature	0°C to 40°C		
Storage temperature	-40°C to +70°C		
Humidity	Long-term operating humidity: 5% to 85%, non-condensing Short-term operating humidity: 0% to 95%, non-condensing		
Altitude	Operating altitude: <1800 m Storage altitude: <5000 m		
MTBF	34.93 years	34.55 years	32.68 years
MTTR	1 hour		
MTTF	34.93 years	34.55 years	32.68 years
Availability	0.9999971695	0.9999971582	0.9999970745

1: For details, see the product specifications in the specification list.

## Safety and Regulatory Compliance

The following table lists the safety and regulatory compliance of CloudEngine 16800 series switches.

Certification Category	Description
Safety	<ul style="list-style-type: none"> <li>EN 60950-1</li> <li>EN 60825-1</li> </ul>

Certification Category	Description
	<ul style="list-style-type: none"> <li>• EN 60825-2</li> <li>• UL 60950-1</li> <li>• CSA-C22.2 No. 60950-1</li> <li>• IEC 60950-1</li> <li>• AS/NZS 60950-1</li> <li>• GB4943</li> </ul>
Electromagnetic Compatibility (EMC)	<ul style="list-style-type: none"> <li>• EN 300386</li> <li>• EN 55032: CLASS A</li> <li>• EN 55024</li> <li>• IEC/EN 61000-3-2</li> <li>• IEC/EN 61000-3-3</li> <li>• FCC 47CFR Part15 CLASS A</li> <li>• ICES-003: CLASS A</li> <li>• CISPR 32: CLASS A</li> <li>• CISPR 24</li> <li>• AS/NZS CISPR32</li> <li>• VCCI- CISPR32: CLASS A</li> <li>• GB9254 CLASS A</li> </ul>
Environment	<ul style="list-style-type: none"> <li>• 2011/65/EU EN 50581</li> <li>• 2012/19/EU EN 50419</li> <li>• (EC) No.1907/2006</li> <li>• GB/T 26572</li> <li>• ETSI EN 300 019-1-1</li> <li>• ETSI EN 300 019-1-2</li> <li>• ETSI EN 300 019-1-3</li> <li>• ETSI EN 300 753 GR63</li> </ul>

Note

EMC: electromagnetic compatibility

CISPR: International Special Committee on Radio Interference

EN: European Standard

ETSI: European Telecommunications Standards Institute

CFR: Code of Federal Regulations

FCC: Federal Communication Commission

IEC: International Electrotechnical Commission

AS/NZS: Australian/New Zealand Standard

VCCI: Voluntary Control Council for Interference

UL: Underwriters Laboratories

CSA: Canadian Standards Association

## Optical Transceivers and Cable

For details about optical transceivers and cables, visit

<https://e.huawei.com/en/material/networking/dcsswitch/f6d91cf16df0474998087676a33fd41e>.

## Ordering Information

Basic Configuration	
GX-RACK-01	A812-20 Assembly Rack(800 mm x 1200 mm x 2000 mm))
CE16804-AH	CE16804 AC/HVDC assembly chassis
CE16804-DC	CE16804 DC assembly chassis
CE16808-AH	CE16808 AC/HVDC assembly chassis
CE16808-DC	CE16808 DC assembly chassis
CE16816-AH	CE16816 AC/HVDC assembly chassis
CE16816-DC	CE16816 DC assembly chassis

Main Processing Unit	
CE-MPUD-HALF	CE16800 Main Processing Unit D (half-width)
CE-MPUE-HALF	CE16800 Main Processing Unit E (half-width)
CE-MPUD-FULL	CE16800 Main Processing Unit D (full-width)
CE-MPUE-FULL	CE16800 Main Processing Unit E (full-width)
CE-MPUD-HALF2	CE16800 Main Processing Unit D-2 (half-width)
CE-MPUD-FULL2	CE16800 Main Processing Unit D-2 (full-width)
CE-MPUD-HALF3E	CE16800 Main Processing Unit D-3E (half-width)
CE-MPUD-FULL3E	CE16800 Main Processing Unit D-3E (full-width)

Switch Fabric Unit	
CE-SFU04J-P	CE16804 Switch Fabric Unit J-P
CE-SFU04H-P	CE16804 Switch Fabric Unit H-P
CE-SFU08J-P	CE16808 Switch Fabric Unit J-P
CE-SFU08H-P	CE16808 Switch Fabric Unit H-P
CE-SFU16J-P	CE16816 Switch Fabric Unit J-P
CE-SFU04F-G	CE16804 Switch Fabric Unit F
CE-SFU04G-G	CE16804 Switch Fabric Unit G
CE-SFU08F-G	CE16808 Switch Fabric Unit F
CE-SFU08G-G	CE16808 Switch Fabric Unit G

Switch Fabric Unit	
CE-SFU16F-G	CE16816 Switch Fabric Unit F
CE-SFU16G-G	CE16816 Switch Fabric Unit G
CE-SFU04F-A	CE16804 Switch Fabric Unit F-A
CE-SFU04G-A	CE16804 Switch Fabric Unit G-A
CE-SFU08F-A	CE16808 Switch Fabric Unit F-A
CE-SFU08F-H1E	CE16808 Switch Fabric Unit F-H1E
CE-SFU04G-L	CE16804 Switch Fabric Unit G-L
CE-SFU08G-L	CE16808 Switch Fabric Unit G-L
CE-SFU16G-L	CE16816 Switch Fabric Unit G-L

25GE Interface Card	
CEL72XSHGA-P	72-port 25GE and 6-port 100GE interface card (HG-P, SFP28, QSFP28)(CM)
CEL72XSHGA-P0	72-port 25GE and 6-port 100GE interface card with RTU bundle 0 (HGA-P, SFP28, QSFP28, 24*10GE+2*100GE port RTU, 3*(24*25GE capacity upgrade RTU))
CEL72XS-SAN0	72-port 25GE and 6-port 100GE SAN interface card with RTU bundle 0 (HGA-P, SFP28, QSFP28, 24*10GE+2*100GE port RTU, 3*(24*25GE capacity upgrade RTU))
CEL72XS-SAN1	72-port 25GE and 6-port 100GE interface card with optical module bundle 1 (HGA-P, SFP28, QSFP28, 24*10GE+2*100GE port RTU, 3*(24*25GE capacity upgrade RTU), 72*25GBase-SR-MP)
CE-L48YSAJ	48-port 25GE and 4-port 100GE interface card (L-J, SFP28, QSFP28)
CE-L48YSJ	48-port 25GE interface card (L-J, SFP28)

40GE Interface Card	
CEL24LQFD-G	24-port 40GE interface card (FD-G, QSFP+)
CEL36LQFD-G	36-port 40GE interface card (FD-G, QSFP+)
CEL36LQFD-A	36-port 40GE interface card (FD-A, QSFP+)
CE-L36LQJ	36-port 40GE interface card (L-J, QSFP+)

100GE Interface Card	
CEL48CQHG-P	48-port 100GE interface card (HG-P, QSFP28)(CM)
CEL48CQHG-P0	48-port 100GE interface card with RTU bundle 0 (HG-P, QSFP28, 8*2*100GE port RTU)
CEL48CQ-SAN0	48-port 100GE SAN interface card with RTU bundle 0 (HG-P, QSFP28, 8*2*100GE port RTU)
CEL48CQ-SAN1	48-port 100GE SAN interface card with optical module bundle 1 (HG-P, QSFP28, 8*2*100GE port RTU, 48*100GBase-SR4-MP)
CEL18CQFD-G	18-port 100GE interface card (FD-G, QSFP28)

100GE Interface Card	
CEL36CQFD-G	36-port 100GE interface card (FD-G, QSFP28)
CEL18CQFD-A	18-Port-100GE Interface Card(FD, QSFP28)
CE-L36CQJ	36-port 100GE interface card (L-J, QSFP28)
CE-L24CQJ	24-port 100GE interface card (L-J, QSFP28)

400GE Interface Card	
CEL36DQHG-P	36-port 400GE interface card (HG-P, QSFP-DD)
CEL48DQHG-P	48-port 400GE interface card (HG-P, QSFP-DD)
CEL36DQHG-X	36-port 400GE interface card (HG-X, QSFP-DD)

Power	
PAH-3000WA	3000W Dual Inputs AC&HVDC Power Module
PDC-2200WC	PSU-PDC-2200WC-PM2200W DC Power Supply

Hardware RTU	
CE168-RTU-2CQ	2*100G Port Right-to-Use for CloudEngine 16800
CE168-RTU-24S	24*10G Port Right-to-Use for CloudEngine 16800
CE168-RTU-U24S	24*10G Upgrade to 24*25G Right-to-Use for CloudEngine16800

Software	
N1-CE168LIC-CFFD	N1-CloudFabric Foundation SW License for CloudEngine 16800
N1-CE168CFFD-SnS1Y	N1-CloudFabric Foundation SW License for CloudEngine 16800-SnS-Year
N1-CE168LIC-CFAD	N1-CloudFabric Advanced SW License for CloudEngine 16800
N1-CE168CFAD-SnS1Y	N1-CloudFabric Advanced SW License for CloudEngine 16800-SnS-Year
N1-CE168LIC-CFPM	N1-CloudFabric Premium SW License for CloudEngine 16800
N1-CE168CFPM-SnS1Y	N1-CloudFabric Premium SW License for CloudEngine 16800-SnS-Year
N1-CE168LIC-AFRD-2	N1-CloudEngine 16800 AI Fabric RDMA Application Acceleration Function 2
N1-CE168AFRD2-SnS1Y	N1-CloudEngine 16800 AI Fabric RDMA Application Acceleration Function 2-SnS-Year
N1-CE168LIC-NOF+	N1-CloudEngine 16800 AI Fabric Value-added Package for the NOF+ Storage Scenarios
N1-CE168NOF+-SnS1Y	N1-CloudEngine 16800 AI Fabric Value-added Package for the NOF+ Storage Scenarios-SnS-Year
N1-CE168LIC-SEC	N1-CloudEngine 16800 Security Function
N1-CE168SEC-SnS1Y	N1-CloudEngine 16800 Security Function-SnS-Year

Software	
N1-CE-M-LIC-MDCA	N1-CloudEngine Data Center Switch Multi-cloud Multi-DC Value-added Package - Modular
N1-CEMMDCASnS1Y	N1-CloudEngine Data Center Switch Multi-cloud Multi-DC Value-added Package, Per Modular device -SnS-Year
N1-CE168UPG-F-A	N1-CloudEngine 16800 Upgrade SW License:Foundation to Advanced
N1-CE168UGFA-SnS1Y	N1-CloudEngine 16800 Upgrade SW License:Foundation to Advanced-Subscription and Support-Year
N1-CE168UPG-A-P	N1-CloudEngine 16800 Upgrade SW License:Advanced to Premium
N1-CE168UGAP-SnS1Y	N1-CloudEngine 16800 Upgrade SW License:Advanced to Premium-Subscription and Support-Year

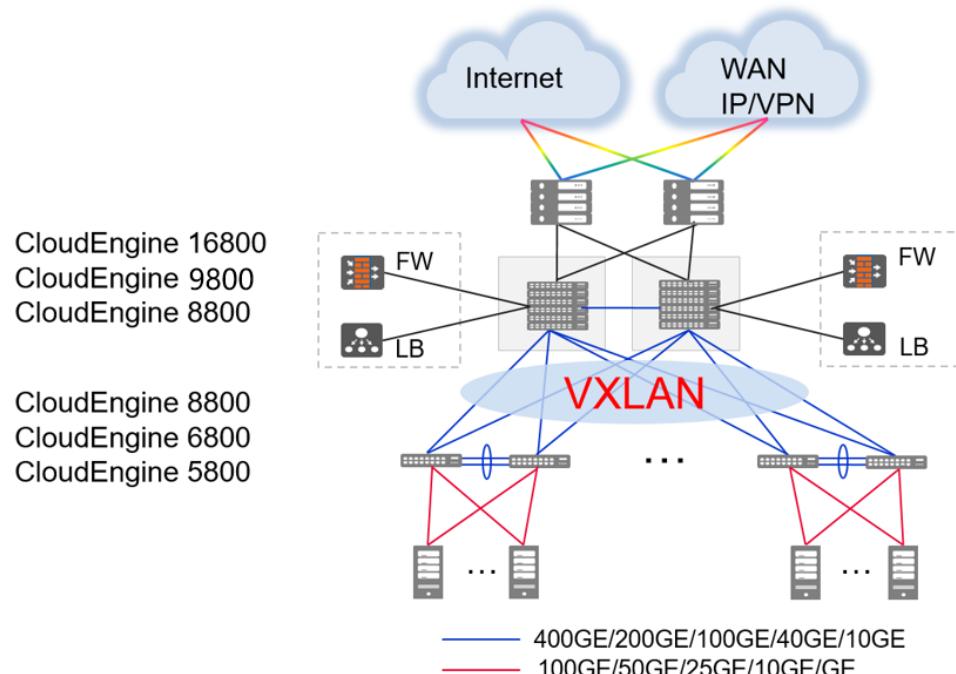
Note:

- The active and standby MPU slots must have MPUs of the same model installed.
- HG-P series LPUs are used with H-P and J-P series SFUs.
- FD-G series line cards (V2) and HG-P/SAN series line cards (V3) cannot be installed on the same device, F/G-G series SFUs (V2) and H/J-P series SFUs (V3) cannot be installed on the same device.

## Networking and Application

### Data Center Applications

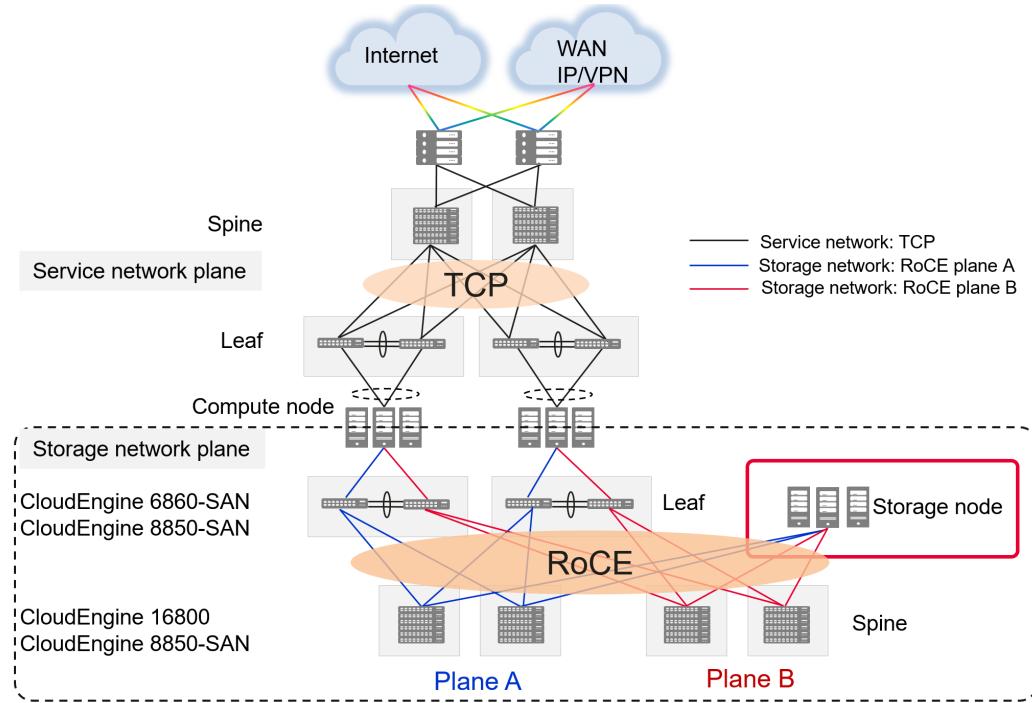
On a typical data center network, CloudEngine 16800 switches work as core switches, whereas CloudEngine 8800/CloudEngine 6800/CloudEngine 5800 switches work as TOR switches and connect to the core switches using 400GE/200GE/100GE/40GE/10GE ports. The core and TOR switches use fabric technologies such as VXLAN to build a non-blocking large Layer 2 network, which allows for 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.

## Typical Applications on the All-Flash Storage DCN

On the all-flash storage DCN, CloudEngine 16800 or CloudEngine 8850-SAN series are used as the core spine switches, and CloudEngine 6860-SAN or CloudEngine 8850-SAN series are used as the leaf switches. The leaf and spine switches are fully interconnected through 100GE or 400GE links. These switches provide a lineup of features such as iNOF, AI ECN, PFC, and PFC deadlock detection to implement plug-and-play of the storage system. In this manner, links are automatically established in seconds based on the network plan once computing and storage nodes are brought online. This helps build a network environment featuring zero packet loss, low latency, and high throughput for RoCEv2 traffic, meeting high-performance requirements of RoCEv2 applications.



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

No part of this document may be reproduced or transmitted 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 contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase 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 518129 People's Republic of China

Website: [www.huawei.com](http://www.huawei.com)