CloudEngine 6860 Series Data Center Switches

HUAWEI TECHNOLOGIES CO., LTD.

HUAWEI

CloudEngine 6860 Series Data Center Switches

Product Overview

Huawei CloudEngine 6860 (CE6860 for short) series switches are next-generation 25G Ethernet switches designed for data centers and high-end campus networks, providing high-performance, high-density 25GE ports, and low latency. The CE6860 has an advanced hardware architecture with 100GE uplink ports and the industry's highest density of 25GE access ports.

Using the Huawei VRP8 software platform, CE6860 switches provide extensive data center service features and high stacking capability. In addition, the airflow direction (front-to-back or back-to-front) can be changed. CE6860 switches can work with CE12800 switches to build an elastic, virtualized, high-quality 100GE fully-connected fabric that meets the requirements of cloud-computing data centers.

CE6860 switches provide high-density 10GE/25GE access to help enterprises and carriers build a scalable data center network platform in the cloud computing era. They can also be used as aggregation or core switches for enterprise campus networks.

Product Appearance

The CE6860 comes In two models.

CE6860-48S8CQ-EI

48*25GE SFP28 ports , 8*100GE QSFP28 ports

CE6865-48S8CQ-EI

48*25GE SFP28 ports , 8*100GE QSFP28 ports

Product Characteristics

High-Density 25GE Access

- The CE6860 is the industry's highest-performing 1 U ToR switch. It provides 3200mpps forwarding performance and supports L2/L3 line-rate forwarding.
- The CE6860 provides up to 48*25GE ports, the highest 25GE port density among 1U ToR switches, allowing for high density 10G/25G server access and smooth evolution.
- The CE6860 provides eight 100GE QSFP28 ports. Each QSFP28 port can be used as one 40GE QSFP+ port, four 25GE SFP28 ports, or four 10GE SFP+ ports, providing flexibility in networking. The 100GE uplink ports can be connected to CE12800 switches to build a non-blocking network platform.

Highly Reliable, High-Performance Stacking

- The industry's first 16-member stack system
  - » A stack system of 16 member switches has a maximum of 768*25GE access ports that provide high-density server access in a data center.
  - » Multiple switches in a stack system are virtualized into one logical device, making it possible to build a scalable, easy-to-manage data center network platform.
  - » A stack system separates the control plane from the data plane. This eliminates the risk of single points of failure and greatly improves system reliability.
- Long-distance, highly reliable stacking
  - » The CE6860 can use service ports as stack ports. A stack system can be established with switches in the same rack or different racks, and even over long distances.
  - » Service and stack bandwidths can be allocated based on the network's scale so that network resources can be used more efficiently.

Inter-device Link Aggregation, High Efficiency and Reliability

- The CE6860 supports multichassis 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 system 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, TRILL, VXLAN, and IP networks, allowing for flexible networking.
- With the industry's most comprehensive inter-device link aggregation technology, the device networking coupling relationship evolves from stacking at the control plane to the use of M-LAG and then finally to coupling-free M-LAG Lite. This achieves active-active server access and zero interruption of services when upgrading switches.

Large-Scale Routing Bridge, On-Demand Scaling

- The CE6860 supports the IETF Transparent Interconnection of Lots of Links (TRILL) protocol and can connect to 25G and 10G servers simultaneously. CE6850 switches can establish a large Layer 2 TRILL network with more than 500 nodes, enabling flexible service deployments and large-scale Virtual Machine (VM) migrations.
- The TRILL protocol uses a routing mechanism similar to IS-IS and sets a limited time to live(TTL) value in packets to prevent Layer 2 loops. This significantly improves network stability and speeds up network convergence.
- On a TRILL network, all data flows are forwarded quickly using Shortest Path First (SPF) and Equal-cost Multi-path (ECMP) routing. SPF and ECMP avoid the suboptimal path selection problem in STP and increase link bandwidth efficiency to 100 percent.
- The CE6860 supports TRILL-based Layer 2 equal-cost paths, greatly improving links' load balancing capabilities. The network has a fat-tree architecture that enhances expansion.

Hardware Overlay Gateway Achieves Fast Service Deployment

- The CE6860 can work with a mainstream virtualization platform. As the high-performance, hardware gateway of an overlay network (VXLAN), the CE6860 can support more than 16 million tenants.
- The hardware gateway deployment enables fast service deployment without changing the customer network, providing investment protection.
- The CE6860 supports Border Gateway Protocol - Ethernet VPN (BGP-EVPN), which can run as the VXLAN control plane to simplify VXLAN configuration within and between data centers.

Converged Enhanced Ethernet, Allowing for Data, Storage, and Computing Traffic on One Network

- CE6860 series switches support Fibre Channel over Ethernet (FCoE), which permits storage, data, and computing services to be transmitted on one network, reducing the costs of network construction and maintenance.
- CE6860 series switches support centralized FCoE gateway deployment, which makes network O&M simpler.
- Various CE6860 series switches support multiple data center features: Priority-based Flow Control (PFC), Enhanced Transmission Selection (ETS) and Data Center Bridging eXchange (DCBX). These features ensure low latency and zero packet loss for FC storage and high-speed computing services.

Full Openness and Programmability, Flexible Customization

- The CE6860 uses the Open Programmability System (OPS) embedded in the VRP8 software platform 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) and third-party controllers. 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 CE6860 supports CE modules for Ansible, which enables unified provisioning of physical and virtual networks.
- CE6860 switches can seamlessly integrate with systems of F5, an industry-leading application delivery network provider, to build an active-active data center network.
- The OPS provides seamless integration of data center service and network in addition to a service-oriented, software-defined networking (SDN).

Zero Touch Provisioning, Automatic O&M

- The CE6860 supports Zero Touch Provisioning (ZTP). ZTP enables the CE6860 to automatically obtain and load version files from a USB flash drive or file server, freeing network engineers from onsite configuration or deployment. ZTP reduces labor costs and improves device deployment efficiency.

- ZTP provides built-in scripts for users through open APIs. Data center personnel can use the programming language they are familiar with, such as Python, to provide unified configuration of network devices.
- ZTP decouples configuration time of new devices from device quantity and area distribution, which improves service provisioning efficiency.

Intelligent O&M with the FabricInsight Solution

- The CE6860 provides proactive path detection on the entire network. It periodically checks sample flows to determine connectivity of all paths on the network and locates failure points, enabling you to know the network health in real time.
- The CE6860 supports visualization of all flows and congestion, improving service experience.
- The CE6860 supports global, precise time synchronization based on IEEE 1588v2, detecting delay with sub-microsecond accuracy.

Flexible Airflow Design, High Energy Efficiency

- Flexible front-to-back/back-to-front airflow design
  - » The CE6860 uses a front-to-back/back-to-front airflow design that isolates cold air channels from hot air channels. This design meets heat dissipation requirements in data center equipment rooms.
  - » Air can flow from front to back, or back to front when different fans and power modules are used.
  - » Redundant power modules and fans can be configured to ensure uninterrupted service transmission.
- Energy-saving technology
  - » The CE6860 series switches have energy-saving chips and can measure system power consumption in real time. Fan speed can be adjusted dynamically based on system consumption. These energy-saving technologies reduce O&M costs and contribute to a greener data center.

AI Fabric, Improving Reliability of High-Performance Computing

- Automatic buffer configuration, eliminating packet loss
  - » The packet buffer of the forwarding chip is automatically configured at boot time.
  - » The threshold for the buffer of the forwarding chip is configurable, ensuring the optimal application performance for specific network topologies and traffic models.
- Dynamic traffic priority adjustment
  - » The scheduling priority of latency-sensitive microflows can be dynamically adjusted to preferentially schedule them, guaranteeing the performance of latency-sensitive applications.
- Traffic congestion control
  - » Fast ECN is supported. When any packet leaves a queue, the ECN flag is set according to the congestion status of the queue. This shortens the delay caused by the queue depth.
  - » Fast CNP is supported. A switch directly sends the CNP to the NIC of the source server, shortening the CNP feedback path.

» The dynamic ECN threshold is supported. The ECN threshold can be dynamically adjusted for a queue according to traffic changes.

- Dynamic load balancing (DLB)
  - » The ECMP and LAG support the DLB function.
- Monitoring prioritized lossless traffic under control
  - » The PFC can be mapped based on DSCP.

Clear Indicators, Simple Maintenance

- Clear indicators
  - » Port indicators clearly show port status and port speeds. The 100GE port indicators can show the states of all ports derived from the 100GE ports.
  - » State and stack indicators on both the front and rear panels enable operators to maintain the switch from either side.
  - » CE6860 series switches support remote positioning. Operators can turn on remote positioning indicators on the switches they want to maintain, so that they can find switches easily in an equipment room full of devices.
- Simple maintenance
  - » The management port, fans, and power modules are on the front panel, which facilitates device maintenance.
  - » Data ports are located at the rear, facing servers. This simplifies cabling.

Product Specifications

Item	CE6860-48S8CQ-EI	CE6865-48S8CQ-EI
25GE SFP28 ports	48	
100GE QSFP28 ports	8	
Switching capacity	4 Tbit/s (Switching capacity after stacking: 64 Tbit/s)	
Forwarding rate	3200 mpps	2000 mpps
Airflow design	Front-to-back or back-to-front	
Device virtualization	iStack[1]	
	M-LAG	
Network virtualization	TRILL	
	VXLAN routing and bridging	
	BGP-EVPN	
	QinQ access VXLAN	

1 For details about the configuration, please see: http://support.huawei.com/onlinetoolsweb/virtual/en/dc/stack_index.html?dcb

Item	CE6860-48S8CQ-EI	CE6865-48S8CQ-EI
Data center interconnect	VXLAN mapping, implementing interconnection between multiple DC networks at Layer 2	
SDN	Agile Controller	
	VMware NSX	
Network convergence	FCoE	
	DCBX, PFC, ETS	
programmability	OPS	
	OpenFlow	
	CE modules for Ansible released on open source websites	
Traffic analysis	NetStream	
	sFlow	
VLAN	Adding access, trunk, and hybrid interfaces to VLANs	
	Default VLAN	
	QinQ	
	MUX VLAN	
	GVRP	
ACL	Ingress 2750, Egress 1k	ingress 7662, egress 2k
MAC address table	Maximum:136k	Maximum: 288k
	Dynamic learning and aging of MAC addresses	
	Static, dynamic, and blackhole MAC address entries	
	Packet filtering based on source MAC addresses	
	MAC address limiting based on ports and VLANs	
ARP	Maximum: 84k	Maximum: 168k
ND	Maximum: 32k	
IPv4 FIB	Maximum: 192k	Maximum: 380k
IP routing	IPv4 routing protocols, such as RIP, OSPF, BGP, and IS-IS	
	IPv6 routing protocols, such as RIPng, OSPFv3, IS-ISv6, and BGP4+	
IPv6 FIB	Maximum: 64k	Maximum: 256k
IPv6	IPv6 Neighbor Discovery (ND)	
	IPv6 VXLAN over IPv4	
	Path MTU Discovery (PMTU)	
	TCP6, ping IPv6, tracert IPv6, socket IPv6, UDP6, and Raw IP6	

Item	CE6860-48S8CQ-EI	CE6865-48S8CQ-EI
Multicast FIB	Maximum: 8k	
Multicast	IGMP, PIM-SM, PIM-DM, MSDP, and MBGP	
	IGMP snooping	
	IGMP proxy	
	Fast leaving of multicast member interfaces	
	Multicast traffic suppression	
	Multicast VLAN	
	Multicast VXLAN	
MPLS	Multi-Protocol Label Switching	
Reliability	LACP	
	STP, RSTP, VBST, and MSTP	
	BPDU protection, root protection, and loop protection	
	Smart Link and multi-instance	
	DLDP	
	ERPS (G.8032)	
	VRRP, VRRP load balancing, and BFD for VRRP	
	BFD for BGP/IS-IS/OSPF/Static route	
	BFD for VXLAN	
QoS	Traffic classification based on Layer 2 headers, Layer 3 protocols, Layer 4 protocols, and 802.1p priority	
	Actions of ACL, CAR, re-marking, and scheduling	
	Queue scheduling algorithms, including PQ, WRR, DRR, PQ+WRR, and PQ+DRR	
	Congestion avoidance mechanisms, including WRED and tail drop	
	Traffic shaping	
O&M	Network-wide path detection	
	IEEE 1588v2[2]	
	Telemetry	
	INT and ERSPAN+[3]	
	Statistics on the buffer microburst status	
	VXLAN OAM: VXLAN ping, VXLAN tracert	

2 Supported only by the CE6865

3 Supported only by the CE6865

Item	CE6860-48S8CQ-EI	CE6865-48S8CQ-EI
AI Fabric[4]	VIQ	
	Fast ECN	
	Fast CNP	
	Dynamic load balancing (DLB)	
Configuration and maintenance	Console, Telnet, and SSH terminals	
	Network management protocols, such as SNMPv1/v2c/v3	
	File upload and download through FTP and TFTP	
	BootROM upgrade and remote upgrade	
	802.3az Energy Efficient Ethernet (EEE)	
	Hot patches	
	User operation logs	
	ZTP	
Security and management	802.1x authentication	
	Command line authority control based on user levels, preventing unauthorized users from using commands	
	DoS, ARP, and ICMP attack defenses	
	Port isolation, port security, and sticky MAC	
	Binding of the IP address, MAC address, interface number, and VLAN ID	
	Authentication methods, including AAA, RADIUS, and HWTACACS	
	Remote Network Monitoring (RMON)	
Dimensions (W x D x H)	442 mm x 420 mm x 43.6 mm	
Weight (fully loaded)	8.8 kg (19.4 lb)	
Environmental parameters	Operating temperature: 0°C to 40°C (32°F to 104°F) (0 m to 1,800 m) Storage temperature: -40°C to +70°C (-40°F to 158°F) Relative humidity: 5% RH to 95% RH, non-condensing	
Operating voltage	AC: 90 V to 290 V DC: -38.4 V to -72 V	
Maximum power consumption	336 W	470W

4 Supported only by the CE6865

Ordering Information

Mainframe	
CE6860-48S8CQ-EI	CE6860-48S8CQ-EI Switch (48-Port 25GE SFP28, 8*100GE QSFP28, Without Fan and Power Module)
CE6860-EI-F-B0B	CE6860-48S8CQ-EI Switch (48-Port 25GE SFP28, 8*100GE QSFP28, 2*AC Power Module, 2*FAN Box, Port-side Exhaust)
CE6860-EI-B-B0B	CE6860-48S8CQ-EI Switch (48-Port 25GE SFP28, 8*100GE QSFP28, 2*AC Power Module, 2*FAN Box, Port-side Intake)
CE6865-48S8CQ-EI	CE6865-48S8CQ-EI Switch(48-Port 25GE SFP28,8*100GE QSFP28,Without Fan and Power Module)
CE6865-EI-F-B0B	CE6865-48S8CQ-EI Switch(48-Port 25GE SFP28,8*100GE QSFP28,2*AC Power Module,2*FAN Box,Port-side Exhaust)
CE6865-EI-B-B0B	CE6865-48S8CQ-EI Switch(48-Port 25GE SFP28,8*100GE QSFP28,2*AC Power Module,2*FAN Box,Port-side Intake)

Fan box		
Model	Description	Applicable Product
FAN-40HA-F	Fan box (HA, Front to Back, FAN panel side intake)	CE6860-48S8CQ-EI and CE6865-48S8CQ-EI
FAN-40HA-B	Fan box (HA, Back to Front, FAN panel side exhaust)	CE6860-48S8CQ-EI and CE6865-48S8CQ-EI

Power		
Model	Description	Applicable Product
PAC-600WA-F	600W AC Power Module (Front to Back, Power panel side intake)	CE6860-48S8CQ-EI and CE6865-48S8CQ-EI
PAC-600WA-B	600W AC Power Module (Back to Front, Power panel side exhaust)	CE6860-48S8CQ-EI and CE6865-48S8CQ-EI
PDC-350WA-F	350W DC Power Module (Front to Back, Power panel side intake)	CE6860-48S8CQ-EI and CE6865-48S8CQ-EI
PDC-350WA-B	350W DC Power Module (Back to Front, Power panel side exhaust)	CE6860-48S8CQ-EI and CE6865-48S8CQ-EI

Software	
CE68-LIC-VXLAN	CloudEngine 6800 VXLAN Function
CE68-LIC-FCF16	CloudEngine 6800 FCF 16 Ports
CE68-LIC-FCFAL	CloudEngine 6800 FCF All Ports
CE68-LIC-BUN01	CE6800 Function License Bundle 1
CE6800-LIC-NPV	CloudEngine 6800 FCOE NPV Function
CE68-LIC-TLM	CE6800 Telemetry Function
CE68-LIC-LLETH	CloudEngine 6800 Lossless Ethernet Fabric Function
CE68-LLETH-SYS1Y	CloudEngine 6800 Lossless Ethernet Fabric Function-SnS-1 Year
CE68-LIC-BASE	CE6800 Basic Software Function
CE68-LIC-PTP	CE6800 Precision Time Protocol Function

Networking and Applications

Data Center Applications

On a typical data center network, CE12800/CE8800/CE7800 switches work as core switches, and CE6860 switches work as ToR switches. CE6860 switches use 100GE ports to set up 100GE full connections with CE12800/CE8800 core switches, or use 40GE ports to connect to CE7800 switches. The core and ToR switches use fabric technology such as TRILL or VXLAN to establish a non-blocking large Layer 2 network, which allows large-scale VM migrations and flexible service deployments.

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

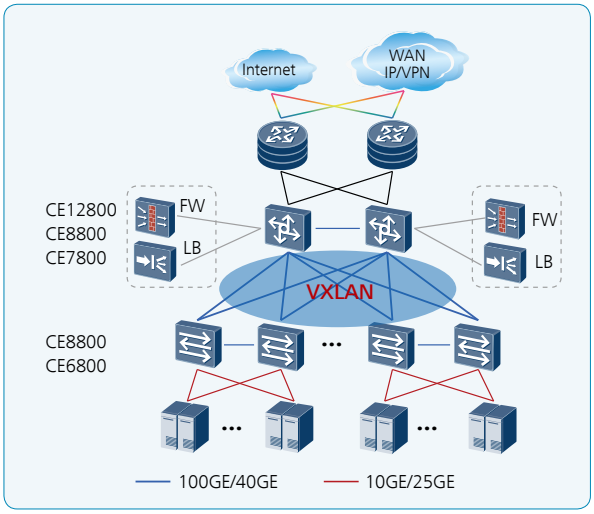

Campus Network Applications

CE6860 switches can be used as aggregation or core switches on a campus network. Their high-density, line-rate 10GE/25GE ports and high stacking capability can meet the ever-increasing demand for network bandwidth. CE6860 switches are cost-effective campus network switches, thanks to their extensive service features and innovative energy-saving technologies.

On a typical campus network, multiple CE12800/CE8800/CE7800 switches are virtualized into a logical core switch using CSS or iStack technology. Multiple CE8800/CE7800/CE6800 switches at the aggregation layer form a logical switch using iStack technology. CSS and iStack improve network reliability and simplify network management. At the access layer, CE6800/CE5800 switches are virtualized with CloudFabric technology, such as iStack or M-LAG (vertical virtualization), to provide high-density line-rate ports.

Note: CSS, iStack, and M-LAG are also widely used in data centers to facilitate network management.

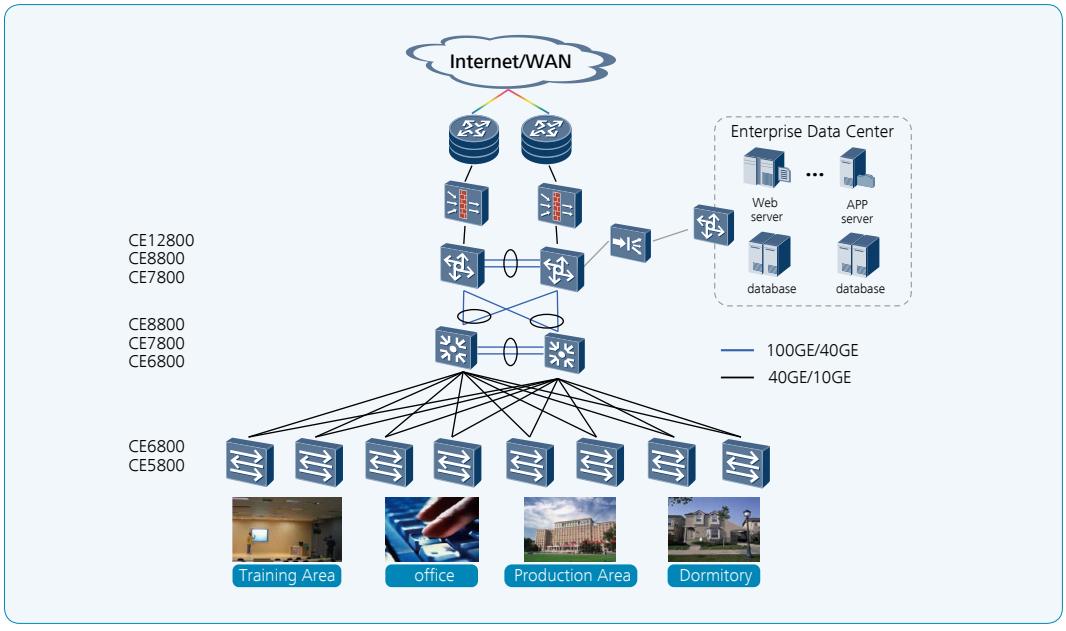

HUAWEI TECHNOLOGIES CO.,LTD.
Huawei Industrial Base
Bantian Longgang
Shenzhen 518129,P.R.China
Tel: +86 755 28780808

www.huawei.com