CloudEngine 12800 Series
Data Center Core Switches

HUAWEI TECHNOLOGIES CO., LTD.

HUAWEI

CloudEngine 12800 Series Data Center Core Switches

Product Overview

The CloudEngine 12800 (CE12800) series switches are next-generation, high-performance core switches designed for data center networks and high-end campus networks. Using Huawei's next-generation VRP8 software platform, CE12800 series switches provide stable, reliable, secure, high-performance L2/L3 switching capabilities to help build an elastic, virtualized, and high-quality network.

CE12800 series switches use advanced hardware architecture design and have the highest performance of any currently available core switches. The CE12800 provides as much as 160 Tbit/s switching capacity and high-density line-speed ports. Each switch has up to 576*100GE, 576*40GE, 2,304*25GE, or 2,304*10GE ports.

The CE12800 series switches use the Clos architecture and provide comprehensive virtualization capabilities along with data center service features. The switches use innovative energy-saving technologies to greatly reduce power consumption. In addition, a front-to-back airflow design provides industrial-grade reliability.

Product Appearance

The CE12800 is available in six models: CE12804S, CE12808S, CE12804, CE12808, CE12812, and CE12816. The CE12800 series uses interchangeable components to reduce costs on spare parts. This design ensures device scalability and protects customers' investment.

Product Characteristics

Next-Generation Core Engine Provides the Industry's Highest Performance

160Tbit/s Switching Capacity

- The CE12800 provides 10 Tbit/s per-slot bidirectional bandwidth (scalable to 20 Tbit/s) and a maximum of 160 Tbit/s switching capacity (scalable to more than 320 Tbit/s). This capacity can support sustainable development of cloud-computing data centers for the next 10 years.

- The CE12800, together with the CE8800/CE7800/CE6800/CE5800 series of Top-of-Rack (TOR) switches, can implement the largest non-blocking switching network in the industry. This network can provide access for up to tens of thousands of 25GE or 10GE servers.

4Terabit High-Density Line Cards

- The forwarding capacity of a line card can reach up to 1.2 Tbit/s.
- The CE12800 line cards provide the industry's highest port densities, from 36*40GE to 36*100GE/144*25GE/144*10GE.
- The CE12800 provides as many as 576*100GE, 576*40GE, 2,304*25GE, or 2,304*10GE line-speed ports.

Super-Large Buffer of 18 GB

- All service ports (100GE/40GE/10GE/GE) support a super-large buffer.
- The distributed buffer mechanism on inbound interfaces can effectively handle incast traffic loads in data centers.
- The line card provides up to 18 GB buffer, which is dynamically shared by interfaces to improve usage efficiency.

Back-to-Back Non-Blocking System

- The CE12816 is the industry's first data center core switch to support a non-blocking system. Two CE12816s can be upgraded to the CE12832 through back-to-back Cluster Switch Systems (CSSs) of Switch Fabric Units (SFUs). The new device provides 32 service slots.
- The CE12832 builds a strict non-blocking system using Clos. All traffic between two CE12816 chassis can be forwarded without occupying any service interface.
- The CE12832 can be upgraded through a CE12816 in-service upgrade. During the upgrade, services are not affected. This ensures continuous evolution and expansion of the customer service system.

Comprehensive Virtualization Capabilities Implement Simple, Efficient Networking

VS Implements On-demand Resource Sharing

- Highest device virtualization capability: The CE12800 uses Virtual System (VS) technology to provide an industry-leading virtualization capability that enables one switch to be virtualized into as many as sixteen logical switches. This 1:16 ratio enables one core switch to manage services for an enterprise's multiple service areas such as production, office, and DMZ, or for multiple tenants.
- Higher security and reliability: VS technology divides a network into separate logical areas for service isolation. The failure of one virtual switch does not affect other virtual switches, enhancing network security.
- Lower CAPEX: VS technology improves the use efficiency of physical devices by implementing on-demand resource allocation. This ensures network scalability and reduces investment in devices.
- Lower OPEX: Using one physical device to implement multiple logical devices saves space in a data center equipment room and reduces the cost of device maintenance.

CSS Simplifies Network Management

- The CE12800 uses industry-leading CSS technology, which can virtualize multiple physical switches into one logical switch to facilitate network management and improve reliability.
- The CE12800 provides the dedicated system inter-connect port and separates the control channel from the service channel, improving reliability.
- The CE12800 provides a cluster bandwidth of 1.6 Tbit/s. This super-high bandwidth prevents traffic bottlenecks on data center networks.

- The CE12800 switches establish a cluster using service ports with distances of up to 80 km between cluster member switches.
- The CE12800 uses CSS+VS synergy technology to turn a network into a resource pool so that network resources can be allocated on demand. On-demand resource allocation is ideal for the cloud-computing service model.

Dual Management and Control Planes of M-LAG Guarantee High Reliability Service

- The management and control planes on one Multichassis Link Aggregation Group (M-LAG) node are independent from that on the other, which substantially improves system reliability.
- The two nodes of an M-LAG can be upgraded independently from each other. During the upgrade of one node, the other node takes over forwarding the services on the first node, ensuring that the services remain uninterrupted.
- M-LAG is able to seamlessly collaborate with CSS, thus enabling highly reliable 4-to-1 virtualization.

Large-Scale Routing Bridge Supports Flexible Service Deployment

- The CE12800 series switches support Transparent Interconnection of Lots of Links (TRILL), a standard IETF protocol. The TRILL protocol helps build a large Layer 2 network with more than 500 nodes, which permits flexible service deployments and Virtual Machine (VM) migrations. A TRILL network can use 10GE/GE servers.
- 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 problem of suboptimal path selection in the Spanning Tree Protocol (STP) and increase link bandwidth efficiency to 100 percent.
- The CE12800 supports up to 32 TRILL-based Layer 2 equal-cost paths, greatly improving links' load-balancing capabilities. The network's fat-tree architecture supports easy expansion.

EVN Supports Resource Sharing Across Data Centers

- The Ethernet Virtual Network (EVN) implements inter-data center Layer 2 interconnection across the IP WAN, and integrates multiple data centers into a large IT resource pool. VMs can migrate between data centers. EVN supports Layer 2 interconnection of a maximum of 32 data centers, which enable scalability that is 5 times higher than the industry. EVN combines the advantages of the Border Gateway Protocol (BGP) and Virtual Extensible LAN (VXLAN) to provide high scalability and highly efficient use of bandwidth.
- High scalability: Based on BGP at the control plane, the CE12800 supports millions of MAC addresses and routes, 32K tenants, and 256K VMs.
- Highly efficient bandwidth usage: The forwarding plane uses VXLAN encapsulation. Flow-based load balancing is implemented on the entire network, which optimizes bandwidth usage.

Agile Controller Implements Fast VM Migration

- The CE12800 works with Huawei's Agile Controller to permit network policies to be dynamically deployed on the CE12800. Agile Controller also supports online VM migration.
- Agile Controller delivers network policies through high-speed RADIUS interfaces. Its online VM migration is 10 to 20 times the rate of other industry platforms, enabling large-scale VM migrations.
- Agile Controller is based on open APIs and is compatible with all major virtualization platforms including VMware.

Fully Programmable Switch Permits Agile Service Provisioning

OPS Implements Programmability at the Control Plane

- The CE12800 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 OPS provides seamless integration of data center service and network in addition to a service-oriented, Software-Defined Network (SDN).

Virtualized Gateway Achieves Fast Service Deployment

- The CE12800 can work with a mainstream virtualization platform. As the high-performance, hardware gateway of an overlay network (VXLAN), a CE series switch can support more than 16M tenants.
- The CE12800 can connect to a cloud platform through an open API to provide unified management of software and hardware networks.
- This function implements fast service deployment without changing the customer network. It also protects customer investments.

ZTP Implements Zero-Configuration Deployment

- The CE12800 supports Zero Touch Provisioning (ZTP). ZTP enables the CE12800 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.

Advanced Architecture Ensures Industry-Leading Network Quality

High-Performance, Non-blocking Switching Architecture

- The CE12800's non-blocking switching architecture includes an orthogonal switch fabric design, Clos architecture, cell switching, Virtual Output Queuing (VOQ), and a super-large buffer.
- Orthogonal switch fabric design: CE12800 service line cards and switch fabric units (SFUs) use an orthogonal design in which service traffic between line cards is directly sent to the SFUs through orthogonal connectors. This approach reduces backplane cabling and minimizes signal attenuation. The orthogonal design can support signal rates as high as 25 Gbit/s per Serdes, which is 2.5 times the industry average. This design greatly improves system bandwidth and evolution capabilities, enabling the system switching capacity to scale to more than 100 Tbit/s.
- Clos architecture: The CE12800'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 easily copes with complex, volatile traffic in data centers.

- VOQ: The CE12800 supports 96,000 VOQ queues that implement fine-grained Quality of Service (QoS) based on the switch fabrics. With the VOQ mechanism and super-large buffer on inbound interfaces, the CE12800 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

- Hot backup of five key components: Main Processing Units (MPUs) and Centralized Monitoring Unit (CMUs) work in 1+1 hot backup mode. SFUs work in N+M hot backup mode. Power supplies support dual inputs and N+N backup and have their own fans. Both fan trays work in 1+1 backup mode; each fan tray has two counter-rotating 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.
- Independent triple-plane design: The independent control, data, and monitoring planes of the CE12800 improve system reliability and ensure service continuity.

High-Performance VRP8 Software Architecture

- The CE12800 takes advantage of Huawei's next-generation VRP8, a high-performance, highly reliable modular software platform that provides continuous services.
- Fine-grained distributed architecture: VRP8, the industry's high-end software platform, uses a fine-grained, fully distributed architecture that can process network protocols and services concurrently using multiple instances. This architecture takes full advantage of multi-core/multi-CPU processes to maximize performance and reliability.
- Highly reliable In-Service Software Upgrade (ISSU): VRP8 supports ISSU.

Pioneering Energy-saving Technologies

Strict Front-to-Back Airflow Design

- The CE12800 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.
- Line cards and SFUs use independent airflow channels, which solve the problems of mixing hot and cold air and cascade heating, and effectively reduce energy consumption in equipment rooms.
- Each fan tray has two counter-rotating fans, ensuring efficient heat dissipation.
- 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 CE12800 uses innovative energy saving technologies. . The port power consumption is merely half of the industry average. It greatly reduces power consumption in the data center equipment room.
- Miercom has performed a series of strict tests for the CE12800, proving its low power consumption.

Efficient, Intelligent Power Supply System

- The CE12800 incorporates the industry's most efficient digital power modules, which provide power efficiency of 96 percent.
- 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 CE12800 can save energy dynamically by adjusting the power consumption of components to adapt to changes in service traffic volume.

Product Specifications

Item	CE12804S	CE12808S	CE12804	CE12808	CE12812	CE12816
Switching capacity	40/80 Tbit/s	80/160 Tbit/s	40/80 Tbit/s	80/160 Tbit/s	120/240 Tbit/s	160/320 Tbit/s
Forwarding performance	17,280 Mpps	34,560 Mpps	17,280 Mpps	34,560 Mpps	51,840 Mpps	69,120 Mpps
Service slots	4	8	4	8	12	16
Switching fabric module slots	2	4	6	6	6	6
Fabric architecture	Clos architecture, cell switching, VoQ, and distributed large buffer					
Airflow design	Strict front-to-back					
Device virtualization	Virtual System (VS)					
	Cluster Switch System (CSS)					
	Super Virtual Fabric (SVF)					
Network virtualization	M-LAG					
	TRILL					
	VXLAN routing and bridging					
	EVPN					
VM awareness	Agile Controller					
Network convergence	FCoE					
	DCBX, PFC, ETS					
Data center interconnection	EVN supports inter-DC Layer 2 network interconnections					
Programmability	OpenFlow					
	OPS					
	Puppet, and OVSDB plugins released on open source websites					
	Linux container for open source and customization programming					
Traffic analysis	NetStream					
	sFlow					
VLAN	Adding access, trunk, and hybrid interfaces to VLANs					
	Default VLAN					
	QinQ					
	MUX VLAN					
MAC address	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					
IP routing	IPv4 dynamic routing protocols, such as RIP, OSPF, IS-IS, and BGP					
	IPv6 dynamic routing protocols, such as RIPng, OSPFv3, ISISv6, and BGP4+					

Item	CE12804S	CE12808S	CE12804	CE12808	CE12812	CE12816
IPV6	IPv6 Neighbor Discovery (ND)					
	Path MTU Discovery (PMTU)					
	TCP6, ping IPv6, tracert IPv6, socket IPv6, UDP6, and Raw IP6					
Multicast	IGMP,PIM-SM,PIM-DM,MSDP,MBGP					
	IGMP snooping					
	IGMP proxy					
	Fast leave of multicast member interfaces					
	Multicast traffic suppression					
	Multicast VLAN					
MPLS	Basic MPLS functions					
	MPLS VPN/VPLS/VPLS over GRE					
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					
	In-Service Software Upgrade (ISSU)					
QoS	Traffic classification based on Layer 2, Layer 3, Layer 4, and priority information					
	Actions include ACL, CAR, and re-marking					
	Queue scheduling modes such as PQ, WFQ, and PQ+WRR					
	Congestion avoidance mechanisms, including WRED and tail drop					
	Traffic shaping					
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					
	Hot patches					
	User operation logs					
	Zero Touch Provisioning (ZTP)					
Security and management	802.1x authentication					
	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					
	Ping and traceroute					
	Remote Network Monitoring (RMON)					

Item	CE12804S	CE12808S	CE12804	CE12808	CE12812	CE12816
Dimensions (W x D x H)	442 mm x 620 mm x 352.8 mm (8U)	442 mm x 620 mm x 708.4 mm (16U)	442 mm x 813 mm x 486.15 mm (11 U)	442 mm x 813 mm x 752.85 mm (17 U)	442 mm x 813 mm x 975.1 mm (22 U)	442 mm x 905 mm x 1597.4 mm (36 U)
Chassis weight (empty)	< 60Kg	< 100Kg	< 110 kg	< 150 kg	< 190 kg	< 290 kg
Operating voltage	AC: 90 V to 290 V DC: -38.4 V to -72 V HVDC: 240V					
Maximum power supply	6000W	12000W	5400 W	10800 W	16200 W	27000 W

Ordering Information

Mainframe	
Basic Configuration	
CE-RACK-A01	FR42812 Assembly Rack(800x1200x2000mm)
CE12804S-AC	CE12804S Assembly Chassis(with Fans)
CE12808S-AC	CE12808S Assembly Chassis(with Fans)
CE12804S-DC	CE12804S DC Assembly Chassis(with Fans)
CE12808S-DC	CE12808S DC Assembly Chassis(with Fans)
CE12804-AC	CE12804 AC Assembly Chassis(with CMUs and Fans)
CE12808-AC	CE12808 AC Assembly Chassis(with CMUs and Fans)
CE12812-AC	CE12812 AC Assembly Chassis(with CMUs and Fans)
CE12816-AC	CE12816 AC Assembly Chassis(with CMUs and Fans)
CE12804-DC	CE12804 DC Assembly Chassis(with CMUs and Fans)
CE12808-DC	CE12808 DC Assembly Chassis(with CMUs and Fans)
CE12812-DC	CE12812 DC Assembly Chassis(with CMUs and Fans)
CE12816-DC	CE12816 DC Assembly Chassis(with CMUs and Fans)
Main Processing Unit	
CE-MPU-S	CE12800S Main Processing Unit
CE-MPU	Main Processing Unit
Switch Fabric Unit	
CE-SFU-S	CE12800S Switch Fabric
CE-SFU04	CE12804 Switch Fabric
CE-SFU08	CE12808 Switch Fabric
CE-SFU12	CE12812 Switch Fabric
CE-SFU16	CE12816 Switch Fabric
GE BASE-T Interface Card	

CE-L48GT	48-Port 10/100/1000BASE-T Interface Card(RJ45)
GE BASE-X Interface Card	
CE-L48GS	48-Port 100/1000BASE-X Interface Card(SFP)
10GBASE-T Interface Card	
CE-L48XT	48-port 100M/1000M/10G BASE-T Interface Card(RJ45)
10GBASE-X Interface Card	
CE-L12XS	12-Port 10GBASE-X Interface Card(SFP/SFP+)
CE-L24XS	24-Port 10GBASE-X Interface Card(SFP/SFP+)
CE-L48XS	48-Port 10GBASE-X Interface Card(SFP/SFP+)
40GE Interface Card	
CE-L06LQ	6-Port 40G Interface Card(QSFP+)
CE-L12LQ	12-Port 40G Interface Card(QSFP+)
CE-L24LQ	24-Port 40G Interface Card(QSFP+)
CE-L36LQ	36-Port 40G Interface Card(QSFP+)
100GE Interface Card	
CE-L04CF	4-Port 100G Interface Card(CFP)
CE-L08CC	8-Port 100G Interface Card(CXP)
CE-L12CF	12-Port 100G Interface Card(CFP2)
CE-L36CQ	36-Port 100G Interface Card(QSFP28)
VAS (Value-Added Service) Card	
CE-FWA	CE12800 NGFW Module A
CE-IPSA	CE12800 IPS Module A
Power	
PHD-3000WA	3000W HVDC Power Module
PAC-2700WA	2700W AC Power Supply
PDC-2200WA	2200W DC Power Supply
Software	
CE128-LIC-B	CloudEngine 12800 Basic SW
CE128-LIC-TRILL	TRILL Function License
CE128-LIC-MPLS	MPLS Function License
CE128-LIC-VS	Virtual System Function License
CE128-LIC-IPV6	IPV6 Function License
CE128-LIC-EVN	EVN Function License
CE128-LIC-FCFAL	CloudEngine 12800 FCF All Ports
CE128-LIC-FCF48	CloudEngine 12800 FCF 48 Ports
Document	
CE128-DOC	CloudEngine 12800 Series Switches Product Documentation

Networking and Application

Data Center Applications

On a typical data center network, CE12800/CE8800/CE7800 switches function as core switches, and CE6800/CE5800 switches function as TOR switches. CE6800/CE5800 switches connect to CE12800/CE8800/CE7800 switches through 40GE/10GE ports. The CE12800/CE8800/CE7800 and CE6800/CE5800 switches use the TRILL protocol to build a non-blocking Layer 2 network, which allows large-scale VM migrations and flexible service deployments.

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

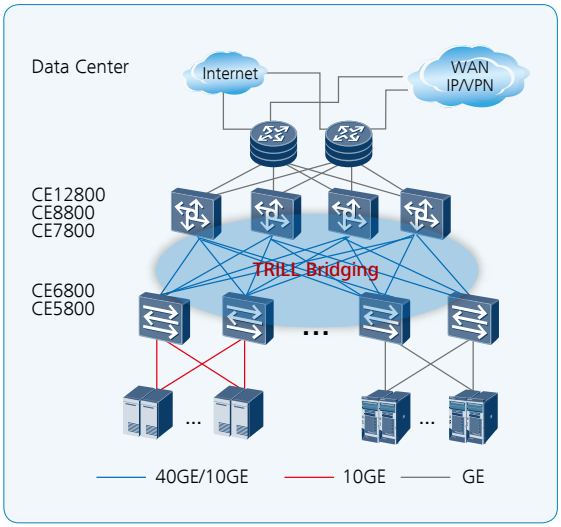

Campus Network Applications

On a typical campus network, two CE12800/CE8800/CE7800 switches are virtualized into a logical core switch using CSS or iStack technology. Multiple 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, CE5800 switches are virtualized with SVF to provide high-density line-speed ports.

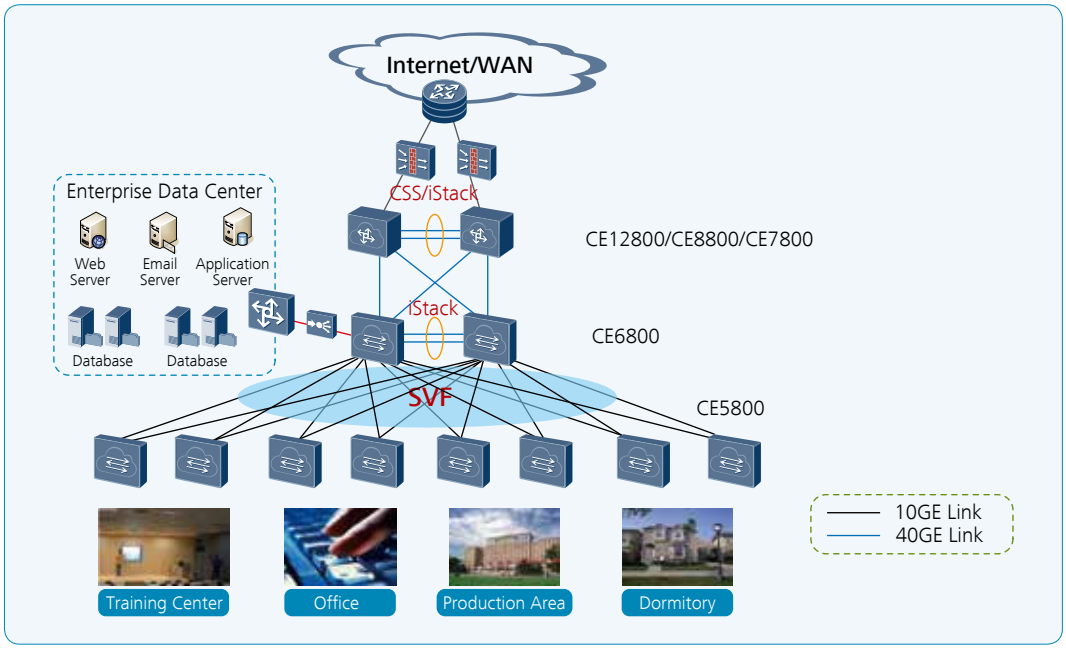

HUAWEI TECHNOLOGIES CO.,LTD.
Huawei Industrial Base
Bantian Longgang
Shenzhen 518129,P.R.China
Tel: +86 755 28780808

www.huawei.com