DONEWIN
Global ICT Solutions Provider

Huawei CloudEngine 16800 Switch Datasheet

Product Overview

Huawei CloudEngine 16800 is the data center switch built for the AI era. It has an embedded AI chip and uses innovative iLossless algorithm to learn and train network-wide traffic, implement zero packet loss and E2E µs-level latency, and achieve the maximum throughput. Using Huawei's VRP8 software platform, 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 100GE, 576 x 40GE, 2,304 x 25GE, or 2,304 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.

VS Implements On-demand Resource Sharing

- High device virtualization capability: The CloudEngine 16800 uses Virtual System (VS) technology to provide an industry-leading virtualization capability that enables one switch to be virtualized into as many as 16 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 CloudEngine 16800 uses industry-leading Cluster Switch System (CSS) technology, which can virtualize two physical switches into one logical switch to facilitate network management and improve reliability.
- The CloudEngine 16800 provides a cluster bandwidth of 3.2 Tbit/s. This super-high bandwidth prevents traffic bottlenecks on data center networks.
- The CloudEngine 16800 combines CSS and VS technologies to turn a network into a resource pool, enabling network resources to be allocated on demand. This on-demand resource allocation is ideal for the cloud-computing service model.

FabricInsight-based Intelligent O&M

- The CloudEngine 16800 provides 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.

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) 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). 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 High Performance

1548 Tbps Switching Capacity

- The CloudEngine 16800 supports GE, 10GE, 40GE, and 100GE 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 servers.

Tbps-Level High-Density Line Cards

- Each line card provides up to 3.6 Tbps forwarding capacity.

- The CloudEngine 16800 supports up to 36 x 100GE, 36 x 40GE, or 144 x 10GE line cards, which provide line-rate forwarding.

- The CloudEngine 16800 provides up to 576 x 100GE, 576 x 40GE, or 2304 x 10GE line-rate ports.

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+N or N+1 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.

High-Performance VRP8 Software Architecture

- The CloudEngine 16800 leverages Huawei's next-generation VRP8, a high-performance, highly reliable modular software platform that provides continuous services.

- Fine-grained distributed architecture: VRP8, a high-end software platform, uses a fine-grained, fully distributed architecture that can process network protocols and services concurrently using multiple instances. This architecture fully leverages multi-core/multi-CPU processes to maximize performance and reliability.

Intelligently Adjusting the ECN Threshold to Implement Non-Blocking Network Transmission

- The CloudEngine 16800 intelligently adjusts the ECN thresholds of different queues based on real-time network traffic changes.

- The CloudEngine 16800 intelligently predicts network congestion and provides quick feedback to reduce retransmission and throughput decrease caused by network packet loss, and to make a compromise between latency and throughput to the maximum degree.

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 1	Multi-cloud Multi-DC Value-added Package
	Basic software	√	√	√		
	IPv6	√	√	√		
	VXLAN	√	√	√		
	Lossless upgrade	√	√	√		
	VS	√	√	√		
	Telemetry	√	√	√		
	NSH		√	√		
	LLETH				√	
	AI ECN				√	
iMaster NCE-Fabric controller	SDN automation	√	√	√		
	Basic intent functions (simulation and verification, and network-wide configuration rollback)			√		
	Multi-cloud multi-DC function					√
iMaster NCE-FabricInsight analyzer	Basic network analysis functions of telemetry	√	√	√		
	Network health ("1-3-5" intelligent O&M)		√	√		
	Value-added functions for			√		

Product	Feature	N1 Mandatory Software Packages			N1 Add-on Software Packages	
	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 Advanced software package.			It is used with the Foundation, Advanced, or Premium software package.	
For details about product function differences, refer to the product documentation.						

Note: V200R005C20, V200R019C10, and later versions support N1 business model.

Note: For detailed information of Huawei CloudFabric N1 business model, visit
https://e.huawei.com/en/material/networking/dcswitch/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 16804	CloudEngine 16808	CloudEngine 16816
Switching capacity (Tbit/s)	43/387[1]	86/774[1]	173/1548[1]
Forwarding rate (Mpps)	11,280	22,560	45,120
Service slots	4	8	16
Switching fabric module slots	6 (scalable to 9 for future expansion)		
Fabric architecture	Clos architecture, cell switching, VoQ		
Airflow design	Strict front-to-back		
Device virtualization	Virtual System (VS)		
	Cluster Switch System (CSS)		
Network virtualization	VXLAN routing and bridging		
	EVPN		
	QinQ access VXLAN		
SDN	iMaster NCE-Fabric		
DCI	BGP-EVPN		
	VXLAN mapping, implementing interconnection between multiple DCI networks at Layer 2		
Network convergence	PFC and ECN		
	RDMA and RoCE (RoCEv1 and RoCEv2)		
Programmability	OPS programming		

Item	CloudEngine 16804	CloudEngine 16808	CloudEngine 16816
	Ansible-based automatic configuration and open-source module release		
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 address entries		
	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 routing protocols, such as RIP, OSPF, IS-IS, and BGP		
	IPv6 routing protocols, such as RIPng, OSPFv3, IS-ISv6, and BGP4+		
	IP packet fragmentation and reassembly		
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	IGMP, PIM-SM, PIM-DM, MSDP, and MBGP		
	IGMP snooping		
	IGMP proxy		
	Fast leaving of multicast member interfaces		
	Multicast traffic suppression		
	Multicast VLAN		
Reliability	Fine-grained microsegmentation isolation		
	Link Aggregation Control Protocol (LACP)		
	M-LAG		
	ESI		
	STP, RSTP, VBST, and MSTP		
	BPDU protection		
	Smart Link and multi-instance		
	Device Link Detection Protocol (DLDP)		
	Hardware-based Bidirectional Forwarding Detection (BFD)		

Item	CloudEngine 16804	CloudEngine 16808	CloudEngine 16816
	VRRP, VRRP load balancing, and hardware-based BFD for VRRP		
	NSF/GR for OSPF/BGP/IS-IS		
	Hardware-based BFD for BGP/IS-IS/OSPF/static route		
	BFD for VXLAN		
	BFD for M-LAG		
NSH	IETF-defined NSH		
RoCE	Scheduling of RoCE traffic and other IP traffic by proportions, ensuring zero packet loss of RoCE traffic		
	RoCE network quality monitoring		
	Adaptive adjustment of traffic thresholds, ensuring that the port bandwidth utilization of RoCE services reaches 90% and no packet loss occurs in various traffic models		
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, PQ+DRR, SP, WRR, SP+WRR, and CBWFQ (or other similar protocols)		
	Congestion avoidance mechanisms, including WRED and tail drop		
	802.1p/DSCP priority marking and re-marking		
	Traffic shaping		
O&M	iPCA		
	Telemetry		
	1588v2		
	ERSPAN+		
	Network-wide path detection		
	Statistics on the buffer microburst status		
	VXLAN OAM: VXLAN ping and VXLAN tracert		
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 patch		
	Passive backplane, plug-and-play components		
	User operation logs		
	Zero Touch Provisioning (ZTP)		
Security and management	RADIUS and HWTACACS authentication for login users		
	Command line authority control based on user levels, preventing unauthorized users from using commands		

Item	CloudEngine 16804	CloudEngine 16808	CloudEngine 16816
	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)		

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

Performance and Scalability

Item	CloudEngine 16800
Maximum number of MAC address entries	256K
Maximum number of routes (FIB IPv4/IPv6)	256K/80K
ARP table size	256K
Maximum number of VRFs	8192
IPv6 ND table size	80K
Maximum number of multicast routes (multicast FIB IPv4/IPv6)	32K/2K
Maximum number of VRRP groups	4096
Maximum number of ECMP paths	128
Maximum number of broadcast domains	32K
Maximum number of BDIF interfaces	32K
Maximum number of VXLAN tunnel endpoints (VTEPs)	2K
Maximum number of LAGs	1024
Maximum number of links in a LAG	128
Maximum number of MSTIs	64
Maximum number of VLANs where VBST can be configured	500

Note: This specification may vary between different scenarios and line cards. Contact Huawei for details.

Hardware Specifications

Item	CloudEngine 16804	CloudEngine 16808	CloudEngine 16816
Dimensions (W x D x H)	483 mm x 985 mm x 438 mm (10 U)	483 mm x 985 mm x 703 mm (16 U)	483 mm x 1144 mm x 1436 mm (32 U)
Chassis weight (empty)	87.8 kg (193.7 lb)	123.9 kg (273.4 lb)	252.8 kg (557.8 lb)
Operating voltage	AC: 176 V to 290 V DC: –40 V to –72 V HVDC: 188 V to 288 V or 260 V to 400 V		
Hot swappable (power modules)	Yes		
Maximum power supply (W)	18,000	30,000	60,000

Item	CloudEngine 16804	CloudEngine 16808	CloudEngine 16816
Operating temperature	0ºC to 40ºC		
Storage temperature	-40ºC to +70ºC		
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	32.55 years
MTTR	1 hour		
MTTF	34.93 years	34.55 years	32.55 years
Availability	0.9999989924	0.9999916778	0.9999913594

Safety and Regulatory Compliance

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

Certification Category	Description
Safety	- EN 60950-1 - EN 60825-1 - EN 60825-2 - UL 60950-1 - CSA-C22.2 No. 60950-1 - IEC 60950-1 - AS/NZS 60950-1 - GB4943
Electromagnetic Compatibility (EMC)	- EN 300386 - EN 55032: CLASS A - EN 55024 - IEC/EN 61000-3-2 - IEC/EN 61000-3-3 - FCC 47CFR Part15 CLASS A - ICES-003: CLASS A - CISPR 32: CLASS A - CISPR 24 - AS/NZS CISPR32 - VCCI- CISPR32: CLASS A - GB9254 CLASS A
Environment	2011/65/EU EN 50581 2012/19/EU 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

Certification Category	Description
	• ETSI EN 300 753 GR63

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 the optical transceivers and cables information, visit
https://e.huawei.com/en/material/networking/dcswitch/f6d91cf16df0474998087676a33fd41e.

Ordering Information

Basic Configuration	
GX-RACK-01	Assembly Cabinet,GX-RACK-01,A812-20 Assembly Rack(800x1200x2000mm)
CE16804-AH	CE16804 AC/HVDC assembly chassis
CE16804-DC	CE16804 DC assembly chassis
CE16804A-B00	CE16804 AC Bundle0 (AC/HVDC Assembly Chassis, 2* MPUD,3*SFU04F,2*3000W AC & HVDC, full Fans)
CE16804A-B01	CE16804 AC Bundle1 (AC/HVDC Assembly Chassis, 2* MPUD,4*SFU04F,4*3000W AC & HVDC, full Fans)
CE16804A-B02	CE16804 AC Bundle2 (AC/HVDC Assembly Chassis, 2* MPUD,5*SFU04G,4*3000W AC & HVDC, full Fans)
CE16808-AH	CE16808 AC/HVDC assembly chassis
CE16808-DC	CE16808 DC assembly chassis
CE16808A-B00	CE16808 AC Bundle0 (AC/HVDC Assembly Chassis, 2* MPUD,3*SFU08F,2*3000W AC & HVDC, full Fans)
CE16808A-B01	CE16808 AC Bundle1 (AC/HVDC Assembly Chassis, 2* MPUD,4*SFU08F,4*3000W AC & HVDC, full Fans)
CE16808A-B02	CE16808 AC Bundle2 (AC/HVDC Assembly Chassis, 2* MPUD,5*SFU08G,4*3000W AC & HVDC, full Fans)
CE16816-AH	CE16816 AC/HVDC assembly chassis

Basic Configuration	
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-HALF2	CE16800 Main Processing Unit D-2 (half-width)
CE-MPUD-FULL2	CE16800 Main Processing Unit D-2 (full-width)

Switch Fabric Unit	
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
CE-SFU16F-G	CE16816 Switch Fabric Unit F
CE-SFU16G-G	CE16816 Switch Fabric Unit G

10GBASE-X Interface Card	
CEL48XSFD-G	48-port 10GE interface card (FD-G, SFP+)

40GE Interface Card	
CEL24LQFD-G	24-port 40GE interface card (FD-G, QSFP+)
CEL36LQFD-G	36-port 40GE interface card (FD-G, QSFP+)

100GE Interface Card	
CEL18CQFD-G	18-port 100GE interface card (FD-G, QSFP28)
CEL36CQFD-G	36-port 100GE interface card (FD-G, QSFP28)

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

Software	
N1-CE168LIC-CFFD	N1-CloudFabric Foundation SW License for CloudEngine 16800
N1-CE168CFFD-SnS1Y	N1-CloudFabric Foundation SW License for CloudEngine 16800-SnS-1 Year
N1-CE168LIC-CFAD	N1-CloudFabric Advanced SW License for CloudEngine 16800
N1-CE168CFAD-SnS1Y	N1-CloudFabric Advanced SW License for CloudEngine 16800-SnS-1 Year
N1-CE168LIC-AFRD-1	N1-CloudEngine 16800 AI Fabric RDMA Application Acceleration Function 1
N1-CE168AFRD1-SnS1Y	N1-CloudEngine 16800 AI Fabric RDMA Application Acceleration Function 1-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-CE-M-LIC-MDCA	N1-CloudEngine Data Center Switch Multi-cloud Multi-DC Value-added Package - Modular
N1-CEMMDCA-SnS1Y	N1-CloudEngine Data Center Switch Multi-cloud Multi-DC Value-added Package, Per Modular device -SnS-Year

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 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.

CloudEngine 16800
CloudEngine 12800
CloudEngine 8800

CloudEngine 8800
CloudEngine 6800
CloudEngine 5800

Internet

WAN IP/VPN

FW

LB

FW

LB

VXLAN

. . .

. . .

. . .

—— 400GE/200GE/100GE/40GE/10GE
—— 50GE/25GE/10GE/GE
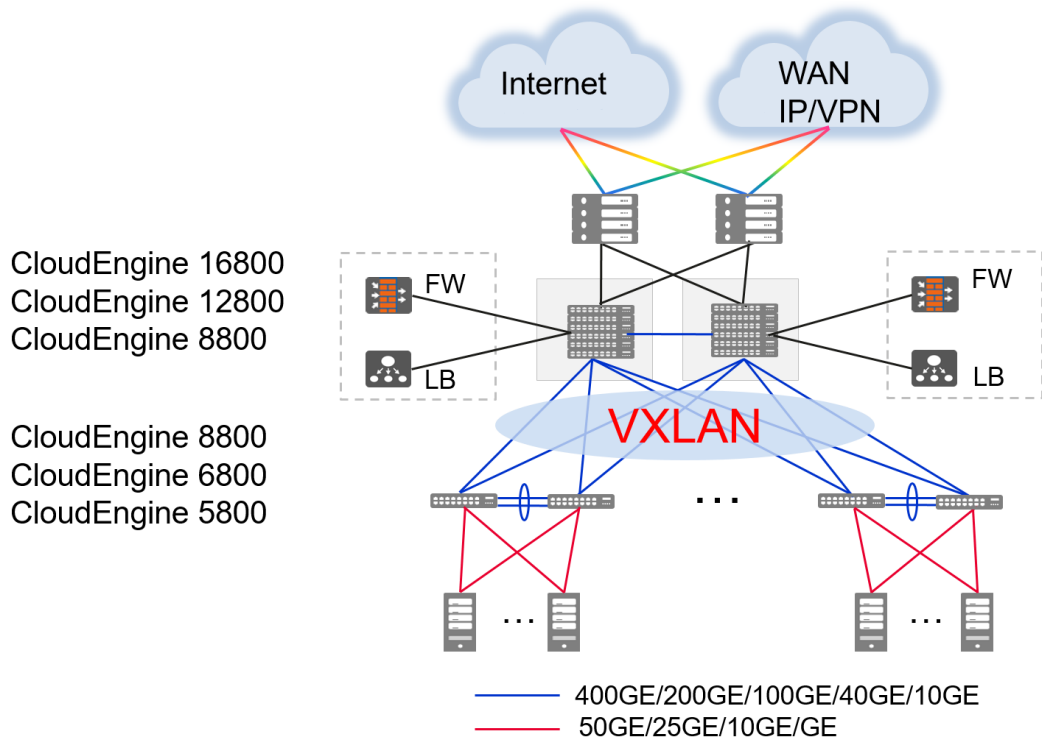
Note: VXLAN can also be used on campus networks to support flexible service deployment in different service areas.

Contact Us

Skype: wendycisco

WhatsAPP: +852-57008326

E-mail: wendy@donewin.com.hk

Website：https://www.uritprice.com

https://www.uritprice.com