

业界首款面向AI时代的数据中心交换机

Industry's First Data Center Switch Built for the AI Era

CloudEngine 16800



CloudEngine

# Making data centers AI-ready with an industry-leading cloud engine

By Leon Wang, General Manager of Huawei Data Center Network Domain

**A**s part of its full-stack, all-scenario AI strategy, Huawei has embedded artificial intelligence into its network products and solutions to create a fully connected, intelligent future. What are the key solutions for realizing this future?

## A cloudy history

On August 8, 2012, Huawei launched CloudEngine 12800, a data center switch built for the cloud computing era. It led the world in design and technological trends for high-density 100G platform data center switches, enjoying the fastest growth for six consecutive years with a CAGR of 82 percent.

On January 9, 2019, Huawei defined three characteristics of data center switches for the AI era: embedded AI chips, 48 x 400GE high-density ports per slot, and the capability to evolve to an Autonomous Driving Network. Huawei also unveiled the industry's first data center switch built for the AI era, CloudEngine 16800, once again setting a new benchmark for the industry.

## Data centers face challenges with AI

Driven by AI, the fourth industrial revolution is leading us into a new era where everything senses, everything is connected, and everything is intelligent. According to Huawei's *Global Industry Vision (GIV)*, the amount of global data

will rise to 180 ZB by 2025. Moreover, 95 percent of unstructured data, such as voice and video, will depend on AI processing. And enterprises will begin harnessing AI for decision-making, reshaping business models and ecosystems, and rebuilding the customer experience, with 86 percent of organizations having adopted AI.

However, while the evolution of data centers from the cloud era to the AI era is inevitable, current data centers face three major challenges:

A packet loss of 0.1 percent on traditional Ethernet limits AI computing power to 50 percent: To boost AI's operating efficiency, AI systems will use flash storage to reduce latency by

more than 100 times. They'll also use GPUs and even dedicated AI chips for computing to increase data processing ability again by the same amount. In this context, network communication latency has emerged as a critical weakness. AI computing power affected by the performance of data center networks is now a worrying bottleneck for the commercial application of AI. High-performance data center clusters are extremely sensitive to network packet loss.

Existing 100GE networks will be unable to handle the data flood over the next five years: The amount of global data is predicted to surge from 10 ZB in 2018 to 180 ZB in 2025. Existing 100GE-based data center networks will be unable to support this data flood. New services, such as enterprise AI, are driving the evolution of data center servers from 10G to 25G and even 100G, which necessitates switches that can support 400G interfaces.

With the deep integration of computing and storage networks, manually locating network problems takes several hours. In recent years, data center architecture has changed dramatically, with the number of servers per data center increasing from dozens to tens of thousands. Moreover, computing, storage, and data networks are converging, and the amount of analyzed traffic has increased many thousands of times.

Locating service faults takes several

hours using traditional manual troubleshooting O&M methods, which is no longer viable.

## Data center switches in the AI era

To tackle these challenges, data centers will require autonomous high-performance networks to improve AI computing power and help customers speed up AI service operations. Therefore, Huawei's three characteristics for data center switches in the AI era responds to this.

### The industry's first data center switch with an embedded AI chip for 100 percent AI computing power

CloudEngine 16800 is the first data center switch in the industry to harness the power of an embedded high-performance AI chip. It uses the iLossless algorithm for auto-sensing and auto-optimization on the traffic model, thereby realizing lower latency and higher throughput based on zero packet loss. CloudEngine 16800 overcomes computing power limitations caused by packet loss on traditional Ethernet, boosting AI computing power from 50 percent to 100 percent and improving data storage IOPS by 30 percent.

### The industry's highest density 48 x 400GE ports per slot, meeting requirements for traffic growth in the AI era

CloudEngine 16800 boasts an upgraded hardware switching platform. Its orthogonal architecture solves multiple technical challenges, including high-speed signal transmission, heat dissipation, and efficient power supply. These advantages enable it to provide the industry's highest density 48-port 400GE line card per slot and largest 768-port 400GE switching capacity (five times the industry average), meeting traffic multiplication requirements in the AI era. In addition, its power consumption per bit is reduced by 50 percent.

### Enables autonomous driving network, identifies faults in seconds, and automatically locates faults in minutes

The CloudEngine 16800 is embedded with an AI chip, substantially enhancing the intelligence of devices deployed at the network edge and enabling the switch to implement local inference and rapid decision-making in real time. With CloudEngine 16800's local intelligence and the centralized network analyzer FabricInsight, the distributed AI O&M architecture identifies faults in seconds and automatically locates them in minutes, helping to accelerate evolution to autonomous driving networks. Additionally, it provides root cause analysis of more than 72 types of typical faults in seconds using the iNetOps smart O&M algorithm, boosting the automatic fault location

rate to 90 percent. Furthermore, the distributed AI O&M architecture dramatically enhances the flexibility and deployability of O&M systems.

### 3 major breakthroughs for the hardware exchange platform

CloudEngine 16800 supports the high-speed smooth evolution of high-density ports from 10GE to 40GE, 100GE, 400GE, and even 800GE. It slashes the number of core layer devices, simplifies the network, and improves management efficiency. CloudEngine 16800 delivers revolutionary technological breakthroughs in three areas:

#### SuperFast: Ultra-high-speed interconnection

When evolving from 100GE to high-density 400GE, the first challenge is implementing the high-speed intra-switch signal transmission capability. Each time the signal frequency doubles, the signal attenuation of the PCB increases by more than 20 percent. As traditional PCBs are made from copper foil using traditional manufacturing techniques, transmission loss and high-frequency interference become more severe when the signal transmission rate increases. This is the main bottleneck limiting the switch capacity of switches. Huawei employs techniques like sub-micron lossless materials and polymer bonding to improve

signal transmission efficiency by 30 percent, thus supporting full-lifecycle compatibility and evolution from 100GE to 400GE and even higher port speeds.

#### SuperPower: Efficient power supply

Based on a traditional design, a high-density 400GE interface core switch like CloudEngine 16800 would require 40 power modules, which would take up over one-third of the entire chassis alone. Huawei has developed the industry's first power module with independent dual inputs and intelligent switching. It utilizes magnetic blowout and large exciter technology to realize fast switching in milliseconds and ensure high reliability. As such, 21 of these new power modules can achieve the same power supply capability and reliability as 40 single-input regular power modules, using 50 percent less space. Line cards use a magnetic matrix and high-frequency magnetic technologies to provide 1600 W power supply capabilities in a space the size of two thumbs, improving power supply efficiency in the space of a single unit by 90 percent.

#### SuperCooling: Powerful heat dissipation

For an ultra-high-density switch, heat dissipation is an important reflection of the engineering capability of the entire system. The CloudEngine 16800 switch's cooling system provides

both card-level and system-level heat dissipation for true energy efficiency.

When it comes to card-level heat dissipation design, evenly exporting the chip-generated heat out of the card and dissipating it is key. CloudEngine 16800 leverages a unique carbon nanotube thermal pad and VC phase-change radiator technology for 4x better cooling capability than the industry average, improving the entire system's reliability by 20 percent.

In terms of system-level cooling, Huawei uses mixed-flow fans, an industry first, to achieve the best heat dissipation efficiency of an entire system in the industry. The average power consumption of each bit of data is 50 percent lower than the industry average, producing savings equivalent to 320,000 kWh and reducing carbon emissions by more than 250 tons per year per switch. The unique magnetic permeability motor and the mute deflector ring reduce noise by 6 dB, making the data center quieter.

Equipped with a high-performance AI chip and featuring the industry's highest switching capacity, CloudEngine 16800 will enable the switch over from cloud-era to AI-era data center switches, lead data centers into the AI era, and help customers succeed in the new AI future. 