

CloudEngine Is the Foundation of the Intent-driven Network

Huawei CloudEngine Series Switches Technical Presentation



Contents

1

CloudEngine Switch Overview


2

CloudEngine Switch Highlights

3

CloudEngine Switch Market Progress

Switches Are the Cornerstone for Transforming Data Centers from Service Centers to Value Centers



Cloud computing



Big Data



Distributed storage AI



SDN

NFV

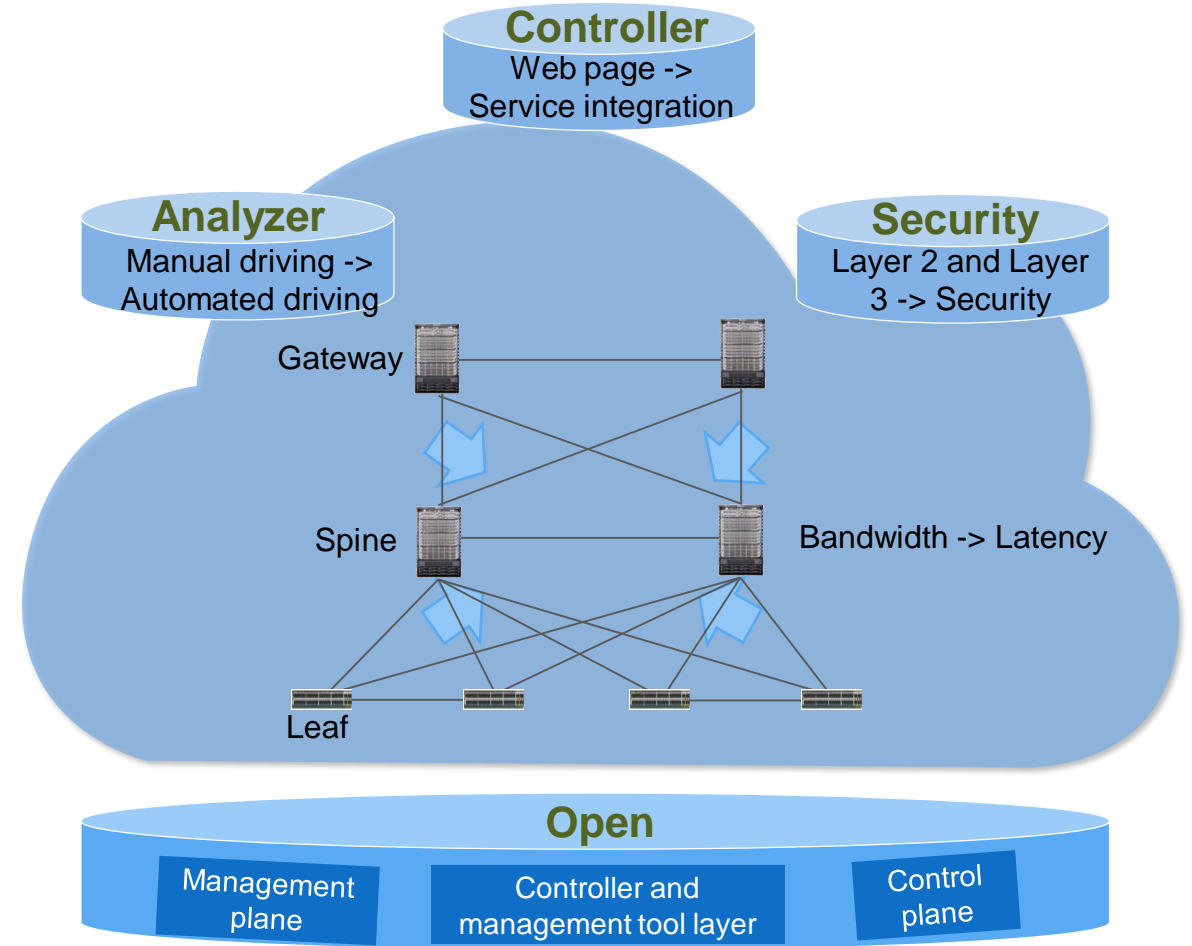
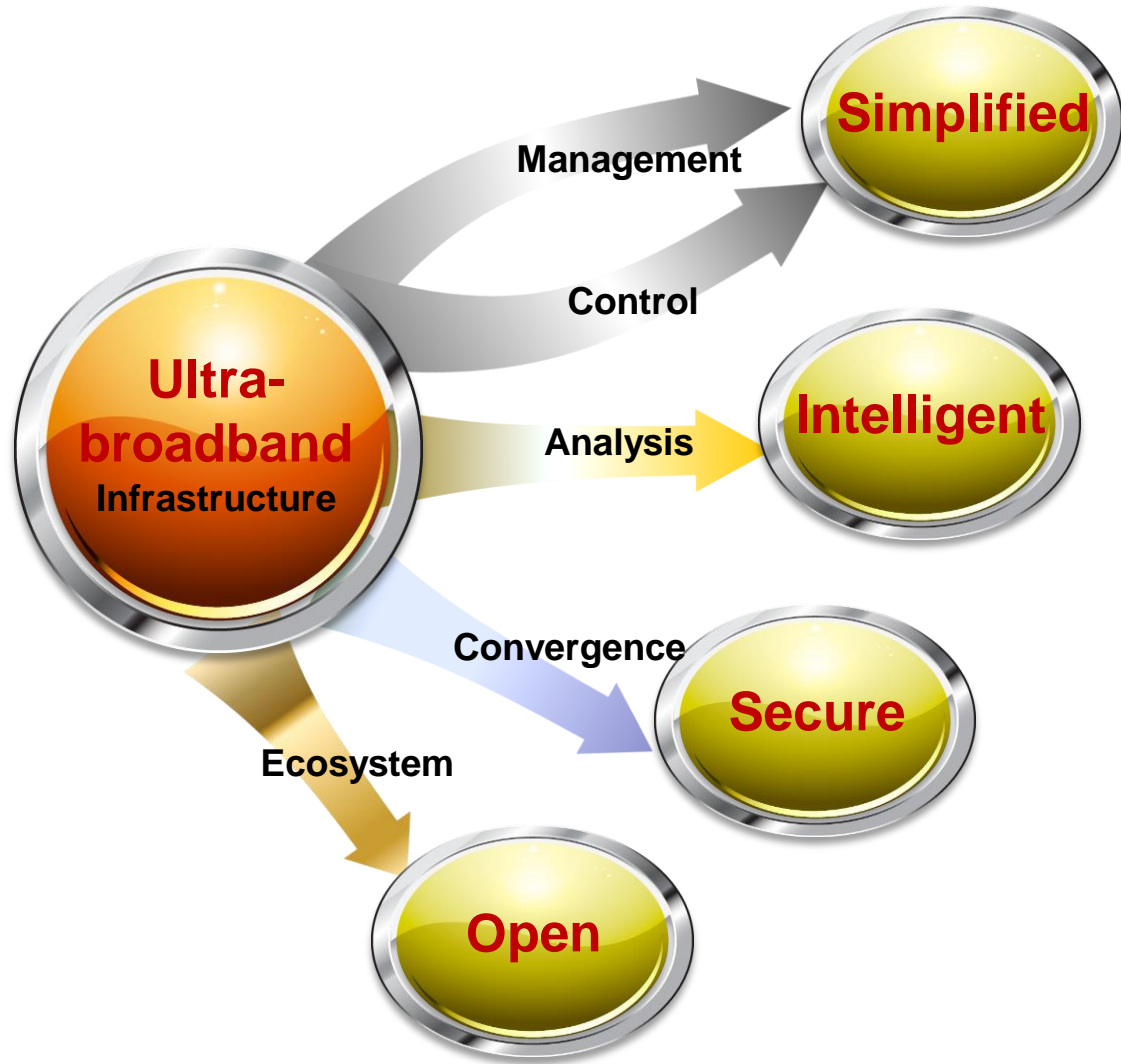
Intent-driven network

Metcalfe's Law: The effect of a telecommunications network is proportional to the square of the number of connected users of the system.

-Robert Metcalfe who invented Ethernet, founder of 3Com

In the future, even if all hardware network devices will disappear, data center switches used as the buses connecting to servers, will always exist. The Ethernet helps release the value of data.

Characteristics of the Future of DC Switches: Ultra-broadband, Simplified, Intelligent, Secure, and Open




CloudEngine Series Data Center Switches Portfolio(1)

Core Switches

CloudEngine 16800 (New)



CloudEngine 16816



CloudEngine 16808



CloudEngine 16804

Access Switches

10GE ToR switch



CloudEngine 6881-48S6CQ (New)


25GE ToR switch




CloudEngine 6863-48S6CQ (New)

CloudEngine Series Data Center Switches Portfolio(2)

Core Switches




Virtual Switches




Access Switches

ToR switch with flexible cards




100GE switch



40GE switch



10GE ToR switch



25GE ToR switch



10GE large-buffer ToR switch




GE ToR switch



High-Quality CloudEngine: High-Quality Architecture Creates a Green and Stable Network

Orthogonal architecture




No cabling of the backplane

Increased bandwidth of the entire system

Strict front-to-back airflow design

Patent No.: CN201110339954.1


CE12800



Independent front-to-back airflow

Even heat dissipation, suitable for data centers

Non-blocking switching



Cell switching and VoQ

Traffic balancing, improving bandwidth utilization

Counter-rotating fans

Leading energy-conserving design



Counter-rotating and turbo fans

Highly efficient heat dissipation

- Industry-leading architecture design and high quality: orthogonal SFU design, Clos architecture, cell switching, and Virtual Output Queue (VOQ) mechanism
- Innovative energy-saving technologies and green pioneer: strict front-to-back airflow and counter-rotating fans

Contents

1 CloudEngine Switch Overview

2 CloudEngine Switch Highlights

- **Ultra-broadband Cloud Engine**
- **Simplified Cloud Engine**
- **Intelligent Cloud Engine**
- **Secure Cloud Engine**
- **Open Cloud Engine**

3 CloudEngine Switch Market Progress

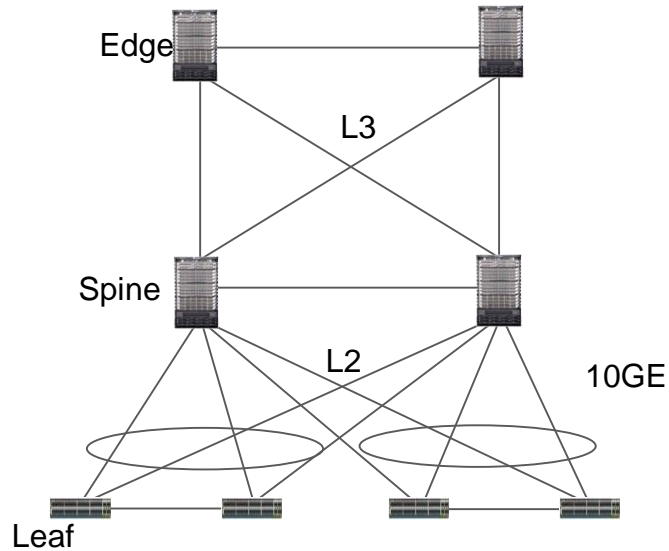
Clos Theory: Cluster Scale Is the Driving Force for Data Center Network Architecture Evolution

Difficulties in the non-blocking Clos architecture: The convergence ratio and packet loss ratio cannot be compromised.

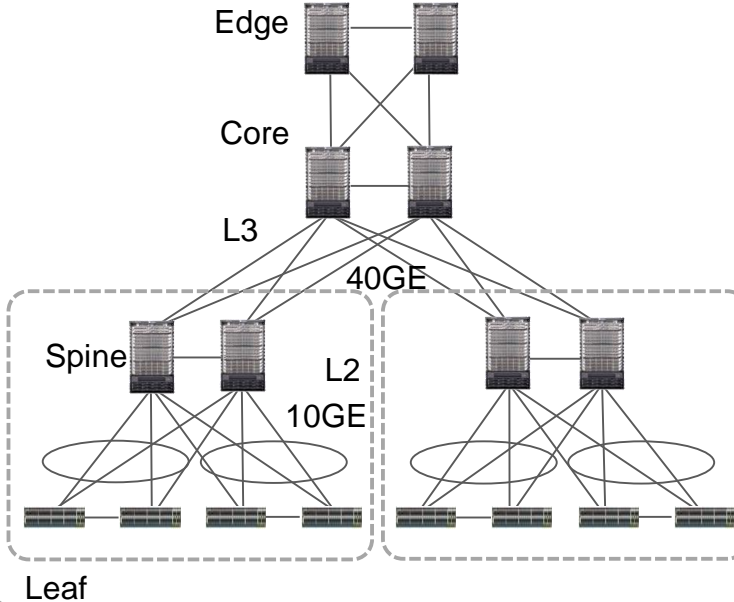
- Data center network architecture: A fat-tree topology is used and the capacity of the root node determines the server cluster scale.
- Evolution direction: Add network layers, and increase the quantity and capacity of spine or core nodes.
- Network congestion control: Increase the buffer and optimize load balancing.

- The port capacity of cards increases continuously, and CE switches can provide 36*100GE ports.
- To avoid HASH polarization, CE switches provide 128 ECMP paths.
- Network congestion control: CE switches provide large buffer, and split a single flow into multiple ones to load balance them.

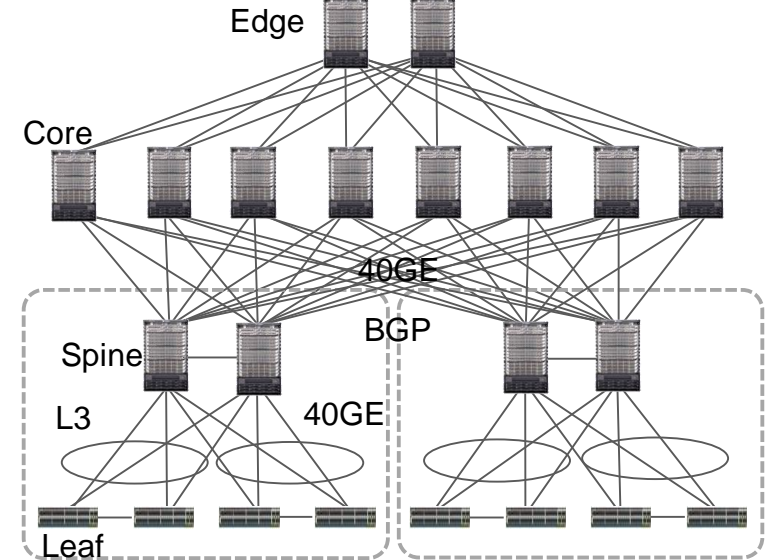
First generation: 3K GE servers



Second generation: 10K GE servers




Third generation: 20K 10GE servers



CE switches' buffer is 80 times higher than the industry average, implementing zero packet loss for microburst traffic. The switches' performance is 1032 Tbit/s. CE switches can connect to over 50,000 servers with no blocking.

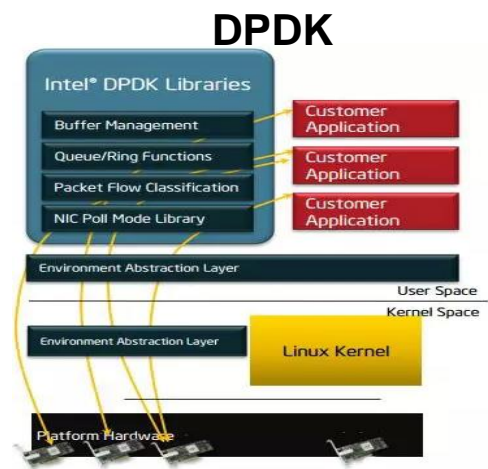

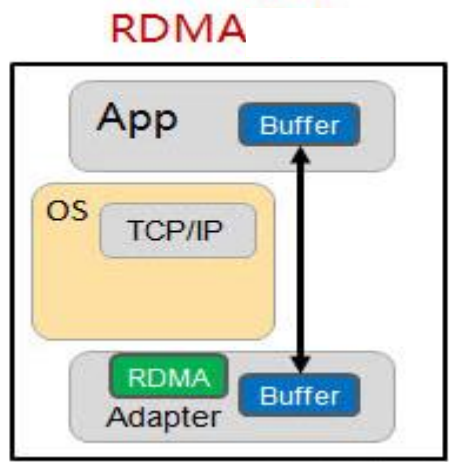
Larger Interface Rate: The Rise of 25GE Interfaces Balance the Cost and Efficiency

Development of Ethernet:

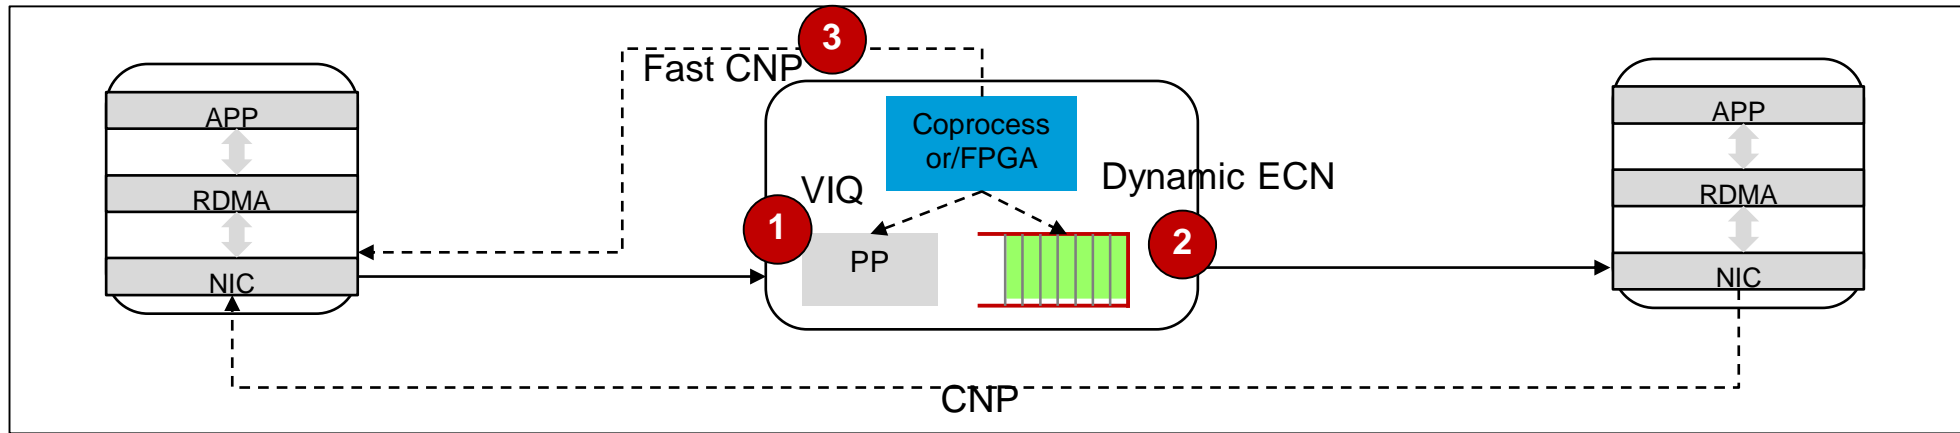


In the past two years, why are 25GE interfaces used?

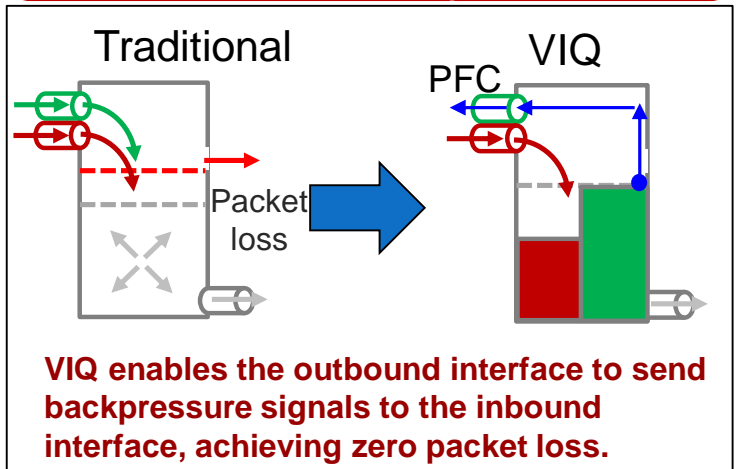
- The 25GE interface can better match the SerDes rate: 1.25 Gbit/s -> 3.125 Gbit/s -> 6.25 Gbit/s -> 10.3125 Gbit/s -> 25 Gbit/s -> 56 Gbit/s
- Compared with the 40GE NIC, the 25GE NIC has higher use efficiency of the PCIe channel. $(40G+40G)/8G*16= 62.5\%$; $25G*2/(8G*8) = 78\%$
- Lower cabling costs for 25GE interfaces: The SFP28 module is used. Because only single-channel connections are used, the SFP28 module is compatible with LC optical fibers in the 10GE era, without cabling.
- The bandwidth between NICs has exceeded 10 Gbit/s: As technologies such as RDMA, SR-IOV, and DPDK develop, the bandwidth between NICs has exceeded 10 Gbit/s.



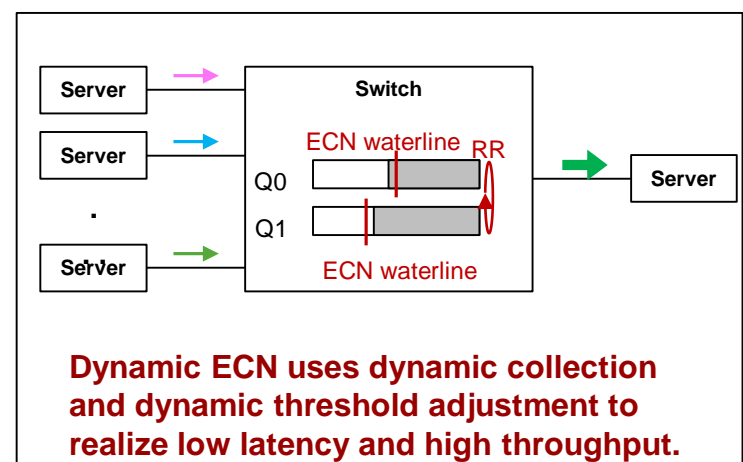
AI Fabric: Intelligent Lossless Data Center Network Solution Provides Low latency and Zero Packet Loss



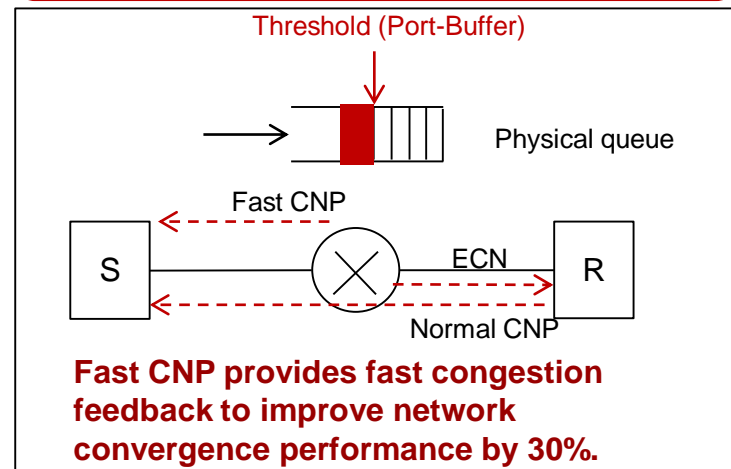
VIQ: eliminates packet loss inside chips



Dynamic ECN




Fast CNP



Device Virtualization: Easy-to-Manage, High-Performing, Highly Reliable Virtual Systems


Cluster Switch System (CSS)



Dual-active detection (DAD) channel
Unique three-channel separated cluster, control plane coupling

- Independent forwarding, control, and detection, and 3-channel cluster design
- Four dedicated GE interfaces are used as cluster control channels.
- A maximum of 3.2 Tbit/s cluster bandwidth is supported.

Multi-Chassis LAG (M-LAG)



Independent control plane, protocol-level coupling

- The control plane runs independently and synchronizes a small amount of information about interface status entries.
- Devices in the DFS group can be upgraded independently, without interrupting services.
- When the peer-link is faulty or the M-LAG master device fails twice, the M-LAG backup device can still work properly.


M-LAG Lite



Independent control plane without coupling

- The control plane runs independently and does not have synchronization information.
- Two switches are configured with the same gateway IP address and MAC address.
- Two links of the server NIC are configured to send broadcast packets simultaneously.

VS (Virtual System)



On-demand VS allocation, improving resource utilization

- Provide a maximum of **1:16** virtualization capabilities in port and port group mode.

Exclusive resources in VSs and highest specifications

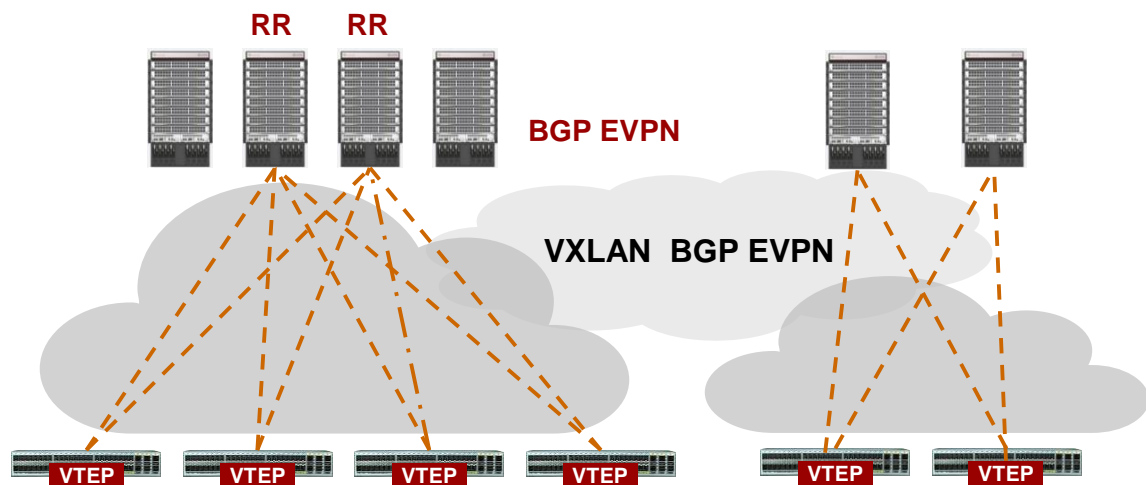
- Exclusive CPU, memory, and MAC/VLAN/FIB entries

Fault isolation between VSs, improving security

- VSs are isolated from each other and cannot communicate with each other. VSs run different protocols without affecting each other.

Layer 2 Boundary Extension: Build a Large-Scale Network Resource Pool Based on BGP EVPN

Layer 2 large-scale horizontal expansion in the data center and extension to the remote DC



BGP EVPN acts as the VXLAN control plane to provide the following functions:

- Triggers automatic VXLAN tunnel setup between VTEPs to avoid the need to manually configure full-mesh tunnels.
- Advertises host routes and MAC address table, prevents unknown traffic flooding, and optimizes packet forwarding.
- Implements Layer 2 interconnection between data centers in different networking.

Protocol vitality: open interconnection and interworking between devices from different vendors


At Interop Tokyo 2016, Huawei CE8860 and CE6851 data center switches received the Best of Show Award Grand Prix in the ShowNet Product category.

ShowNet is a comprehensive demonstration of the Interop and has always been recognized as an important occasion to verify interoperability between multi-vendor equipment. Exhibitors from all over the world demonstrated their latest technologies and products, and tested their solutions in a real network environment. Huawei showcased the open SDN architecture composed of CE8860 and CE6800 at the exhibition, and verified the Virtual Extensible LAN (VXLAN) and Ethernet VPN (EVPN) features of the switches, which can be used to establish SDN networks within and between data centers. During the demonstration, Huawei switches also interoperated with products from other vendors to establish EVPN networks successfully.

グランプリ



華為技術日本 (株) (ファーウェイ・ジャパン)
CE8860
CE6851
NE40E




Network Automation: Interconnection with Third-Party Management Tools, Controllers, Virtualization Management Platforms, and Cloud Platforms

Scenario 1
Interconnection
with a management
tool

Scenario 2
Interconnection with a
virtualization
management
platform

Scenario 3
Interconnection
with a third-party
controller

Scenario 4
Interconnection with
a cloud platform



Scenario 1: traditional network management

- **Interconnection with a third-party management tool:** CE switches can interconnect with a third-party management tool such as Ansible to implement automatic network configuration.

Scenario 2: network and computing association

- **Interconnection with a virtualization management platform:** CE switches are connected to the Agile Controller-DCN, and the Agile Controller-DCN is associated with the third-party computing management platform.

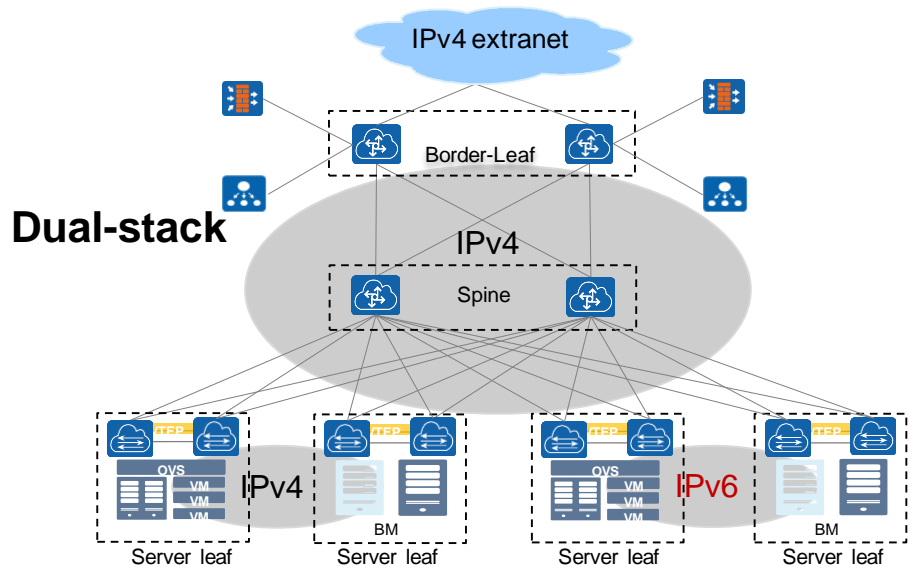
Scenario 3: third-party management on the overlay

- **Interconnection with a third-party controller:** The CE switch functions as the VXLAN Layer 2 VTEP and is managed by the NSX.

Scenario 4: cloud-network integration

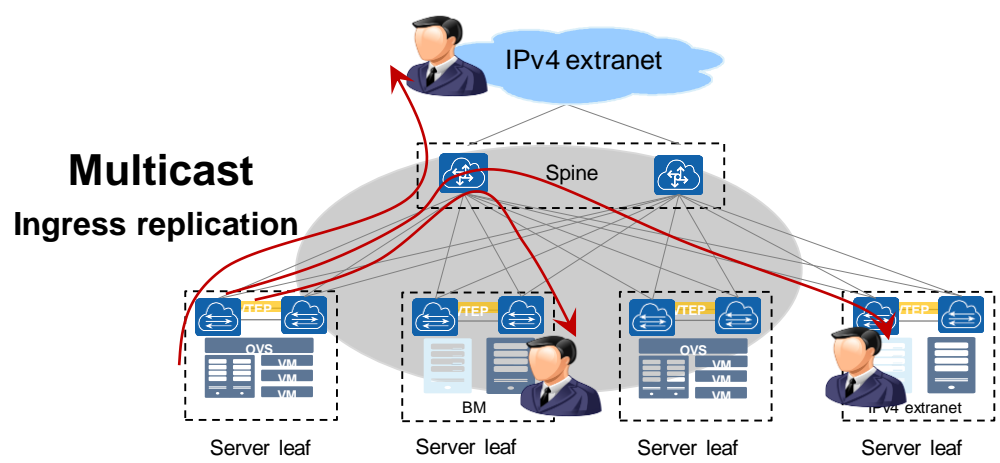
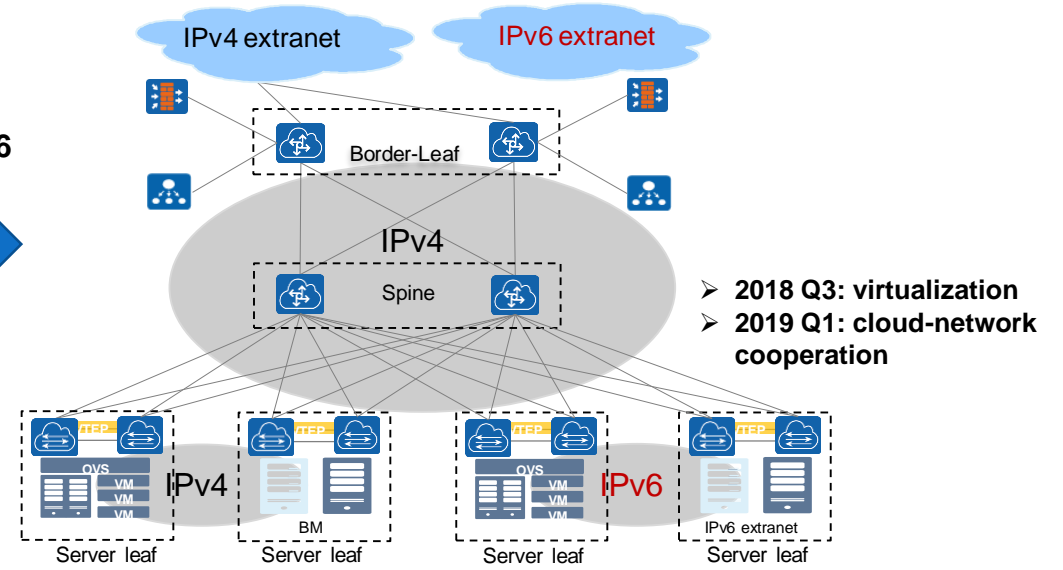
- **Interconnection with a cloud platform through the Agile Controller-DCN:** CE switches are connected to the Agile Controller-DCN, and the Agile Controller-DCN connects to the third-party cloud platform.

Simplified Deployment: IPv4 and IPv6, Unicast and Overlay Multicast, and Rollout of Full-stack Services Within Minutes



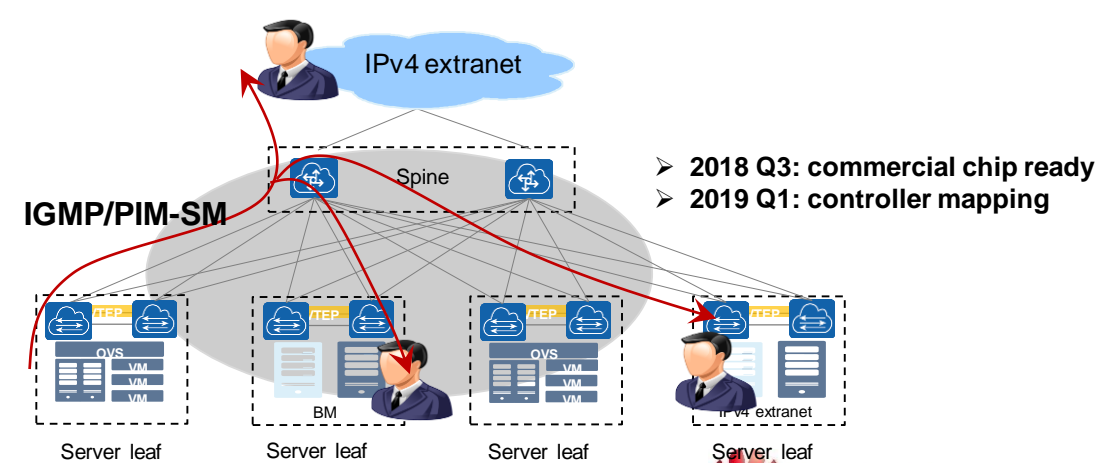
Service-centered IPv6 evolution mode

Replicate IPv4 O&M experiences

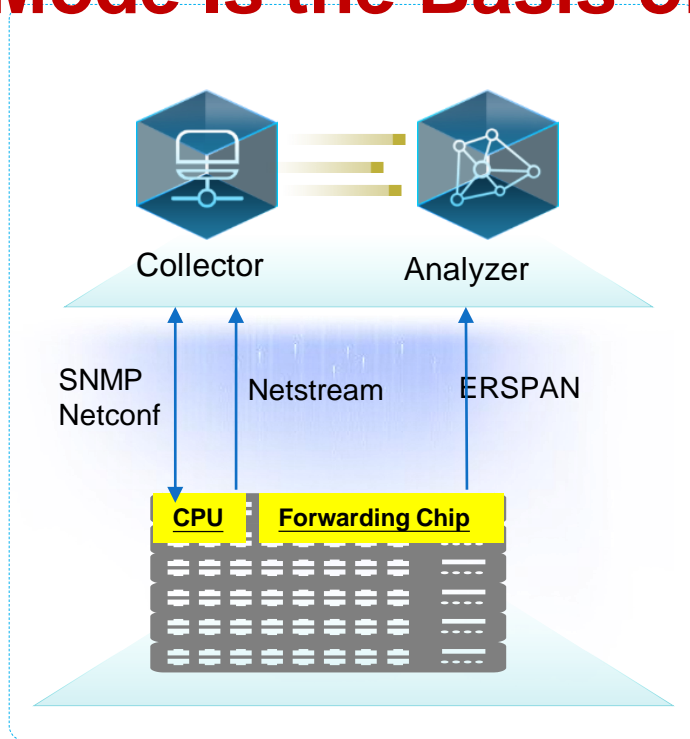


Overlay multicast

Save bandwidth

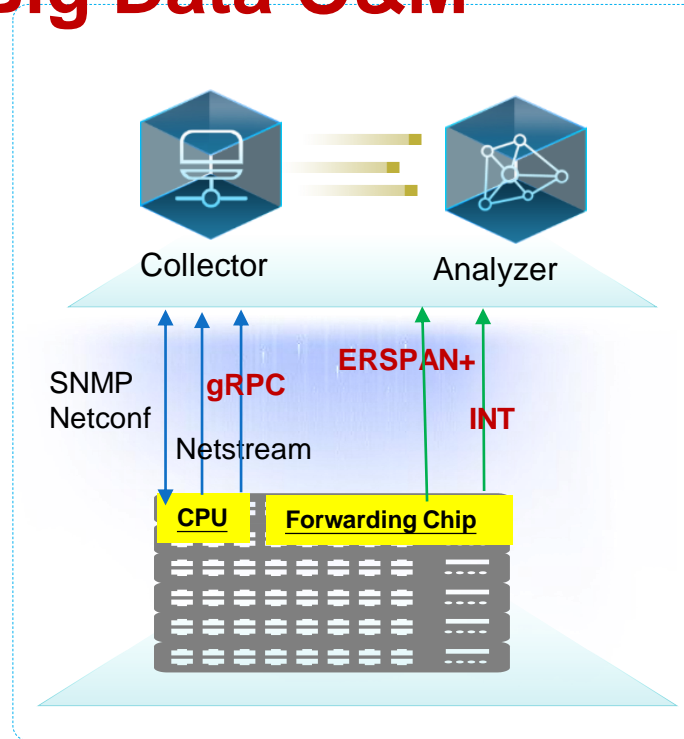


Telemetry Capability: Transformation of the Data Collection Mode Is the Basis of Big Data O&M



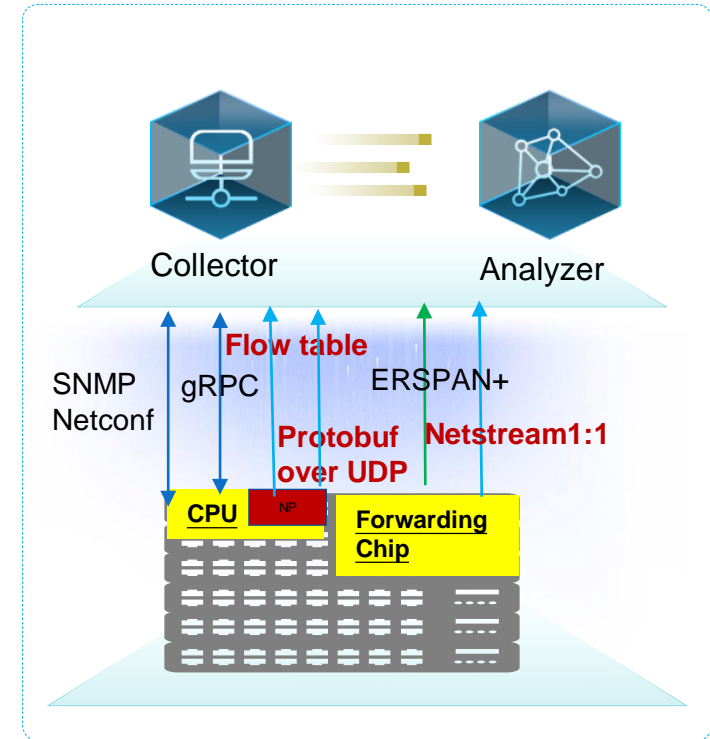
Historical Capabilities

- SNMP or NETCONF uses the query/response mechanism, minute-level reporting, and XML or text encoding, which is inefficient.
- NetStream uses the flow sampling mechanism and requires CPU participation, which has low performance and is inaccurate.



Current Capabilities

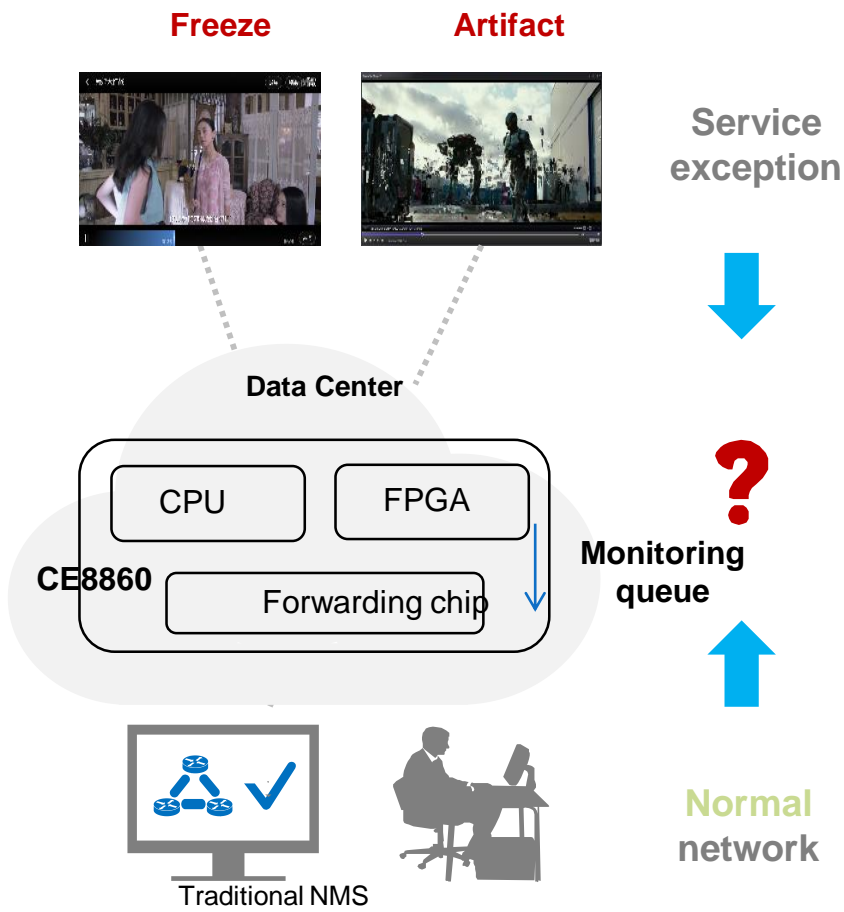
- gRPC uses the subscription/reporting mechanism, **subsecond-level reporting**, protobuf coding, and HTTP transmission, **which has a high efficiency**.
- ERSPAN+ adds **ingress and egress ports or timestamps** of original flows to calculate the **flow path and delay**.
- INT supports in-line path or quality detection.



Future Evolution

- Protobuf over UDP is used to encode and transmit forwarding plane information, **which is efficient and does not affect CPU performance**.
- Small NP intelligent analysis algorithm is used to perform in-depth analysis of abnormal flows to learn in-depth information such as the **latency, jitter, packet loss ratio, and packet loss location**.

Microburst Detection Capability: Millisecond-level Buffer Monitoring and Subscription Collection, Which Are Visible and Clear



- The collection period is too long, which may ignore network details.

SNMP request and response

...

5-minute polling period Multiple requests for a single task

- The detection interval is too long, so device details may be incomplete.

100%

50%

5 10 → s

Microsecond-level buffer monitoring

- Visible:** subsecond-level subscription data collection

gRPC subscription

Content feedback

Subsecond-level collection One request for multiple tasks

- Clear:** high-precision data monitoring

100%

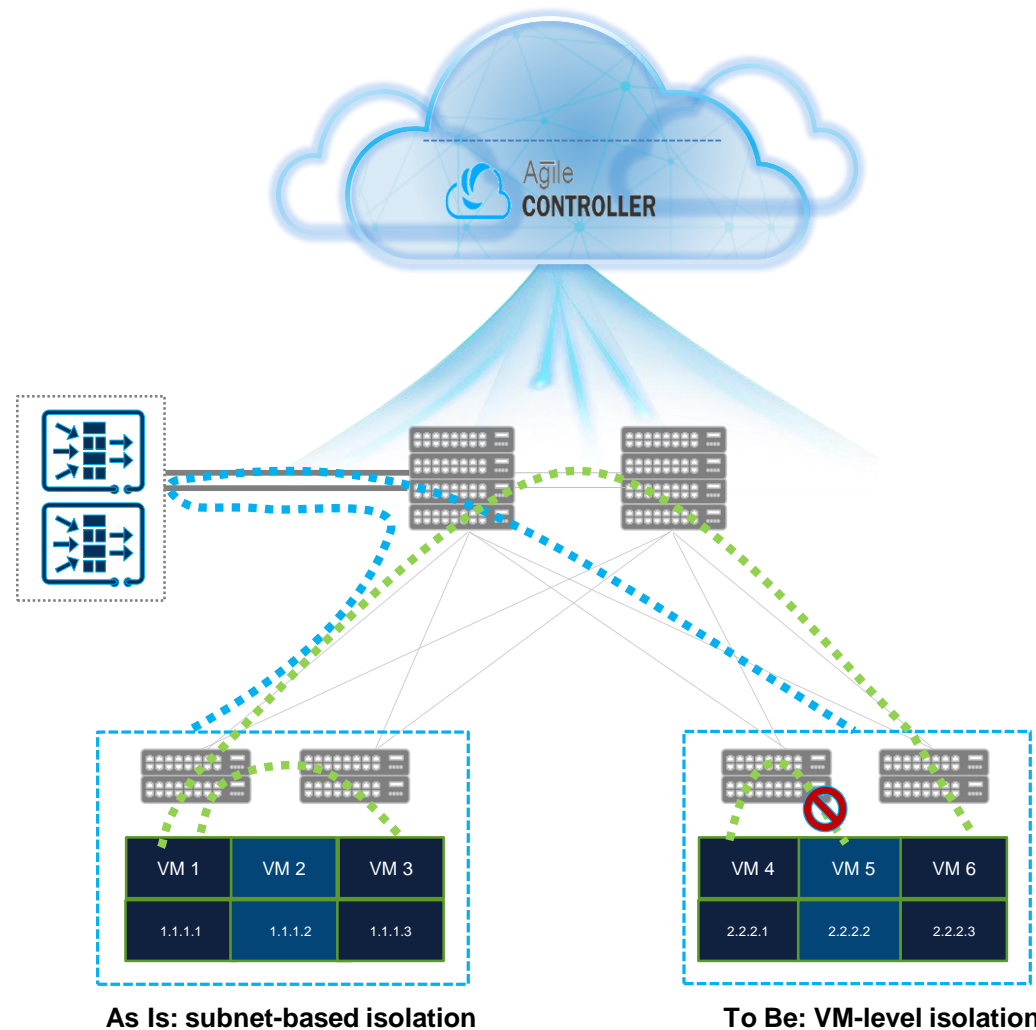
50%

5 10 → us

2-ms buffer monitoring

Note: The CE8860 supports this function.

Use Microsegmentation to Achieve Fine-grained Isolation and Service Security



Fine-grained Defense

Defining applications based on VM names and discrete IP addresses, with finer granularity and wider dimensions

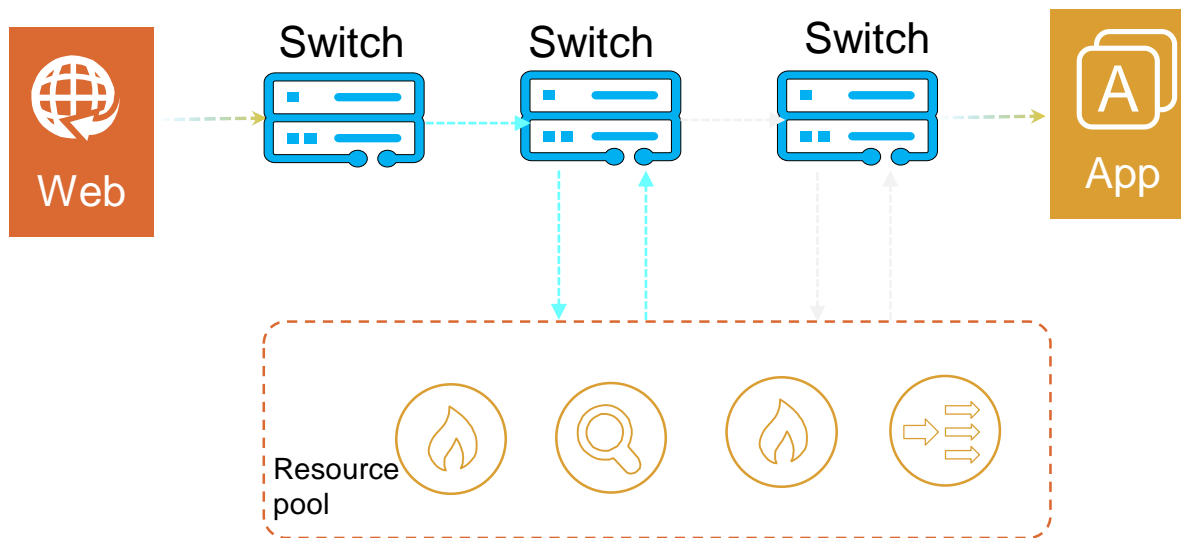
Distributed Security

Traffic of access switches is filtered nearby and east-west isolation is implemented without using firewalls.

Flexible Deployment

Defining services based on application groups and decoupling from subnets to achieve flexible deployment

Agile CloudEngine: Supporting NSH Service Chains, Providing Easier VAS Orchestration



Simplified deployment

The SDN controller defines service chains through drag-and-drop operations.

Flexible orchestration


Decouple the VAS function from Fabric, providing flexible orchestration.

Efficient forwarding

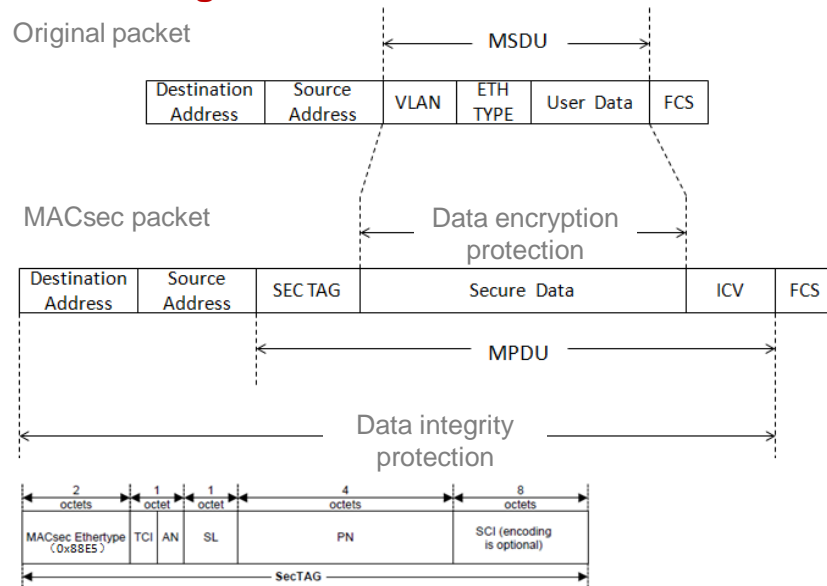
Provide traffic diversion for one time, simple configuration, service traffic forwarding, and secure monitoring.

MacSec at the Link Layer: IP Layer 3 Features such as Encryption Are Introduced to the MAC Link Layer

Networking



Networking



Scenario

- In scenarios that require high data confidentiality, such as government, military, and finance scenarios, interconnection is required between data centers or between different modules of data centers across buildings.
- The CE6875 uplink port (100GE), and CEL16CQFD (16*100GE) and CEL08CFFG1 (8*200GE) cards of the CE12800 can be used.

Definition

- Media Access Control Security (MACsec) ensures secure communication within LANs in compliance with IEEE 802.1AE and 802.1X. It provides identity authentication, data encryption, integrity check, and replay protection to protect Ethernet frames and prevent devices from processing attack packets.

Open Ecosystem: Huawei Joins Hands with 20+ Industry Chain Partners to Perform System Integration

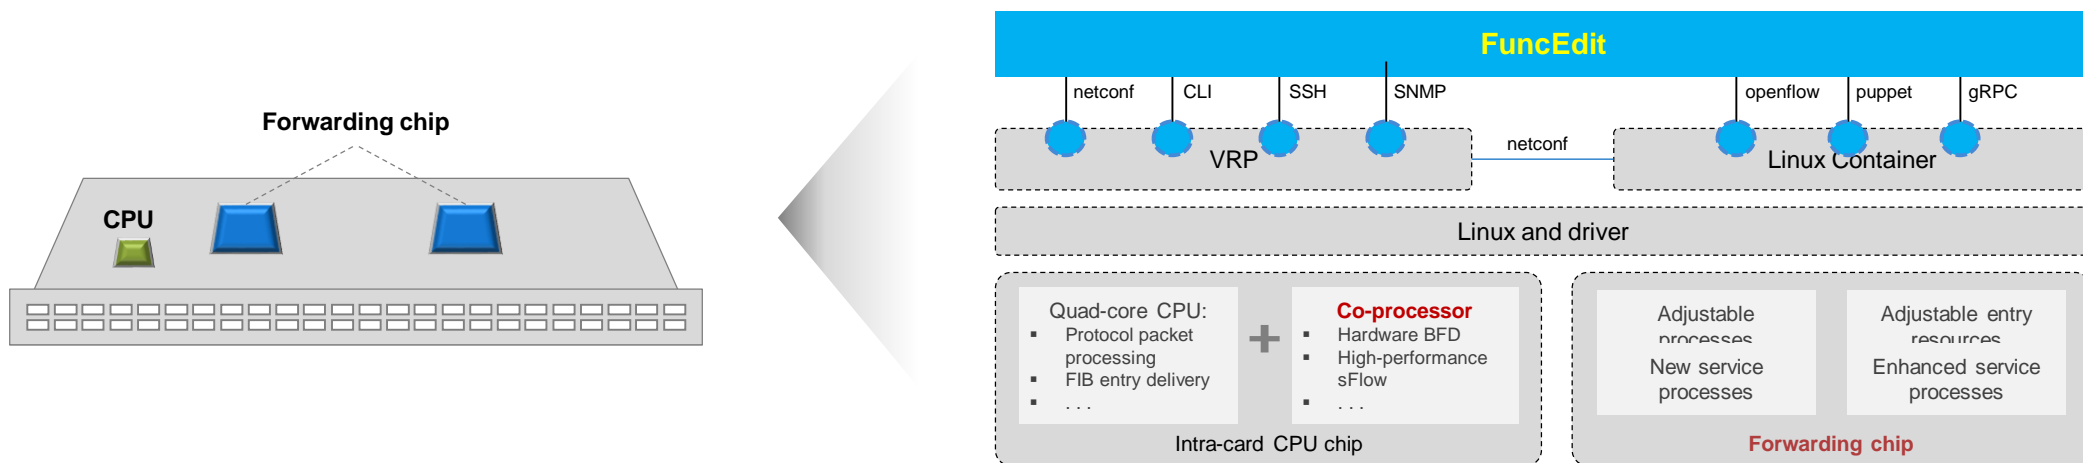
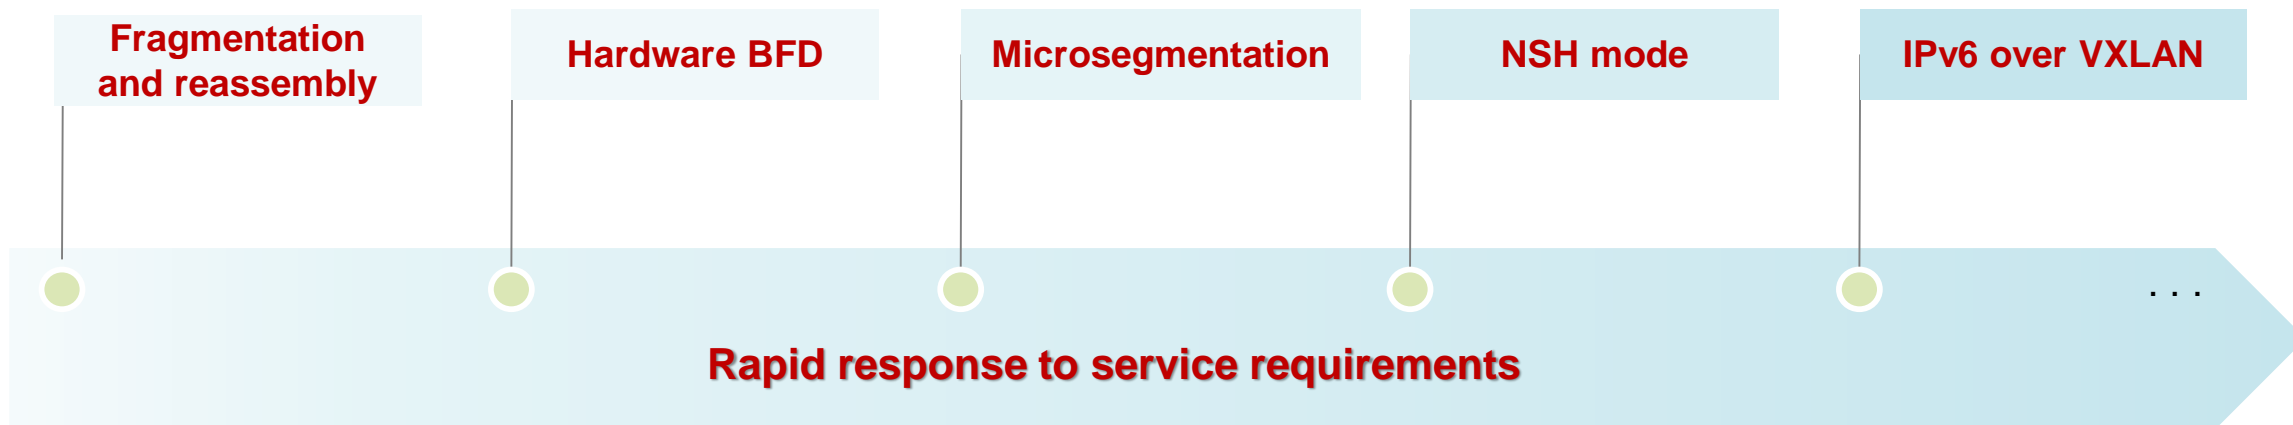
Open ecosystem: fast integration and simplified management



System integration: 10+ OpenLabs in the globe



Open architecture, Flexible Business Innovation



CloudEngine High-Performance Cloud Switches

Ultra-broadband: higher interface rate, more even load balancing, larger buffer, and lower latency



- Higher interface rate: 25GE interfaces and larger buffer cope with traffic surge in N:1 scenarios.
- Flowlet&DLB: One flow is load balanced among multiple links.
- AI Fabric intelligent lossless data center network solution: low latency and zero packet loss

Simplified: automatic deployment of full-stack services and service rollout within minutes



- Multiple virtualization technologies: CSS, M-LAG, M-LAG Lite, and VS
- VXLAN + BGP EVPN: intra-DC and inter-DC virtualization
- SDN controller: deployment in drag-and-drop mode, IPv4 and IPv6, rollout of unicast and multicast full-stack services in minutes

Intelligent: enabling service agility



- Telemetry capability
- Microburst detection
- Edge analysis capability

Secure: best quality in the industry and pioneering energy-saving technology



- Microsegmentation used to isolate east-west traffic on switches (east-west traffic is isolated on firewalls originally)
- SFC used to divert traffic from the control plane to the data plane
- MACsec hardware encryption, providing high security and reliability

Open: easy integration and timely response to services



- Open API
- Interconnection with third-party management tools: Ansible
- Interconnection with third-party management tools or controllers: VMaare NSX

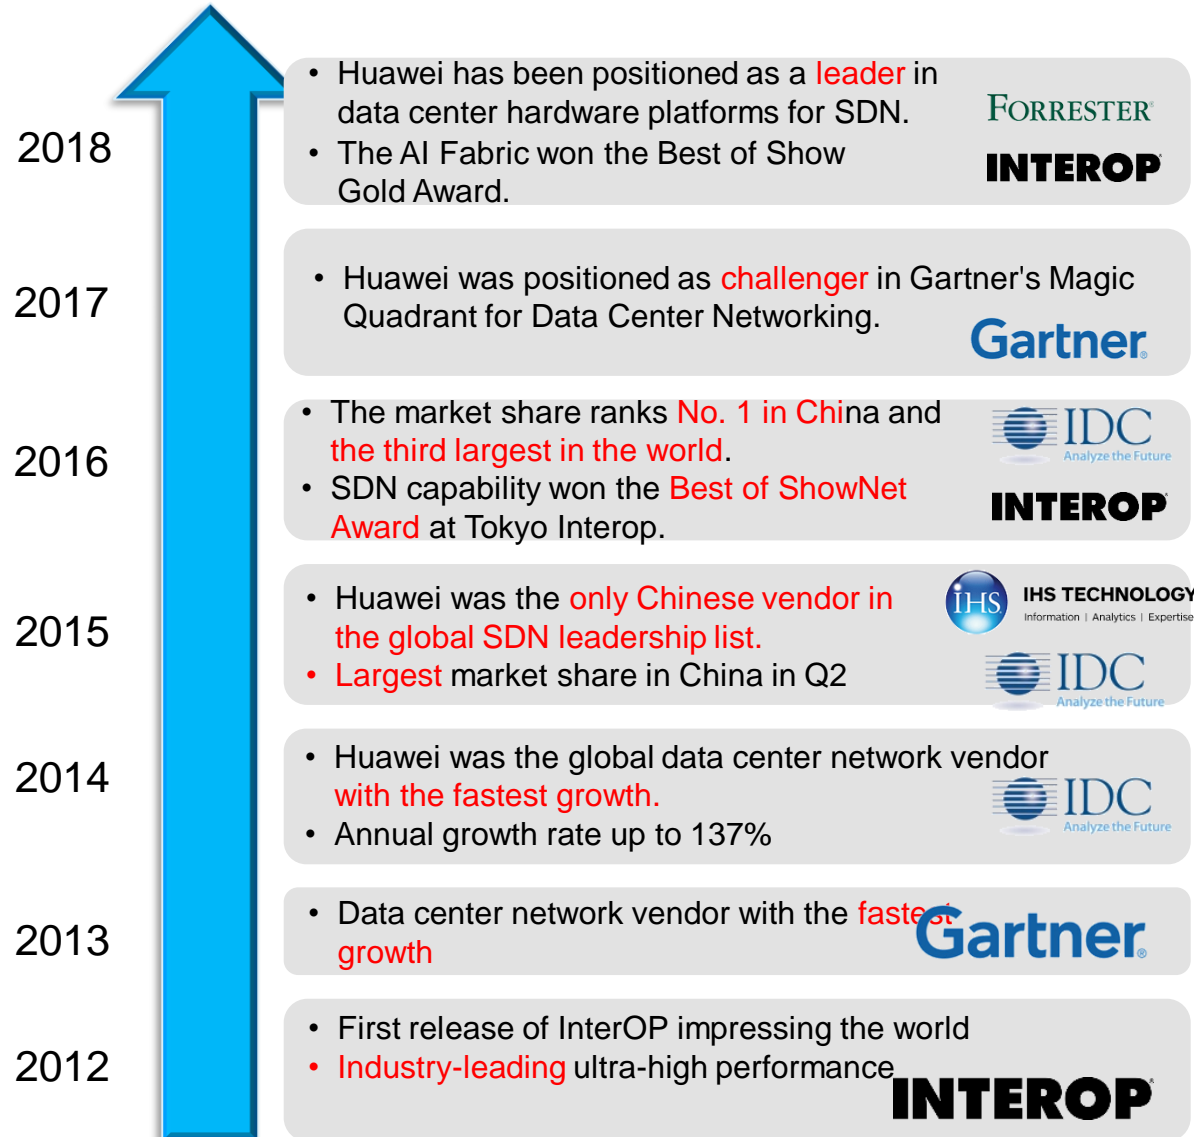
Contents

1 CloudEngine Switch Overview

2 CloudEngine Switch Highlights

3 CloudEngine Switch Market Progress

China's No.1 and One of World's Top 3 DCN Vendors



IHS TECHNOLOGY
Information | Analytics | Expertise

**Global SDN
Authoritative
Report of Leading
Vendors**

Vendors	Summary
BROADCOM	Data center merchant silicon pioneer providing SDN-aware chips and developer tools
BROCADE	Data center fabric innovator providing open source SDN controller distribution
CISCO	Ethernet switch market leader providing custom silicon-based SDN underlay
cumulus networks	Ethernet switch OS innovator provider disaggregating switch hardware from software
Hewlett Packard Enterprise	Early market pioneer with over 30 million OpenFlow-capable ports shipped
HUAWEI	Comprehensive SDN portfolio innovating with programmability in software and silicon
JUNIPER NETWORKS	4th largest data center network vendor driving SDN innovation in silicon and software
NEC	SDN pioneer with successful large-scale and geographically-distributed deployments
nuagenetworks	Early market entrant innovating BGP-based SDN network virtualization overlay
vmware	SDN virtualization innovator applying software abstraction techniques to network control

Source: IHS [“2015 Infonetics Data Center and Enterprise SDN Vendor Leadership Analysis”](#)

Highly Recognized Performance

InterOP Awards



In 2013, the CloudEngine 12800 won the Best of Show Award at Interop, which is the highest exhibition in the IT industry. Huawei is the first Chinese provider that wins the position.

In 2016, the CE8860 and CE6851 won the Best of Show Award at Interop.



Huawei's AI Fabric Intelligent Lossless Data Center Network Solution Takes Home Interop Tokyo Best of Show Award



Awards and Certifications



Award of Excellent Product Trusted by CIO



Award of the Most Competitive Product



Award of Excellent Product in Big Data



Award of Annual Excellent Technology



China SDN SDN Best Practice Award



Preferred Brand of Cloud Computing and Network Solution

CloudEngine Series Switches Serve 7800+ Global Customers

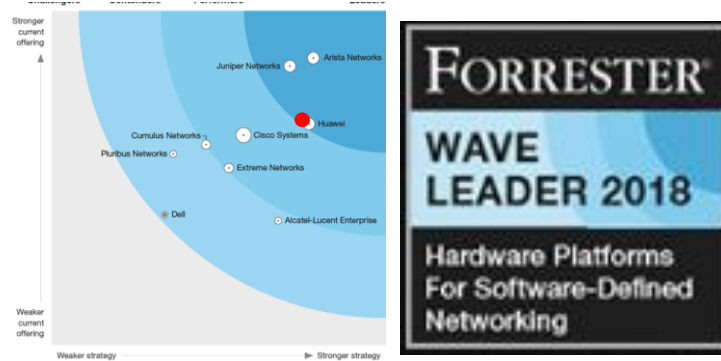
- The market share is **No.1** in China and **No.3** in the world.
- The global market share growth rate is **No. 1** for **four consecutive years**.



- **Over 32,000** CE12800 switches have been sold around the world, serving **7800+** customers in **120+** countries.



- 2018 Approaching the **Leaders Quadrant**
- 2017 **Challenger**



DC SDN
SDN hardware
platform **leader**



Gartner Peer Insights
Customers' Choice
for Data Center Networking

Thank you.

把数字世界带入每个人、每个家庭、
每个组织，构建万物互联的智能世界。

Bring digital to every person, home and
organization for a fully connected,
intelligent world.

**Copyright©2018 Huawei Technologies Co., Ltd.
All Rights Reserved.**

The information in this document may contain predictive statements including, without limitation, statements regarding the future financial and operating results, future product portfolio, new technology, etc. There are a number of factors that could cause actual results and developments to differ materially from those expressed or implied in the predictive statements. Therefore, such information is provided for reference purpose only and constitutes neither an offer nor an acceptance. Huawei may change the information at any time without notice.

